UNIVERSITÀ DEGLI STUDI DI PISA
DIPARTIMENTO DI INFORMATICA

DOTTORATO DI RICERCA IN INFORMATICA
UNIVERSITÀ DI PISA

PH.D. THESIS: TD-2/01

Soft Constraint Solving and Programming: a general framework.

Stefano Bistarelli

March 2001

ADDR: Corso Italia 40, 56125 Pisa, Italy. TEL:+39-50-887268. FAX: +39-50-887226.
E-MAIL: bista@di.unipi.it. HOME-PAGE: http://www.di.unipi.it/~bista/.

UNIVERSITÀ DEGLI STUDI DI PISA
DIPARTIMENTO DI INFORMATICA

DOTTORATO DI RICERCA IN INFORMATICA
UNIVERSITÀ DI PISA

PH.D. THESIS: TD-2/01

Soft Constraint Solving and Programming: a general framework.

Stefano Bistarelli

Abstract. The *Soft Constraints* idea is able to capture many real life situations that cannot be represented and solved with the classical crisp constraint framework. Several extensions were proposed in the last years, but we show that we can use

One Ring to rule them all,
One Ring to find them,
One Ring to bring them all and in the darkness bind them.
(from *"The Lord of the Rings"* of Tolkien, J.R.R.)

March 2001

ADDR: Corso Italia 40, 56125 Pisa, Italy. TEL:+39-50-887268. FAX: +39-50-887226.
E-MAIL: bista@di.unipi.it. HOME-PAGE: http://www.di.unipi.it/~bista/.

Thesis supervisor:
Prof. Ugo Montanari

Abstract

In this thesis we study the soft constraint solving and programming framework based on the notion of semiring. We show how the semiring carrier and operations can be instantiated to capture many *non-crisp* features representing fuzziness, optimization, probability, hierarchies, and others. The solution of each Soft Constraint Satisfaction Problem (CSP) is computed by using the appropriate $\times$ and $+$ semiring operations, which model constraint projection and combination respectively.

We show, moreover, its ability to represent several non-crisp constraint systems and how the properties of the semiring operations can give hints on the applicability of several solution and preprocessing methods. We give some sufficient conditions for their applicability and efficiency. In particular:

- we show that using an idempotent $\times$ operator the classical local consistency (and also dynamic programming) techniques can be used to reduce the complexity of the problem without modifying its solution;

- we adapt to the soft framework partial local consistency and labeling techniques, that perform fewer pruning steps of the domain. This means that, although they are able to remove fewer non-optimal solutions than are classical algorithms, partial local consistency algorithms can be beneficial because they are faster and easier to be implemented;

- we extend general local consistency algorithms to the soft framework that use several pruning rules until fix-point. In this case only a monotonic set of rules is needed.

Moreover, in the last part of the thesis, we discuss constraint programming: as classical constraint solving has been embedded into Constraint Logic Programming (CLP) systems, so too can a generalized notion of soft constraint solving be handled within a logic language, thus giving rise to new instances of the CLP scheme. This doesn't only give new meanings to constraint solving in CLP, but it also allows one to treat optimization problem solving within CLP in a uniform way, without the need to resort to *ad hoc* methods. In fact, we show that it is possible to generalize the semantics of CLP programs to consider the chosen semiring and thus solve problems according to the semiring operations. This is done by associating an element of the semiring to each ground atom and by using the two semiring operations to combine goals.

Acknowledgments

The year I spent working on my thesis was not only rich of technical and scientific discussion but also of human interaction and friendships. I'm happy to have the occasion to say some word to all the people that I met during my phd thesis work.

I wish to thank first of all Ugo, my supervisor; I spent much time discussing about constraints (and sometimes also about concurrency because he likes it more!) in the office in Pisa but also, and with more pleasure, in some nice restaurants around the world.

Many thanks are due to my external reviewers, Philippe Codognet and Pascal van Hentenryck, for their careful reading of a preliminary version of this thesis, and for their useful comments and suggestions.

I want also to thank Francesca: she shared with me the passion for constraints and I'm in debt with her for having helped develop the ideas collected in this thesis. She was (and is) also an important friend for me: she was able to stay close to me during the gray moments of my life in the same way that she likes to enjoy sharing with me my happy hours. She was for me a guide.

A special thanks also to my roommate friends at the office; the people of the "268-dream" (268 is the telephone number of the office :-)) are the most workaholic people that I met. To share a room with them means to find someone to talk (and to work) with at any time in the day (and night!). I was forgetting to write down the names: thanks to Alessio, Diego, Emilio, Paolo, Roberto (and also to the friends that shared the office with me only for a while: Mercedes and Fabio).

Some thanks and affectionate kisses to the women of the burocratic office of the department; they helped me in the "daily" problems of the department! Thanks to Letizia, Marilisa, Paola and Rita.

A thanks also to the people of room 266. I was their guest all the times that I had no desire to work (yes, sometimes it happens)! Thanks to Fabio, Chiara and Simone for the coffees.

Thanks also to Elvinia, Giampaolo, Rossella and Yan: I was so happy to have met them, and enjoyed writing some interesting papers together.

Thanks (grazie :-)) to my english speaking friends Emily, MariaTeresa and Olga: they were so kind to read all the conversational parts of my thesis to cut away some typing errors and making it just a little nicer!

Thanks again to some roommate friends; this time I thanks the friend that shared with me the tea of 5 pm: Giuseppe, Laura, Nadia, Valentina and Vincenzo.

. . . **and some personal thanks** . . .

It is now time to thank also the (many) friends I met not inside the wall of the department. First of all the ”sport-making” friends.

Thanks to all the people that I met each Wednesday in Livorno to play football. I met them only one hour each week but playing football with them was exciting (and helped me to relax from the office work). Here are some names: Ale, il Belli (1 and 2), Gambo (or Cambo? I never understand at all: his name was pronounced with a Livorno dialect!), Yuri, ...

Thanks to the friends who enjoyed spending their time working out and to our wonderful teacher (sometimes also this helped me :-)): thanks to Eleonora (yes, you guessed well, she’s the teacher), Marco, Simone and Primo.

Thanks also to the ”salsa-dancer” friends. I love to dance salsa, and the time spent with them was fabulous! Thanks to the friends that started learning the first step with me (Roberta, Alessandra, and Paolo) and then to all the other salsa dancers: thanks to Angela, Sergio, Laura, Rosa, and Monica.

I wish to thanks (before forgetting to do!) my ”fratellini”; Giampaolo, Leonardo and Sandro (with Simona too) are more than friends! No word could be enough to say all my affection.

And speaking of brothers (”fratellini”is the italian word for brothers) I wish to hug my brothers and their wives: Andrea, Miriam, Riccardo and Susanna. Il vostro affetto e la vostra amicizia mi ha tante volte aiutato. Grazie di cuore.

A special thanks also to my father and to my mother. Non sempre è stato un rapporto facile. Ma questo capita sempre nei rapporti che contano.

Thanks also to all my friends of the town council of Livorno. I spent lots of time with them. A special thanks is for Marco and Angela. They was (and are) the people that I feel closer to me in all these years. A thanks also to Sabrina and Fabio, the ”beach-volley” friends.

Un grazie anche agli amici di Pistrino. Avere un posto tranquillo dove tornare a riposarsi ogni tanto e passare delle serate in tranquillità mi è servito. Grazie a Barbara, Deborah, Fabiano, Marco, Oscar, Sandro, Sauro; Grazie a Lorenzo (il mio cuginetto); grazie a Luca e alle sue bimbe ;-)!

A special thank also to Francesco, Andrea and Giuseppe (I spent fantastic weeks with them in several states of America), Daniela and her friends (she was very close to me when I needed it in the last two years). A thanks also to Mimmo, Roberto, Michele, Silvina, Paolo and Fabio. We made the best parties together.

At the end to be sure to not forget anyone I wish to thank all the people that enjoyed with me the “cocomerata” party. I organized it for 5 years and each time I made some new friend. Thanks to Sandro, Catia, Salvo, Cinzia, Katia, Claudia, Fabio, Sara, . . . ,... I do not put any other name, but I’m sure: pick a name, I will find a person with that name who will come to the “cocomerata” party next year.

Finally, a special thanks to Olga and Laura: they helped me a lot to manage some problems I had during the writing time of this thesis.

Contents

List of Figures

List of Tables

Introduction

Gandalf
Tolkien, J.R.R. The Lord of the Rings

Motivations and thesis contributions

This thesis is centered around the notion of constraint solving and programming. The interest in constraint satisfaction problems can be easily justified, since they represent a very powerful, general, and declarative knowledge representation formalism. In fact, many real-life situations can be faithfully described by a set of objects, together with some constraints among them. Examples of such situations can be found in several areas, like VLSI, graphics, typesetting, scheduling, planning, as well as CAD, decision-support systems and robotics.

The constraint idea comes from the early 1960's when Sutherland introduced Sketchpad [Sut63], the first interactive graphical interface that solved geometric constraints. After that came Fikes' REF-ARF [Fik70] and at the beginning of 1970's Montanari described fundamental properties of the constraints when applied to picture processing [Mon71]. Another study in finite domain constraint satisfaction was done at the end of 1970's in Laurière's ALICE [Lau78], a system developed to solve prediction/detection problems in geology. After these, several constraint languages have been proposed in the literature: the language of Steele ([Ste80]), CONSTRAINTS [SS80] of Sussman & Steele, Thinglab ([Bor81]) and Bertrand ([Lel88]).

From Sketchpad until now, a lot of research has been done and improvements made, and the classical constraint satisfaction problems (CSPs) [Mon74, Mac92] have been shown to be a very expressive and natural formalism to specify many kinds of real-life problems.

Today, the use of the constraint programming idea to solve many real-life problem is reality. Many important companies develop tools based on the constraint technology to solve *assignment*, *network management*, *scheduling*, *transport* and many other problems:

Assignment problems Assignment problems were one of the first type of industrial applications that were solved with the CLP technology. These problems usually have to handle two types of resources, and constraints among them, and try to assign one resource of the first kind to one of the second kind such that all constraints are satisfied.

An example is the stand allocation for airports, where aircrafts (the first kind of resources) must be parked on the available stands (the second kind of resources) during their stay at the airport. The first industrial CLP application was developed for the HIT container harbor in Hong Kong [Per91], using the language CHIP: the objective was to allocate berths to container ships in the harbor, in such a way that resources and stacking space for the containers is available. Other Hong Kong applications are at the airport, where a CHIP-based system is used to solve the counter allocation problem [CP97], and another constraint-based system, which uses the ILOG libraries, is used for the stand allocation problem since mid-1998 [CCTY99]. Another system, called APACHE [DS91], was a demonstrator for stand allocation at Roissy airport in Paris: the objective was to replan the allocation when a change of arrival/departure times occurred.

Personnel assignment Personnel assignment problems are a special case of assignment problems where one resource type consists of humans. This peculiarity makes them specific enough to be considered separately. In fact, changing work rules and regulations impose difficult constraints which evolve over time. Also, user preferences often lead to over-constrained problems, which have no solution satisfying all constraints. Another important aspect is the requirement to balance the work among different persons, which leads to hard optimization problems.

The Gymnaste system [CHV98] produces rosters for nurses in hospitals, and is used in many hospitals in France. Around 250 technicians and journalists of the French TV and radio station RFO are scheduled with the OPTI-SERVICE system [Col96], which generates weekly workplans from the individual activities which must be covered by different persons with different qualifications. The personnel for the TGV high-speed train bar and restaurant services is scheduled with the EPPER application [DJ96]: all services for a month are covered by a team of 250 people. Recently a distributed CLP system has been used to tackle a workforce management problem within British Telecom [KG98]. Also Telecom Italia is using a constraint-based system to schedule all its technical tasks (about 100.000 tasks involving 20.000 technicians) over its territory [GCM99]; the system which controls the whole workforce management has a scheduling module, called ARCO, which dispatches activities to technicians, and which makes extensive use of constraint propagation techniques.

Network management Another application domain for finite domain CLP is network management, where many different problems can be addressed and solved using CLP.

The LOCARIM system was developed by COSYTEC for France Telecom: starting from an architectural plan of a building, it proposes a cabling of the telecommunication network of the building. The PLANETS system, developed by the University of Catalonia in Barcelona for the Spanish electricity company, is a tool for electrical power network reconfiguration which allows to schedule maintenance operations by isolating network segments without disrupting customer services. The company Icon in Italy produces a load-balancing application which is controlling network flow for the inter-banking system in Italy. The Esprit project CLOCWiSe (IN 10316I) is using CHIP for the management and operational control of water systems. The planning of wireless digital networks for mobile communication has been tackled by the system POPULAR [FB97], written in ECLiPSe and then in CHR: the main advantages with respect to other approaches to this same problem (using traditional imperative programming) are a good balance among flexibility, efficiency, and rapid prototyping.

Scheduling problems Perhaps the most successful application domain for finite domain CLP are scheduling problems. Given a set of resources with given capacities, a set of activities with given durations and resource requirements, and a set of temporal constraints between activities, a "pure" scheduling problem consists of deciding when to execute each activity, so that both temporal constraints and resource constraints are satisfied.

A typical example of a constraint-based scheduling application is ATLAS [SC95], which schedules the production of herbicides at the Monsanto plant in Antwerp. The PLANE system [BCP92] is used by Dassault Aviation to plan the production of the military Mirage 2000 jet and the Falcon business jet. The objective is to minimize changes in the production rate, which has a high set-up cost, while finishing the aircraft just in time for delivery. The MOSES application was developed by COSYTEC for an animal feed producer in the UK: it schedules the production of compound food for different animal species, eliminating contamination risk and satisfying costumer demand with minimal cost. The FORWARDC system is a decision support system, based on CHIP, which is used in three oil refineries in Europe to tackle all the scheduling problems occurring in the process of crude oil arrival, processing, finished product blending and final delivery [GR97]. Recently, Xerox has adopted a constraint-based system for scheduling various tasks in reprographic machines (like pothocopiers, printers, fax machines, etc.); the role of the constraint-based scheduler is to determine the sequence of print making and to coordinate the time-sensitive activities of the various hardware modules that make up the machine configuration [FGS95]. Recent results on the tractability of classes of constraint problems have shown that such scheduling problems are indeed tractable, and thus amenable for an efficient solution [PJ99].

Transport problems A variety of transport problems have been tackled using constraints. These problems are often very complex due to their size, the number and variety of constraints, and the presence of complex constraints on possible routes. Moreover, often these problems have a personnel allocation problem as a sub-aspect, usually with complex safety and union regulations.

The COBRA system [SC98] generates diagrams representing workplans for train drivers of North Western Trains in the UK. For each week, around 25.000 activities must be scheduled in nearly 3.000 diagrams, taking a complex route network into account. The DAYSY Esprit project (8402) and the SAS-Pilot program [BKC94] considers the operational re-assignment of airline crews to flights. This same problem is tackled also by another system [FLMM97], which uses a combination of CLP and OR techniques. A recently developed system uses the ILOG constraint libraries to produce and optimize train operating plans for a freight railway company, by creating transport movements for rolling stock between sources, hubs and sinks while satisfying transport requirements [MR98].

All this constraint application, however, have some obvious limitations, mainly due to the fact that they do not appear to be very flexible when trying to represent real-life scenarios where the knowledge is neither completely available nor crisp. In fact, in such situations, the ability to state whether an instantiation of values to variables is allowed or not, is not sufficient or sometimes not even possible. For these reasons, it is natural to try to extend the CSP formalism in this direction.

For example [RHZ76, DFP93, Rut94], CSPs have been extended with the ability to associate with each tuple, or to each constraint, a level of preference, and with the possibility of combining constraints using min-max operations. This extended formalism was called Fuzzy CSPs (FCSPs). Other extensions concern the ability to model incomplete knowledge of the real problem [FL93], to solve over-constrained problems [FW92], and to represent cost optimization problems [Bel57, Bis94].

In this thesis we describe a constraint solving framework where most of these extensions, as well as classical CSPs, can be cast. The main idea is based on the observation that a semiring (that is, a domain plus two operations satisfying certain properties) is all that is needed to describe many constraint satisfaction schemes. In fact, the domain of the semiring provides the levels of consistency (which can be interpreted as either cost, degrees of preference, probabilities, or others), and the two operations define a way to combine constraints together. More precisely, we define the notion of constraint solving over any semiring. Specific choices of the semiring will then give rise to different instances of the framework, which may correspond to known or new constraint solving schemes.

With the definition of the new framework, we also need to extend the classical AI local consistency techniques [Fre78, Fre88, Mac92, Mac77, Mon74, MR91] that have proven to be very effective when approximating the solution of a problem in order to be applied in the semiring framework. We may say that all of them are variants or extensions of two basic algorithms, which started the field. One is the *arc consistency* algorithm (used for the first time in [Wal75] but studied and named later in [Mac77], where we may say that inconsistency is eliminated by each unary constraint (i.e., involving only one variable); or, that each unary constraint (i.e., connecting only one variable, say v) is propagated to all variables attached to the variable v by some binary constraint. The other pioneering algorithm is the *path-consistency* algorithm (which was proposed in [Mon74]), where inconsistency is eliminated by each binary constraint; or, in other words, where each binary constraint connected to variables v and v' is propagated to all other variables attached to v and v' by some binary constraint. Then, a more general one is the *k-consistency* algorithm scheme [Fre78], where inconsistency is eliminated by each set of constraints involving k variables (it is easy to see that arc-consistency is 2-consistency and that path-consistency is 3-consistency).

We generalize these techniques to our framework, and we provide sufficient conditions over the semiring operations which assure that they can also be fruitfully applied to the considered scheme. By “fruitfully applicable” we mean that 1) the algorithm terminates and 2) the resulting problem is equivalent to the given one and it does not depend on the nondeterministic choices made during the algorithm. The advantage of the SCSP framework is that one can hopefully see one’s own constraint solving paradigm as an instance of SCSP over a certain semiring, and can inherit the results obtained for the general scheme. In particular, one can immediately see whether a local consistency and/or a dynamic programming technique can be applied.

Moreover, the study of generalized versions of AI algorithms for the SCSP framework seems to be very important. In particular, the *partial local consistency* algorithms seem to be efficient and useful. More precisely, we study the behavior of a class of algorithms able to eliminate from a semiring-based constraint problem some solutions that are not *interesting* (where *interesting* might mean optimal) and we will study their application to SCSPs.

In order to solve an SCSP problem in a faster way, we also try to ”simplify” it by changing the structure of the semiring. The solutions obtained in the easier semiring have to be mapped back and used to solve the original problem. In order to capture some interesting properties we use a notion of *abstract interpretation* to represent this mapping. The idea is to take an SCSP problem P and map it over a new problem P' which is easier to handle and solve. As soon as we have solved (or partially solved) P' we have to map back the solutions (or the partial solutions) over the original problem. Depending on the properties of the semiring operations, we can perform some simplification steps over the original problem.

We next demonstrate how to extend the logic programming paradigm to deal with semiring based constraints. One of the ways to program with constraint is to dip the constraints in the logic programming theory obtaining the CLP framework. Programming in CLP means choosing a constraint system for a specific class of constraints (for example, linear arithmetic constraints, or finite domain constraints) and embedding it into a logic programming engine. This approach is very flexible because one can choose among many constraint systems without changing the overall programming language, and it has been shown to be very successful in specifying and solving complex problems in terms of constraints of various kinds. It can, however, only handle classical constraint solving. Thus, it is natural to try to extend the CLP formalism in order to be able to also handle SCSP problems. We will call such an extension SCLP (for Semiring-based CLP).

In passing from CLP to SCLP languages, we replace classical constraints with the more general SCSP constraints. By doing this, we also have to modify the notions of interpretation, model, model intersection, and others, since we have to take into account the semiring operations and not the usual CLP operations. For example, while CLP interpretations associate a truth value (either *true* or *false*) to each ground atom, here, ground atoms must be given one of the elements of the semiring. Also, while in CLP the value associated to an existentially quantified atom is the *logical or* among the truth values associated to each of its instantiations, here we have to replace the *or* with another operation which refers to one of the semiring ones. The same reasoning applies to the value of a conjunction of atomic formulas: instead of using the logical *and* we have to use a specific operation of the semiring.

After describing the syntax of SCLP programs, we will define three equivalent semantics for such languages: model-theoretic, fix-point, and operational. These semantics are conservative extensions of the corresponding ones for Logic Programming (LP), since by choosing a particular semiring (the one with just two elements, *true* and *false*, and the logical *and* and *or* as the two semiring operations) we get exactly the LP semantics. The extension is in some cases predictable but it possesses some crucial new features. We also show the equivalence of the three semantics. In particular, we prove that, given the set of all refutations starting from a given goal, it is possible to derive the declarative meaning of both the existential closure of the goal and its universal closure. Moreover, we will investigate the decidability of the semantics of SCLP programs and we will prove that if a goal has a semiring value greater than a certain value in the semiring, then we can discover this in finite time.

Moreover, we study the classical Operational Research (OR) problems and their embedding in the SCLP framework. This seems to be able to 1) give a new semantics w.r.t. the algorithmic one and 2) give some hints towards an extension of the SCLP framework to handle multilevel and multisorted problems. More precisely, the "levels of preference" together with a constraint framework can be used fruitfully to describe and solve optimization problems. In fact, solving a problem does not mean finding "a solution", but finding the best one (in the sense of the semiring). This intuition suggests we investigate around the operational research field in order to find out how to describe and solve classical problems that usually are solved using OR techniques.

We also give suggestions on the use of the SCSP framework to discovery security problems in the network while a protocol session is executed. To do this we first model the policy of the protocol and the initial state of the network as SCSPs, and then we compare the policy SCSP with the initial one, modified by the protocol session.

Outline of the work and main results

In Chapter 1 we introduce some background material: some basic notions related to Constraint Satisfaction Problems (CSPs) and to the Constraint Logic Programming (CLP) framework.

In Chapter 2 we introduce the Soft CSP (SCSP) framework. The framework is based on the notion of semiring, that is, a set plus two operations. Different instantiations of the two operations give rise to different existing frameworks (Fuzzy CSP, Probabilistic, Classical, VCSP, etc.). Each of these frameworks satisfies specific properties that are useful in testing the applicability of the constraint preprocessing and solving techniques described in Chapter 3.

In Chapter 3 the properties studied in the previous chapter are applied to the different frameworks in order to prove the applicability of soft local consistency and dynamic programming techniques. These techniques are defined starting from the classical ones, by using the semiring operations. We also generalize the local consistency technique by reducing the needed properties of the semiring operators and by using different functions in the local consistency steps. Some ideas on the efficient applicability of these techniques are given, by looking at the graph structure of the problem.

In Chapter 4 classical techniques of abstraction are applied to SCSPs. The problem is mapped in an ”easier” one, some solution related information is found and mapped back to the original problem. The soft local consistency techniques are also proven to be useful in the abstract framework to find the set of optimal solutions or one of its approximations.

In Chapter 5 a more structured view of the SCSP framework is given: domains and semirings are used as basic elements for building high order functions representing solution or preprocessing algorithms.

In chapter 6 the CLP framework is extended to deal with soft constraints. We define the Soft CLP (SCLP) framework and we give an operational, model-theoretic and fixpoint semantics. Some equivalence and properties of the semantics are proven.

In Chapter 7 we show how the ability of the CLP language to deal with soft constraints can be used to represent and solve in a declarative fashion OR problems. As an example we consider general shortest path problems over partially ordered criteria and we show how the SCLP framework can easily represent and solve them. Moreover, we give a semantics to a class of SCLP programs by using a classical shortest-path algorithm.

In Chapter 8 we try to enlarge the applicability domain of the constraint framework, by defining a mapping between a network of agents executing a protocol and an SCSP. Some security properties of the network are translated into specific properties of the SCSP solutions.

The Origin of the Chapters

Many chapters of this thesis are based on already published papers, jointly with Giampaolo Bella, Philippe Codognet, Helene Fargier, Rossella Gennari, Yan Georget, Ugo Montanari, Elvinia Riccobene, Francesca Rossi, Thomas Schiex and Gerard Verfaillie. In particular:

- The basic definitions and properties of the soft constraints' framework described in Chapter 2 appear in [BMR95, BMR97b]; moreover the mapping between SCSPs and VCSPs [SFV95] is based on the ideas developed in [BFM$^+$96, BFM$^+$99];
- The solutions and preprocessing techniques described in Chapter 3 are instead introduced in [BMR95, BMR97b] and extended in [BR98, BCGR00b, BGR00];
- Chapter 4 describe the abstraction framework and develops the ideas presented in [BCGR00a, BCR00];
- The language to program with soft constraints of Chapter 6 is introduced in [BMR97a, BR97b]; an extended version of the work appears in [BMR01];
- Chapter 7 is based on the ideas developed in [BMR99, BMR00b].

Chapter 5 and Chapter 8 are instead based on some preliminary work presented respectively in [BMR00a] and in [BB00, BB01].

Chapter 1

Background

> *"Constraint programming represents one of the closest approaches computer science has yet made to the Holy Grail of programming: the user states the problem, the computer solves it."*
> Eugene C. Freuder, Constraints, April 1997

Abstract

Constraint programming is an emergent software technology for declarative description and effective solving of large, particularly combinatorial, problems especially in areas of planning and scheduling. It has recently emerged as a research area that combines researchers from a number of fields, including Artificial Intelligence, Programming Languages, Symbolic Computing and Computational Logic. Constraint networks and constraint satisfaction problems have been studied in Artificial Intelligence starting from the seventies. Systematic use of constraints in programming has started in the eighties. In constraint programming the programming process consists of the generation of requirements (constraints) and solution of these requirements, by specialized constraint solvers.

Constraint programming has been successfully applied in numerous domains. Recent applications include computer graphics (to express geometric coherence in the case of scene analysis), natural language processing (construction of efficient parsers), database systems (to ensure and/or restore consistency of the data), operations research problems (like optimization problems), molecular biology (DNA sequencing), business applications (option trading), electrical engineering (to locate faults) and circuit design (to compute layouts).

In this chapter we review some definitions and notions which will be fundamental for our work and we make some considerations, which are the basis for the motivation of this thesis. The chapter is divided into two parts: In the first one we define the Constraint Satisfaction Problems (CSPs), we introduce the classical solving techniques and we give a state of the art of the "non-crisp" extensions that are developed; In the second one, we introduce instead the Constraint Logic Programming (CLP) framework. There is also some other background material that will be used later in the thesis. That material, however, will be introduced only when needed.

1.1 Constraint Satisfaction Problems

Constraint Satisfaction Problems (CSPs) are a very powerful and general knowledge representation formalism, since many real situations can be faithfully described by a set of objects, together with some constraints among them. They were a subject of research in Artificial Intelligence for many years starting from the 1970s. From that date several classes of constraints were studied (e.g., temporal constraint problems, constraint problems over intervals and over reals, linear constraints) but in this thesis we will concentrate on the specific class of *constraints over Finite Domain*.

We believe that finite domain constraint problems [Mon74, DP88a, Fre88, Mac77, MR91], i.e., constraint problems where the objects may assume a finite number of distinct configurations, are of special interest, since they are significantly simpler to analyze than general constraint problems, while still being able to express, either faithfully or through a certain degree of abstraction, the main features of several classes of real life problems.

Before providing the formal definition of the CSP and its solution methods, let's look at some examples [Rut98], also to illustrate the variety of the potential application fields.

EXAMPLE 1.1.1 (8-QUEENS AND GRAPH COLOURING) Famous test-problem popular also in the CSP world is the 8-queens problem: place 8 queens on the chess board such that they do not attack each other. In order to formulate this problem as a CSP, the location of the queens should be given by variables, and the "do not attack each other" requirement should be expressed in terms of a number of constraints. A simple way to do this is to assign a variable to each queen.

As the 8 queens must be placed in 8 different columns, we can identify each queen by its column, and represent its position by a variable which indicates the row of the queen in question. Let x_i stand for the row of the queen in the i-th column. The domain of each of the variables $x_1, \ldots, x_8$ is $\{1, 2, ...8\}$. For any two different variables the following two constraints must hold, expressing that the queens should be in different rows and on different diagonals:

$$x_i \neq x_j$$
$$|x_i - x_j| \neq |i - j|$$

In this formulation of the problem, we have to find a solution out of the total possible instantiations of the variables, which is 8^8. This formulation, though seems natural, does contain a trick: a part of the requirements of the problem is reflected in the representation, not in the constraints. We could have used the most straightforward representation, namely identifying the squares of the chess board by the $1, 2, \ldots, 64$ numbers, and having

8 variables for the 8 queens all with the domain $\{1, 2, \ldots, 64\}$. In this case, the "different columns" requirement should be expressed too by constraints, and all the three types of constraints become more intrinsic to formulate. The total number of possible arrangements becomes as large as 64^{64}, containing a configuration of queens multiple times due to the identification of the 8 queens. So we have many reasons to prefer the first representation over the second one. It is true in general that a problem can be formulated as a CSP in a number of ways. The resulting CSPs may differ significantly considering the number and complexity of the constraints and the number of the possible instantiations of the variables, and thus may require very different amount of time and memory to be dealt with. Hence when modelling a problem as a CSP, one has to pay attention to different possibilities, and try to commit to the one which will be the easiest to cope with. The in-depth analysis of the different solution methods and of the characteristics of the CSPs may provide a basis to make a good choice. Several cases have been reported when a notoriously difficult problem could be solved finally as a result of change of the representation. Both representations of the 8-queens problem are pleasantly regular: the domain of all the variables is the same, all the constraints refer to 2 variables, and for each pair of variables the same type of constraints are prescribed. Hence the 8-queens problem is not appropriate as a test case for solution algorithms developed to solve general CSPs. In spite of this intuitive observation, earlier the 8-queens had been a favourite test problem: there had been a race to develop search algorithms which were able to solve the problem for bigger and bigger n. (It is possible to construct a solution analytically.) This practice was stopped by two discoveries. On the one hand, Sosic [SG91] came up with a polynomial-time search algorithm, which was heavily exploiting the above mentioned special characteristics of the problem. On the other hand, by analysing the search space of the n-queens problem, it was shown that the general belief that "the bigger the n the more difficult the problem is" does not hold - in fact, the truth is just the opposite [Mor92]. Another, equally popular test problem is graph colouring: colour the vertices of a given graph using k colours in such a way that connected vertices get different colours. It is obvious how to turn this problem into a CSP: there are as many variables as vertices, and the domain for each variable is $\{1, 2, \ldots, k\}$, where k is the allowed number of colours to be used. If there is an edge between the vertices represented by the variables x_i and x_j, then there is a constraint referring to these two variables, namely: $x_i \neq x_j$. Though for the first sight graph colouring may seem to be just as a toy problem as the n-queens, there are basic differences between the two problems. First of all, graph colouring is known to be NP-complete, so one does not expect a polynomial-time search algorithm to be found. Secondly, it is easy to generate a great number of test graphs with certain parameters, which are more or less difficult to be coloured, so the family of graph colouring problems is appropriate to test algorithms thoroughly. Finally, many practical problems, like ones from the field of scheduling and planning, can be expressed as an appropriate graph colouring problem. □

Let's now provide a formal definition of constraint satisfaction problem:

DEFINITION 1.1.2 (CONSTRAINT SATISFACTION PROBLEM) *A Constraint Satisfac-*

tion Problem is a tuple $\langle V, D, C, \mathrm{con}, \mathrm{def}, a\rangle$ *where*

- V *is a finite set of variables, i.e.,* $V = \{v_1, \ldots, v_n\}$*;*
- D *is a set of values, called the domain;*
- C *is a finite set of constraints, i.e.,* $C = \{c_1, \ldots, c_m\}$*.* C *is ranked, i.e.* $C = \bigcup_k C_k$*, such that* $c \in C_k$ *if* c *involves* k *variables;*
- con *is called the connection function and it is such that*

  $$\mathrm{con} : \bigcup_k (C_k \to V^k),$$

  where $\mathrm{con}(c) = \langle v_1, \ldots, v_k\rangle$ *is the tuple of variables involved in* $c \in C_k$*;*
- def *is called the definition function and it is such that*

  $$\mathrm{def} : \bigcup_k (C_k \to \wp(D^k)),$$

  where $\wp(D^k)$ *is the powerset of* D^k*, that is, all the possible subsets of* k*-tuple in* D^k*;*
- $a \subseteq V$*, and represent the* distinguished *variables of the problem.*

In words, function con describes which variables are involved in which constraint, while function def specifies which are the domain tuples permitted by the constraint. The set a is used to point out the variables of interest in the given CSP, i.e., the variables for which we want to know the possible assignments, compatibly with all the constraints.

Note that other classical definitions of constraint problems do not have the notion of distinguished variables, and thus it is as if all variables are of interest. We choose this definition, however, for two reasons: first, the classical definition can be trivially simulated by having a set of distinguished variables containing all the variables, and second, we think that this definition is much more realistic.

In fact, it is reasonable to think that the CSP representation of a problem contains many details (in terms of constraints and/or variables) which are needed for a correct specification of the problem but are not important as far as the solution of the problem is concerned. In other words, the non distinguished variables play the role of existential variables: we want to assure that they can be assigned to values consistently with the constraints, but we do not care to know the assignment.

The solution $\mathrm{Sol}(P)$ of a CSP $P = \langle V, D, C, \mathrm{con}, \mathrm{def}, a\rangle$ is defined as the set of all instantiations of the variables in a which can be extended to instantiations of all the variables which are consistent with all the constraints in C.

DEFINITION 1.1.3 (TUPLE PROJECTION AND CSP SOLUTION) *Given a tuple of domain values* $\langle v_1, \ldots, v_n \rangle$*, consider a tuple of variables* $\langle x_{i1}, \ldots, x_{im} \rangle$ *such that* $\forall j = 1, \ldots, m$*, there exists a* k_j *with* $k_j \in \{1, \ldots, n\}$ *such that* $x_{ij} = x_{k_j}$*. Then the projection of* $\langle v_1, \ldots, v_n \rangle$ *over* $\langle x_{i1}, \ldots, x_{im} \rangle$*, written* $\langle v_1, \ldots, v_n \rangle_{|\langle x_{i1}, \ldots, x_{im} \rangle}$*, is the tuple of values* $\langle v_{i1}, \ldots, v_{im} \rangle$*. The solution* $\mathtt{Sol}(P)$ *of a CSP* $P = \langle V, D, C, \mathtt{con}, \mathtt{def}, a \rangle$ *is defined as*

$$\{\langle v_1, \ldots, v_n \rangle_{|a} \text{ such that } \begin{cases} v_i \in D \text{ for all } i; \\ \forall c \in C, \langle v_1, \ldots, v_n \rangle_{|\mathtt{con}(c)} \in \mathtt{def}(c). \end{cases} \}$$

The solution to a CSP is hence an assignment of a value from its domain to every variable, in such a way that every constraint is satisfied. We may want to find just one solution, with no preference as to which one, or all solutions.

To give a graphical representation of a CSP problem, we use a labelled hypergraph which is usually called a *constraint graph* ([DP88a]).

DEFINITION 1.1.4 (LABELLED HYPERGRAPH) *Given a set of labels* L*, a hypergraph labelled over* L *is a tuple* $\langle N, H, c, l \rangle$*, where* N *is a set of nodes,* H *is a set of hyperarcs,* $c : \bigcup_k (H_k \to N^k)$*, and* $l : \bigcup_k (H_k \to \wp(L^k))$*. I.e.,* c *gives the tuple of nodes connected by each hyperarc, and* l *gives the label of each hyperarc.*

DEFINITION 1.1.5 (FROM CSPS TO LABELLED HYPERGRAPHS) *Consider a CSP* $P = \langle V, D, C, \mathtt{con}, \mathtt{def}, a \rangle$*. Then the labelled hypergraph corresponding to* P *, written* $G(P)$*, is defined as the hypergraph* $G(P) = \langle V, C, \mathtt{con}, \mathtt{def} \rangle$ *labelled over* D*.*

In the hypergraph corresponding to a CSP, the nodes represent the variables of the problem, and the hyperarcs represent the constraints. In particular, each constraint c can be represented as a hyperarc connecting the nodes representing the variables in $\mathtt{con}(c)$. Constraint definitions are instead represented as labels of hyperarcs. More precisely, the label of the hyperarc representing constraint c will be $\mathtt{def}(c)$.

The representation of a CSP by a hypergraph turns out to be very useful when focusing on properties of the CSP which involve notions of sparseness and/or locality. In particular, locality can be defined in term of subgraphs, where a subgraph of a given graph consists of a subset S of the nodes together with some of the hyperarcs connecting subsets of S. For example, the theory of *local consistency* techniques for discrete CSPs (which we will review later), whose aim is to remove local inconsistencies, can be given in terms of operations on hypergraphs.

As for the notion of solution of a finite domain CSP, consider for example the CSP depicted in Figure 1.1, where each arc is labeled by the definition of the corresponding constraint, given in terms of tuples of values of the domain, and the distinguished variables are marked with a *. The solution of this problem is the set $\{\langle a, b \rangle\}$.

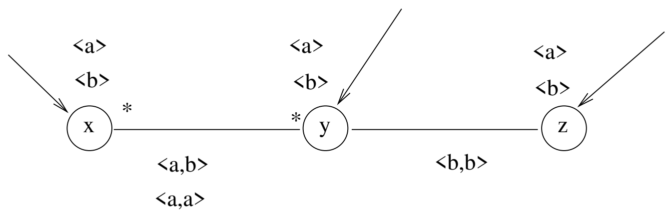

Figure 1.1: A CSP which is not solved.

1.1.1 Solution techniques

The methods to generate a solution for a CSP fall into three classes [Rut98]. In the first class are the variants of backtracking search. These algorithms construct a solution by extending a partial instantiation step by step, relying on different heuristics and using more or less intelligent backtracking strategies to recover from dead ends. The reduction of a problem is advantageous, resulting in a smaller solution space to be searched. The so-called constraint propagation algorithms do eliminate some non-solution elements from the search space. In general, these algorithms do not eliminate all the non-solution elements, hence, they do not produce a solution on their own. They are used either to preprocess the problem before another type of algorithm is applied, or interwoven with steps of another kind of - e.g. backtracking search - algorithm to boost its performance. Finally, the structure-driven algorithms exploit the structure of the primary or dual graph of the problem. There are very different algorithms in this class, including ones which decompose the initial CSP into loosely-coupled subproblems, which can be solved by methods from the previous two classes. Hence, structure-based methods can be also coupled with some other types of algorithms. All the algorithms from the above three classes investigate the solution space systematically. Hence all those algorithms of the above classes which are meant to find a solution, in theory really do the job as long as there is a solution. These algorithms are:

- *sound*, that is if they terminate with a complete instantiation of the variables than it is a solution;
- *complete*, that is capable to investigate the entire search space and hence find all the solutions.

These requirements seem to be very essential, however, often one has to be satisfied with algorithms which do not fulfill one or both of them. The systematic search algorithms require exponential time for the most difficult CSPs. A CSP is difficult if (almost) the entire search space has to be investigated before finding a solution or concluding that the problem has none. If the search space is large, then it may take days or weeks to run a complete and sound algorithm. This can be forbidding in case of applications where a solution can be used only if provided within a short time. In such cases a compromise is made, by using an algorithm which provides an answer fast, but the answer is not guaranteed to be a solution. However, it is "good enough" in the sense that not all the constraints are satisfied, but the number of non-satisfied constraints the and degree of violations can be accepted. Though such an algorithms cannot be used to generate all the (good) solutions for sure, usually it is possible to generate several quite different "almost solutions" (if they exist). The so-called local stochastic search algorithms have in common that they explore the solution space in a non-systematic way, stepping from one complete instantiation to another, based on random choices, and may navigate on the basis of heuristics, often adopted from systematic search methods. In the recent years such algorithms have been used with success to solve large practical problems, and they are suitable to handle CSPs extended with some objective function to be optimised.

Constraint Propagation

By eliminating redundant values from the problem definition, the size of the solution space decreases. Reduction of the problem can be done once, as pre-processing step for another algorithm, or step by step, interwoven with the exploration of the solution space by a search algorithm. In the latter case, subsets of the solution space are cut off, saving the search algorithm the effort of systematically investigating the eliminated elements, which otherwise would happen, even repeatedly. If as a result of reduction any domain becomes empty, then it is known immediately that the problem has no solution. One should be careful with not spending more effort on reduction than what will "pay off" in the boosted performance of the search algorithm to be used to find a solution of the reduced problem. The reduction algorithms eliminate values by propagating constraints. The amount of constraint propagation is characterised by the consistency level of the problem, hence these algorithms are also called *consistency-algorithms*. The iterative process of achieving a level of consistency is sometimes referred to as the relaxation process, which should not be mixed up with relaxation of constraints. Constraint propagation has a long tradition in CSP research. Below we introduce the most well-known and widely used algorithms.

Node- and arc-consistency A CSP is *node-consistent*, if all the unary constraints hold for all the elements of the domains. The straightforward node-consistency algorithm (NC), which removes the redundant elements by checking the domains one after the other has $O(dn)$ time complexity, where d is the maxim of the size of the domains and n is the number of the variables. A CSP is arc-consistent, if for any u value from the domain of any variable x, any binary constraints which refers to x can be satisfied.

The very first AC-1 algorithm has been followed by the algorithms AC-2, ..., AC-7, each improving the bookkeeping of the changed domains and the administration of the iterations. The time complexity of AC-3 is $O(d^3e)$, memory complexity $O(e+nd)$, where e is the number of binary constraints. AC-4 [MH86] improves AC-3 with $O(d^2e)$ time and space complexity. Another improvement of AC-3 is the algorithm AC-5 [HDT92], which exploits the semantics of special types of constraints. The AC-7 algorithm uses the symmetry of binary constraints.

k**-consistency** Arc-consistency can be also understood as telling something about how far a partial solution can always be extended. Namely, any partial solution containing only one instantiated variable can be extended by instantiating any second variable to a properly chosen value. Applying the same principle for more variables, we arrive at the concept of k-consistency.

DEFINITION 1.1.6 (k-CONSISTENCY) *A CSP is k-consistent, if any consistent instantiation of any $k-1$ variables can be extended by instantiating any one of the remaining variables.*

For binary CSPs there is yet another type of consistency used:

DEFINITION 1.1.7 (PATH CONSISTENCY) *A binary CSP is path-consistent, if for any path in its constraint graph it holds that if the assignments of the starting and ending variables are consistent, then this can be extended to a consistent partial instantiation by assigning the remaining variables along the path.*

Similarly to the family of arc-consistency algorithms, there is a family of algorithms to achieve path-consistency of a binary CSP. These algorithms also assure node- and arc-consistency. The best of them, PC-4 which is the analogue of AC-4, has $O(d^3n^3)$ time- and space-complexity. It is important to understand clearly the significance of consistency. A CSP which is k-consistent, is not necessarily solvable, and in the other way around, the solvability of a problem does not imply any level of consistency, not alone 1-consistency. The consistency as a feature of a problem does guarantee that certain values and value h-tuples which are not in the projection of the solution set have been removed form the domains and constraints. The level of consistency, k indicates for what h-values has this been done.

In general, the main idea is to improve the brute force backtracking by cutting down the search space. This is done by first modifying the original constraint network to a more *explicit* one (by eliminating redundancy), and then to run the search.

In a finite domain CSP, redundancy can be found mainly in the definitions of the constraints. In other words, a tuple of domain values in a constraint definition may be considered to be redundant if its elimination does not change the solution of the problem.

A sufficient condition for tuple redundancy is that the tuple be inconsistent with (all the tuples in) some other constraint in the CSP. In fact, if this condition holds, then the tuple will never be used in any solution of the CSP (otherwise the other constraint would be violated), and thus the solution will not change if the tuple is deleted. Consider, for example, a CSP which, among other variables and constraints, contains the variables x, y, and z, and the two constraints as depicted in Figure 1.2.

Then, tuple $\langle a, a\rangle$ of the constraint between x and y is inconsistent with the constraint between y and z. In fact, if x and y are both instantiated to a, then there is no way to instantiate z so as to satisfy the constraint between y and z. Therefore, tuple $\langle a, a\rangle$ is redundant and may be eliminated from the CSP, yielding a new CSP which is obviously equivalent to the old one. More precisely, given a CSP P_1 , the aim is to obtain a new CSP

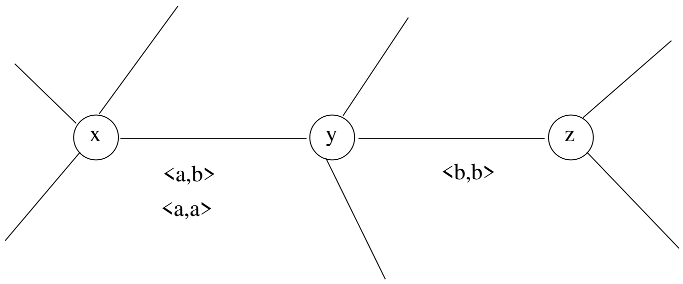

Figure 1.2: Tuple redundancy.

P_2 such that P_1 and P_2 have the same solution and the same constraint graph, and P_2 has smaller constraint definitions.

Given that in order to check the redundant status of a tuple we look at some constraints which are adjacent to the one containing that tuple (because constraints not sharing variables with the given one would not help), usually this kind of redundancy is also called *local inconsistency*. In fact, we do not look at the whole CSP, but only at some small portion of it (which contains the constraint with the tuple under consideration). Local inconsistency obviously implies global inconsistency, and this is why the condition of local inconsistency is sufficient for redundancy.

Search Algorithms

A CSP can be solved by systematically exploring the solution space by an uninformed search. The algorithms instantiate variables one after the other in such a way that the partial instantiation is always consistent. If this is not possible for a variable on turn, that is, all the possible values are in conflict with some earlier assignment, then backtracking takes place.

Chronological backtracking The simplest search scenario is when the order of variables as well as the order of values to be considered next is fixed, and if a dead end occurs, then the latest instantiation is reconsidered. In this uninformed, depth-first search no heuristics used for any of the choices made during exploration.

The worst case is when the algorithm has to generate and check all the complete instantiations to find the only solution or to conclude that there is none. This requires $O(d^m)$ steps and $O(d^n)$ memory, where d is the cardinality of the domain, m is the number of constraints and n is the number of variables.

The CB algorithm can be improved at different points in the following ways:

1. Whenever a choice has to be made, some heuristic is to be used instead of the fixed lexicographical ordering given by the description of the problem.

2. After each variable instantiation, a part of the search space not containing solutions is identified and discarded from further investigation.

3. Analysing the structure of the problem, the chronological backtracking is replaced by some intelligent backtracking.

4. The algorithm learns while exploring the search space: both negative and positive conclusions on partial instantiations are remembered and reused later if appropriate, instead of redoing the necessary checks again.

Ordering heuristics variable instantiations should be done in such an order that we try to avoid deep backtracking in course of the search. A deep backtracking means that a big part of the variable instantiations have to be undone, often repeatedly many times. (This annoying phenomena in search got a distinct name: thrashing.) The reason for this is that the variable, which cannot be instantiated properly for a big set of partial instantiations, is dealt with too late. Hence, intuitively, one would deal with the critical variables first.

Different heuristics are used to judge how critical variables are, based on the number of possible values for the variable and the already satisfied and remaining constraints. The heuristic-based selection, whatever heuristic is used, replaces the uninformed pick_one_variable process. If the order of variables is set in advance, then we talk about static variable ordering. Otherwise, in the case of dynamic variable ordering, the next variable to be instantiated depends on the current state of the search.

The *fail first principle (FFP)*, which is a general search heuristic, in our case of variable selection means that variables with the least possible values should be instantiated first. This principle can be applied in a static way, if the initial domain of the variable is taken into account, or in a dynamic way, if the domains of uninstantiated variables may get reduced due to the already committed instantiations.

The *minimum width ordering (MWO)* heuristic is based on the idea that if a variable does not depend on many earlier instantiated variables, then it will be easier to assign an appropriate value to it. Hence the smaller the width of an ordering is, the fewer backtracking can be expected. The heuristics takes a static ordering of the variables, namely one with minimal width. If the width of a CSP is w, then taking an ordering of the variables with this minimal width, for each variable on turn at most $w-1$ earlier instantiated variables are to be taken into account.

The *maximum cardinality ordering (MCO)* heuristic, which takes the variables in the decreasing order of cardinality, can be considered an approximation of the MWO heuristic. This cheaper heuristic has been used with success.

Another heuristic is based on the idea that variables which constraint each other should be close to each other in the order of instantiations. This principle is reflected in the *minimum bandwidth ordering heuristic (MBO)*, which takes a variable ordering with minimal bandwidth. The bandwidth of an ordering is the maximum of the bandwidth of the vertices in the ordering. The bandwidth of a vertex in an ordering is the maximum distance of the vertex from all the vertices in the ordering which are constrained by the given vertex.

We have seen that it is advantageous to consider the most critical variables to be instantiated first. When deciding about the order of assigning values to a variable on turn, the situation is just the opposite: one would like to take a value which will not have to be reconsidered later on, due to backtracking. Hence, the "most promising" values should be tried first. "Most promising" in the sense that the partial instantiation resulting by assigning the value to the variable on turn, the partial instantiation can be extended further to a solution. A common measure to judge the "promise" of a value is based on the number of values not in conflict with the value in question, for each uninstantiated variable. The value which leaves the most options for each uninstantiated variable is the best. This heuristic is known as *minimum conflict first (MCF)* heuristic.

Both for variable and value ordering there is the question of breaking ties, that is, selecting one from more equally good candidates. This is done either in a deterministic way, based on lexicographic ordering or applying another principle to differentiate further between the candidates; or by choosing randomly from the equally good candidates. It has been observed that random choice is better uninformed deterministic choice.

Intelligent backtracking We illustrate two weaknesses of uninformed search by an example. Let us assume that the CSP to be solved has the variables x, y and z, each with the domain $\{1, 2, 3\}$. The constraints to be satisfied are: $y \leq z$, $xz \geq 4$. Let us assume that in the search process the $x = 1$, $y = 1$ instantiations have taken place. Now when trying to instantiate z, it turns out that for all the possible values for z the $xz \geq 4$ constraint is violated. Uninformed backtracking will reconsider the variable y in vain, ending up in the same dead-end situation, without noticing that the cause of the dead end is obviously in the value assigned to x, the very variable involved in the violated constraint. On the other hand, when after having backtracked to x, y is re-instantiated again to 1, and a value for z is to be chosen, the $y \leq z$, constraint will be checked again, though earlier it was found that if $y = 1$ then none of the variables from the domain of z were in conflict with the value of y, hence this constraint should not be checked again if $y = 1$. These shortcomings, due to chronological backtracking and blind constraint checking can be avoided by:

- backtracking to such a variable which can be the reason of the detected dead-end situation;
- when extending a partial solution after backtracking, some earlier proven partial instantiations which are not effected by recent changes are re-used instead of re-generating and testing the assignments again.

Depending on the policy used to pick the variable to backtrack to, and on the amount of learning taking place during the exploration of the solution space, different intelligent backtracking methods are defined (backjumping, conflict-directed backjumping, backchecking, backmarking).

Lookahead algorithms Let us assume that when searching for a solution, the variable x is given a value which excludes all the possible values for the variable y. In case of uninformed search this will only turn out when y will be considered to be instantiated.

Moreover, in case of chronological backtracking, thrashing will occur: the search tree will be expanded again and again till y, as long as the level of backtracking does not reach x. Both anomalies could be avoided by recognising that the chosen value for x cannot be part of a solution as there is no value for y which is compatible with it. Lookahead algorithms do this, by accepting a value for the variable on turn only if after having looked ahead, it could not be seen that the instantiation would lead to a dead-end. When checking this, problem reduction can also take place, by removing the values from the domain of the future variables which are not compatible with the current instantiation. The algorithms differ in how far they look ahead and how much reduction they perform.

Forward checking (FC) checks the satisfiability of the binary constraints, and removes the values which are not compatible with the current variable's instantiation.

The *full lookahead* or *arc-consistency lookahead* extends this with assuring arc-consistency of the remaining problem to be solved. That is, the domain of the uninstantiated variables is reduced such that the problem becomes arc-consistent. In contrast to its name, full lookahead does not perform complete possible lookahead, only assuring arc-consistency. In principle, there are lookahead algorithms possible which achieve higher consistency for the remaining problem. However, such algorithms are not used in practice due to the high time-complexity of the consistency-algorithms.

In-between the forward checking and full lookahead algorithms is *partial lookahead* or *directed arc-consistency lookahead*, which assure directed arc-consistency for the remaining problem, assuming a fixed ordering of the variables.

1.1.2 Non-crisp constraints

The constraints in classical constraint satisfaction problems are crisp, i.e., they either allow a tuple (of values of involved variables) or not. In some real-life applications, this is not a desired feature and, therefore, some alternative approaches to constraint satisfaction were proposed to enable non-crisp constraints, probabilities or weights respectively. These extensions can be used to solve optimization problems as well as over-constrained problems as they allow relaxing of constraints.

Partial Constraint Satisfaction Problems

The Partial Constraint Satisfaction (PCSP) [FW92] scheme of Freuder and Wallace is an interesting extension of CSP, which allows the relaxation and optimization of problems via the weakening of the original CSP.

A theory of Partial Constraint Satisfaction Problems (PCSPs) were developed to weaken systems of constraints which have no solutions (over-constrained systems), or for which finding a solution would take too long. Instead of searching for a solution to a complex problem, we consider searching for a simpler problem that we know we can solve.

Definition 1.1.8 (Partial Constraint Satisfaction Problems) *The notion of a partial CSP is formalized by considering three components:* $\langle (P, U)$, $(PS, \leq)$,

$(M, (N, S))\rangle$ *where:*

- $(PS, \leq)$ *is a problem space with* PS *a set of problems and* $\leq$ *a partial order over problems,*
- $P \in PS$ *is a constraint satisfaction problem (CSP),* U *is a set of 'universes' i.e. a set of potential values for each of the variables in* P*,*
- M *is a distance function over the problem space, and* (N, S) *are* N*ecessary and* S*ufficient bounds on the distance between the given problem* P *and some solvable member of the problem space* PS

A solution to a $PCSP$ *is a problem* P' *from the problem space and its solution, where the distance between* P *and* P' *is less than* N*. If the distance between* P *and* P' *is minimal, then this solution is optimal.*

Constraint Hierarchies

Constraint hierarchies [BFBW92] were introduced for describing over-constrained systems of constraints by specifying constraints with hierarchical strengths or preferences. It allows one to specify declaratively not only the constraints that are required to hold, but also weaker, so called soft constraints at an arbitrary but finite number of strengths. Weakening the strength of constraints helps in finding a solution of a previously over-constrained system of constraints. Intuitively, the hierarchy does not permit the weakest constraints to influence the result. Moreover, constraint hierarchies allow "relaxing" of constraints with the same strength by applying, for example a weighted-sum, least-squares or similar comparators.

DEFINITION 1.1.9 (CONSTRAINT HIERARCHIES) *In constraint hierarchies, one uses* labeled constraints *which are constraints labeled with a strength or preference. Then, a* constraint hierarchy *is a (finite) multiset of labeled constraints. Given a constraint hierarchy* H*, let* H_0 *denote the set of required constraints in* H*, with their labels removed. In the same way, we define the sets* $H_1, H_2, \ldots$ *for levels* $1, 2, \ldots$.

A solution to a constraint hierarchy H *will consist of instantiations for variables in* H*, that satisfies the constraints in* H *respecting the hierarchy. Formally (*U*,* V *are instantiations):*

$$S_{H,0} = \{U \mid \textit{for each } c \textit{ in } H_0\textit{, } c \textit{ is satisfied by } U\}$$
$$S_H = \{U \mid U \in S_{H,0} \textit{ and } \forall V \in S_{H,0}\textit{, } V \textit{ is not better than } U \textit{ respecting } H\}.$$

The set $S_{H,0}$ *contains instantiations satisfying required constraints in* H *and* S_H*, called a* solution set *for* H*, is a subset of* $S_{H,0}$ *containing all instantiations which are not worse than other instantiations in* $S_{H,0}$.

Fuzzy, Probabilistic and Valued CSPs

The Fuzzy [RHZ76, DFP93, Rut94], Probabilistic [FL93] and Valued CSPs [SFV95] use values to express fuzzyness, probability or several preference values in the problem. In particular, Fuzzy CSPs are a very significant extension of CSPs since they are able to model partial constraint satisfaction (PCSP [FW92]) and also prioritised constraints (Constraint hierarchies [BFBW92]), that is, constraints with different levels of importance .

We will give a detailed description of the fuzzy, probabilistic and VCSP framework in the next chapter, when will also show how to represent all this non-crisp extension in our framework.

1.2 Constraint Logic Programming

As far as we know, even though some languages for handling special kinds of constraints were developed in the past (like ThingLab [Bor81], Bertrand [Lel88], and ALICE ([Lau78]), no general linguistic support were given to constraint solving and constraint propagation until the development of the constraint logic programming (CLP) scheme. This scheme can be seen as a logic programming shell, supported by an underlying constraint system which may handle and solve any kind of constraints. For this reason the CLP scheme can be denoted by CLP(X), where X specifies the class of constraints we want to deal with.

In this section we describe the basic concepts and the operational semantics of the constraint logic programming framework. For a longer and more detailed treatment, which also includes other kinds of semantics, we refer to [JL86, JL87, JM94]. We assume the readers to be already familiar with the logic programming formalism, as defined in [Llo87]. The main idea in CLP is to extend logic programming, in such a way that

- substitutions are replaced by constraints,
- unification is replaced by a constraint solution, and
- all the semantic properties of logic programming (mainly, the existence of declarative, denotational, and procedural semantics and the equivalence of these three objects) still hold.

Figure 1.3 describes informally the CLP framework. As we can see, the logic programming shell and the underlying constraint system communicate by means of consistency checks, which are requested by the shell and performed by the constraint system.

In constraint logic programs the basic components of a problem are stated as *constraints*. The problem as a whole is then represented by putting the various constraints together by means of *rules*.

DEFINITION 1.2.1 (CONSTRAINT RULE, GOAL, PROGRAMS) *Consider a particular instance of the CLP(X) framework.*

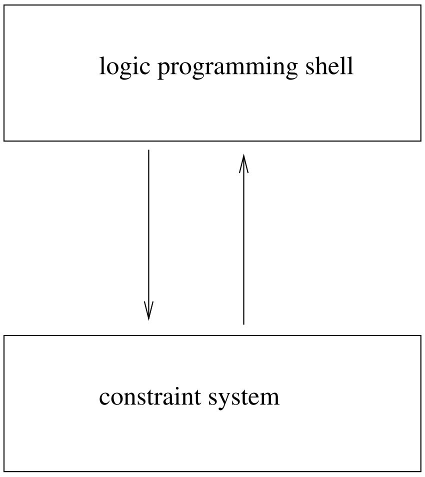

Figure 1.3: The CLP framework.

- *A constraint rule is of the form*

  $$A : -c, B_1, \ldots, B_n.$$

  where $n \geq 0$, c *is a constraint over the class* X, *and* A *and all* B_i *s are atoms over the Herbrand Universe.* A *is called the head of the rule, and* $c, B_1, \ldots, B_n$ *the body. If* $n = 0$ *then such a constraint rule can also be called a fact.*

- *A goal is a constraint rule without the head, i.e of the form :* $: -c.$ *or* $: -c, B_1, \ldots, B_n$.

- *A constraint logic program is a finite set of constraint rules.*

Let us now describe how the CLP system works to decide if a given goal is a logical consequence of the program; i.e., how the operational semantics of logic programming, which we assume is given by SLD-resolution, has to be extended in order to deal with constraints instead of (or, better, besides) substitutions.

DEFINITION 1.2.2 (DERIVATION STEP) *Given a CLP(*X*) program P, a derivation step takes a goal*

$$: -(c, A_1, \ldots, A_n)$$

and a (variant of a) constraint rule

$$A : -c', B_1, \ldots, B_m.$$

in P. If $(c \cup c') \cup (A_i = A)$ *is satisfiable, then the derivation step returns a new goal*

$$: -((c \cup c') \cup (A = A_i), A_1, \ldots, A_{i-1}, B_1, \ldots B_m, A_{i+1}, \ldots A_n)$$

DEFINITION 1.2.3 (DERIVATION, ANSWER CONSTRAINT) *A derivation is a sequence of derivation steps. Derivation sequences are* successful *(and are called refutations) when the last goal therein contains only constraints. These* answer constraints *constitute the output of a CLP program.*

The satisfiability of $A_i = A$ is the problem of finding, if there is one, an mgu, and thus, it is a task which can be accomplished by the logic programming engine that underlies the CLP system. On the contrary, to check the satisfiability of $(c \cup c')$ we need to use the particular constraint solver corresponding to the chosen structure X. If, for example, we are considering arithmetic inequality constraints over the reals, then the constraint solver could be a simplex algorithm. Thus, at each derivation step two constraint solvers are needed: one for the Herbrand universe and one for the chosen class of other constraints.

The fact that CLP can be interpreted as the possibility of using logic programming as a logic "shell" where to formally reason about constraints and solve constraint problems intuitively implies that logic programming is not only an instance of CLP, but it is always "present" in the CLP framework. More precisely, since a constraint is simply a relation which must hold among a collection of objects, then term equalities, which are handled in logic programming while performing unification, are a special case of constraints. Therefore, it is safe to say that any logic program is a CLP program where the only allowed constraints are term equalities. On the other hand, it is also true that any CLP instance needs to perform unifications, and thus term equalities must always be among the allowed constraints.

Chapter 2

Soft Constraint Satisfaction Problems

Three Rings for the Elven-kings under the sky,
Seven for the Dwarf-lords in their halls of stone,
Nine for the Mortal Men doomed to die,
One for the Dark Lord on his dark throne.
In the Land of Mordor where the Shadows lie.
One Ring to rule them all, One Ring to find them,
One Ring to bring them all and in the darkness bind them.
In the Land of Mordor where the Shadows lie.
Tolkien, J.R.R. The Lord of the Rings

Abstract

We introduce a general framework for constraint satisfaction and optimization where classical CSPs, fuzzy CSPs, weighted CSPs, partial constraint satisfaction, and others can be easily cast. The framework is based on a semiring structure, where the carrier set of the semiring specifies the values to be associated with each tuple of values of the variable domain, and the two semiring operations ($+$ and $\times$) model constraint projection and combination respectively.

Classical constraint satisfaction problems (CSPs) [Mon74, Mac92] are a very expressive and natural formalism to specify many kinds of real-life problems. In fact, problems ranging from map coloring, vision, robotics, job-shop scheduling, VLSI design, and so on, can easily be cast as CSPs and solved using one of the many techniques that have been developed for such problems or subclasses of them [Fre78, Fre88, MF85, Mac77, Mon74].

They also, however, have evident limitations, mainly due to the fact that they do not appear to be very flexible when trying to represent real-life scenarios where the knowledge is neither completely available nor crisp. In fact, in such situations, the ability of stating whether an instantiation of values to variables is allowed or not, is not enough or sometimes not even possible. For these reasons, it is natural to try to extend the CSP formalism in this direction.

For example, in [RHZ76, DFP93, Rut94] CSPs were extended with the ability to associate with each tuple, or to each constraint, a level of preference, and with the possibility of combining constraints using min-max operations. This extended formalism were called Fuzzy CSPs (FCSPs). Other extensions concern the ability to model incomplete knowledge of the real problem [FL93], to solve over-constrained problems [FW92], and to represent cost optimization problems [Bis94].

In this chapter we describe a constraint solving framework where all such extensions, as well as classical CSPs, can be cast. We do not, however, relax the assumption of a finite domain for the variables of the constraint problems. The main idea is based on the observation that a semiring (that is, a domain plus two operations satisfying certain properties) is all that is needed to describe many constraint satisfaction schemes. In fact, the domain of the semiring provides the levels of consistency (which can be interpreted as cost, or degrees of preference, or probabilities, or others), and the two operations define a way to combine constraints together. More precisely, we define the notion of constraint solving over any semiring. Specific choices of the semiring will then give rise to different instances of the framework, which may correspond to known or new constraint solving schemes.

The notion of semiring for constraint solving was used also in [GDL92]. However, the use of such a notion is completely different from ours. In fact, in [GDL92] the semiring domain (hereby denoted by A) is used to specify the domain of the variables, while here we always assume a finite domain (hereby denoted by D) and A is used to model the values associated with the tuples of values of D.

The chapter is organized as follows. Section 2.1 defines c-semirings and their properties. Then Section 2.2 introduces constraint problems over any semirings and the associated notions of solution and consistency. Then Section 2.3 studies several instances of the SCSP framework, and in particular studies the differences between the framework presented in [SFV95] and ours.

The basic definitions and properties of the soft constraints' framework described in this chapter appear in [BMR95, BMR97b]; moreover the mapping between SCSPs and VCSPs [SFV95] is based on the ideas developed in [BFM$^+$96, BFM$^+$99].

2.1 C-semirings and their properties

We extend the classical notion of constraint satisfaction to allow also for 1) non-crisp statements and 2) a more general interpretation of the operations on such statements. This allows us to model both classical CSPs and several extensions of them (like fuzzy CSPs [DFP93, Rut94], partial constraint satisfaction [FW92], and so on). To formally do that, we associate a semiring with the standard notion of constraint problem, so that different choices of the semiring represent different concrete constraint satisfaction schemes. In fact, such a semiring will give us both the domain for the non-crisp statements and also the allowed operations on them.

Definition 2.1.1 (semiring) *A semiring is a tuple* $\langle A, sum, \times, 0, 1\rangle$ *such that*

- A *is a set and* $0, 1 \in A$*;*
- sum*, called the additive operation, is a commutative (i.e.,* $\mathrm{sum}(a, b) = \mathrm{sum}(b, a)$*) and associative (i.e.,* $\mathrm{sum}(a, \mathrm{sum}(b, c)) = \mathrm{sum}(\mathrm{sum}(a, b), c)$*) operation with* 0 *as its unit element (i.e.,* $\mathrm{sum}(a, 0) = a = \mathrm{sum}(0, a)$*);*
- $\times$*, called the multiplicative operation, is an associative operation such that* 1 *is its unit element and* 0 *is its absorbing element (i.e.,* $a \times 0 = 0 = 0 \times a$*);*
- $\times$ *distributes over* sum *(i.e., for any* $a, b, c \in A$*,* $a \times \mathrm{sum}(b, c) = \mathrm{sum}((a \times b), (a \times c))$*).*

In the following we will consider semirings with additional properties of the two operations. Such semirings will be called *c-semiring*, where “c” stands for “constraint”, meaning that they are the natural structures to be used when handling constraints.

DEFINITION 2.1.2 (C-SEMIRING) *A c-semiring is a tuple* $\langle A, +, \times, 0, 1\rangle$ *such that*

- A *is a set and* $0, 1 \in A$*;*
- $+$ *is defined over (possibly infinite) sets of elements of* A *as follows*[1]*:*
  - *for all* $a \in A$*,* $\sum(\{a\}) = a$*;*
  - $\sum(\emptyset) = 0$ *and* $\sum(A) = 1$*;*
  - $\sum(\bigcup A_i, i \in I) = \sum(\{\sum(A_i), i \in I\})$ *for all sets of indices* I *(*flattening property*).*
- $\times$ *is a binary, associative and commutative operation such that* 1 *is its unit element and* 0 *is its absorbing element;*
- $\times$ *distributes over* $+$ *(i.e., for any* $a \in A$ *and* $B \subseteq A$*,* $a \times \sum(B) = \sum(\{a \times b, b \in B\})$*).*

Operation $+$ is defined over any set of elements of A, also over infinite sets. This will be useful later in proving Theorem 2.1.9. The fact that $+$ is defined over *sets* of elements, and not *pairs* or *tuples*, automatically makes such an operation commutative, associative, and idempotent. Moreover, it is possible to show that 0 is the unit element of $+$; in fact, by using the flattening property we get $\sum(\{a, 0\}) = \sum(\{a\} \cup \emptyset) = \sum(\{a\}) = a$. This means that a c-semiring is a semiring (where the sum operation is $+$) with some additional properties.

It is also possible to prove that 1 is the absorbing element of $+$. In fact, by flattening and by the fact that we set $\sum(A) = 1$, we get $\sum(\{a, 1\}) = \sum(\{a\} \cup A) = \sum(A) = 1$.

[1]When $+$ is applied to a two-element set, we will use the symbol $+$ in infix notation, while in general we will use the symbol $\sum$ in prefix notation.

Let us now consider the advantages of using c-semirings instead of semirings. First, the idempotency of the $+$ operation is needed in order to define a partial ordering $\leq_S$ over the set A, which will enable us to compare different elements of the semiring. Such partial order is defined as follows: $a \leq_S b$ iff $a + b = b$. Intuitively, $a \leq_S b$ means that b is "better" than a, or, from another point of view, that, between a and b, the $+$ operation chooses b. This ordering will be used later to choose the "best" solution in our constraint problems.

THEOREM 2.1.3 ($\leq_S$ IS A PARTIAL ORDER) *Given any c-semiring $S = \langle A, +, \times, 0, 1\rangle$, consider the relation $\leq_S$ over A such that $a \leq_S b$ iff $a + b = b$. Then $\leq_S$ is a partial order.*

PROOF. We have to prove that $\leq_S$ is reflexive, transitive, and antisymmetric:

- Since $+$ is idempotent, we have that $a + a = a$ for any $a \in A$. Thus, by definition of $\leq_S$, we have that $a \leq_S a$. Thus $\leq_S$ is reflexive.
- Assume $a \leq_S b$ and $b \leq_S c$. This means that $a + b = b$ and $b + c = c$. Thus $c = b + c = (a + b) + c = a + (b + c) = a + c$. Note that here we also used the associativity of $+$. Thus $\leq_S$ is transitive.
- Assume that $a \leq_S b$ and $b \leq_S a$. This means that $a + b = b$ and $b + a = a$. Thus $a = b + a = a + b = b$. Thus $\leq_S$ is antisymmetric.

$\square$

The fact that 0 is the unit element of the additive operation implies that 0 is the minimum element of the ordering. Thus, for any $a \in A$, we have $0 \leq_S a$.

It is important to note that both the additive and the multiplicative operations are monotone on such an ordering.

THEOREM 2.1.4 ($+$ AND $\times$ ARE MONOTONE OVER $\leq_S$) *Given any c-semiring $S = \langle A, +, \times, 0, 1\rangle$, consider the relation $\leq_S$ over A. Then $+$ and $\times$ are monotone over $\leq_S$. That is, $a \leq_S a'$ implies $a + b \leq_S a' + b$ and $a \times b \leq_S a' \times b$.*

PROOF. Assume $a \leq_S a'$. Then, by definition of $\leq_S$, $a + a' = a'$.

Thus, for any b, $a' + b = a + a' + b$. By idempotency of $+$, we also have $a' + b = a + a' + b = a + a' + b + b$, which, by commutativity of $+$, becomes $a' + b = (a + b) + (a' + b)$. By definition of $\leq_S$, we have $a + b \leq_S a' + b$.

Also, from $a + a' = a'$ derives that, for any b, $a' \times b = (a' + a) \times b$ =(by distributiveness) $(a' \times b) + (a \times b)$. This means that $(a \times b) \leq_S (a' \times b)$. $\square$

The commutativity of the $\times$ operation is desirable when such an operation is used to combine several constraints. In fact, were it not commutative, it would mean that different orders of the constraints would give different results.

Since 1 is also the absorbing element of the additive operation, then $a \leq_S 1$ for all a. Thus 1 is the maximum element of the partial ordering. This implies that the $\times$ operation is *intensive*, that is, that $a \times b \leq_S a$. This is important since it means that combining more constraints leads to a worse (w.r.t. the $\leq_S$ ordering) result.

THEOREM 2.1.5 ($\times$ IS INTENSIVE) *Given any c-semiring $S = \langle A, +, \times, 0, 1\rangle$, consider the relation $\leq_S$ over A. Then $\times$ is intensive, that is, $a, b \in A$ implies $a \times b \leq_S a$.*

PROOF. Since 1 is the unit element of $\times$, we have $a = a \times 1$. Also, since 1 is the absorbing element of $+$, we have $1 = 1 + b$. Thus $a = a \times (1 + b)$. Now, $a \times (1 + b)$ = { by distributiveness of $\times$ over $+$ } $(a \times 1) + (a \times b) = \{1$ unit element of $\times\}$ $a + (a \times b)$. Thus we have $a = a + (a \times b)$, which, by definition of $\leq_S$, means that $(a \times b) \leq_S a$. □

In the following we will sometimes need the $\times$ operation to be closed on a certain finite subset of the c-semiring.

DEFINITION 2.1.6 (AD-CLOSED) *Given any c-semiring $S = \langle A, +, \times, 0, 1\rangle$, consider a finite set $AD \subseteq A$. Then $\times$ is AD-closed if, for any $a, b \in AD$, $(a \times b) \in AD$.*

We will now show that c-semirings can be assimilated to complete lattices. Moreover, we will also sometimes need to consider c-semirings where $\times$ is idempotent, which we will show equivalent to distributive lattices. See [DP90] for a deep treatment of lattices.

DEFINITION 2.1.7 (LUB, GLB, (COMPLETE) LATTICE [DP90]) *Consider a partially ordered set S and any subset I of S. Then we define the following:*

- *an upper bound (resp., lower bound) of I is any element x such that, for all $y \in I$, $y \leq x$ (resp., $x \leq y$);*
- *the least upper bound (lub) (resp., greatest lower bound (glb)) of I is an upper bound (resp., lower bound) x of I such that, for any other upper bound (resp., lower bound) x' of I, we have that $x \leq x'$ (resp., $x' \leq x$).*

A lattice is a partially ordered set where every subset of two elements has a lub and a glb. A complete lattice is a partially ordered set where every subset has a lub and a glb.

We will now prove a property of partially ordered sets where every subset has the lub, which will be useful in proving that $\langle A, \leq_S\rangle$ is a complete lattice. Notice that when every subset has the lub, then also the empty set has the lub. Thus, in partial orders with this property, there is always a global minimum of the partial order (which is the lub of the empty set).

LEMMA 2.1.8 (LUB $\Rightarrow$ GLB) *Consider a partial order $\langle A, \leq\rangle$ where there is the lub of every subset I of A. Then there exists the glb of I as well.*

PROOF. Consider any set $I \subseteq A$, and let us call $LB(I)$ the set of all lower bounds of S. That is, $LB(I) = \{x \in A \mid \text{for all } y \in I, x \leq_S y\}$. Then consider $a = lub(LB(I))$. We will prove that a is the glb of I.

Consider any element $y \in I$. By definition of $LB(I)$, we have that all elements x in $LB(I)$ are smaller than y, thus y is an upper bound of $LB(I)$. Since a is by definition the smallest among the upper bounds of $LB(I)$, we also have that $a \leq_S y$. This is true for all elements y in I. Thus a is a lower bound of I, which means that it must be contained in $LB(I)$. Thus we have found an element of A which is the greatest among all lower bounds of I. □

THEOREM 2.1.9 ($\langle A, \leq_S \rangle$ IS A COMPLETE LATTICE) *Given a c-semiring $S = \langle A, +, \times, 0, 1 \rangle$, consider the partial order $\leq_S$. Then $\langle A, \leq_S \rangle$ is a complete lattice.*

PROOF. To prove that $\langle A, \leq_S \rangle$ is a complete lattice it is enough to show that every subset of A has the lub. In fact, by Lemma 2.1.8 we would get that each subset of A has both the lub and the glb, which is exactly the definition of a complete lattice.

We already know by Theorem 2.1.3 that $\langle A, \leq_S \rangle$ is a partial order. Now, consider any set $I \subseteq A$, and let us set $m = \sum(I)$ and $n = lub(I)$. Given any element $x \in I$, we have that $x + m$ =(by flattening) $\sum(\{x\} \cup I) = \sum(I) = m$. Therefore, by definition of $\leq_S$, we have that $x \leq_S m$. Thus also $n \leq_S m$, since m is an upper bound of I and n by definition is the least among the upper bounds.

On the other hand, we have that
$m + n =$ (by definition of sum)
$\sum(\{m\} \cup \{n\}) =$ (by flattening and since $m = \sum(I)$)
$\sum(I \cup \{n\}) =$ (by giving an alternative definition of the set $I \cup \{n\}$)
$\sum(\bigcup_{x \in I}(\{x\} \cup \{n\})) =$ (by flattening)
$\sum(\{\sum(\{x\} \cup \{n\}), x \in I\}) =$ (since $x \leq_S n$ and thus $\sum(\{x\} \cup \{n\}) = n$)
$\sum(\{n\}) = n$.
Thus we have proven that $m \leq_S n$, which, together with the previous result (that $n \leq_S m$) yields $m = n$. In other words, we proved that $\sum(I) = lub(I)$ for any set $I \subseteq A$. Thus every subset of A, say I, has a least upper bound (which coincides with $\sum(I)$). Thus $\langle A, \leq_S \rangle$ is a lub-complete partial order. □

Note that the proof of the previous theorem also says that the $+$ operation coincides with the lub of the lattice $\langle A, \leq_S \rangle$.

THEOREM 2.1.10 ($\times$ IDEMPOTENT) *Given a c-semiring $S = \langle A, +, \times, 0, 1 \rangle$, consider the corresponding complete lattice $\langle A, \leq_S \rangle$. If $\times$ is idempotent, then we have that:*

1. *$+$ distributes over $\times$;*

2. *$\times$ coincides with the glb operation of the lattice;*

3. *$\langle A, \leq_S \rangle$ is a distributive lattice.*

PROOF.

1. $(a + b) \times (a + c) =$ {since $\times$ distributes over $+$}
   $((a + b) \times a) + ((a + b) \times c)) =$ {same as above}
   $((a \times a) + (a \times b)) + ((a + b) \times c)) =$ {by idempotency of $\times$}
   $(a + (a \times b)) + ((a + b) \times c)) =$ {by intensivity of $\times$ and definition of $\leq_S$}
   $a + ((a + b) \times c)) =$ {since $\times$ distributes over $+$}
   $a + ((c \times a) + (c \times b)) =$ {by intensivity of $\times$ and definition of $\leq_S$}
   $a + (c \times b)$.

2. Assume that $a \times b = c$. Then, by intensivity of $\times$ (see Theorem 2.1.5), we have that $c \leq_S a$ and $c \leq_S b$. Thus c is a lower bound for a and b. To show that it is a glb, we need to show that there is no other lower bound c' such that $c \leq_s c'$. Assume that such c' exists. We now prove that it must be $c' = c$:
   $c' = \{\text{since } c \leq_S c'\}$
   $c' + c = \{\text{since } c = a \times b\}$
   $c' + (a \times b) = \{\text{since } + \text{ distributes over } \times\text{, see previous point}\}$
   $(c' + a) \times (c' + b) = \{\text{since we assumed } c' \leq_S a \text{ and } c' \leq_S b\text{, and by definition of } \times\}$
   $a \times b = \{\text{by assumption}\}$
   c.

3. This comes from the fact that $+$ is the lub, $\times$ is the glb, and $\times$ distributes over $+$ by definition of semiring (the distributiveness in the other direction is given by Lemma 6.3 in [DP90], or can be seen in the first point above).

□

Note that, in the particular case in which $\times$ is idempotent and $\leq_S$ is total, we have that $a + b = max(a, b)$ and $a \times b = min(a, b)$.

2.2 Constraint systems and problems

We will now define the notion of constraint system, constraint, and constraint problem, which will be parametric w.r.t. the notion of c-semiring just defined. Intuitively, a constraint system specifies the c-semiring $\langle A, +, \times, 0, 1\rangle$ to be used, the set of all variables and their domain D.

DEFINITION 2.2.1 (CONSTRAINT SYSTEM) *A constraint system is defined as a tuple* $CS = \langle S, D, V\rangle$*, where* S *is a c-semiring,* D *is a finite set, and* V *is an ordered set of variables.*

Now, a constraint over a given constraint system specifies the involved variables and the "allowed" values for them. More precisely, for each tuple of values (of D) for the involved variables, a corresponding element of A is given. This element can be interpreted as the tuple's weight, or cost, or level of confidence, or other.

DEFINITION 2.2.2 (CONSTRAINT) *Given a constraint system* $CS = \langle S, D, V\rangle$*, where* $S = \langle A, +, \times, 0, 1\rangle$*, a constraint over* CS *is a pair* $\langle def, con\rangle$*, where*

- $con \subseteq V$*, it is called the* type *of the constraint;*

- $def : D^k \rightarrow A$ *(where* k *is the cardinality of* con*) is called the* value *of the constraint.*

A constraint problem is then just a set of constraints over a given constraint system, plus a selected set of variables (thus a *type*). These are the variables of interest in the problem, i.e., the variables of which we want to know the possible assignments' compatibly with all the constraints.

DEFINITION 2.2.3 (CONSTRAINT PROBLEM) *Consider any constraint system* $\mathrm{CS} = \langle S, D, V\rangle$. *A constraint problem* P *over* CS *is a pair* $P = \langle C, \mathtt{con}\rangle$, *where* C *is a set of constraints over* CS *and* $\mathtt{con} \subseteq V$. *We also assume that* $\langle \mathtt{def}_1, \mathtt{con}'\rangle \in C$ *and* $\langle \mathtt{def}_2, \mathtt{con}'\rangle \in C$ *implies* $\mathtt{def}_1 = \mathtt{def}_2$. *In the following we will write SCSP to refer to such constraint problems.*

Note that the above condition (that there are no two constraints with the same type) is not restrictive. In fact, if in a problem we had two constraints $\langle \mathtt{def}_1, \mathtt{con}'\rangle$ and $\langle \mathtt{def}_2, \mathtt{con}'\rangle$, we could replace both of them with a single constraint $\langle \mathtt{def}, \mathtt{con}'\rangle$ with $\mathtt{def}(t) = \mathtt{def}_1(t) \times \mathtt{def}_2(t)$ for any tuple t. Similarly, if C were a multiset, e.g. if a constraint $\langle \mathtt{def}, \mathtt{con}'\rangle$ had n occurrences in C, we could replace all its occurrences with the single constraint $\langle \mathtt{def}', \mathtt{con}'\rangle$ with

$$\mathtt{def}'(t) = \underbrace{\mathtt{def}(t) \times \ldots \times \mathtt{def}(t)}_{n \text{ times}}$$

for any tuple t. This assumption, however, implies that the operation of union of constraint problems is not just set union, since it has to take into account the possible presence of constraints with the same type in the problems to be combined (at most one in each problem), and, in that case, it has to perform the just-described constraint replacement operations.

When all variables are of interest, like in many approaches to classical CSP, $\mathtt{con}$ contains all the variables involved in any of the constraints of the given problem, say P. Such a set, called $V(P)$, is a subset of V that can be recovered by looking at the variables involved in each constraint. That is, $V(P) = \bigcup_{\langle \mathtt{def}, \mathtt{con}'\rangle \in C} \mathtt{con}'$.

As for classical constraint solving, SCSPs as defined above also can be graphically represented via labeled hypergraphs where nodes are variables, hyperarcs are constraints, and each hyperarc label is the definition of the corresponding constraint (which can be seen as a set of pairs $\langle$ tuple, value $\rangle$). The variables of interest can then be marked in the graph.

Note that the above definition is parametric w.r.t. the constraint system CS and thus w.r.t. the semiring S. In the following we will present several instantiations of such a framework, and we will show them to coincide with known and also new constraint satisfaction systems.

In the SCSP framework, the values specified for the tuples of each constraint are used to compute corresponding values for the tuples of values of the variables in $\mathtt{con}$, according to the semiring operations: the multiplicative operation is used to combine the values of the tuples of each constraint to get the value of a tuple for all the variables, and the additive operation is used to obtain the value of the tuples of the variables in the type of the problem. More precisely, we can define the operations of *combination* ($\otimes$) and *projection* ($\Downarrow$) over constraints. Analogous operations were originally defined for fuzzy relations in [Zad65], and used for fuzzy CSPs in [DFP93]. Our definition is, however, more general since we do not consider a specific c-semiring (like that which we will define for fuzzy CSPs later) but a general one.

To define the operators of combination ($\otimes$) and projection ($\Downarrow$) over constraints we need the definition of tuple projection (which is given here in a more general way w.r.t. that given in Definition 1.1.3 where it was used only for the definition of CSP solutions).

DEFINITION 2.2.4 (TUPLE PROJECTION) *Given a constraint system* $CS = \langle S, D, V \rangle$ *where* V *is totally ordered via ordering* $\prec$*, consider any* k*-tuple*[2] $t = \langle t_1, \ldots, t_k \rangle$ *of values of* D *and two sets* $W = \{w_1, \ldots, w_k\}$ *and* $W' = \{w'_1, \ldots, w'_m\}$ *such that* $W' \subseteq W \subseteq V$ *and* $w_i \prec w_j$ *if* $i \leq j$ *and* $w'_i \prec w'_j$ *if* $i \leq j$*. Then the projection of* t *from* W *to* W'*, written* $t \downarrow^{W}_{W'}$*, is defined as the tuple* $t' = \langle t'_1, \ldots, t'_m \rangle$ *with* $t'_i = t_j$ *if* $w'_i = w_j$.

DEFINITION 2.2.5 (COMBINATION) *Given a constraint system* $CS = \langle S, D, V \rangle$*, where* $S = \langle A, +, \times, 0, 1 \rangle$*, and two constraints* $c_1 = \langle def_1, con_1 \rangle$ *and* $c_2 = \langle def_2, con_2 \rangle$ *over* CS*, their combination, written* $c_1 \otimes c_2$*, is the constraint* $c = \langle def, con \rangle$ *with*

$$con = con_1 \cup con_2$$

and

$$def(t) = def_1(t \downarrow^{con}_{con_1}) \times def_2(t \downarrow^{con}_{con_2}).$$

Since $\times$ is both commutative and associative, so too is $\otimes$. Thus this operation can be easily extended to more than two arguments, say $C = \{c_1, \ldots, c_n\}$, by performing $c_1 \otimes c_2 \otimes \ldots \otimes c_n$, which we will sometimes denote by $\bigotimes C$.

DEFINITION 2.2.6 (PROJECTION) *Given a constraint system* $CS = \langle S, D, V \rangle$*, where* $S = \langle A, +, \times, 0, 1 \rangle$*, a constraint* $c = \langle def, con \rangle$ *over* CS*, and a set* I *of variables (*$I \subseteq V$*), the projection of* c *over* I*, written* $c \Downarrow_I$*, is the constraint* $\langle def', con' \rangle$ *over* CS *with*

$$con' = I \cap con$$

and

$$def'(t') = \Sigma_{\{t \mid t \downarrow^{con}_{I \cap con} = t'\}} def(t).$$

PROPOSITION 2.2.7 ($c \Downarrow_I = c \Downarrow_{I \cap con}$) *Given a constraint* $c = \langle def, con \rangle$ *over a constraint system* CS *and a set* I *of variables (*$I \subseteq V$*), we have that* $c \Downarrow_I = c \Downarrow_{I \cap con}$.

[2]Given any integer k, a k-tuple is just a tuple of length k. Also, given a set S, an S-tuple is a tuple with as many elements as the size of S.

PROOF. Follows trivially from Definition 2.2.6. □

A useful property of the projection operation is the following.

THEOREM 2.2.8 ($c \Downarrow_{S_1} \Downarrow_{S_2} = c \Downarrow_{S_2}$) *Given a constraint* c *over a constraint system* CS, *we have that* $c \Downarrow_{S_1} \Downarrow_{S_2} = c \Downarrow_{S_2}$ *if* $S_2 \subseteq S_1$.

PROOF. To prove the theorem, we have to show that the two constraints $c_1 = c \Downarrow_{S_1} \Downarrow_{S_2}$ and $c_2 = c \Downarrow_{S_2}$ coincide, that is, they have the same con and the same def. Assume $c = \langle def, con \rangle$, $c_1 = \langle def_1, con_1 \rangle$, and $c_2 = \langle def_2, con_2 \rangle$. Now, $con_1 = S_2 \cap (S_1 \cap con)$. Since $S_2 \subseteq S_1$, we have that $con_1 = S_2 \cap con$. Also, $con_2 = S_2 \cap con$, thus $con_1 = con_2$. Consider now def_1. By Definition 2.2.6, we have that $def_1(t_1) = \Sigma_{\{t' | t' \downarrow^{S_1}_{S_2} = t_1\}} \Sigma_{\{t | t \downarrow^{con}_{S_1} = t'\}} def(t)$, which, by associativity of $+$, is the same as $\Sigma_{\{t | t \downarrow^{con}_{S_2} = t_1\}} def(t)$, which coincides with $def_2(t_1)$ by Definition 2.2.6. □

We will now prove a property which can be seen as a generalization of the distributivity property in predicate logic, which we recall is $\exists x.(p \wedge q) = (\exists x.p) \wedge q$ if x not free in q (where p and q are two predicates). The extension we prove for our framework is given by the fact that $\otimes$ can be seen as a generalized $\wedge$ and $\Downarrow$ as a variant[3] of $\exists$. This property will be useful later in Section 3.6.

THEOREM 2.2.9 *Given two constraints* $c_1 = \langle def_1, con_1 \rangle$ *and* $c_2 = \langle def_2, con_2 \rangle$ *over a constraint system* CS, *we have that* $(c_1 \otimes c_2) \Downarrow_{(con_1 \cup con_2) - x} = c_1 \Downarrow_{con_1 - x} \otimes c_2$ *if* $x \cap con_2 = \emptyset$.

PROOF. Let us set $c = (c_1 \otimes c_2) \Downarrow_{(con_1 \cup con_2) - x} = \langle def, con \rangle$ and $c' = c_1 \Downarrow_{con_1 - x} \otimes c_2 = \langle def', con' \rangle$. We will first prove that $con = con'$: $con = (con_1 \cup con_2) - x$ and $con' = (con_1 - x) \cup con_2$, which is the same as con if $x \cap con_2 = \emptyset$, as we assumed. Now we prove that $def = def'$. By definition,

$$def(t) = \Sigma_{\{t' | t' \downarrow^{con_1 \cup con_2}_{con_1 \cup con_2 - x} = t\}} (def_1(t' \downarrow^{con_1 \cup con_2}_{con_1}) \times def_2(t' \downarrow^{con_1 \cup con_2}_{con_2})).$$

Now, $def_2(t' \downarrow^{con_1 \cup con_2}_{con_2})$ is independent from the summation, since x is not involved in c_2. Thus it can be taken out. Also, $t' \downarrow^{con_1 \cup con_2}_{con_2}$ can be substituted by $t \downarrow^{con_1 \cup con_2}_{con_2}$, since t' and t must coincide on the variables different from x. Thus we have:

$$(\Sigma_{\{t' | t' \downarrow^{con_1 \cup con_2}_{con_1 \cup con_2 - x} = t\}} def_1(t' \downarrow^{con_1 \cup con_2}_{con_1})) \times def_2(t \downarrow^{con_1 \cup con_2 - x}_{con_2}).$$

Now, the summation is done over those tuples t' that involve all the variables and coincide with t on the variables different from x. By observing that the elements of the summation are given by $def_1(t' \downarrow^{con_1 \cup con_2}_{con_1})$, and thus they only contain variables in con_1, we can conclude that the result of the summation does not change if it is done over the tuples involving only the variables in con_1 and still coinciding with t over the variables in con_1 that are different from x. Thus we get:

$$(\Sigma_{\{t_1 | t_1 \downarrow^{con_1}_{con_1 - x} = t \downarrow^{con_1 \cup con_2 - x}_{con_1 - x}\}} def_1(t_1)) \times def_2(t \downarrow^{con_1 \cup con_2 - x}_{con_2}).$$

It is easy to see that this formula is exactly the definition of def'. □

[3] Note, however, that the operator corresponding to $\exists x$ is given by $\Downarrow^{W}_{W - \{x\}}$.

Using the properties of $\times$ and $+$, it is easy to prove that: $\otimes$ is associative and commutative; $\otimes$ is monotone over $\sqsubseteq_S$. Moreover, if $\times$ is idempotent $\otimes$ is idempotent.

Using the operations of combination and projection, we can now define the notion of solution of an SCSP. In the following we will consider a fixed constraint system $CS = \langle S, D, V \rangle$, where $S = \langle A, +, \times, 0, 1 \rangle$.

DEFINITION 2.2.10 (SOLUTION) *Given a constraint problem* $P = \langle C, con \rangle$ *over a constraint system* CS*, the solution of* P *is a constraint defined as* $Sol(P) = (\bigotimes C) \Downarrow_{con}$.

In words, the solution of an SCSP is the constraint induced on the variables in con by the whole problem. Such constraint provides, for each tuple of values of D for the variables in con, a corresponding value of A. Sometimes, it is enough to know just the best value associated with such tuples. In our framework, this is still a constraint (over an empty set of variables), and will be called the best level of consistency of the whole problem, where the meaning of "best" depends on the ordering $\leq_S$ defined by the additive operation.

DEFINITION 2.2.11 (BEST LEVEL OF CONSISTENCY) *Given an SCSP* $P = \langle C, con \rangle$*, we define* $blevel(P) \in S$ *such that* $\langle blevel(P), \emptyset \rangle = (\bigotimes C) \Downarrow_{\emptyset}$*. Moreover, we say that* P *is consistent if* $0 <_S blevel(P)$*, and that* P *is* α*-consistent if* $blevel(P) = \alpha$.

Informally, the best level of consistency gives us an idea of how much we can satisfy the constraints of the given problem. Note that $blevel(P)$ does not depend on the choice of the distinguished variables, due to the associative property of the additive operation. Thus, since a constraint problem is just a set of constraints plus a set of distinguished variables, we can also apply function $blevel$ to a set of constraints only. More precisely, $blevel(C)$ will mean $blevel(\langle C, con \rangle)$ for any con.

Note that $blevel(P)$ can also be obtained by first computing the solution and then projecting such a constraint over the empty set of variables, as the following proposition shows.

PROPOSITION 2.2.12 *Given an SCSP* P*, we have that* $Sol(P) \Downarrow_{\emptyset} = \langle blevel(P), \emptyset \rangle$.

PROOF. $Sol(P) \Downarrow_{\emptyset}$ coincides with $((\bigotimes C) \Downarrow_{con}) \Downarrow_{\emptyset}$ by definition of $Sol(P)$. This coincides with $(\bigotimes C) \Downarrow_{\emptyset}$ by Theorem 2.2.8, thus it is the same as $\langle blevel(P), \emptyset \rangle$ by Definition 2.2.11. $\square$

Another interesting notion of solution, more abstract than the one defined above, but sufficient for many purposes, is the one that does not consider those tuples whose associated value is worse (w.r.t. $\leq_S$) than that of other tuples.

DEFINITION 2.2.13 (ABSTRACT SOLUTION) *Given an SCSP* $P = \langle C, con \rangle$*, consider* $Sol(P) = \langle def, con \rangle$*. Then the abstract solution of* P *is the set* $ASol(P) = \{\langle t, v \rangle \mid def(t) = v$ *and there is no* t' *such that* $v \leq_S def(t')\}$.

Note that, when the $\leq_S$ ordering is a total order (that is, when, for any a and b, $a+b$ is either a or b), then the abstract solution contains only those tuples whose associated value coincides with $\mathit{blevel}(P)$. In general, instead, an incomparable set of tuples is obtained, and thus $\mathit{blevel}(P)$ is just an upper bound of the values associated with the tuples.

By using the ordering $\leq_S$ over the semiring, we can also define a corresponding ordering on constraints with the same type.

DEFINITION 2.2.14 (CONSTRAINT ORDERING) *Consider two constraints c_1, c_2 over CS, and assume that $con_1 = con_2$ and $|con_1| = k$. Then $c_1 \sqsubseteq_S c_2$ if and only if, for all k-tuples t of values from D, $def_1(t) \leq_S def_2(t)$.*

Notice that, if $c_1 \sqsubseteq_S c_2$ and $c_2 \sqsubseteq_S c_1$, then $c_1 = c_2$.

THEOREM 2.2.15 ($\sqsubseteq_S$ IS A PARTIAL ORDER) *The relation $\sqsubseteq_S$ over the set of constraints over CS is a partial order.*

PROOF. By definition, such a relation is antisymmetric. Also, it is easy to see that it is reflexive and transitive. □

The notion of constraint ordering, and the fact that the solution is a constraint, can also be useful to define an ordering on problems.

DEFINITION 2.2.16 (PROBLEM ORDERING AND EQUIVALENCE) *Consider two SCSPs $P_1 = \langle C_1, con\rangle$ and $P_2 = \langle C_2, con\rangle$ over CS. Then $P_1 \sqsubseteq_P P_2$ if $Sol(P_1) \sqsubseteq_S Sol(P_2)$. If $P_1 \sqsubseteq_P P_2$ and $P_2 \sqsubseteq_P P_1$, then they have the same solution. Thus we say that P_1 and P_2 are equivalent and we write $P_1 \equiv P_2$.*

THEOREM 2.2.17 ($\sqsubseteq_P$ IS A PREORDER AND $\equiv$ IS AN EQUIVALENCE) *The relation $\sqsubseteq_P$ over the set of constraint problems over CS is a preorder. Moreover, $\equiv$ is an equivalence relation.*

PROOF. It is trivial to see that $\sqsubseteq_P$ is reflexive and transitive, due to the definition of constraint ordering $\sqsubseteq_S$. Thus $\sqsubseteq_P$ is a preorder. From this it derives that $\equiv$ is reflexive and transitive as well. Moreover, it is trivially symmetric. □

The ordering $\sqsubseteq_P$ can also be used for ordering sets of constraints, since a set of constraint is just a problem where con contains all the variables.

It is interesting now to note that, as in the classical CSP case, also the SCSP framework is monotone. That is, if some constraints are added, the solution of the new problem precedes that of the old one in the ordering $\sqsubseteq_S$. In other words, the new problem precedes the old one w.r.t. the preorder $\sqsubseteq_P$.

THEOREM 2.2.18 (MONOTONICITY) *Consider two SCSPs $P_1 = \langle C_1, con\rangle$ and $P_2 = \langle C_1 \cup C_2, con\rangle$ over CS. Then $P_2 \sqsubseteq_P P_1$ and $\mathit{blevel}(P_2) \leq_S \mathit{blevel}(P_1)$.*

PROOF. $P_2 \sqsubseteq_P P_1$ follows from the intensivity of $\times$ and the monotonicity of $+$. The same holds also for $blevel(P_2) \leq_S blevel(P_1)$, since both Sol and $blevel(P)$ are obtained by projecting over a subset of the variables (which is always empty in the case of $blevel(P)$). $\square$

When one is interested in the abstract solution rather than in the solution of an SCSP, then the above notion of monotonicity is lost. In fact, it is possible that by adding some constraints the best level of consistency gets worse, while the set of tuples in the abstract solution grows. See for example Figure 2.1, where there is a problem P_1 with one constraint and both variables as type, and a problem P_2 with the same constraint plus another one, and the two leftmost variables as type. Assume also that $\times$ is min and that $+$ is max. As we will see later, this amounts to considering the fuzzy constraint satisfaction framework (FCSP) (with the set of reals as carrier set instead of the set $[0, 1]$. Now, it is easy to see that $blevel(P_1) = 3$ and $blevel(P_2) = 2$. The abstract solution of P_1 contains the tuples $\langle a, a\rangle$ and $\langle b, a\rangle$, while that of P_2 contains also $\langle a, b\rangle$ and $\langle b, b\rangle$. Thus, adding one constraint in this case makes the best level of consistency worse while enlarging the set of tuples in the abstract solution.

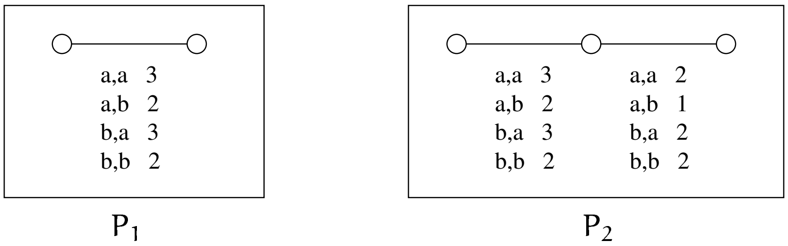

Figure 2.1: Two SCSPs.

Another notion of monotonicity holds, however, but only when the best level is not changed by the addition of some constraints and the order $\leq_S$ is total.

THEOREM 2.2.19 (SUBSET-MONOTONICITY VS. ASOL) *Consider two SCSPs* $P_1 = \langle C_1, con\rangle$ *and* $P_2 = \langle C_1 \cup C_2, con\rangle$ *over* CS, *and assume that* $\leq_S$ *is total and that* $blevel(P_1) = blevel(P_2)$. *Then* $Asol(P_2) \subseteq Asol(P_1)$.

PROOF. If $\leq_S$ is total, then the abstract solution is a set of tuples with associated value exactly the best level. Take a tuple t in $ASol(P_2)$. This means that the value associated with t in P_2 is $blevel(P_2)$. Then, by Theorem 2.2.18, the value associated with t in P_1, say v, is such that $blevel(P_2) \leq_S v$. By assumption, $blevel(P_1) = blevel(P_2)$. If $v \neq blevel(P_2)$, then this assumption is violated. Thus it must be $v = blevel(P_2) = blevel(P_1)$. Thus t is also in $Asol(P_1)$. Therefore $Asol(P_2) \subseteq Asol(P_1)$. $\square$

2.3 Instances of the framework

We will now show how several known, and also new, frameworks for constraint solving may be seen as instances of the SCSP framework. More precisely, each of such frameworks corresponds to the choice of a specific constraint system (and thus of a semiring).

This means that we can immediately know whether one can inherit the properties of the general framework by just looking at the properties of the operations of the chosen semiring, and by referring to the theorems in the previous sections. This is interesting for known constraint solving schemes, because it puts them into a single unifying framework and it justifies in a formal way many informally taken choices, but it is especially significant for new schemes, for which one does not need to prove all the properties that it has (or not) from scratch.

Note that the constraint systems of different instances differ only in the choice of the semiring. Therefore, in the following we will only specify the semiring that has to be chosen to obtain a particular instance of the SCSP framework.

2.3.1 Classical CSPs

As described in Section 1.1 a classical CSP problem [Mon74, Mac92] is just a set of variables and constraints, where each constraint specifies the tuples that are allowed for the involved variables. Assuming the presence of a subset of distinguished variables, the solution of a CSP consists of a set of tuples that represent the assignments of the distinguished variables, which can be extended to total assignments (for all the values) while satisfying all the constraints.

Since constraints in CSPs are crisp, that is, they either allow a tuple or not, we can model them via a semiring domain with only two values, say 1 and 0: allowed tuples will have associated the value 1, and not-allowed ones the value 0. Moreover, in CSPs, constraint combination is achieved via a join operation among allowed tuple sets. This can be modeled here by assuming the multiplicative operation to be the logical *and* (and interpreting 1 as true and 0 as false). Finally, to model the projection over some of the variables (for example the distinguished ones), as the k-tuples (assuming we project over k variables) for which there exists a consistent extension to an n-tuple, it is enough to assume the additive operation to be the logical *or*. Therefore, a CSP is just an SCSP where the semiring is

$$S_{CSP} = \langle \{0, 1\}, \vee, \wedge, 0, 1 \rangle.$$

THEOREM 2.3.1 $S_{CSP} = \langle \{0, 1\}, \vee, \wedge, 0, 1 \rangle$ *is a c-semiring.*

PROOF. It is easy to see that it is a semiring. Thus, we will only prove the additional properties needed in a c-semiring:

- the additive operation must be idempotent; it is trivially true since the additive operation is $\vee$, and $a \vee a = a$ for any $a \in \{0, 1\}$;
- the multiplicative operation must be commutative; trivially from the definition of $\wedge$;
- 1 is the absorbing element of the additive operation, that is $1 \vee a = 1$, which is trivially true.

The ordering $\leq_S$ here reduces to $0 \leq_S 1$. $\square$

EXAMPLE 2.3.2 Consider the CSP P in Figure 2.2a), where the distinguished variables are marked and the constraints are arcs labeled by the allowed tuples. The solution of such a CSP is the singleton set containing tuple $\langle a, a \rangle$ (which means $x = a$ and $y = a$). This CSP can be seen as an SCSP P' over the semiring S_{CSP} in Figure 2.2b). It can then be shown, by using the semiring operations, that the solution of P and that of P' coincide (modulo the semiring domain value). To compute the solution of P', we have to assign a value to each tuple for the distinguished variables. This is done by first considering all 3-tuples and assigning a value to them, and then by projecting (via the $\vee$ operation) such values over the distinguished variables. To assign a value to a 3-tuple, one has to combine (via the $\wedge$ operation) the values of all the subtuples corresponding to subsets of variables which are connected by some constraint. In Figure 2.3, for each 3-tuple we have written the values of its subtuples: subtuples are pointed out by underlining them, and subtuples of one element only are not indicated, since they all have value 1. Then the value for the tuple itself can be seen to the right as a combination of the subtuple values (again, for simplicity of the picture one-value subtuples are not considered since they always have value 1), and the value for each 2-tuple which is a projection of some 3-tuples onto the variables x and y can be seen to the left, as a combination of the values of the two tuples which have the same projection over x and y. It is easy to see that a 3-tuple gets the value 1 only when it satisfies all the constraints (due to the definition of $\wedge$), and that a 2-tuple gets a value 1 when there exists a possible extension to a 3-tuple whose value is 1 (due to the definition of $\vee$). From Figure 2.3, one can see that the only group of tuples that gets value 1 is the first one, thus the solution of the CSP is the projection of this group of tuples onto x and y, that is, $x = a$ and $y = a$.

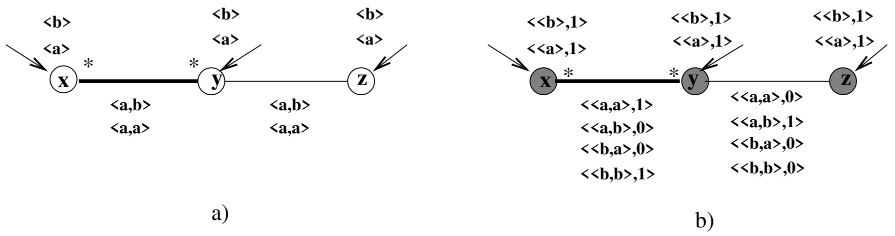

Figure 2.2: A CSP and the corresponding SCSP.

□

It is easy to see, also from the example above, that the chosen c-semiring allows us to faithfully represent any given CSP. That is, taken the given CSP P and the SCSP P' obtained by applying a suitable transformation to P, the solutions of P and P' are the same modulo a similar transformation.

It is possible to check that an alternative way to represent CSPs in the SCSP framework is by using the following c-semiring: $\langle \wp(\{a\}), \bigcup, \bigcap, \emptyset, \{a\} \rangle$, where $\wp(S)$ is the powerset of a set S, and a is any value.

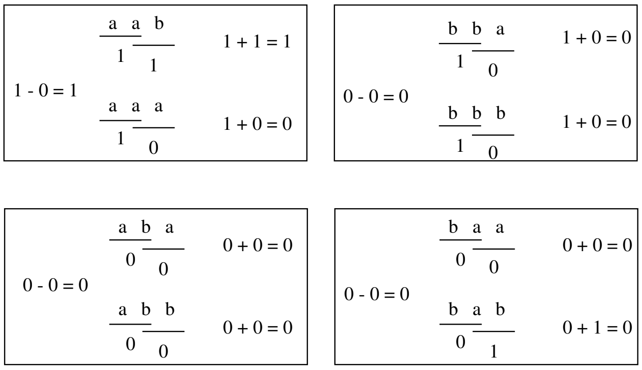

Figure 2.3: Combination and projection in classical CSPs.

2.3.2 Fuzzy CSPs (FCSPs)

Fuzzy CSPs (FCSPs) [RHZ76, DFP93, Rut94] extend the notion of classical CSPs by allowing non-crisp constraints, that is, constraints which associate a preference level with each tuple of values. Such level is always between 0 and 1, where 1 represents the best value (that is, the tuple is allowed) and 0 the worst one (that is, the tuple is not allowed). The solution of a fuzzy CSP is then defined as the set of tuples of values (for all the variables) which have the maximal value. The way they associate a value with an n-tuple is by minimizing the values of all its subtuples. The reason for such a max-min framework relies on the attempt to maximize the value of the least preferred tuple.

Fuzzy CSPs are already a very significant extension of CSPs. In fact, they are able to model partial constraint satisfaction (as described in Section 1.1.2), so as to get a solution even when the problem is overconstrained, and also prioritised constraints, that is, constraints with different levels of importance (see Section 1.1.2).

DEFINITION 2.3.3 (FUZZY CSPS) *A Fuzzy Constraint Satisfaction Problem (FCSP) consist of:*

- *a set of variables* $X = \{x_1, \ldots, x_n\}$,
- *for each variable* x_i, *a finite set* D_i *of possible values (its domain),*
- *a set of constraints; each constraint* c *is defined by a function fuzzy level* $fl(c, A)$, *that assigns a real number between* 0 *and* 1 *to each tuple* A *of domain values.*

A solution to a FCSP is an assignment of a value from its domain to every variable such that the expression $\min\{fl(c, A) \mid c$ *is a constraint in FCSP*$\}$ *is maximized among all possible assignments* A *of domain values.*

It is easy to see that fuzzy CSPs can be modeled in our framework by choosing the semiring

$$S_{FCSP} = \langle \{x \mid x \in [0, 1]\}, max, min, 0, 1 \rangle.$$

THEOREM 2.3.4 S_{FCSP} *is a c-semiring.*

PROOF. It is easy to see that it is a semiring. Thus we will only prove the additional properties that are needed in a c-semiring:

- the additive operation must be idempotent; this is trivially true since the additive operation is max, and $max(a, a) = a$ for any a;
- the multiplicative operation must be commutative; trivially from the definition of min;
- 1 is the absorbing element of the additive operation, that is $max(1, a) = 1$, which is true since $0 \leq a \leq 1$.

The ordering $\leq_S$ here reduces to the $\leq$ ordering on reals. It is also easy to see that any FCSP P can be rewritten as an SCSP P' over the semiring S_{FCSP} such that $sol(P) = sol(P')$ (modulo a suitable transformation between the problems and the solutions). □

Fuzzy CSPs are much closer to our framework than classical CSPs, since they already introduce the notion of preference levels. This means that they have to generalize the notions of constraint combination and projection as well, so as to get the appropriate definition of solution. We are, however, obviously much more general in that we do not make any assumption on the semiring operations.

2.3.3 Probabilistic CSPs (Prob-CSPs)

Probabilistic CSPs (Prob-CSPs) [FL93] were introduced to model those situations where each constraint c has a certain probability $p(c)$, independent from the probability of the other constraints, to be part of the given problem (actually, the probability is not of the constraint, but of the situation which corresponds to the constraint: saying that c has probability p means that the situation corresponding to c has probability p of occurring in the real-life problem).

This allows one to also reason about problems which are only partially known. The probability levels on constraints then gives to each instantiation of all the variables, a probability that it is a solution of the real problem. This is done by first associating with each subset of constraints the probability that it is in the real problem (by multiplying the probabilities of the involved constraints), and then by summing up all the probabilities of the subsets of constraints where the considered instantiation is a solution. Alternatively, the probability associated with an n-tuple t can also be seen as the probability that all constraints that t violates are indeed in the real problem. This is just the product of all $1 - p(c)$ for all c violated by t. Finally, the aim is to get those instantiations with the maximum probability.

DEFINITION 2.3.5 (PROBABILISTIC CSPS) *A Probabilistic Constraint Satisfaction Problem (Prob-CSP) consists of:*

- *a set of variables* $X = \{x_1, \ldots, x_n\}$,

- *for each variable* x_i*, a finite set* D_i *of possible values (its domain),*
- *a set of constraints restricting the values that the variables can simultaneously take; each constraint* c *is also assigned a certain probability* $p(c)$ *to be part of the given problem.*

There are two alternative approaches to define the solution to a Prob-CSP. Again, the solution is an assignment of a value from its domain to every variable such that

1. *the expression Sum*$\{$*Product*$\{p(c) \mid c$ *is a constraint in* $S\} \mid S$ *is a subset of the set of all constraints satisfied by* $A\}$ *is maximized among all possible assignments* A *of values, or*
2. *the expression Product*$\{(1 - p(c)) \mid c$ *is a constraint in FCSP violated by* $A\}$ *is maximized among all possible assignments* A *of values.*

The relationship between Prob-CSPs and SCSPs is complicated by the fact that Prob-CSPs contain crisp constraints with probability levels, while SCSPs contain non-crisp constraints. That is, we associate values with tuples, and not to constraints. It is still possible however to model Prob-CSPs, by using a transformation which is similar to that proposed in [DFP93] to model prioritised constraints via soft constraints in the FCSP framework. More precisely, we assign probabilities to tuples instead of constraints. Consider any constraint c with probability $p(c)$, and let t be any tuple of values for the variables involved in c. Then $p(t) = 1$ if t is allowed by c, otherwise $p(t) = 1 - p(c)$. The reasons for such a choice are as follows: if a tuple is allowed by c and c is in the real problem, then t is allowed in the real problem; this happens with probability $p(c)$; if instead c is not in the real problem, then t is still allowed in the real problem, and this happens with probability $1-p(c)$. Thus, t is allowed in the real problem with probability $p(c)+1-p(c) = 1$. Consider instead a tuple t which is not allowed by c. Then it will be allowed in the real problem only if c is not present; this happens with probability $1-p(c)$.

To give the appropriate value to an n-tuple t, given the values of all the smaller k-tuples, with $k \leq n$ and which are subtuples of t (one for each constraint), we just perform the product of the value of such subtuples. By the manner in which values have been assigned to tuples in constraints, this coincides with the product of all $1 - p(c)$ for all c violated by t. In fact, if a subtuple violates c, then by construction its value is $1 - p(c)$; if instead a subtuple satisfies c, then its value is 1. Since 1 is the unit element of $\times$, we have that $1 \times a = a$ for each a. Thus we get $\Pi(1 - p(c))$ for all c that t violates.

As a result, the semiring corresponding to the Prob-CSP framework is

$$S_{prob} = \langle \{x \mid x \in [0, 1]\}, max, \times, 0, 1 \rangle.$$

THEOREM 2.3.6 S_{prob} *is a c-semiring.*

PROOF. It is a semiring. To be also a c-semiring, we need the following:

- the additive operation must be idempotent; this is trivially true since the additive operation is max, and $max(a, a) = a$ for any a;
- the multiplicative operation must be commutative; trivially from the definition of $\times$;
- 1 is the absorbing element of the additive operation, that is $max(1, a) = 1$, which is true since $0 \leq a \leq 1$.

The associated ordering $\leq_S$ here reduces to $\leq$ over reals. $\square$

It is easy to see that any Prob-CSP P can be rewritten as a n SCSP P' over the semiring S_{prob}, such that $sol(P) = sol(P')$ (modulo a suitable transformation between the problems and the solutions). Note that the fact that P' is α-consistent means that in P there exists an n-tuple which has probability α to be a solution of the real problem.

2.3.4 Weighted CSPs

While fuzzy CSPs associate a level of preference with each tuple in each constraint, in weighted CSPs (WCSPs) tuples come with an associated cost. This allows one to model optimization problems where the goal is to minimize the total cost (time, space, number of resources, and so on) of the proposed solution. Therefore, in WCSPs the cost function is defined by summing up the costs of all constraints (intended as the cost of the chosen tuple for each constraint). Thus, the goal is to find the n-tuples (where n is the number of all the variables) which minimize the total sum of the costs of their subtuples (one for each constraint).

Another way to understand the difference between WCSPs and FCSPs is that FCSPs have an egalitarianistic approach to optimization problems (that is, they aim at maximizing the overall level of consistency while balancing the levels of all constraints), while WCSPs have an utilitarianistic approach (that is, they aim at getting the minimum cost globally, even though some constraints may be neglected and thus present a big cost) [Mou88]. We believe that both approaches present advantages and drawbacks, and thus one may be preferred to the other, depending on the real-life situation to be modeled.

According to the informal description of WCSPs given above, the associated semiring is

$$S_{WCSP} = \langle \mathcal{R}^+, \text{min}, +, +\infty, 0 \rangle.$$

THEOREM 2.3.7 *S_{WCSP} is a c-semiring.*

PROOF. It is easy to see that it is a semiring. To be a c-semiring, we need:

- the additive operation must be idempotent; this is trivially true since the additive operation is min, and $min(a, a) = a$ for any a;
- the multiplicative operation must be commutative; trivially from the definition of $+$;

- 0 is the absorbing element of the additive operation, that is $min(0, a) = 0$, which is true since $a \geq 0$.

The associated ordering $\leq_S$ reduces here to $\geq$ over the reals. This means that a value is preferred to another one if it is smaller. $\square$

Note that the same properties hold also for the semirings $\langle \mathcal{Q}^+, min, +, +\infty, 0\rangle$ and $\langle \mathcal{Z}^+, min, +, +\infty, 0\rangle$ (which can be proved to be c-semirings).

2.3.5 Egalitarianism and Utilitarianism

As noted above, the FCSP and the WCSP systems can be seen as two different approaches to giving a meaning to the notion of optimization. In fact, the two models correspond on two definitions of social welfare in utility theory [Mou88]: *egalitarianism*, which maximizes the minimal individual utility, and *utilitarianism*, which maximizes the sum of the individual utilities. FCSPs are based on the egalitarian approach, while WCSPs are based on utilitarianism.

In this section we will show how our framework also allows for the combination of these two approaches. In fact, we will construct an instance of the SCSP framework where the two approaches coexist, and allow us to discriminate among solutions which otherwise would result indistinguishable. More precisely, we can first compute the solutions according to egalitarianism (that is, using a max-min computation as in FCSPs), and then discriminate more among them via utilitarianism (that is, using a max-sum computation as in WCSPs).

The resulting semiring is the following:

$$S_{ue} = \langle \{\langle l, k\rangle \mid l \in [0,1] \text{ and } k \in \mathcal{R}^-\}, \underline{max}, \underline{min}, \langle 0, -\infty\rangle, \langle 1, 0\rangle\rangle$$

where $\underline{max}$ and $\underline{min}$ are defined as follows:

- $\langle l_1, k_1\rangle \underline{max} \langle l_2, k_2\rangle = \begin{cases} \langle l_1, max(k_1, k_2)\rangle & \text{if } l_1 = l_2 \\ \langle l_1, k_1\rangle & \text{if } l_1 > l_2 \end{cases}$

- $\langle l_1, k_1\rangle \underline{min} \langle l_2, k_2\rangle = \begin{cases} \langle l_1, k_1 + k_2\rangle & \text{if } l_1 = l_2 \\ \langle l_2, k_2\rangle & \text{if } l_1 > l_2 \end{cases}$

That is, the domain of the semiring contains pairs of values: the first element is used to reason via the max-min approach, while the second one is used to further discriminate via the max-sum approach. The operation $\underline{min}$ (which is the multiplicative operation of the semiring, and thus is used to perform the constraint combination) takes two elements of the semiring, say $a = \langle a_1, a_2\rangle$ and $b = \langle b_1, b_2\rangle$) and returns the one with the smallest first element (if $a_1 \neq b_1$), otherwise it returns the pair that has the sum of the second elements as its second element (that is, $\langle a_1, a_2 + b_2\rangle$). The operation $\underline{max}$ performs a max on the first elements if they are different (thus returning the pair with the maximum first element), otherwise (if $a_1 = b_1$) it performs a max on the second elements (thus returning a pair which has such a max as second element, that is, $\langle a_1, max(a_2, b_2)\rangle$). More abstractly, we can say that, if the first elements of the pairs differ, then the $\underline{max}$-$\underline{min}$ operations behave like a normal max-min, otherwise they behave like max-sum. This can be interpreted as the fact that, if the first elements coincide, it means that the max-min criterion cannot discriminate enough, and thus the max-sum criterion is used.

One can show that S_{ue} is a c-semiring.

2.3.6 Set-based CSPs

An interesting class of instances of the SCSP framework that are based on set operations like union and intersection is the one which uses the c-semiring

$$S_{set} = \langle \wp(A), \bigcup, \bigcap, \emptyset, A\rangle$$

where A is any set. It is easy to see that S_{set} is a c-semiring. Also, in this case the order $\leq_{S_{set}}$ reduces to set inclusion (in fact, $a \leq b$ iff $a \cup b = b$), and therefore it is partial in general. Furthermore, $\times$ is $\bigcap$ in this case, and thus it is idempotent.

A useful application of such SCSPs based on a set occurs when the variables of the problem represent processes, the finite domain D is the set of possible states of such processes, and A is the set of all time intervals over reals. In this case, a value t associated to a pair of values $\langle d_1, d_2\rangle$ for variables x and y can be interpreted as the set of all time intervals in which process x can be in state d_1 and process y in state d_2. Therefore, a solution of any SCSP based on the c-semiring S_{set} consists of a tuple of process states, together with (the set of) the time intervals in which such system configuration can occur. This description of a system can be helpful, for example, to discover if there are time intervals in which a deadlock is possible.

2.3.7 Valued Constraint Problems

The framework [SFV95] described in this subsection is based on an ordered monoid structure (that is, an ordered domain plus one operation satisfying some properties). The values of the domain are interpreted as levels of violation (which can be interpreted as cost, or degrees of preference, or probabilities, or others) and can be combined using the monoid operator. Specific choices of the monoid will then give rise to different instances of the framework.

Elements of the set of the monoid, called valuations, are associated with each constraint and the monoid operation (denoted $\circledast$) is used to assign a valuation to each assignment, by combining the valuations of all the constraints violated by the assignment. The order on the monoid is assumed to be total and the problem considered is always to minimize the combined valuation of all violated constraints. The reader should note that the choice of associating one valuation to each constraint rather than to each tuple (as in the SCSP framework) is done only for the sake of simplicity and is not fundamentally different.

In this section, a classical CSP is defined by a set $V = \{v_1, \ldots, v_n\}$ of variables, each variable v_i having an associated finite domain d_i. A constraint $c = (V_c, R_c)$ is defined by a set of variables $V_c \subseteq V$ and a relation R_c between the variables of V_c *i.e.*, a subset of the Cartesian product $\prod_{v_i \in V_c} d_i$. A CSP is denoted by $\langle V, D, C \rangle$, where D is the set of the domains and C the set of the constraints. A solution of the CSP is an assignment of values to the variables in V such that all the constraints are satisfied: for each constraint $c = (V_c, R_c)$, the tuple of the values taken by the variables of V_c belongs to R_c.

Valuation structure

In order to deal with over-constrained problems, it is necessary to be able to express the fact that a constraint may eventually be violated. To achieve this, we annotate each constraint with a mathematical item which we call a *valuation*. Such valuations will be taken from a set E equipped with the following structure:

DEFINITION 2.3.8 (VALUATION STRUCTURE) *A* valuation structure *is defined as a tuple* $\langle E, \circledast, \succ \rangle$ *such that:*

- E *is a set, whose elements are called valuations, which is totally ordered by* $\succ$*, with a maximum element noted* $\top$ *and a minimum element noted* $\bot$*;*
- $\circledast$ *is a commutative, associative closed binary operation on* E *that satisfies:*
  - Identity*:* $\forall a \in E, a \circledast \bot = a$*;*
  - Monotonicity*:* $\forall a, b, c \in E, (a \succcurlyeq b) \Rightarrow \big((a \circledast c) \succcurlyeq (b \circledast c)\big)$*;*
  - Absorbing element*:* $\forall a \in E, (a \circledast \top) = \top$.

Valued CSP

A valued CSP is then simply obtained by annotating each constraint of a classical CSP with a valuation denoting the impact of its violation or, equivalently, of its rejection from the set of constraints.

DEFINITION 2.3.9 (VALUED CSP) *A* valued CSP *is defined by a classical CSP* $\langle V, D, C \rangle$*, a valuation structure* $S = (E, \circledast, \succ)$*, and an application* φ *from* C *to* E*. It is denoted by* $\langle V, D, C, S, \varphi \rangle$. $\varphi(c)$ *is called the valuation of* c.

An assignment A of values to some variables $W \subset V$ can now be simply evaluated by combining the valuations of all the violated constraints using $\circledast$:

DEFINITION 2.3.10 (VALUATION OF ASSIGNMENTS) *In a Valued CSP* $\mathcal{P} = \langle V, D, C, S, \varphi \rangle$ the valuation of an assignment A *of the variables of* $W \subset V$ *is defined by:*

$$\mathcal{V}_{\mathcal{P}}(A) = \underset{\substack{c \in C, V_c \subset W \\ A \text{ violates } c}}{\circledast} [\varphi(c)]$$

The semantics of a VCSP is a distribution of valuations on the assignments of V (potential solutions). The problem considered is to find an assignment A with a *minimum* valuation. The valuation of such an optimal solution will be called the CSP valuation. It provides a *gradual* notion of inconsistency, from $\perp$, which corresponds to consistency, to $\top$, for complete inconsistency.

From SCSPs to VCSPs

The SCSP and the VCSP framework are similar, but, however, they are not completely equivalent: only if one assumes a total order on the semiring set, it will be possible to define appropriate mappings to pass from one of them to the other.

In this section we compare the two approaches. Some more details related to the two approaches can be found in [BFM+96, BFM+99].

We will consider SCSPs $\langle C, con\rangle$ where *con* involves all variables. Thus we will omit it in the following. An SCSP is, therefore, just a set of constraints C over a constraint system $\langle S, D, V\rangle$, where S is a c-semiring $S = \langle A, +, \times, 0, 1\rangle$, and D is the domain of the variables in V. Moreover, we will assume that the $+$ operation induces an order $\leq_S$ which is total. This means that $+$ corresponds to max, or, in other words, that it always chooses the value which is closer to 1.

Given an SCSP, we will now show how to obtain a corresponding VCSP, where by correspondence we mean that they associate the same value with each variable assignment, and that they have the same solution.

DEFINITION 2.3.11 *Given an SCSP P with constraints C over a constraint system $\langle S, D, V\rangle$, where S is a c-semiring $S = \langle A, +, \times, 0, 1\rangle$, we obtain the VCSP $P' = \langle V, D, C', S', \varphi\rangle$, where $S' = \langle E, \circledast, \succ\rangle$, where $E = A$, $\circledast = \times$, and $\prec = >_S$. (Note that $\top = 0$ and $\perp = 1$.)*

For each constraint $c = \langle def, con\rangle \in C$, we obtain a set of constraints $c'_1, \dots, c'_k$, where k is the cardinality of the range of def. That is, $k = |\{a \text{ s.t. } \exists t \text{ with } def(t) = a\}|$. Let us call $a_1, \dots, a_k$ such values, and let us call T_i the set of all tuples t such that $def(t) = a_i$. All the constraints c_i involve the same variables, which are those involved in c. Then, for each $i = 1, \dots, k$, we set $\varphi(c'_i) = a_k$, and we define c'_i in such a way that the tuples allowed are those not in T_i. We will write $P' = sv(P)$. Note that, by construction, each tuple t is allowed by all constraints $c_1, \dots, c_k$ except the constraint c_i such that $\varphi(c_i) = def(t)$.

EXAMPLE 2.3.12 Consider an SCSP which contains the constraint $c = \langle \mathtt{def}, \mathtt{con}\rangle$, where $\mathtt{con} = \{x, y\}$, and $\mathtt{def}(\langle a, a\rangle) = l_1$, $\mathtt{def}(\langle a, b\rangle) = l_2$, $\mathtt{def}(\langle b, a\rangle) = l_3$, $\mathtt{def}(\langle b, b\rangle) = l_1$. Then, the corresponding VCSP will contain the following three constraints, all involving x and y:

- c_1, with $\varphi(c_1) = l_1$ and allowed tuples $\langle a, b\rangle$ and $\langle b, a\rangle$;
- c_2, with $\varphi(c_2) = l_2$ and allowed tuples $\langle a, a\rangle$, $\langle b, a\rangle$ and $\langle b, b\rangle$;

- c_3, with $\varphi(c_3) = l_3$ and allowed tuples $\langle a, a\rangle$, $\langle a, b\rangle$ and $\langle b, b\rangle$.

Figure 2.4 shows both c and the three constraints c_1, c_2, and c_3.

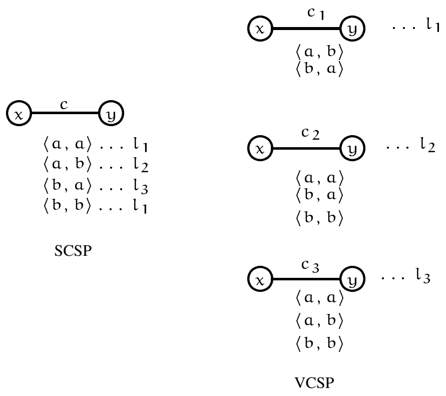

Figure 2.4: From SCSPs to VCSPs.

□

First we need to make sure that the structure we obtain via the definition above is indeed a valued CSP. Then we will prove that it is equivalent to the given SCSP.

THEOREM 2.3.13 (FROM C-SEMIRING TO VALUATION STRUCTURE) *If we consider a c-semiring* $S = \langle A, +, \times, 0, 1\rangle$ *and the structure* $S' = \langle E, \circledast, \succ\rangle$*, where* $E = A$*,* $\circledast = \times$*, and* $\prec = >_S$ *(obtained using the transformation in Definition 2.3.11), then* S' *is a valuation structure.*

PROOF. First, $\succ$ is a total order, since we assumed that $\leq_S$ is total and that $\prec = >_S$. Moreover, $\top = 0$ and $\perp = 1$, from the definition of $\leq_S$. Then, since $\circledast$ coincides with $\times$, it is easy to see that it is commutative, associative, monotone, and closed. Moreover, 1 (that is, $\perp$) is its unit element and 0 (that is, $\top$) is its absorbing element. □

THEOREM 2.3.14 (EQUIVALENCE BETWEEN SCSP AND VCSP) *Consider an SCSP* P *and the corresponding VCSP problem* P'*. Consider also an assignment* t *to all the variables of* P*, with associated value* $a \in A$*. Then* $\mathcal{V}_{P'}(t) = a$*.*

PROOF. Note first that P and P' have the same set of variables. In P, the value of t is obtained by multiplying the values associated with each subtuple of t, one for each constraint of P. Thus, $a = \prod\{def_c(t_c)$, for all constraints $c = \langle def_c, con_c\rangle$ in C and such that t_c is the projection of t over the variables of $c\}$. Now, $\mathcal{V}_{P'}(t) = \prod\{\varphi(c')$ for all $c' \in C'$ such that the projection of t over the variables of c' violates $c'\}$. It is easy to see that the number of values multiplied in this formula coincides with the number of constraints in C, since, as noted above, each tuple violates only one of the constraints in C' that have been generated because of the presence of the constraint $c \in C$. Thus, we just have to show that, for each $c \in C$, $def_c(t_c) = \varphi(c_i)$, where c_i is the constraint violated by t_c. But this is easy to show by what we have noted above. In fact, we have defined the translation from SCSP to VCSP in such a way that the only constraint of the VCSP violated by a tuple is exactly the one whose valuation coincides with the value associated with the tuple in the SCSP. $\square$

Note that SCSPs that do not satisfy the restrictions imposed at the beginning of the section, that is, that *con* involves all the variables and that $\leq_S$ is total, do not have a corresponding VCSP.

COROLLARY 2.3.15 (SAME SOLUTION) *Consider an SCSP* P *and the corresponding VCSP problem* P'. *Then* P *and* P' *have the same solution.*

PROOF. It follows from Theorem 2.3.14, from the fact that P uses $+ = max$ which goes towards 1, that P' uses min which goes towards $\perp$, that $1 = \perp$, and that a solution is just one of the total assignments. $\square$

From VCSPs to SCSPs

Here we will define the opposite translation, which allows one to get an SCSP from a given VCSP.

DEFINITION 2.3.16 *Given the VCSP* $P = \langle V, D, C, S, \varphi \rangle$, *where* $S = \langle E, \circledast, \succ \rangle$, *we will obtain the SCSP* P' *with constraints* C' *over the constraint system* $\langle S', D, V \rangle$, *where* S' *is the c-semiring* $\langle E, +, \circledast, \top, \perp \rangle$, *and* $+$ *is such that* $a + b = a$ *iff* $a \preccurlyeq b$. *It is easy to see that* $\geq_S = \preccurlyeq$. *For each constraint* $c \in C$ *with allowed set of tuples* T, *we define the corresponding constraint* $c' = \langle con', def' \rangle \in C'$ *such that* con' *contains all the variables involved in* c *and, for each tuple* $t \in T$, *def*$'(t) = \perp$, *otherwise* $def'(t) = \varphi(c)$. *We will write* $P' = vs(P)$.

EXAMPLE 2.3.17 Consider a VCSP which contains a binary constraint c connecting variables x and y, for which it allows the pairs $\langle a, b \rangle$ and $\langle b, a \rangle$, and such that $\varphi(c) = l$. Then, the corresponding SCSP will contain the constraint $c' = \langle con, def \rangle$, where *con* $= \{x, y\}$, $def(\langle a, b \rangle) = def(\langle b, a \rangle) = \perp$, and $def(\langle a, a \rangle) = def(\langle b, b \rangle) = l$. Figure 2.5 shows both c and the corresponding c'.

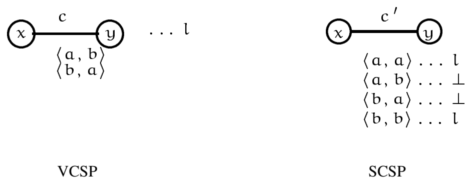

Figure 2.5: From VCSP to SCSP.

$\square$

Again, we need to make sure that the structure we obtain via the definition above is indeed a semiring-based CSP. Then we will prove that it is equivalent to the given VCSP.

THEOREM 2.3.18 (FROM VALUATION STRUCTURE TO C-SEMIRING) *If we consider a valuation structure $\langle E, \circledast, \succ \rangle$ and the structure $S = \langle E, +, \circledast, \top, \bot \rangle$, where $+$ is such that $a + b = a$ iff $a \preccurlyeq b$ (obtained using the transformation in Definition 2.3.16), then S is a c-semiring.*

PROOF. Since $\succ$ is total, $+$ is closed. Moreover, $+$ is commutative by definition, and associative because of the transitivity of the total order $\succ$. Furthermore, 0 is the unit element of $+$, since it is the top element of $\succ$. Finally, $+$ is idempotent because of the reflexivity of $\succ$, and 1 is the absorbing element of $+$ since $1 = \bot$. Operation $\times$ of S coincides with $\circledast$. Thus it is closed, associative, and commutative, since $\circledast$ is so. Also, $\top$ is its absorbing element and $\bot$ is its identity (from corresponding properties of $\circledast$). The distributivity of $\circledast$ over $+$ can easily be proved. For example, consider $a, b, c \in E$, and assume $b \preccurlyeq c$. Then $a \circledast (b + c) = a \circledast b$ (by definition of $+$) = $(a \circledast b) + (a \circledast c)$ (by the definition of $+$ and the monotonicity of $\circledast$). The same reasoning applies to the case where $c \preccurlyeq b$. □

THEOREM 2.3.19 (EQUIVALENCE BETWEEN VCSP AND SCSP) *Consider a VCSP problem P and the corresponding SCSP P'. Consider also an assignment t to all the variables of P. The value associated with such an assignment is $A = \mathcal{V}_P(t) = \circledast\{\varphi(c)$ for all $c \in C$ such that the projection of t over the variables of c violates $c\}$. Instead, the value associated with the same assignment in P' is $B = \circledast\{\mathtt{def}_{c'}(t_{c'})$, for all constraints $c' = \langle \mathtt{def}_{c'}, \mathtt{con}_{c'} \rangle$ in C' and such that $t_{c'}$ is the projection of t over the variables of $c'\}$. Then, $A = B$.*

PROOF. The values multiplied to produce A are as many as the constraints violated by t; instead, the values multiplied to produce B are as many as the constraints in C'. By construction, however, each tuple t_c involving the variables of a constraint $c \in C$ has been associated, in P', with a value that is either $\varphi(c)$ (if t_c violates c), or $\bot$ (if t_c satisfies c). Thus, the contribution of t_c to the value of B is important only if t_c violated c in P, because $\bot$ is the unit element for $\circledast$. Thus A and B are obtained by the same number of significant values. Now we have to show that such values are the same. But this is easy, since we have defined the translation in such a way that each tuple for the variables of c is associated with the value $\varphi(c)$ exactly when it violates c. □

Normal Forms and Equivalences

Note that, while passing from an SCSP to a VCSP the number of constraints in general increases, in the opposite direction the number of constraints remains the same. This can also be seen in Example 2.3.12 and 2.3.17. This means that, in general, going from an SCSP P to a VCSP P' and then from the VCSP P' to the SCSP P'', we do not get $P = P''$. In fact, for each constraint c in P, P'' will have in general several constraints $c_1, \ldots, c_k$ over the same variables as c. It is easy to see, however, that $c_1 \otimes \cdots \otimes c_k = c$, and thus P and P'' associate the same value with each variable assignment.

EXAMPLE 2.3.20 Figure 2.6 shows how to pass from an SCSP to the corresponding VCSP (this part is the same as in Example 1), and then again to the corresponding SCSP. Note that the starting SCSP and the final one are not the same. In fact, the latter has three constraints between variables x and y, while the former has only one constraint. One can see, however, that the combination of the three constraints yields the starting constraint.

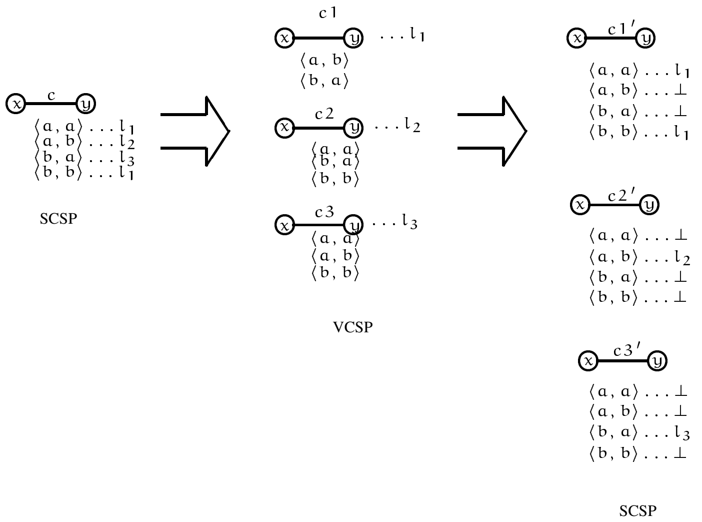

Figure 2.6: From SCSP to VCSP and back to SCSP again.

□

Consider now the opposite cycle, that is, going from a VCSP P to an SCSP P' and then from P' to a VCSP P''. In this case, for each constraint c in P, P'' has two constraints: one is c itself, and the other one is a constraint with associated value $\perp$. This means that violating such a constraint has cost $\perp$, which, in other words, means that this constraint can be eliminated without changing the behaviour of P'' at all.

EXAMPLE 2.3.21 Figure 2.7 shows how to pass from a VCSP to the corresponding SCSP (this part is the same as in Example 2), and then again to the corresponding VCSP. Note that the starting VCSP and the final one are not the same. In fact, the latter one has two constraints between variables x and y. One is the same as the one in the starting VCSP, while the other one has the value $\perp$ associated with it. This means that violating such constraint yields a cost of value $\perp$.

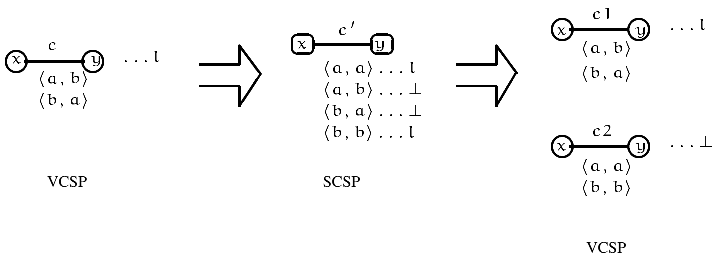

Figure 2.7: From VCSP to SCSP, and to VCSP again.

□

Let us now define normal forms for both SCSPs and VCSPs, as follows. For each VCSP P, its normal form is the VCSP $P' = nfv(P)$ which is obtained by deleting all constraints c such that $\varphi(c) = \bot$. It is easy to see that P and P' are equivalent.

DEFINITION 2.3.22 *Consider* $P = \langle V, D, C, S, \varphi \rangle$*, a VCSP where* $S = \langle E, \circledast, \succ \rangle$*. Then* P *is said to be in normal form if there is no* $c \in C$ *such that* $\varphi(c) = \bot$*. If* P *in not in normal form, then it is possible to obtain a unique VCSP* $P' = \langle V, D, C - \{c \in C \mid \varphi(c) = \bot\}, S, \varphi \rangle$*, denoted by* $P' = nfv(P)$*, which is in normal form.*

THEOREM 2.3.23 (NORMAL FORM) *For any VCSP* P*,* P *and* $nfv(P)$ *are equivalent.*

PROOF. The theorem follows from the fact that $\forall a, (\bot \circledast a) = a$ and from the definitions of $\mathcal{V}_P(A)$ and $\mathcal{V}_{P'}(A)$. □

Also, for each SCSP P, its normal form is the SCSP $P' = nfs(P)$, which is obtained by combining all constraints involving the same set of variables. Again, this is an equivalent SCSP.

DEFINITION 2.3.24 *Consider any SCSP* P *with constraints* C *over a constraint system* $\langle S, D, V \rangle$*, where* S *is a c-semiring* $S = \langle A, +, \times, 0, 1 \rangle$*. Then,* P *is in normal form if, for each subset* W *of* V*, there is at most one constraint* $c = \langle \mathtt{def}, con \rangle \in C$ *such that* $con = W$*. If* P *is not in normal form, then it is possible to obtain a unique SCSP* P'*, as follows. For each* $W \subseteq V$*, consider the set* $C_W \subseteq C$*, which contains all the constraints involving* W*. Assume* $C_W = \{c_1, \ldots, c_n\}$*. Then, replace* C_W *with the single constraint* $c = \bigotimes C_W$*.* P'*, denoted by* $nfs(P)$*, is in normal form.*

THEOREM 2.3.25 (NORMAL FORMS) *For any SCSP* P*,* P *and* $nfs(P)$ *are equivalent.*

PROOF. It follows from the associative property of $\times$. □

Even though, as noted above, the transformation from an SCSP P to the corresponding VCSP P' and then again to the corresponding SCSP P'' does not necessarily yield $P = P''$, we will now prove that there is a strong relationship between P and P''. In particular, we will prove that the normal forms of P and P'' coincide. The same holds for the other cycle, where one passes from a VCSP to an SCSP and then to a VCSP again.

THEOREM 2.3.26 (SAME NORMAL FORM 1) *Given any SCSP* P *and the corresponding VCSP* $P' = sv(P)$*, consider the SCSP* P'' *corresponding to* P'*, that is,* $P'' = vs(P')$*. Thus* $nfs(P) = nfs(P'')$*.*

PROOF. We will consider one constraint at a time. Take any constraint c of P. With the first transformation (to the VCSP P'), we get as many constraints as the different values associated with the tuples in c. Each of the constraints, say c_i, is such that $\varphi(c_i)$ is equal to one of such values, say l_i, and allows all tuples that do not have value l_i in c. With the second transformation (to the SCSP P''), for each of the c_i, we get a constraint c_i', where tuples which are allowed by c_i

have value $\perp$ while the others have value l_i. Now, if we apply the normal form to P'', we combine all the constraints c'_i, getting one constraint that is the same as c, since, given any tuple t, it is easy to see that t is forbidden by exactly one of the c_i. Thus, the combination of all c'_i will associate with t a value, which is the one associated with the unique c_i that does not allow t. $\square$

THEOREM 2.3.27 (SAME NORMAL FORM 2) *Given any VCSP problem P and the corresponding SCSP $P' = vs(P)$, consider the VCSP P'' corresponding to P', that is, $P'' = sv(P')$. Then we have that nfv(P) =nfv(P'').*

PROOF. We will consider one constraint at a time. Take any constraint c in P, and assume that $\varphi(c) = l$ and that c allows the set of tuples T. With the first transformation (to the SCSP P'), we get a corresponding constraint c' where tuples in T have value $\perp$ and tuples not in T have value l. With the second transformation (to the VCSP P''), we get two constraints: c_1, with $\varphi(c_1) = \perp$, and c_2, with $\varphi(c_2) = l$ and which allows the tuples of c' with value $\perp$. It is easy to see that $c_2 = c$. Now, if we apply the normal form to both P and P'', which implies the deletion of all constraints with value $\perp$, we get exactly the same constraint. This reasoning applies even if the starting constraint has value $\perp$. In fact, in this case the first transformation will give us a constraint where all tuples have value $\perp$, and the second one gives us a constraint with value $\perp$, which will be deleted when obtaining the normal form. $\square$

The statements of the above two theorems can be summarized by the two diagrams represented in Figure 2.8. Note that in such diagrams each arrow represents one of the transformations defined above, and all problems in the same diagram are equivalent (by the theorems proved previously in this section).

$$\begin{tikzcd}
\mathrm{VCSP} \arrow[r, "vs"] \arrow[d, "nfv"] & \mathrm{SCSP} \arrow[d, "sv"] \\
\mathrm{VCSP} & \mathrm{VCSP} \arrow[l, "nfv"]
\end{tikzcd}$$

$$\begin{tikzcd}
\mathrm{SCSP} \arrow[r, "sv"] \arrow[d, "nfs"] & \mathrm{VCSP} \arrow[d, "vs"] \\
\mathrm{SCSP} & \mathrm{SCSP} \arrow[l, "nfs"]
\end{tikzcd}$$

Figure 2.8: Diagrams representing the thesis of Theorem 2.3.26 and 2.3.26.

2.3.8 N-dimensional c-semirings

Choosing an instance of the SCSP framework means specifying a particular c-semiring. This, as discussed above, induces a partial order, which can be interpreted as a (partial) guideline for choosing the "best" among different solutions. In many real-life situations, however, one guideline is not enough, since, for example, it could be necessary to reason with more than one objective in mind, and thus choose solutions which achieve a good compromise w.r.t. all such goals.

Consider, for example, a network of computers, where one would like to both minimize the total computing time (thus the cost) and also to maximize the work of the least used computers. Then, in our framework, we would need to consider two c-semirings, one for cost minimization (see Section 2.3.4 on weighted CSPs), and another one for work maximisation (see Section 2.3.2 on fuzzy CSPs). Then one could work first with one of these c-semirings and then with the other one, trying to combine the solutions which are the best for each of them. A much simpler approach, however, consists of combining the two c-semirings and only then work with the resulting structure. The nice property is that such a structure is a c-semiring itself, thus all the techniques and properties of the SCSP framework can be used for such a structure as well.

More precisely, the way to combine several c-semirings and getting another c-semiring simply consists of vectorizing the domains and operations of the combined c-semirings.

DEFINITION 2.3.28 (COMPOSITION OF C-SEMIRINGS) *Given* n *c-semirings* $S_i = \langle A_i, +_i, \times_i, 0_i, 1_i \rangle$, *for* $i = 1, \ldots, n$, *let us define the structure* $\mathrm{Comp}(S_1, \ldots, S_n) = \langle \langle A_1, \ldots, A_n \rangle, +, \times, \langle 0_1, \ldots, 0_n \rangle, \langle 1_1 \ldots 1_n \rangle \rangle$. *Given* $\langle a_1, \ldots, a_n \rangle$ *and* $\langle b_1, \ldots, b_n \rangle$ *such that* $a_i, b_i \in A_i$ *for* $i = 1, \ldots, n$, $\langle a_1, \ldots, a_n \rangle + \langle b_1, \ldots, b_n \rangle = \langle a_1 +_1 b_1, \ldots, a_n +_n b_n \rangle$, *and* $\langle a_1, \ldots, a_n \rangle \times \langle b_1, \ldots, b_n \rangle = \langle a_1 \times_1 b_1, \ldots, a_n \times_n b_n \rangle$.

We will now prove that by composing c-semirings we get another c-semiring. This means that all the properties of the framework defined in the previous sections hold.

THEOREM 2.3.29 (A C-SEMIRING COMPOSITION IS A C-SEMIRING) *Given* n *c-semirings* $S_i = \langle A_i, +_i, \times_i, 0_i, 1_i \rangle$, *for* $i = 1, \ldots, n$, *the structure* $\mathrm{Comp}(S_1, \ldots, S_n)$ *is a c-semiring.*

PROOF. It is easy to see that all the properties required for being a c-semiring hold for $\mathrm{Comp}(S_1, \ldots, S_n)$, since they directly follow from the corresponding properties of the component semirings S_i. □

According to the definition of the ordering $\leq_S$ (in Section 2.1), such an ordering for $S = \mathrm{Comp}(S_1, \ldots, S_n)$ is as follows. Given $\langle a_1, \ldots, a_n \rangle$ and $\langle b_1, \ldots, b_n \rangle$ such that $a_i, b_i \in A_i$ for $i = 1, \ldots, n$, we have $\langle a_1, \ldots, a_n \rangle \leq_S \langle b_1, \ldots, b_n \rangle$ if and only if $\langle a_1 +_1 b_1, \ldots, a_n +_n b_n \rangle = \langle b_1, \ldots, b_n \rangle$. Since the tuple elements are completely independent, $\leq_S$ is in general a partial order, even if each of the $\leq_{S_i}$ is a total order. This means (see Section 2.2) that the abstract solution of a problem over such a semiring in general contains an incomparable set of tuples, none of which has $\mathrm{blevel}(P)$ as its associated value. Therefore, if one wants to reduce the number of "best" tuples (or to get just one), one has to specify some priorities among the orderings of the component c-semirings.

Notice that, although the c-semiring discussed in Section 2.3.5 may seem a composition of two c-semirings, as defined in Definition 2.3.28, this is not so, since the behaviour of each of the two operations on one of the two elements of each pair (we recall that that semiring is a set of pairs) is not independent from the behaviour of the same operation on the other element of the pair. Thus, it is not possible to define two independent operations (that is, $+_1$ and $+_2$, or $\times_1$ and $\times_2$), as needed by Definition 2.3.28. For example, operation $\underline{\mathtt{max}}$ returns the maximum of the second elements of two pairs *only when* the first elements of the pairs are equal.

2.4 Conclusions

In this chapter we have introduced the SCSP framework and we have shown how it can be instantiated to represent several *non-crisp* constraint formalisms. We have described the different semirings that are needed to represent fuzziness, probabilities and optimizations. Furthermore, we have shown how to build new semirings by combining the existing ones.

The structure that we have defined in this framework will be the starting point for the results of the next chapters. In particular, in the next chapter we will review several solution techniques used in the classical frameworks, and we will prove their applicability to SCSPs, by looking only to the properties of the semiring operations.

Chapter 3

Towards SCSPs Solutions

Abstract

Local consistency algorithms, as usually used for classical CSPs, can be exploited in the SCSP framework as well, provided that certain conditions on the semiring operations are satisfied. We show how the SCSP framework can be used to model both old and new constraint solving and optimization schemes, thus allowing one to both formally justify many informally taken choices in existing schemes, and to prove that local consistency techniques can also be used in newly defined schemes. We generalize to soft constraints the approximation techniques usually used for local consistency in classical constraint satisfaction and programming. The theoretical results show that this is indeed possible without losing the fundamental properties of such techniques (and the experimental results (on partial arc-consistency) [Geo99] show that this work can help develop more efficient implementations for logic-based languages working with soft constraints). Then, we consider dynamic programming-like algorithms, and we prove that these algorithms can always be applied to SCSPs, and have a linear time complexity when the given SCSPs can be provided with a parsing tree of bounded size. Finally, we provide several instances of SCSPs which show the generality and also the expressive power of the framework.

In classical CSPs, so-called local consistency techniques [Fre78, MF85, Fre88, Mac92, Mac77, Mon74, MR91] have been proven to be very effective when approximating the solution of a problem. In this chapter we study how to generalize these techniques to our framework, and we provide sufficient conditions over the semiring operations which assure that they can also be fruitfully applied to the considered scheme. Here, by "fruitfully applicable" we mean that 1) the algorithm terminates and 2) the resulting problem is equivalent to the given one and it does not depend on the nondeterministic choices made during the algorithm. In particular, such conditions rely mainly on having an idempotent operator (the $\times$ operator of the semiring).

The advantage of our framework, that we call SCSP (for Semiring-based CSP), is that one can hopefully see one's own constraint solving paradigm as an instance of SCSP over a certain semiring, and can inherit the results obtained for the general scheme. In particular, one can immediately see whether a local consistency technique can be applied. In fact, our sufficient conditions, which are related to the chosen semiring, guarantee that the above basic properties of local consistency hold. Therefore, if they are satisfied, local consistency can safely be applied. Otherwise, it means that we cannot be sure, in general, that local consistency will be meaningful in the chosen instance.

In this chapter we consider several known and new constraint solving frameworks, casting them as instances of SCSP, and we study the possibility of applying local consistency algorithms. In particular, we confirm that CSPs enjoy all the properties needed to use such algorithms, and that they can also be applied to FCSPs. Moreover, we consider probabilistic CSPs [FL93] and we see that local consistency algorithms might be meaningless for such a framework, as well as for weighted CSPs (since the above-cited sufficient conditions do not hold).

We also define a suitable notion of dynamic programming over SCSPs, and we prove that it can be used over any instance of the framework, regardless of the properties of the semiring. Furthermore, we show that when it is possible to provide a whole class of SCSPs with parsing trees of bounded size, then the SCSPs of the class can be solved in linear time.

To complete our study of the local consistency algorithms, we extend to the soft case some of the partial solving techniques that have proved useful for classical constraints. In fact, such techniques are widely used in current constraint programming systems like `clp(FD)` [CD96], Ilog Solver [Pug94] and CHIP[AB93].

To do this, we define a class of Partial Local Consistency (PLC) algorithms, which perform less pruning than (complete) local consistency but can, nevertheless, be rather useful in reducing the search. Some instances of these algorithms are able to reduce the domains of the variables in constant time, and can work with a convenient representation of the domains of the variables. Using PLC together with labeling, the aim is to find the solutions of a problem with a better practical complexity bound. Moreover we study the termination condition in the SCSP framework of the general local consistency algorithms whose rules have to be only monotone and extensive (we eliminate the hypothesis of idempotency).

The chapter is organized as follows. Section 3.1 introduces the concept of local consistency for SCSPs and gives sufficient conditions for the applicability of the local consistency algorithms. Section 3.2 proves the applicability, or not, of the local consistency algorithms to several instances. Then in Section 3.3 we show some properties of the soft version of the Arc Consistency, and in Section 3.4 we describe some approximated algorithms of local consistency. Moreover Section 3.5 describes local consistency algorithms as a set of rules and gives sufficient conditions for the termination of the soft version of the GI schema [Apt99a, Apt99b]. Finally Section 3.6 describes a dynamic programming algorithm to solve SCSPs that can be applied to any instance of our framework, without any condition.

The firsts two section of this chapter appeared as part of [BMR95, BMR97b]. Moreover Section 3.4, 3.3 and 3.5, already appeared respectively in [BR98], [BCGR00b] and [BGR00].

3.1 Soft local consistency

Computing any of the previously defined notions (the best level of consistency, the solution, and the abstract solution) is an NP-hard problem. Thus, it can be convenient in many cases to approximate such notions. In classical CSPs, this is done using the so-called local consistency techniques [Fre78, Fre88, Kum92, MR91]. The main idea is to choose some subproblems in which to eliminate local inconsistency, and then iterate such elimination in all the chosen subproblems until stability. The most widely know local consistency algorithms are arc-consistency [Mac77], where subproblems contain just one constraint, and path-consistency [Mon74], where subproblems are triangles (that is, they are complete binary graphs over three variables). The concept of k-consistency [Fre78], where subproblems contain all constraints among subsets of k chosen variables, generalizes them both. All subproblems are, however, of the same size (which is k). In general, one may imagine several algorithms where, instead, the considered subparts of the SCSP have different sizes. This generalization of the k-consistency algorithms were described in [MR91] for classical CSPs. Here we follow the same approach as in [MR91], but we extend the definitions and results presented there to the SCSP framework, and we show that all the properties still hold, provided that certain properties of the semiring are satisfied.

Applying a local consistency algorithm to a constraint problem means making explicit some implicit constraints, thus possibly discovering inconsistency at a local level. In classical CSPs, this is crucial, since local inconsistency implies global inconsistency. We will now show that such a property also holds for SCSPs.

DEFINITION 3.1.1 (LOCAL INCONSISTENCY) *Consider an SCSP* $P = \langle C, con \rangle$. *Then we say that* P *is locally inconsistent if there exists* $C' \subseteq C$ *such that* $\mathtt{blevel}(C') = 0$.

THEOREM 3.1.2 (NECESSITY OF LOCAL CONSISTENCY) *Consider an SCSP* P *which is locally inconsistent. Then it is not consistent.*

PROOF. We have to prove that $\mathtt{blevel}(P) = 0$. We know that P is locally inconsistent. That is, there exists $C' \subseteq C$ such that $\mathtt{blevel}(C') = 0$. By Theorem 2.2.18, we have that $\mathtt{blevel}(P) \leq_S \mathtt{blevel}(C')$, thus $\mathtt{blevel}(P) \leq_S 0$. Since 0 is the minimum in the ordering $\leq_S$, then we immediately have that $\mathtt{blevel}(P)$ must be 0 as well. □

In the SCSP framework, we can be even more precise, and relate the best level of consistency of the whole problem (or, equivalently, of the set of all its constraints) to that of its subproblems, even though such a level is not 0. In fact, it is possible to prove that if a problem is α-consistent, then all its subproblems are β-consistent, where $\alpha \leq_S \beta$.

THEOREM 3.1.3 (LOCAL AND GLOBAL α-CONSISTENCY) *Consider a set of constraints* C *over* CS, *and any subset* C' *of* C. *If* C *is* α*-consistent, then* C' *is* β*-consistent, with* $\alpha \leq_S \beta$.

PROOF. The proof is similar to the previous one. If C is α-consistent, it means, by definition, that $\mathtt{blevel}(C) = \alpha$. Now, if we take any subset C' of C, by Theorem 2.2.18 we have that $\mathtt{blevel}(C) \leq_S \mathtt{blevel}(C')$. Thus $\alpha \leq_S \mathtt{blevel}(C')$. □

In order to define the subproblems to be considered by a local consistency algorithm, we use the notion of local consistency rules. The application of such a rule consists of solving one of the chosen subproblems. To model this, we need the additional notion of *typed location*. Informally, a typed location is just a location l (as in ordinary store-based programming languages) which has a set of variables con as type, and thus can only be assigned a constraint $c = \langle def, con\rangle$ with the same type. In the following we assume to have a location for every set of variables, and thus we identify a location with its type. Note that at any stage of the computation the store will contain a constraint problem with only a finite number of constraints, and thus the relevant locations will always be in a finite number.

DEFINITION 3.1.4 (TYPED LOCATION) *A typed location l is a set of variables.*

DEFINITION 3.1.5 (VALUE OF A LOCATION) *Given a CSP $P = \langle C, con\rangle$, the value $[l]_P$ of the location l in P is defined as the constraint $\langle def, l\rangle \in C$ if it exists, as $\langle 1, l\rangle$ otherwise. Given n locations $l_1, \ldots, l_n$, the value $[\{l_1, \ldots, l_n\}]_P$ of this set of locations in P is defined instead as the set of constraints $\{[l_1]_P, \ldots, [l_n]_P\}$.*

DEFINITION 3.1.6 (ASSIGNMENT) *An assignment is a pair $l := c$ where $c = \langle def, l\rangle$. Given a CSP $P = \langle C, con\rangle$, the result of the assignment $l := c$ is the problem $[l := c](P)$ defined as:*

$$[l := c](P) = \langle\{\langle def', con'\rangle \in C \mid con' \neq l\} \cup c, con\rangle.$$

Thus an assignment $l := c$ is seen as a function from constraint problems to constraint problems, which modifies a given problem by changing just one constraint, the one with type l. The change consists in replacing such a constraint with c. If there is no constraint of type l, then constraint c is added to the given problem.

DEFINITION 3.1.7 (LOCAL CONSISTENCY RULE) *Given an SCSP $P = \langle C, con\rangle$, a local consistency rule r for P is defined as $r = l \leftarrow L$, where L is a set of locations, l is a location, and $l \notin L$.*

Applying r to P means assigning to location l the constraint obtained by solving the subproblem of P containing the constraints specified by the locations in $L \cup \{l\}$.

DEFINITION 3.1.8 (RULE APPLICATION) *Given a local consistency rule $r = l \leftarrow L$ and a constraint problem P, the result of applying r to P is*

$$[l \leftarrow L](P) = [l := Sol(\langle [L \cup \{l\}]_P, l\rangle)](P).$$

Since a rule application is defined as a function from problems to problems, the application of a sequence S of rules to a problem is easily provided by function composition. Thus we have that $[r_1; S](P) = [S]([r_1](P))$.

Notice for example that $[l \leftarrow \emptyset](P) = P$.

In other words, a *local consistency rule*, written also r_l^L, is a function r_l^L which, taken any problem P and , returns $r_l^L(P) = [l := Sol(\langle [L]_P, l\rangle)](P)$.

THEOREM 3.1.9 (EQUIVALENCE FOR RULES) *Given an SCSP* P *and a rule* r *for* P, *we have that* $P \equiv [r](P)$ *if* $\times$ *is idempotent.*

PROOF. Assume that $P = \langle C, con\rangle$, and let $r = l \leftarrow L$. Then we have $P' = [l \leftarrow L](P) = [l := Sol(\langle [L \cup \{l\}]_P, l\rangle)](P) = \langle \{\langle def', con'\rangle \in C \mid con' \neq l\} \cup Sol(\langle [L \cup \{l\}]_P, l\rangle), con\rangle$. Let now $C(r) = [L \cup \{l\}]_P$, and $C' = C - C(r)$. Then P contains the constraints in the set $C' \cup C(r)$, while P' contains the constraints in $C' \cup (C(r) - c) \cup ((\bigotimes C(r)) \Downarrow_l)$, where $c = \langle def, l\rangle = [l]_P$. In fact, by definition of solution, $Sol(\langle [L \cup \{l\}]_P, l\rangle)$ can be written as $(\bigotimes C(r)) \Downarrow_l$. Since the set C' is present in both P and P', we will not consider it, and we will instead prove that the constraint $c_{pre} = \bigotimes C(r)$ coincides with $c_{post} = (\bigotimes(C(r) - \{c\})) \otimes ((\bigotimes C(r)) \Downarrow_l)$. In fact, if this is so, then $P \equiv P'$.

First of all, it is easy to see that c_{pre} and c_{post} have the same type $\bigcup(L \cup \{l\})$. Let us now consider the definitions of such two constraints. Let us set $c_i = [l_i]_P = \langle def_i, l_i\rangle$ for all $l_i \in L$, and assume $L = \{l_1, \ldots, l_n\}$. Thus, $C(r) = \{c, c_1 \ldots, c_n\}$. We have that, taken any tuple t of length $|\bigcup(L \cup \{l\})|$, $def_{pre}(t) = def(t\downarrow_l) \times (\Pi_i def_i(t\downarrow_{l_i}))$, while $def_{post}(t) = (\Pi_i def_i(t\downarrow_{l_i})) \times \Sigma_{t'|t'\downarrow_l = t\downarrow_l}(def(t'\downarrow_l) \times (\Pi_i def_i(t'\downarrow_{l_i})))$. Now, since the sum is done over all t' such that $t'\downarrow_l = t\downarrow_l$, we have that $def(t'\downarrow_l) = def(t\downarrow_l)$. Thus $def(t'\downarrow_l)$ can be taken out from the sum since it appears the same in each factor. Thus we have $def_{post}(t) = (\Pi_i def_i(t\downarrow_{l_i})) \times def(t\downarrow_l) \times \Sigma_{t'|t'\downarrow_l = t\downarrow_l}(\Pi_i def_i(t'\downarrow_{l_i}))$. Consider now $\Sigma_{t'|t'\downarrow_l = t\downarrow_l}(\Pi_i def_i(t'\downarrow_{l_i}))$: one of the t' such that $t'\downarrow_l = t\downarrow_l$ must be equal to t. Thus the sum can be written also as $(\Pi_i def_i(t\downarrow_{l_i}) + \Sigma_{t'|t'\downarrow_l = t\downarrow_l, t' \neq t}(\Pi_i def_i(t'\downarrow_{l_i}))$. Thus we have $def_{post}(t) = (\Pi_i def_i(t\downarrow_{l_i})) \times def(t\downarrow_l) \times (\Pi_i def_i(t\downarrow_{l_i}) + \Sigma_{t'|t'\downarrow_l = t\downarrow_l, t' \neq t}(\Pi_i def_i(t'\downarrow_{l_i})))$.

Let us now consider any two elements a and b of the semiring. Then it is easy to see that $a \times (a + b) = a$. In fact, by intensivity of $\times$ (see Theorem 2.1.5), we have that $a \times c \leq_S a$ for any c. Thus, if we choose $c = a + b$, $a \times (a + b) \leq_S a$. On the other hand, since $a + b$ is the lub of a and b, we have that $a + b \geq_S a$. Thus, by monotonicity of $\times$, we have $a \times (a + b) \geq_S a \times a$. Now, since we assume that $\times$ is idempotent, we have that $a \times a = a$, thus $a \times (a + b) \geq_S a$. Therefore $a \times (a + b) = a$. This result can be used in our formula for $def_{post}(t)$, by setting $a = (\Pi_i def_i(t\downarrow_{l_i}))$ and $b = \Sigma_{t'|t'\downarrow_l = t\downarrow_l, t' \neq t}(\Pi_i def_i(t'\downarrow_{l_i}))$. In fact, by replacing $a \times (a + b)$ with a, we get $def_{post}(t) = (\Pi_i def_i(t\downarrow_{l_i})) \times def(t\downarrow_l)$, which coincides with $def_{pre}(t)$. □

DEFINITION 3.1.10 (STABLE PROBLEM) *Given an SCSP* P *and a set* R *of local consistency rules for* P, P *is said to be stable w.r.t.* R *if, for each* $r \in R$, $[r](P) = P$.

A local consistency algorithm consists of the application of several rules to the same problem, until stability of the problem w.r.t. all the rules. At each step of the algorithm, only one rule is applied. Thus the rules will be applied in a certain order, which we call a strategy.

DEFINITION 3.1.11 (STRATEGY) *Given a set* R *of local consistency rules for an SCSP, a strategy for* R *is an infinite sequence* $S \in R^\infty$. *A strategy* S *is fair if each rule of* R *occurs in* S *infinitely often.*

DEFINITION 3.1.12 (LOCAL CONSISTENCY ALGORITHM) *Given an SCSP* P*, a set* R *of local consistency rules for* P*, and a fair strategy* S *for* R*, a local consistency algorithm* $lc(P, R, S)$ *applies to* P *the rules in* R *in the order given by* S*. Thus, if the strategy is* $S = s_1 s_2 s_3 \ldots$*, the resulting problem is*

$$P' = [s_1; s_2; s_3; \ldots](P)$$

The algorithm stops when the current SCSP is stable w.r.t. R.

Note that this formulation of local consistency algorithms extends the usual one for k-consistency algorithms [Fre78], which can be seen as local consistency algorithms where all rules in R are of the form $l \leftarrow L$, where $L = \{l_1, \ldots, l_n\}$ and $| \bigcup_{i=1,\ldots,n} l_i | = | l | = k - 1$. That is, exactly $k - 1$ variables are involved in the locations[1] in L.

Arc-consistency [HDT92, Bes94] (AC) is an instance of local consistency where one propagates only the domains of the variables. With our notations, this type of algorithm deals only with rules of the form $\{x\} \leftarrow \{\{x, y_1, \ldots, y_n\}, \{y_1\}, \ldots, \{y_n\}\}$. In fact, an arc-consistency rule considers a constraint, say over variables $x, y_1, \ldots, y_n$, and all unary constraints over these variables, and combines all these constraints to get some information (by projecting) over one of the variables, say x[2]. Let us call AC this set of rules. In the following, we will write $ac(P)$ to denote $lc(P, AC, S)$.

When a local consistency algorithm terminates[3], the result is a new problem that has the graph structure of the initial one plus, possibly, new arcs representing newly introduced constraints (we recall that constraints posing no restrictions on the involved variables are usually not represented in the graph structure), but where the definition of some of the constraints has been changed. More precisely, assume $R = \{r_1, \ldots, r_n\}$, and $r_i = l_i \leftarrow L_i$ for all $i = 1, \ldots, n$. Then the algorithm uses, among others, n typed locations l_i of type con_i. Assume also that each of such typed locations has value def_i when the algorithm terminates. Consider also the set of constraints $C(R)$ that have type con_i. That is, $C(R) = \{\langle def, con \rangle$ such that $con = con_i$ for some i between 1 and $n\}$. Informally, these are the constraints that may have been modified (via the typed location mechanism) by the algorithm. Then, if the initial SCSP is $P = \langle C, con \rangle$, the resulting SCSP is $P' = lc(P, R, S) = \langle C', con \rangle$, where $C' = C - C(R) \cup (\bigcup_{i=1,\ldots,n} \langle def_i, con_i \rangle)$.

In classical CSPs, any local consistency algorithm enjoys some important properties. We now will study these same properties in the SCSP framework, and point out the corresponding properties of the semiring operations that are sufficient for them to hold. The desired properties are as follows:

1. any local consistency algorithm returns a problem which is equivalent to the given one;

[1] We recall that locations are just sets of variables.

[2] Actually, this is a generalized form of arc-consistency, since originally arc-consistency was defined for binary constraints only [Mac77].

[3] We will consider the termination issue later.

2. any local consistency algorithm terminates in a finite number of steps;

3. the strategy, if fair, used in a local consistency algorithm does not influence the resulting problem.

THEOREM 3.1.13 (EQUIVALENCE) *Consider an SCSP* P *and the problem* $P' = lc(P, R, S)$. *Then* $P \equiv P'$ *if the multiplicative operation of the semiring* ($\times$) *is idempotent.*

PROOF. By Definition 3.1.12, a local consistency algorithm is just a sequence of applications of local consistency rules. Since by Theorem 3.1.9 we know that each rule application returns an equivalent problem, by induction we can conclude that the problem obtained at the end of a local consistency algorithm is equivalent to the given one. □

THEOREM 3.1.14 (TERMINATION) *Consider any SCSP* $P = \langle C, con \rangle$ *over the constraint system* $CS = \langle S, D, V \rangle$ *and the set* $AD = \bigcup_{\langle def, con \rangle \in C} R(def)$, *where* $R(def) = \{a \mid \exists t$ *with* $def(t) = a\}$. *Then the application of a local consistency algorithm to* P *terminates in a finite number of steps if* AD *is contained in a set* I *which is finite and such that + and* $\times$ *are I-closed.*

PROOF. Each step of a local consistency algorithm may change the definition of one constraint by assigning a different value to some of its tuples. Such value is strictly worse (in terms of $\leq_S$) since $\times$ is intensive. Moreover, it can be a value which is not in AD but in $I - AD$. If the state of the computation consists of the definitions of all constraints, then at each step we get a strictly worse state (in terms of $\sqsubseteq_S$). The sequence of such computation states, until stability, has finite length, since, by assumption, I is finite and thus the value associated with each tuple of each constraint may be changed at most $|I|$ times. □

An interesting special case of the above theorem occurs when the chosen semiring has a finite carrier set A. In fact, in that case the hypotheses of the theorem hold with $I = A$.

COROLLARY 3.1.15 *Consider any SCSP* P *over the constraint system* $CS = \langle S, D, V \rangle$, *where* $S = \langle A, +, \times, 0, 1 \rangle$, *and* A *is finite. Then the application of a local consistency algorithm to* P *terminates in a finite number of steps.*

PROOF. Easily holds with $I = AD$ by Theorem 3.1.14. □

We will now prove that, if $\times$ is idempotent, no matter which strategy is used during a local consistency algorithm, the result is always the same problem.

THEOREM 3.1.16 (ORDER-INDEPENDENCE) *Consider an SCSP* P *and two different applications of the same local consistency algorithm to* P, *producing respectively the SCSPs* $P' = lc(P, R, S)$ *and* $P'' = (P, R, S')$. *Then* $P' = P''$ *if the multiplicative operation of the semiring* ($\times$) *is idempotent.*

PROOF. Each step of the local consistency algorithm, which applies one of the local consistency rules in R, say r, were defined as the application of a function $[r]$ that takes an SCSP P and returns another problem $P' = [r](P)$, and that may change the definition of the constraint connecting the variables in l (if the applied rule is $r = l \leftarrow L$). Thus the whole algorithm may be seen as the repetitive application of the functions $[r]$, for all $r \in R$, until no more change can be done. If we can prove that each $[r]$ is a closure operator [DP90], then classical results on chaotic iteration [Cou77] allow us to conclude that the problem resulting from the whole algorithm does not depend on the order in which such functions are applied. Closure operators are just functions which are idempotent, monotone, and intensive. We will now prove that each $[r]$ enjoys such properties. We remind that $[r]$ just combines a constraint with the combination of other constraints.

If $\times$ is idempotent, then $[r]$ is idempotent as well, that is, $[r]([r](P)) = [r](P)$ for any P. In fact, combining the same constraint more than once does not change the problem by idempotency of $\times$. Also, the monotonicity of $\times$ (see Theorem 2.1.4) implies that of $[r]$, that is, $P \sqsubseteq P'$ implies $[r](P) \sqsubseteq [r](P')$. Note that, although $\sqsubseteq$ has been defined only between constraints, here we use it among constraint problems, meaning that such a relationship holds among each pair of corresponding constraints in P and P'. Finally, the intensivity of $\times$ implies that $[r]$ is intensive as well, that is, $[r](P) \sqsubseteq P$. In fact, P and $[r](P)$ just differ in one constraint, which, in $[r](P)$, is just the corresponding constraint of P combined with some other constraints. $\square$

Note that the above theorems require that $\times$ is idempotent for a local consistency algorithm to be meaningful. There is, however, no requirement over the nature of $\leq_S$. More precisely, $\leq_S$ can be partial. This means that the semiring operations $+$ and $\times$ can be different from max and min respectively (see note at the end of Section 2.1).

Note also that, by definition of rule application, constraint definitions are changed by a local consistency algorithm in a very specific way. In fact, even in the general case in which the ordering $\leq_S$ is partial, the new values assigned to tuples are always smaller than or equal to the old ones in the partial order. In other words, local consistency algorithms do not “jump” in the partially ordered structure from one value to another one that is unrelated to the first one. More precisely, the following theorem holds.

THEOREM 3.1.17 (LOCAL CONSISTENCY AND PARTIAL ORDERS) *Given an SCSP P, consider any value v assigned to a tuple in a constraint of such problem. Also, given any set R of local consistency rules and strategy S, consider $P' = lc(P, R, S)$, and the value v' assigned to the same tuple of the same constraint in P'. Then we have that $v' \leq_S v$.*

PROOF. By definition of rule application (see Definition 3.1.8) the formula defining the new value associated to a tuple of a constraint can be assimilated to $(v \times a) + (v \times b)$, where v is the old value for that tuple and a and b are combinations of values for tuples of other constraints in the graph of the rule. Now, we have $(v \times a) + (v \times b) = v + (a \times b)$ by distributivity of $+$ over $\times$ (see Theorem 2.1.10). Also, $(v \times c) \leq_S v$ for any c by intensivity of $\times$ (see Theorem 2.1.5), thus $v + (a \times b) \leq_S v$. $\square$

One could imagine other generalizations with the same desired properties as the ones proven above (that is, termination, equivalence, and order-independence). For example, one could design an algorithm which avoids (or compensates) the effect of cumulation coming from a non-idempotent $\times$. Or also, one could generalize the equivalence property, by allowing the algorithm to return a non-equivalent problem (but, nonetheless, being in a certain relationship with the original problem). Finally, the order-independence property states that any order is fine, as far as the resulting problem is concerned. Another desirable property, related to this, could be the existence of an ordering that makes the algorithm belong to a specific complexity class.

3.2 Applying local consistency to the instances

3.2.1 Classical CSPs

Let us now consider the application of a local consistency algorithm to CSPs. As predictable, we will show that all the classical properties hold. First, $\wedge$ is idempotent. Thus, by Theorem 3.1.13, the problem returned by a local consistency algorithm is equivalent to the given one, and from Theorem 3.1.16, the used strategy does not matter. Also, since the domain of the semiring is finite, by Corollary 3.1.15 any local consistency algorithm terminates in a finite number of steps.

3.2.2 Fuzzy CSPs

Let us now consider the properties that hold in the FCSP framework. The multiplicative operation (that is, min) is idempotent. Thus, local consistency algorithms on FCSPs do not change the solution of the given problem (by Theorem 3.1.13), and do not care about the strategy (by Theorem 3.1.16). Moreover, min is AD-closed for any finite subset AD of [0,1]. Thus, by Theorem 3.1.14, any local consistency algorithm on FCSPs terminates.

Thus FCSPs, although providing a significant extension to classical CSPs, can exploit the same kind of algorithms. Their complexity will, of course, be different due to the presence of more than two levels of preference. It can be proven. however, that if the actually used levels of preference are p, then the complexity of a local consistency algorithm is just $O(p)$ times greater than that of the corresponding algorithm over CSPs (as also discussed above and in [DFP93]).

An implementation of arc-consistency, suitably adapted to be used over fuzzy CSPs, is given in [Sch92]. No formal properties of its behavior, however, are proven there. Thus our result can be seen as a formal justification of [Sch92].

3.2.3 Probabilistic CSPs

The multiplicative operation of S_{prob} (that is, $\times$) is not idempotent. Thus the result of Theorem 3.1.13 cannot be applied to Prob-CSPs. That is, by applying a local consistency algorithm one is not guaranteed to obtain an equivalent problem. Theorem 3.1.16 cannot be applied as well. Therefore the strategy used by a local consistency algorithm may matter. Also, $\times$ is not closed on any finite superset of any subset of $[0, 1]$. Hence the result of Theorem 3.1.14 cannot be applied to Prob-CSPS. That is, we cannot be sure that a local consistency algorithm terminates.

As a result of these observations, local consistency algorithms should make no sense in the Prob-CSP framework. As a result, we are not sure that their application has the desired properties (termination, equivalence, and order-independence). However, the fact that we are dealing with a c-semiring implies that, at least, we can apply Theorem 3.1.3: if a Prob-CSP problem has a tuple with probability α to be a solution of the real problem, then any subproblem has a tuple with probability at least α to be a solution of a subproblem of the real problem. This can be fruitfully used when searching for the best solution. In fact, if one employs a branch-and-bound search algorithm, and in one branch we find a partial instantiation with probability smaller than α, then we can be sure that such a branch will never lead to a global instantiation with probability α. Thus, if a global instantiation with such a probability has already been found, we can cut this branch. Similar (and possibly better) versions of branch-and-bound for non-standard CSPs may be found in [FW92, SFV95].

3.2.4 Weighted CSPs

The multiplicative operation of S_{WCSP} (that is, +) is not idempotent. Consequently the results of Section 3.1 on local consistency algorithms cannot be applied to WCSPs. As in Prob-CSPs, however, the fact that we are dealing with a c-semiring implies that, at least, we can apply Theorem 3.1.3: if a WCSP problem has a best solution with cost α, then the best solution of any subproblem has a cost smaller than α. This can be fruitfully used when searching for the best solution in a branch-and-bound search algorithm.

Note that the same properties hold also for the semirings $\langle \mathcal{Q}^+, \mathtt{min}, +, +\infty, 0 \rangle$ and $\langle \mathcal{Z}^+, \mathtt{min}, +, +\infty, 0 \rangle$ (which can be proved to be c-semirings).

3.2.5 Egalitarianism and Utilitarianism

One can show that S_{ue} is a c-semiring. However, since it uses a combination of the operations of the semirings S_{FCSP} and S_{WCSP}, and since the multiplicative operation of S_{WCSP} is not idempotent, also the multiplicative operation of S_{ue}, that is, $\underline{\mathtt{min}}$, is not idempotent. This means that we cannot guarantee that local consistency algorithms can be used meaningfully in this instance of the framework.

Note that also the opposite choice (that is, first perform a max-sum and then a max-min) can be made, by using a similar c-semiring. Due however to the same reasoning as above, again the multiplicative operation would not be idempotent, thus we cannot be sure that the local consistency techniques possess the properties of Section 3.1.

3.2.6 Set-based SCSPs

The set-based SCSPs are represented by the c-semiring

$$S_{set} = \langle \wp(A), \bigcup, \bigcap, \emptyset, A \rangle$$

where A is any set. It is easy to see that S_{set} is a c-semiring. Also, in this case the order $\leq_{S_{set}}$ reduces to set inclusion (in fact, $a \leq b$ iff $a \cup b = b$), and, therefore, it is partial in general. Furthermore, $\times$ is $\bigcap$ in this case, and thus it is idempotent. Therefore, the local consistency algorithms possess all the properties stated in Section 3.1 and, therefore, can be applied. We recall that when a local consistency algorithm is applied to an SCSP with a partially ordered semiring, its application changes the constraints such that the new constraints are smaller than the old ones in the ordering $\sqsubseteq_S$.

3.3 About arc-consistency in SCSPs

In this section we study the properties and possible advantages of the extension to SCSPs of the most successful of the local consistency techniques, called *arc-consistency* [Wal75, Mac77], although we claim that all our results can be extended to higher levels of consistency. What arc-consistency (AC) does over a CSP is considering each pair of variables and checking whether there are values for one of these variables, which are inconsistent with all the values for the other one. If so, such inconsistent values are deleted from the CSP, since they cannot participate in any solution. Semiring-based arc-consistency (SAC) works similarly over SCSPs, considering, however, also the semiring values, and the final effect is not a deletion, but a worsening of the semiring value assigned to some domain element of a variable.

Here we study the relationship between AC and SAC over SCSPs, showing that they discover the same inconsistencies. That is, the domain elements whose semiring value gets replaced by the worst value (that is, the minimum of the semiring) are exactly those elements which are deleted by AC in the corresponding CSP, obtained by the given SCSP by forgetting the semiring values. Thus sometimes it may be reasonable to apply just a classical AC instead of SAC, since AC is more efficient, although less informative. This can lead to an original methodology where, instead of preprocessing a given SCSP via SAC and then solving it, we can first obtain the corresponding CSP, then preprocess this CSP via AC (or some other form of local consistency), and then bring back to the initial SCSP the information obtained by AC over the inconsistencies. In particular, if AC discovered an inconsistency, we can safely assign to such an element the worst semiring value. Note that we could also make this method more flexible, by replacing the given SCSP not with a CSP but with a simpler SCSP, that is, an SCSP over a simpler semiring. In this way, there could be a whole range of possible simplifications, in the line of [SdGS97], and the results of SAC over the simpler problem could be brought back to the given problem via established techniques like abstract interpretation [CC92]. In Chapter 4 we will show how the usual Galois connection of abstract interpretation is used to pass from an SCSP to an "easier" one.

Moreover, we consider SCSPs with visible and hidden variables, and we show that, if such SCSPs are SAC, then some hidden variables can be removed without changing the set of solutions and associated semiring values. In particular, we show that all hidden variables of degree 1 can be removed. The reason behind this result is that, by achieving SAC, all the information provided by the variables of degree 1 is put into the other (connected) variables. If such variables are not hidden, then we can still gain from the application of SAC, because they can be instantiated without the need for any backtracking. Moreover, since in a tree all leaves have degree 1, and the other variables get degree 1 once the leaves are removed, tree-shaped SCSPs can be solved without backtracking, provided that we follow an instantiation order that visits the tree top-down. For SCSPs which are not tree-shaped, we can still use this result, on the parts which are tree-shaped. Therefore, such SCSPs can be solved by searching (possibly with backtracking) over the part with cycles, and then by continuing the instantiation process over the tree-shaped parts without backtracking.

These results are extensions of similar results already existing for the CSP case [Ros95, Fre88]; however, here they are particularly interesting both because of the generality of the SCSP framework, and also because SCSPs usually express constraint optimization, and not just satisfaction, problems.

The Section is organized as follows. Subsection 3.3.1 compares the power of classical arc-consistency with that of semiring-based arc-consistency. Then, Subsection 3.3.2 studies the possibility of removing hidden variables in an arc-consistent SCSP, subsection 3.3.3 shows that tree-shaped SCSPs can be solved easily, and Section 3.3.4 extends this result to SCSPs with cycles with bounded size.

3.3.1 Classical arc-consistency vs. semiring-based one

Consider an SCSP $\mathtt{P} = \langle \mathtt{C}, \mathtt{con} \rangle$. Let us now construct a corresponding CSP $\mathtt{cut(P)}$ by simply forgetting the semiring values associated with all the tuples in all constraints of $\mathtt{P}$, except for those values which are the minimum of the semiring. That is, if a tuple has a semiring value different from the minimum we say that such tuple satisfies the constraint in $\mathtt{cut(P)}$, otherwise we say that it does not satisfy it. For example, the CSP in Figure 3.1 is the CSP corresponding to the FCSP of Figure 3.2 according to function $\mathtt{cl}$.

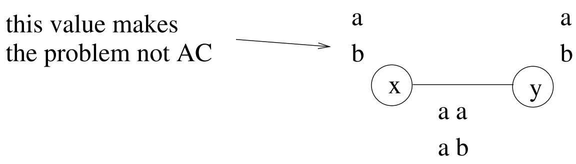

Figure 3.1: A CSP which is not AC.

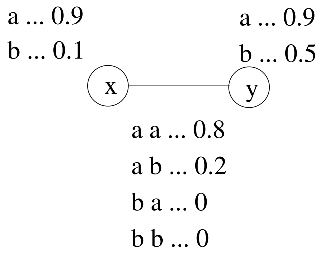

Figure 3.2: A fuzzy CSP.

Then, let us apply SAC over P and AC over $\mathtt{cut}(P)$. As a result of that, some element in the domain of some variables in P will get a smaller semiring value, possibly the minimum of the semiring, and some domain elements of some variables in $\mathtt{cut}(P)$ will be eliminated. Then, the main result is the following one:

THEOREM 3.3.1 (SAC VS. AC) *In an SAC-consistent SCSP P, if the order $\leq_S$ induced by the $+$ operation is total, then the domain elements which will get the minimum of the semiring in P by applying an SAC algorithm are exactly those that will be eliminated in $\mathtt{cut}(P)$ by applying an AC algorithm.*

In fact, consider what happens during the SAC algorithm: for each pair of variables x and y, to see which value has to be given to the domain element t_x for x, that now has value a_x, we have to consider also all the pairs t_{xy} that extend t_x over y. Let us assume there are n of them, and let us call them $t^1_{xy}, \ldots, t^n_{xy}$. Assume also that each t^i_{xy} is given the semiring value a^i_{xy}. Finally, consider all domain elements of y, say $t^1_y, \ldots, t^n_y$, and assume that they now have the semiring values $a^1_y, \ldots, a^n_y$. What we have to do now to perform SAC over this constraint is to compute the following: $\sum_{i=1,\ldots,n}(a_x \times a^i_{xy} \times a^i_y)$. This is the semiring value to be associated with t_x. That is, if t_x has a different semiring value, the problem is not SAC-consistent.

Now the question is: when does this value coincide with the minimum of the semiring? In general we don't know. We know, however, that if the $\leq_S$ order is total, then $a \times b$ is always equal to either a or b, for any a, b in the semiring, and the same is for $a + b$. In fact, the glb and the lub of two elements in a total order is always one of the two elements.

Thus, the only way that the above formula coincides with the minimum of the semiring is such a minimum appears somewhere in the formula. That is, either a_x, or a^i_{xy} for some i, or a^i_y for some i, is already the minimum of the semiring. In other words, if none of these values coincides with the minimum of the semiring, then we can be sure that the result of the above formula does not coincide with the minimum of the semiring.

Now, by the way we have constructed $\mathtt{cut}(P)$, tuples that have the minimum of the semiring do not appear in $\mathtt{cut}(P)$. That is, they do not satisfy the corresponding constraint. Therefore, by the way AC works, they will create a local inconsistency and they will cause the deletion of an element from the domain of a variable. More precisely, in $\mathtt{cut}(P)$, to see whether element t_x of the domain of x has to be eliminated, we have to look at all pairs t_{xy} that extend t_x over y and that satisfy the constraint between x and y, and for each of such pairs we have to see whether its second element (that is, its projection over y) is in the domain of y or not. If all such elements are not in the domain of y, then t_x has to be eliminated by the domain of x. This means that for all pairs t_{xy} that extend t_x over y, either t_{xy} does not satisfy the constraint between x and y, or the projection of t_{xy} onto y is not in the domain of y. But if this is so, then, by the way $\mathtt{cut}(P)$ was constructed, the element t_x in P will get the minimum of the semiring by SAC. The opposite is also easily provable: if t_x gets the minimum of the semiring in P via SAC, then it will be eliminated in $\mathtt{cut}(P)$ via AC.

Figure 3.3 shows the FCSP obtained by applying SAC onto the FCSP of Figure 3.2 and the CSP obtained by applying AC onto the CSP of Figure 3.1.

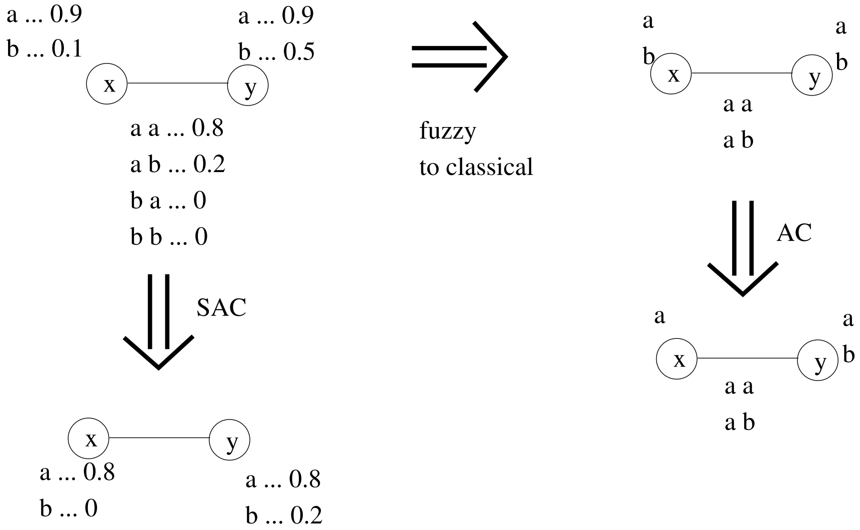

Figure 3.3: An SAC-consistent FCSP and the corresponding AC-consistent CSP.

Why is this result interesting? Because applying AC over $\mathtt{cut}(P)$ is more efficient, both in space and in time, than applying SAC over P. In fact, domains could have a more compact representation, and each step of the algorithm could deal with a smaller set of tuples (only those that have semiring value *true*).

Therefore, one could use AC instead of SAC if he/she is interested only in discovering local inconsistencies. Of course some information will be lost, because the semiring values that are present in the problem after AC and which are different from the minimum are not as accurate, but in some cases this could be a reasonable trade-off between pruning power and space/time efficiency.

In fact, this approach could also be seen as the extreme case of a flexible abstraction technique, where, given an SCSP over a certain semiring S, we replace SAC over such a semiring with SAC over a simpler semiring, not necessarily S_{CSP} (as we proposed in this section) which is the simplest semiring of all. More details in this direction are given in Chapter 4 and in [BCGR00a, BCR00] where the usual Galois connection of abstract interpretation is used to pass from an SCSP to an "easier" one.

3.3.2 Removing hidden variables in SCSPs

An SCSP, as defined above, is a pair $\langle C, \mathtt{con} \rangle$, where $\mathtt{con}$ is a subset of the variables, and specifies the variables of interest as far as solutions are concerned. Variables in $\mathtt{con}$ can then be called the *visible* variables, while all the others are *hidden*.

Since the solution set of such problems is defined as $(\otimes C) \Downarrow_{\mathtt{con}}$, this means that a solution is an assignment of the visible variables, with an associated semiring value, such that it can be extended to the hidden variables in a way that the overall assignment (to all the variables) has the same semiring value. Since, however, we do not really care to know the values for the hidden variables, but just that there is a way to instantiate them, in some cases some of these variables can be found to be redundant. This means that their elimination would not change the solutions of the overall problem. This, of course, can be very convenient in terms of the search for one or the best solution, because the depth of the search tree coincides with the number of variables. In particular, this section shows the following:

THEOREM 3.3.2 (HIDDEN VARIABLE REMOVAL) *If we have an SCSP which is SAC-consistent, and if the order $\leq_S$ is total, then all hidden variables with degree 1 (or 0) are redundant and thus can be eliminated, without changing the set of solution-value pairs for the problem.*

This is an extension to SCSP of a result which has been shown already for CSPs [Ros95]. For example, Figure 3.4 shows the case of a CSP where, assuming that the problem is AC, variable v is redundant. Here is, however, more interesting, since we are dealing with optimization and not satisfaction problems.

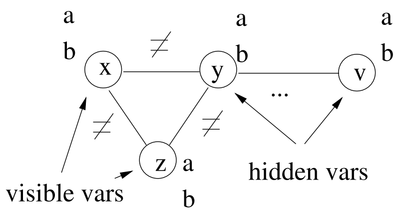

Figure 3.4: Redundant hidden variables in CSPs.

Let us now prove the statement of the above theorem. Consider an SCSP that is SAC-consistent, and assume that there is a hidden variable, say z, with degree 1. Figure 3.5 shows an example of such an SCSP.

What we have to prove is that, given an instantiation of x and y, which satisfies the constraints between them, for example $\langle x = a, y = a \rangle$, which has value 0.8, there is an extension to z, say d, such that $\langle x = a, y = a, z = d \rangle$ has value 0.8 and all other extensions have a worse value. In fact, if this is the case, then we can safely delete z without changing the solution set.

More generally, let us consider an instantiation of the visible variables with a corresponding semiring value, which is given by combining all constraints involving subsets of

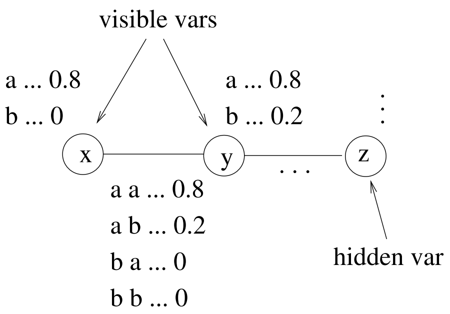

Figure 3.5: Redundant hidden variables in SCSPs.

the visible variables: say it is $\langle d_1, \ldots, d_n \rangle$ (there are n visible variables) with semiring value v. To show that z is redundant, we have to show that such a tuple can be extended to z without making its semiring value worse. In other words, there is a value for z, say d_z, such that the tuple $\langle d_1, \ldots, d_n, d_z \rangle$ has semiring value v.

To compute the semiring value for this tuple, once we have the semiring value v for the shorter tuple, we need to multiply (by using the $\times$ operator of the semiring) v by the semiring value given by c_{xz} (we assume that x is the i-th visible variable) to $\langle x = d_i, z = v \rangle$, say v_{xz}, and also by the semiring value given by c_z to $z = v$, say v_z. Notice that if z is indeed connected to x, this is the only constraint connected to z, otherwise we will just consider a null constraint which assigns the maximum of the semiring to each tuple. Thus, we get $v \times v_{xz} \times v_z$, and we have to prove that this coincides with v. Figure 3.6 shows the three components of this new value for the example in Figure 3.5, to be compared to v.

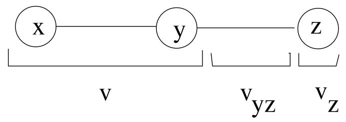

Figure 3.6: New semiring value for the extended tuple.

Now, by the extensivity of $\times$ (that is, $a \times b \leq a$ for any a, b), we know that $(v \times v_{xz} \times v_z) \leq v$. We will now prove that this disequality is in fact an equality, that is, $(v \times v_{xz} \times v_z) = v$. We know that any SAC algorithm enforces the formula $c_x = (\otimes\{c_y, c_{xy}, c_x\}) \Downarrow_x$ for each pair of variables x and y. Thus, this also applies to the pair of variables x and z of our example. Therefore, the semiring values associated with elements of the domain of x, such as d_i, have been computed during the SAC algorithm by also considering c_{xz} and c_z. In other words, the information that we are now adding to v, that is, $v_{xz} \times v_z$, is already in v, because of SAC. Of course there are many ways to instantiate z, but if the order $\leq_S$ is total, one of them, that is, the best one, has been inserted in v. By choosing that value for z, by the idempotency of the $\times$ operation, which has to be assumed if we want to be sure that local consistency has the desired properties, we get that $v = v \times v_{xz} \times v_z$. Therefore, there exists a way to extend $\langle d_1, \ldots, d_n \rangle$ to z such that its semiring value is not changed.

Notice also that once some variables are eliminated the degree of other variables may be reduced and thus satisfy the redundancy condition. Thus these other variables can be eliminated as well, until no variable has degree 1.

Notice also that it is not necessary to achieve SAC all over the problem to eliminate some hidden variable, but that it is enough to achieve it on the neighborhood (just two variables) of the variable to eliminate. Also, in such a neighborhood, SAC could be achieved just in one direction, because we have to make sure that a solution involving the visible variables can be extended to the hidden variables, but we do not care about the opposite direction.

Again, we claim that the result of this section can easily be generalized to a higher level of consistency, allowing the elimination of variables of higher degree. That is, the generalization should say that, if an SCSP is k-consistent, then all hidden variables with degree up to $k - 1$ are redundant.

3.3.3 Tree-shaped SCSPs

The result of the previous section says that, if we have an SCSP with a total order over the semiring values, and if the SCSP is SAC-consistent, then any tuple involving the visible variables can be extended to any hidden variable with degree 1 without changing the semiring value of the tuple.

Consider now the same situation, except that the variable with degree 1 is not hidden, but visible. Still, the extension works just the same, and the only thing that is different here is that we have to constructively find that extension, and we cannot be satisfied with just knowing that it exists. Let us see what we have to do to find the (or *a*) domain element for z (the variable with degree 1), which does not change the semiring value of the given tuple. Since the difference between the semiring value of the new tuple and that of the old one depends only on the constraint connecting z, say c_{xz}, and on the domain of z, say c_z, we just have to look at all pairs of values for x and z such that $x = d_i$ (that is, the value that x has in the short tuple), and for each of such pairs we have to multiply the value given by c_{xz} by that given by c_z. Now, the best of such values is the one to choose. That is, we choose $z = d_z$ if $c_{xz}(d_i, d_z) \times c_z(d_z)$ is the best among all $c_{xz}(d_i, d'_z) \times c_z(d'_z)$ for all d'_z other values for z (where given a constraint $c = \langle def, con \rangle$, $c(d)$ means the application of the definition function to the domain value d4, that is the value $def(d)$).

This operation is linear in the domain size, since for each element of the domain of z, say d_z, we just need to compute $c_{xz}(d_i, d_z) \times c_z(d_z)$. More important, if the tuple we started with (that is, $\langle d_1, \dots, d_n \rangle$) was the best one for the first n variables, then the new tuple, $\langle d_1, \dots, d_n, d_z \rangle$ is still the best one for the n initial variables plus z. In fact, it is impossible to find a tuple which has a better semiring value for such $n + 1$ variables. Otherwise the assumption that the initial tuple was the best one for the n variables will be contradicted.

This discussion leads to the following theorem, which again is an extension of a result already known for classical CSPs [Fre88]:

THEOREM 3.3.3 (LINEAR TIME FOR TREES) *If we have an SCSP with a total order* *over the semiring values, and such an SCSP is SAC-consistent and tree-shaped, then the optimal solution can be found in time linear in the number of its variables.*

In fact, in a tree-shaped SCSP, all variables have either degree 1 (the leaves) or can get degree 1 when the current leaves are eliminated. Therefore, the extension result just discussed can be applied at every step of the search, provided that we choose an instantiation order that visits the tree from top to bottom. Figure 3.7 shows a tree-shaped SCSP and a possible top-down instantiation order.

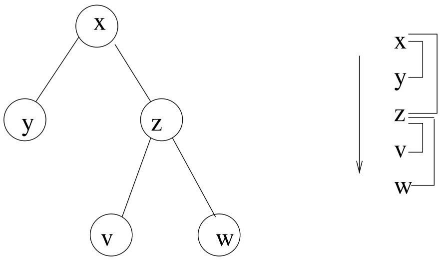

Figure 3.7: A tree-shaped SCSP and a top-down instantiation order.

That is, at every step during the search, we have a partial solution involving only the previous variables, and we have to extend it to the next variable. By what we said before, this extension can be done without changing the semiring value of the solution, and also without the need to come back to check whether there is a better solution. Therefore, a complete best solution can be found without backtracking.

An interesting remark is that the semiring value for the best solution can be found without even finding the best solution. In fact, since each step does not change the semiring value of the current incomplete solution, the best value of a complete solution coincides with the best value for the first variable in the order. In fact, by choosing the best value for the first variable, we know that we can complete the solution maintaining that value.

PROPOSITION 3.3.4 (BEST VALUE) *Given an SCSP with a total order over the semiring values, which is SAC-consistent and tree-shaped, the semiring value for the optimal solutions is the best value among those assigned to the domain elements of the first variable in a top-down instantiation order for this problem.*

Even for this result we can easily see that SAC is not really necessary, we just need a directional version of SAC that goes bottom-up in the tree.

3.3.4 Cycle-cutsets in SCSPs

By stretching the discussion of the previous section even more, we can consider SCSPs, which in general do not have a tree shape, where, however, we identify in some way the part of the SCSP that creates the cycles. In such problems, we can first find the best solution for the difficult part (the one with cycles), and after that we can complete the solution on the rest of the variables, which is tree-shaped. Therefore, the overall search for the best solution will possibly backtrack over the first part of the problem, to find the best solution for that part, but then will not backtrack at all on the rest of the problem. Figure 3.8 shows an SCSP with a cycle over variables x, y, and z, and then a tree shape over the rest of the problem (not completely drawn), and Figure 3.9 shows the search tree for this SCSP.

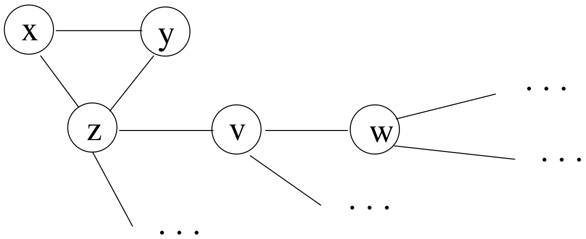

Figure 3.8: An SCSP with a cycle over x, y, z and a tree shape over the other variables.

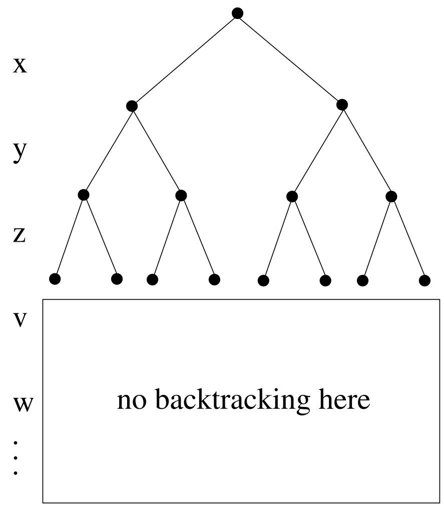

Figure 3.9: The search tree for the SCSP of Figure 3.8.

In classical constraint solving this technique is called *cycle-cutset* [DP88a], since the first part of the problem (that possibly involves some backtracking) is identified by a subset of the variables whose elimination allows to remove all cycles of the problem. Here we can use the same technique and get similar results, which, however, as noted before, deal with optimization instead of satisfaction. Therefore, here we can say something about the time complexity to find not *a solution*, but the *best solution*. Moreover, as in the previous section, we have an additional result which allows us to know the semiring value associated with the best solution *before* we actually compute the solution itself. In fact, once all variables in the cycle-cutset have been instantiated and we have found the best solution for them, the semiring value for this partial solution will coincide with the semiring value for the best complete solution.

More interesting results can be obtained if one applies the SAC algorithm not only before the search but also *during* the search, after each instantiation. In fact, instantiating a variable is like eliminating that variable for the problem to be solved, and thus new variables could become of degree 1 with respect to the set considered at the beginning of the search. By maintaining SAC at each step during the search, we could, therefore, exploit the presence of this variables to dynamically augment the set of variables over which the search does not need any backtracking. Again, this idea has been generated by recent results on classical CSPs [SF97], which, however, we claim can be extended to SCSPs as well.

3.4 Labeling and Partial Local Consistency for SCSPs

The aim of this section is to extend the local consistency algorithms, defined in the previous sections for SCSPs, in order to capture some approximated algorithms which have been proven useful for classical constraints, so much so that they are widely used in current constraint programming systems like `clp(FD)` [CD96], Ilog Solver [Pug94] and CHIP[AB93].

The SCSP framework comes with a class of local consistency (LC) algorithms, which are useful to perform a simplification of the problem, cutting away some tuples of domain values that are shown to be useless w.r.t. the solutions of the problem.

This scheme is in theory very useful but experiments show that even the simplest nontrivial LC algorithm (which corresponds to Arc-Consistency [HDT92, Bes94] extended to SCSPs) has a heavy complexity for most applications. Moreover, one has to apply it several times if used during the labeling phase. Hence, the reduction of its complexity is a crucial point.

To go in this direction, we define a class of Partial Local Consistency (PLC) algorithms, which perform less pruning than (complete) local consistency but can nevertheless be rather useful in reducing the search. Some instances of these algorithms are able to reduce the domains of the variables in constant time, and can work with a convenient representation of the domains of the variables. Using PLC together with labeling, the aim is to find the solutions of a problem with a better practical complexity bound.

From this point of view, our work is related to the one of K. Apt [Apt98]. In fact, he studied a class of filtering functions for classical CSPs, which can express both local consistency and also less pruning notions.

Summarizing, here are our results that appear also in [BCGR00b]:

1. the extension of the definitions and properties of the labeling procedure from classical CSPs to SCSPs;

2. a general scheme for approximating local consistency in SCSPs, and properties that are sufficient for the correctness of the approximations;

3. some implementations issues in the particular case of arc-consistency;

4. experimental results for partial arc-consistency.

The rest of the section is organized as follows. In Section 3.4.1 we introduce some properties about labelings, we describe their interaction with usual filtering algorithms, and we define the generalization of the labeling procedure to SCSPs. Then, we introduce the notion of rule approximations and study the correctness of such rules in Section 3.4.2 and we discuss some implementations issues in Section 3.4.3.

Another operation. Besides combination and projection, we need, in this section, another operator, which we will call the *disjunction* operator. Given two constraints $c_1 = \langle def_1, con_1 \rangle$ and $c_2 = \langle def_2, con_2 \rangle$ over CS, their *disjunction* $c_1 \oplus c_2$ is the constraint $c = \langle def, con \rangle$ with $con = con_1 \cup con_2$ and $def(t) = def_1(t \downarrow^{con}_{con_1}) + def_2(t \downarrow^{con}_{con_2})$.

The informal meaning of the disjunction operator is to give more possibility to (or to enhance the level of preference of) certain instantiations. If we consider the semiring $\mathtt{Bool}$ which corresponds to classical CSPs, the meaning of having the disjunction of two constraints C_1 and C_2 is to have the possibility of choosing not only the tuples permitted by C_1 but also those permitted by C_2. Figure 3.10 shows an example using the semiring Fuzzy, where, we recall, the operation $+$ is the maximum.

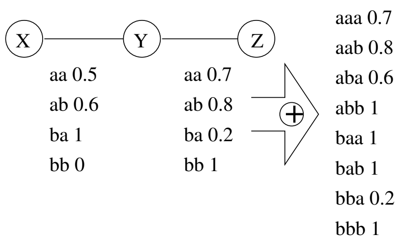

Figure 3.10: Disjunction example.

Using the properties of $\times$ and $+$, it is easy to prove that: $\oplus$ is associative, commutative and idempotent; $\otimes$ distributes over $\oplus$; $\oplus$ is monotone over $\sqsubseteq_S$. Moreover, if $\times$ is idempotent: $\oplus$ distributes over $\otimes$.

Let us also give some new definition that will be useful in the following of this section.

DEFINITION 3.4.1 (CONSTRAINT TYPE) *Given a constraint* $c = \langle def, con \rangle$, *we call* $type(c)$ *the set* con *of the variables of* c. *Given a constraint set* C, $type(C) = \{type(c)/c \in C\}$ *and* $V(C) = \bigcup_{c \in C} type(c)$ *is the set of all the variables appearing in the constraints of* C.

3.4.1 Labeling in SCSPs

Local consistency algorithms, when applied to an SCSP P, can reduce the search space to find its solutions. Solving the resulting problem P', however, is still NP-hard. This scenario is similar to what happens in the classical CSP case, where, after applying a local consistency algorithm, the solution is found by replacing P' with a set of subproblems $\{P_1, \ldots, P_n\}$ where some of the variables have been instantiated such that P' is "equivalent" to $\{P_1, \ldots, P_n\}$, in the sense that the set of solutions of P' coincides with the union of the sets of solutions of $P_1, \ldots, P_n$. Usually, problems P_i are obtained by choosing a variable, say x, and instantiating it to its i-th domain value. Then, the local consistency (usually arc-consistency) algorithm is applied again [Hen89]. By doing this, one hopes to detect the inconsistency of some of the subproblems by detecting their local inconsistency (via the local consistency algorithm). When all the variables of the initial problem are instantiated, arc-consistency becomes complete in the sense that if the problem is not consistent then it is not arc-consistent. Therefore, at that point arc-consistency implies global consistency.

Let us consider, for example, the following problem:

$$X \in [1,3], Y \in [1,2], Z \in [1,2], X \neq Y, Y \neq Z, Z \neq X.$$

Figure 3.11 shows the interaction between AC and the labeling procedure, since each node of the search tree (except the root) shows the CSP obtained after the application of an arc-consistency algorithm. In this particular case we have two solutions (the dashed line indicates the second solution found by the labeling).

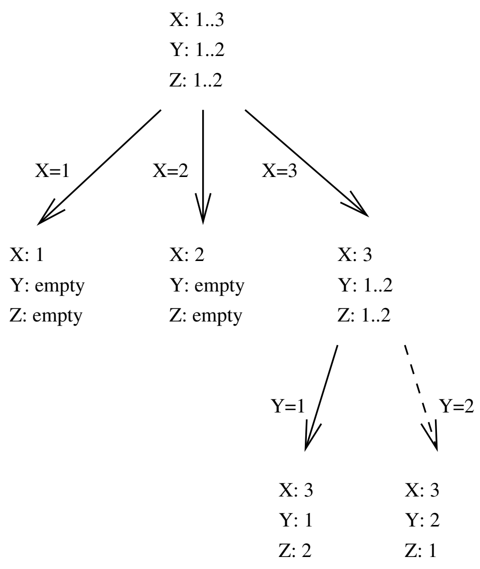

Figure 3.11: AC and labeling.

We will now study the properties of the labeling procedure when applied to SCSPs.

In the following we assume to work with a constraint system $CS = \langle S, D, V \rangle$, where $S = \langle A, +, \times, \mathbf{0}, \mathbf{1} \rangle$.

Given a set $\mathcal{P} = \{P_1, \cdots, P_n\}$ of SCSPs, the *solution of* $\mathcal{P}$ is the constraint defined as $Sol(\mathcal{P}) = Sol(P_1) \oplus \cdots \oplus Sol(P_n)$.

Given $x \in V$, $d \in D$ and $v \in A$, we call $c_{x,d,v}$ the unary constraint $\langle \mathrm{def}, \{x\}\rangle$ where $\mathrm{def}(d) = v$ and $\mathrm{def}(d') = \mathbf{0}$ if $d' \neq d$. In practice, $c_{x,d,v}$ instantiates x to d with semiring value v. We call this constraint a *simple instantiation constraint*.

Given a total order $\prec$ on V, let $W = \{x_1, \ldots, x_n\}$ be a subset of V, and assume that $x_1 \prec \cdots \prec x_n$. Given an n-tuple of domain values $d = \langle d_1, \ldots, d_n\rangle$ and an n-tuple of semiring values $v = \langle v_1, \ldots, v_n\rangle$, we define: $I_W^{d,v} = \{c_{x_i,d_i,v_i} \mid i \in [1,n]\}$. We write I_W^d for $I_W^{d,\mathbf{1}}$, and we call it a W*-instantiation set*. In practice, I_W^d gives the best semiring value ($\mathbf{1}$) to the assignment of d_i to variable x_i. Instead, $I_W^{d,v}$ gives semiring value v_i to the assignment of d_i to variable x_i.

Given a problem $P = \langle C, \mathrm{con}\rangle$, let W be a subset of $V(C)$ and t be a $|W|$-tuple. We call a W*-labeling* of the problem P the problem $P_W^t = \langle C \cup I_W^t, \mathrm{con}\rangle$. In words, a W*-labeling* of P adds to P some additional constraints which identify a specific partial labeling for P (partial because it involves only the variables in W). A *complete labeling* of P is thus written $P_{V(C)}^t$. Also, $\mathcal{L}_W(P) = \{P_W^t \mid t \in D^{|W|}\}$ is the set of all W-labelings of P. We call it *the* W-labeling of P.

Figure 3.12 presents a fuzzy SCSP and shows its solution, while Figure 3.13 shows the $\{x, y\}$-labeling corresponding to the original problem and the set of possible partial solutions, whose disjunction gives rise to the complete solution. Notice that the problem set solution coincides with the solution of the problem in Figure 3.12.

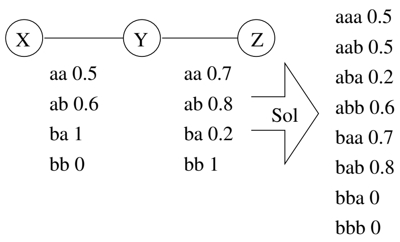

Figure 3.12: A fuzzy CSP and its solution.

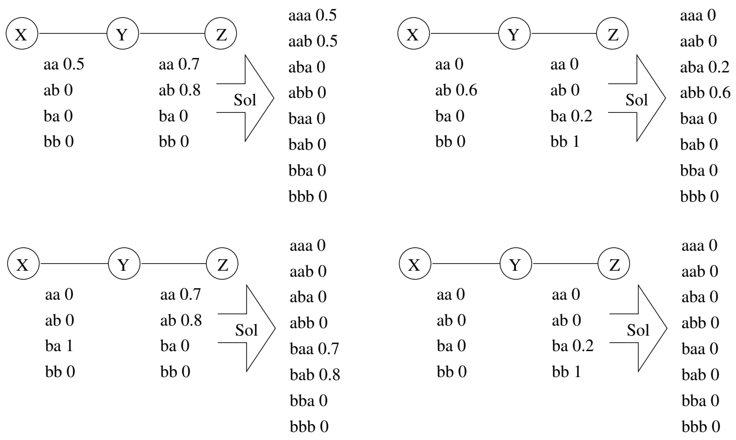

Figure 3.13: The set of the $\{x, y\}$-labelings corresponding to the problem.

This is true also in general; therefore, we can compute the solution of an SCSP by computing the disjunction of the solutions associated with the labeling. This is important, since it means that we can decompose an SCSP and solve it by solving its sub-problems. The following theorem formalizes this informal statement.

THEOREM 3.4.2 (DISJUNCTION CORRECTNESS) *Given an SCSP* $P = \langle C, con \rangle$, *a subset* W *of* $V(C)$, *and a set* $\{c_1, \dots, c_n\}$ *of constraints over* W *such that* $\bigoplus_i c_i = \langle 1, W \rangle$, *we have that* $Sol(P) = \bigoplus_i Sol(\langle C \cup \{c_i\}, con \rangle)$. *Moreover,* $P \equiv \mathcal{L}_W(P)$.

PROOF. By the assumptions, and by the distributivity of $\times$ over $+$, for any tuple t of domain values for all the variables, we get:

We note V the set $V(C)$ and p the size of W. We assume $\bigotimes C = \langle def_C, V \rangle$, $c_i = \langle def_i, W \rangle$ and $\bigoplus_i Sol(\langle C \cup c_i, con \rangle) = \langle def, con \rangle$. We have:

$$def(t) = \sum_i \sum_{\{t' | t' \downarrow_{con}^{V \cup W} = t\}} def_i(t' \downarrow_W^{V \cup W}) \times def_C(t' \downarrow_V^{V \cup W})$$

Since $W \subseteq V$, we have:

$$def(t) = \sum_i \sum_{\{t' | t' \downarrow_{con}^{V} = t\}} def_i(t' \downarrow_W^{V}) \times def_C(t')$$

Now, by distributivity of $\times$ over $+$, we get:

$$def(t) = \sum_{\{t' | t' \downarrow_{con}^{V} = t\}} def_C(t') \times (\sum_i def_i(t' \downarrow_W^{V}))$$

Since $\bigoplus_i c_i = \langle 1, W \rangle$, we have $\sum_i def_i(t' \downarrow_W^V) = 1$. Thus:

$$def(t) = \sum_{\{t' | t' \downarrow_{con}^{V} = t\}} def_C(t')$$

that coincides with $Sol(P)$. □

This result allows us to compute the solution of an SCSP using the labeling procedure. Moreover, if all the variables are instantiated, i.e. when the labeling $\mathcal{L}$ is complete ($W = V(P)$), then the values found for the variables contain all the informations needed to compute (very easily) the exact solution of the problem:

THEOREM 3.4.3 (SOLUTION OF A COMPLETE LABELING) *We suppose given a problem* $P = \langle C, con \rangle$ *where* $C = \{c_i | i \in [1, n]\}$ *and* $\forall i \in [1, n], c_i = \langle def_i, con_i \rangle$. *Let* $P^t_{V(C)} = \langle C \cup I^t_{V(C)}, con \rangle$ *be a complete labeling of* P. *Then* $Sol(P^t_{V(C)}) = \langle def, con \rangle$ *with:*

- $def(t \downarrow_{con}^{V(C)}) = \prod_{i \in [1,n]} def_i(t \downarrow_{con_i}^{V(C)})$,
- $def(t') = 0$ *if* $t' \neq t \downarrow_{con}^{V(C)}$.

Moreover, $\mathtt{blevel}(P^t_{V(C)}) = \prod_{i\in[1,n]} \mathtt{def}_i(t \downarrow^{V(C)}_{con_i})$.

PROOF. We denote by p the size of $V(C)$. By definition of a complete labeling problem:

$$\mathtt{Sol}(P^t_{V(C)}) = ((\bigotimes_{j\in[1,n]} c_j) \otimes (\bigotimes_{k\in[1,p]} c_{x_k,t_k,1})) \Downarrow_{con}$$

By associativity and idempotency of $\otimes$, we have:

$$\mathtt{Sol}(P^t_{V(C)}) = (\bigotimes_{j\in[1,n]} (c_j \otimes (\bigotimes_{k\in[1,p]} c_{x_k,t_k,1})) \Downarrow_{con}$$

Let $\langle \mathtt{def}', V(C)\rangle = c_j \otimes (\bigotimes_{k\in[1,p]} c_{x_k,t_k,1})$. Then:

- $\mathtt{def}'(t) = \mathtt{def}_j(t \downarrow^{V(C)}_{con_j})$,
- $\mathtt{def}'(t') = 0$ if $t' \neq t$.

The statement of the theorem easily follows. $\square$

As in the CSP case ([HDT92]), it is sufficient to maintain AC during the labeling procedure to completely solve the problem (after a complete labeling has been obtained).

THEOREM 3.4.4 (CORRECTNESS OF AC) *Given an SCSP* $P = \langle C, con\rangle$ *and a total order* $\prec$ *between the variables, assume that* $V(C) = \{x_1, \ldots, x_p\}$ *with* $x_1 \prec \cdots \prec x_p$, *and let* $P^t_{V(C)} = \langle C \cup I^t_{V(C)}, con\rangle$ *be a complete labeling of* P. *Also, let* $P' = ac(P^t_{V(C)})$ *be the arc-consistent problem obtained starting from* $P^t_{V(C)}$. *Then:*

- *for each* i *in* $[1,p]$, *the value* $[\{x_i\}]_{P'}$ *of location* x_i *in problem* P' *is a simple instantiation constraint* c_{x_i,t_i,v_i},
- $\mathtt{blevel}(P^t_{V(C)}) = \prod_{i\in[1,p]} v_i$.

PROOF. It is easy to see that the value $[\{x_i\}]_{P'}$ of location x_i in problem P' is still a simple instantiation constraint of the form c_{x_i,t_i,v_i}. Let us now prove that $\mathtt{blevel}(P^t_{V(C)}) = \prod_{i\in[1,p]} v_i$. One should first notice that $\forall i \in [1,p], \mathtt{blevel}(P^t_{V(C)}) \leq_S v_i$. Thus $\mathtt{blevel}(P^t_{V(C)}) \leq \prod_{i\in[1,p]} v_i$. We will now prove that $\mathtt{blevel}(P^t_{V(C)}) \geq \prod_{i\in[1,p]} v_i$. With the notations of Theorem 3.4.3, we have to prove that $\prod_{j\in[1,n]} \mathtt{def}_j(t \downarrow^{V(C)}_{con_j}) \geq \prod_{i\in[1,p]} v_i$. It is easy to see that $\forall i \in [1,p], \forall j \mid x_i \in con_j, v_i \leq \mathtt{def}_j(t \downarrow^{V(C)}_{con_j})$. Thus, we have $\forall i \in [1,p], v_i \leq \prod_{j|x_i\in con_j} \mathtt{def}_j(t \downarrow^{V(C)}_{con_j})$. The statement of the theorem follows from the idempotency of $\times$. $\square$

3.4.2 Partial Local Consistency

In this section we introduce the notion of *approximate local consistency algorithms*, and we give some sufficient conditions for such algorithms to terminate and be correct.

Notice that a local consistency algorithm is already an approximation in itself. In fact, because of its incompleteness, it approximates a complete solution algorithm. Usually, however, local consistency algorithms (as we have defined them in this section and as they are usually used in practice) replace a constraint by the *solution* of a set of other constraints. Here we want to approximate this aspect of such algorithms, by allowing them to replace a constraint by *a superset of the solution* of a set of other constraints. To preserve the nice properties of local consistency, however, we must assure that such a superset is not too big, so as not to introduce any additional solution of the whole problem.

Approximated Rules

The idea of an approximation function is to replace the real computation (that is, the solution of the subproblem defined by a rule) with a simpler one. The following definition states that a "correct" and "complete" approximation function should not enlarge the domains and also it should not lose any solution. Given a rule $r = l \leftarrow L$, an *approximation function* ϕ for r is a function from $\mathtt{type}^{-1}(L) \times \mathtt{type}^{-1}(l)$ to $\mathtt{type}^{-1}(l)$ such that:

- for all constraint set C of type L, for all constraint c of type l, $\mathtt{Sol}(\langle C \cup \{c\}, l\rangle) \sqsubseteq_S \phi(C, c) \sqsubseteq_S c$,
- for all constraint sets $\{c_1, \ldots, c_p\}$, $\{c'_1, \ldots, c'_p\}$ of type L, for all constraints c, c' of type l, $(\forall i \in [1, p], c_i \sqsubseteq_S c'_i, c \sqsubseteq_S c') \Rightarrow (\phi(\{c_1, \ldots, c_p\}, c) \sqsubseteq_S \phi(\{c'_1, \ldots, c'_p\}, c'))$.

As trivial examples, one can consider the approximation function that does no approximation (let us call it ϕ_{best}) and also the one that does no domain reduction (ϕ_{worst}).

EXAMPLE 3.4.5 Let $r = l \leftarrow L$ be a rule. Let ϕ_{best} such that $\forall C \in \mathtt{type}^{-1}(L), \forall c \in \mathtt{type}^{-1}(l)$, $\phi_{\mathrm{best}}(C, c) = \mathtt{Sol}(\langle C \cup \{c\}, l\rangle)$. It is easy to see that ϕ_{best} is an approximation function for r. It is the "best" approximation function (in fact, there are no rule that compute an approximation closer to the solution). □

EXAMPLE 3.4.6 Let $r = l \leftarrow L$ be a rule. Let ϕ_{worst} such that $\forall C \in \mathtt{type}^{-1}(L)$, $\forall c \in \mathtt{type}^{-1}(l)$, $\phi_{\mathrm{worst}}(C, c) = c$, It is easy to see that ϕ_{worst} is an approximation function for r. It is the "worst" approximation function (in fact, the rule does not change any information over the constraints). □

Given a rule $r = l \leftarrow L$ and an approximation function ϕ for r, the *approximation of rule* r *by function* ϕ is defined as $r_\phi = l \leftarrow_\phi L$. Given an approximated rule $r_\phi = l \leftarrow_\phi L$, the application of r_ϕ to problem P is defined by $[l \leftarrow_\phi L](P) = [l := \phi([L]_P, [l]_P)](P)$. That is, we apply to P the approximated rules instead of the original ones.

This notion of rule approximation can be used to modify the definition of a constraint without changing the solution of the problem. In fact, as the following theorem states, applying the approximated rule leads to a problem which has better values than the problem obtained by applying the rule. Thus, we cannot lose any solution, since we obtain an SCSP that is "between" the original problem and the problem obtained by using the non-approximated rules, which, we know by definition of local consistency, do not change the solution set.

THEOREM 3.4.7 $[l \leftarrow L](P) \sqsubseteq_P [l \leftarrow_\phi L](P) \sqsubseteq P$.

PROOF. By definition of rule application $[l \leftarrow L](P) = [l := Sol(\langle [L \cup \{l\}]_P, l\rangle)](P)$ and $[l \leftarrow_\phi L](P) = [l := \phi([L]_P, [l]_P)](P)$. Now, since the problem $[l \leftarrow L](P)$ and the problem $[l \leftarrow_\phi L](P)$ differ only for the constraint over the variable set l and since, by definition of an approximation function, $Sol(\langle C \cup \{c\}, l\rangle) \sqsubseteq_S \phi(C, c)$, the statement obviously holds. □

It is easy to verify that ϕ_{best} computes the same problem as the rule itself while ϕ_{worst} does not change the problem.

THEOREM 3.4.8 (EQUIVALENCE FOR RULES) *Given a problem* P *and an approximated rule* r_ϕ, $P \equiv [r_\phi](P)$.

PROOF. From $[r](P) \sqsubseteq_P [r_\phi](P) \sqsubseteq_P P$ and $[r](P) \equiv P$, it follows that $P \equiv [r_\phi](P)$. □

EXAMPLE 3.4.9 Let $r = l \leftarrow L$ be a rule. We have $[l \leftarrow_{\phi_{best}} L](P) = [l \leftarrow L](P)$ and $[l \leftarrow_{\phi_{worst}} L](P) = P$. □

We remind the reader that in order to fruitfully apply the local consistency scheme the $\times$ operation of the semiring has to be idempotent ([BMR95, BMR97b]) so in every statement regarding local consistency in the following we assume to have an idempotent $\times$ operation.

Since the application of an *approximated rule* does not change the solution of the problem, we can define a correct and complete semiring-based *partial local consistency* algorithm, by applying, following a given strategy, several approximated rules to a given SCSP.

Given a set of rules R, an *approximation scheme* Φ for R associates to each rule $r = l \leftarrow L$ of R an approximation $\Phi(r) = l \leftarrow_\phi L$. Given a problem P, a set R of local consistency rules, an approximation scheme Φ for R, and a fair strategy S for R, a local consistency algorithm applies to P the Φ-approximations of the rules in R in the order given by S. Thus, if the strategy is $S = s_1 s_2 s_3 \ldots$, the resulting problem is $P' = [\Phi(s_1); \Phi(s_2); \Phi(s_3); \ldots](P)$. The algorithm stops when the current problem is stable w.r.t. R. In that case, we denote by $plc(P, R, S, \Phi)$ the resulting problem.

PROPOSITION 3.4.10 (EQUIVALENCE) $P \equiv plc(P, R, S, \Phi)$.

PROOF. Trivially follows from Theorem 3.4.8 and by transitivity of $\equiv$. □

THEOREM 3.4.11 (TERMINATION) *Consider an SCSP P over a finite semiring. Then the application of a partial local consistency algorithm over* P *terminates in a finite number of steps.*

PROOF. Trivially follows from the intensivity of $\times$ and finiteness of A.

□

THEOREM 3.4.12 (ORDER-INDEPENDENCE) *Consider an SCSP* P *and two different applications of the same partial local consistency algorithm to* P*, producing respectively the problem* $P' = \mathtt{plc}(P, R, S, \Phi)$ *and* $P'' = \mathtt{plc}(P, R, S', \Phi)$*. Then,* $P' = P''$*. Thus, we will write* $P' = \mathtt{plc}(P, R, \Phi)$.

PROOF. It is easy to verify that $[\Phi(r)]$ is intensive and monotone, which is sufficient to apply the chaotic iteration theorem [Cou77]. □

Partial Arc-Consistency

Given an SCSP P and an approximation scheme Φ for AC, we will write $\mathtt{pac}(P, \Phi)$ instead of $\mathtt{plc}(P, AC, \Phi)$.

Let ϕ be an approximation function for $r = \{x\} \longleftarrow \{\{x, y_1, \ldots, y_n\}, \{y_1\}, \ldots, \{y_n\}\}$ (which, we recall, is the scheme for the AC rules). Let $Y = \{y_1, \ldots, y_n\}$ and $V = \{x\} \cup Y$. We say that ϕ is *instantiation-correct* if, for all constraints c of type V, for all $I_V^{d,\nu} = \{c_x\} \cup I_Y^{d',\nu'}$, we have

$$\phi(\{c\} \cup I_Y^{d',\nu'}, c_x) = \mathtt{Sol}(\langle \{c\} \cup I_V^{d,\nu}, \{x\} \rangle).$$

Informally, ϕ is instantiation-correct if the rule $\Phi(r)$ performs an exact computation as soon as all the variables appearing in such a rule are instantiated. For example, $\phi_{\mathtt{best}}$ is instantiation-correct, but $\phi_{\mathtt{worst}}$ is not.

For approximation functions satisfying the previous definition, partial AC performs the same domain reductions as AC, as soon as all the variables of the problem are instantiated.

THEOREM 3.4.13 (CORRECTNESS OF PAC) *Given an SCSP* $P = \langle C, \mathtt{con} \rangle$*, and an approximation scheme* Φ *for* AC *that is instantiation-correct, assume that* $V(C) = \{x_1, \ldots, x_p\}$ *with* $x_1 \prec \cdots \prec x_p$*. Let* $P^t_{V(C)} = \langle C \cup I^t_{V(C)}, \mathtt{con} \rangle$ *be a complete labeling of* P*. Let* $P'' = \mathtt{pac}(P^t_{V(C)}, \Phi)$ *and* $P' = \mathtt{ac}(P^t_{V(C)})$*. Then, for each* i *in* $[1, p]$*,* $[\{x_i\}]_{P'} = [\{x_i\}]_{P''}$.

PROOF. We will study the computation of the two algorithms (ac and pac). Because of the property of order-independence, we can assume that the rules are applied with the same strategy. Φ is correct w.r.t. to the instantiation. Hence, for each step (rule application and approximation of a rule application), the two algorithms compute the same domains (since all the variables are instantiated). Finally, for each i in $[1, p]$, $[\{x_i\}]_{P'} = [\{x_i\}]_{P''}$. □

COROLLARY 3.4.14 *With the notations of Theorem 3.4.13, for each* i *in* $[1, p]$*, the value* $[\{x_i\}]_{P''}$ *of location* x_i *in problem* P'' *is a simple instantiation constraint* c_{x_i, t_i, ν_i}*. Moreover,* $\mathtt{blevel}(P^t_{V(C)}) = \prod_{i \in [1,p]} \nu_i$.

PROOF. Easily follows from the previous theorem and from Theorem 3.4.4. □

3.4.3 Domains

We have seen that it is possible to approximate the computation of AC. We are now going to restrict our attention to certain kinds of approximations for AC, for which there exists a "good" representation for the variable domains. Then, we will adapt the computation of the approximations to these representations.

In the following we will always assume to have a constraint system $CS = \langle S, D, V \rangle$ where D is totally ordered. This means that there must be a way to compare any two elements of the domain D.

We call a *domain* any unary constraint. Given a set of domains E, we are looking for approximations schemes that build domains in E. Because of the properties of approximations, the set E has to satisfy the following three properties:

- for each domain, there must be a bigger domain in E;
- for each domain, there must be a smaller domain in E;
- E has to contain instantiation constraints (by instantiation-correctness); or, equivalently: $E \supseteq \{\langle 1, \{x\} \rangle | x \in V\} \cup \{c_{x,i,v} | x \in V, i \in D, v \in S\}$.

We will call an *approximation domain set* a set satisfying the above properties. The following proposition states that the notion of approximation domain is sufficient for the purpose of this work.

PROPOSITION 3.4.15 *Given an approximation domain set* E*, there exists an approximation scheme* Φ *such that the application of an approximation function (corresponding to* Φ*) builds domains in* E.

PROOF. To prove this proposition, just take Φ such that, for each rule $r = \{x\} \leftarrow \{\{x, y_1, \ldots, y_n\}, \{y_1\}, \ldots, \{y_n\}\}$, for each $c \in type^{-1}(\{x, y_1, \ldots, y_n\})$[4], for each $i \in [1, n]$, for each $c_{y_i} \in type^{-1}(\{y_i\})$, for each $c_x \in type^{-1}(\{x\})$, we have $\Phi(r)(\{c, c_{y_1}, \ldots, c_{y_n}\}, c_x) = Sol(\langle \{c, c_x, c_{y_1}, \ldots, c_{y_n}\}, \{x\} \rangle)$ if c_x and c_{y_i} $(\forall i \in [1, n])$ are instantiations constraints, and $\Phi(r)(\{c, c_{y_1}, \ldots, c_{y_n}\}, c_x) = c_x$ otherwise. □

The choice of an approximation domain set is very open. The approximation domains, however, should have a fixed memory requirement. Here we will consider some generalizations of the min-max scheme usually used in current constraint programming.

An *up-down-stair* is a domain $d = (d(1), \ldots, d(n))$ such that there exists $k \in [1, n]$ s.t. $d(1) \leq_S \ldots d(k) \geq_S \ldots d(n)$. We call m-*up-down-stair* any *up-down-stair* domain such that $|d(D)| - 1 \leq m$. We note $S(m)$ the set of m-up-down-stairs.

Note that the 1-up-down-stair corresponds exactly to the min-max scheme in the case where the semiring is the boolean one. Figure 3.14 presents an example of a 3-up-down-stair in the case of the fuzzy semiring.

[4]The function $type^{-1}(V)$ gives the constraint(s) c s.t. $type(c) = V$

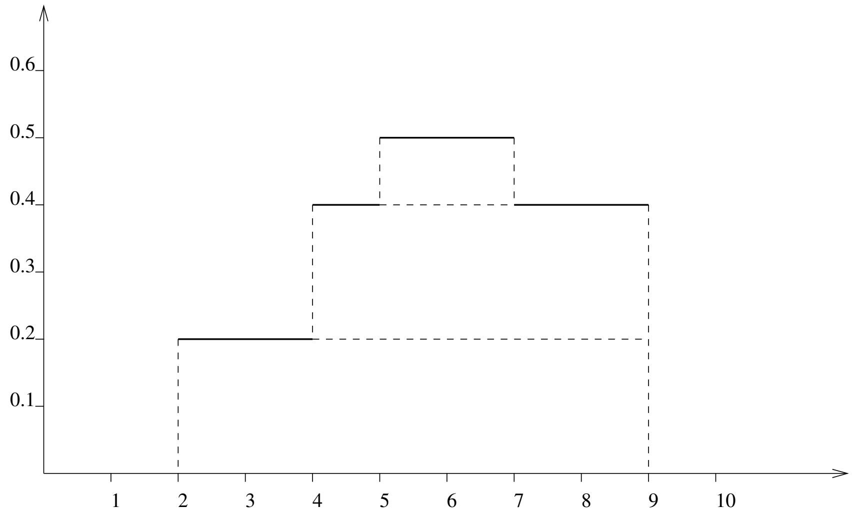

Figure 3.14: A 3-up-down-stair.

Note that an m-up-down-stair can be stored using $2m$ integers and m semiring values. A m-up-down-stair is made of m rectangles of decreasing width. Then, for each k between 1 and m, one has to store the interval $[inf_k, sup_k]$ that defines the kth rectangle and the total height of the k first rectangles. For the example of Figure 3.14, we get: $([2,9], 0.2); ([4,9], 0.4); ([5,7], 0.5)$.

Let us now slightly restrict our definition of approximation functions in order to give some insights for their construction. We suppose m given and we will focus on m-up-down-stair domains. Let us consider a constraint c over $V = \{x\} \cup Y$ where $Y = \{y_1, \ldots, y_n\}$ and $y_1 \prec \cdots \prec y_n$. Let us also consider one of the associated AC rules: $\{x\} \leftarrow \{V, \{y_1\}, \ldots, \{y_n\}\}$. We are looking for an approximation function $\phi : type^{-1}(\{V, \{y_1\}, \ldots, \{y_n\}\}) \times type^{-1}(\{x\}) \to type^{-1}(\{x\})$.

Since we do not consider general local consistency but only AC, only the domains can change; this is the reason why the approximation function should not take the constraint as a parameter. Moreover, because of their presence in many applications, and also for implementation issues, we will focus on approximations functions that intersect the new domain with the old one (more precisely, we will consider the case of a unique intersection function).

Hence, we will focus on the case where ϕ is defined by:

$$\phi(\{c, c_{y_1}, \ldots, c_{y_n}\}, c_x) = \chi(\psi(c_{y_1}, \ldots, c_{y_n}), c_x)$$

where:

- $\psi : S(m) \times \ldots \times S(m) \to S(m)$ computes a new domain (m-up-down-stair) for variable x;
- $\chi : S(m) \to S(m)$ intersects[5] the new domain and the old one.

[5]In the general case, $\otimes$ is not a function from $S(m) \times S(m)$ to $S(m)$ (this is true only in the case where

In the following, we will assume that the domain of the variable D has the form $[0, \infty]$ where ∞ is the maximal value an integer can take. As a practical and simple example, let us consider the case where $m = 1$. Here $\otimes$ is stable over $S(m)$, hence we take $\chi = \otimes$, which can be implemented by:

$$\chi(([i_1, s_1], v_1), ([i_1', s_1'], v_1')) = ([\max(i_1, i_1'), \min(s_1, s_1')], v_1 \times v_1').$$

Using 1-up-down-stairs, it is also easy to implement the rules for usual constraints.

EXAMPLE 3.4.16 (CONSTRAINT $X \leq Y$) For example, the two rules corresponding to constraint $x \leq y$ can be implemented by:

$$\psi_x(([i_y, s_y], v_y)) = ([0, s_y], v_y)$$

and

$$\psi_y(([i_x, s_x], v_x)) = ([s_x, \infty], v_x).$$

□

EXAMPLE 3.4.17 (CONSTRAINT $X = Y + C$) As another example, the two rules corresponding to constraint $x = y + c$ can be implemented by:

$$\psi_x(([i_y, s_y], v_y)) = ([i_y + c, s_y + c], v_y)$$

and

$$\psi_y(([i_x, s_x], v_x)) = ([i_x - c, s_x - c], v_x).$$

□

In the case where $m > 1$, it is also possible to implement χ and the rules associated with usual constraints. This implementation is more complicated since there are different representations corresponding to the same up-down-stair. Hence, one has to define a function that chooses a representation among all the possible ones.

3.5 Constraint Propagation: generalization and termination conditions

As described in Section 3.1 the propagation techniques usually used for classical CSPs have been extended and adapted to deal with soft constraints, provided that certain conditions are met. This has led to a general framework for soft constraint propagation, where at each step a subproblem is solved, as in classical constraint propagation. By studying the properties of this schema, it is proved (see Theorem 3.1.9 and 3.1.16) that such steps can be seen as applications of functions that are monotonic, inflationary, and idempotent over a certain partial order.

$m = 1$), hence we have to use χ instead.

On an orthogonal line of research, the concept of constraint propagation over classical constraints has been studied in depth in [Apt99a, Apt99b], and a general algorithmic schema (called GI) has been developed. In such a schema, constraint propagation is achieved whenever we have a set of functions that are monotonic and inflationary over a partial order with a bottom.

By studying these two frameworks and comparing them, we noticed that the GI schema can be applied to soft constraints (see Section 3.5.4), since the functions and the order used for soft constraints have all the necessary properties for GI. This is proven in this section by extending the partial order over soft constraint problems defined in Section 2.2.

By analyzing the features of the GI algorithm, we also realized (see Section 3.5.4) that indeed soft constraint propagation can be extended to deal with functions that are not necessarily idempotent. Notice that this is a double generalization: we don't require any longer that each step has to solve a subproblem (it could do some other operation over the problem), nor that it is idempotent. This allows us to model several forms of "approximate" constraint propagation, which were instead not modeled in [BMR97b]. Examples are: bounds-consistency for classical constraints [MS98], and partial soft arc-consistency for soft constraints [BCGR00b] described in Section 3.4.

These two results allow us to use the GI algorithm schema for performing a generalized form of soft constraint propagation. What is important to study, at this point, is when the resulting GI schema terminates. In fact, if we work with classical constraints over finite domains, it is easy to see that the GI algorithm always terminates. When moving to soft constraints over a semiring, however, even if the variable domain is finite, we could have an infinite behavior due to an infinite number of elements in the semiring. For example, fuzzy constraints have a semiring containing all reals between 0 and 1, and the semiring of weighted constraints contains all the reals, or all the naturals.

In Section 3.5.5 we identify some sufficient conditions for the termination of the GI algorithm over soft constraints. The first, predictable, condition that we consider is the well-foundness of the partial order over soft constraint problems: if the partial order over which the GI algorithm works has chains of finite length, since constraint propagation never goes from one chain to another one, obviously the whole algorithm terminates.

The second condition is in some sense more precise, although less general: when the propagation steps are defined via the two semiring operations, then we can just consider the sub-order over semiring elements obtained by taking the elements initially appearing in the given problem, and closing it under the two operations. In fact, in this case the GI algorithm cannot reach other elements. Therefore, if such a set (or a superset of it) is well-found, the GI algorithm terminates.

Each of these two conditions is sufficient for termination; however, they could be difficult to check, unless the partial order has a well-known structure of which we know the well-foundness. Nevertheless, in a special case we can formally prove that there exists a well-founded set of the shape required by the second condition above, and thus we can automatically deduce termination. This special case is related to the idempotency of the multiplicative operation of the semiring, the one that we use to combine constraints:

if this operation is idempotent, then GI terminates. For example, in classical constraints the multiplicative operation is logical and, and in fuzzy constraints it is the minimum, thus we can formally prove that the algorithm GI over *any* classical or fuzzy constraint problem *always* terminates, provided that the functions are defined via the two semiring operations.

We believe that the generalizations and termination conditions that we have developed and proven will make soft constraints more widely applicable, and soft constraint propagation more practically usable.

3.5.1 Some useful orderings over semiring constraints

We now review and modify some of the orderings among semiring elements, constraints, and problems, which have been introduced in the previous chapter; moreover, we also define new orderings that will be used in the next sections.

All the orderings we will consider in this section are derived from the partial order $\leq_S$ over semiring elements, which, we recall, is defined as follows: $a \leq_S b$ iff $a + b = b$. This intuitively means that b is "better" than a.

DEFINITION 3.5.1 *Consider any partial ordering $\langle D, \sqsubseteq \rangle$ and the component-wise ordering $\langle D^n, \sqsubseteq_n \rangle$, with $n \geq 1$, where $\langle d_1, \ldots, d_n \rangle \sqsubseteq_n \langle d'_1, \ldots, d'_n \rangle$ iff $d_i \sqsubseteq d'_i$ for each $i = 1, \ldots, n$. Let f be a function from D^n to D. Then:*

- *f is monotonic iff $\langle d_1, \ldots, d_n \rangle \sqsubseteq_n \langle d'_1, \ldots, d'_n \rangle$ implies $f(\langle d_1, \ldots, d_n \rangle) \sqsubseteq f(\langle d'_1, \ldots, d'_n \rangle)$;*

- *f is inflationary w.r.t. $\sqsubseteq$ iff $d_i \sqsubseteq f(\langle d_1, \ldots, d_n \rangle)$ for every $i = 1, \ldots, n$.*

Given the definition above, it is easy to see that the following results hold when D is the semiring set A and the order considered is $\leq_S$:

- $\leq_S$ is a partial order;

- 0 is the minimum;

- 1 is the maximum.

- if $\times$ is idempotent, then $\langle A, \leq_S \rangle$ is a distributive lattice where $+$ is the lub and $\times$ is the glb.

- $+$ and $\times$ are monotonic with respect to $\leq_S$;

- $+$ is inflationary with respect to $\leq_S$; instead $\times$ is inflationary with respect to $\geq_S$.

Constraint orders

From the ordering $\leq_S$ over A, we can also define a corresponding order between constraints (that extend the ordering already defined in the previous chapter). Before introducing the new order we define its domain, namely the set of all possible constraints over a constraint system.

DEFINITION 3.5.2 *Given a semiring* $S = \langle A, +, \times, 0, 1\rangle$ *and a constraint system* $CS = \langle S, D, V\rangle$*, we define the* Constraint Universe *related to the constraint system* CS *as follows:* $\mathcal{C}_{CS} = \bigcup_{con \subseteq V} \{\langle def, con\rangle \mid def : D^{|con|} \to A\}$.

We will write $\mathcal{C}$ (instead of $\mathcal{C}_{\mathcal{CS}}$) when the constraint system CS is clear from the context.

DEFINITION 3.5.3 *Consider two constraints* c_1, c_2 *over a constraint system* CS*; assume that* $con_1 \supseteq con_2$ *and* $|con_1| = k$*. Then we write* $c_1 \sqsubseteq_S c_2$ *if and only if, for all* k*-tuples* t *of values from* D*,* $def_1(t) \leq_S def_2(t \downarrow^{con_1}_{con_2})$.

Loosely speaking, a constraint c_1 is smaller than c_2 in the order $\sqsubseteq_S$ iff it constrains possibly more variables and assigns to each tuple a smaller value with respect to $\leq_S$ than c_2 does.

THEOREM 3.5.4 ($\sqsupseteq_S$ IS A PO) *Given a semiring* $S = \langle A, +, \times, 0, 1\rangle$ *with* $\times$ *idempotent and a constraint system* $CS = \langle S, D, V\rangle$*, we have the following:*

- *the relation* $\sqsupseteq_S$ *is a partial order over the set* $\mathcal{C}_{CS}$*;*
- *its bottom is* $\langle 1, \emptyset\rangle$*, where the* 0*-arity function* $1 : \emptyset \to A$ *is the constant* 1 *of the semiring.*

PROOF. We prove our first claim. We need to demonstrate that $\sqsubseteq_S$ is a reflexive, antisymmetric and transitive relation. Reflexivity holds trivially. To prove antisymmetry, suppose that $c_1 \sqsubseteq_S c_2$ and $c_2 \sqsubseteq_S c_1$; this yields that $con_1 = con_2$. Now, for all $t \in D^{|con_1|}$, we have both $def_1(t) \leq_S def_2(t)$ and $def_2(t) \leq_S def_1(t)$, hence $def_1(t) = def_2(t)$ and so $c_1 = c_2$. The transitivity of $\sqsubseteq_S$ follows from the transitivity of $\leq_S$. The other claim immediately follows from the definition of $\sqsubseteq_S$. □

We can easily extend the order $\sqsubseteq_S$ over constraints to a new order over *constraint sets* as follows.

DEFINITION 3.5.5 *Consider two sets of constraints* C_1, C_2 *over a constraint system* CS*. Suppose furthermore that* $C_1 = \{c^1_i : i \in I\}$*,* $C_2 = \{c^2_j : j \in J\}$*,* $J \subseteq I$ *and that, for every* $j \in J$*, the relation* $c^1_j \sqsubseteq_S c^2_j$ *holds. Then we write* $C_1 \sqsubseteq_C C_2$.

The intuitive reading of $C_1 \sqsubseteq_C C_2$ is that C_1 is a problem generally "more constraining" than C_2 is, because C_1 has (possibly) a larger number of "more restrictive" constraints than C_2 has.

THEOREM 3.5.6 ($\sqsupseteq_C$ IS A PARTIAL ORDER) *Given a semiring $S = \langle A, +, \times, 0, 1 \rangle$, and a constraint system $CS = \langle S, D, V \rangle$, we have that:*

- *the relation $\sqsupseteq_C$ is a partial order over $\wp(\mathcal{C})$;*
- *the bottom of the relation is $\emptyset$.*

PROOF. We only prove the first claim, the other one being straightforward. Reflexivity trivially holds. As far as antisymmetry is concerned, suppose that $C_1 = \{c_i^1\}_{i \in I}$, $C_2 = \{c_j^2\}_{j \in J}$ and both $C_1 \sqsubseteq_C C_2$ and $C_2 \sqsubseteq_C C_1$ hold; this means that $I = J$. Moreover, the following relations hold for every $i \in I$: $c_i^1 \sqsubseteq_S c_i^2$ and $c_i^2 \sqsubseteq_S c_i^1$. Hence $c_i^1 = c_i^2$ for every $i \in I$, because $\sqsubseteq_S$ is a partial order relation, cf. Theorem 3.5.4. Transitivity follows similarly, by exploiting the transitivity of $\sqsubseteq_S$. □

So far, we have introduced two partial orders: one between constraints ($\sqsubseteq_S$) and another one between constraint sets ($\sqsubseteq_C$). Local consistency algorithms, however, take constraint problems as input; therefore, we need an ordering relation between problems if we want the `GI` algorithm to be used for soft local consistency.

First we define the set of all problems that can be built over a constraint system CS.

DEFINITION 3.5.7 *Given a semiring $S = \langle A, +, \times, 0, 1 \rangle$ and a constraint system $CS = \langle S, D, V \rangle$, we define the* Problem Set Universe *related to the constraint system CS as $\mathcal{P}_{CS} = \{\langle C, con \rangle \mid C \subseteq \mathcal{C}_{CS}, con \subseteq var(C)\}$. When no confusion can arise, we shall simply write $\mathcal{P}$ instead of $\mathcal{P}_{CS}$.*

DEFINITION 3.5.8 *Given a constraint system CS, consider two problems $P_1 = \langle C_1, con_1 \rangle$ and $P_2 = \langle C_2, con_2 \rangle$ in $\mathcal{P}_{CS}$. We write $P_1 \sqsubseteq_{CP} P_2$ iff $C_1 \sqsubseteq_C C_2$ and $con_2 \subseteq con_1$.*

We now need to define a partially ordered structure that contains all SCSPs that can be generated by enforcing local consistency, starting from a given problem.

DEFINITION 3.5.9 *Consider a constraint system CS and an SCSP P over it. The* up-closure of P, *briefly $P \uparrow$, is the class of all problems P' on CS such that $P \sqsupseteq_{CP} P'$.*

PROPOSITION 3.5.10 *Consider a constraint system CS, an SCSP P over it and its up-closure $P \uparrow$. Then the following statements hold:*

1. *if $P_1 \sqsupseteq_{CP} P_2$ and $P_1 \in P \uparrow$, then $P_2 \in P \uparrow$;*
2. *if $P_1 \sqsupseteq_{CP} P_2$, then $P_2 \uparrow \subseteq P_1 \uparrow$.*

PROOF. The proof of the above proposition is an immediate consequence of the previous definition and of the fact that $\sqsubseteq_C$ is transitive, cf. Theorem 3.5.6. □

THEOREM 3.5.11 ($\sqsupseteq_{CP}$ IS A PO) *Given a constraint system $CS = \langle S, D, V \rangle$ and a problem P on it, we have:*

- *the relation* $\sqsupseteq_{CP}$ *is a partial order over* $\mathcal{P}_{CS}$*;*

- *in particular* $\langle P \uparrow, \sqsupseteq_{CP|P\uparrow}\rangle$ *is a partial ordering, where* $\sqsupseteq_{CP|P\uparrow}$ *is the restriction of* $\sqsupseteq_{CP}$ *to* $P \uparrow$*; when no confusion can arise, we simply write* $\langle P \uparrow, \sqsupseteq_{CP}\rangle$*;*

- *the bottom* $\bot_{CS}$ *of* $\langle P \uparrow, \sqsupseteq_{CP}\rangle$ *is* P.

Proof. We prove the first claim, the other ones following immediately from the definition of $P \uparrow$ and Proposition 3.5.10. As usual, we only prove that the relation is antisymmetric, because transitivity can be proven similarly and reflexivity trivially holds. Hence, suppose that both $P_1 \sqsubseteq_{CP} P_2$ and $P_2 \sqsubseteq_{CP} P_1$ hold. This means that we have the following relations: $con_2 \subseteq con_1$, $C_1 \sqsubseteq_C C_2$, $con_1 \subseteq con_2$, $C_2 \sqsubseteq_C C_1$. From the two previous relations and Theorem 3.5.6, it follows that $con_1 = con_2$ and $C_1 = C_2$; hence $P_1 = P_2$. □

3.5.2 Order-related properties of soft local consistency rules

We remind, from Section 2.2, that two problems P_1 and P_2, which share the same set of variables, are equivalent if they have the same solution set, and we write $P_1 \equiv_P P_2$.

Now we can list some useful properties of soft local consistency rules, which are related to equivalence and to problem ordering. Here we assume that we are given a constraint system C and a rule r on CS:

- (equivalence) $P \equiv_P r_l^L(P)$ if $\times$ is idempotent.

- (inflationarity) $P \sqsupseteq_{CP} r_l^L(P)$. This means that the new semiring values assigned to tuples by the rule application are always smaller than or equal to the old ones with respect to $\leq_S$.

- (monotonicity) Consider two SCSPs $P_1 = \langle C_1, con_1\rangle$ and $P_2 = \langle C_2, con_2\rangle$ over CS. If $P_1 \sqsubseteq_{CP} P_2$, then $r(P_1) \sqsubseteq_{CP} r(P_2)$.

It is easy to prove that all the results about local consistency rules hold also for a whole local consistency algorithm. Moreover, we can also prove that the strategy does not influence the result, if it is fair (see Section 3.1).

3.5.3 The Generic Iteration algorithm

In [Apt99b] the Generic Iteration ($\mathtt{GI}$) algorithm is introduced to find the least fixpoint of a finite set of functions defined on a partial ordering with bottom. This was then used as an algorithmic schema for classical constraint propagation: each step of constraint propagation was seen as the application of one of these functions. Our idea is to compare this schema with the one used for soft constraints, with the aim of obtaining a new schema which is the most general (that is, it can be applied both to classical and to soft constraints) and has the advantages of both of them.

Given a partial ordering with bottom, say $\langle D, \sqsubseteq, \bot \rangle$, consider now a set of functions $F := \{f_1, \ldots, f_k\}$ on D. The following algorithm can compute the least common fix point of the functions in F.

GENERIC ITERATION ALGORITHM (GI)

$d := \bot$;
$G := F$;
while $G \neq \emptyset$ **do**
    choose $g \in G$;
    $G := G - \{g\}$;
    $G := G \cup \mathtt{update}(G, g, d)$;
    $d := g(d)$
od

where for all G, g, d the set of functions $\mathtt{update}(G, g, d)$ from F is such that:

A. $\{f \in F - G \mid f(d) = d \wedge f(g(d)) \neq g(d)\} \subseteq \mathtt{update}(G, g, d)$;

B. $g(d) = d$ implies $\mathtt{update}(G, g, d) = \emptyset$;

C. $g(g(d)) \neq g(d)$ implies $g \in \mathtt{update}(G, g, d)$.

Assumption **A** states that $\mathtt{update}(G, g, d)$ at least contains all the functions from $F - G$ for which d is a fix point but $g(d)$ is not. So at each loop iteration such functions are added to the set G. In turn, assumption **B** states that no functions are added to G in case the value of d did not change. Note that, even though after the assignment $G := G - \{g\}$ we have $g \in F - G$, still $g \notin \{f \in F - G \mid f(d) = d \wedge f(g(d)) \neq g(d)\}$ holds. So assumption **A** does not provide any information when g is to be added back to G. This information is provided in assumption **C**. On the whole, the idea is to keep in G at least all functions f for which the current value of d is not a fix point.

We now recall the results which state the (partial) correctness of the GI algorithm, cf. [Apt99b]:

i. Every terminating execution of the GI algorithm computes in d a common fixpoint of the functions from F.

ii. Suppose that all functions in F are monotonic. Then every terminating execution of the GI algorithm computes in d the least common fixpoint of all the functions from F.

iii. Suppose that all functions in F are inflationary and that D is finite. Then every execution of the GI algorithm terminates.

3.5.4 Generalized local consistency for SCSPs via algorithm GI

In this section we will try to combine the two formalisms described so far (soft constraints and the GI algorithm). Our goal is to exploit the `GI` algorithm to perform local consistency over soft constraint problems.

GI for standard local consistency over soft constraints

The functions that `GI` needs in input are defined on a partial ordering with bottom. In the case of local consistency rules for SCSPs, the partial ordering is $\langle \mathsf{P} \uparrow, \sqsupseteq_{\mathsf{CP}} \rangle$, and the bottom is the problem P itself, cf. Theorem 3.5.11. Moreover, the local consistency rules (and also the more general local consistency functions) have all the "good" properties that GI needs. Namely, those functions are monotonic and inflationary. Thus, algorithm GI can be used to perform constraint propagation over soft constraint problems. More precisely, we can see that algorithm GI and the local consistency algorithm schema for soft constraints obtain the same result.

THEOREM 3.5.12 (GI FOR SOFT LOCAL CONSISTENCY RULES) *Given an SCSP* P *over a constraint system* CS*, consider the SCSP* $\mathsf{lc}(\mathsf{P}, \mathsf{R}, \mathsf{S})$ *obtained by applying to* P *a local consistency algorithm using the rules in* R *and with a fair strategy* S*. Consider also the partial order* $\langle \mathsf{P} \uparrow, \sqsupseteq_{\mathsf{CP}} \rangle$*, and the set of functions* R*, and apply algorithm GI to such input. Then the output of* GI *coincides with* $\mathsf{lc}(\mathsf{P}, \mathsf{R}, \mathsf{S})$.

PROOF. Since we already know that $\mathsf{lc}(\mathsf{P}, \mathsf{R}, \mathsf{S})$ and GI terminate (by the results appeared in [Apt99a, BMR97b], and since the computed fixpoint is the same (by previous item *ii.*), their output obviously coincide. □

GI for generalized local consistency over soft constraints

While all local consistency rules are idempotent (since they solve a subproblem), algorithm GI does not need this property. This means that we can define a generalized notion of local consistency rules for soft constraints, by dropping idempotency.

DEFINITION 3.5.13 (LOCAL CONSISTENCY FUNCTIONS) *Consider an SCSP* P *over a semiring* S*. A* local consistency function *for* P *is a function* $\mathsf{f} : \mathsf{P} \uparrow \rightarrow \mathsf{P} \uparrow$ *which is monotonic and inflationary over* $\sqsupseteq_{\mathsf{CP}}$.

With this definition of a local consistency function we relax two conditions about a local consistency step:

- that it must solve a subproblem;
- that it must be idempotent.

The second generalization has been triggered by the results about the GI algorithm, which have shown that idempotency is not needed for the desired results. Moreover, many practical local consistency algorithms do not exactly solve subproblems, but generate an approximation of the solution (see for example the definition of *bounds consistency* in [MS98] or the notion of partial soft arc-consistency in [BCGR00b]). Thus, the first extension allows one to model many more practical propagation algorithms.

THEOREM 3.5.14 (GI FOR SOFT LOCAL CONSISTENCY FUNCTIONS) *Given a constraint system CS and an SCSP* P *on it, let us apply the* GI *algorithm to the partial order* $\langle \mathsf{P}\uparrow, \sqsupseteq_{\mathsf{CP}}\rangle$ *and a finite set* R *of local consistency functions. Then every terminating execution of the* GI *algorithm computes in the output problem* P' *the least common fixpoint of all the functions from* R.

PROOF. Easily follows by item *i.* of the previous section. □

What is now important to investigate is when the algorithm terminates. This is particularly crucial for soft constraints, since, even when the variable domains are finite, the semiring may contain an infinite number of elements, which is obviously a source of possible non-termination.

3.5.5 Termination of the GI algorithm over soft constraints

As noted above, the presence of a possibly infinite semiring may lead to a constraint propagation algorithm that does not terminate. In the following we will give several independent conditions which guarantee termination in some special cases.

The first condition is a predictable extension of the one given in Section 3.5.3: instead of requiring the finiteness of the domain of computation, we just require that its chains have finite length, since it is easy to see that constraint propagation moves along a chain in the partial order.

THEOREM 3.5.15 (TERMINATION 1) *Given a constraint system CS and an SCSP* P *on it, let us instantiate the* GI *algorithm with the po* $\langle \mathsf{P}\uparrow, \sqsupseteq_{\mathsf{CP}}\rangle$ *and a finite set* R *of local consistency functions. Suppose that the order* $\sqsupseteq_{\mathsf{CP}}$ *restricted to* $\mathsf{P}\uparrow$ *is well founded. Then every execution of the* GI *algorithm terminates.*

PROOF. Easily follows from Theorem 1 in [Apt99b]. Note that there the author says to need a finite (and not a well found) partial order, but the proof do not uses this restriction. □

This theorem can be used to prove termination in many cases. For example, classical constraints over finite domains generate a partial order that is finite (and thus trivially well-found), so the above theorem guarantees termination. Another example occurs when dealing with weighted soft constraints, where we deal with the naturals. Here the semiring is $\langle \mathsf{N}, \mathtt{min}, +, 0, +\infty\rangle$. Thus we have an infinite order, but well-found.

There are also many interesting cases, however, in which the ordering $\langle \mathsf{P}\uparrow, \sqsupseteq_{\mathsf{CP}}\rangle$ is not well-found. Consider for instance the case of fuzzy or probabilistic CSPs. For fuzzy

CSPs, the semiring is $\langle [0,1], max, min, 0, 1\rangle$. Thus, the partially ordered structure containing all problems that are smaller than the given one, according to the semiring partial order, is not well-found, since we have all the reals between 0 and a certain element in $[0,1]$. Thus, the above theorem cannot say anything about termination of GI in this case. This does not mean, however, that GI does not terminate, but only that the theorem above cannot be applied. In fact, later we will give another sufficient condition that will guarantee termination in the last two cases as well.

In fact, if we restrict our attention to local consistency functions defined via $+$ and $\times$, we can define another condition on our input problem that guarantees the termination of the GI algorithm; this condition exploits the fact that the local consistency functions are defined by means of the two semiring operations, and the properties of such operations.

DEFINITION 3.5.16 (SEMIRING CLOSURE) *Consider a constraint system* CS *with semiring* $S = \langle A, +, \times, \mathbf{0}, \mathbf{1}\rangle$*, and an SCSP* P *on* CS*. Consider also the set of semiring values appearing in* P*:* $Cl(P) = \bigcup_{\langle def', con'\rangle \in C}\{def'(d) \mid d \in D^{|con'|}\}$*. Then, a* semiring closure *of* P *is any set* B *such that:* $Cl(P) \subseteq B \subseteq A$*;* B *is closed with respect to* $+$ *and* $\times$*;* $<_S$ *restricted to* B *is well founded.*

THEOREM 3.5.17 (TERMINATION 2) *Consider a constraint system* CS *with semiring* $S = \langle A, +, \times, \mathbf{0}, \mathbf{1}\rangle$*, an SCSP* P *on it and a finite set of local consistency functions* R *defined via* $+$ *and* $\times$*. Assume there also exists a semiring closure of* P*. Then every execution of the* `GI` *algorithm terminates.*

PROOF. The proof is similar to the one in [Apt99b] for the termination theorem; just replace the order $\sqsubseteq_{CP}$ with $\leq_S$ and accordingly the set D with B, where B is a semiring closure of P. $\square$

Notice that this theorem is similar to the one in Section 3.1 about termination; however there we force the set B to be finite in order to guarantee the termination of a local consistency algorithm; a hypothesis that is implied by ours.

If we have a fuzzy constraint problem, then we can take B as the set of all semiring values appearing in the initial problem. In fact, this set is closed with respect to min and max, which are the two semiring operations in this case. Moreover, it is well-found, since it is finite. Another example is constraint optimization over the reals: if the initial problem contains only natural numbers, then the set B can be the set of all naturals, which is is a well-founded subset of the reals and it is closed w.r.t. $+$ (min) and $\times$ (sum).

Therefore, by using Theorem 3.5.17 we can also prove that constraint propagation over fuzzy constraint problems always terminates, provided that each step of the algorithm uses a local consistency function, which is defined in terms of the two semiring operations only.

It is not, however, always easy to find a semiring closure of a given SCSP P, mainly because we should check that the order restricted to a tentative set B, closed and containing $Cl(P)$, is well-found. Nevertheless, there is a special case in which we do not have to find such a set, because we can prove that it always exists (which is what Theorem 3.5.17 requires). This special case occurs when the multiplicative operation of the semiring is idempotent. In fact, we can prove that in this case there always exists a finite (and thus well-found) semiring closure of any given problem over that semiring. This obviously is very convenient, since it provides us with an easy way to check whether Theorem 3.5.17 can be applied.

THEOREM 3.5.18 (IDEMPOTENCY OF $\times$ AND TERMINATION) *Consider a constraint system $\mathtt{CS}$, an SCSP $\mathtt{P}$ on it and a finite set of local consistency functions $\mathtt{R}$ defined via $+$ and $\times$. Assume also that $\times$ is idempotent. Then there exists a finite semiring closure of $\mathtt{P}$, and thus every execution of the GI algorithm terminates.*

PROOF. Consider the set $\mathtt{Cl}(\mathtt{P})$ of all semiring elements appearing in $\mathtt{P}$. If we combine any subset of them via the $+$ operation, we generate a set of elements, which contains $\mathtt{Cl}(\mathtt{P})$, that we denote by $\mathtt{Cl}+$. Notice that this set is finite, because the number of subsets of $\mathtt{Cl}(\mathtt{P})$ is finite.

Let us now combine any subset of elements of $\mathtt{Cl}+$ via the $\times$ operation: in this way we generate a larger set of elements, containing $\mathtt{Cl}+$, which we denote by $\mathtt{Cl}+\times$. Again, this set is finite. Therefore $\mathtt{Cl}+\times$ is a finite semiring closure of $\mathtt{P}$. □

Consider again the fuzzy constraint example. Here $\times$ is min, thus it is idempotent. Therefore, by Theorem 3.5.18, GI over such problems always terminates. This is an alternative, and easier, way (to Theorem 3.5.17) to guarantee that soft constraint propagation over fuzzy constraints terminates. In fact, we do not have to find a semiring closure of the problem, but just check that the multiplicative operation is idempotent.

Considering all the above results, we can devise the following steps towards proving the termination of algorithm GI on a soft constraint problem $\mathtt{P}$ over a semiring S:

- If the local consistency functions are defined via the two operations of S, and the multiplicative operation of S is idempotent, then GI terminates (by Theorems 3.5.18).

- If instead $\times$ is not idempotent, but we still have local consistency functions defined via the two semiring operations, we can try to find a semiring closure of $\mathtt{P}$. If we find it, then GI terminates (by Theorem 3.5.17).

- If we cannot find a semiring closure of $\mathtt{P}$, or the local consistency functions are more general, then we can try to prove that the partial order of problems is well-found. If it is so, the GI terminates (by Theorem 3.5.15).

While Theorem 3.5.18 applies in a special case of the hypothesis of Theorem 3.5.17, it is interesting to investigate the relationship between the hypothesis of Theorem 3.5.15 and 3.5.17. What can be proven is that these two conditions, namely, the well-foundness of the partial order of problems and the existence of a semiring closure, are independent. In other words, there are cases in which one holds and not the other one, and vice versa. To prove this result, we need the following definition.

DEFINITION 3.5.19 *Let* $S = \langle A, +, \times, 0, 1\rangle$ *be a semiring and* B *a subset of* A*. The set* B *is a* down-set *(or an* order ideal*) if, whenever* $a \in B$*,* $a' \in A$ *and* $a' \leq_S a$*, then* $a' \in B$*. Given any subset* B *of* A*, the downward closure of* B *is*

$$B \downarrow := \{d' \in A \,:\, \exists d\,(d \in B \textit{ and } d' \leq_S d)\}.$$

Observe that the class $\mathcal{F}$ of down-sets containing a subset B of A is not empty, since A itself is such a set. Moreover, it is easy to check that the downward closure of B is the smallest down-set of $\mathcal{F}$; hence the downward closure of a set is well defined.

Given a semiring $S = \langle A, +, \times, 0, 1\rangle$, the following result links the upward closure of a problem P with the downward closure of $Cl(P)$, thus allowing us to compare the conditions in Theorem 3.5.15 and 3.5.17.

PROPOSITION 3.5.20 *Given a constraint system* CS *and a problem* P *defined on it, consider the set* $B := \{def(t) \in A \,:\, \exists P' \in P\uparrow\ (c := \langle def, con\rangle \in P',\ t \in D^{|con|})\}$*. Then* $B = Cl(P)\downarrow$*.*

PROOF. It follows immediately from the definition of $\sqsupseteq_{CP}$, of $P\uparrow$ and $Cl(P)\downarrow$. □

Now we can notice that, given a subset B of a semiring, if $<_S$ restricted to B is well founded, then so is $<_S$ restricted to $\hat{B} := B \cup \{0, 1\}$. Furthermore, if B is finite, so is $\hat{B}$. In fact it is sufficient to check that the following identities hold because of the fact that $S = \langle A, +, \times, 0, 1\rangle$ is a c-semiring: if $a \in B$ then $a + 0 = a \in B$; if $a \in B$ then $a \times 1 = a \in B$; if $a \in B$ then $a + 1 = 1 \in \hat{B}$; if $a \in B$ then $a \times 0 = 0 \in \hat{B}$. Hence, in Theorem 3.5.17, we can replace the hypothesis "$Cl(P) \subseteq B$ and B a semiring closure of $Cl(P)$" with the condition that "Cl(P) is a subset of a well founded sub-c-semiring B of §$S = \langle A, +, \times, 0, 1\rangle$".

Moreover, $+$ is the least upper bound operation and, if $\times$ is idempotent, $\times$ is the greatest lower bound operation (see chapter 2 and [BMR97b]). Hence a sub-c-semiring is also a sub-lattice of $S = \langle A, +, \times, 0, 1\rangle$ and vice versa if $\times$ is idempotent. Thus, a subset B of a semiring can be a down-set and yet it may be not closed with respect to $\times$ and $+$. For instance, the set of negative real numbers augmented with $-\infty$ is a down-set in the lattice $\mathbf{R}^\infty$ of reals extended with $\{+\infty, -\infty\}$ and the usual linear ordering; however, it is not a sub-lattice itself. Vice versa, there are sub-lattices of $\mathbf{R}^\infty$ - hence sets that are closed with respect to the least upper bound and the greatest lower bound operations - that are not down-sets. For instance, the extended interval $[0, 1] \cup \{+\infty, -\infty\}$. Therefore, the two conditions that guarantee the termination of algorithm GI in Theorems 3.5.15 and 3.5.17 are independent.

3.6 Dynamic programming for SCSPs

Dynamic programming [Bel57, BD62] can be used to solve a problem by solving some subproblems of it and then combining their solutions to obtain the solution of the whole problem (see for example [GMM81] for its use in graphs and [MR86, MR91] for its adoption in classical CSPs where is called *perfect relaxation*). In the SCSP framework, a suitably instantiated version of dynamic programming can be fruitfully used as well: at each step, a subset of constraints is chosen and solved, and its solution (which, we recall, is a constraint) replaces the whole subset of constraints.

Before defining the algorithm, we need to describe in a tree-like way the SCSP to be solved, so that the algorithm can then follow the tree structure in a bottom-up way.

DEFINITION 3.6.1 (PARSING TREE) *Given an SCSP* $P = \langle C, \mathtt{con} \rangle$*, a parsing tree of* P *is a sequence of local consistency rules* $S = r_1; \ldots ; r_n$*, where* $r_i = (l_i \leftarrow L_i)$*. Let*

- $L = \{l_i, i = 1, \ldots, n\} \cup \bigcup_{i=1,\ldots,n} L_i$ *(that is,* L *is the set of all locations occurring in all rules in* S*);*
- $l_i \prec l$ *iff* $l \in L_i$ *(that is, rule* r_i *"generates" location* l*);*
- $l_i \prec v$ *iff* $v \notin l_i$ *and* $v \in l'$ *with* $l' \in L_i$ *(that is, rule* r_i *"generates" variable* v*);*
- *for any* x *location or variable,* $\mathtt{prec}(x) = \{l \mid l \prec x\}$.

Then we require that

- (tree-like) *for any* x *variable or location, if* $x = l_n$ *or* $x \in l_n$*, then* $\mathtt{prec}(x) = \emptyset$*, else* $\mathtt{prec}(x)$ *is a singleton;*
- (cover) $\mathtt{con} = l_n$*, and* $\langle \mathtt{def}, \mathtt{con}' \rangle \in C$ *implies* $\mathtt{con}' \in L$*;*
- (bottom-up parsing) $l_i \prec l_j$ *implies* $j < i$.

In words, the above definition provides a sequence of rules which cover the whole problem and are connected among them as a tree. Moreover, the sequence gives a bottom-up visit of the tree. It is now easy to define a dynamic programming algorithm where, at each step, one rule is processed, according to the given sequence. Note also that the rules we use to explore the tree in a dynamic programming fashion have the same structure of the rule used to perform a step of relaxation of the problem. This means that dynamic programming can be seen as a special case of relaxation algorithm.

DEFINITION 3.6.2 (DYNAMIC PROGRAMMING ALGORITHM) *Given an SCSP* P*, consider a parsing tree* $S = r_1; \ldots ; r_n$ *of* P*. Then compute* $[S](P)$.

We will now prove that the value of location l_n when the algorithm terminates coincides with the solution of the given problem P.

THEOREM 3.6.3 (SOLUTION) *Given an SCSP* $P = \langle C, l \rangle$ *and a parsing tree* S *for* P*, we have that* $\mathtt{Sol}(P) = [l]_{[S](P)}$.

PROOF. We will prove the statement of the theorem by induction on the length of the parsing tree. For the base case, consider a parsing tree with just one rule, that is, $l \leftarrow L$. It is easy to see that $[l]_{[r](P)}$ is the solution of P. In fact, applying this rule means exactly solving the whole problem. Assume now that the statement holds for all parsings with $n-1$ rules, and consider a parsing S_1 with n rules $r_1, \ldots, r_n$, where $r_i = l_i \leftarrow L_i$ for all $i = 1, \ldots, n$. Consider now the first rule of this parsing, that is, $r_1 = l_1 \leftarrow L_1$, and take another rule r_i such that $l_1 = l_i$ or $l_1 \in L_i$, and there is no other rule r_j with $j < i$ and such that $l_1 = l_j$ or $l_1 \in L_j$. That is, r_i is the first rule after r_1 which has l_1 either in its left hand side or in its right hand side. Note that such a rule must exist by definition of parsing tree. In fact, either $l_1 = l_n$, and in this case l_1 must appear at least in the left hand side on r_n, or $l_1 \neq l_n$, in which case it must be generated by one rule (r_i). We will now consider the two separate cases: that $l_1 = l_i$ and that $l_1 \in L_i$.

Assume that $l_1 \in L_i$. Consider then the sequence $S_2 = r_2, \ldots, r_{i-1}, r_1, r_i, \ldots, r_n$. It is easy to see that this is still a parsing for the given problem. Moreover, the constraint of type l_n resulting after applying the rules in S_2 coincides with that obtained after applying the rules in S_1. In fact, since l_1 can be generated by just one rule (by definition of parsing), all rules in $\{r_2, \ldots, r_{i-1}\}$ cannot change the definition of l_1 and thus applying r_1 before or after them cannot change the resulting constraint. Consider now another sequence S_3 of $n-1$ rules $r_2, \ldots, r_{i-1}, r'_i, r_{i+1}, \ldots, r_n$, where $r'_i = l_i \leftarrow L_i \cup L_1$. It is easy to see that this sequence of rules is again a parsing tree for P, thus by induction hypothesis the constraint with type l_n after the application of S_2 is the solution of the problem. Now we have to show that such a constraint coincides with that obtained at location l_n after applying sequence S_2 (which we already know to coincide with that obtained at location l_n after sequence S_1). What makes S_2 and S_3 differ is the fact that rule r_1 of S_2 has been "merged" with rule r_i to give rule r'_i in S_3. More precisely, the application of rule r'_i combines all constraints specified by $L_1 \cup L_i$, while the application of first rule r_1 and then rule r_i combines first the constraint specified by L_1 and then combines the resulting constraint with those constraints specified by L_i. We will show that these two methods yield the same final constraint.

In the following of this proof, for ease of readability, we will write l instead of $[l]_P$. On one side (rule r'_i) we have the constraint $(\bigotimes(L_i \cup L_1 \cup \{l_i\})) \Downarrow_{l_i}$, and on the other side (rule r_1 and then r_i) we have the constraint $(\bigotimes((L_i - \{l_1\}) \cup \{l_i\} \cup \{(\bigotimes(L_1 \cup \{l_1\})) \Downarrow_{l_1}\})) \Downarrow_{l_i}$. Thus we have to prove that:

$$(\bigotimes(L_i \cup L_1 \cup \{l_i\})) \Downarrow_{l_i} = (\bigotimes((L_i - \{l_1\}) \cup \{l_i\} \cup \{(\bigotimes L_1 \cup \{l_1\}) \Downarrow_{l_1}\})) \Downarrow_{l_i} .$$

We will start from the left hand side of the formula and try to reach the right hand side. We have:
$(\bigotimes(L_i \cup L_1 \cup \{l_i\})) \Downarrow_{l_i} =$
$\{$by separating $(L_i \cup \{l_i\}) - \{l_1\}$ and $L_1 \cup \{l_1\}$ which are disjoint$\}$
$((\bigotimes((L_i \cup \{l_i\}) - \{l_1\})) \otimes (\bigotimes(L_1 \cup \{l_1\}))) \Downarrow_{l_i} =$
$\{$by Theorem 2.2.8, and letting $V_i = \bigcup(L_i \cup \{l_i\}$, that is the set of all variables involved in rule r_i
$\}$
$((\bigotimes((L_i \cup \{l_i\}) - \{l_1\})) \otimes (\bigotimes(L_1 \cup \{l_1\}))) \Downarrow_{V_i} \Downarrow_{l_i} =$
$\{$by Theorem 2.2.9, since the variables not in V_i are not involved in $L_i \cup \{l_i\} - \{l_1\}\}$
$((\bigotimes((L_i \cup \{l_i\}) - \{l_1\})) \otimes (\bigotimes(L_1 \cup \{l_1\}) \Downarrow_{V_i})) \Downarrow_{l_i} =$
$\{$by Proposition 2.2.7, since $V_i \cap V_1 = l_1$, where $V_1 = (\bigcup L_1) \cup l_1\}$
$((\bigotimes((L_i \cup \{l_i\}) - \{l_1\})) \otimes (\bigotimes(L_1 \cup \{l_1\})) \Downarrow_{l_1}) \Downarrow_{l_i}$. which coincides with the right hand side of the formula.

Let us now consider the second case, in which $l_1 = l_i$. Then consider the sequence $S_2 =$

$r_2, \ldots, r_{i-1}, r_1, r_i, \ldots, r_n$. With a reasoning similar to above, it is easy to see that S_2 is still a parsing tree for P and that the constraint of type l_n obtained via S_2 is the same as the one obtained via S_1, since by assumption location l_1 does not appear in rules $r_2, \ldots, r_{i-1}$. Consider now $S_3 = r_2, \ldots, r_{i-1}, r'_i, r_{i+1}, \ldots, r_n$ where $r'_i = l_1 \leftarrow L_i \cup L_1$. Again, S_3 is a parsing tree for P. Moreover, it has $n-1$ rules, thus the constraint of type l_n after the application of S_3 coincides with the solution of P. Now we will show that this is the same as that obtained after the application of S_2. More precisely, we have to show that

$$(\bigotimes(L_1 \cup L_i \cup \{l_1\})) \Downarrow_{l_1} = (\bigotimes(L_i \cup (\bigotimes(L_1 \cup \{l_1\})) \Downarrow_{l_1}) \Downarrow_{l_1} .$$

This can be easily proven by using reasoning similar to that of the first case.

□

Note that this dynamic programming algorithm can be applied to any instance of the SCSP framework, even when $\times$ is not idempotent (and thus the local consistency algorithms cannot safely be applied).

Obviously any SCSP has a parsing tree[6], but not all classes of SCSPs have *convenient* parsing trees for each SCSP of the class (see [MM72], where it is shown that the class of rectangular lattices does not have such a property). Here, by convenient we mean that the size of each rule is smaller than some bound N, which is fixed and the same for the whole class of considered SCSPs. In such a case, each step of the algorithm solves an SCSP with bounded size. Thus, the complexity of this step may be exponential in the size of such an SCSP, but constant w.r.t. the size of the overall problem. Therefore, the complexity of the algorithm is linear in the number of rules of the parsing tree, which is linear in the number of variables of the problem. Thus, the overall algorithm is linear in the number of variables of the problem. When applied to standard CSPs, this algorithm reduces to the *perfect relaxation* algorithm of [MR86, MR91].

DEFINITION 3.6.4 (N-BOUNDED PARSING TREE) *Given an SCSP $P = \langle C, con \rangle$ and an integer N, consider a parsing tree $S = \{l_1 \leftarrow L_1; \ldots; l_n \leftarrow L_n\}$ for P. Then S is N-bounded for P if*

1. *for all $i = 1, \ldots, n$, $| \, l_i \cup \bigcup_{l \in L_i} l \, | \leq N$ (that is, the number of variables of a rule is bounded by N);*
2. *for all $i = 1, \ldots, n$, there is a variable v such that $l_i \prec v$ (that is, each rule generates at least a variable);*
3. *$V_P = \bigcup_{\langle def, con' \rangle \in C} con' = \bigcup_{l \in L} l$ (that is, all the variables occurring in the rules are present in the constraint problem P).*

THEOREM 3.6.5 (LINEAR ALGORITHM WHEN CONVENIENT PARSING) *Given an integer N, consider the class of all SCSPs which have an N-bounded parsing tree. Consider now any SCSP P of this class, and let n be the number of its variables. Then in the worst case P can be solved in time $O(n)$.*

[6] Just take the parsing tree with just one rule.

PROOF. Just apply the dynamic programming algorithm which follows an N-bounded parsing tree for P, which exists by assumption. Each step of the algorithm applies one of the rules. This is in the worst case $O(|D|^N \times 2^N)$, where D is the domain of each variable and N is an upper bound to the number of variables of the rules. In fact, the number of tuples of values for the variables of the rule is exponential in the number of variables, which by assumption is bounded by N, and for each tuple one has to check a number of constraints equal to the number of locations, which again can be exponential in the number of variables of the rule. Thus, the worst case time complexity of a rule application is constant, since both N and D are fixed. There are as many steps as rules in the parsing tree. Since each rule by assumption generates at least one variable, the number of rules is bounded by the number of variables n of the whole problem. Therefore, in the worst case the number of steps of the algorithm is $O(n)$. □

Conditions *2* and *3* in Definition 3.6.4 are not restrictive for Theorem 3.6.5. In fact, if we have a parsing for some problem, which satisfies condition *1* but where some rule, say $r = l \leftarrow L$, does not eliminate any variable, then it is always possible to obtain another parsing for the same problem that still satisfies condition *1* but where each rule eliminates at least a variable: just eliminate rule r and replace every occurrence of l in the other rules with L, and do similarly for all rules that do not eliminate variables. If instead there are rules containing variables that do not appear in the problem, then one can always obtain another parsing where such variables are not there: just remove all locations involving such "ghost" variables. By doing this, we still have the cover property of the parsing, since a location involving one or more ghost variables cannot correspond to any constraint of the given problem.

Theorem 3.6.5 states that in some cases there is an algorithm which is $O(n)$. One could wonder whether n is a good measure of the size of the given SCSP, and consider instead the number m of its constraints as more significant. In fact, in general the number of constraints may be much larger than the number of variables (for example, in binary CSPs, m is $O(n^2)$). For the classes of SCSPs that have an N-bounded parsing tree, however, it is possible to show that m is $O(n)$. In fact, consider an N-bounded parsing tree for a problem P of the class. Then, each rule may have at most N variables, thus it will contain at most 2^N locations, and there are at most n rules. Thus the total number of locations in the parsing is $n \times 2^N$. Also, since a parsing tree for P covers P, the number of locations is either the same or greater than the number of constraints in P. Thus $m \leq (n \times 2^N)$. Therefore we have that m is $O(n)$.

In order to find a parsing tree for a given SCSP, one could adapt the studies on the *secondary problem* in dynamic programming (see for example [BB72]). In some cases, however, the problem is generated in a way that the parsing tree is already explicit, like in the case of constraint logic programming (CLP) languages [JL87]. In fact, during the execution of a CLP program, the use of the clauses of the program builds a constraint problem with a parsing tree where each node of the tree directly corresponds to one of the used clauses.

A characterization of what dynamic programming is and when it can be applied, which is similar to the one given in this section, can be found in [She91, She94]. There, *valuation-based systems* are defined as systems based on variables, valuations, and two operations, called *combination* and *marginalization*. These two operations are very related, respectively, to our notions of *combination* and *projection*. In valuation-based systems, three axioms are required for the correct application of a dynamic programming algorithm. These three axioms are indeed satisfied by our framework as well. In fact, they correspond, respectively, to 1) commutativity and associativity of $\otimes$ (see comment after Definition 2.2.5), 2) Theorem 2.2.8, and 3) Theorem 2.2.9. Note that the proof of Theorem 3.6.3 relies just on these three properties. In fact, our dynamic programming algorithm, defined in Definition 3.6.2, can be seen as an extension of the *fusion algorithm* defined in [She91, She94], where the extension consists in the fact that more than one variable at a time can be eliminated. In fact, in our algorithm all variables in the right-hand side of a rule are considered at a time.

Moreover, a related algorithm that solves a constraint problem with a tree-like shape in a bottom-up way, as in Definition 3.6.2, has been described in [DDP90] for optimization problems and in [DD88] for belief maintenance. In those papers, the idea is to consider either the constraint graph, if it is acyclic, or the dual graph of a constraint problem (where nodes are constraints and arcs are associated with variables shared among constraints), and to use techniques like cycle-cutset [DP88a] or tree-clustering [DP88b] to provide such a dual graph with a tree-like shape. In [DDP90], however, constraints are combined via the usual *and* operator, and the value associated to each tuple is computed in a completely independent way, via a given utility function. Thus, apart from the tree-structure, our approach and that in [DDP90] are very different, since we also generalize the way constraints are combined, via a general notion of combination and projection, of which *and* and *or* are just instances.

Our notion of parsing tree is more general than that of *hinge-trees* presented in [GJC94]. In fact, we do not assume anything about the structure of each node of the tree, which may be a generic graph. On the contrary, in [GJC94] they consider only nodes that cannot be further decomposed into hinge-trees.

3.7 Conclusions

In this chapter we have generalized the classical solution and preprocessing techniques used for crisp CSPs so they could also be fruitfully used in SCSPs. In particular, we have shown that sufficient conditions for the effective applicability of local consistency and of dynamic programming algorithms could be checked just looking at the semiring properties. We have reviewed in this sense the different non-crisp frameworks and we have given a detailed description of the applicability of these techniques. Sufficient conditions for the termination of these (and more general) techniques have been given by just looking at the partial order structure represented by the constraint system or by the semiring.

Moreover, we have described some methodologies that could be applied to speed up the search for solutions (by applying labeling and/or partial local consistency) and some topological properties of the constraint graph that have to be checked to perform special cuts in the search tree.

The techniques of local consistency (or partial local consistency) will also be used in the next chapter where SCSPs' abstraction will be introduced. The idea is that if we are dealing with a framework where the local consistency steps cannot be safely applied (like in the WCSPs) we can, nonetheless, reduce the search space. In fact, we can abstract the problem, perform safe local consistency in the abstract framework, and then take back the collected information over the concrete problem.

Moreover, these same techniques will be the object of the studies presented in Chapter 5. There, in fact, a small language will be introduced and each of the steps used in the local consistency techniques will be seen as a function plus an assignment operation that will change the semiring value associated with the problem.

Chapter 4

SCSP Abstraction

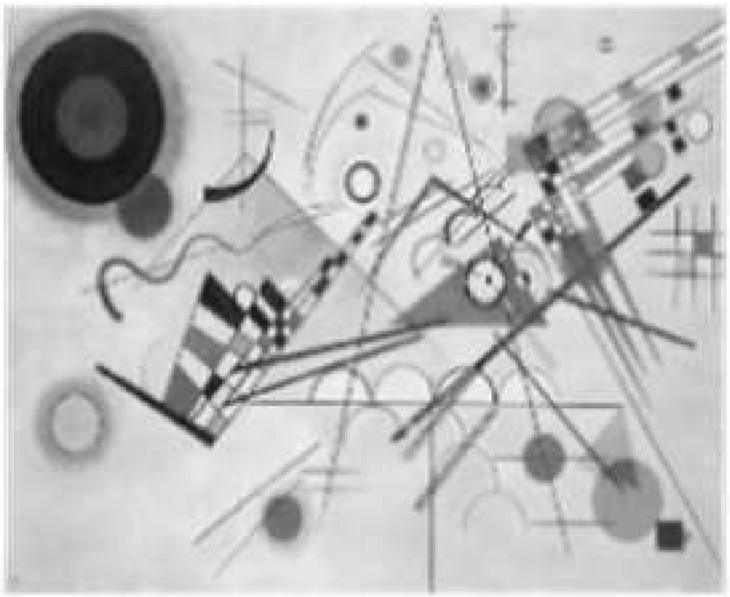

(Abstraction) Composition VIII 1923; Oil on canvas.
Wassily Kandinsky

Abstract

We propose an abstraction scheme for soft constraint problems and we study its main properties. Processing the abstracted version of a soft constraint problem can help us in many ways: for example, to find good approximations of the optimal solutions, or also to provide us with information that can make the subsequent search for the best solution easier. Moreover, we show how the abstraction framework can be used to import constraint propagation algorithms from the abstract scenario to the concrete one. This may be useful when we don't have any (or any efficient) propagation algorithm in the concrete setting.

Although it is obvious that SCSPs are much more expressive than classical CSPs, they are also more difficult to process and to solve. Therefore, sometimes it may be too costly to find all, or even only one, optimal solution. Also, although classical propagation techniques, like arc-consistency [Mac77], can be extended to SCSPs (see Chapter 3), even such techniques can be too costly to be used, depending on the size and structure of the partial order associated to the SCSP.

For these reasons, it may be reasonable to work on a simplified version of the given problem, trying, however, to not lose too much information. We propose to define this simplified version by means of the notion of abstraction, which takes an SCSP and returns a new one that is simpler to solve. Here, as in many other works on abstraction [WG92b, WG92a], "simpler" may mean many things, like the fact that a certain solution algorithm finds a solution, or an optimal solution, in a fewer number of steps, or also that the abstracted problem can be processed by machinery that is not available in the concrete context.

There are many formal proposals to describe the process of abstracting a notion, be it a formula, or a problem [WG92b], or even a classical [Fre91] or a soft CSP [SFV95]. Among these, we chose to use one based on Galois insertions [CC77], mainly to refer to a well-know theory, with many results and properties that can be useful for our purposes. This made our approach compatible with the general theory of abstraction in [WG92b]. Then, we adapted it to work on soft constraints: given an SCSP (the *concrete* one), we get an abstract SCSP by just changing the associated semiring, and relating the two structures (the concrete and the abstract one) via a *Galois insertion*. Note that this way of abstracting constraint problems does not change the structure of the problem (the set of variables remains the same, as well as the set of constraints), but just the semiring values to be associated to the tuples of values for the variables in each constraint.

Once we get the abstracted version of a given problem, we propose to

1. process the abstracted version: this may mean either solving it completely, or also applying some incomplete solver, which may derive some useful information from the abstract problem;

2. bring back to the original problem some (or possibly all) of the information gained in the abstract context;

3. continue the solution process on the transformed problem, which is a concrete problem equivalent to the given one.

All this process has the main aim of finding an optimal solution, or an approximation of it, for the original SCSP, within the resource bounds we have. The hope is that, by following the above three steps, we get to the final goal faster than just solving the original problem.

A deep study of the relationship between the concrete SCSP and the corresponding abstract one allows us to prove some results that can help in deriving useful information from the abstract problem and then take some of the derived information back to the concrete problem. In particular, we can prove the following:

- If the abstraction satisfies a certain property, all optimal solutions of the concrete SCSP are also optimal in the corresponding abstract SCSP (see Theorem 4.3.8). Thus, in order to find an optimal solution of the concrete problem, we could find all the optimal solutions of the abstract problem, and then just check their optimality on the concrete SCSP.

- Given any optimal solution of the abstract problem, we can find upper and lower bounds for an optimal solution for the concrete problem (see Theorem 4.3.10). If we are satisfied with these bounds, we could just take the optimal solution of the abstract problem as a reasonable approximation of an optimal solution for the concrete problem.

- If we apply some constraint propagation technique over the abstract problem, say P, obtaining a new abstract problem, say P', some of the information in P' can be inserted into P, obtaining a new concrete problem, which is closer to its solution and thus easier to solve (see Theorem 4.3.14 and 4.3.17).

  This, however, can be done only if the semiring operation that describes how to combine constraints on the concrete side is idempotent (see Theorem 4.3.14).

- If instead this operation is not idempotent, we can still bring back some information from the abstract side. In particular, we can bring back the inconsistencies (that is, tuples associated with the worst element of the semiring), since we are sure that these same tuples are also inconsistent also in the concrete SCSP (see Theorem 4.3.17).

In both the last two cases, the new concrete problem is easier to solve, in the sense, for example, that a branch-and-bound algorithm would explore a smaller (or equal) search tree before finding an optimal solution.

The results of this chapter, already appeared in [BCGR00a, BCR00].

4.1 Abstraction

Abstract interpretation [BL65, CC77, CC79] is a theory developed to reason about the relation between two different semantics (the *concrete* and the *abstract* semantics). The idea of approximating program properties by evaluating a program on a simpler domain of descriptions of "concrete" program states goes back to the early' 70's. The inspiration was that of approximating properties from the exact (concrete) semantics into an approximate (abstract) semantics, that explicitly exhibits a structure (e.g., ordering) which is somehow present in the richer concrete structure associated to program execution.

The guiding idea is to relate the concrete and the abstract interpretation of the calculus by a pair of functions, the *abstraction* function α and the *concretization* function γ, which form a Galois connection.

Let $(\mathcal{C}, \sqsubseteq)$ (concrete domain) be the domain of the concrete semantics, while $(\mathcal{A}, \leq)$ (abstract domain) be the domain of the abstract semantics. The partial order relations reflect an approximation relation. Since in approximation theory a partial order specifies the precision degree of any element in a poset, it is obvious to assume that if α is a mapping associating an abstract object in $(\mathcal{A}, \leq)$ for any concrete element in $(\mathcal{C}, \sqsubseteq)$, then the following holds: if $\alpha(x) \leq y$, then y is also a correct, although less precise, abstract approximation of x. The same argument holds if $x \sqsubseteq \gamma(y)$. Then y is also a correct approximation of x, although x provides more accurate information than $\gamma(y)$. This gives rise to the following formal definition.

DEFINITION 4.1.1 (GALOIS INSERTION) *Let $(\mathcal{C}, \sqsubseteq)$ and $(\mathcal{A}, \leq)$ be two posets (the concrete and the abstract domain). A Galois connection $\langle \alpha, \gamma \rangle : (C, \sqsubseteq) \rightleftharpoons (A, \leq)$ is a pair of maps $\alpha : \mathcal{C} \to \mathcal{A}$ and $\gamma : \mathcal{A} \to \mathcal{C}$ such that*

1. *α and γ are monotonic,*
2. *for each $x \in \mathcal{C}$, $x \sqsubseteq \gamma(\alpha(x))$ and*
3. *for each $y \in \mathcal{A}$, $\alpha(\gamma(y)) \leq y$.*

Moreover, a Galois insertion (of $\mathcal{A}$ in $\mathcal{C}$) $\langle \alpha, \gamma \rangle : (\mathcal{C}, \sqsubseteq) \rightleftharpoons (\mathcal{A}, \leq)$ is a Galois connection where $\gamma \cdot \alpha = \mathrm{Id}_{\mathcal{A}}$.

Property 2 is called *extensivity* of $\alpha \cdot \gamma$. The map α (γ) is called the *lower* (*upper*) *adjoint* or *abstraction* (*concretization*) in the context of abstract interpretation.

The following basic properties are satisfied by any Galois insertion:

1. γ is injective and α is surjective.
2. $\alpha \cdot \gamma$ is an upper closure operator in $(\mathcal{C}, \sqsubseteq)$.
3. α is additive and γ is co-additive.
4. Upper and lower adjoints uniquely determine each other. Namely,
$$\gamma = \lambda y. \bigcup_{\mathcal{C}} \{x \in \mathcal{C} \mid \alpha(x) \sqsubseteq y\}, \alpha = \lambda x. \bigcap_{\mathcal{A}} \{y \in \mathcal{A} \mid x \leq \gamma(y)\}$$
5. α is an isomorphism from $(\gamma\alpha)(\mathcal{C})$ to $\mathcal{A}$, having γ as its inverse.

An example of a Galois insertion can be seen in Figure 4.1. Here, the concrete lattice is $\langle [0, 1], \leq \rangle$, and the abstract one is $\langle \{0, 1\}, \leq \rangle$. Function α maps all real numbers in $[0, 0.5]$ into 0, and all other integers (in $(0.5, 1]$) into 1. Function γ maps 0 into 0.5 and 1 into 1.

One property that will be useful later relates to a precise relationship between the ordering in the concrete lattice and that in the abstract one.

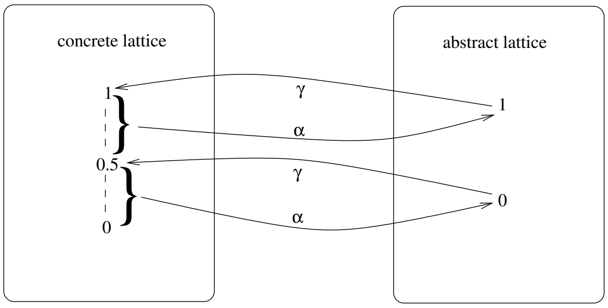

Figure 4.1: A Galois insertion.

THEOREM 4.1.2 (TOTAL ORDERING) *Consider a Galois insertion from* $(\mathcal{C}, \sqsubseteq)$ *to* $(\mathcal{A}, \leq)$. *Then, if* $\sqsubseteq$ *is a total order, also* $\leq$ *is so.*

PROOF. It easily follows from the monotonicity of α (that is, $x \sqsubseteq y$ implies $\alpha(x) \leq \alpha(y)$, and from its surjectivity (that is, there is no element in $\mathcal{A}$ which is not the image of some element in $\mathcal{C}$ via α). □

Usually, both the concrete and the abstract lattice have some operators that are used to define the corresponding semantics. Most of the times it is useful, and required, that the abstract operators show a certain relationship with the corresponding concrete ones. This relationship is called *local correctness*.

DEFINITION 4.1.3 (LOCAL CORRECTNESS) *Let* $f : \mathcal{C}^n \to \mathcal{C}$ *be an operator over the concrete lattice, and assume that* $\tilde{f}$ *is its abstract counterpart. Then* $\tilde{f}$ *is locally correct w.r.t.* f *if* $\forall x_1, \ldots, x_n \in \mathcal{C}. f(x_1, \ldots, x_n) \sqsubseteq \gamma(\tilde{f}(\alpha(x_1), \ldots, \alpha(x_n)))$.

4.2 Abstracting soft CSPs

Given the notions of soft constraints and abstraction, recalled in the previous sections, we now want to show how to abstract soft constraint problems. The main idea is very simple: we just want to pass, via the abstraction, from an SCSP P over a certain semiring S to another SCSP $\tilde{P}$ over the semiring $\tilde{S}$, where the lattices associated to $\tilde{S}$ and S are related by a Galois insertion as shown above.

DEFINITION 4.2.1 (ABSTRACTING SCSPS) *Consider the* concrete *SCSP* $P = \langle C, con \rangle$ *over semiring* S, *where*

- $S = \langle A, +, \times, 0, 1 \rangle$ *and*
- $C = \{c_0, \ldots, c_n\}$ *with* $c_i = \langle con_i, def_i \rangle$ *and* $def_i : D^{|con_i|} \to A$*;*

we define the abstract *SCSP* $\tilde{P} = \langle \tilde{C}, con \rangle$ *over the semiring* $\tilde{S}$, *where*

- $\tilde{S} = \langle \tilde{A}, \tilde{+}, \tilde{\times}, \tilde{0}, \tilde{1} \rangle$;
- $\tilde{C} = \{\tilde{c_0}, \ldots, \tilde{c_n}\}$ *with* $\tilde{c_i} = \langle con_i, \tilde{def_i} \rangle$ *and* $\tilde{def_i} : D^{|con_i|} \rightarrow \tilde{A}$;
- *if* $L = \langle A, \leq \rangle$ *is the lattice associated to* S *and* $\tilde{L} = \langle \tilde{A}, \tilde{\leq} \rangle$ *the lattice associated to* $\tilde{S}$, *then there is a Galois insertion* $\langle \alpha, \gamma \rangle$ *such that* $\alpha : L \rightarrow \tilde{L}$;
- $\tilde{\times}$ *is locally correct with respect to* $\times$.

Notice that the kind of abstraction we consider in this chapter does not change the structure of the SCSP. That is, C and $\tilde{C}$ have the same number of constraints, and c_i and $\tilde{c_i}$ involve the same variables. The only thing that is changed by abstracting an SCSP is the semiring. Thus, P and $\tilde{P}$ have the same graph topology (variables and constraints), but different constraint definitions, since if a certain tuple of domain values in a constraint of P has semiring value a, then the same tuple in the same constraint of $\tilde{P}$ has semiring value $\alpha(a)$. Notice also that α and γ can be defined in several different ways, but all of them have to satisfy the properties of the Galois insertion, from which it derives, among others, that $\alpha(\mathbf{0}) = \tilde{\mathbf{0}}$ and $\alpha(\mathbf{1}) = \tilde{\mathbf{1}}$.

EXAMPLE 4.2.2 As an example, consider any SCSP over the semiring for optimization

$$S_{WCSP} = \langle \mathcal{R}^- \cup \{-\infty\}, \max, +, -\infty, 0 \rangle$$

(where costs are represented by negative reals) and suppose we want to abstract it onto the semiring for fuzzy reasoning

$$S_{FCSP} = \langle [0, 1], max, min, 0, 1 \rangle.$$

In other words, instead of computing the maximum of the sum of all costs, we just want to compute the maximum of the minimum of all costs, and we want to normalize the costs over $[0..1]$. Notice that the abstract problem is in the FCSP class and it has an idempotent $\times$ operator (which is the min). This means that in the abstract framework we can perform local consistency over the problem in order to find inconsistencies. As noted above, the mapping $\alpha : \langle \mathcal{R}^-, \leq_{WCSP} \rangle \rightarrow \langle [0, 1], \leq_{FCSP} \rangle$ can be defined in different ways. For example one can decide to map all the reals below some fixed real x onto 0 and then to map the reals in $[x, 0]$ into the reals in $[0, 1]$ by using a normalization function, for example $f(r) = \frac{x-r}{x}$. □

EXAMPLE 4.2.3 Another example is the abstraction from the fuzzy semiring to the classical one:

$$S_{CSP} = \langle \{0, 1\}, \vee, \wedge, 0, 1 \rangle.$$

Here function α maps each element of $[0, 1]$ into either 0 or 1. For example, one could map all the elements in $[0, x]$ onto 0, and all those in $(x, 1]$ onto 1, for some fixed x. Figure 4.1 represents this example with $x = 0.5$. □

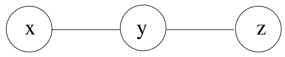

a a ... 0	a a ... 0.3
a b ... 1	a b ... 0.1
b a ... 0.5	b a ... 1
b b ... 0.7	b b ... 1

Solutions:

a a a ... min(0,0.3) = 0
a a b ... min(0,0.1) = 0
a b a ... min(1,1) = 1
a b b ... min(1,1) = 1
b a a ... min(0.5,0.3) = 0.3
b a b ... min(0.5,0.1) = 0.1
b b a ... min(0.7,1) = 0.7
b b b ... min(0.7,1) = 0.7

Best solutions:
a b a ... 1
a b b ... 1

Figure 4.2: A fuzzy CSP and its solutions.

In Figure 4.2 we show an example of fuzzy CSP and its solutions.

We have defined Galois insertions between two lattices $\langle L, \leq_S \rangle$ and $\langle \tilde{L}, \tilde{\leq}_{\tilde{S}} \rangle$ of semiring values. For convenience, however, in the following we will often use Galois insertions between lattices of problems $\langle PL, \sqsubseteq_S \rangle$ and $\langle \tilde{PL}, \tilde{\sqsubseteq}_{\tilde{S}} \rangle$ where PL contains problems over the concrete semiring and $\tilde{PL}$ over the abstract semiring. This does not change the meaning of our abstraction, we are just upgrading the Galois insertion from semiring values to problems. Thus, when we will say that $\tilde{P} = \alpha(P)$, it will mean that $\tilde{P}$ is obtained by P via the application of α to all the semiring values appearing in P.

An important property of our notion of abstraction is that the composition of two abstractions is still an abstraction. This allows to build a complex abstraction by defining several simpler abstractions to be composed.

THEOREM 4.2.4 (ABSTRACTION COMPOSITION) *Consider an abstraction from the lattice corresponding to a semiring S_1 to that corresponding to a semiring S_2, denoted by the pair $\langle \alpha_1, \gamma_1 \rangle$. Consider now another abstraction from the lattice corresponding to the semiring S_2 to that corresponding to a semiring S_3, denoted by the pair $\langle \alpha_2, \gamma_2 \rangle$. Then the pair $\langle \alpha_1 \cdot \alpha_2, \gamma_2 \cdot \gamma_1 \rangle$ is an abstraction as well.*

PROOF. We first have to prove that $\langle \alpha, \gamma \rangle = \langle \alpha_1 \cdot \alpha_2, \gamma_2 \cdot \gamma_1 \rangle$ satisfies the four properties of a Galois insertion:

- since the composition of monotone functions is again a monotone function, we have that both α and γ are monotone functions;
- given a value $x \in S_1$, from the first abstraction we have that $x \sqsubseteq_1 \gamma_1(\alpha_1(x))$. Moreover, for any element $y \in S_2$, we have that $y \sqsubseteq_2 \gamma_2(\alpha_2(y))$. This holds also for $y = \alpha_1(x)$, thus by monotonicity of γ_1 we have $x \sqsubseteq_1 \gamma_1(\gamma_2(\alpha_2(\alpha_1(x))))$.
- a similar proof can be used for the third property;
- the composition of two identities is still an identity .

To prove that $\times_3$ is locally correct w.r.t. $\times_1$, it is enough to consider the local correctness of $\times_2$ w.r.t. $\times_1$ and of $\times_3$ w.r.t. $\times_2$, and the monotonicity of γ_1. $\square$

4.3 Properties and advantages of the abstraction

In this section we will define and prove several properties that hold on abstractions of soft constraint problems. The main goal here is to point out some of the advantages that one can have in passing through the abstracted version of a soft constraint problem instead of directly solving the concrete version.

4.3.1 Relating a soft constraint problem and its abstract version

Let us consider the scheme depicted in Figure 4.3. Here, and in the following pictures, the left box contains the lattice of concrete problems, and the right one the lattice of abstract problems. The partial order in each of these lattices is shown via dashed lines. Connections between the two lattices, via the abstraction and concretization functions, is shown via directed arrows. In the following, we will call $S = \langle A, +, \times, 0, 1 \rangle$ the concrete semiring and $\tilde{S} = \langle \tilde{A}, \tilde{+}, \tilde{\times}, \tilde{0}, \tilde{1} \rangle$ the abstract one. Thus, we will always consider a Galois insertion $\langle \alpha, \gamma \rangle : \langle A, \leq_S \rangle \rightleftharpoons \langle \tilde{A}, \leq_{\tilde{S}} \rangle$.

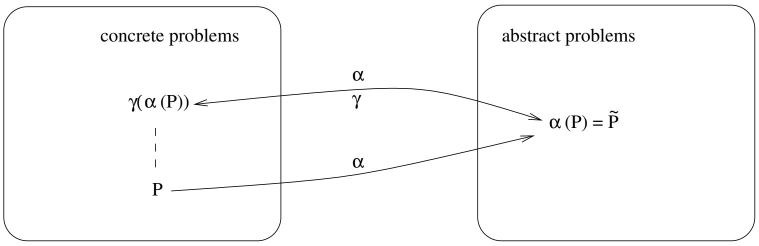

Figure 4.3: The concrete and the abstract problem.

In Figure 4.3, P is the starting SCSP. Then with the mapping α we get $\tilde{P} = \alpha(P)$, which is an abstraction of P. By applying the mapping γ to $\tilde{P}$, we get the problem $\gamma(\alpha(P))$. Let us first notice that these two problems (P and $\gamma(\alpha(P))$) are related by a precise property, as stated by the following theorem.

THEOREM 4.3.1 *Given an SCSP P over S, we have that $P \sqsubseteq_S \gamma(\alpha(P))$.*

PROOF. Immediately follows from the properties of a Galois insertion, in particular from the fact that $x \leq_S \gamma(\alpha(x))$ for any x in the concrete lattice. In fact, $P \sqsubseteq_S \gamma(\alpha(P))$ means that, for each tuple in each constraint of P, the semiring value associated to such a tuple in P is smaller (w.r.t. $\leq_S$) than the corresponding value associated to the same tuple in $\gamma(\alpha(P))$. □

Notice that this implies that, if a tuple in $\gamma(\alpha(P))$ has semiring value 0, then it must have value 0 also in P. This holds also for the solutions, whose semiring value is obtained by combining the semiring values of several tuples.

COROLLARY 4.3.2 *Given an SCSP P over S, we have that $\mathrm{Sol}(P) \sqsubseteq_S \mathrm{Sol}(\gamma(\alpha(P)))$.*

PROOF. We recall that $\mathrm{Sol}(P)$ is just a constraint, which is obtained as $\bigotimes(C) \Downarrow_{con}$. Thus the statement of this corollary comes from the monotonicity of $\times$ and $+$. $\square$

Therefore, by passing from P to $\gamma(\alpha(P))$, no new inconsistencies are introduced: if a solution of $\gamma(\alpha(P))$ has value 0, then this was true also in P. It is possible, however, that some inconsistencies are forgotten (that is, they appear to be consistent after the abstraction process).

EXAMPLE 4.3.3 Consider the abstraction from the fuzzy to the classical semiring, as described in Figure 4.1. Then, if we call P the fuzzy problem in Figure 4.2, Figure 4.4 shows the concrete problem P, the abstract problem $\alpha(P)$, and its concretization $\gamma(\alpha(P))$. It is easy too see that, for each tuple in each constraint, the associated semiring value in P is lower than or equal to that in $\gamma(\alpha(P))$.

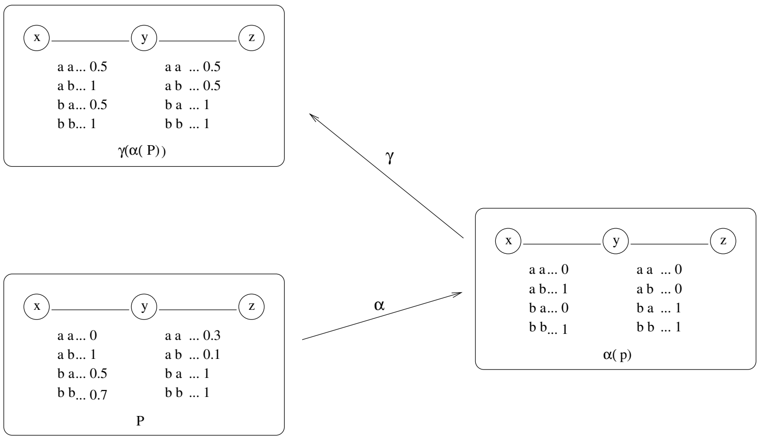

Figure 4.4: An example of the abstraction fuzzy-classical.

$\square$

If the abstraction preserves the semiring ordering (that is, applying the abstraction function and then combining gives elements that are in the same ordering as the elements obtained by combining only), then there is also an interesting relationship between the set of optimal solutions of P and that of $\alpha(P)$. In fact, if a certain tuple is optimal in P, then this same tuple is also optimal in $\alpha(P)$. Let us first investigate the meaning of the *order-preserving* property.

DEFINITION 4.3.4 (ORDER-PRESERVING ABSTRACTION) *Consider two sets* I_1 *and* I_2 *of concrete elements. Then an abstraction is said to be order-preserving if*

$$\tilde{\prod_{x \in I_1}} \alpha(x) \leq_{\tilde{S}} \tilde{\prod_{x \in I_2}} \alpha(x) \Rightarrow \prod_{x \in I_1} x \leq_S \prod_{x \in I_2} x$$

where the products refer to the multiplicative operations of the concrete ($\prod$) *and the abstract* ($\tilde{\prod}$) *semirings.*

In words, this notion of order-preservation means that if we first abstract and then combine, or we combine only, we get the same ordering (but on different semirings) among the resulting elements.

EXAMPLE 4.3.5 An abstraction that is not order-preserving can be seen in Figure 4.5. Here, the concrete and the abstract sets, as well as the additive operations of the two semirings, can be seen from the picture. For the multiplicative operations, we assume they coincide with the glb of the two semirings.

In this case, the concrete ordering is partial, while the abstract ordering is total. Functions α and β are depicted in the figure by arrows going from the concrete semiring to the abstract one and vice versa. Assume that the concrete problem has no solution with value 1. Then all solutions with value a or b are optimal. Suppose a solution with value b is obtained by computing $b = 1 \times b$, while we have $a = 1 \times a$. Then the abstract counterparts will have values $\alpha(1) \times' \alpha(b) = 1 \times' 0 = 0$ and $\alpha(1) \times' \alpha(a) = 1 \times' 1 = 1$. Therefore, the solution with value a, which is optimal in the concrete problem, is not optimal anymore in the abstract problem. □

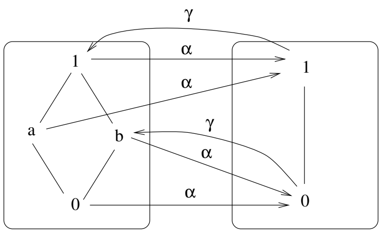

Figure 4.5: An abstraction which is not order-preserving.

EXAMPLE 4.3.6 The abstraction in Figure 4.1 is order-preserving. In fact, consider two abstract values that are ordered, that is $0 \leq' 1$. Then $1 = 1 \times' 1 = \alpha(x) \times' \alpha(y)$, where both x and y must be greater than 0.5. Thus their concrete combination (which is the min), say v, is always greater than 0.5. On the other hand, 0 can be obtained by combining either two 0's (therefore the images of two elements smaller than or equal to 0.5, whose minimum is smaller than 0.5 and thus smaller than v), or by combining a 0 and a 1, which are images of a value greater than 0.5 and one smaller than 0.5. Also in this case, their combination (the min) is smaller than 0.5 and thus smaller than v. Thus the order-preserving property holds. □

EXAMPLE 4.3.7 Another abstraction that is not order-preserving is the one that passes from the semiring $\langle N \cup \{+\infty\}, min, sum, 0, +\infty \rangle$, where we minimize the sum of values over the naturals, to the semiring $\langle N \cup \{+\infty\}, min, max, 0, +\infty \rangle$, where we minimize the maximum of values over the naturals. In words, this abstraction maintains the same domain of the semiring, and the same additive operation (min), but it changes the multiplicative operation (which passes from sum to max). Notice that the semiring orderings are the opposite as those usually used over the naturals: if i is smaller than j then $j \leq_S i$, thus, the best element is 0 and the worst is $+\infty$. The abstraction function α is just the identity, and also the concretization function γ.

In this case, consider in the abstract semiring two values and the way they are obtained by combining other two values of the abstract semiring: for example, $8 = max(7, 8)$ and $9 = max(1, 9)$. In the abstract ordering, 8 is higher than 9. Then, let us see how the images of the combined values (the same values, since α is the identity) relate to each other: we have $sum(7, 8) = 15$ and $sum(1, 9) = 10$, and 15 is lower than 10 in the concrete ordering. Thus the order-preserving property does not hold. □

Notice that, if we reduce the sets I_1 and I_2 to singletons, say x and y, then the order-preserving property says that $\alpha(x) \leq_{\tilde{S}} \alpha(y)$ implies that $x \leq_S y$. This means that if two abstract objects are ordered, then their concrete counterparts will be ordered as well, and in the same way. Of course they could never be ordered in the opposite sense, otherwise α would not be monotonic; but they could be incomparable. Therefore, if we choose an abstraction where incomparable objects are mapped by α onto ordered objects, then we do not have the order-preserving property. A consequence of this is that if the abstract semiring is totally ordered, and we want an order-preserving abstraction, then the concrete semiring will be totally ordered as well.

On the other hand, if two abstract objects are not ordered, then the corresponding concrete objects can be ordered in any sense, or they can also be not comparable. Notice that this restriction of the order-preserving property to singleton sets always holds when the concrete ordering is total. In fact, in this case, if two abstract elements are ordered in a certain way, then it is impossible that the corresponding concrete elements are ordered in the opposite way, because, as we said above, of the monotonicity of the α function.

THEOREM 4.3.8 *Consider an abstraction that is order-preserving. Given an SCSP P over S, we have that $Opt(P) \subseteq Opt(\alpha(P))$.*

PROOF. Let us take a tuple t that is optimal in the concrete semiring S, with value v. Then v has been obtained by multiplying the values of some subtuples. Suppose, without loss of generality, that the number of such subtuples is two (that is, we have two constraints): $v = v_1 \times v_2$. Let us then take the value of this tuple in the abstract problem, that is, the abstract combination of the abstractions of v_1 and v_2: this is $v' = \alpha(v_1) \times' \alpha(v_2)$. We have to show that if v is optimal, then also v' is optimal.

Suppose then that v' is not optimal, that is, there exists another tuple t'' with value v'' such that $v' \leq_{S'} v''$. Assume $v'' = v''_1 \times' v''_2$. Now let us see the value of tuple t'' in P. If we set $v''_i = \alpha(\bar{v_i})$, then we have that this value is $\bar{v} = \bar{v_1} \times \bar{v_2}$. Let us now compare v with $\bar{v}$. Since $v' \leq_{\tilde{S}} v''$, by order-preservation we get that $v \leq_S \bar{v}$. But this means that v is not optimal, which was our initial assumption. Therefore v' has to be optimal. □

Therefore, in case of order-preservation, the set of optimal solutions of the abstract problem contains all the optimal solutions of the concrete problem. In other words, it is not allowed that an optimal solution in the concrete domain becomes non-optimal in the abstract domain. Some non-optimal solutions, however, could become optimal by becoming incomparable with the optimal solutions.

EXAMPLE 4.3.9 Consider again the previous example. The optimal solutions in P are the tuples $\langle a, b, a\rangle$ and $\langle a, b, b\rangle$. It is easy to see that these tuples are also optimal in $\alpha(P)$. In fact, this is a classical constraint problem where the solutions are tuples $\langle a, b, a\rangle$, $\langle a, b, b\rangle$, $\langle b, b, a\rangle$, and $\langle b, b, b\rangle$. □

Thus, if we want to find an optimal solution of the concrete problem, we could find all the optimal solutions of the abstract problem, and then use them on the concrete side to find an optimal solution for the concrete problem. Assuming that working on the abstract side is easier than on the concrete side, this method could help us find an optimal solution of the concrete problem by looking at just a subset of tuples in the concrete problem.

Another important property, which holds for any abstraction, concerns computing bounds that approximate an optimal solution of a concrete problem. In fact, any optimal solution, say t, of the abstract problem, say with value $\tilde{v}$, can be used to obtain both an upper and a lower bound of an optimum in P. In fact, we can prove that there is an optimal solution in P with value between $\gamma(\tilde{v})$ and the value of t in P. Thus, if we think that approximating the optimal value with a value within these two bounds is satisfactory, we can take t as an approximation of an optimal solution of P.

THEOREM 4.3.10 *Given an SCSP P over S, consider an optimal solution of $\alpha(P)$, say t, with semiring value $\tilde{v}$ in $\alpha(P)$ and v in P. Then there exists an optimal solution $\bar{t}$ of P, say with value $\bar{v}$, such that $v \leq \bar{v} \leq \gamma(\tilde{v})$.*

PROOF. By local correctness of the multiplicative operation of the abstract semiring, we have that $v \leq_S \gamma(\tilde{v})$. Since v is the value of t in P, either t itself is optimal in P, or there is another tuple that has value better than v, say $\bar{v}$. We will now show that $\bar{v}$ cannot be greater than $\gamma(\tilde{v})$.

In fact, assume by absurd that $\gamma(\tilde{v}) \leq_S \bar{v}$. Then, by local correctness of the multiplicative operation of the abstract semiring, we have that $\alpha(\bar{v})$ is smaller than the value of $\bar{t}$ in $\alpha(P)$. Also, by monotonicity of α, by $\gamma(\tilde{v}) \leq_S \bar{v}$ we get that $\tilde{v} \leq_{\tilde{S}} \alpha(\bar{v})$. Therefore, by transitivity we obtain that $\tilde{v}$ is smaller than the value of $\bar{t}$ in $\alpha(P)$, which is not possible because we assumed that $\tilde{v}$ was optimal.

Therefore, there must be an optimal value between v and $\gamma(\tilde{v})$. □

Thus, given a tuple t with optimal value $\tilde{v}$ in the abstract problem, instead of spending time to compute an exact optimum of P, we can do the following:

- compute $\gamma(\tilde{v})$, thus obtaining an upper bound of an optimum of P;
- compute the value of t in P, which is a lower bound of the same optimum of P;
- If we think that such bounds are close enough, we can take t as a reasonable approximation (to be precise, a lower bound) of an optimum of P.

Notice that this theorem does not need the order-preserving property in the abstraction, thus any abstraction can exploit its result.

EXAMPLE 4.3.11 Consider again the previous example. Now take any optimal solution of $\alpha(P)$, for example tuple $\langle b, b, b\rangle$. Then the above result states that there exists an optimal solution of P with semiring value v between the value of this tuple in P, which is 0.7, and $\gamma(1) = 1$. In fact, there are optimal solutions with value 1 in P. □

4.3.2 Working on the abstract problem

Consider now what we can do on the abstract problem, $\alpha(P)$. One possibility is to apply an abstract function $\tilde{f}$, which can be, for example, a local consistency algorithm or also a solution algorithm. In the following, we will consider functions $\tilde{f}$, which are always intensive, that is, which bring the given problem closer to the bottom of the lattice. In fact, our goal is to solve an SCSP, thus going higher in the lattice does not help in this task, since solving means combining constraints and thus getting lower in the lattice. Also, functions $\tilde{f}$ will always be locally correct with respect to any function f_{sol} that solves the concrete problem. We will call such a property *solution-correctness*. More precisely, given a problem P with constraint set C, $f_{sol}(P)$ is a new problem P' with the same topology as P whose tuples have semiring values possibly lower. Let us call C' the set of constraints of P'. Then, for any constraint $c' \in C'$, $c' = (\bigotimes C) \Downarrow_{var(c')}$. In other words, f_{sol} combines all constraints of P and then projects the resulting global constraint over each of the original constraints.

DEFINITION 4.3.12 *Given an SCSP P over S, consider a function $\tilde{f}$ on $\alpha(P)$. Then $\tilde{f}$ is solution-correct if, given any f_{sol} which solves P, $\tilde{f}$ is locally correct w.r.t. f_{sol}.*

We will also need the notion of safeness of a function, which just means that it maintains all the solutions.

DEFINITION 4.3.13 *Given an SCSP P and a function $f : PL \to PL$, f is safe if $Sol(P) = Sol(f(P))$.*

It is easy to see that any local consistency algorithm, as defined in [BMR97b], can be seen as a safe, intensive, and solution-correct function.

From $\tilde{f}(\alpha(P))$, applying the concretization function γ, we get $\gamma(\tilde{f}(\alpha(P)))$, which therefore is again over the concrete semiring (the same as P). The following property says that, under certain conditions, P and $P \otimes \gamma(\tilde{f}(\alpha(P)))$ are equivalent. Figure 4.6 describes such a situation. In this figure, we can see that several partial order lines have been drawn:

- on the abstract side, function $\tilde{f}$ takes any element closer to the bottom, because of its intensiveness;
- on the concrete side, we have that
  - $P \otimes \gamma(\tilde{f}(\alpha(P)))$ is smaller than both P and $\gamma(\tilde{f}(\alpha(P)))$ because of the properties of $\otimes$;

– $\gamma(\tilde{f}(\alpha(P)))$ is smaller than $\gamma(\alpha(P))$ because of the monotonicity of γ;

– $\gamma(\tilde{f}(\alpha(P)))$ is higher than $f_{sol}(P)$ because of the solution-correctness of $\tilde{f}$;

– $f_{sol}(P)$ is smaller than P because of the way $f_{sol}(P)$ is constructed;

– if $\times$ is idempotent, then it coincides with the glb, thus we have that $P \otimes \gamma(\tilde{f}(\alpha(P)))$ is higher than $f_{sol}(P)$, because by definition the glb is the highest among all the lower bounds of P and $\gamma(\tilde{f}(\alpha(P)))$.

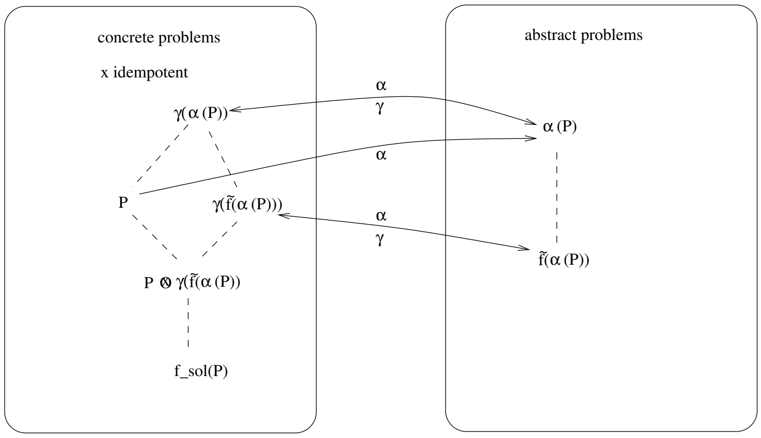

Figure 4.6: The general abstraction scheme, with $\times$ idempotent.

THEOREM 4.3.14 *Given an SCSP P over S, consider a function $\tilde{f}$ on $\alpha(P)$ which is safe, solution-correct, and intensive. Then, if $\times$ is idempotent,* $Sol(P) = Sol(P \otimes \gamma(\tilde{f}(\alpha(P))))$.

PROOF. Take any tuple t with value v in P, which is obtained by combining the values of some subtuples, say two: $v = v_1 \times v_2$. Let us now consider the abstract versions of v_1 and v_2: $\alpha(v_1)$ and $\alpha(v_2)$. Function $\tilde{f}$ changes these values by lowering them, thus we get $\tilde{f}(\alpha(v_1)) = v'_1$ and $\tilde{f}(\alpha(v_2)) = v'_2$.

Since $\tilde{f}$ is safe, we have that $v'_1 \times' v'_2 = \alpha(v_1) \times' \alpha(v_2) = v'$. Moreover, $\tilde{f}$ is solution-correct, thus $v \leq_S \gamma(v')$. By monotonicity of γ, we have that $\gamma(v') \leq_S \gamma(v'_i)$ for $i = 1, 2$. This implies that $\gamma(v') \leq_S (\gamma(v'_1) \times \gamma(v'_2))$, since $\times$ is idempotent by assumption and thus it coincides with the glb. Thus we have that $v \leq_S (\gamma(v'_1) \times \gamma(v'_2))$.

To prove that P and $P \otimes \gamma(\tilde{f}(\alpha(P)))$ give the same value to each tuple, we now have to prove that $v = (v_1 \times \gamma(v'_1)) \times (v_2 \times \gamma(v'_2))$. By commutativity of $\times$, we can write this as $(v_1 \times v_2) \times (\gamma(v'_1) \times \gamma(v'_2))$. Now, $v_1 \times v_2 = v$ by assumption, and we have shown that $v \leq_S \gamma(v'_1) \times \gamma(v'_2)$. Therefore $v \times (\gamma(v'_1) \times \gamma(v'_2)) = v$. □

This theorem does not say anything about the power of $\tilde{f}$, which could make many modifications to $\alpha(P)$, or it could also not modify anything. In this last case, $\gamma(\tilde{f}(\alpha(P))) = \gamma(\alpha(P)) \sqsupseteq P$ (see Figure 4.7), so $P \otimes \gamma(\tilde{f}(\alpha(P))) = P$, which means

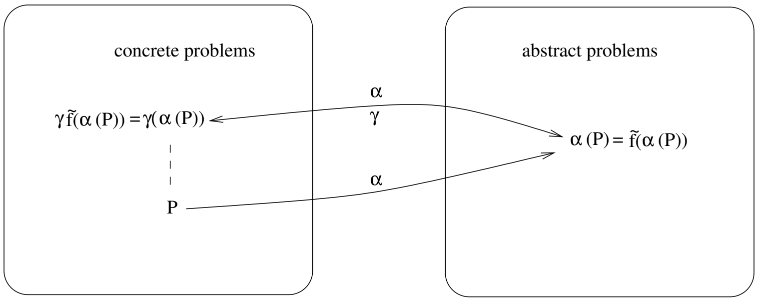

Figure 4.7: The scheme when $\tilde{f}$ does not modify anything.

that we have not gained anything in abstracting P. We can, however, always use the relationship between P and $\alpha(P)$ (see Theorem 4.3.8 and 4.3.10) to find an approximation of the optimal solutions and of the inconsistencies of P.

If instead $\tilde{f}$ modifies all semiring elements in $\alpha(P)$, then if the order of the concrete semiring is total, we have that $P \otimes \gamma(\tilde{f}(\alpha(P))) = \gamma(\tilde{f}(\alpha(P)))$ (see Figure 4.8), and thus we can work on $\gamma(\tilde{f}(\alpha(P)))$ to find the solutions of P. In fact, $\gamma(\tilde{f}(\alpha(P)))$ is lower than P and thus closer to the solution.

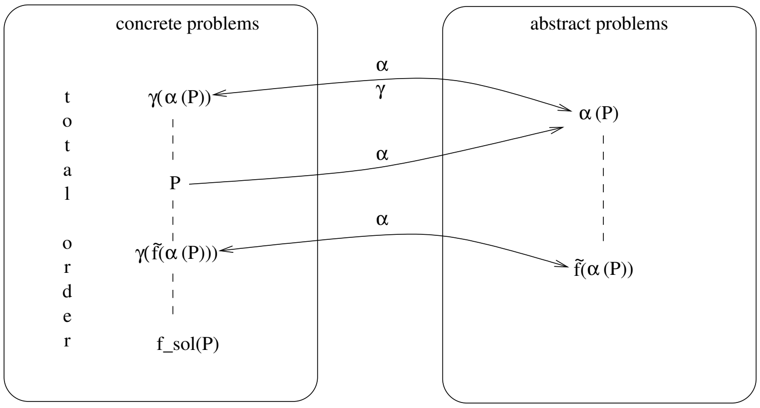

Figure 4.8: The scheme when the concrete semiring has a total order.

THEOREM 4.3.15 *Given an SCSP P over S, consider a function $\tilde{f}$ on $\alpha(P)$ which is safe, solution-correct, and intensive. Then, if $\times$ is idempotent, $\tilde{f}$ modifies every semiring element in $\alpha(P)$, and the order of the concrete semiring is total, we have that $P \sqsupseteq_S \gamma(\tilde{f}(\alpha(P))) \sqsupseteq_S f_{sol}(P)$.*

PROOF. Consider any tuple t in any constraint of P, and let us call v its semiring value in P and v_{sol} its value in $f_{sol}(P)$. Obviously, we have that $v_{sol} \leq_S v$. Now take $v' = \gamma(\alpha(v))$. By monotonicity of α, we cannot have $v <_S v'$. Also, by solution-correctness of $\tilde{f}$, we cannot have $v' <_S v_{sol}$. Thus, we must have $v_{sol} \leq_S v' \leq_S v$, which proves the statement of the theorem. □

EXAMPLE 4.3.16 Figure 4.9 uses the abstraction in Figure 4.1 and shows a concrete problem and the result of the construction of Figure 4.6 over it.

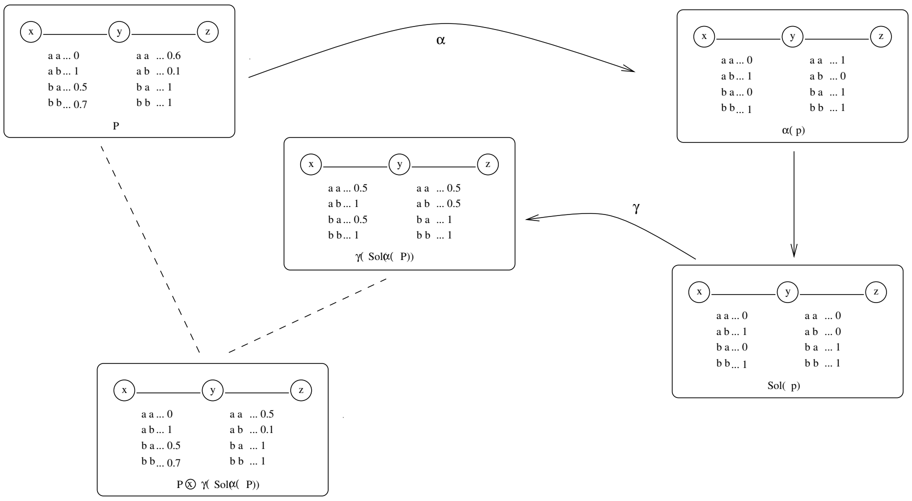

Figure 4.9: An example with $\times$ idempotent.

□

Notice that we need the idempotency of the $\times$ operator for Theorem 4.3.14 and 4.3.15. If instead $\times$ is not idempotent, then we can prove something weaker. Figure 4.10 shows this situation. With respect to Figure 4.6, we can see that the possible non-idempotency of $\times$ changes the partial order relationship on the concrete side. In particular, we do not have the problem $P \otimes \gamma(\tilde{f}(\alpha(P)))$ any more, nor the problem $f_{sol}(P)$, since these problems would not have the same solutions as P and thus are not interesting to us. We have instead a new problem P', which is constructed in such a way as to "insert" the inconsistencies of $\gamma(\tilde{f}(\alpha(P)))$ into P. P' is obviously lower than P in the concrete partial order, since it is the same as P with the exception of some more 0's, but the most important point is that it has the same solutions as P.

THEOREM 4.3.17 *Given an SCSP P over S, consider a function $\tilde{f}$ on $\alpha(P)$ which is safe, solution-correct and intensive. Then, if $\times$ is not idempotent, consider P' to be the SCSP which is the same as P except for those tuples which have semiring value 0 in $\gamma(\alpha(\tilde{f}(P)))$: these tuples are given value 0 also in P'. Then we have that $\mathrm{Sol}(P) = \mathrm{Sol}(P')$.*

PROOF. Take any tuple t with value v in P, which is obtained by combining the values of some subtuples, say two: $v = v_1 \times v_2$. Let us now consider the abstract versions of v_1 and v_2: $\alpha(v_1)$ and $\alpha(v_2)$. Function $\tilde{f}$ changes these values by lowering them, thus we get $\tilde{f}(\alpha(v_1)) = v'_1$ and $\tilde{f}(\alpha(v_2)) = v'_2$.

Since $\tilde{f}$ is safe, we have that $v'_1 \times' v'_2 = \alpha(v_1) \times' \alpha(v_2) = v'$. Moreover, $\tilde{f}$ is solution-correct, thus $v \leq_S \gamma(v')$. By monotonicity of γ, we have that $\gamma(v') \leq_S \gamma(v'_i)$ for $i = 1, 2$. Thus we have that $v \leq_S \gamma(v'_i)$ for $i = 1, 2$.

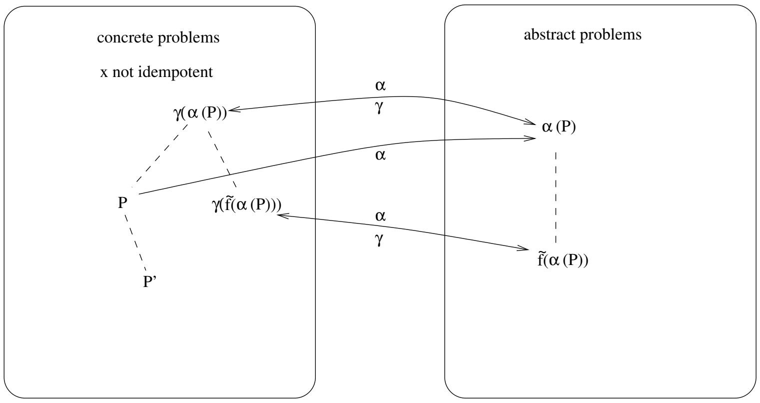

Figure 4.10: The scheme when $\times$ is not idempotent.

Now suppose that $\gamma(\nu'_1) = 0$. This implies that also $\nu = 0$. Therefore, if we set $\nu_1 = 0$, again the combination of ν_1 and ν_2 will result in ν, which is 0. □

EXAMPLE 4.3.18 Consider the abstraction from the semiring $S = \langle Z^- \cup \{-\infty\}, max, +, -\infty, 0\rangle$ to the semiring $S' = \langle Z^- \cup \{-\infty\}, max, min, -\infty, 0\rangle$, where α and γ are the identity. This means that we perform the abstraction just to change the multiplicative operation, which is min instead of +. Then Figure 4.11 shows a concrete problem over S, and the construction shown in Figure 4.10 over it.

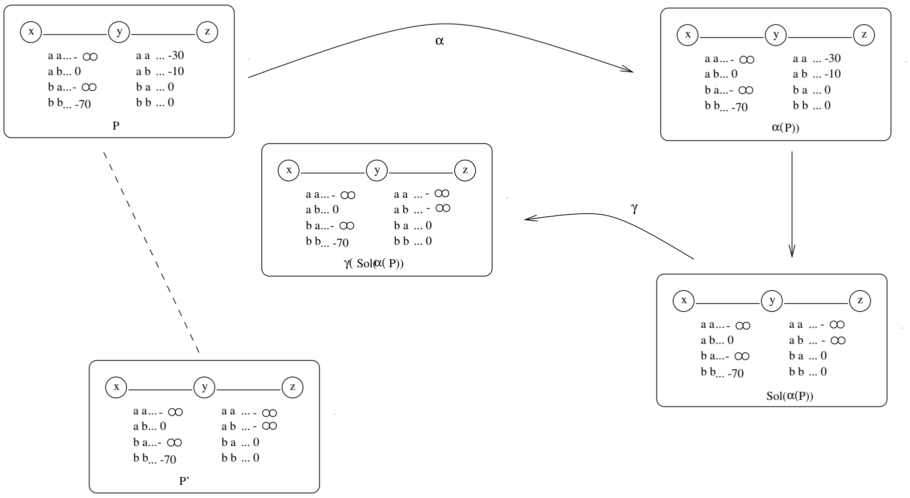

Figure 4.11: An example with $\times$ not idempotent.

□

Summarizing, the above theorems can give us several hints on how to use the abstraction scheme to make the solution of P easier: If $\times$ is idempotent, then we can replace P with $P \otimes \gamma(\alpha(\tilde{f}(P)))$, and get the same solutions (by Theorem 4.3.14). If instead $\times$ is not idempotent, we can replace P with P' (by Theorem 4.3.17). In any case, the point in passing from P to $P \otimes \gamma(\alpha(\tilde{f}(P)))$ (or P') is that the new problem should be easier to solve than P, since the semiring values of its tuples are more explicit, that is, closer to the values of these tuples in a completely solved problem.

More precisely, consider a branch-and-bound algorithm to find an optimal solution of P. Then, once a solution is found, its value will be used to cut away some branches, where the semiring value is worse than the value of the solution already found. Now, if the values of the tuples are worse in the new problem than in P, each branch will have a worse value and thus we might cut away more branches. For example, consider the fuzzy semiring (that is, we want to maximize the minimum of the values of the subtuples): if the solution already found has value 0.6, then each partial solution of P with value smaller than or equal to 0.6 can be discarded (together with all its corresponding subtree in the search tree), but all partial solutions with value greater than 0.6 must be considered; if instead we work in the new problem, the same partial solution with value greater than 0.6 may now have a smaller value, possibly also smaller than 0.6, and thus can be disregarded. Therefore, the search tree of the new problem is smaller than that of P.

Another point to notice is that, if using a greedy algorithm to find the initial solution (to use later as a lower bound), this initial phase in the new problem will lead to a better estimate, since the values of the tuples are worse in the new problem and thus close to the optimum. In the extreme case in which the change from P to the new problem brings the semiring values of the tuples to coincide with the value of their combination, it is possible to see that the initial solution is already the optimal one.

Notice also that, if $\times$ is not idempotent, a tuple of P' has either the same value as in P, or $\mathbf{0}$. Thus, the initial estimate in P' is the same as that of P (since the initial solution must be a solution), but the search tree of P' is again smaller than that of P, since there may be partial solutions, which in P have value different from $\mathbf{0}$ and in P' have value $\mathbf{0}$, and thus the global inconsistency may be recognized earlier.

The same reasoning used for Theorem 4.3.8 on $\alpha(P)$ can also be applied to $\tilde{f}(\alpha(P))$. In fact, since $\tilde{f}$ is safe, the solutions of $\tilde{f}(\alpha(P))$ have the same values as those of $\alpha(P)$. Thus, also the optimal solution sets coincide. Therefore, we have that $\mathtt{Opt}(\tilde{f}(\alpha(P)))$ contains all the optimal solutions of P if the abstraction is order-preserving. This means that, in order to find an optimal solution of P, we can find all optimal solutions of $\tilde{f}(\alpha(P))$ and then use such a set to prune the search for an optimal solution of P.

THEOREM 4.3.19 *Given an SCSP P over S, consider a function $\tilde{f}$ on P which is safe, solution-correct and intensive, and let us assume the abstraction is order-preserving. Then we have that* $\mathtt{Opt}(P) \subseteq \mathtt{Opt}(\tilde{f}(\alpha(P)))$.

PROOF. Easy follows from Theorem 4.3.8 and from the safeness of $\tilde{f}$. □

Theorem 4.3.10 can be adapted to $\tilde{f}(\alpha(P))$ as well, thus allowing us to use an optimal solution of $\tilde{f}(\alpha(P))$ to find both a lower and an upper bound of an optimal solution of P.

4.4 Some abstraction mappings

In this section we will list some semirings and several abstractions between them, in order to provide the reader with a scenario of possible abstractions that he/she can use, starting from one of the semirings considered here. Some of these semirings and/or abstractions have been already described in the previous sections of the chapter; however, here we will re-define them to make this section self-contained. Of course many other semirings could be defined, but here we focus on the ones for which either it has been defined, or it is easy to imagine, a system of constraint solving. The semirings we will consider are the following ones:

- the classical one, which describes classical CSPs via the use of logical and and logical or:

$$S_{CSP} = \langle \{T, F\}, \vee, \wedge, F, T \rangle$$

- the fuzzy semiring, where the goal is to maximize the minimum of some values over $[0, 1]$:

$$S_{fuzzy} = \langle [0, 1],\ \max,\ \min, \mathbf{0}, \mathbf{1} \rangle$$

- the extension of the fuzzy semiring over the naturals, where the goal is to maximize the minimum of some values of the naturals:

$$S_{fuzzyN} = \langle N \cup \{+\infty\},\ \max,\ \min, 0, +\infty \rangle$$

- the extension of the fuzzy semiring over the positive reals:

$$S_{fuzzyR} = \langle R^+ \cup \{+\infty\},\ \max,\ \min, 0, +\infty \rangle$$

- the optimization semiring over the naturals, where we want to maximize the sum of costs (which are negative integers):

$$S_{optN} = \langle Z^- \cup \{-\infty\}, \max, +, -\infty, 0 \rangle$$

- the optimization semiring over the negative reals:

$$S_{optR} = \langle R^- \cup \{-\infty\}, \max, +, -\infty, 0 \rangle$$

- the probabilistic semiring, where we want to maximize a certain probability which is obtained by multiplying several individual probabilities. The idea here is that each tuple in each constraint has associated with it the probability of being allowed in the real problems we are modeling, and different tuples in different constraints have independent probabilities (so that their combined probability is just the multiplication of their individual probabilities) [FL93]. The semiring is:

$$S_{prob} = \langle [0,1],\ \max, \times, 0, 1 \rangle$$

- the subset semiring, where the elements are all the subsets of a certain set, the operations are set intersection and set union, the smallest element is the empty set, and the largest element is the whole given set:

$$S_{sub} = \langle \mathcal{P}(A), \bigcup, \bigcap, \emptyset, A \rangle.$$

We will now define several abstractions between pairs of these semirings. The result is drawn in Figure 4.12, where the dashed lines denote the abstractions that have not been defined but can be obtained by abstraction composition.

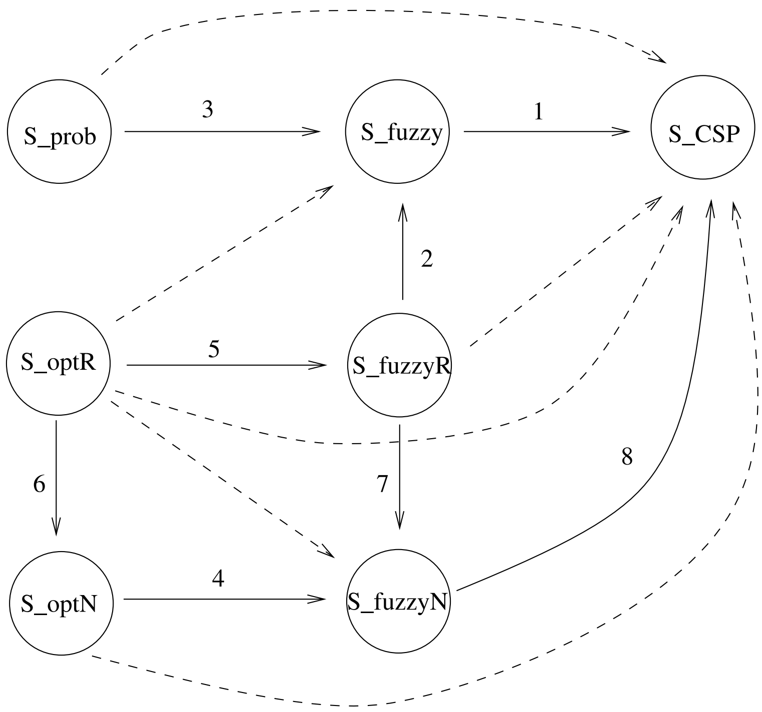

Figure 4.12: Several semirings and abstractions between them.

In reality, each line in the figure represents a whole family of abstractions, since each $\langle \alpha, \gamma \rangle$ pair makes a specific choice which identifies a member of the family. Moreover, by defining one of this families of abstractions we do not want to say that there do not exist other abstractions between the two semirings.

It is easy to see that some abstractions focus on the domain, by passing to a given domain to a smaller one, others change the semiring operations, and others change both:

1. from fuzzy to classical CSPs: this abstraction changes both the domain and the operations. The abstraction function is defined by choosing a threshold within the interval $[0,1]$, say x, and mapping all elements in $[0,x]$ to F and all elements in $(x,1]$ to T. Consequently, the concretization function maps T to 1 and F to x. See Figure 4.1 as an example of such an abstraction. We recall that all the abstractions in this family are order-preserving, so Theorem 4.3.8 can be used.

2. from fuzzy over the positive reals to fuzzy CSPs: this abstraction changes only the domain, by mapping the whole set of positive reals into the $[0,1]$ interval. This means that the abstraction function has to set a threshold, say x, and map all reals above x into 1, and any other real, say r, into $\frac{r}{x}$. Then, the concretization function will map 1 into $+\infty$, and each element of $[0,1)$, say y, into $y \times x$. It is easy to prove that all the members of this family of abstractions are order-preserving.

3. from probabilistic to fuzzy CSPs: this abstraction changes only the multiplicative operation of the semiring, that is, the way constraints are combined. In fact, instead of multiplying a set of semiring elements, in the abstracted version we choose the minimum value among them. Since the domain remains the same, both the abstraction and the concretization functions are the identity (otherwise they would not have the properties required by a Galois insertion, like monotonicity). Thus this family of abstractions contains just one member.

   It is easy to see that this abstraction is not order-preserving. In fact, consider for example the elements 0.6 and 0.5, obtained in the abstract domain by $0.6 = min(0.7, 0.6)$ and $0.5 = min(0.9, 0.5)$. These same combinations in the concrete domain would be $0.7 \times 0.6 = 0.42$ and $0.9 \times 0.5 = 0.45$, thus resulting in two elements which are in the opposite order with respect to 0.5 and 0.6.

4. from optimization-N to fuzzy-N CSPs: here the domain remains the same (the negative integers) and only the multiplicative operation is modified. Instead of summing the values, we want to take their minimum. As noted in a previous example, these abstractions are not order-preserving.

5. from optimization-R to fuzzy-R CSPs: similar to the previous one but on the negative reals.

6. from optimization-R to optimization-N CSPs: here we have to map the negative reals into the negative integers. The operations remain the same. A possible example of abstraction is the one where $\alpha(x) = \lceil x \rceil$ and $\gamma(x) = x$. It is not order-preserving.

7. from fuzzy-R to fuzzy-N CSPs: again, we have to map the positive reals into the naturals, while maintaining the same operations. The abstraction could be the same

as before, but in this case it is order-preserving (because of the use of min instead of sum).

8. from fuzzy-N to classical CSPs: this is similar to the abstraction from fuzzy CSPs to classical ones. The abstraction function has to set a threshold, say x, and map each natural in $[0, x]$ into F, and each natural above x into T. The concretization function maps T into $+\infty$ and F into x. All such abstractions are order-preserving.

9. from subset CSPs to any of the other semirings: if we want to abstract to a semiring with domain A, we start from the semiring with domain $\mathcal{P}(A)$. The abstraction mapping takes a set of elements of A and has to choose one of them by using a given function, for example min or max. The concretization function will then map an element of A into the union of all the corresponding sets in $\mathcal{P}(A)$. For reasons similar to those used in Example 3, some abstractions of this family may be not order-preserving.

4.5 Abstraction vs. local consistency

It is now interesting to consider the relationship between our abstraction framework and the concept of local consistency.

In fact, it is possible to show that, given an abstraction $\langle \alpha, \gamma \rangle$ between semirings S and $\bar{S}$ and any propagation rule r in $\bar{S}$, the function $\gamma(r(\alpha(P))) \otimes P$ is a propagation rule for problem P over S. This can be convenient when S does not have any, or any efficient, propagation algorithms. In fact, in such cases, we can resort to the propagation algorithms of $\bar{S}$ to perform propagation also over S.

Notice, however, that, when S has a non-idempotent multiplicative operator, function $\gamma(r(\alpha(P))) \otimes P$ could change the solution of P. To avoid this problem, we just have to follow the same reasoning as in the previous section, that is, to replace such a function with a function that just inserts into P the inconsistencies of $\gamma(r(\alpha(P))$. We will denote such a function by using a different combination operator: $\otimes_0$. Thus, the function to be used in these cases is $\gamma(r(\alpha(P)) \otimes_0 P$. Notice that $\otimes_0$ is a non-commutative operator, since it inserts into the right operand the zeros of the left operand.

This results, however, hold only when the abstraction is order-preserving. We recall that this means that applying the abstraction function and then combining gives elements that are in the same ordering as the elements obtained by combining only. In particular, if two abstract elements $\alpha(x)$ and $\alpha(y)$ are ordered, then also x and y are ordered as well, and in the same direction.

THEOREM 4.5.1 *Given an order-preserving abstraction $\langle \alpha, \gamma \rangle$ between semiring S and $\bar{S}$, assume that S has an idempotent multiplicative operation and consider any propagation rule r in $\bar{S}$ and any problem P over S. Then the function $f(P) = \gamma(r(\alpha(P))) \otimes P$ is a propagation rule for P.*

PROOF. By definition, a propagation rule is an intensive, monotone, and idempotent function that takes a problem and returns an equivalent problem over the same semiring. Since $\otimes$ is intensive, also f is so. Moreover, by monotonicity of γ, r, α, and $\otimes$, also f is monotone.

For proving idempotency of f, we need the order-preserving property of the abstraction. In fact, consider what happens when applying function f to P: some tuple values in $\alpha(P)$, say $\bar{v} = \alpha(v)$, will not be changed by r, while others will receive a lower value, say $\bar{v}' = r(\bar{v})$. By order-preservation, the new tuples values in the concrete semiring (that is, $\gamma(r(\alpha(v))) \times v$), are equal or lower than the original values. Let us now apply function f again. Function α will bring these new concrete values to either $\bar{v}$ (if we start from v) or $\bar{v}'$ (if we start from $\gamma(r(\alpha(v)))$). In any case, r will bring such values to $\bar{v}'$, for Lemma 4.5.2 (see below). Thus f is idempotent. Finally, f returns an equivalent problem by the theorem depicted in Figure 4.6. □

LEMMA 4.5.2 *Consider any SCSP P over S and any propagation rule r for P, with $r(P) = P'$. Then, taken any SCSP P'' such that $P' \leq_S P'' \leq_S P$, we have $r(P'') = P'$.*

PROOF. Any rule r solves a subproblem $\langle C, con\rangle$ and changes the values of the tuples connecting the variables in con. Thus, the result of applying r is a new constraint over con: $(C \otimes C_{con}) \Downarrow_{con}$, where C_{con} is the original constraint connecting the variables in con. This can also be written as $C_{con} \otimes C \Downarrow_{con}$. Let us now take any C''_{con} such that $(C_{con} \otimes C \Downarrow_{con}) \leq_S C''_{con} \leq_S C_{con}$. We can now multiply all these three constraints by $C \Downarrow_{con}$, obtaining: $(C_{con} \otimes C \Downarrow_{con} \otimes C \Downarrow_{con}) \leq_S (C''_{con} \otimes C \Downarrow_{con}) \leq_S (C_{con} \otimes C \Downarrow_{con})$. By idempotency of $\otimes$, we get: $(C_{con} \otimes C \Downarrow_{con}) \leq_S (C''_{con} \otimes C \Downarrow_{con}) \leq_S (C_{con} \otimes C \Downarrow_{con})$. Thus we have that $(C_{con} \otimes C \Downarrow_{con}) = (C''_{con} \otimes C \Downarrow_{con})$. □

We can now prove a similar result for the case of a non-idempotent multiplicative operation in the concrete semiring. As noted above, however, we cannot combine the new problem with the old one, but we can just insert the zeroes of the new problem into the old one.

THEOREM 4.5.3 *Given an order-preserving abstraction $\langle \alpha, \gamma\rangle$ between semiring S and $\bar{S}$, assume that S has a non-idempotent multiplicative operation and consider any propagation rule r in $\bar{S}$ and any problem P over S. Then the function $f(P) = \gamma(r(\alpha(P))) \otimes_0 P$ is a propagation rule for P, where $\otimes_0$ inserts the zeroes of its left operand into the right one.*

PROOF. The proof of this theorem is similar to the previous one, and for the equivalence it refers also to the theorem depicted in Figure 4.10. □

By taking several propagation rules in the abstract semiring, we can thus obtain an equal number of propagation rules over the concrete semiring. This set of rules can then be used to perform constraint propagation over a concrete problem. Notice, however, that, while an idempotent multiplicative operation in the concrete semiring allows us to use such rules until stability, with all the desired properties (equivalence, uniqueness, and termination), in the case of a non-idempotent multiplicative operation we can just apply the various propagation rules once each to insert several zeroes into the original problem (with the same properties as above).

4.6 Related work

We will compare here our work to other abstraction proposals, more or less related to the concepts of constraints.

Abstracting valued CSPs. The only other abstraction scheme for soft constraint problems we are aware of is the one in [SdGS97], where *valued CSPs* (see Section 2.3.7) are abstracted in order to produce good lower bounds for the optimal solutions. The concept of valued CSPs is similar to our notion of SCSPs. In fact, in valued CSPs, the goal is to minimize the value associated to a complete assignment. In valued CSPs, each constraint has one associated element, not one for each tuple of domain values of its variables; however, our notion of soft CSPs and that in valued CSPs are just different formalizations of the same idea, since one can pass from one formalization to the other one without changing the solutions, provided that the partial order is total (see Section 2.3.7). However, our abstraction scheme is different from the one in [SdGS97]. In fact, we are not only interested in finding good lower bounds for the optimum, but also in finding the exact optimal solutions in a shorter time. Moreover, we do not define *ad hoc* abstraction functions but we follow the classical abstraction scheme devised in [CC77], with Galois insertions to relate the concrete and the abstract domain, and locally correct functions on the abstract side. We think that this is important in that it allows to inherit many properties that have already been proven for the classical case. It is also worth noticing that our notion of an order-preserving abstraction is related to their concept of aggregation compatibility, although generalized to deal with partial orders.

Abstracting classical CSPs. Other work related to abstracting constraint problems proposed the abstraction of the domains [Fre91, FS95, SM92], or of the graph topology (for example to model a subgraph as a single variable or constraint) [Ell93]. We did not focus on these kinds of abstractions for SCSPs in this chapter, but we believe that they could be embedded into our abstraction framework: we just need to define the abstraction function in such a way that not only we can change the semiring but also any other feature of the concrete problem. The only difference will be that we cannot define the concrete and abstract lattices of problems by simply extending the lattices of the two semirings.

A general theory of abstraction. A general theory of abstraction has been proposed in [WG92b]. The purpose of this work is to define a notion of abstraction that can be applied to many domains: from planning to problem solving, from theorem proving to decision procedures. Then, several properties of this notion are considered and studied. The abstraction notion proposed consists of just two formal systems Σ_1 and Σ_2 with languages L_1 and L_2 and an effective total function $f : L_1 \rightarrow L_2$, and it is written as $f : \Sigma_1 \Rightarrow \Sigma_2$. Much emphasis is placed in [WG92b] onto the study of the properties that are preserved by passing from the concrete to the abstract system. In particular, one property that appears to be very desirable, and present in most abstraction frameworks, is that *what is a* *theorem in the concrete domain, remains a theorem in the abstract domain* (called the **TI** property, for *Theorem Increasing*).

It is easy to see that our definition of abstraction is an instance of this general notion. Then, to see whether our concept of abstraction has this property, we first must say what is a theorem in our context. A natural and simple notion of a theorem could be an SCSP which has at least one solution with a semiring value different from the 0 of the semiring. We can be more general than this, however, and say that a theorem for us is *an SCSP which has a solution with value greater than or equal to* k*, where* $k \geq 0$. Then we can prove our version of the TI property:

THEOREM 4.6.1 (OUR TI PROPERTY) *Given an SCSP* P *which has a solution with value* $v \geq k$*, then the SCSP* $\alpha(P)$ *has a solution with value* $v' \geq \alpha(k)$.

PROOF. Take any tuple t in P with value $v > k$. Assume that $v = v_1 \times v_2$. By abstracting, we have $v' = \alpha(v_1) \times' \alpha(v_2)$. By solution correctness of $\times'$, we have that $v \leq_S \gamma(v')$. By monotonicity of α, we have that $\alpha(v) \leq_{S'} \alpha(\gamma(v')) = v'$. again by monotonicity of α, we have $\alpha(k) \leq_{S'} \alpha(v)$, thus by transitivity $\alpha(k) \leq_S v'$. □

Notice that, if we consider the boolean semiring (where a solution has either value true or false), this statement reduces to saying that if we have a solution in the concrete problem, then we also have a solution in the abstract problem, which is exactly what the TI property says in [WG92b]. Thus, our notion of abstraction, as defined in the previous sections, on the one hand can be cast within the general theory proposed in [WG92b], while on the other hand it generalizes it to concrete and abstract domains, which are more complex than just the boolean semiring. This is predictable, because, while in [WG92b] formulas can be either true (thus theorems) or false, here they may have any level of satisfaction, which can be described by the given semiring.

Notice also that, in our definition of abstraction of an SCSP, we have chosen to have a Galois insertion between the two lattices $\langle A, \leq \rangle$ (which corresponds to the concrete semiring S) and $\langle \tilde{A}, \tilde{\leq} \rangle$ (which corresponds to the abstract semiring $\tilde{S}$). This means that the ordering in the two lattices coincide with those of the semirings. We could have chosen differently: for example, that the ordering of the lattices in the abstraction be the opposite of those of those in the semirings. In that case, we would not have had property **TI**. We would, however, have the dual property (called **TD** in [WG92b]), which states that abstract theorems remain theorems in the concrete domain. It has been shown that such a property can be useful in some application domains, such as databases.

4.7 Conclusions

In this chapter we have described how to apply *abstraction* to SCSPs. The main motivation to apply such a transformation is to obtain an easier representation of the problem, or (and many times) to obtain an SCSP whose semiring satisfied the needed properties for preprocessing algorithms applications.

Chapter 5

Higher Order Semiring-based Constraints

Abstract

Semiring-based constraint problems (SCSPs), as described in Chapter 2, extend classical constraint problems (CSPs) by allowing preferences, costs, priorities, probabilities, and other soft features. They are based over the notion of a semiring, that is, a set plus two operations. In this chapter we introduce a uniform, abstract presentation of (soft) constraint satisfaction concepts and constructions. Moreover, the soft constraint environment is enriched with the operation of function abstraction and application that can suitably be used to give modularity and compositionality to the framework. Finally, a small language is defined and used to represent and give semantics to general local propagation and dynamic programming algorithms.

In this chapter the basic structures of the framework are represented as functions, and this gives modularity and compositionality. Many of the algorithms defined over a constraint system, like constraint propagation and dynamic programming techniques, can easily be expressed using expressions and commands of a suitable language.

The basic elements that we need to consider are the *domains* and the *semirings*; the whole structure of an SCSP will be represented as a function between these elements. As soon as we define constraints and SCSPs as specific functions, we can use abstraction, application and composition to model SCSP solution, dynamic programming, and local consistency preprocessing algorithms.

The representation of constraints as functions gives the possibility to represent a lot of useful features present in real problems. By performing a step-by-step application of the definitions of the constraints to the resolution function, we can model environments where the constraints are specified at run time by the user, or we can compute (partial) solutions for partially instantiated problems.

The chapter is organized as follows. First, in Section 5.1 we define the domain and semiring structure and we illustrate how constraints and SCSPs are represented as specific functions; in Section 5.2 we define a small language useful to represent algorithms over SCSPs; and in Section 5.3 we define the solution and the preprocessing algorithms as specific expressions and commands of the language. Finally in Section 5.4 we give some possible extension of the framework.

This chapter is based on some preliminary work presented in [BMR00a].

5.1 Domains and semirings

All the SCSP's elements (constraints and problems) are defined here as functions over the set of domains D and semirings S. A domain can be a set D_0 (representing all the values that can be assigned to variables) or can be constructed using the cartesian product. Note that the domain of the variables may be different, so we have a collection $D_1, \ldots, D_n$ of variable domains.

DEFINITION 5.1.1 (DOMAINS) *The class of domains is denoted* D *and it is defined as*

$$D ::= D_1 \mid \cdots \mid D_n \mid D \times D$$

Intuitively, a value $d \in D$ represents a domain value for a variable in a given SCSP, while a value $\langle d_1, d_2 \rangle \in D_1 \times D_2$ and a value $\langle d_1, \ldots, d_k \rangle \in D_1 \times \cdots \times D_k$ respectively represents a pair of domain values for a binary constraint and a k-tuple of domain values for a k-ary constraint.

The semiring structure will be represented as S; note that S represents all the primitive instances (for example the structures used for *fuzzy-CSPs*, *probabilistic-CSPs*, *weighted-CSPs*, etc.) and also the semiring obtained by using some special operators like the cartesian product (described in Section 2.3.8). In this case there is a collection of basic semiring $S_1, \ldots, S_n$ (see Chapter 2) and their cartesian product.

DEFINITION 5.1.2 (SEMIRING) *The class of semirings is denoted by* S *and is defined as*

$$S ::= S_1 \mid \cdots \mid S_n \mid S \times S$$

By using the properties of $\times$, $+$ and of the cartesian product, we can write down some basic relationships between semiring elements; In Table 5.1 symbol $\star$ denotes any of the semiring operators $\times$ and $+$. The table shows how the semiring operations of the cartesian product is defined using the corresponding operations in the primitive semiring.

By using functions from D to S we can now represent (soft) constraints.

DEFINITION 5.1.3 (CONSTRAINTS) *The set of constraints (over a semiring* S*) is denoted by* C *and it is defined as*

$$C ::= D \rightarrow S$$

$$\frac{a_1 \in S_1, a_2 \in S_2}{\langle a_1, a_2 \rangle \in S_1 \times S_2}$$

$$\frac{\langle a_1, a_2 \rangle, \langle b_1, b_2 \rangle \in S_1 \times S_2}{\langle a_1, a_2 \rangle \star_{S_1 \times S_2} \langle b_1, b_2 \rangle = \langle a_1 \star_{S_1} b_1, a_2 \star_{S_2} b_2 \rangle \in S_1 \times S_2}$$

Table 5.1: Relationships between semiring elements.

Note that the syntactic category D represents the set of domains of all the constraints of the problem. For unary constraints (that is, variable domain), there will be functions $D_i \rightarrow S$ for some i, and for the non unary constraint over $k \geq 1$ variables there will be a function $D_1 \times \cdots \times D_k \rightarrow S$.

We can now describe entire constraint problems:

DEFINITION 5.1.4 (CONSTRAINT PROBLEM) *Fixed a semiring S and the number n of constraints, a constraint problem P is defined as*

$$P ::= C^n \rightarrow C.$$

where C^n is obtained as the cartesian product of n constraints, that is $C^n = \underbrace{C \times \cdots \times C}_{n\ times}$

EXAMPLE 5.1.5 Consider the fuzzy CSP represented in Figure 5.1. In this case the semiring for the problem is

$$S_{FCSP} = \langle [0, 1], max, min, 0, 1 \rangle.$$

The domain is the same for all the variables that is $D = \{a, b\}$. The constraints c_x, c_y, c_z are functions $\{a, b\} \rightarrow S_{FSCP}$,and constraints c_1 and c_2 are functions $\{a, b\}^2 \rightarrow S_{FSCP}$.

The problem represented in Figure 5.1 is instead a function $c_x \times c_y \times c_z \times c_1 \times c_2 \rightarrow (D^2 \rightarrow S_{FSCSP})$.

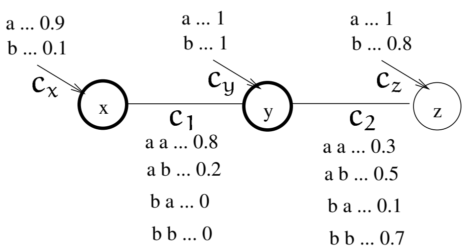

Figure 5.1: A 3-variable fuzzy CSP.

□

5.1.1 Constraint problems and solutions

Since both constraints and constraint problems are defined as functions, we are now able to use function composition, abstraction and application to define in a general and easy way many solution or preprocessing algorithms. The definition of a constraint problem P can be split in two parts: the first one is called *constraint graph* and describes its *topology*, and the second one is called *solution* and represents its *value*. Note that problems with different topology can have the same value, and that many times the same topology corresponds to different values.

To define the constraint graph in a concise way we need to introduce some special operators and properties:

DEFINITION 5.1.6 (NEEDED OPERATORS AND PROPERTIES)

- *The* restriction *operator* ν *will be useful to capture the variables of interest of a problem:*

  $$\frac{f \in D \to S}{\nu f = \sum_{x \in D} f(x)}$$

  We will write $\nu x.M$ *instead of* $\nu(\lambda x.M)$ *and* $\nu\langle x, y\rangle$ *instead of* $\nu x.\nu y$*, with the meaning of restricting over pairs (or tuples) of domains instead of single domains.*

- *some of its important properties*

  $$\nu x.(t_1 \star t_2) = (\nu x.t_1) \star t_2 \text{ if } x \notin FV(t_2) \qquad (\nu\text{-distributivity})$$
  $$\nu x.(t_1 +_S t_2) = (\nu x.t_1) +_S (\nu x.t_2) \qquad (\nu\text{-associativity})$$

Let us define now the *constraint graph* and the *problem solution*.

We also remind (see Section 2.2) that a a constraint system is defined as a tuple $CS = \langle S, D, V\rangle$, where S is a c-semiring, D is a finite set, and V is an ordered set of variables. A constraint over CS is a pair $\langle def, con\rangle$, where

- $con \subseteq V$, it is called the *type* of the constraint;
- $def : D^k \to A$ (where k is the cardinality of con) is called the *value* of the constraint.

DEFINITION 5.1.7 (CONSTRAINT GRAPH) *Given an SCSP* $P = \langle C, con\rangle$ *over a constraint system* $CS = \langle S, D, V\rangle$ *with* $C = \{c_1, \ldots, c_n\}$*,* $c_i = \langle def_i, con_i\rangle$*,* $def_i : D^{|con_i|} \to S$ *and* $V = \{x_1, \ldots, x_k\}$*, a* constraint graph *is a function defined as follows:*

$$P = \lambda\langle c_1, \ldots, c_n\rangle.\lambda\langle con\rangle.\nu\langle V - con\rangle. \prod_{i=1\ldots n} c_i(con_i)$$

The meaning of this function is to take the constraints $\{c_1, \dots, c_n\}$ over the variables in V and to perform their combination $\prod_{S\,i=1\dots n} c_i(con_i)$. Since we are interested only to a subset of the variable $con \in V$, we need to use the ν operator to sum up over the variable in $V - con$. This is not completely correct because con and $V - con$ are sets of variables and not tuples. Anyway, since we fix a lexicographic order over the variables in V it is easy to pass from sets to tuples.

The *value* of an SCSP is obtained applying the constraint graph to the definition of its constraints. The result of this application represents the *problem solution*.

DEFINITION 5.1.8 (PROBLEM SOLUTION) *Given an SCSP* $P = \langle C, con\rangle$ *over a constraint system* $CS = \langle S, D, V\rangle$ *with* $C = \{c_1, \dots, c_n\}$, $c_i = \langle def_i, con_i\rangle$, $def_i : D^{|con_i|} \to S$, $V = \{x_1, \dots, x_k\}$ *and the corresponding constraint graph*

$$P = \lambda\langle c_1, \dots, c_n\rangle.\lambda\langle con\rangle.\nu\langle V - con\rangle. \prod_{i=1\dots n} c_i(con_i),$$

the function solution is defined as

$$Sol(P) = P(DEF_P) = \lambda\langle con\rangle.\nu\langle V - con\rangle. \prod_{i=1\dots n} def_i(con_i)$$

where DEF_P *represents the ordered list of all the definitions of the constraints in* P.

EXAMPLE 5.1.9 Consider again the fuzzy CSP represented in Figure 5.1. The constraint graph is represented by the function

$$\lambda\langle c_1, c_2, c_x, c_y, c_z\rangle.\lambda\langle x, y\rangle.\nu z.c_1(x, y) \times_S c_2(y, z) \times_S c_x(x) \times_S c_y(y) \times_S c_z(z).$$

The structure of this term gives an explicit way to represent an SCSP with 5 constraints and 3 variables, and shows how constraints c_1 and c_2 share one of the variables. Note that in this term, the variables x, y and z, and the constraints c_1,c_2,c_x,c_y,c_z are bound by λ-abstraction, so any term obtained from this one via α-renaming represents the same SCSP (i.e., α-renamed terms represent all the SCSPs with this shape). Note also that the $\times_S$ operator in the body of the function is not instantiated to the corresponding operator of the (fuzzy, in this case) semiring. This means that only the shape of the problem is represented and not its value. In other words, the semantical meaning of the $\times$ operator (that is the min) is not used.

To compute the solution of the problem, we first need to apply the definition of the constraints to the constraint graph, that is, to compute $P(DEF_P)$. In this way we obtain

$$\lambda\langle x, y\rangle.\nu z.def_1(x, y) \times_S def_2(y, z) \times_S def_x(x) \times_S def_y(y) \times_S def_z(z).$$

Now, if we want to compute the value of this SCSP we have to instantiate the variable and to compute the value of the solution function. As an example of solution, consider the pair $x = a$ and $y = b$. We obtain the function

$$\nu z.0.2 \times_S def_2(b, z) \times_S 0.9 \times_S 1 \times_S def_z(z) =$$

by instantiating the $\times_S$ operator to $\mathtt{min}$

$$\nu z.\mathtt{min}(0.2, \mathtt{def}_2(b, z), 0.9, 1, \mathtt{def}_z(z)).$$

By definition of ν we obtain

$$\begin{aligned}
&\mathtt{max}_{z\in\{a,b\}}\mathtt{min}(0.2, \mathtt{def}_2(b, z), 0.9, 1, \mathtt{def}_z(z)) = \\
&\mathtt{max}(\mathtt{min}(0.2, \mathtt{def}_2(b, a), 0.9, 1, \mathtt{def}_z(a)), \mathtt{min}(0.2, \mathtt{def}_2(b, b), 0.9, 1, \mathtt{def}_z(b))) = \\
&\mathtt{max}(\mathtt{min}(0.2, 0.1, 0.9, 1, 1), \mathtt{min}(0.2, 0.7, 0.9, 1, 0.8)) = \\
&\mathtt{max}(0.1, 0.2) = 0.2
\end{aligned}$$

□

The possibility to write the problem solution as a function, give as the possibility to compute partial solution for partially instantiated problem. This situation often arise in reality specially when the problem involve several data and the user knows only some of them. Another situation that can be represented by using partial instantiated SCSP is when the data of the problem are sent sequentially in several steps. In this case we can try to approximate the solution of the problem also before having all the problem data.

As an example let us consider same problem of Example 5.1.9 but by considering the case we do not know any information for constraint c_y.

EXAMPLE 5.1.10 Consider again the fuzzy CSP represented in Figure 5.1, but suppose we do not know the detail of constraint c_y.

The solution of the problem is represented by the function

$$\lambda\langle x, y\rangle.\nu z.\mathtt{def}_1(x, y) \times_S \mathtt{def}_2(y, z) \times_S \mathtt{def}_x(x) \times_S \mathtt{def}_y(y) \times_S \mathtt{def}_z(z).$$

Let suppose we want as before compute the value of the solution for the pair $x = a$ and $y = b$. We obtain the function

$$\nu z.0.2 \times_S \mathtt{def}_2(b, z) \times_S 0.9 \times_S \mathtt{def}_y(b) \times_S \mathtt{def}_z(z) =$$

by instantiating the $\times_S$ operator to $\mathtt{min}$

$$\nu z.\mathtt{min}(0.2, \mathtt{def}_2(b, z), 0.9, \mathtt{def}_y(b), \mathtt{def}_z(z)).$$

By definition of ν we obtain

$$\begin{aligned}
&\mathtt{max}_{z\in\{a,b\}}\mathtt{min}(0.2, \mathtt{def}_2(b, z), 0.9, \mathtt{def}_y(b), \mathtt{def}_z(z)) = \\
&\mathtt{max}(\mathtt{min}(0.2, \mathtt{def}_2(b, a), 0.9, \mathtt{def}_y(b), \mathtt{def}_z(a)), \mathtt{min}(0.2, \mathtt{def}_2(b, b), 0.9, \mathtt{def}_y(b), \mathtt{def}_z(b))) = \\
&\mathtt{max}(\mathtt{min}(0.2, 0.1, 0.9, \mathtt{def}_y(b), 1), \mathtt{min}(0.2, 0.7, 0.9, \mathtt{def}_y(b), 0.8))
\end{aligned}$$

At this point we can try to make some consideration:

- $\mathtt{Sol}(P)(a, b) = \mathtt{def}_y(b)$ if $\mathtt{def}_y(b) \leq_S 0.2$
- $\mathtt{Sol}(P)(a, b) = 0.2$ else

So, if we statistically know that the value of $\mathtt{def}_y(b)$ is always lower than 0.2 we can give a good approximation of the solution by declaring $\mathtt{sol}(P)(a, b) = 0.2$. □

5.2 A small language to program with soft constraints

We will now define a small language that will be useful in representing and solving SCSPs, using local consistency and dynamic programming algorithms. To define such algorithms, we need to explicitly define the *typed locations*, which we remind are just a set of variables, that will be used to identify a subproblem of a given SCSP.

DEFINITION 5.2.1 (TYPED LOCATION) *A* typed location l *is a set of variables. Given a constraint system* $\langle S, D, V\rangle$*, the set of locations* $\mathcal{L}$ *is defined as* $\mathcal{L} = \wp(V)$.

Locations represent variables of type *constraint*; that is, each location can contain a constraint (i.e. a function $D \rightarrow S$). Let us now define the *(Constraint) Store*.

DEFINITION 5.2.2 (CONSTRAINT STORE) *A constraint store is a function* σ *such that* $\sigma l : D^{|l|} \rightarrow S$*, where we have* $\sigma l\langle x_1, \dots, x_{|l|}\rangle = \mathbf{1}$ *for all except a finite number of locations.*

Given a constraint Problem $P = \langle C, con\rangle$, we can build the corresponding constraint store σ_C defined as follows:

$$\sigma_C l = \begin{cases} def_i \text{ if } c_i = \langle def_i, l\rangle \in C \\ \lambda\langle v_1, \dots, v_{|l|}\rangle.\mathbf{1} \text{ otherwise.} \end{cases}$$

In the following we will write σ instead of σ_C if the corresponding problem $P = \langle C, con\rangle$ is easy to be recognized.

Let us now define the syntactic structure of our language: the *expressions* Expr and the *commands* Com.

DEFINITION 5.2.3 (EXPRESSIONS) *The expressions we can build with the language are:*

$$\mathsf{Expr} ::= c \mid l \mid \mathsf{Sol}(l) \mid \mathsf{Sol}(l, L) \mid \phi(l, L)$$

The element c represents a constraint (that is, a basic piece of information); $\mathsf{Sol}(l)$ represents the solution of a constraint problem whose variables of interest are the variables in l. $\mathsf{Sol}(l, L)$ represents instead the solution of a subproblem (whose constraints are those connecting variables in $L \cup \{l\}$) and will be useful to define local consistency and dynamic programming algorithms. Finally, $\phi(l, L)$ is a generic monotone and extensive function that will be useful to represent generic soft local consistency algorithms.

Before giving a semantics to the expressions, let us introduce the syntactic category of commands.

DEFINITION 5.2.4 (COMMANDS) *The commands we can build with the language are defined by the following grammar:*

$$\mathsf{Com} ::= l := \mathsf{Expr} \mid \mathsf{Com}; \mathsf{Com}$$

Note that the assignment command “$\mathtt{l} := \mathtt{Expr}$” can be performed only if the type of the expression $\mathtt{Expr}$ is compatible with the location $\mathtt{l}$.

The assignment command is useful to change the value of a location present in the constraint store, and the ”;” operator is useful to perform several assignments sequentially.

To formally compute the *semantics* of the expressions and of the commands, we need to define two semantic functions: $\mathcal{E}$ for the expressions and $\mathcal{C}$ for the commands.

Function $\mathcal{E}$ evaluates an expression on a constraint store and gives as a result a constraint (that is, an element $c \in C$ that assigns to each tuple of elements in D a value in S).

DEFINITION 5.2.5 (FUNCTION $\mathcal{E}$) *The semantic function*

$$\mathcal{E} : \mathtt{Expr} \to \sigma \to C$$

is defined by structural recursion as follows:

$$\begin{aligned}
\mathcal{E}[\![\, \mathtt{Sol}(l) \,]\!]\sigma &= \lambda\langle l\rangle.\nu\langle V - l\rangle. \prod\nolimits_{S}{}_{l' \in \mathcal{L}} \sigma l'\langle l'\rangle \\
\mathcal{E}[\![\, \mathtt{Sol}(l, L) \,]\!]\sigma &= \lambda\langle l\rangle.\nu\langle L\rangle. \prod\nolimits_{S}{}_{l' \in (L \cup \{l\})} \sigma l'\langle l'\rangle \\
\mathcal{E}[\![\, c \,]\!]\sigma &= \hat{c} \\
\mathcal{E}[\![\, l \,]\!]\sigma &= \sigma l \\
\mathcal{E}[\![\, \phi(l, L) \,]\!]\sigma &= \hat{\phi}(l, L)
\end{aligned}$$

The element $\hat{c}$ and $\hat{\phi}(l, L)$ are basic elements representing respectively a constraint definition (that is, its $\mathtt{def}$ *function) and a generic monotone and extensive function $\hat{\phi}(l, L) : (\prod_{l' \in (L \cup \{l\})} D^{|l'|} \to S) \to (D^{|l|} \to S)$. The function $\mathtt{Sol}(l, L)$ is a specific $\hat{\phi}(l, L)$, and the function $\mathtt{sol}(l)$ is obtained by performing $\mathtt{Sol}(l, \mathcal{L})$ where $\mathcal{L}$ is the set of all the locations of the problem.*

Since the language we are going to define uses assignments, we need to define the meaning of store modification:

DEFINITION 5.2.6 (STORE MODIFICATIONS) *Given a constraint store σ, and a constraint $c : D^{|l|} \to S$, the store modification operator is defined as follows:*

$$\sigma[c/l]l' = \begin{cases} c \textit{ if } l' = l \\ \sigma l' \textit{ otherwise.} \end{cases}$$

Let us now evaluate the commands using the function $\mathcal{C}$.

DEFINITION 5.2.7 (THE FUNCTION $\mathcal{C}$) *The semantic function*

$$\mathcal{C} : \mathtt{Com} \rightarrow \sigma \rightarrow \sigma$$

is defined by structural recursion as follows:

$$\mathcal{C}[\![\, l := e \,]\!]\sigma = \sigma[\mathcal{E}[\![\, e \,]\!]\sigma/l]$$
$$\mathcal{C}[\![\, \mathtt{com}_1; \mathtt{com}_2 \,]\!]\sigma = \mathcal{C}[\![\, \mathtt{com}_2 \,]\!](\mathcal{C}[\![\, \mathtt{com}_1 \,]\!]\sigma)$$

The assignment command $l := e$ change the constraint defined over the variable in l by using the constraint e. This type of modification of the store usually happens when a preprocessing algorithm is applied to the problem. In fact, in this case, the value of the constraints is changed (usually is lowered) by also maintaining the equivalence with the original problem.
The sequentialisation operator is instead needed since we want to put together several assignment operation and build several preprocessing (or solving) techniques.

5.3 Solving SCSPs

In this section we will show how the dynamic programming and local consistency algorithms can be represented as specific expressions of the language defined in Section 5.2.

5.3.1 Dynamic programming techniques

Using dynamic programming techniques, the problem is not solved in one step, but it is partitioned in subproblems which form a *parsing tree* S, and then it is solved by solving such subproblems in a bottom-up visit of S.

A subproblem in a parsing tree S is represented by a rule $r_i = (l_i \leftarrow L_i)$ where l_i represents the location of interest of the problem, and L_i represents all the variables and constraints involved in the subproblem (see Section 3.1). The parsing tree $S = r_1; \ldots; r_n$, as defined in [BMR97b], represents a correct way to decompose the global problem, and also gives a computation order of the solutions of the subproblems. In our syntax the solution of a subproblem defined by a rule $l \leftarrow L$ is defined as the term $\mathtt{Sol}(l, L)$.

THEOREM 5.3.1 (DYNAMIC PROGRAMMING SOLVING) *Consider a problem $P = \langle C, \mathtt{con} \rangle$ over the constraint system $CS = \langle S, D, V \rangle$, the corresponding constraint store σ and one of its parsing trees $S = r_1; \ldots; r_n$, where $r_i = (l_i \leftarrow L_i)$, we have:*

$$\mathcal{E}[\![\, \mathtt{Sol}(\mathtt{con}) \,]\!]\sigma = \mathcal{E}[\![\, l_n \,]\!](\mathcal{C}[\![\, l_1 := \mathtt{Sol}(l_1, L_1); \ldots; l_n := \mathtt{Sol}(l_n, L_n) \,]\!]\sigma)$$

PROOF. The proof follows the lines of the Theorem 3.6.3, where the store σ takes the place of P and the semantics of subproblem solution $\mathcal{C}[\![\, l := \mathtt{sol}(l, L) \,]\!]\sigma$ takes the place of the rule application $[l \leftarrow L]P$. □

EXAMPLE 5.3.2 Let us review some steps of the dynamic programming technique applied to the running example of the chapter represented in Figure 5.1.

Suppose to have the strategy $S = r_1, r_2$ with $r_1 = \{y\} \leftarrow \{\langle y, z\rangle, z\}$ and $r_2 = \{\langle x, y\rangle\} \leftarrow \{x, y\}$. So we need to show that

$$\mathcal{E}[\![\, \mathtt{Sol}(\langle x, y\rangle) \,]\!]\sigma = \mathcal{E}[\![\, \mathtt{Sol}(\langle x, y\rangle, \{x, y\}) \,]\!](\mathcal{C}[\![\, \{y\} := \mathtt{Sol}(\{y\}, \{\langle y, z\rangle, z\}) \,]\!]\sigma).$$

By solving the subproblem represented by the first rule, we obtain the constraint problem represented in Figure 5.2. Let us call σ'_C the constraint store represented in the figure.

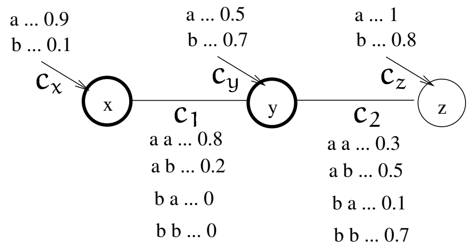

Figure 5.2: The fuzzy CSP after the application of rule r_1.

By definition of the ";" operator, we now need to solve the problem

$$\mathcal{E}[\![\, \mathtt{Sol}(\langle x, y\rangle, \{x, y\}) \,]\!]\sigma'_C.$$

It is now easy to check that the solution computed in this way is equivalent to the solution computed by giving semantics to $\mathcal{E}[\![\, \mathtt{Sol}(\langle x, y\rangle) \,]\!]\sigma$. □

5.3.2 Local consistency techniques

In this section we show how several soft local consistency algorithms can be captured and described in a uniform way in the framework described in this chapter. In section 3.5 we already described in a general way constraint propagation algorithms using monotone functions (extending in this way the techniques described in Section 3.1 with the results of Apt in [Apt99a, Apt99b]). Here we first represent classical soft local consistency rules as commands in our language, and then we show how the language is also able to capture more extended versions of the local propagation algorithms as depicted in Section 3.4.

For each local consistency rule that we have in our strategy, we define a suitable assignment that modifies a piece of the constraint store related to the consistency rule itself. We can claim that the semantics of $\mathtt{Sol}(con)$ is the same if we compute it in the constraint store σ or in a constraint store $\mathcal{C}[\![\, \mathtt{Sol}(l, L) \,]\!]\sigma$ for any l, L, and σ.

THEOREM 5.3.3 (EQUIVALENCE FOR $\mathtt{Sol}(l, L)$) *Consider a problem* $P = \langle C, con\rangle$ *and a local consistency rule* $[l \leftarrow L]$. *Then we have that*

$$\mathcal{E}[\![\, \mathtt{Sol}(con) \,]\!]\sigma = \mathcal{E}[\![\, \mathtt{Sol}(con) \,]\!](\mathcal{C}[\![\, l := \mathtt{Sol}(l, L) \,]\!]\sigma).$$

PROOF. Easily follows from Theorem 3.1.9. □

5.3.3 Extending local propagation rules

In the classical CSP environment, a local propagation function is defined (following the approach of [Apt99a, Apt99b]) as a monotone and extensive function that transforms the original problem into a new (possibly equivalent) one. This means that not only a solution over a set of locations $L \cup \{l\}$ and a projection operation over the variables in l are performed but *any* monotone and extensive transformation can be applied to the location l using the locations $L \cup \{l\}$.

For this purpose we have introduced in the language the function $\phi(l, L)$ with the meaning to compute any modification over l by using the constraints over $l \cup L$.

More generally, a local propagation function can be seen as a function that transforms a problem into a new one (possibly by modifying only a subset of it).

DEFINITION 5.3.4 (LOCAL PROPAGATION) *Given a problem $P = \langle C, \mathtt{con} \rangle$ and the corresponding store σ, a local propagation function $\mathtt{lc}$ is a monotone and extensive (w.r.t. the semiring order $\leq_S$) function that transforms (some of) its constraints obtaining a (possibly equivalent) problem:*

$$\mathtt{lc} : \sigma \rightarrow \sigma$$

The definition of $\mathtt{lc}$ is so general that it could match with any of the previously defined local consistency algorithms. We want to characterize it in a more precise way, by using a set of *general local propagation rules* by extending the rules of Section 3.1.

DEFINITION 5.3.5 (GENERAL LOCAL PROPAGATION RULE) *A general local propagation rule is defined by a subproblem and a monotone and extensive (w.r.t. the semiring order $\leq_S$) function ϕ. To indicate the application of a function ϕ to a subproblem defined by a rule $r = l \leftarrow L$, we will write $\phi(l, L)$.*

In general since the function $\phi(l, L)$ is monotone and extensive we have:

THEOREM 5.3.6 *Consider a problem $P = \langle C, \mathtt{con} \rangle$ and a general local consistency rule $\phi(l, L)$. Then, for any domain tuple l' we have that*

$$\mathcal{C}[\![\, l := \phi(l, L) \,]\!]\sigma l'\langle l' \rangle \leq_S \sigma l'\langle l' \rangle.$$

PROOF. Easily follows by the hypothesis of monotonicity of ϕ. □

5.4 Constraint Problem as Semiring

In this section we illustrate a possible approach to extend the notion of semiring to constraint problems. This means that we want the semiring to capture non only the levels of preference of each instantiated constraint but also the constraint itself.

To do this we define a possible semiring representing Constraint Problems. We need first to define the operator that represents the combination of two constraints. Note that since the combination of constraints has an operational behaviour that depends on the variables that are shared between the two constraints, we do not have a unique combination operator, but a collection of them.

Moreover we need also a notion of *type* for each constraint. In the following we will write $c : (D_0 \times \ldots \times D_n) \to S$ if the constraint c involve variable $x_0, \ldots, x_n$.

DEFINITION 5.4.1 (COMBINATION: $\otimes_D$)

$$\frac{c_1 : (D_0 \times D_1) \to S, c_2 : (D_0 \times D_2) \to S}{c_1 \otimes_{D_0} c_2 : (D_0 \times D_1 \times D_2) \to S = \lambda\langle x_0, x_1, x_2\rangle.c_1(x_0, x_1) \times_S c_2(x_0, x_2)}$$

The $\oplus$ operator can be instead defined as follows:

DEFINITION 5.4.2 (DISJUNCTION: $\oplus_D$)

$$\frac{c_1 : (D_0 \times D_1) \to S, c_2 : (D_0 \times D_2) \to S}{c_1 \oplus_{D_0} c_2 : D_0 \to S = \lambda x_0.\nu\langle x_1, x_2\rangle.c_1(x_0, x_1) +_S c_2(x_0, x_2)}$$

We can now claim that the set of constraints C with the combination and disjunction operations is an *enhanced*[1] semiring.

To define the unit element of the two operation we need the definition of constant function over a given set of variables:

DEFINITION 5.4.3 (CONSTANT FUNCTION: I_a) *Consider a set of variables* I. *We define the function* $I_a : D_I \to A$ *as the function that assigns to each* I*-tuple* t *the semiring value* a.

THEOREM 5.4.4 ($\langle C, \oplus, \otimes, V_0, \emptyset_1\rangle$ IS AN ENHANCED SEMIRING) $\langle C, \oplus, \otimes, V_0, \emptyset_1\rangle$ *is an enhanced semiring.*

PROOF SKETCH. It is enough to check all the properties by using the fact that the same properties holds for the original semiring S. □

5.5 Conclusions

In this chapter we have described how the *semiring* framework can be used to both embed the constraint structure and topology in a suitable semiring of functions. We have also described a small language to "program" with soft constraints. Starting from basic objects (domains and semirings), we have built functions able to represent single steps of the local consistency and dynamic programming algorithms. By using this language we could give a procedural view of the solving techniques already described in the previous chapters.

This chapter concludes our study of the soft constraint solving framework; in the next one we will study how to program with soft constraints using a declarative language.

[1] We use the adjective enhanced because the $\otimes$ and $\oplus$ operators are *overloaded*. In fact, each of them represents a class of operators $\otimes_D$ and $\oplus_D$ (by varying the domain D).

Chapter 6

Soft CLP

Abstract

The framework presented in the previous chapters shows how the soft constraint idea can give us an easy way to model non-crisp problems. To program applications we need, nevertheless, a language where we can use soft constraints. This is why we extend the Constraint Logic Programming (CLP) formalism in order to handle semiring-based constraints.

This allows us to perform in the same language both constraint solving and optimization. In fact, constraints based on semirings are able to model both classical constraint solving and more sophisticated features like uncertainty, probability, fuzziness, and optimization. We then provide this class of languages with three equivalent semantics: model-theoretic, fix-point, and proof-theoretic, in the style of classical CLP programs.

Constraint logic programming (CLP) [JL87] languages extend logic programming (LP) by replacing term equalities with constraints and unification with constraint solving. Programming in CLP means choosing a constraint system for a specific class of constraints (for example, linear arithmetic constraints, or finite domain constraints) and embedding it into a logic programming engine. This approach is very flexible since one can choose among many constraint systems without changing the overall programming language, and has shown to be very successful in specifying and solving complex problems in terms of constraints of various kind [Wal96]. It can, however, only handle classical constraint solving. Thus, it is natural to try to extend the CLP formalism in order to also be able to handle soft constraints. In fact, this new programming paradigm, which we will call SCLP (for Semiring-based CLP, or also Soft CLP), has the advantage of treating in a uniform way, and with the same underlying machinery, all constraints that can be seen as instances of the semiring-based approach: from optimization to satisfaction problems, from fuzzy to probabilistic, prioritised, or uncertain constraints, and also multi-criteria problems, without losing the ability to treat and solve classical hard constraints. This leads to a high-level declarative programming formalism where real-life problems involving constraints of all these kinds can be easily modeled and solved.

In passing from CLP to SCLP languages, we will replace classical constraints with the more general SCSP constraints. By doing this, from a technical point of view, we have to modify the notions of interpretation, model, model intersection, and others, since we have to take into account the semiring operations and not the usual CLP operations. For example, while CLP interpretations associate a truth value (either *true* or *false*) to each ground atom, in this case ground atoms must be given one of the elements of the semiring. Furthermore, whereas in CLP the value associated with an existentially quantified atom is the *logical or* among the truth values associated to each of its instantiations, here we have to replace the *or* with another operation that refers to one of the semiring operations.

After describing the syntax of SCLP programs, we will define three equivalent semantics for such languages: model-theoretic, fix-point, and operational. These semantics are conservative extensions of the corresponding ones for LP, since by choosing a particular semiring (the one with just two elements, *true* and *false*, and the logical *and* and *or* as the two semiring operations) we get exactly the LP semantics. The extension is in some cases predictable but it possesses some crucial new features. For example, the presence of a partial order among the semiring elements (and not a *total* order like it is in the LP/CLP case, where we just have two comparable elements) brings some conceptual complexity in some aspects of the semantics. In fact, in the operational semantics, there could be two refutations for a goal leading to different semiring elements that are not comparable in the partial order. In this case, these elements have to be combined in order to get the solution corresponding to the given goal, and their combination could not be reachable by any derivation path in the search tree. This means that any constructive way to get such a solution by visiting the search tree would have to follow all the incomparable paths before being able to find the correct answer. In practice, however, classical branch and bound techniques can be adapted to this framework to cut some useless branches.

We also show the equivalence of the three semantics. In particular, we prove that, given the set of all refutations starting from a given goal, it is possible to derive the declarative meaning of both the existential closure of the goal and its universal closure.

Additionally, we investigate the decidability of the semantics of SCLP programs, obtaining an interesting semi-decidability result: if a goal has a semiring value greater than, or greater than or equal to, a certain value in the semiring, then we can discover this in finite time. Moreover, for SCLP programs without functions, the problem is completely decidable: the semantics of a goal can be computed in finite and bounded time. In fact, in this case we can consider only a finite number of finite and bounded-length refutations (see Section 6.6): infinite refutations do not bring more information due to the properties of the semiring operations. Notice that the absence of functions is obviously a restriction, however, not all sources of infiniteness are taken away, since nothing is said about the semiring, which could still be infinite.

The chapter is organized as follows. Section 6.1 defines the syntax of SCLP programs. Afterwards, sections 6.2, 6.3, and 6.4 provide SCLP programs with a model-theoretic, a fix-point, and an operational semantics, respectively. Then, Section 6.5 presents a semi-decidability result for SCLP programs, and Section 6.6 adds some more decidability results that hold for programs without functions. Section 6.7, instead, presents an abstract model for the operational semantics of the SCLP language without functions, using the *Gurevich's Abstract State Machines* (ASMs). Finally, Section 6.8 concludes the chapter by discussing the relationship with related work. The SCLP framework described in this chapter appeared in [BMR97a, BMR01].

The chapter is based on the results introduced in [BMR97a, BR97b] and extended in [BMR01].

6.1 Syntax of SCLP Programs

For readers familiar with Constraint Logic Programming (CLP) programs, we can say that SCLP(S) programs (also written SCLP when the semiring is obvious or not important) are just CLP programs [JL87] where constraints are defined over a certain c-semiring $S = \langle A, +, \times, \mathbf{0}, \mathbf{1} \rangle$.

As usual, a program is a set of *clauses*. Each clause is composed by a *head* and a *body*. The head is just an *atom* and the body is either a collection of atoms, or a value of the semiring, or a special symbol ($\square$) to denote that it is empty. Clauses where the body is empty or it is just a semiring element are called *facts* and define predicates which represent constraints. When the body is empty, we interpret it as having the best semiring element (that is, $\mathbf{1}$).

Atoms are n-ary predicate symbols followed by a tuple of n *terms*. Each term is either a constant or a variable or an n-ary function symbol followed by n terms. *Ground* terms are terms without variables. Finally, a *goal* is a collection of atoms.

The BNF for this syntax follows in table 6.1.

P ::	$CL \mid CL, P$
CL ::	$H :- B$
H ::	AT where AT is the category of atoms
LAT ::	$\square \mid LAT'$
LAT′ ::	$AT \mid AT, LAT'$
B ::	$LAT \mid a$ where $a \in A$
G ::	$:- LAT$

Table 6.1: BNF for the SCLP syntax.

As an example, consider the following SCLP(S) program, represented in Table 6.2 where the semiring is $S = \langle [0, 1], max, min, 0, 1 \rangle$. We recall that this is the fuzzy semiring, where tuples of values are given values between 0 and 1, and constraints are combined via the min operator and compared via the max operator. Note that the ordering $\leq_S$ in this semiring coincides with the $\leq$ ordering over the reals in $[0, 1]$.

Our example is a generalization of the usual n-queens problem, which can be found for example in [Hen89]. The classical formulation requires that n queens are placed on a $n \times n$ chessboard in such a way that they do not attack each other. In our formulation, we allow also attacking queens but we give a higher preference to solutions where queens attacking each other are farther apart. Thus, if there are solutions where no queens attack each other, these solutions will remain the best ones. But there are also other solutions, and among these additional solutions the best ones are those where the queens attacking each other are as far apart as possible.

We assume here that the reader is familiar at least with the logic programming [Llo87] concepts and notation.

```
myqueens([],N).
myqueens([X|Y],N) :-
     noattack(X,Y,N), myqueens(Y,N).
noattack(X,Xs,N) :-
     noattack(X,Xs,N,1).
noattack(X,[],N,Nb).
noattack(X,[Y|Ys],N,Nb) :-
     row(X,Y,N,Nb), diag1(X,Y,N,Nb), diag2(X,Y,N,Nb),
     noattack(X,Ys,N,Nb+1).
row(X,X,N,Nb) :- Nb/N.
row(X,Y,N,Nb) :- different(X,Y).
diag1(X,X+Nb,N,Nb) :- Nb/N.
diag1(X,Y,N,Nb) :- different(Y,X+Nb).
diag2(X,X-Nb,N,Nb) :- Nb/N.
diag2(X,Y,N,Nb) :- different(Y,X-Nb).
```

Table 6.2: Soft n-queens problem.

In this program, the n queens are represented by a list of n variables, say `[X1, ..., Xn]`, where variable `Xi` represents the queen in column i and its value represents the row index where this queen is located. This is one of the usual formulations of the n-queens problem, and it is worth noticing that by adopting this formulation we assume that different queens are in different columns.

The first predicate, `myqueens/2`, traverses the whole list of n variables and sets the constraints between any pair of queen, by using the predicate `noattack/4`, which sets the constraints between any queen and all the queens in subsequent columns. The actual constraints are set by predicates `row/4`, `diag1/4`, and `diag2/4`. In the classical (hard) formulation, these predicates basically say that different queens cannot stay in the same row, nor in the same diagonal. In our case, these constraints are made soft by also accepting situations where different queens are on the same row or diagonal. For example, for the row constraint, we give the semiring value Nb/N to two queens in the same row, where Nb is the distance between the two queens (that is, the number of columns between them) and N is the total number of columns. This means that the farther apart the two queens are, the higher this value will be. The same reasoning holds also for the two diagonals, except that the presence of both queens X and Y on the same diagonal corresponds to having $Y = X + Nb$ or $Y = X - Nb$.

Notice that each solution of this generalized n-queens problem has a semiring value which is obtained by minimizing the semiring values of all its constraints. This comes from the choice of the fuzzy semiring, where the multiplicative operation is the min. Therefore, if a solution contains three pairs of attacking queens, each of such pairs will have a semiring value given by one of the clauses defining predicates `row/4`, `diag1/4`, or `diag2/4` (proportional to the distance between the two queens), and then the value of this solution will be the minimum among such three values. Different solutions are then ordered using the other semiring operation, which in this case is the max. Note that this same program can also be used with a different semiring, obtaining a different way of computing a solution and a different ordering. For example, we could have chosen the semiring $\{\mathcal{R} \cup +\infty, min, +, +\infty, 0\}$, where the value of each solution would have been obtained by summing the values of each attacking pair, and solutions would have been compared using the min operator.

To use this program over a specific value for N, we need to add some clauses to set the domain for the constraint variables (from 1 to N) and to define predicate `different/2`. For example, for $N = 5$, we have to write the clauses represented in Table 6.3:

```
fivequeens([X1,X2,X3,X4,X5]) :-
     domain5(X1), domain5(X2), domain5(X3),
     domain5(X4), domain5(X5),
     myqueens([X1,X2,X3,X4,X5],5).

domain5(a) :- 1.
```
(for all $a \in \{1, \dots, 5\}$)

```
different(a,b) :- 1.
```
(for all $a, b \in \{1, \dots, 5\}$ such that $a \neq b$)

Table 6.3: Clauses for the 5-queens problem.

We now anticipate the behavior of this SCLP program, which obeys the semantic development of the future sections.

- Given the goal `:- fivequeens(L).`, the program will instantiate L to be a list of 5 values, indicating the row positions of the 5 queens. Since there are solutions where queens do not attack each other, the solution returned by the program will be one of these, with semiring value 1.

- Given the goal `:- fivequeens([1,3,X3,X4,X5]).`, the program will return a position for each queen such that the first queen is in row 1 and the second one in row 3. Given these additional constraints (i.e., the positions of the first two queens), it is possible that all solutions will have some attacking pairs of queens. Then, among these possible solutions, the program will return one with the highest level of preference, which means that the attacking queens are as far apart as possible.

- Assume to delete one or more of the facts defining predicate `domain5`. Then, if we give the goal `:- fivequeens(L).`, it means that we have five queens (thus five columns) but a smaller number of domain elements (thus rows). Even in this case, it is possible that all solutions will have some attacking pairs of queens. Thus the program will return one where the attacking queens are as far apart as possible.

While this example of an SCLP program uses semiring values to be able to find the best quasi-solutions in an otherwise over-constrained problem, other examples exploit the power of the chosen semiring to express features that are intrinsic to the considered problem. Consider the problem that can arise when a client in a restaurant wants to select items from the menu in a way that his/her preferences over the combinations of drinks and dishes are satisfied in the best way. In this example, each combination (for example, beer and pizza, or white wine and fish) is associated with a level of preference, according to the taste of the client. Then, solving the problem means finding the menu with the highest level of preference. As in the previous example, here the semiring that is used is the fuzzy one, but its role is definitely different. This menu problem is one of those that have been used to show the expressive power of the clp(fd,S) system [GC98], a very general implementation of SCLP programming[1].

Notice that, by just changing the semiring, but maintaining the same program structure of the menu example, we could model rather different situations. Consider for example the situation in which a conference organizer has to decide the menu for the conference dinner, trying to satisfy the participants as much as possible. Of course the organizer cannot ask for their preferences, but can reason with probabilities. Therefore, he/she can associate, with each combination of drink and dish, the estimated probability that it will please the conference attendees. Then, solving the problem means finding the menu which has the highest probability to please the participants. To model this situation, it is enough to keep the same program above (modulo the new semiring values), but choose the semiring $\langle [0, 1], \mathit{max}, \times, 0, 1 \rangle$, which can represent probabilities combined via $\times$ (assuming their independence) and compared via the max operator.

6.2 Model-theoretic Semantics

In this section we will generalize the usual development of the model-theoretic semantics in logic programming [Llo87] to be able to deal correctly with semiring values. The

[1] We will give a short description of the `SCLP(FD,S)` framework in Section 6.8.1.

main generalization will involve the assignment of semiring values to atoms and formulas, instead of truth values. As usual, we will just consider Herbrand interpretations in the following, which, however, we will call just interpretations for sake of conciseness.

DEFINITION 6.2.1 (PRE-INTERPRETATION) *A pre-interpretation maps each ground term in a program into a chosen domain. More precisely, it consists of a domain* D *plus a mapping from each constant to an element of* D *and, for each n-ary function, a mapping from* D^n *to* D.

Notice that the domain D of the pre-interpretations contains the domain of the constraints.

DEFINITION 6.2.2 (INTERPRETATION) *An interpretation* I *is a pre-interpretation plus a function, which takes a predicate and an instantiation of its arguments (that is, a ground atom), and returns an element of the semiring:*

$$I : \bigcup_n (P_n \to (D^n \to A)),$$

where P_n *is the set of n-ary predicates and* A *is the carrier set of the semiring.*

This notion of interpretation can now be extended and used to associate elements of the semiring also to formulas that are more complex than ground atoms. In the following, this extension of an interpretation I will be called an *interpretation* and denoted by I as well, since the extension is uniquely determined. When we will want to consider the restriction of an interpretation I to ground atoms we will sometimes write $GA(I)$.

- The value associated with a formula of the form $F = \exists x.F'(x)$ is computed by considering the lub of the values associated with all ground formulas $F'(x/d)$, where d is any domain element. That is, $I(F) = lub\{I(F'(d))$, for all $d \in D\}$. Formulas of this kind occur in SCLP languages since variables appearing in the body of a clause but not in its head are considered to be existentially quantified. For example, in the special case of logic programming the clause `p(a) :- q(X,a)` is just a shorthand for the formula $p(a) \leftarrow \exists x.q(x, a)$.

- The value associated with a formula of the form $F = \forall x.F'(x)$ is computed by considering the greatest lower bound (glb) of the values associated with all the ground formulas $F'(x/d)$, where d is any domain element. That is, $I(F) = glb\{I(F'(d))$,for all $d \in D\}$. Formulas of this kind occur when a variable appears in the head of a clause. In fact, for example, in logic programming, a clause like `p(X) :- q(X,a)` is a shorthand for the formula $\forall x.(p(x) \leftarrow q(x, a))$.

- The value associated with a conjunction of atomic formulas of the form (A, B) is the semiring product of the values associated to A and B: $I(A, B) = I(A) \times I(B)$. Such formulas appear in the body of the clauses, when the body contains more than one atom.

- For any semiring element a, $I(a) = a$. Such elements appear in the body of the facts.

Note that the meaning associated to formulas by function I coincides with the usual logic programming interpretation [Llo87] when considering constraints over the semiring $S_{CSP} = \langle\{\mathtt{true}, \mathtt{false}\}, \vee, \wedge, \mathtt{false}, \mathtt{true}\rangle$. In fact, in this case the ordering $\leq_S$ is defined by $\mathtt{false} \leq_S \mathtt{true}$, the lub operation of the lattice $\langle\{\mathtt{true}, \mathtt{false}\}, \leq_S\rangle$ is $\vee$, and the glb is $\wedge$. Thus, for example, $I(\exists x.A(x)) = lub\{I(A(d))$, for all $d \in D\} = \bigvee\{I(A(d))$, for all $d \in D\}$. Thus, it is enough that one of the $A(d)$ is assigned the value *true* that the value associated to the whole formula $\exists x.A(x)$ is *true*. Note also that in this special instance the lub and glb of the lattice coincide with the two semiring operations, but this is not true in general for the multiplicative operation.

DEFINITION 6.2.3 (CLAUSE SATISFACTION) *Given a clause of the form* $H : -B$ *and an interpretation* I*, we say that the clause is satisfied in* I *if and only if, for any ground instantiation of* H*, say* $H\theta$*, we have that* $I(H\theta) \geq_S I(\exists B\theta)$.

Note that the existential quantification over the body $B\theta$ is needed, since there may be variables in B that do not appear in H. Thus $B\theta$ could be not ground. This definition of clause satisfiability is consistent with the usual treatment of clauses in logic programming, where a clause is considered to be satisfied if the body logically implies the head, and by noting that logical implication in the semiring S_{CSP} coincides with the ordering $\leq_{S_{CSP}}$.

Our running example As our running example in this chapter, we will use the program represented in Table 6.4, which is not as expressive as the program in the previous section but is simple enough to be analyzed in detail. This is an SCLP(S) program over the semiring $S = \langle N \cup \{+\infty\}, min, +, +\infty, 0\rangle$, where N is the set of non-negative integers. This semiring allows to model constraint optimization problems where each tuple of values is assigned an integer, to be interpreted as its cost, constraints are combined by summing their costs, and are compared by using the min operator. Note that the ordering $\leq_S$ in this semiring coincides with the $\geq$ ordering over integers.

```
s(X)   :- p(X,Y).
p(a,b) :- q(a).
p(a,c) :- r(a).
q(a)   :- t(a).
t(a)   :- 2.
r(a)   :- 3.
```

Table 6.4: Our running example.

In this program, the constraints are represented by predicates `t` and `r`. The intuitive meaning of a semiring value like 3 associated to the atom $r(a)$ is that $r(a)$ costs 3 units.

Thus the set $N \cup \{+\infty\}$ contains all possible costs, and the choice of the two operations min and $+$ implies that we intend to minimize the sum of the costs. This gives us the possibility to select the atom instantiation that gives the minimal cost overall. $\square$

As an example of clause satisfiability, consider the following four clauses:

- the clause `p(a) :- q(b)` is satisfied in I if $I(p(a)) \geq_S I(q(b))$;
- the clause `p(X) :- q(X,a)` is satisfied if $\forall x.(I(p(x)) \geq_S I(q(x, a)))$;
- the clause `p(a) :- q(X,a)` is satisfied if $I(p(a)) \geq_S I(\exists x.q(x, a))$;
- the clause `p(X) :- q(X,Y)` is satisfied if $\forall x.(I(p(x)) \geq_S I(\exists y.q(x, y)))$.

As in logic programming, an interpretation I is a *model* for a program P if all clauses of P are satisfied in I. Given a program and all its models, one would like to identify a unique single model as the representative one. In logic programming this is done by considering the minimal model [Llo87], which is obtained by intersecting all the models of the program. This works because models in logic programming are assimilable to sets of ground atoms, those with associated value *true*. Here we follow the same approach, but we have to generalize the notion of intersection of two models, written as "$\circ$", as their greatest lower bound in the lattice $\langle A, \leq_S \rangle$.

DEFINITION 6.2.4 (MODEL INTERSECTION) *Consider an SCLP program over the c-semiring $\langle A, +, \times, 0, 1 \rangle$, the corresponding ordering $\leq_S$ and the lattice $\langle A, \leq_S \rangle$. For every ground atomic formula F and a family of models $\{M_i\}_{i \in I}$, we define $\circ_{i \in I} M_i(F) = glb_{i \in I}\{M_i(F)\}$, where glb is the greatest lower bound over the lattice $\langle A, \leq_S \rangle$.*

THEOREM 6.2.5 (MODEL INTERSECTION) *Consider a family of models $\{M_i\}_{i \in I}$ for a CLP(S,D) program P. Then $\circ_{i \in I} M_i$ is a model for P as well.*

PROOF. Since M_i is a model for P for all $i \in I$, it must be that, for every clause $H :- B$, and for all θ such that $H\theta$ is ground, $M_i(\exists B\theta) \leq_S M_i(H\theta)$. Consider now the model $M = \circ_{i \in I} M_i$. We need to prove that, for all $H :- B$, and for all θ such that $H\theta$ is ground, also $M(\exists B\theta) \leq_S M(H\theta)$ holds.

Without loss of generality, assume that $B = A_1, A_2$. Thus, for all $i \in I$, $M_i(\exists B\theta) = lub\{M_i(B\theta\theta')$, for all θ' such that $B\theta\theta'$ is ground$\}$ = $lub\{ M_i(A_1\theta\theta') \times M_i(A_2\theta\theta')$, for all θ' such that $B\theta\theta'$ is ground$\}$. Moreover, $M(\exists B\theta) = lub\{ M(B\theta\theta')$, for all θ' such that $B\theta\theta'$ is ground$\}$ = $lub\{M(A_1\theta\theta') \times M(A_2\theta\theta')$, for all θ' such that $B\theta\theta'$ is ground$\}$= $lub\{glb_{i \in I}\{ M_i(A_1\theta\theta')\} \times glb_{i \in I}\{ M_i(A_2\theta\theta')\}$ for all $\theta'\}$. Also, $M(H\theta) = glb_{i \in I}\{M_i(H\theta)\}$.

Consider any model M_j with $j \in I$. Since $glb_{i \in I}\{M_i(A_1\theta\theta')\} \leq_S M_j(A_1\theta\theta')$ and $glb_{i \in I}\{ M_i(A_2\theta\theta')\} \leq_S M_j(A_2\theta\theta')$ by definition of glb, and recalling that $\times$ is monotone, we have that $M(\exists B\theta) \leq_S M_j(\exists B\theta)$. By transitivity of $\leq_S$, we thus get $M(\exists B\theta) \leq_S M_j(H\theta)$. Since $M(H\theta)$ is the glb of all $M_i(H\theta)$ for $i \in I$, and since the glb of a set of elements is the greatest among the elements which are smaller than all of them, we have that $M(\exists B\theta) \leq_S M(H\theta)$. $\square$

It is easy to see that the operation of model intersection is associative, idempotent, and commutative.

DEFINITION 6.2.6 (MINIMAL MODEL) *Given a program* P *and the set of all its models, its* minimal model *is obtained by intersecting all models:* $M_P = \circ(\{M \mid M$ *is a model for* $P\})$. *The* model-theoretic semantics *of a program* P *is its minimal model,* M_P.

Consider our running example program P. The minimal model M_P for such a program must assign an integer to each formula, and when restricted to ground atoms it is the following function: $M_P(t(a)) = 2$, $M_P(q(a)) = 2$, $M_P(r(a)) = 3$, $M_P(p(a,c)) = 3$, $M_P(p(a,b)) = 2$, $M_P(s(a)) = min(2,3) = 2$. For each atom different from the ones considered above, M_P returns $+\infty$. To explain why function M_P returns these values, we give some examples:

- Any model must assign to $t(a)$ a semiring value smaller (that is, better) than 2, because of the clause `t(a) :- 2`. Since M_P is the minimal model, it must assign to $t(a)$ the glb (that is, the max) of all such values, that is, 2.

- The value assigned to $p(a,c)$ must be smaller than that of $r(a)$, which in turn must be smaller than 3. Being in the minimal model, we have that the value of $r(a)$ is exactly 3, and the value of $p(a,c)$ is again 3. The same reasoning holds also for $p(a,b)$, whose value is 2. Instead, $p(a,v)$, for any $v \neq b, c$, gets the value $+\infty$, because a model can give any value to $p(a,v)$, since there is no clause about it, and thus the minimal model gives to it the glb (that is, the max) of all values, that is, the worst element of the semiring.

- For $s(a)$, we know that every model must assign to it a value smaller than the value assigned to $\exists y.p(a,y)$. Now, for any model M, $M(\exists y.p(a,y))$ is the lub (that is, the min) of all the values assigned by M to $p(a,v)$ for any v in the domain. We know that $p(a,b)$ has value 2, $p(a,c)$ has value 3, and any $p(a,v)$, with $v \neq b, c$, has value $+\infty$. Therefore the lub of all such values is 2. Thus any model must assign to $s(a)$ a value smaller than 2, and the minimal model M_P must give it value 2.

For the same program, it is also useful to notice which semiring value is assigned to formulas like $\forall y.p(a,y)$ by the function M_P. In fact, this is one of the kinds of formulas we will consider when studying the relationship between the operational and the model-theoretic semantics, in Section 6.4. By definition, to get $M_P(\forall y.p(a,y))$ we must compute the glb of all the semiring values assigned by M_P to the ground atoms of the form $p(a,v)$ where v is any element of the domain D. We know, by the paragraph above, that $M_P(p(a,c)) = 3$, $M_P(p(a,b)) = 2$, and $M_P(p(a,v)) = +\infty$ if v is different from both b and c. Thus the glb (that is, the max) of all these values is $+\infty$. Therefore $M_P(\forall y.p(a,y)) = +\infty$.

6.3 Fix-point Semantics

In order to describe the fix-point semantics, we need to define the operator T_P which extends the one used in logic programming [Llo87]. We will do that by following the same approach as in the previous section. The resulting operator maps interpretations into interpretations, that is, $T_P : IS_P \rightarrow IS_P$, where IS_P is the set of all interpretations for P.

DEFINITION 6.3.1 (T_P OPERATOR) *Given an interpretation* I *and a ground atom* A, *assume that program* P *contains* k *clauses defining the predicate in* A. *Clause* i *is of the form* $A :- B^i_1, \ldots, B^i_{n_i}$. *Then*

$$T_P(I)(A) = \sum_{i=1}^{k}(\prod_{j=1}^{n_i} I(B^i_j)).$$

This function coincides with the usual immediate consequence operator of logic programming (see [Llo87]) when considering the semiring S_{CSP}.

Consider now an ordering $\preceq$ among interpretations which respects the semiring ordering.

DEFINITION 6.3.2 (PARTIAL ORDER OF INTERPRETATIONS) *Given a program* P *and the set of all its interpretations* IS_P, *we define the structure* $\langle IS_P, \preceq\rangle$, *where for any* $I_1, I_2 \in IS_P$, $I_1 \preceq I_2$ *if* $I_1(A) \leq_S I_2(A)$ *for any ground atom* A.

It is easy to see that $\langle IS_P, \preceq\rangle$ is a complete partial order, whose greatest lower bound coincides with the glb operation in the lattice A (suitably extended to interpretations). It is also possible to prove that function T_P is monotone and continuous over the complete partial order $\langle IS_P, \preceq\rangle$.

By using these properties, classical results on partial orders [Tar55] allow us to conclude the following:

- T_P has a least fix-point, $lfp(T_P)$, which coincides with $glb(\{I \mid T_P(I) \preceq I\})$;
- the least fix-point of T_P can be obtained by computing $T_P \uparrow \omega$. This means starting the application of T_P from the bottom of the partial order of interpretations, called I_0, and then repeatedly applying T_P until a fix-point.

Consider again our running example program. We recall that in this specific case the semiring is $S = \langle N \cup \{+\infty\}, min, +, +\infty, 0\rangle$ and $D = \{a, b, c\}$. Thus function T_P is:

$$T_P(I)(A) = min\{\sum_{j=1}^{n_1} I(B^1_j), \ldots, \sum_{j=1}^{n_k} I(B^k_j)\}.$$

In this semiring the bottom interpretation I_0 is the interpretation that maps each semiring element into itself and each ground atom into the bottom of the lattice associated to the semiring, that is, $+\infty$. Note that we slightly abused the notation since interpretations are functions whose domain contains only ground atoms (see Section 6.2), while here we also included semiring elements. This simplifies the definition of I_0; however, it is possible to obtain the same result with a more complex definition of I_0 which satisfies the definition of interpretation. Given I_0, we obtain I_1 by applying function T_P above. For example, $I_1(r(a)) = +3$. Instead, $I_1(p(a,c)) = +\infty$, and $I_2(p(a,c)) = I_1(r(a)) = +3$. The Table 6.5 gives the value associated by the interpretation I_i to each ground atom. Some of the atoms[2] are not listed because each interpretation I_i gives them value $+\infty$. All interpretation I_i with $i > 4$ coincide with I_4, thus I_4 is the fix-point of T_P.

	I_1	I_2	I_3	I_4
t(a)	2	2	2	2
r(a)	3	3	3	3
q(a)	$+\infty$	2	2	2
p(a,c)	$+\infty$	3	3	3
p(a,b)	$+\infty$	$+\infty$	2	2
s(a)	$+\infty$	$+\infty$	3	2
s(b)	$+\infty$	$+\infty$	$+\infty$	$+\infty$
s(c)	$+\infty$	$+\infty$	$+\infty$	$+\infty$

Table 6.5: The T_P operator applied at the running example.

The most interesting case is the computation of the value associated to $s(a)$. In fact, $I_3(s(a)) = min\{I_2(p(a,a)), I_2(p(a,b)), I_2(p(a,c))\} = min\{+\infty, +\infty, 3\} = 3$. Instead, $I_4(s(a)) = min\{I_3(p(a,a)), I_3(p(a,b)), I_3(p(a,c))\} = min\{+\infty, 2, 3\} = 2$. Note that the clause `s(X) :- p(X,Y)` is considered equivalent to all its instantiations. In particular, when $x = a$, we have the three clauses `s(a) :- p(a,a)`, `s(a) :- p(a,b)`, and `s(a) :- p(a,c)`. These are the clauses to consider when computing $I(s(a))$.

We will now prove that the least fix-point of function T_P coincides with the minimal model of program P, when restricted to ground atoms. To do that, we need an intermediate result that shows that the ground models of a given program P are the solutions of the equation $T_P(I) \preceq I$.

THEOREM 6.3.3 (MODELS AND T_P) *Given any interpretation* I *for a program* P*,* I *is a ground model for* P *if and only if* $T_P(I) \preceq I$.

PROOF. Consider any ground atom H and assume there are two clauses with H as their head: $H : -B_1$ and $H : -B_2$. By definition of model, each clause $H : -B_i$ is satisfied in I. Thus $I(H) \geq_S I(\exists B_i)$. Now, function T_P assigns to H the sum of the values assigned by I to $\exists B_1$ and $\exists B_2$, thus $T_P(I)(H) = I(\exists B_1) + I(\exists B_2)$. But the + operation coincides with the lub of the semiring, thus any value of the semiring which is greater than both $I(\exists B_1)$ and $I(\exists B_2)$ is also greater than their sum. Therefore $T_P(I)(H) \leq I(H)$. A similar reasoning works also for proving that if $T_P(I)(H) \leq I(H)$ for any ground atom H then I is a model. □

[2] Actually, an infinite number of them.

THEOREM 6.3.4 (GROUND MODELS AND FIX-POINT SEMANTICS) *Given an SCLP program* P*, we have that* $GA(M_P) = lfp(T_P)$[3].

PROOF. By definition of ground minimal model, $GA(M_P) = glb(\{I \mid I \text{ is a model for } P\})$. By the previous theorem, we get $GA(M_P) = glb(\{I \mid T_P(I) \preceq I\})$. By the classical results cited above [Tar55], this coincides with the least fix-point of T_P. □

Notice that the result of Theorem 6.2.5 (that is, the glb of two models is a model) can be proven also by using Theorem 6.3.3 (which states that a model is a prefixpoint of T_P) and by easily showing that the glb of two prefixpoints of a monotone operator (as T_P is) is a prefixpoint as well.

6.4 Proof-theoretic Semantics

We will define here a proof-theoretic semantics in the style of CLP [JL87]. We first, however, rewrite the program into a form that allows us to make the semantics treatment more uniform and simpler.

First, we rewrite each clause so that the head is an atom whose arguments are different variables. This means that we must explicitly specify the substitution that was written in the head, by inserting it in the body. That is, given a clause of the form $p(t_1, \ldots, t_n) : -B$, we transform it into $p(x_1, \ldots, x_n) : -\langle B, \theta \rangle$ where $\theta = \{x_1/t_1, \ldots, x_n/t_n\}$. Thus, bodies now have the following syntax: $B1 :: \langle B, \theta \rangle$. We recall that B can be either a collection of atoms or a value of the semiring. To give a uniform representation to bodies, we can define them as triples containing a collection of atoms (possibly empty), a substitution, and a value of the semiring (possibly, **1**). Thus, bodies are now of the form $B2 :: \langle LAT, \theta, a \rangle$. If we have a body belonging to the syntactic category $B1$ of the form $\langle a, \theta \rangle$, we get $\langle \Box, \theta, a \rangle$. If instead we have $\langle C, \theta \rangle$, where C is a collection of atoms, we get $\langle C, \theta, \mathbf{1} \rangle$. Therefore, clauses have now the syntax $CL1 :: H : -B2$.

Initial goals need to be transformed as well: given a goal $G = (: -C)$, where C is a collection of atoms, we get the goal $G' = (: -\langle C, \varepsilon, \mathbf{1} \rangle)$. The reason why we write the value **1** of the semiring is that this element is the unit element w.r.t. the operation we want to perform on it, that is, constraint combination.

In summary, given a SCLP program, we get a program in an intermediate language, whose syntax is as follows:

$$\begin{aligned} B2 &:: \langle LAT, \theta, a \rangle \\ CL1 &:: H : -B2 \\ P1 &:: CL1 \mid CL1, P1 \\ G1 &:: B2 \end{aligned}$$

Consider again our running example. The transformed program is then

[3] We recall that, for any interpretation I, $GA(I)$ is the restriction of I to ground atoms.

```
s(X)   :-  ⟨ p(X,Y), ε,0⟩.
p(X,Y) :-  ⟨ q(a), {X=a,Y=b}, 0⟩.
p(X,Y) :-  ⟨ r(a), {X=a,Y=c}, 0⟩.
q(X)   :-  ⟨ t(a), {X=a}, 0⟩.
t(X)   :-  ⟨□, {X=a}, 2 ⟩.
r(X)   :-  ⟨□, {X=a}, 3 ⟩.
```

Once we have transformed the given SCLP program into a program in the syntax just given, we can apply the following semantic rule. This rule defines the transitions of a nondeterministic transition system whose states are goals (according to the syntactic category G1).

$$\frac{\begin{array}{c} C = A, Cr \\ Cl = (A' : -\langle C_1, \theta_1, a_1\rangle) \text{ is a variant of a clause or a fact} \\ \theta' = mgu(A\theta, A'\theta_1) \end{array}}{\langle C, \theta, a\rangle \xrightarrow{Cl,\theta'} \langle (C_1, Cr), (\theta \circ \theta' \circ \theta_1), a \times a_1\rangle}$$

If the current goal contains an atom that unifies with the head of a clause, then we can replace that atom with the body of the considered clause, performing a step similar to the resolution step in CLP. The main difference here is that we must update the third element of the goal, that is, the semiring value associated to the goal: if before the transition this value is a and the transition uses a clause whose body has value a_1, then the value associated to the new goal is $a \times a_1$. The reason for using the $\times$ operation of the semiring is that this is exactly the operation used when accumulating constraints in the SCSP framework. Notice that we must use a *variant* of the clause involved in the rule, because we need fresh variables to avoid confusion between the variables of the clause and those of the current goal.

DEFINITION 6.4.1 (DERIVATIONS AND REFUTATIONS) *A derivation is a finite or infinite sequence of applications of the above rule. A refutation is a finite derivation whose final goal is of the form* $\langle \square, \theta, a\rangle$.

Note that in this chapter we give a simplified view of the solver for soft constraints, where the solver is *implemented* by some clauses in the program (see the last two clauses of our running example). These clauses have a special shape, since they are ground facts. Also, they are in a finite number, since our constraints have finite domains. Therefore, when executing such a program, the soft constraints (like `t` and `r` in our example) are *solved* using these clauses and, thus, they do not appear in the resulting final goal. This allows us to simplify the usual CLP operational semantic rules (like those in [JM94]), so that, instead of having computation states made by a current goal and a store, we can just have a goal, a substitution for the variables (this would be in the store in CLP) and a semiring value. Therefore, our operational semantic rule and states are consistent with those of CLP, in the case that we have 1) finite domain variables in the constraints, 2) solver specified by program clauses (this again can be done in this straightforward way only because we have finite domains), and 3) soft constraints. While the first two points are special cases, the third is an extension of CLP. Notice also that we can have just one rewriting rule (instead of several rules, as in [JM94]) because we first rewrite the program into a syntax that makes clauses and facts uniform. A last point to notice is that we do not test at each step if the constraint store is consistent (that is, if the semiring value collected is greater than 0) because the notion of inconsistency is not so strong in soft constraints, and we perform this check only at the end.

DEFINITION 6.4.2 (COMPACT REFUTATION SET) *Given an SCLP program* P*, its compact refutation set* $S(P)$ *is defined as follows:*

$$S(P) = \langle C, \theta_{|var(C)}, a\rangle \mid \langle C, \varepsilon, 1\rangle \rightarrow^* \langle \square, \theta, a\rangle\}.$$

$S(P)$ contains all triples representing all refutations for the given program. Note that we only record the part of θ which involves the variables in C. Notice also that a triple $\langle C, \theta, a\rangle$ may represent more than one refutation, but all these refutations start from the same goal C, build the same substitution θ, and generate the same semiring value a. Thus, we do not mind making them indistinguishable.

A property of the set $S(P)$ is that for any triple $\langle C, \theta, a\rangle$ in this set, there is also the triple $\langle C\theta, \varepsilon, a\rangle$. Notice that this is a generalized version of the lifting lemma of [Llo87].

THEOREM 6.4.3 (GOAL SPECIALIZATION) *Given an SCLP program* P *and the corresponding set* $S(P)$*, let us consider any triple* $\langle C, \theta, a\rangle$ *in* $S(P)$*. Then also the triple* $\langle C\theta, \varepsilon, a\rangle$ *is in* $S(P)$*.*

PROOF. Triple $\langle C, \theta, a\rangle$ says that there is a refutation, which starts from C, builds θ, and obtains the semiring value a. If we start from $C\theta$, we can construct a refutation which follows exactly the same steps as the previous one, and thus obtains the same semiring value, since θ is compatible with all these steps. $\square$

Our goal now is to study the correspondence between the operational semantics of a goal and its model-theoretic meaning. In particular, given a goal C, we will define the operational semantics of C in two different ways, to model the meaning of both $\forall C$ and $\exists C$.

6.4.1 Universal Closure

Among all refutations represented by the triples in $S(P)$, there are some that have the same first element, say C, and which build the empty substitution, ε, during the computation. These refutations have to be merged by the operational semantics since they represent different alternative branches for goal C, which lead to possibly different semiring values and that hold for any value of the variables of C. Thus, they naturally correspond to the intuitive meaning of $\forall C$.

When merging such elements of $S(P)$, the respective semiring values must be merged as well. This is done by performing the + operation. Actually, we should also merge those refutations that start from different goals, say C and C', such that C' is more instantiated than C. That is, such that there exists a substitution θ such that $C' = C\theta$. In fact, the semiring value obtained from C holds also for all goals that are more instantiated than C. For example, if we have the triples $\langle p(x), \varepsilon, v_1\rangle$ and $\langle p(a), \varepsilon, v_2\rangle$, the first refutation says that for all values of x, $p(x)$ gets the semiring value v_1. Thus, also $p(a)$ will get this value. On the other hand, the other refutation says that $p(a)$ gets the value v_2. Therefore, the value to be assigned to $p(a)$ is the lub between v_1 and v_2, which is v_1+v_2 by definition of $+$. It is possible, however, to show that if there is the triple $\langle C, \varepsilon, v\rangle$ in $S(P)$, then there is also the triple $\langle C\theta, \varepsilon, v\rangle$ for any θ. In other words, if we have a refutation starting from C, building the empty substitution and obtaining the semiring value v, then there is also a refutation starting from any goal more instantiated than C (that is, $C\theta$), which follows the same steps as the refutation for C and thus obtains the same semiring value. Therefore, we just need to merge those refutations that start from the same goal.

THEOREM 6.4.4 (MORE ABOUT GOAL SPECIALIZATIONS) *Given an SCLP program* P *and the corresponding set* $S(P)$*, let us consider any triple* $\langle C, \varepsilon, a\rangle$ *in* $S(P)$*. Then, for any* θ*, also the triple* $\langle C\theta, \varepsilon, a\rangle$ *is in* $S(P)$.

PROOF. The refutation represented by the triple $\langle C, \varepsilon, a\rangle$ does not bind any variable present in C. By starting from a goal $C\theta$, which is more instantiated than C, we can follow exactly the same steps as in the refutation for C. In fact, at each step we can use the same clause as before, since we know that such a step did not bind the variables in C. □

Let us now define a function $OS1_P$ (for Operational Semantics), which, given a goal, returns a value of the semiring by looking at all refutations for that goal that build the empty substitution.

DEFINITION 6.4.5 (FUNCTION $OS1_P$) *Given an SCLP program* P*, function* $OS1_P$: $LAT \to A$*, where* LAT *is the set of conjunctions of atoms and* A *is the semiring set, is defined as follows:* $OS1_P(C) = \sum_{\langle C,\varepsilon,a\rangle\in S(P)} a$.

Note that, when there is no triple $\langle C, \varepsilon, a\rangle \in S(P)$, function $OS1_P(C)$ returns the unit element of the $+$ operation, that is, $\mathbf{0}$. This is reasonable, since the absence of such triples means that there is no refutation starting from C and builds the empty substitution.

Notice that, if $OS1_P(C) = a$, there is not necessarily a refutation starting from C and obtaining the semiring value a. In fact, as noted above, the value a may have been obtained by combining several refutations, which may be incomparable with respect to $\leq_S$.

Let us now consider some examples of goal refutations and their operational semantics via function $OS1$. By considering the goal $\langle s(a), \{\varepsilon\}, 0\rangle$ in our running example, we get two refutations, which end respectively with the goals $\langle \Box, \{\varepsilon\}, 2\rangle$ and $\langle \Box, \{\varepsilon\}, 3\rangle$. Thus $S(P)$ contains the triples $\langle s(a), \{\varepsilon\}, 2\rangle$ and $\langle s(a), \{\varepsilon\}, 3\rangle$. Therefore $OS1_P(s(a)) =$

$\mathrm{min}(2,3) = 2$. Instead, $OS1_P(s(b)) = +\infty$ (which is the bottom of the semiring), since there is no refutation for $\langle s(b), \{\varepsilon\}, 0\rangle$. Consider now the goal $\langle s(x), \varepsilon, 0\rangle$. In this case, we get two refutations with the same final goals as above, and thus $S(P)$ contains the same elements as above. More precisely, $S(P)$ does not contain any triple of the form $\langle s(x), \varepsilon, \mathbf{a}\rangle$, and therefore $OS1_P(s(x)) = +\infty$.

A more complex example is related to the goal $\langle p(a,y), \{\varepsilon\}, 0\rangle$, which has two refutations ending with the goals $\langle \Box, \{y = b\}, 2\rangle$ and $\langle \Box, \{y = c\}, 3\rangle$, which are therefore represented in $S(P)$ by the triples $\langle p(a,y), \{y = b\}, 2\rangle$ and $\langle p(a,y), \{y = c\}, 3\rangle$. Therefore, the operational semantics of $p(a,y)$ is $OS1_P(p(a,y)) = +\infty$, even if there are refutations starting from $p(a,y)$ and obtaining a semiring value smaller than $+\infty$. Formally, this is due to the fact that such refutations all build a substitution different from ε, and thus are not considered by function $OS1_P$. Intuitively, the fact that the operational semantics assigns to $p(a,y)$ the worst semiring value, even if there are refutations that lead to better values, can be explained by considering that by only looking at the refutations, as the operational semantics does, there will be other domain values, say d, for which $p(a,d)$ does not hold. We will come back to this issue in the following paragraph, when we will formally compare the operational and the model-theoretic semantics. We can, however, already notice that this behavior leads to the same result (that is, the same semiring value) we computed via the model-theoretic tools (that is, function M_P) at the end of Section 6.2 for the formula $\forall y.p(a,y)$.

We now prove formally that in general the operational semantics of SCLP programs, as just defined via function $OS1_P$, for each goal C, computes the same semiring value as the model-theoretic semantics of $\forall C$, defined in Section 6.2, and thus also as the fix-point semantics of Section 6.3 when applied to any ground instantiation of C.

Before stating and proving the main theorem, however, it is useful to prove two lemmas, which will be used in the proof of the theorem. The first lemma basically proves that the operational semantics of a collection of ground atoms is the same as the multiplication of the operational semantics of the single atoms, while the second one uses the fact that we have an infinite domain, but a finite number of domain elements in the program, to prove that the operational semantics of a goal coincides with the glb of the operational semantics of all its ground instances.

LEMMA 6.4.6 (DISTRIBUTION) *Given an SCLP program* P, *consider two collections of atoms* C_1 *and* C_2. *Then we have that* $OS1_P(C_1, C_2) = OS1_P(C_1) \times OS1_P(C_2)$.

PROOF. By definition of $OS1_P$, we have that $OS1_P(C_1) \times OS1_P(C_2) = (\sum_{\langle C_1,\epsilon,\mathbf{a}\rangle \in S(P)} \mathbf{a}) \times (\sum_{\langle C_2,\epsilon,\mathbf{b}\rangle \in S(P)} \mathbf{b})$. Let us assume that we have n elements of the form $\langle C_1, \epsilon, \mathbf{a}\rangle$ and m elements of the form $\langle C_2, \epsilon, \mathbf{b}\rangle$ in $S'(P)$. Therefore, we will consider the semiring values $\mathbf{a}_1, \ldots, \mathbf{a}_n$ and $\mathbf{b}_1, \ldots, \mathbf{b}_m$. Notice that n and m may be infinite; however, since we work with complete lattices, the lub (+) and glb are defined also in the infinite case. Thus, we have $OS1_P(C_1) \times OS1_P(C_2) = (\sum_{i=1,\ldots,n} \mathbf{a}_i) \times (\sum_{j=1,\ldots,m} \mathbf{b}_j)$. Now, since $\times$ distributes over $+$ (by the definition of semiring), this can be rewritten as $\sum_{i=1,\ldots,n,j=1,\ldots,m}(\mathbf{a}_i \times \mathbf{b}_j)$. Consider now any refutation for the goal (C_1, C_2) which does not instantiate any variable, and assume, without loss of generality, that its selection strategy first "consumes" C_1 and then C_2. Then, after C_1 has been "consumed", we have a goal of the form $\langle C_2, \epsilon, a_i \rangle$. Then the refutation continues by "consuming" C_2, leading to a final goal of the form $\langle \Box, \epsilon, a_i \times b_j \rangle$. Now, $OS1_P(C_1, C_2)$ just sums all the semiring values of all such simple refutations, therefore, we have that $OS1_P(C_1, C_2) = \sum_{i,j}(a_i \times b_j)$, which is the same as $OS1_P(C_1) \times OS1_P(C_2)$ for what we have said before. $\Box$

LEMMA 6.4.7 (INFINITE DOMAIN) *Given an SCLP program* P*, consider any collection of atoms* C*. Then* $OS1_P(C) = \mathrm{glb}\{OS1_P(C\theta)$*, for all* θ *such that* $C\theta$ *is ground*$\}$.

PROOF. Given a goal C, function $OS1_P$ returns the sum of all the semiring values obtained by all those refutations of C which build the empty substitution. Let us call the set of such refutations $E(C)$. Since $C\theta$ is more instantiated than C, then the refutations of $C\theta$ will be able to use more clauses than those for C. In fact, if the head of a clause is as instantiated as a subgoal of $C\theta$ but more instantiated than the corresponding subgoal of C, then a refutation of $C\theta$ will be able to use it without building a non-empty substitution, while any refutation of C using such a clause will have to add a non-empty substitution to the substitution of the current goal. Therefore, we have that $E(C) \subseteq E(C\theta)$, which, by definition of $OS1_P$ and by the fact that the $+$ operation is the lub, means that $OS1_P(C) \leq_S OS1_P(C\theta)$. Now, this holds for any θ, but it is possible to show that for some θ we actually have $OS1_P(C) = OS1_P(C\theta)$. To do that, it is enough to consider a θ which binds all variables of C to domain elements not present in P. In this way, the refutations of $C\theta$ will not be able to use more clauses than those of C. Therefore, the glb of all $OS1_P(C\theta)$ for all θ is actually equal to $OS1_P(C)$, since we know that $OS1_P(C)$ is smaller than (or equal to) all elements $OS1_P(C\theta)$ and that it is actually equal to one of them. $\Box$

THEOREM 6.4.8 (OPERATIONAL MEANING OF $\forall C$) *Given a SCLP program* P*, consider a collection of atoms* C*. Then we have that* $M_P(\forall C) = OS1_P(C)$.

PROOF. If $OS1_P(C) = a$, it means that $S(P)$ contains triples of the form $\langle C, \varepsilon, a_i \rangle$, for $i = 1, \ldots, n$, such that $\sum_{i=1,\ldots,n} a_i = a$. Again, we recall that n may be infinite. We work with complete lattices, however, and thus both infinite lub (that is, sum) and infinite glb are defined.

We will prove the statement of the theorem by induction on the length of the longest of the refutations corresponding to such triples.

- **base case:** If all refutations $\langle C, \varepsilon, a_i \rangle$ have length 1, it means that there are n facts in the program of the form $C' : -\langle \Box, \theta_i, a_i \rangle$, where C' contains only variables and $C'\theta_i$ is equally or less instantiated than C. In fact, if the head of such facts would be more instantiated than C, then we would have a substitution different from ε in the refutation.

  By definition of model-theoretic semantics, $M_P(\forall C)$ is the greatest lower bound of the values given to each ground instantiation of C by function M_P. Now, consider any of such ground instantiations, say $C\theta$. Since $C'\theta_i$ is equally or less instantiated than C, it is also equally or less instantiated than $C\theta$. Therefore $M_P(C\theta) = \sum_{i=1,\ldots,n} a_i + \sum_{j=1,\ldots,m} b_j$, where the semiring values b_j are present in other m facts whose head is less instantiated than $C\theta$ but more instantiated than C. Now we must compute the glb of all such $M_P(C\theta)$ for all θ. But it is possible to see that there is a θ such that $M_P(C\theta)$ involves only the n facts with semiring values a_i: just choose a substitution, say θ', which binds all variables in C to domain elements not present in P. Therefore $M_P(C\theta') = \sum_{i=1,\ldots,n} a_i$. This means that the glb of all $M_P(C\theta)$ is exactly $M_P(C\theta')$, that is, a.

- **inductive case:** Let us assume that the statement of the theorem holds when the longest of the refutations of the form $\langle C, \varepsilon, \mathbf{a_i}\rangle$ has length n. Let us now consider a collection of atoms C such that the longest refutations for C have length $n+1$. We have two cases to consider: one occurs when C contains just one atom, and the other one when C has two or more atoms.

  - **1st case (one atom):** The reasoning we will use in this case is very similar to that of the base case. To consider the refutation $\langle A, \varepsilon, \mathbf{a}\rangle$, let us assume to start from the goal $\langle A, \varepsilon, \mathbf{1}\rangle$, where A is an atom. Consider now the set of n clauses contained in the program P of the form $A' : -\langle B_i, \theta_i, \mathbf{0}\rangle$, for $i = 1, \ldots, n$, where B_i is a collection of atoms and A' contains only variables and the same predicate symbol as A. To be used to expand our goal, it must be that $A'\theta_i$ is equally or less instantiated than A. In fact, if the head of such clauses were more instantiated than A, then we would have a substitution different from ε in the refutation. By definition of $OS1_P$, we have that $OS1_P(A) = \sum_{i=1,\ldots,n} OS1_P(B_i\theta_i)$.
  By definition of model-theoretic semantics, $M_P(\forall A)$ is the greatest lower bound of the values given to each ground instantiation of A by function M_P. Now, consider any of such ground instantiations, say $A\theta$. Since $A'\theta_i$ is equally or less instantiated than A, it is also equally or less instantiated than $A\theta$. Therefore $M_P(A\theta) = \sum_{i=1,\ldots,n} \mathbf{a_i} + \sum_{j=1,\ldots,m} \mathbf{b_j}$, where $\mathbf{a_i} = M_P(\forall B_i\theta_i)$ for $i = 1, \ldots, n$, and $\mathbf{b_j} = M_P(\forall B_j\theta_j)$ for $j = 1, \ldots, m$, with $A' : -\langle B_j, \theta_j, \mathbf{0}\rangle$, for $j = 1, \ldots, m$ other clauses where $A'\theta_j$ is more instantiated than A but less instantiated than $A\theta$.
  Now we must compute the glb of all such $M_P(A\theta)$ for all θ. But it is possible to see that there is a θ such that $M_P(A\theta)$ involves only the n clauses with semiring values $\mathbf{a_i}$: just choose a substitution, say θ', which binds all variables in A to domain elements not present in P. Therefore $M_P(A\theta') = \sum_{i=1,\ldots,n} \mathbf{a_i} = \mathbf{a}$. This means that the glb of all $M_P(A\theta)$ is exactly $M_P(A\theta')$, that is, $\sum_{i=1,\ldots,n} M_P(\forall B_i\theta_i)$. Now, by inductive hypothesis, this coincides with $\sum_{i=1,\ldots,n} OS1_P(B_i\theta_i)$, which is $OS1_P(A)$ as we stated before.
  - **2nd case (two or more atoms):** Let us assume that we start from the goal $\langle C, \varepsilon, \mathbf{1}\rangle$, and that $C = A, C'$. Then we have that:
  $M_P(\forall C)$ = {by definition of C }
  $M_P(\forall(A, C'))$ = {by definition of interpretation over a universally quantified formula}
  $glb_\theta\{M_P((A, C')\theta)\}$ =
  $glb_\theta\{M_P((A\theta, C'\theta)\}$ = {by definition of interpretation over a conjunction of formulas}
  $glb_\theta\{M_P(A\theta) \times M_P(C'\theta)\}$ = {by inductive hypothesis and by the assumption that C' is not empty}
  $glb_\theta\{OS1_P(A\theta) \times OS1_P(C'\theta)\}$ = {by 6.4.6}
  $glb_\theta\{OS1_P(A\theta, C'\theta)\}$ =
  $glb_\theta\{OS1_P((A, C')\theta)\}$ = {by 6.4.7}
  $OS1_P(A, C')$ = {by definition of C}
  $OS1_P(C)$.

□

6.4.2 Existential Closure

Theorem 6.4.8 relates the operational semantics of a goal C to the model-theoretic meaning of its universal quantification, that is, $\forall C$. We will now show that also the declarative meaning of the existential quantification of C, that is, $\exists C$, can be computed operationally. First we need to define a function $OS2_P(C)$ which combines all the refutations for C by summing all the corresponding semiring values.

DEFINITION 6.4.9 (FUNCTION $OS2_P$) *Given an SCLP program* P*, function* $OS2_P : LAT \to A$*, where* LAT *is the set of conjunctions of atoms and* A *is the semiring set, is defined as follows:* $OS2_P(C) = \sum_{\langle C,\theta,a\rangle \in S(P)} a$.

An intuitive explanation of why $OS2_P$ is defined this way is the following: since we are interested in describing the operational meaning of $\exists C$, we do not care which substitution is built during a refutation starting from C. Therefore, we take all refutations starting from C. Then, we choose the best among all such refutations by using the $+$ operator.

It is worth noting at this point that this definition of $OS2_P$ coincides with the following one: $\sum_{\langle C\theta,\varepsilon,a\rangle \in S(P)} a$. First we need the following lemma.

LEMMA 6.4.10 *Given an SCLP program* P *and a conjunction of atoms* C*, we have that, if* $\langle C\theta, \varepsilon, a\rangle$ *is in* $S(P)$*, then there exists* θ' *such that* $\langle C, \theta', a\rangle$*, with* $\theta' \leq \theta$ *is in* $S(P)$ *as well.*

PROOF. If we start a refutation from C instead of $C\theta$, we can follow the same steps as in refutation $\langle C\theta, \varepsilon, a\rangle$, because we have a more general goal. The only difference is that we may need to build a substitution, say θ', different from ε. Such a substitution, however, will be compatible with θ since we followed the same steps. Moreover, θ' cannot be more specific than θ, otherwise the refutation starting from $C\theta$ would have built a substitution different from ε. □

THEOREM 6.4.11 *Given an SCLP program* P *and a conjunction of atoms* C*, we have that* $\sum_{\langle C,\theta,a\rangle \in S(P)} a = \sum_{\langle C\theta,\varepsilon,a\rangle \in S(P)} a$.

PROOF. We will prove this theorem by showing that the set S_1 of semiring values reached by refutations of the form $\langle C, \theta, a\rangle$ coincides with the set S_2 of values reached by refutations of the form $\langle C\theta, \varepsilon, a\rangle$. In fact, if we show this, then, since $+$ is idempotent, the two sums coincide as well.

One direction of the proof has already been proven in Theorem 6.4.3: if $\langle C, \theta, a\rangle$ is a refutation, then $\langle C\theta, \varepsilon, a\rangle$ is a refutation as well. Therefore we have that $S_1 \subseteq S_2$. Now we have to prove that $S_2 \subseteq S_1$. In general, it is not true that, if $\langle C\theta, \varepsilon, a\rangle$ is a refutation, then $\langle C, \theta, a\rangle$ is a refutation as well. By Lemma 6.4.10, however, we have that, if $\langle C\theta, \varepsilon, a\rangle$ is a refutation, then there exists θ' such that $\langle C, \theta', a\rangle$, with $\theta' \leq \theta$, is a refutation as well. This is enough to prove that $S_2 \subseteq S_1$, since we are proving the equality between sets of semiring values and not sets of refutations. □

Another alternative way to define $OS2_P$ is by using the definition of $OS1_P$, as follows.

THEOREM 6.4.12 *Given an SCLP program* P *and a conjunction of atoms* C*, we have that* $OS2_P(C) = \sum_{C\theta \text{ ground}} OS1_P(C\theta)$.

PROOF. The statement of the theorem comes from the definition of $OS1_P$, the associativity and idempotency of $+$, and Theorem 6.4.4. □

Now we formally show that the result of the application of function $OS2_P$ over any goal C is a semiring value that coincides with the model-theoretic meaning of $\exists C$.

THEOREM 6.4.13 (OPERATIONAL MEANING OF $\exists C$) *Given an SCLP program* P*, consider a collection of atoms* C*. Then we have that* $M_P(\exists C) = OS2_P(C)$.

PROOF. By definition of interpretation, $M_P(\exists C) = lub\{M_P(C\theta)$ for all θ s.t. $C\theta$ is ground$\} = \sum_{C\theta \text{ ground}} M_P(C\theta)$. Since $C\theta$ is ground, Theorem 6.4.8 says that $M_P(C\theta) = OS1_P(C\theta)$. Therefore, we have that $M_P(\exists C) = \sum_{C\theta \text{ ground}} OS1_P(C\theta)$, which is exactly $OS2_P(C)$ by the Theorem 6.4.12. □

6.5 A semi-decidability result

In logic programming, looking for an answer of a given goal is a semi-decidable problem: if the goal is satisfiable, then a refutation for such a goal can be found in finite time; but if the goal is not satisfiable, then there are cases in which we can go on forever without detecting such an unsatisfiability.

We will now show that a similar semi-decidability result holds also for SCLP programs. More precisely, if the semantics of an SCLP goal is a semiring value greater than (or greater than or equal to) a certain semiring value k, then we can discover this in finite time.

The main idea, as in the logic programming case, is to visit the derivation tree of the given goal in a breadth-first way. We recall that the derivation tree of a goal is the tree whose root is the given goal, each node is a state in a derivation, and each path from the root to a leaf represents a derivation for that goal. If we visit this tree in a breadth-first way, it means that after k steps (that is, after examining k levels) we have seen all refutations of length k or less, and the first k steps of all derivations/refutations longer than k.

During this visit, we construct a sum of semiring values: starting from the $\mathbf{0}$ of the semiring at the root level, when we are at level k we add to the current sum the sum of all the semiring values associated to all refutations of length k. From the fact that the sum always leads to better values, as we go on we get new values that are better than or equal to the old ones.

Now, the crucial point is that, given any semiring value representing a partial sum, even though the values better than it are infinite (because a semiring can be infinite), the number of times that this value can be improved is finite. Thus, after examining at most a finite number of refutations, we will have computed the semiring value of the goal.

To formally prove this result, we will basically define an ordering among partial sums, and we will show that such an ordering is well-founded. Any partial sum will be represented by a suitable set of semiring values, obtained from the refutations examined so far.

We first need to give some formal definitions.

DEFINITION 6.5.1 (C-PRODUCTS) *Let* $C = \{c_i\}_{i=1,\dots,k}$ *be a finite set of elements of a semiring* S. *A (symbolic) product* $p = \prod_{i=1,\dots,k} c_i^{n_i}$ *is called a C-product and its value in* S *is denoted by* $[\![p]\!]$.

Informally, a C-product represents the semiring value obtained by a refutation. In fact, such a value is obtained by multiplying all the semiring values of the various clauses used by the refutation. Such values belong to a specific set: the set of all semiring values that appear in the given program, which we call C here. The exponents in the product are needed because each semiring value in C may be used several times during a refutation, and all these occurrences have to be considered, since in general the multiplicative operation is not idempotent.

DEFINITION 6.5.2 (PARTIAL ORDER OF C-PRODUCTS) *We define a partial ordering* $\sqsubseteq$ *on C-products as* $\prod_{i=1,\dots,k} c_i^{n_i} \sqsubseteq \prod_{i=1,\dots,k} c_i^{n_i'}$ *iff* $n_i \geq n_i'$ *for all* $i = 1, \dots, k$.

Given two C-products p and p', if $p \sqsubseteq p'$ it means that all the semiring values of C have more (or the same number of) occurrences in p than in p'. Thus, the value of p is worse than that of p' in the semiring: multiplying more items leads to worse results, by the intensivity of $\times$. Formally: $p \sqsubseteq p'$ implies $[\![p]\!] \leq_S [\![p']\!]$. Thus, we can say that p is "dominated" by p'.

DEFINITION 6.5.3 (SATURATION OF A SET OF C-PRODUCTS) *Given a finite or countably infinite set* P *of C-products, its saturation* $\overline{P}$ *is defined as* $\overline{P} = \{p \mid \exists p' \in P. p \sqsubseteq p'\}$.

By saturating a set of C-products P, we basically add to P all those other C-products which are dominated by some element in P. Notice that $[\![P]\!] = [\![\overline{P}]\!]$, where we extended the use of the semantic parenthesis $[\![]\!]$ from C-products to sets of C-products: $[\![Q]\!] = \sum_{p \in Q} [\![p]\!]$.

In our method to compute the semantics of a goal, every time we add an element to the current partial sum, such an element is a C-product. If the semiring value of this element is dominated by another one already in the partial sum, in reality the sum value does not change, because of the properties of the $+$ operation: if $a \leq_S b$ then $a + b = b$.

The main result of this section is that such a chain of partial sums has a finite number of distinct elements, and that, thus, after such a finite number of steps we have computed the semantics of a goal in an SCLP program. But before stating this results we need a lemma.

LEMMA 6.5.4 (WELL-FOUNDEDNESS) *Let $\mathcal{P}_C$ be the set of all saturated (i.e., $\overline{P} = P$) sets of C-products. This set is well-founded under the inverse proper inclusion relation $\supset$, i.e. all the chains $P_0 \subset P_1 \subset \ldots$ of elements of $\mathcal{P}_C$ are of finite length.*

PROOF. We first consider only chains where $P_j = \overline{R_j}$, with R_j finite. Also, without loss of generality, we assume $P_0 = R_0 = \emptyset$ and $R_j = R_{j-1} \cup \{p_j\}$, with of course $\forall p \in R_{j-1}, p_j \not\sqsubseteq p$.

We prove the property by mathematical induction on the number k of constants in C. If $k = 1$ the property is trivial, since $P_j = \{c^n \mid n \geq \overline{n_j}\}$ for some $\overline{n_j}$, and thus $P_j \subset P_{j+1}$ means $\overline{n_j} > \overline{n_{j+1}}$.

Let us now assume that the property holds for $k-1$ and prove it for k. We work by absurd, assuming that an infinite chain exists for k and constructing an infinite chain also for $k-1$. Let us now decompose p_j as $p_j = p'_j \times c_k^{n_k^j}$, with $p'_j = \prod_{i=1,\ldots,k-1} c_i^{n_i^j}$. Now, since we must have $p_h \not\sqsubseteq p_j$ when $j < h$, we either have $p'_h \not\sqsubseteq p'_j$ or $n_k^j > n_k^h$. We now use a colored graph method to help us in the proof. The nodes of the graph are the indices of the chain, and we draw a red arc from i to h in the latter case (that is, $n_k^j > n_k^h$) or when both conditions hold, a black arc otherwise. We now construct an infinite subsequence with no red arcs, i.e. where $p'_h \not\sqsubseteq p'_j$ for $h > j$.

Notice first that there is no infinite red path: in fact our graph is acyclic and exponent n_k^j cannot be indefinitely decreased. Now assume that we have already examined a finite initial segment of the chain, and we have already constructed a set of indexes I, initially empty. We assume inductively that there is no red arc outgoing from the indexes in I to the indexes of the rest of the chain. To find a new index to add to I, let us consider the first index of the rest of the chain. If it has no outgoing red arc, we are done. Otherwise, we follow any outgoing red arc, and we repeat the above procedure until an index without outgoing red arcs is found. Since there is no infinite red path, the procedure must terminate.

Finally, it is easy to see that for every P in $\mathcal{P}_C$ we have $P = \overline{R}$, for some finite R. In fact, let $R = \{p \in P \mid \forall p' \in P.p \not\sqsubseteq p'\}$. If R were infinite (but of course countable: $R = \{p_j\}_{j=1,2,\ldots}$), the chain $\emptyset \subset \overline{\{p_1\}} \subset \overline{\{p_1, p_2\}} \subset \ldots \subset \overline{\{p_j \mid j \leq h\}} \subset \ldots$ would be infinite. □

THEOREM 6.5.5 (FINITE CHAINS) *Let $P = \{p_j\}_{j=1,2,\ldots}$ be a finite or countably infinite set of C-products. Then there is a natural number N such that $\sum_{j=1,\ldots,N} [\![p_j]\!] = \sum_{j=1,2,\ldots} [\![p_j]\!]$.*

PROOF. The statement is trivial if P is finite. Thus, let us assume that P is infinite. According to Lemma 6.5.4, the chain $Q_1 \subseteq Q_2 \subseteq \ldots \subseteq Q_n \subseteq \ldots$ with $Q_h = \overline{\{p_j \mid j \leq h\}}$, is finite, i.e. there is a natural number N such that $Q_r = Q_N$ for $r \geq N$. Since $[\![Q_h]\!] = \sum_{j=1,\ldots,h} [\![p_j]\!]$ by definition, this proves the theorem. □

Summarizing, the formal developments of this section show that, by examining the refutations of a given SCLP goal while visiting the search tree in a breadth-first way, the semantics of the goal can be computed in a finite number of steps (that is, after examining a finite number of refutations). This result, together with the property that $+$ is the lub of the lattice, leads us to the following semi-decidability statement:

> It is semi-decidable to decide whether the semantics of a goal is in relation R with a certain semiring value $\mathbf{k}$, where $R \in \{>, \geq\}$.

In fact, if it is in such a relation, then the theorem of this section tells us that after a finite number of steps we have computed such semantics. Thus, after a smaller or equal number of steps the current partial sum will have value k or more, at which point we can stop and say that the goal has a semantics in relation R with k. If instead it is not in such a relation, then we do not have a method to know this in finite time. If there were, however, a method to semi-decide whether the semantics of a goal is in relation $\mathtt{not}(R)$ with a certain semiring value k, where $\mathtt{not}(R) \in \{<, =, \leq, \not\equiv\}$, with $\not\equiv$ meaning "incompatible with", then by Post's theorem (if a property and its complement are semi-decidable, then the property is decidable) we would conclude that it is decidable to know whether the semantics of a goal is equal to a certain semiring value k, which we know is not true for the special case of logic programming.

6.6 SCLPs with no functions

In the previous section we have proven that, if a goal of an SCLP program has a semiring value greater than, or greater than or equal to, some k, then it is possible to discover this in finite time. In this section we will show that, for the special class of SCLP programs with no functions, we have the additional result that once the program is fixed, the time for computing the value of any goal for this program is finite and bounded by a constant (see Theorem 6.6.4 later in this section). Thus, the semantics of SCLP programs without functions is decidable.

This result is based on the observation that, in SCLP programs without functions, we just have to consider a finite subclass of refutations, called in the following *simple* refutations, with a bounded length. After having considered all these refutations up to that bounded length, we have finished computing the semiring value of the given goal.

Notice that, while the absence of functions in SCLP programs is obviously a restriction, the underlying semiring could, in general, contain an infinite number of elements. Thus, this result is not so obvious as it may appear, since not all source of infiniteness are taken away.

First we define an alternative representation for refutations, which is based on trees.

DEFINITION 6.6.1 (REFUTATION TREE) *Given a refutation* r *as follows:*

$$\langle A, \varepsilon, \mathbf{1}\rangle \overset{Cl_1,\theta'_1}{\rightarrow} \langle C_1, \theta_1, a_1\rangle \overset{Cl_2,\theta'_2}{\rightarrow} \ldots \overset{Cl_n,\theta'_n}{\rightarrow} \langle \Box, \theta_n, \mathbf{a_n}\rangle$$

its refutation tree *is a labeled tree where each node is labeled by a clause instantiation of the form* $\langle Cl_i, \sigma_i\rangle$*, for* $i \in \{1, \ldots, n\}$*, where* $\sigma_i = \theta'_i$ *restricted onto the variables of the head of* Cl_i*. More precisely, starting from a single non-labeled node, the entire tree can be built as follows:*

```
for k:=1 to n do
```

> *Select the first non-labeled node (which is a leaf) in a depth-first visit of the current tree, and label it with* $\langle Cl_k, \sigma_k\rangle$*.*
> *If* Cl_k *is* $H :- B_1, \ldots, B_m$*, attach to this node* m *children.*

`end-for.`

DEFINITION 6.6.2 (SIMPLE REFUTATION) *Given a refutation tree, a path from the root to a leaf is called* simple *if all its nodes have different labels up to variable renaming. A refutation is a simple refutation if all paths from the root to a leaf in its refutation tree are simple.*

Thus, in a simple refutation it is not possible to use the same clause, instantiated in the same way, more than once on atoms that depend on each other.

We will now show that, if we delete all non-simple refutations from $S(P)$, we do not change the value computed by $OS1_P(C)$ for any C.

THEOREM 6.6.3 (SIMPLE REFUTATIONS ONLY) *Consider an SCLP program P with no functions and its compact refutation set $S(P)$, and let us call $S'(P)$ the subset of $S(P)$ containing all its simple refutations. Then, for any goal C, we have that $\sum_{\langle C,\varepsilon,a\rangle \in S'(P)} a = OS1_P(C)$.*

PROOF. Given any non-simple refutation r_1 for a goal, we can obtain a simple refutation r_2 for the same goal: just take the refutation tree of the non-simple refutation and delete the part of the tree between any two nodes with the same label up to variable renaming. It is easy to see that this new tree still represents a refutation, and that such a refutation is simple.

Now we can notice that refutation r_2 has fewer steps than r_1, and all the steps in r_2 are present in r_1. Therefore, by the extensivity of $\times$, the semiring value computed by r_2 is better than or equal to the one computed by r_1.

Now, $OS1_P(C)$ sums all the semiring values computed by all refutations in $S(P)$ that start from C and build ε. Therefore, considering what we have said before, the sum of the semiring values associated with r_1 and r_2 gives the semiring value of r_2, that is, of the simple refutation. Thus, by "forgetting" all non-simple refutations we do not change the result of the sum of all the semiring values, which is exactly $OS1_P(C)$. □

We will now prove that, given a goal C, there is only a finite number of simple refutations starting from C and building the empty substitution.

THEOREM 6.6.4 (FINITE SET OF SIMPLE REFUTATIONS) *Given an SCLP program P with no functions and a goal C, consider the set $SR(C)$ of simple refutations starting from C and building the empty substitution. Then $SR(C)$ is finite. Moreover, each refutation in $SR(C)$ has length at most N, where $N = \sum_{i=1}^{c} b^i$, c is the number of all clauses with head instantiated in all possible ways, and b is the greatest number of atoms in the body of a clause in program P.*

PROOF. Since a program contains a finite number of domain elements and of clauses, the number of different labels in a path of a simple refutation tree is finite. Thus the number of different refutation trees for simple refutations of C is finite. Therefore $SR(C)$ is finite.

Moreover, each simple refutation tree has the property that no two labels in a path from the root to a leaf are the same. Since such labels are clauses plus head instantiations, each path is as long as at most the number of all possible clause head instantiations, say c. Furthermore, the branching factor of a simple refutation tree depends on the number of atoms in the bodies of the clauses used in the refutation. Thus, the number of nodes of a simple refutation tree, and thus of steps in a simple refutation, is bounded by $\sum_{i=1}^{c} b^i$. □

Therefore, computing $OS1_P(C)$ involves looking at a finite number of bounded-length simple refutations. Thus, $OS1_P(C)$ can be computed in a finite number of steps.

COROLLARY 6.6.5 (FINITE NUMBER OF STEPS FOR $OS1_P$) *Given an SCLP program* P *with no functions, consider a collection of atoms* C. *Then we have that* $OS1_P(C)$ *can be computed in a finite number of steps.*

PROOF. The statement follows directly from the results of Theorem 6.6.3 (we can forget about the non-simple refutations) and Theorem 6.6.4 (the number of simple refutations is finite, and the length of each simple refutation is bounded by a constant N). In fact, we can consider only those refutations with length at most N, and among these we take only the simple refutations and we sum their semiring values. □

It is also easy to prove that $OS2_P(C)$ can be computed in a finite number of steps as well. Moreover, the T_P operator need to be applied only a finite number of times before reaching the fixpoint and thus computing the fixpoint semantics, as defined in Section 6.3.

6.7 An operational model for the SCLP Language using ASM

In this section we develop a simple interpreter for SCLP programs without functions using the *Gurevich's Abstract State Machines* (ASMs) (previously called *evolving algebras* [Gur95]). This formalism has been applied successfully to specify real programming languages and architectures, to validate standard language implementations, to verify real-time protocols, etc. (see annotated ASM bibliography in [B̈94]).

We specify an operational semantic model that directly reflects the intuitive procedural understanding of SCLP programs, but is formulated at the level of abstract search spaces. This combination of procedural and abstract features, provides a tool for mathematical description and analysis of the design decisions for the language, in a machine and proof system independent manner. Moreover, it lays the ground for provably correct stepwise refinement, through a hierarchy of specifications at lower level, down to implementations. Therefore, the ASM model we propose can be viewed as the first step versus a real implementation of the language: according to the result presented in [BBD+95], that can be reached through a sequence of (correctly) refined ASMs down to a model that can be mechanically translated into an executable code.

The model also suggests a possible strategy to optimize the parallel computation of all possible solutions, formalizing useful techniques that allow to avoid infinite and/or non optimal computations.

In Subsection 6.7.1 we briefly introduce the reader to the Gurevich's Abstract State Machine and then in Subsection 6.7.2 the operational semantics of the language in terms of ASMs is given.

6.7.1 The Gurevich's Abstract State Machine

We assume that the reader is familiar with the semantics of the Abstract State Machine defined in [Gur95], and we quote here only the essential definitions.

Given a program *Prog* (consisting of a finite number of *transition rules*) and a (class of) *initial state*(s) S_0, a *Gurevich's Abstract State Machine* $\mathcal{A}$ models the operational behavior of a real dynamic system $\mathcal{S}$ in terms of state transitions. A *state* is a first-order structure representing the instantaneous configuration of $\mathcal{S}$.

The basic form of a transition rule is the following *function update*

$$f(t_1, \dots, t_n) := t$$

where f is an arbitrary n-ary function and $t_1, \dots, t_n, t$ are first-order terms. To fire this rule to a state S_i, $i \geq 0$, evaluate all terms $t_1, \dots, t_n, t$ at S_i and update the function f to t on parameters $t_1, \dots, t_n$. This produces another state S_{i+1} which differs from S_i only in the new interpretation of the function f.

There are some rule constructors.

- The *conditional constructor* which produces "guarded" transition rules of the form:
  $$\texttt{if}\ g\ \texttt{then}\ R_1\ \texttt{else}\ R_2$$
  where g is a ground term (the *guard* of the rule) and R_1, R_2 are transition rules. To fire that new rule to a state S_i, $i \geq 0$, evaluate the guard; if it is true, then execute R_1, otherwise execute R_2. The `else` part may be omitted.

- The *parallel constructor* which produces the "parallel" synchronous rule of form:
  $$\texttt{var}\ x\ \texttt{ranges over}\ U$$
  $$R(x)$$
  where $R(x)$ is a generalized basic transition rule with a variable x ranging over the universe U.

  The operational meaning of the constructor consists in the execution of the new rule $R(x)$ for every $x \in U$.

We also make use of the following construct to let universes grow:

$$\texttt{extend}\ U\ \texttt{by}\ x_1, \dots, x_n\ \texttt{with}\ Updates\ \texttt{endextend}$$

where *Updates* may (and should) depend on the x_i and are used to define certain functions for (some of) the new objects x_i of the resulting universe U.

State transitions of $\mathcal{A}$ may be influenced in two ways: *internally*, through the rules of the program *Prog*, or *externally* through the modifications of the environment. A non static function f that is not updated in any transition rule is called *oracle function* if its values are given non-deterministically by an oracle.

A computation of $\mathcal{S}$ is modeled through a finite or infinite sequence $S_0, S_1, \ldots, S_n, \ldots$ of states of $\mathcal{A}$, where S_0 is an initial state and each S_{n+1} is obtained from S_n by firing simultaneously all of the rules of *Prog* to S_n.

6.7.2 The abstract operational model of SCLP

Signature

This section explains the basic data types (domains and functions) which are used in the transition rules of Section 6.7.2.

A SCLP computation can be seen as a parallel search for all solutions (final goals) of an initially given query, to be found in a space structured as a tree. Each query's computation is represented by a branch of the tree and the breadth search strategy is similar to the OR-Parallel model for Prolog.

The operational result of a computation is the (semiring) sum of the values of all computed solutions; since some branches might lead to infinite computations as well as to useless solutions (final goal whose semiring value a is worse than the sum of the values associated with the already computed solutions), we introduce some optimization techniques which allow to cut away all 'bad' branches, thus avoiding useless computations.

We represent the search space by a set *NODE* of nodes and the tree structure on *NODE* is formalized by a partial function $\mathtt{parent} : \mathsf{NODE} \to \mathsf{NODE}$ (undefined on $\mathtt{root}$), to be dynamically updated through rules.

All possible SCLP *computation states* intuitively correspond to the levels of the tree, therefore they are represented by the subset of *NODE* made up of all *active* nodes (see below) at that state.

Each element *n* of *NODE* has to carry the relevant information for a complete description – at this abstraction level – of the (sub)computation in the subtree rooted at *n*. This information consists of the current goal $G = \langle C, \theta, a \rangle$, where C is the collection of atoms which still have to be computed, θ is the substitution computed so far and a is the value of the semiring representing the *cost* (or level of preference, certainty, probability, etc.) accumulated performing the previous (part of the) computation.

Hence, we introduce the function

$$\mathtt{goal} : \mathsf{NODE} \to \mathsf{GOAL}$$

associating with each node the goal which has still to be computed.

GOAL is the set of SCLP goals, and it yields *GOAL* = $AT^* \times SUB \times A$, where *AT* denotes the universe of atoms, *SUB* denotes the set of substitutions, *A* is the set of the semiring $S = < A, +, \times, 0, 1 >$.

GOAL comes with the following functions yielding the three components of a goal:

$$\mathtt{listatom} : \mathsf{GOAL} \to \mathsf{AT}^*$$
$$\mathtt{sub} : \mathsf{GOAL} \to \mathsf{SUB}$$
$$\mathtt{cost} : \mathsf{GOAL} \to \mathsf{A}$$

SUB is a set of (not further specified) *substitutions* and comes together with three abstract functions

$$\mathtt{mgu} : TERM \times TERM \rightarrow SUB \cup \{\mathtt{undef}\}$$

associating with two terms either their most general unifier substitution or the answer that there is not any; the substitution applying function

$$\mathtt{apply} : TERM \times SUB \rightarrow TERM$$

yielding the result of applying the given substitution to a given term; substitution concatenation

$$\circ : SUB \times SUB \rightarrow SUB$$

TERM is the set of generic SCLP terms (atoms or collections of atoms).

CLAUSE denotes the set of program statements and comes with auxiliary functions yielding *head* and *body* of SCLP clauses

$$clhead : CLAUSE \rightarrow AT$$
$$clbody : CLAUSE \rightarrow GOAL$$

The current program is represented by a distinguished element *db* (database) of a universe *PROGRAM*. The candidate clauses, which have to be considered in a call to define the alternatives of a goal computation, are accessed using an abstract function

$$\mathtt{procdef} : AT \times PROGRAM \times \mathrm{N} \rightarrow CLAUSE^*$$

which yields the (renamed) clauses defining, in the given program, the predicate having the same functor as the selected atom. $\mathcal{N}$ represents the set of natural numbers and the following function yields the current renaming index

$$\mathtt{lev} : NODE \rightarrow \mathrm{N}$$

This function is recursively defined as follows:

$$\mathtt{lev}(n) = \begin{cases} 0 & \text{if } n = \mathtt{root} \\ 1 + \mathtt{lev}(\mathtt{parent}(n)) & \text{otherwise} \end{cases}$$

We also allow the use of the functions $\mathtt{head}$ and $\mathtt{tail}$ on lists with obvious meaning, and the function $\mathtt{proj}(i, L)$ to get the i-th element of a list L.

We use the abbreviation

$$\mathtt{act}(n) \equiv \mathtt{head}(\mathtt{listatom}(\mathtt{goal}(n)))$$

to select the current atom to compute at the level of the node n.

In order to control the computation of useless solutions, we define the subset *SUCCLEAF* = $\{n \in NODE \mid \mathtt{goal}(n) = \langle \Box, \theta, a \rangle, \theta \in SUB, a \in A\}$ of *NODE* containing all the leaf nodes of the search tree whose paths represent successful computations already performed. We also introduce the global variable

$$\mathtt{estimate} = \begin{cases} \sum_{n \in SUCCLEAF} \mathtt{cost}(\mathtt{goal}(n)) & \text{if } SUCCLEAF \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$$

which yields, at each state, the *best* semiring value among those associated to the already computed solutions. If the value a of the current goal $\langle C, \theta, a \rangle$ is worse than *estimate*, this goal computation is not to be carried on. We also make use of the predicate

$$\mathtt{admissible} : A \rightarrow \{\mathtt{true}, \mathtt{false}\}$$

such that

$$\mathtt{admissible}(c) = \begin{cases} \mathtt{true} & \text{if } \mathtt{estimate} \leq_S c \\ \mathtt{false} & \text{otherwise} \end{cases}$$

where $\leq_S$ is the partial order relation on the semiring S.

Since the SCLP terms do not contain any function symbol and the initial queries consist of only ground atoms, all infinite computations can be found out by looking for cycles; a cycle occurs along a branch when there are two nodes containing the same goal (without considering variable renaming). This strategy allow us to stop all infinite computations.

Therefore, we introduce the function

$$\mathtt{double} : \mathit{NODE} \rightarrow \{\mathtt{true}, \mathtt{false}\}$$

which is true on a given n if there is an ancestor node of n having the same sequence of atoms to compute. Formally speaking,

$$\mathtt{double}(n) = \begin{cases} \mathtt{true} & \text{if } \exists n' \in \mathtt{ancestor}(n) : \mathtt{comp}(n') \doteq [\mathtt{comp}(n)]\theta, \\ & \text{for some renamed substitution} \theta \\ \mathtt{false} & \text{otherwise} \end{cases}$$

where $\mathtt{ancestor}(n)$ yields the set of all nodes along the path from n to the $\mathtt{root}$, $\doteq$ stands for the syntactical equality, and $\mathtt{comp}(n) \equiv \mathtt{apply}(\mathtt{listatom}(\mathtt{goal}(n)), \mathtt{sub}(\mathtt{goal}(n)))$.

To be able to speak about *termination* we will distinguish between *active* and *non-active* nodes by the function

$$\mathtt{active} : \mathit{NODE} \rightarrow \{\mathtt{true}, \mathtt{false}\}$$

A node is born active and becomes non-active either after performing its own computation step or if one of the following three cases occurs:

- computation failure: when the first atom of the current goal does not unify with the head of the selected clause;
- the current goal will be a useless solution: when its semiring value is worse than *estimate*;
- infinite computation: when the current node already occurs along its path.

To keep information of all computed solutions of the initial query, we define the subset $\mathit{SOL} = \{g \in \mathit{GOAL} \mid g = \langle \Box, \theta, a \rangle\}$ of GOAL.

Transition rules

We now define the rules by which the system, starting from an *initial state*, tries to reach successful execution of the query $\mathtt{Q}$ ($\neq \Box$) by a given program $\mathtt{P}$. The initial state satisfies the conditions: $\mathit{NODE} = \{\mathtt{root}\}$; $\mathtt{goal}(\mathtt{root}) = \langle \mathtt{Q}, \varepsilon, 1 \rangle$; $\mathtt{active}(\mathtt{root}) = \mathtt{true}$; $\mathit{SUCCLEAF} = \emptyset$; $\mathit{SOL} = \emptyset$.

According to the operational semantic rule of a SCLP programs, the basic computation step consists of splitting the current goal $\langle(A, Cr), \theta, a\rangle$ associated with a node n, into m goals of the form $\langle(C_i, Cr), \theta \circ \theta_i' \circ \theta_i, a \times a_i\rangle$, where m is the number of candidate clauses for A, $A_i :- \langle C_i, \theta_i, a_i\rangle$, $1 \leq i \leq m$, is the i-th candidate clause for A and $\theta_i' = mgu(A\theta, A_i\theta_i)$. All the new m goals are computed in parallel. We can imagine the tree representing the computational space grows in breadth at each computation step performed by all the active nodes.

If an active node n has an admissible semiring value (i.e. $admissible(cost(goal(n)) \neq false$), and does not belong to an infinite branch (i.e. $double(n) \neq true$) and its computation has not failed (i.e. $sub(goal(n)) \neq undef$), by the following **Reduction Rule** as many child nodes of n are created as there are the candidate clauses for the activator $act(n)$ of n; and n becomes non-active. Otherwise the active node becomes non-active, as a leaf of the tree, without performing any splitting and stopping the computation along its path. In such a way all the "bad" branches are pruned. This strategy allows one to avoid the computation of non optimal solutions and infinite computations, thus optimizing the computation of the initial query.

Each child node is created *active* and receives as goal the result of the resolution between its parent goal and the relative selected candidate clause. This resolution step is performed by the macro *set_goal*(t_i, n) according to the semantic rule of the SCLP language.

Reduction Rule

```
var n ranges over NODE
    If active(n)
    then if sub(goal(n)) ≠ undef
            ∧ admissible(cost(goal(n)))
            ∧ ¬double(n)
         then let l = length(procdef(act(n), db, lev(n)))
              extend NODE by t1, t2, ... , tl with
                    parent(ti) := n
                    active(ti) := true
                    set_goal(ti, n)
              endextend
         active(n) := false
```

```
set_goal(ti, n) ≡ let clause = proj(i, procdef(act(n), db, lev(n)))
               goal(ti) := ⟨ [listatom(clbody(clause)) | tail(listatom(goal(n)))],
                            mgu(apply(act(n), sub(goal(n))),
                                apply(clhead(clause), sub(clbody(clause)))),
                            cost(goal(n)) × cost(clbody(clause))⟩
```

If the tree cannot grow any more, i.e. if all the nodes are not active, by the following **Collect Rule** all the computed solutions of the initial query are collected into the set *SOL*.

Collect Rule
If $\forall n \in \mathit{NODE}$: $\neg \mathrm{active}(n)$
then
$\mathit{SOL} := \bigcup_{n \in \mathit{SUCCLEAF}} \mathrm{goal}(n)$

According to the OS function definition, we can compute the operational meaning of the program P with initial query Q as $OS_P(Q) = \sum_{g \in \mathit{SOL}} \mathrm{cost}(g)$.

6.8 Related work

In [Fit81] a bilattice structure is used to model the presence of a family of truth values in logic programming, and a fixpoint semantics for this kind of programs is given. In this approach, the *meet* and *join* operators of the bilattice are used as extensions of classical *and* and *or*, while in our approach we use such operators to model the universal and existential quantification, but we adopt a different operator (the multiplicative operation of the semiring, possibly different from the *meet* operator), to extend the *logical and*. The bilattice structure he uses, allows one to always have a negation operator, which in general we do not have. By having less properties for our structure, we can model also situations where the multiplicative operator is not idempotent, like optimization and probabilistic problems.

The approach taken in [Emd86] associates a value to each clause. From such values, taken from $[0, 1]$, a value is also associated to each atom. Then, atoms are combined by using the min and max operators. Thus, this kind of logic programming is similar to what we have when using the fuzzy semiring. He also gives a model-theoretic semantics, a fixpoint semantics, and an operational semantics based on game theory.

A recent approach to multi-valued logic programming [MPS97] uses bilattices with two orderings to model both truth and knowledge levels. The resulting logic programming semantics is just operational and fix-point, while no model-theoretic semantics is presented. Moreover, the presence in our approach of just one ordering (modeling truth levels) is not a restriction, since the vectorization of several semirings is still a semiring (see Section 2.3.8) and thus optimization based on multiple criteria can be cast in our framework as well.

From another point of view, where classical constraints are extended to have several degrees of satisfaction, a related approach is HCLP (Hierarchical CLP) [BMMW89], where each constraint has a level of importance (like *strong, weak, required*), and these levels are used to decide which constraints to satisfy. A constraint, however, can only be satisfied or not, and thus HCLP is a *crisp* formalism. Moreover, their treatment is only algorithmic, and they do not provide their language with a fix-point or a model-theoretic semantics.

6.8.1 The CLP(FD,S) system

In this section we want to give a brief description of the CLP(FD,S) language that represents a real implementation of that defined in this chapter. The language was developed on top of the existing constraint language clp(fd) [CD96] by the LOCO research group at INRIA-Roquencourt. At the URL http://pauillac.inria.fr/~georget/clp_fds.html there is the latest version of this system, and in [GC98, Geo99] there is a description of their implementation.

In this implementation, full semiring-based arc-consistency (as described in Section 3.1) is used to keep the computation states compact and to check them for inconsistencies. In general, each domain element, as well as each constraint tuple, can be assigned a semiring value. They also provide, however, useful built-ins that assign just two values to each arithmetic constraint: one to the tuples which satisfy it, and the other one to those which do not satisfy it. In this way, the modeling is easier and the computation faster, and in some cases it is expressive enough.

Other plug-ins that we can mention are: arithmetic constraints, boolean constraints, meta constraints, optimization constraints, global constraints, predicates for constraints retraction and some facilities to build new constraints.

The kernel of clp(FD,S), is called SFD and is generic with respect to the semiring. Hence, the users are able to generate new languages (new solvers) by specifying semirings, the rest of the implementation being unchanged. For example, clp(FD,S) allows the computation of satisfaction and optimization of: usual CSPs (clp(FD,Bool)), Fuzzy CSPs, hierarchical CSPs (clp(FD,Fuzzy)), hypothetical CSPs (clp(FD,Sets)).

The language has been ported over Solaris, SunOS, GNU-Linux (i86) and requires GNU C (gcc) version 2.4.5 or higher. Some performance statistics may be found in [Geo99].

6.9 Conclusions

In this chapter a programming framework to deal with soft constraints has been introduced. We have described its syntax and we give several (equivalent) semantics: the operational, the model-theoretic and the fix-point one. Moreover, a generalization of the semidecidability result already known for logic programs has been described and proven: we can decide if the semantics of an SCLP goal has semiring value greater than k.

In this chapter a detailed semantics of the language has been given but no real applications to specific problems are given. We will do this in the next chapter where we will use this language to describe and solve shortest path problems.

Chapter 7

SCLP and Generalized Shortest Path Problems

Abstract

In recent years, an increasing number of researchers, from both the Operations Research (OR) and the Artificial Intelligence (AI) communities, have investigated the possibility of integrating methodologies of the two fields in order to obtain better results in solving combinatorial optimization problems. In this chapter we show how the SCLP paradigm could be useful to merge together AI and OR interests.

In fact, in this context constraint programming may be imposed as a suitable environment for performing such an integration, and the level of preference introduced can be used to discriminate between several solutions. In particular, we study the relationship between Constraint Programming (CP) and Shortest Path (SP) problems. More specifically, we show that classical, multicriteria, partially ordered, and modality-based SP problems can be naturally modeled and solved within the Soft Constraint Logic Programming (SCLP) framework, where logic programming is coupled with soft constraints. In this way we provide this large class of SP problems with a high-level and declarative linguistic support whose semantics takes care of both finding the cost of the shortest path(s) and also of actually finding the path(s). On the other hand, some efficient algorithms for certain classes of SP problems can be exploited to provide some classes of SCLP programs with an efficient way to compute their semantics.

Shortest Path (SP) problems [Dre69, TCR90] are among the most studied network optimization problems. They are mainly used to represent and solve transportation problems, where the optimization may involve different criteria, for example cost, time, resources, and so on. Most interesting is the multi-criteria case, where the optimization involves a set of criteria to be all optimized [PS88].

In this chapter, we propose the Soft Constraint Logic Programming (SCLP) framework [BMR97a, GC98] as a linguistic support and a high-level and flexible programming environment in which to model SP problems naturally and solve them efficiently.

Here we consider several versions of SP problems, from the classical one to the multi-criteria case, from partially ordered SP problems to those that are based on modalities, and we show how to model and solve them via SCLP programs. The basic idea is that soft constraints allow to faithfully represent the optimization criteria, and CLP provides a declarative way to describe the given SP problem. Moreover, this way of modeling and solving SP problems allows to associate with such problems both a declarative and an operational semantics.

The main results of the chapter are as follows:

- Both classical, multi-criteria, partially-ordered, and modality-based SP problems are given a modelization as SCLP programs; such programs are able to find both the cost of the shortest path(s) and also the shortest path itself.

- A general methodology is provided to find a non-dominated path for both multi-criteria and partially-ordered SP problems; this methodology is based on a change of semiring, but does not require any change in the underlying SCLP syntax and semantics.

- A new algorithm to obtain the semantics for a particular class of SCLP programs is given, which is obtained by using a generalized version of the Floyd-Warshall algorithm [Flo62] for SP problems.

The chapter (based on the ideas developed in [BMR99, BMR00b]) is organized as follows: Section 7.1 shows the construction to pass from a classical SP problem to a CLP program, while Section 7.2 considers multi-criteria and partially-ordered SP problems, and Section 7.3 deals with modality-based SP problems. Finally, Section 7.4 concludes the chapter by providing a class of SCLP programs with an efficient algorithm to compute their semantics.

7.1 Classical SP Problems

A shortest path (SP) problem can be represented as a directed graph $G = (N, E)$, where each arc $e \in E$ from node p to node q ($p, q \in N$) has associated with a label representing the cost of the arc from p to q.

There are four versions of the problem: the single pair problem, the single source problem, the single sink problem, and the all pair problem. The single source and the single sink problems, however, are directional duals of each other, the single pair problem at least partially solves the single source/sink problem, and one way to solve the all pair problem is to solve n single source/sink problems. For these reasons, the single sink (source) is fundamental and we concentrate on it. Given a set S of nodes and any node v in S (the sink), a solution of the problem consists of finding a path (or a set of paths) between any node of S and v, whose cost is minimal.

Consider for example the SP problem represented in Figure 7.1: each arc has associated with it a label representing the cost (in money, time, space, ...) of that arc. In this example node v is the sink. Thus, given any node, we want to find a path from this node to the sink v (if it exists) that minimizes the cost.

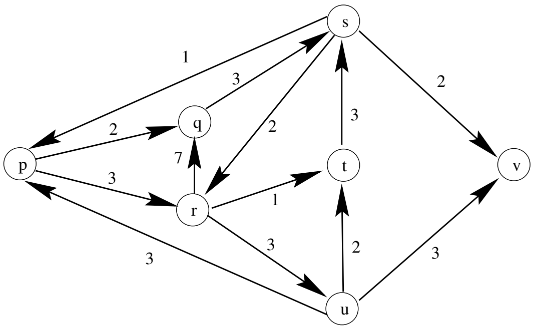

Figure 7.1: An SP problem.

To represent the classical version of SP problems, we consider SCLP programs over the semiring $S = \langle N, \min, +, +\infty, 0\rangle$, which, as noted above, is an appropriated framework to represent constraint problems where one wants to minimize the sum of the costs of the solutions.

From any SP problem we can build an SCLP program as follows. For each arc we have two clauses: one describes the arc and the other one its cost. More precisely, the head of the first clause represents the starting node, and its body contains both the final node and a predicate, say c, representing the cost of the arc. Then, the second clause is a fact associating with predicate c its cost (which is a semiring element). For example, if we consider the arc from p to q with cost 2, we have the clause

$\texttt{p :- } c_{pq}\texttt{, q.}$

and the fact

$c_{pq}\texttt{ :- 2.}$

Finally, we must code that we want v to be the final node of the path. This is done by adding a clause of the form `v :- 0.` Note also that any node can be required to be the final one, not just those nodes without outgoing arcs (like v is in this example). The whole program corresponding to the SP problem in Figure 7.1 can be seen in Table 7.1.

To compute a solution of the SP problem it is enough to perform a query in the SCLP framework; for example, if we want to compute the cost of the path from r to v we have to perform the query `:- r`. For this query, we obtain the value 6, that represents the cost of the best path(s) from r to v.

Notice that to represent classical SP problems in SCLP, we do not need any variable. Thus the resulting program is propositional. This program, however, while giving us the cost of the shortest paths, does not give us any information about the arcs which form such paths. This information could be obtained by providing each predicate with an argument, which represents the arc chosen at each step.

Figure 7.2 shows the same SP problem of Figure 7.1 where the arcs outgoing each node have been labeled with different labels to distinguish them. Such labels can then be

p :- c_{pq}, q.	c_{pq} :- 2.
p :- c_{pr}, r.	c_{pr} :- 3.
q :- c_{qs}, s.	c_{qs} :- 3.
r :- c_{rq}, q.	c_{rq} :- 7.
r :- c_{rt}, t.	c_{rt} :- 1.
r :- c_{ru}, u.	c_{ru} :- 3.
s :- c_{sp}, p.	c_{sp} :- 1.
s :- c_{sr}, r.	c_{sr} :- 2.
s :- c_{sv}, v.	c_{sv} :- 2.
t :- c_{ts}, s.	c_{ts} :- 3.
u :- c_{up}, p.	c_{up} :- 3.
u :- c_{ut}, t.	c_{ut} :- 2.
u :- c_{uv}, v.	c_{uv} :- 3.
v :- 0.	

Table 7.1: The SCLP program representing the SP problem in Figure 7.1.

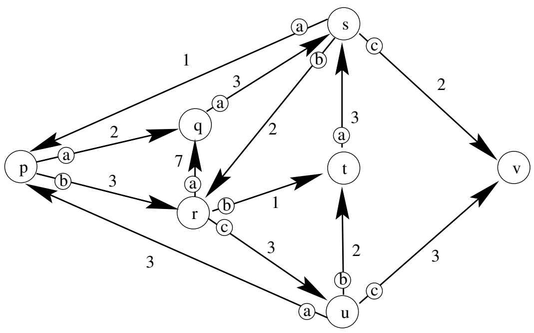

Figure 7.2: An SP problem with labeled arcs.

coded into the corresponding SCLP program to "remember" the arcs traversed during the path corresponding to a solution. For example, clause

`p :-` c_{pq}`, q.`

would be rewritten as

`p(a) :-` c_{pq}`, q(X).`

Here, constant a represents one of the arcs going out of `p`: the one which goes to `q`. If all clauses are rewritten similarly, then the answer to a goal like `:- r(X)` will be both a semiring value (in our case 6) and a substitution for X. This substitution will identify the first arc of a shortest path from r to v. For example, if we have $X = b$, it means that the first arc is the one that goes from r to t. To find a complete shortest path, we just need to compare the semiring values associated with each instantiated goal, starting from r and following the path. For example, in our case (of the goal $\exists X.r(X)$) we have that the answer to the goal will be $X = c$ with semiring value 6. Thus, we know that a shortest path from r to v can start with the arc from r to u. To find the following arc of this path, we compare the semiring values of $u(a)$, $u(b)$, and $u(c)$. The result is that $u(c)$ has the smallest value, which is 3. Hence, the second arc of the shortest path we are constructing is the one from u to v. The path is now finished because we reached v, which is our sink.

Notice that a shortest path could be found even if variables are not allowed in the program, but more work is needed. In fact, instead of comparing different instantiations of a predicate, we need to compare the values associated with the predicates that represent nodes reachable by alternative arcs starting from a certain node, and sum them to the cost of such arcs. For example, instead of comparing the values of $p(a)$ and $p(b)$ (Figure 7.2), we have to compare the values of $q + 2$ and of $r + 3$ (Figure 7.1).

A third alternative to compute a shortest path, and not only its cost, is to use lists: by replacing each clause of the form

`p :-` c_{xy}`, q.`

with the clause

`p([a|T]) :-` c_{xy}`, q(T).`

during the computation we also build the list containing all arcs that constitute the corresponding path. Thus, by giving the goal `:- p(L).`, we would get both the cost of a shortest path and also the shortest path itself, represented by the list L.

An alternative representation, probably more familiar for CLP-ers, of SP problems in SCLP is one where there are facts of the form

```
c(p,q) :- 2.
⋮
c(u,v) :- 3.
```

to model the graph, and the two clauses

```
path(X,Y) :- c(X,Y).
path(X,Y) :- c(X,Z), path(Z,Y).
```

to model paths of length one or more. In this representation the goal `:- path(p,v).` represents the cost of the shortest path from `p` to v. This representation is obviously more compact than the one in Table 7.1, and has equivalent results and properties. In this chapter, however, we will use the simpler representation, used in Table 7.1, where all clauses have, at most, one predicate in the body. The possibility of representing SP problems with SCLP programs containing only such kind of clauses is important, since it will allow us to use efficient algorithms to compute the semantics of such programs (see Section 7.4 for more details).

7.2 Partially-ordered SP Problems

Sometimes, the costs of the arcs are not elements of a totally ordered set. A typical example is obtained when we consider multi-criteria SP problems.

In this case the goal of the problem is to find a solution that is *Pareto-optimal*. A solution is Pareto-optimal if by reallocation you cannot make someone better off without making someone else worse off. In Pareto's words [Par06]:

> We will say that the members of a collectivity enjoy *maximium ophelimity* in a certain position when it is impossible to find a way of moving from that position very slightly in such a manner that the ophelimity enjoyed by each of the individuals of that collectivity increases or decreases. That is to say, any small displacement in departing from that position necessarily has the effect of increasing the ophelimity which certain individuals enjoy, and decreasing that which others enjoy, of being agreeable to some, and disagreeable to others.

Perhaps the best description of Pareto-optimality is the underutilized one coined by Maurice Allais [All43]: an allocation is "Pareto-optimal" if there is an "absence of distributable surplus".

As an example of multi-criteria SP problem we can consider the situation where an arc represents a piece of highway between two cities; in this case, we can label each arc both with the cost and with the time needed to follow this piece. The goal is to find the paths that have both the minimum overall cost and the minimum overall time. In this example, there may be cases in which the labels of two arcs are not compatible, like $\langle \$5, 20' \rangle$ and $\langle \$7, 15' \rangle$. In general, when we have a partially ordered set of costs, it may be possible to have several paths, all of which are not *dominated* by others, but which have different incomparable costs.

Consider, for example, the multi-criteria SP problem shown in Figure 7.3: each arc has associated with it a pair representing the weight of the arc in terms of *cost* and *time*. Given any node, we want to find a path from this node to v (if it exists) that minimizes both criteria.

We can translate this SP problem into the SCLP program in Table 7.2. This program works over the semiring

$$\langle \mathsf{N}^2, \min', +', \langle +\infty, +\infty \rangle, \langle 0, 0 \rangle \rangle,$$

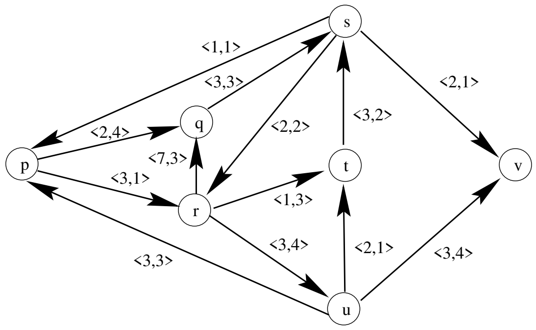

Figure 7.3: A multi-criteria SP problem.

where min' and $+'$ are classical min and $+$, suitably extended to pairs. In practice, this semiring is obtained by putting together, via the Cartesian product, two instances of the semiring $\langle N, \min, +, +\infty, 0\rangle$ (we recall that the Cartesian product of two c-semirings is a c-semiring as well [BMR97b]). One of the instances is used to deal with the cost criteria, the other one is for the time criteria. By working on the combined semiring, we can deal with both criteria simultaneously: the partial order will tell us when a pair $\langle$ cost, time $\rangle$ is preferable to another one, and also when they are not comparable.

To give an idea of another practical application of partially ordered SP problems, just think of network routing problems where we need to optimize according to the following criteria: minimize the time, minimize the cost, minimize the number of arcs traversed, and maximize the bandwidth. The first three criteria correspond to the same semiring, which is $\langle N, min, +, +\infty, 0\rangle$, while the fourth criteria can be characterized by the semiring $\langle [0, U], max, min, 0, U\rangle$, where U is the maximal bandwidth in an arc. In this example, we have to work on a semiring that is obtained by vectorizing all these four semirings. Each of the semirings is totally ordered but the resulting semiring, whose elements are four-tuples, is partially ordered.

Notice that the only difference between the structure of the program in Table 7.2 and the one in the previous section is that here we have costs represented by pairs of values, and, since we have a partial order, two such pairs may possibly be incomparable. This may lead to a strange situation while computing the semantics of a given goal. For example, if we want to compute the cost and time of a best path from p to v, by giving the query `:- p.`, the resulting answer in this case is the value $\langle 7, 7\rangle$. While the semiring value obtained in totally ordered SCLP programs represented the cost of one of the shortest paths, here it is possible that there are no paths with this cost: the obtained semiring value is in fact in general the greatest lower bound of the costs of all the paths from p to v. This behavior comes from the fact that, if different refutations for the same goal have different semiring values, the SCLP framework combines them via the $+$ operator of the semiring (which, in the case of our example, is the min' operator). If the semiring is partially ordered, it may be that $a + b$ is different from both a and b. On the contrary, if we have a total order $a + b$ is always either a or b.

This, of course, is not satisfactory, because usually one does not want to find the great-

p :- c_{pq}, q.	c_{pq} :- < 2,4 >.
p :- c_{pr}, r.	c_{pr} :- < 3,1 >.
q :- c_{qs}, s.	c_{qs} :- < 3,3 >.
r :- c_{rq}, q.	c_{rq} :- < 7,3 >.
r :- c_{rt}, t.	c_{rt} :- < 1,3 >.
r :- c_{ru}, u.	c_{ru} :- < 3,4 >.
s :- c_{sp}, p.	c_{sp} :- < 1,1 >.
s :- c_{sr}, r.	c_{sr} :- < 2,2 >.
s :- c_{sv}, v.	c_{sv} :- < 2,1 >.
t :- c_{ts}, s.	c_{ts} :- < 3,2 >.
u :- c_{up}, p.	c_{up} :- < 3,3 >.
u :- c_{ut}, t.	c_{ut} :- < 2,1 >.
u :- c_{uv}, v.	c_{uv} :- < 3,4 >.
v :- < 0,0 >.	

Table 7.2: The SCLP program representing the multi-criteria SP problem in Figure 7.3.

est lower bound of the costs of all paths from the given node to the sink, but rather prefers to have one of the non-dominated paths. To solve this problem, we can add variables to the SCLP program, as we did in the previous section, and also change the semiring. In fact, we now need a semiring that allows us to associate with each node the set of the costs of all non-dominated paths from there to the sink. In other words, starting from the semiring $S = \langle A, +, \times, 0, 1 \rangle$ (which, we recall, in our case is $\langle N^2, min', +', \langle +\infty, +\infty \rangle, \langle 0, 0 \rangle \rangle$), we now have to work with the semiring $P^H(S) = \langle P^H(A), \uplus, \times^*, \emptyset, A \rangle$, where:

- $P^H(A)$ is the Hoare Power Domain [Smy78] of A, that is, $P^H(A) = \{S \subseteq A \mid x \in S, y \leq_S x \text{ implies } y \in S\}$. In words, $P^H(A)$ is the set of all subsets of A which are downward closed under the ordering $\leq_S$. It is easy to show that such sets are isomorphic to those containing just the non-dominated values. Thus, in the following we will use this more compact representation for efficiency purposes. In this compact representation, each element of $P^H(A)$ will represent the costs of all non-dominated paths from a node to the sink;

- the top element of the semiring is the set A (its compact form is $\{1\}$, which in our example is $\{\langle 0, 0 \rangle\}$);

- the bottom element is the empty set;

- the additive operation $\uplus$ is the *formal union* [Smy78] that takes two sets and obtains their union;

- the multiplicative operation $\times^*$ takes two sets and produces another set obtained by multiplying (using the multiplicative operation $\times$ of the original semiring, in our case +') each element of the first set with each element of the second one;

- the partial order of this semiring is as follows: $a \leq_{P^H(S)} b$ iff $a \uplus b = b$, that is, for each element of a, there is an element in b which dominates it (in the partial order $\leq_S$ of the original semiring).

From the theoretical results in [Smy78], adapted to consider c-semirings, we can prove that $P^H(S)$ and its more compact form are indeed isomorphic. Moreover, we can also prove that given a c-semiring S, the structure $P^H(S)$ is a c-semiring as well.

THEOREM 7.2.1 **($P^H(S)$ IS A C-SEMIRING)** *Given a c-semiring* $S = \langle A, +, \times, 0, 1\rangle$, *the structure* $P^H(S) = \langle P^H(A), \uplus, \times*, \emptyset, A\rangle$ *obtained using the* Power domain of Hoare *operator is a c-semiring.*

PROOF. It easily follows from the properties of the $\times$ operator in the c-semiring S and from the properties (commutativity, associativity, and idempotency) of the formal union $\uplus$ in $P^H(S)$. □

Note that in this theorem we do not need any assumption over the c-semiring S. Thus, the construction of $P^H(S)$ can be done for any c-semiring S. Notice also that, if S is totally ordered, the c-semiring $P^H(S)$ does not give any additional information w.r.t. S. In fact, if we consider together the empty set (with the meaning that there are no paths) and the set containing only the bottom of A (meaning that there exists a path whose cost is 0), it is possible to build an isomorphism between S and $P^H(S)$ by mapping each element p (a set) of $P^H(A)$ onto the element a of A such that $a \in p$ and a dominates all elements in the set p.

The only change we need to make to the program with variables, in order to work with this new semiring, is that costs now have to be represented as singleton sets. For example, clause c_{pq} `:- < 2, 4 >.` will become c_{pq} `:- {< 2, 4 >}..`

Let us now see what happens in our example if we move to this new semiring. First we give a goal like `:- p(X) ..` As the answer, we get a set of pairs, representing the costs of all non-dominated paths from p to v. All these costs are non-comparable in the partial order, thus the user is requested to make a choice. This choice, however, could identify a single cost or also a set of them. In this second case, it means that the user does not want to commit to a single path from the beginning and rather prefers to maintain some alternatives. The choice of one cost of a specific non-dominated path will thus be delayed until later. If instead the user wants to commit to one specific cost at the beginning, say $\langle c_1, c_2\rangle$, he/she then proceeds to find a path that costs exactly $\langle c_1, c_2\rangle$. By comparing the answers for all goals of the form $p(a)$, where a represents one of the arcs out of p, we can see which arc can start a path with cost $\langle c_1, c_2\rangle$. In fact, such an arc will be labeled $\langle c_{1a}, c_{2a}\rangle$ and will lead to a node with an associated set of costs $\langle c_3, c_4\rangle$ such that $\langle c_3, c_4\rangle \times \langle c_{1a}, c_{2a}\rangle = \langle c_1, c_2\rangle$. Suppose it is the arc from p to q, which has cost $\langle 7, 3\rangle$. Now we do the same with goals of the form $q(a)$, to see which is the next arc to follow. We now, however, have to look for the presence of a pair $\langle c_3, c_4\rangle$ such that $\langle c_3, c_4\rangle \times\langle 7, 3\rangle = \langle c_1, c_2\rangle$.

Notice that each time we look for the next arc, we choose just one alternative and we disregard the others. If we used a fixpoint (or any bottom-up) semantics to compute the answer of the initial goal :- p(X)., then all the other answers we need for this method have already been computed, thus the method does not require any additional computational effort to find a non-dominated path.

Notice also that the sets of costs associated with each node are non-dominated. Thus, the size of these sets in the worst case is the size of the maximal ”horizontal slice” of the partial order: if we can have at most N non-dominated elements in the partial order, then each of such sets will have size at most N. Of course, in the worst case N could be the size of the whole semiring (when the partial order is completely ”flat”).

Most classical methods to handle multi-criteria SP problems find the shortest paths by considering each criteria separately, while our method deals with all criteria at once. This allows us to obtain optimal solutions, which are not generated by looking at each single criteria. In fact, some optimal solutions could be non-optimal in each of the single criteria, but still are incomparable in the overall ordering. Thus, we offer the user a greater set of non-comparable optimal solutions. For example, by using the time-cost multi-criteria scenario, the optimal solution w.r.t. time could be 1 minute (with cost of 10 dollars), while the optimal solution w.r.t. cost could be 1 dollar (with time of 10 minutes). By considering both criteria together, we could also obtain the solution with 2 minutes and 2 dollars!

Finally, this method is applicable not only to the multi-criteria case, but to any partial order, giving us a general way to find a non-dominated path in a partially-ordered SP problem. It is important to notice here the flexibility of the semiring approach, which allows us to use the same syntax and computational engine, but on a different semiring, to compute different objects.

7.3 Modality-based SP Problems

Until now we have considered situations in which an arc is labeled by its cost, be it one element or a tuple of elements as in the multi-criteria case. Sometimes, however, it may be useful to also associate with each arc information about the *modality* to be used to traverse the arc.

For example, interpreting arcs as links between cities, we may want to model the fact that we can cover such an arc by *car*, or by *train*, or by *plane*. Another example of a modality could be the *time of the day* in which we cover the arc, like *morning*, *afternoon*, and *night*. In all these cases, the cost of an arc may depend on its modality.

An important thing to notice is that a path could be made of arcs which are not necessarily all covered with the same modality. For example, the path between two cities can be made of arcs, some of which are covered in the morning and others in the afternoon. Moreover, it can be that different arcs have different sets of modalities. For example, from city A to city B we can use both the car or the train, and from city B to city C we can only use the plane. Thus, modalities cannot be simply treated by selecting a subset of arcs (all those with the same modality).

An example of an SP problem with the three modalities representing car (c), plane (p), and train (t) can be seen in Figure 7.4. Here the problem is to find a shortest path from any node to v, and to know both its cost and also the modalities of its arcs.

This SP problem with modalities can be modeled via the SCLP program in Table 7.3. In this program, the variables represent the modalities.

If we ask the query `:-p(c).`, it means that we want to know the smallest cost for a trip from p to v using the *car*. The result of this query in our example is $p(c) = 9$ (using the path $p - r - u - v$).

Notice that the formulation shown in Figure 7.3 puts some possibly undesired constraints on the shortest path to be found. In fact, by using the same variable in all the predicates of a rule, we make sure that the same modality (in our case the same transport mean) is used throughout the whole path. If instead we want to allow different modalities in different arcs of the path, then we just need to change the rules by putting a new variable on the last predicate of each rule. For example, rule
`p(X) :-` c_{pq}`(X), q(X).`
would become
`p(X) :-` c_{pq}`(X), q(Y).`
Now we can use a modality for the arc from p to q, and another one for the next arc. In this new program, asking the query `:-p(c).` means that we want to know the smallest cost for a trip from p to v using the car in the first arc.

The same methods used in the previous sections to find a shortest path, or a non-dominated path in the case of a partial order, can be used in this kind of SCLP programs as well. Thus, we can put additional variables in the predicates to represent alternative arcs outgoing the corresponding nodes, and we can shift to the semiring containing sets of costs to find a non-dominated path. In particular, a clause like
`p(X) :-` c_{pq}`(X), q(Y).`
would be rewritten as
`p(X,a) :-` c_{pq}`(X), q(Y,Z).`

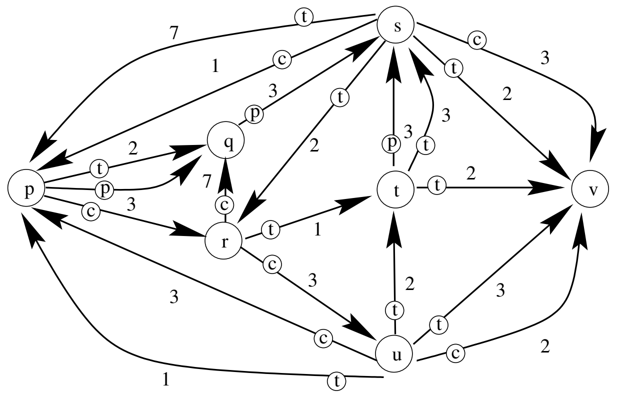

Figure 7.4: An SP problem with modalities.

p(X) :- c_{pq}(X), q(X).	c_{pq}(t) :- 2.
p(X) :- c_{pr}(X), r(X).	c_{pq}(p) :- 3.
q(X) :- c_{qs}(X), s(X).	c_{pr}(c) :- 3.
r(X) :- c_{rq}(X), q(X).	c_{qs}(p) :- 3.
r(X) :- c_{rt}(X), t(X).	c_{rq}(c) :- 7.
r(X) :- c_{ru}(X), u(X).	c_{rt}(t) :- 1.
s(X) :- c_{sp}(X), p(X).	c_{ru}(c) :- 3.
s(X) :- c_{sr}(X), r(X).	c_{sp}(c) :- 1.
s(X) :- c_{sv}(X), v(X).	c_{sp}(t) :- 7.
t(X) :- c_{ts}(X), s(X).	c_{sr}(t) :- 2.
u(X) :- c_{up}(X), p(X).	c_{sv}(t) :- 2.
u(X) :- c_{ut}(X), t(X).	c_{sv}(c) :- 3.
u(X) :- c_{uv}(X), v(X).	c_{ts}(p) :- 3.
v(X) :- 0.	c_{ts}(t) :- 3.
	c_{up}(c) :- 3.
	c_{up}(t) :- 1.
	c_{ut}(t) :- 2.
	c_{uv}(t) :- 3.
	c_{uv}(c) :- 2.

Table 7.3: The SCLP program representing the SP problem with modalities.

7.4 An SP Algorithm for a class of SCLP programs

Summarizing what we did in the previous sections, we can say that the general form of the SCLP programs we use to represent SP problems consists of several clauses for each predicate p_i, where each clause body has one constraint and one other predicate p_j, plus one special clause for the sink predicate (with a 0 in the body), plus the facts defining the costs of the arcs. Table 7.4 shows this general form (only the clauses and without variables). It is important to notice that these are not general SCLP programs, since there is always one predicate symbol in each clause body (since the constraint can be replaced by its cost). In the logic programming literature, these programs are called *linear*, because of the restriction on the number of predicates in the bodies of the clauses.

Given any SCLP program of the form shown in Table 7.4, its semantics can be obtained using classical methodologies (for the CLP literature). We now, however, will show an additional method, based on a classical algorithm for SP problems, suitably adapted to c-semirings, to obtain the same semantics. The algorithm we will use is the one developed by Floyd and Warshall [Flo62]. To put the SCLP program into a shape that can be handled by this algorithm, we first need to transform it into a system of linear equations, one for each predicate: the left hand side of the equation is the chosen predicate, while the right hand side is obtained by combining the bodies of all the clauses defining the predicates via the $+$ operator, and replacing the comma with the $\times$ operator in each body. The result of this transformation, applied to the program of Table 7.4, can be seen in Table 7.5.

From this system of equations we can now build a graph with a node i for each predicate p_i, and where the cost of the arc between nodes i and j is given by predicate c_{ij}. In this way, this system of equations can be interpreted also as the matrix showing, for each pair of nodes, the cost of the arc (if it exists) between such nodes. In fact, each column (and row) can be associated with one node.

Given this matrix, the algorithm works by selecting each triple of nodes, say i, j and k, and performing the following assignment: $c_{ij} := c_{ij} + (c_{ik} \times c_{kj})$, where the $+$ and $\times$ operators in this assignment refer to the operations of the chosen semiring.

The triples are chosen by first selecting a value for index k (which is the intermediate node in the two-step path), and then varying the indices i and j in all possible ways (to consider all two-step paths from i to j). After considering all values for k, that is, after n^3 steps, the value of the element in row i and column v represents the cost of the shortest path from node i to v. In terms of SCLP semantics, this is the semantics of predicate p_i.

This method to obtain the semantics of an SCLP program of the form in Table 7.4 can be used independently of the semiring underlying the given program. In fact, as noted above, the version we use of the original Floyd-Warshall algorithm does not depend on the meaning of the $+$ and $\times$ operators. In particular, it can also be used for partially ordered semirings. If the semiring we consider is totally ordered, however, we can also use any other classical algorithm for solving SP problems.

$$
\begin{array}{l}
p_1 :-\ c_{12}, p_2. \\
\vdots \\
p_1 :-\ c_{1n}, p_n. \\
p_2 :-\ c_{12}, p_1. \\
p_2 :-\ c_{13}, p_3. \\
\vdots \\
p_2 :-\ c_{2n}, p_n. \\
\vdots \\
p_v :-\ 0. \\
\vdots \\
p_n :-\ c_{n1}, p_1. \\
\vdots \\
p_n :-\ c_{n,n-1}, p_{n-1}.
\end{array}
$$

Table 7.4: The general form of an SCLP program representing an SP problem.

$$
\begin{array}{llllll}
p_1 & = & & c_{12} \times p_2 + \ \ldots & & + c_{1n} \times p_n \\
p_2 & = c_{21} \times p_1 \ + & & c_{23} \times p_3 + \ \ldots & & + c_{2n} \times p_n \\
\vdots & & & & & \\
p_v & = & & & & \qquad \mathbf{0} \\
\vdots & & & & & \\
p_n & = c_{n1} \times p_n + \ \ldots & & + c_{nn-1} \times p_{n-1} & &
\end{array}
$$

Table 7.5: The system of equations corresponding to the SCLP program form of Table 7.4.

7.5 Best-Tree Search in AND-OR Graphs

In this section we extend our reasoning to Best-Tree Search over $AND - OR$ graphs. The main idea underlying this extension concerns the use of non-linear clauses in SCLP, that is, clauses which have more than one atom in their body. In fact, in this way we can represent trees instead of paths.

7.5.1 AND-OR Graphs and Best Solution Trees

An AND-OR graph is defined essentially as a hypergraph. Namely, instead of arcs connecting pairs of nodes there are hyperarcs connecting n-tuple of nodes ($n = 1, 2, 3, \ldots$). Hyperarcs are called *connectors* and they must be considered as directed from their first node to all others. Formally an AND-OR graph is a pair $G = (N, C)$, where N is a set of *nodes* and C is a set of connectors

$$C \subseteq N \times \bigcup_{i=1}^{k} N^i.$$

Each k-connector $(n_{i_0}, n_{i_1}, \ldots, n_{i_k})$ is an ordered $(k+1)$-tuple, where n_{i_0} is the *input* node and $n_{i_1}, \ldots, n_{i_k}$ are the *output* nodes. We say that n_{i_0} is the *predecessor* of $n_{i_1}, \ldots, n_{i_k}$ and these nodes are the *successors* of n_{i_0}. Note that when $C \subseteq N^2$ we have a usual graph whose arcs are the 1-connectors. Note that there are also 0-connectors, i.e., connectors with one input and no output node.

In Figure 7.5 we give an example of an AND-OR graph. The nodes are $n_0, \ldots, n_8$. The 0-connectors are represented as a line ending with a square, whereas k-connectors ($k \geq 0$) are represented as k directed lines connected together. For instance, (n_0, n_1) and (n_0, n_5, n_4) are the 1-connector and 2-connector, respectively, with input node n_0.

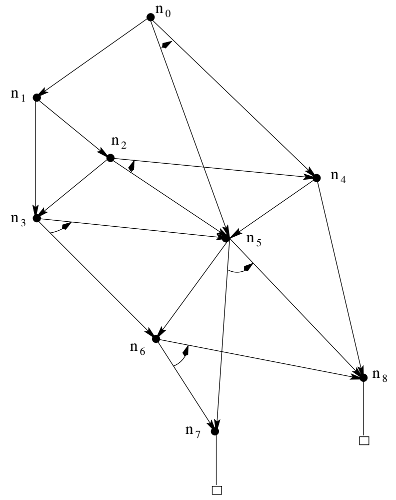

Figure 7.5: An example of an AND-OR graph.

An AND tree is a special case of an AND-OR graph, where every node appears exactly twice in the set of connectors C, once as an input node of some connector, and once as an output node of some other connector. The only exception is a node called *root* which appears only once, as an input node of some connector. The *leaves* of an AND tree are those nodes which are input nodes of a 0-connector. An example of an AND tree is given in Figure 7.6. Here n'_7, n'_8, n''_7 and n''_8 are leaves.

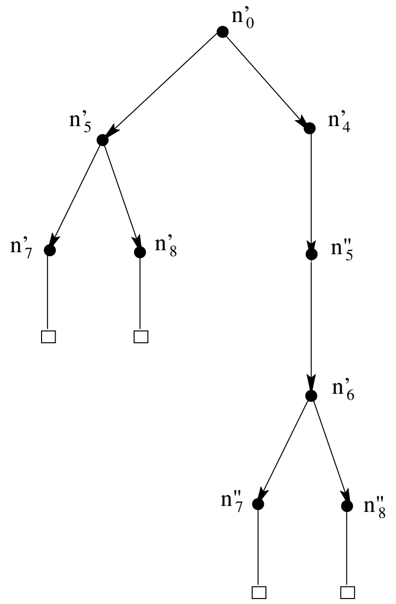

Figure 7.6: An example of an AND tree.

Given an AND-OR graph G, an AND tree H is a *solution tree of* G *with start node* n_0, if there is a function g mapping nodes of H into nodes of G such that:

- the root of H is mapped in n_0; and
- if $(n_{i_0}, n_{i_1}, \ldots, n_{i_k})$ is a connector of H, then $(g(n_{i_0}), g(n_{i_1}), \ldots, g(n_{i_k}))$ is a connector of G.

Informally, a solution tree of an AND-OR graph is analogous to a path of an ordinary graph. It can be obtained by selecting exactly one outgoing connector for each node. For instance, the AND tree in Figure 7.6 is a solution tree of the graph in Figure 7.5 with start node n_0, if we let $g(n'_0) = n_0, g(n'_5) = n_5, g(n'_7) = n_7, g(n'_8) = n_8, g(n'_4) = n_4, g(n''_5) = n_5, g(n''_6) = n_6, g(n''_7) = n_7, g(n''_8) = n_8$. Note that distinct nodes of the tree can be mapped into the same node of the graph.

A *functionally weighted* AND-OR *graph* is an AND-OR graph with a k-adic function over the reals (*cost function*) associated with each k-connector. In particular, a constant is associated with each 0-connector. It is easy to see that if the functionally weighted graph is an AND tree H, a *cost* can be given to it, just evaluating the functions associated with the connectors. Recursively, to every subtree of H with root in node n_{i_0} a cost c_{i_0} is given as follows:

- If n_{i_0} is a leaf, then its cost is the associated constant.
- If n_{i_0} is the input node of a connector $(n_{i_0}, n_{i_1}, \ldots, n_{i_k})$, then its cost is $c_{i_0} = f_r(c_{i_1}, \ldots, c_{i_k})$ where f_r is the function cost associated with the connector, and $c_{i_1}, \ldots, c_{i_k}$ are the costs of the subtrees rooted at nodes $n_{i_1}, \ldots, n_{i_k}$.

The general optimization problem can be stated as follows: *given a functionally weighted* AND-OR *graph, find a minimal cost solution tree with start node* n_0. In general, the function used to assign a value to the input node n_{i_0} of a k-connector $(n_{i_0}, n_{i_1}, \ldots, n_{i_k})$ is of the form $f_r(c_{i_1}, \ldots, c_{i_k}) = a_r + c_{i_1} + \ldots + c_{i_k}$ where a_r is a constant associated to the connector and $c_{i_1}, \ldots, c_{i_k}$ are the costs of the subtrees rooted at nodes $n_{i_1}, \ldots, n_{i_k}$.

In the following we will show that this cost function is only an instantiation of a more general one based on the notion of *c-semiring*.

7.5.2 AND-OR graphs using SCLP

In this section we will show how to represent an AND-OR graph as an SCLP program over a specific c-semiring. As soon as we have described the AND-OR graph as an SCLP program, the problem of finding the *best* tree in the graph rooted at n_i becomes the problem of finding the *best* refutation for the goal n_i, and this is *exactly* what the semantic of the SCLP language does.

To represent the classical problem where the meaning of *best* tree is the tree whose *sum* of the costs of its connectors is *minimum*, we consider an SCLP program over the semiring $S = \langle N, \min, +, +\infty, 0\rangle$, which, as noted above, is an appropriated framework to represent constraint problems where one wants to minimize the sum of the costs of the solutions.

Consider for example the AND-OR graph problem represented in Figure 7.7: each connector $(n_{i_1}, \ldots, n_{i_k})$ has associated a label $f_{i_1,\ldots,i_k}$ representing its cost.

From this AND-OR graph problem we can build an SCLP program as follows. For each connector we have two clauses: one describes the connector and the other one its cost. More precisely, the head of the first clause represents the starting node n_{i_0}, and its body contains both the final nodes and a predicate, say $f_{i_0,\ldots,i_k}$, representing the cost of the connector from node n_{i_0} to nodes $n_{i_1}, \ldots, n_{i_k}$. Then, the second clause is a fact associating to predicate $f_{i_0,\ldots,i_k}$ its cost (which is a semiring element).

For example, if we consider the connector (n_0, n_1, n_2, n_3) with label f_0 in Figure 7.8, this can be represented by the two clauses

```
n0 :- f0, n1, n2, n3.
f0 :- c0.
```

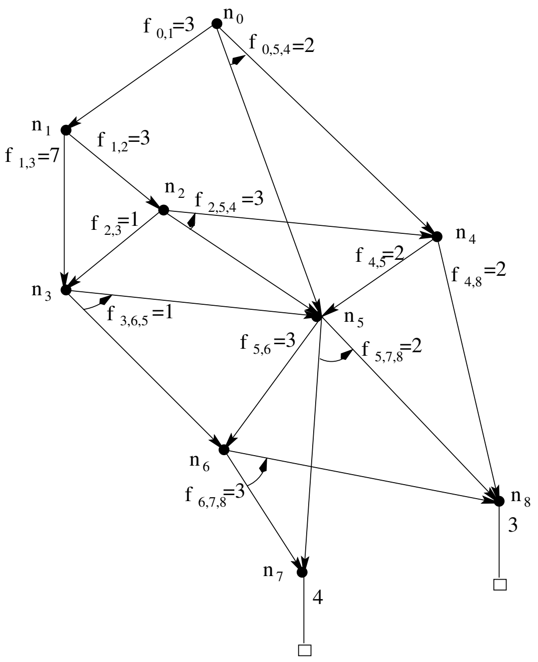

Figure 7.7: A weighted AND-OR graph problem.

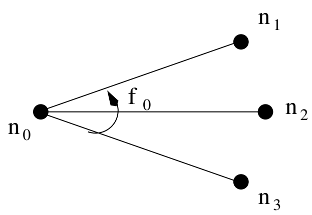

Figure 7.8: A typical connector.

n_0 :- $f_{0,5,4}$, n_5, n_4.	$f_{0,5,4}$:- 2.
n_0 :- $f_{0,1}$, n_1.	$f_{0,1}$:- 3.
n_1 :- $f_{1,2}$, n_2.	$f_{1,2}$:- 3.
n_1 :- $f_{1,3}$, n_3.	$f_{1,3}$:- 7.
n_2 :- $f_{2,3}$, n_3.	$f_{2,3}$:- 1.
n_2 :- $f_{2,5,4}$, n_5, n_4.	$f_{2,5,4}$:- 3.
n_3 :- $f_{3,6,5}$, n_6, N_5.	$f_{3,6,5}$:- 1.
n_4 :- $f_{4,5}$, n_5.	$f_{4,5}$:- 2.
n_4 :- $f_{4,8}$, n_8.	$f_{4,8}$:- 2.
n_5 :- $f_{5,6}$, n_6.	$f_{5,6}$:- 3.
n_5 :- $f_{5,7,8}$, n_7, n_8.	$f_{5,7,8}$:- 2.
n_6 :- $f_{6,7,8}$, n_7, n_8.	$f_{6,7,8}$:- 3.
n_7 :- 4.	n_8 :- 3.

Table 7.6: The SCLP program representing the AND-OR graph problem in Figure 7.7.

The whole program corresponding to the AND-OR graph problem in Figure 7.7 can be seen in Table 7.5.2.

To solve the AND-OR graph problem it is enough to perform a query in the SCLP framework; for example, if we want to compute the cost of the best tree rooted at n_2 we have to perform the query n_2. The operational semantics machinery finds all trees (modulo some cuts due to heuristics) and then combines all the solutions via the additive operation of the semiring, which in this case is min. For this query, we obtain the value 17 that represents the cost of the best tree rooted at n_2.

Notice that to represent classical best tree problems in SCLP, we do not need any variable. Thus the resulting program is propositional. However, this program, while giving us the cost of the best tree, does not give us any information about the connectors which form such a tree. This information could be obtained by providing each predicate with an argument which represents the connector chosen at each step, as we did for shortest path problems in the previous section.

The best tree for the program in Table 7.5.2, with sum of costs 17 as given by the semantics of the program, is represented in Figure 7.9.

In the examples given so far, we were interested in finding the tree with the minimum cost. Thus, the constant associated to each connector, that is, its cost, is in $\Re$ and the costs of the subtrees rooted at the sons of a certain node are *combined* together using the $+$ operator. More generally, the operation used to combine the value of the subtrees can be mapped over the c-semiring structure. More precisely, to describe the operation of combination we will use the $\times$ operator of the semiring; the additive operator will be useful instead to compare different trees, since the partial order $\leq_S$ is induced by $+$ operation of the semiring.

Notice that we can apply to best tree search all the extensions already defined in the

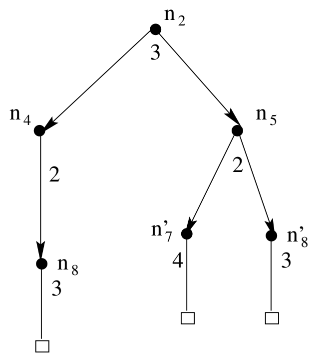

Figure 7.9: The best tree corresponding to the program in Table 7.5.2.

previous sections for shortest paths. This means that we can use also a partially ordered set for the levels of preference of the connectors. This can be useful for multi-criteria best tree problems or for general partially ordered best tree problems. We can also use several classes of connectors and then search for a best tree where the connectors are all in the same class, or also of different classes (modality-based best tree problems).

7.6 Conclusions

In this chapter the shortest path problem taken from the operational research field has been considered and translated in a declarative fashion using the SCLP framework. This, on the one hand, enriches a class of SCLP programs with a semantics given using an ”algorithm”, and on the other hand gives a resolution methodology to the shortest path problems, based upon the semantics of a declarative language.

Moreover, the encoding of the shortest path in a soft constraint framework also provides for the possibility to apply all the methodologies of partial solving described in the previous chapters (that is abstraction, partial local consistency, and so on) when performing the constraint consistency steps of the resolution process.

Chapter 8

SCSPs for Modelling Attacks to Security Protocols

Abstract

Security protocols are executed over an *insecure* network by a finite but unbounded population of agents. We show that any network configuration arising from the execution of a specific protocol can be modeled as a soft constraint satisfaction problem (SCSP). In particular, the *policy SCSP* models the network configuration where all admissible protocol sessions have terminated successfully. A network configuration suffers an attack if the solution of the corresponding SCSP is *lower* than that for the policy SCSP. Different attacks bear more or less drastic consequences, hence soft constraints are necessary. For the sake of demonstration, we model the asymmetric Needham-Schroeder protocol, and the attack discovered by Lowe on that protocol.

Modern computer networks are *insecure* in the sense that the ongoing traffic can be intercepted and eavesdropped on by a malicious attacker, called *spy* below. Agents trying to communicate on an insecure network execute suitable *security protocols* to take advantage of the protocol *goals*. A typical goal is *confidentiality*, which holds for a message that remains undisclosed to the spy. Failure to achieve the claimed goals of a protocol [AN96, Low96, LR97] has motivated a number of approaches for reasoning formally on security protocols (e.g. [Low95, BR97a, Pau98, Bel99]).

Our original contribution to the field is an approach to specify each network configuration arising from the execution of a protocol as a soft constraint satisfaction problem (SCSP), and to detect attacks mounted by the spy. Recall that an SCSP may be viewed as a classical constraint satisfaction problem (CSP) [Mac92, Wal96] where each assignment of values to variables in the constraints is associated with an element taken from a partially ordered set. These elements can then be interpreted as levels of preference, or of certainty, etc.

When modelling security protocols, the partially ordered set contains *security levels*. For example, the Kerberos protocol [BP98] relies on two separate sets of session keys:

one contains the "authorisation keys", the other contains the "service keys". Each authorisation key *authK* is used to encrypt several service keys $servK_i$, hence, if confidentiality is violated on *authK*, it will also be violated on all the $servK_i$. Clearly, different levels of security must be associated with different violations of confidentiality. We demonstrate our approach on the *asymmetric* Needham-Schroeder protocol [NS78], which is one of the most discussed security protocols. Despite its apparent simplicity, the protocol hides a subtle but realistic attack, which was discovered by Lowe [Low95] nearly two decades after the protocol publication.

We assume a basic familiarity with the concepts of encryption and decryption, which are extensively presented elsewhere [Nat77, RSA76]. We indicate encryption by fat braces, so that $\{m\}_K$ represents the ciphertext obtained by encrypting message m under key K. We say that encryption is **perfect** when m can be recovered from $\{m\}_K$ if and only if K^{-1} is available. In this definition, the keys K and K^{-1} are interchangeable. A major feature is that K cannot be obtained from K^{-1} or vice versa. While most cryptographic algorithms are perfect at the specification level, several of their implementations no longer enjoy that desirable property. Designing a protocol that achieves the goals for which it was conceived is not trivial even if perfect cryptography is available, as confirmed by Lowe's attack.

We consider the formalizations of SCSPs based on semirings described in chapter 2, where the semiring specifies the partially ordered set and the appropriate operation that must be used to combine constraints together. Below, we describe how to use them to model attacks to security protocols (§8.1). Then, we describe the asymmetric Needham-Schroeder protocol (§8.2), use our approach to model that protocol (§8.3), and conclude (§8.4).

This chapter is based on some preliminary work presented in [BB00, BB01].

8.1 Using SCSPs for Modelling Protocol Attacks

The intuition underlying the machinery developed below is to express protocol attacks as properties of the solutions for appropriate soft constraint problems. Verifying whether there is an attack on a given network configuration arising from the execution of a protocol will reduce to solving a constraint problem and analyzing its solution. Given a security protocol, we aim at modelling each possible network configuration as a constraint problem over a specific constraint system — hence, a semiring is required. The following, general treatment is demonstrated in §8.3.

8.1.1 The Security Semiring

We define the set L as to containing $unknown$, $private$, $debug$ and $public$. Each of these elements represents a possible *security level* that a protocol associates to messages. The security levels regulate the agents' knowledge of messages. They may be seen as an extension of the two-valued property of known/unknown. For example, $unknown$ will be assigned to those messages that a protocol does not encompass, and obviously to those messages that a given agent does not know; $\mathtt{debug}$ will be assigned to the messages that are created and then exchanged during a protocol session. The remaining levels are self-explanatory.

Figure 8.1 defines a multiplicative operator, $\times_{sec}$, and an additive one, $+_{sec}$. The following theorem introduces the *security semiring*.

$\times_{sec}$	unknown	private	debug	public
unknown	unknown	private	debug	public
private	private	private	debug	public
debug	debug	debug	debug	public
public	public	public	public	public

$+_{sec}$	unknown	private	debug	public
unknown	unknown	unknown	unknown	unknown
private	unknown	private	private	private
debug	unknown	private	debug	debug
public	unknown	private	debug	public

Figure 8.1: The definitions of $\times_{sec}$ and $+_{sec}$.

THEOREM 8.1.1 (SECURITY SEMIRING)
The structure $\mathcal{S}_{sec} = \langle L, +_{sec}, \times_{sec}$,public,unknown$\rangle$ is a semiring.

PROOF. It is easy to see that there is an isomorphism between the semiring $\mathcal{S}_{sec}$ and the fuzzy one: it is enough to change the security levels with the values in $[0, 1]$ and the semiring operations with max and min. □

Since $\times_{sec}$ is idempotent, it is also the $\mathtt{glb}$ operator in the total order of the carrier set L of the secrecy levels. While the current four levels will suffice to model most protocols, it is understood that more complex protocols may require additional ones, such as $\mathtt{notice}$, $\mathtt{warning}$, $\mathtt{error}$, $\mathtt{crit}$. The security semiring can be easily extended by upgrading the definitions of $\times_{sec}$ and $+_{sec}$.

8.1.2 The Network Constraint System

A computer network can be modeled as a constraint problem over a constraint system $CS_n = \langle \mathcal{S}, \mathcal{D}, \mathcal{V} \rangle$ defined as follows:

- $\mathcal{S}$ is the security semiring (§8.1.1);
- $\mathcal{V}$ is a finite but unbounded set of variables, including *SPY*, each representing a network agent;

- $\mathcal{D}$ is a finite but unbounded set including all atomic messages, such as agent names, timestamps, nonces and cryptographic keys, as well as all messages recursively obtained by concatenation and encryption. Note that, for $\mathcal{D}$ to be finite, the number of concatenation and encryption operations that may be applied must be finite. Intuitively, $\mathcal{D}$ represents all agents' possible knowledge.

We name CS_n as *network constraint system*. Note that CS_n does not depend on any protocols, for it merely portrays the topology of a computer network on which any protocol can be implemented. Members of $\mathcal{V}$ will be indicated by capital letters, while members of $\mathcal{D}$ will be in small letters.

8.1.3 The Initial SCSP

Each security protocol $\mathcal{P}$ is associated with a policy that should, at least, state

- which message forms are admissible (e.g. the maximum number of components per compound message);
- which messages are public;
- which messages are private for which agents;
- which messages can be learnt from the traffic, and by which agents;
- which messages must be created by which agents during the execution of $\mathcal{P}$.

It is intuitive to capture these policy rules by our security levels (§8.1.1). Precisely, these rules can be translated into unary and/or binary constraints. For each agent $A \in \mathcal{V}$, we define a unary constraint that states the security levels of A's knowledge as follows. It associates security level $\mathtt{public}$ to all agent names and to timestamps (if $\mathcal{P}$ uses them); level $\mathtt{private}$ to A's keys (namely to *Ka* if $\mathcal{P}$ uses symmetric encryption, or to Ka^{-1} if it uses asymmetric encryption), and to all messages that A creates during the execution of the protocol with any other agent; level $\mathtt{debug}$ to all messages that A either sends or receives from the traffic during the protocol execution; level $\mathtt{unknown}$ to other agents' keys, to all messages A must never learn from the traffic, and to all *remaining* messages. The remaining messages, whose form is admissible by the policy and whose elements are drawn from the domain $\mathcal{D}$, are a finite number. Clearly, we do not need to consider messages whose structure is not admissible by the policy.

This procedure defines what we name *initial SCSP for* $\mathcal{P}$, which portrays a network where $\mathcal{P}$ can be executed but none of its sessions has yet started. The initial SCSP for $\mathcal{P}$ also highlights the agents' initial knowledge, which typically consists of atomic messages such as agent names, timestamps, some keys and some nonces.

Considerations on how official protocol specifications often fail to provide a satisfactory policy [BMPT00] exceed the scope of this chapter. Nevertheless, having to define the initial SCSP for a protocol may pinpoint unknown deficiencies or ambiguity in the policy.

8.1.4 The Policy SCSP

The policy for a protocol $\mathcal{P}$ also establishes which messages must be exchanged during a session between a given pair of agents.

We extend the initial SCSP for $\mathcal{P}$ (§8.1.3) with further constraints. Each step of a session of $\mathcal{P}$ can be translated into two constraints, one of which is unary and the other one is binary. The unary constraint corresponds to the creation of the message mentioned by the given step, while the binary constraint corresponds to the shipment and receipt of the message. Precisely, consider the protocol step whereby A creates a message m and sends it to B. We add a unary constraint on variable A that assigns security level $private$ to m, and level $unknown$ to all other possible domain values for A; we also add a binary constraint between A and B that assigns security level $debug$ to the tuple $\langle\{\}, m\rangle$, and level $unknown$ to all other possible tuples.

We obtain the *policy SCSP for* $\mathcal{P}$, which formalises a network where each agent has successfully terminated a protocol session with every other agent.

8.1.5 The Imputable SCSP

A network history induced by a protocol $\mathcal{P}$ may be viewed as a sequence of message shipments and message receipts. In the real world, however, not all messages that are sent are then received by the intended recipient or received at all. This is due to the malicious activity of the spy. Nevertheless, we can model the configuration of the network at a certain point in a possible history as an SCSP. To do so, we constrain the initial SCSP for $\mathcal{P}$ (§8.1.3) additionally, as stated by the following rule.

> If an agent A sends a message m to an agent B, and an agent c receives m, then add a binary constraint between A and c that assigns security level $debug$ to the tuple $\langle\{\}, m\rangle$, and level $unknown$ to all other possible tuples.

Note that C could either be B or the spy. In particular, should the spy, in the latter case, also deliver m to B, the rule quoted above would apply to the triple *SPY*, m and B. Or, should the spy tamper with m obtaining m', and then deliver it to another agent D, then the rule would apply to the triple *SPY*, m', and D.

We name the originated soft constraint problem as *imputable SCSP* for the network of agents running protocol $\mathcal{P}$. There exist a finite number of imputable SCSPs for $\mathcal{P}$ because the number of agents, the number of protocol steps, and the length of a network history up to a certain point are all finite values.

8.1.6 Agent knowledge as an enhanced *Security Entailment*

The agents' knowledge increases dynamically while they are running a protocol. This can be represented using a language like *cc* [Sar89] suitably extended to deal with soft constraints (that will be a future work of this thesis). The extension relies on the notion of α-consistency, rather than on the notion of consistency/inconsistency. Moreover, the entailment relation $\vdash$ [Sco82], which captures from the store the constraints implied but not explicitly added by the agents, must be modified to deal with the semiring values.

The entailment relation on the *network constraint system* models directly the possible operations — message encryption, decryption, concatenation and splitting — that the agents may perform on the messages that they see. The agents' knowledge increases by means of these operations.

The definition of the entailment relation must be augmented by the following rules, where the *def* function is associated with a generic constraint over a generic agent A.

Encryption

$$\begin{aligned}&\mathit{def}(m_1) = v_1,\ \mathit{def}(m_2) = v_2,\ \mathit{def}(\{m_1\}_{m_2}) = v\\ &\quad\vdash \mathit{def}(\{m_1\}_{m_2}) = (v \times_{sec} (v_1 +_{sec} v_2))\end{aligned}$$

If A knows two messages m_1 and m_2, then she/he also knows the encrypted message $\{m_1\}_{m_2}$. Precisely, if m_1 and m_2 respectively have security level v_1 and v_2, then encryption bears level $(v_1 +_{sec} v_2)$; if A already knows $\{m_1\}_{m_2}$ with security level v, then this must be accounted for by means of $\times_{sec}$.

Decryption

$$\begin{aligned}&\mathit{def}(m_1) = v_1,\ \mathit{def}(m_2) = v_2,\ \mathit{def}(\{m_1\}_{m_2}) = v\\ &\quad\vdash \mathit{def}(m_1) = (v_1 \times_{sec} v),\ \mathit{def}(m_2) = (v_2 \times_{sec} v)\end{aligned}$$

An agent A can decrypt a ciphertext and obtain its body, whose security level comes from the application of $\times_{sec}$ to the levels of the ciphertext and of the encryption key.

Concatenation

$$\begin{aligned}&\mathit{def}(\{m_1, m_2\}) = v,\ \mathit{def}(m_1) = v_1,\ \mathit{def}(m_2) = v_2\\ &\quad\vdash \mathit{def}(\{m_1, m_2\}) = ((v_1 +_{sec} v_2) \times_{sec} v)\end{aligned}$$

The security level for compound messages is computed as is that for encrypted messages.

Splitting

$$\mathit{def}(m_1) = v_1,\ \mathit{def}(\{m_1, m_2\}) = v \quad \vdash \mathit{def}(m_1) = (v_1 \times_{sec} v)$$

Once a component is extracted from a compound message, its security level must be normalised with that already owned by the agent.

Recall that, in the initial SCSP for any protocol, all domain values for all agents are associated with security level $\mathit{unknown}$, which is the unit element of $\times_{sec}$.

8.1.7 Modelling Protocol Attacks

Let us consider a security protocol $\mathcal{P}$. Our aim is to decide whether a given configuration of the network that arises when some agents execute some sessions of $\mathcal{P}$ hides an attack.

First, we build the policy SCSP for $\mathcal{P}$ and the imputable SCSP for the given configuration. Then, we calculate the solutions of both SCSPs. Finally, if the solution of the imputable SCSP is *lower* (that is close to security level $\mathtt{public}$) than that for the policy SCSP, then and only then does the given configuration have an attack.

In particular, the solution of the imputable SCSP being lower (than that for the policy SCSP) signifies that some messages have a lower security level. The attack can be understood by focusing on those messages.

It is straightforward to prove that the above procedure for modeling attacks is sound and complete. Indeed, all attacks rely on the spy's intercepting and learning certain messages, which lowers the security levels of the spy for those messages (typically from $\mathtt{unknown}$ to $\mathtt{debug}$, §8.1.6). Vice versa, if this happens, then the spy has got hold of messages that the policy of the protocol does not allow her/him to. Hence, there is an attack.

8.2 The Asymmetric Needham-Schroeder Protocol

The "asymmetric" attribute signifies that the protocol makes use of an asymmetric cryptographic algorithm (e.g. RSA [RSA76]) rather than a symmetric one (e.g. DES [Nat77]). Each agent A is endowed with a *public* key Ka, which is known to all, and a *private* key Ka^{-1}, which should be known only to A. Note that limiting the knowledge of Ka^{-1} only to A is an assumption explicitly required by the protocol, rather than a property that is enforced.

The protocol version that is presented in figure 8.2 omits the initial phase where a trusted certification authority makes it public that the bit sequences Ka and Kb, are respectively the public keys for agents A and B. That phase is irrelevant to the following discussion. Recall that a *nonce* is a "number that is used only once" [NS78]. The protocol assumes that agents can invent *truly-random* nonces, so that, given a nonce N invented by an agent P, the probability that agents other than P guess N is negligible.

1. $\mathsf{A} \rightarrow \mathsf{B} : \{Na, \mathsf{A}\}_{Kb}$
2. $\mathsf{B} \rightarrow \mathsf{A} : \{Na, Nb\}_{Ka}$
3. $\mathsf{A} \rightarrow \mathsf{B} : \{Nb\}_{Kb}$

Figure 8.2: The asymmetric Needham-Schroeder protocol.

The first step sees an *initiator* A initiate the protocol with a *responder* B. A issues a nonce Na and encrypts it along with her identity under B's public key. Upon reception of that message, B decrypts it and extracts A's nonce. Then, he issues a nonce *Nb* and encrypts it along with *Na* under A's public key. When A receives message 2, she extract *Nb* and sends it back to B, encrypted under his public key.

The goal of the protocol is *authentication*: at completion of a protocol session initiated by A with B, A should get evidence showing she communicated with B and, likewise, B should get evidence that he communicated with A. An informal analysis would probably confirm that the goal is met. Upon reception of *Na* inside message 2, A would conclude that she is interacting with B, the only agent who could retrieve *Na* from message 1, since *Na* is a truly-random nonce and encryption is perfect. In the same fashion, when B receives *Nb* inside message 3, he would conclude that A was at the other end of the network because *Nb* must have been obtained from message 2, and no-one but A could perform this operation. How rigorous is this reasoning?

8.2.1 Lowe's attack to the Needham-Schroeder Protocol

Lowe shows that an informal reasoning of the sort presented above is flawed [Low95] because of the scenario depicted in figure 8.3.

1. $\mathsf{A} \rightarrow \mathsf{C} : \{Na, \mathsf{A}\}_{Kc}$

1′. $\mathsf{C} \rightarrow \mathsf{B} : \{Na, \mathsf{A}\}_{Kb}$

2′. $\mathsf{B} \rightarrow \mathsf{A} : \{Na, Nb\}_{Ka}$

2. $\mathsf{C} \rightarrow \mathsf{A} : \{Na, Nb\}_{Ka}$

3. $\mathsf{A} \rightarrow \mathsf{C} : \{Nb\}_{Kc}$

3′. $\mathsf{C} \rightarrow \mathsf{B} : \{Nb\}_{Kb}$

Figure 8.3: Lowe's attack to the Needham-Schroeder Protocol.

Security protocols are implemented as distributed concurrent programs. A malicious agent C can interleave two of the protocol sessions, provided that some agent A initiates one session with C. Note that C could be a registered user of the network, so no-one could suspect his tampering. Since A initiates with C, she encrypts her nonce and her identity under C's public key. Once these data have been obtained, C initiates another session (indicated by the primes) with another agent B, quoting A's data rather than his own. From this message, B deduces that A is trying to communicate with him. Therefore, B replies to A, quoting her nonce and his own, *Nb*. Since the entire network is under C's control, C intercepts this message before it is delivered to A but cannot decrypt it because encryption is perfect. So, C forwards it to A. The message is of the form that A was expecting, hence A extracts *Nb* and sends it to the agent with whom she had initiated the first session, C. This reveals *Nb* to C, who can use it to complete the session with B by issuing message 3', which is of the form that B was expecting.

As a result, B believes to have communicated with A, while A was communicating with C. In other words, C impersonates A with B: the protocol has failed to achieve authentication. The consequences of the attack are that B will consider every future messages quoting *Nb* as coming from A. If B is a bank and A and C are account holders, C could ask B to transfer money from A's account to C's without A is realising it.

Lowe proposes to quote B's identity in message 2, so that A would address message 3 to the agent mentioned in that message rather than to the one with whom she initiated the session. This upgrade would prevent the attack in figure 8.3 because message 3 would be encrypted under B's public key, so C would not discover *Nb* and could not issue message $3'$ to complete the session with B.

8.3 Modelling Needham-Schroeder

We demonstrate the machinery developed above on the Needham-Schroeder protocol. It is understood that all SCSPs presented below are in fact fragments pertaining to two or three variables.

Figure 8.4 contains the initial SCSP for our protocol. It may be seen that each agent's knowledge contains public information (agent names and their public keys), private information (the specific agent's private key). All other domain elements that the protocol does not use are set to the more secure level $unknown$ (and are omitted in the figure).

Now, we build the imputable SCSP corresponding to the network on which only a protocol session between A and B has initiated and terminated (Figure 8.5). The general procedure (§8.1.3) must be followed: we extend the initial SCSP with the unary constraints that assign security level $private$ to the nonces invented at run-time; we add the binary constraints that set all exchanged messages to $debug$, and all *remaining* messages to $unknown$.

Intuitively, the more steps the protocol has, the more constraints we put in the store, while the bigger is the probability to introduce insecurity (in this case inconsistency) in the network (respectively in the SCSP).

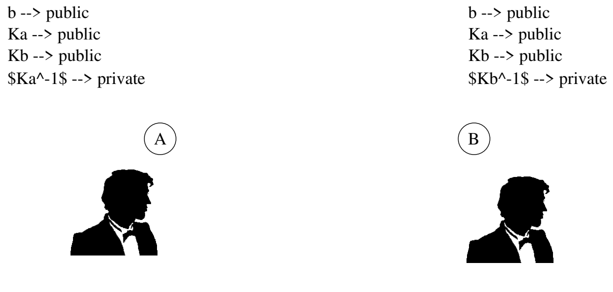

Figure 8.4: Initial SCSP for Needham-Schroeder.

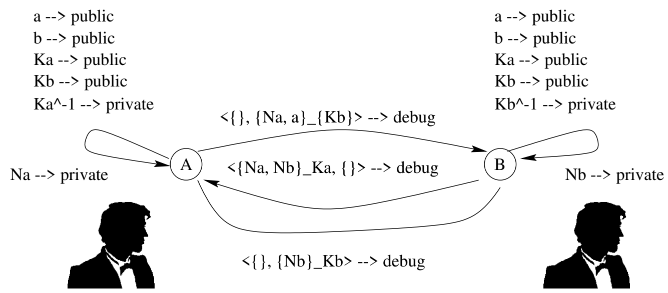

Figure 8.5: Imputable SCSP for completion of protocol session between A and B.

We also build the *policy SCSP* (omitted here), and observe that the solution of our imputable SCSP is not lower than that for the policy SCSP. Hence, we can formally confirm that the completion of a session of Needham-Schroeder between two agents bears no attacks.

8.3.1 The attack

Figure 8.6 shows the imputable SCSP corresponding to Lowe's attack: an agent C is acting as *middle person* between A and B.

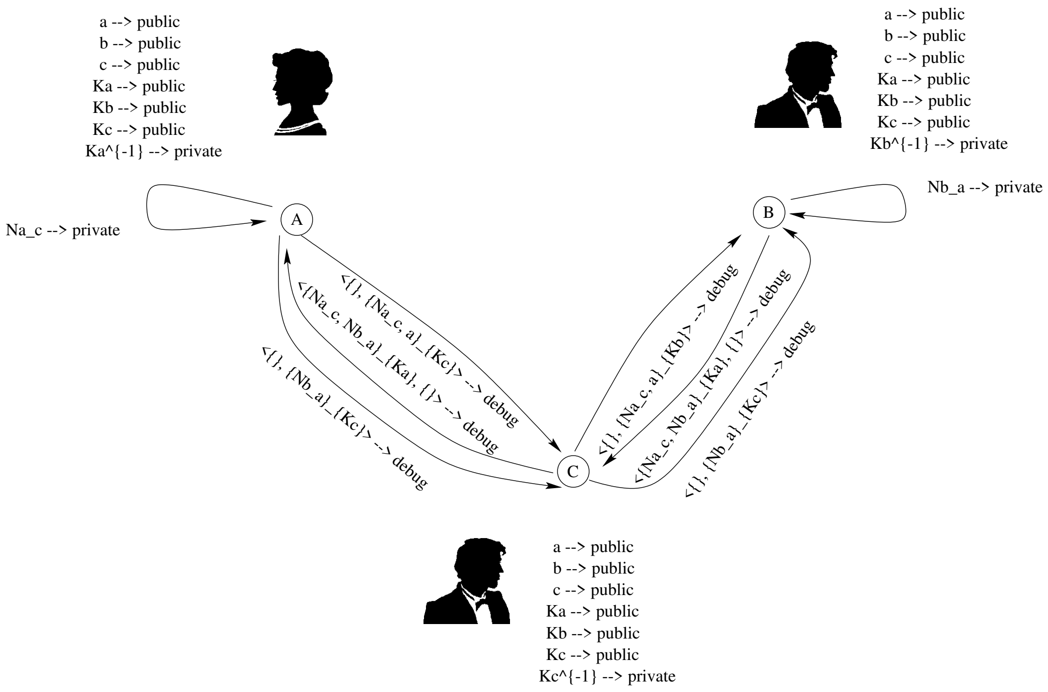

Figure 8.6: Imputable SCSP corresponding to Lowe's attack.

A larger fragment of the initial SCSP is in Figure 8.7, while a suitable fragment of the policy SCSP is given in in Figure 8.8.

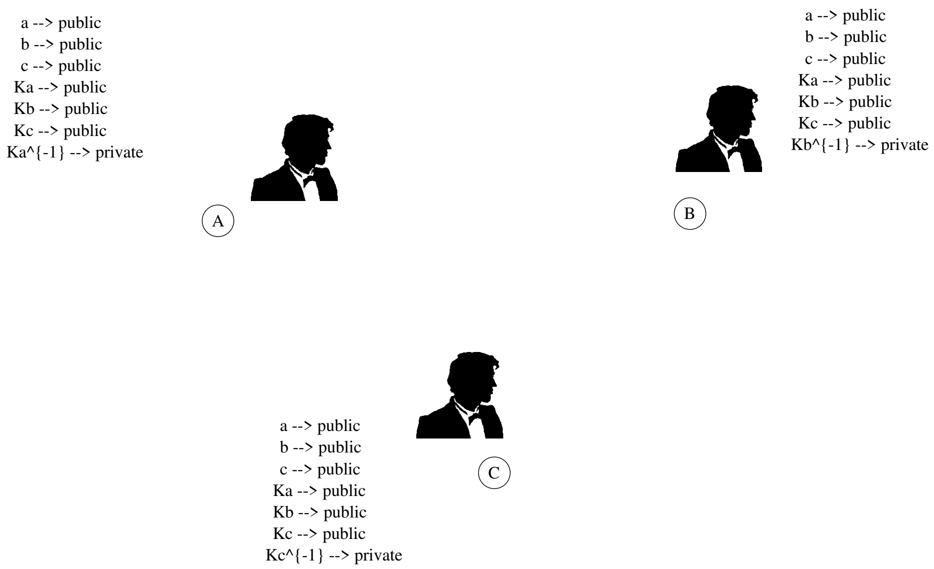

Figure 8.7: Initial SCSP for the Needham-Schroeder protocol.

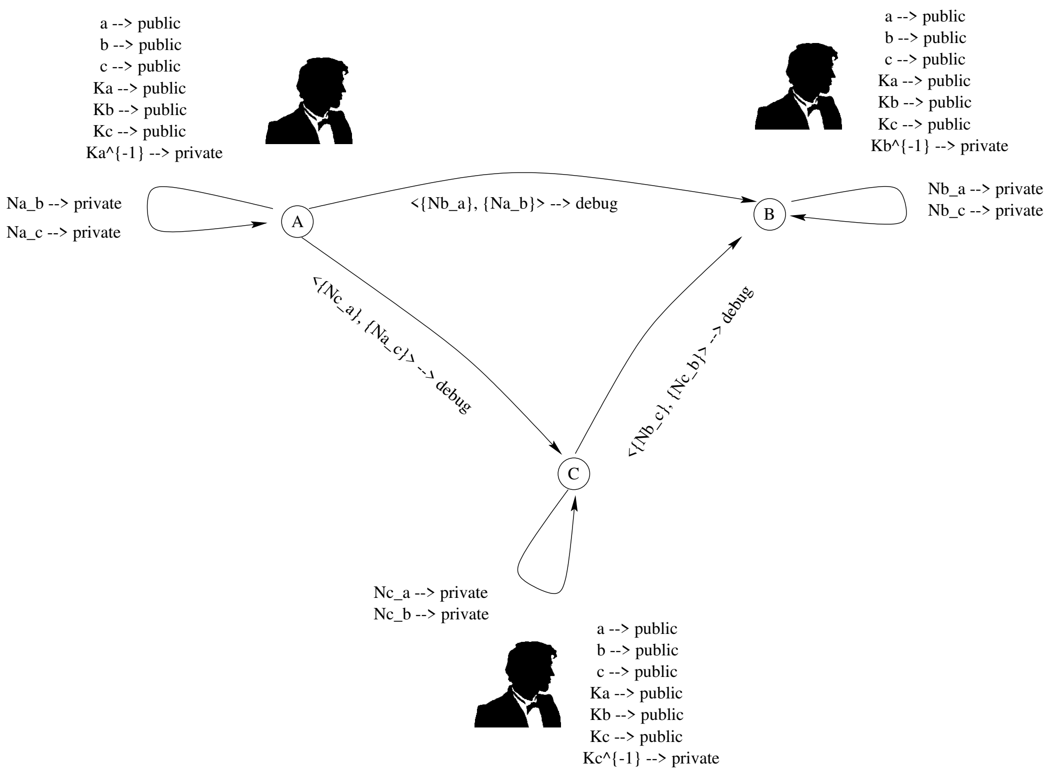

Figure 8.8: Policy SCSP for the Needham-Schroeder protocol.

It is easy to verify that the solution of our SCSP is worse than that for the policy SCSP. By restricting our attention to the security levels assigned to the domain values of the variable C, we observe that they differ. While the domain value Nb_a has level $\mathtt{unknown}$ in the policy SCSP, it has level $\mathtt{debug}$ in the imputable SCSP, which is lower. Hence, the given configuration has an attack because the spy has lowered the security level of Nb_a (note that Nb_a represents the nonce created by B to be sent to A).

8.4 Conclusions

Recent literature (e.g. [Low95, Pau98, Bel99, BR97a]) shows that reasoning formally about security protocols is a challenging task. Our work shows that SCSPs can be tailored to model protocol attacks. We have provided a general guideline to build the initial SCSP, the policy SCSP for any protocol, and the imputable SCSP that corresponds to a given configuration of the network.

An attack is formally defined by comparing the solution of the imputable SCSP with that of the policy SCSP. The lower is the solution of the imputable SCSP w.r.t. that of the policy SCSP, the more significant is the attack on the message that determines the difference. This is one of the original aspects that our approach brings to the field.

Our ongoing research includes formalising a dynamic update of the constraint store in order to model the dynamic evolution of the network traffic while the protocol is being executed.

Conclusions and Future Work

You'll never know how well you're made until all you had is gone.
Then the good will come out and you'll be free.
. . . a friend of mine . . .

In this thesis we explored and added softness to the constraint solving and programming framework. We have been involved in several research topics; in the following points we summarize the main results of the thesis.

A unifying framework for soft constraint handling. We have proposed a general framework for constraint solving where each constraint allows each tuple with a certain level of confidence (or degree, or cost, or other). This allows for a more realistic modelization of real-life problems, but requires a new constraint solving engine that has to take such levels into account. To do this, we used the notion of semiring, which provides both the levels and the new constraint combination operations. Different instantiations of the semiring set and operations give rise to different existing frameworks (Fuzzy CSP, Probabilistic, Classical, VCSP, etc.).

A soft version of the solution and preprocessing techniques. We considered the issue of local consistency in such an extended framework, and we provided sufficient conditions that assure these algorithms work. Then, we considered dynamic programming-like algorithms, and we proved that these algorithms can always be applied to SCSPs (and have a linear time complexity when the given SCSPs can be provided with a parsing tree of bounded size).

We have studied some possible uses of Soft Arc Consistency (SAC) over SCSPs, which could make the search for the best solution faster. In particular, we have analyzed the relationship between AC and SAC. Then we have studied the relationship between the soft constraint formalism and the constraint propagation schema proposed in [Apt99b]. What we have discovered is that the GI algorithm of [Apt99b] can also be used for soft constraints, since soft constraints provide what is needed for the GI algorithm to work correctly: a partial order with a bottom, and a set of monotone functions.

Moreover, in studying this relationship we have also discovered that, in soft constraints, we do not have to restrict ourselves to local consistency functions which solve a subproblem, but we can use *any* monotone function. Of course, in this more general case, the equivalence of the resulting problem and the initial one is not assured any more, and has to be studied on a case-by-case basis.

By passing from classical constraints to soft constraints, although on finite domains, we have more possible sources of non-termination for the GI algorithm, since the semiring can be infinite. Therefore, we have studied in depth the issue of termination for GI over soft constraints, finding some convenient sufficient conditions. In particular, one of them just requires checking, in a certain special case, that the multiplicative operation of the semiring is idempotent. We show that indeed it is possible to treat in a uniform way hard and soft constraints.

An abstraction-concretization framework for SCSPs. We have proposed an abstraction scheme for abstracting soft constraint problems, with the goal of finding an optimal solution in a shorter time. The main idea is to work on the abstract version of the problem and then bring back some useful information to the concrete one, to make it easier to solve. We have also shown how to use it to import propagation rules from the abstract setting to the concrete one.

A higher-order language to program with constraints. We have described how the semiring framework can be used to both embed the constraint structure and topology in a suitable semiring of functions. The introduction of the functions enriches the expressivity of the framework and gives the possibility to express in a general way dynamic programming and general local consistency techniques.

A CLP version able to handle Soft Constraints. We have introduced a framework for constraint programming over semirings. This allows us to use a CLP-like language for both constraint solving and optimization. In fact, constraint systems based on semirings are able to model both classical constraint solving and also more sophisticated features like uncertainty, probability, fuzziness, and optimization.

We have then given this class of languages three equivalent semantics: model-theoretic, fix-point, and proof-theoretic, in the style of classical CLP programs. Finally, we have obtained interesting decidability results for general SCLP programs and also for those SCLP programs without functions.

A declarative way to represent and solve OR problems. We have investigated the relationship between shortest path problems and constraint programming, proposing the soft constraint logic programming framework to model and solve many versions of this problem in a declarative way. In particular, both classical, multi-criteria, partially-ordered, and modality-based SP problems are shown to be expressible in this programming formalism. We also have shown that certain algorithms that solve SP problems can be used to efficiently compute the semantics of a certain class of SCLP programs.

Security constraints We have shown how to use SCSPs for the definition of network of agents executing a security protocol. Some security properties of the network are translated into specific properties of the SCSP solution and in particular the notion of inconsistency is mapped upon the notion of network insecurity.

Future work

The introduction of a general framework to deal with soft constraints could give suggestions for new research topics, only partially studied up now. Here are some ideas we plan to investigate in already explored fields and also in new ones:

Abstraction. Future work will concern the comparison between our properties and the one of [SdGS97], the generalization of our scheme to also abstract domains and graph topologies, and an experimental phase to assess the practical value of our proposal. Moreover, an experimental phase is necessary to check the real practical value of our proposal. We plan to perform such a phase within the `clp(fd,S)` system developed at INRIA [GC98], which can already solve soft constraints in the classical way (branch-and-bound plus propagation via partial arc-consistency). We also plan to investigate the relationship of our notion of abstraction with several other existing notions currently used in constraint solving. For example, it seems to us that many versions of the intelligent backtracking search could be easily modeled via soft constraints, by associating some information about the variables responsible for the failure to each constraint. Then, it should be possible to define suitable abstractions between the more complex of these frameworks and the simpler ones.

High order semirings. Future work will try to study the shapes of SCSPs by starting from their function representation. Moreover, the possibility to also use partially instantiated functions in the SCSP framework makes it possible to represent and solve *parametric problems*.

SCLP. To make SCLP a practical programming paradigm, we plan to investigate efficient techniques to implement their operational semantics, so that an optimal solution can be found efficiently. This will require considering variants of branch and bound methods, and developing intelligent ways to recognize and cut useless search branches. In this respect, we plan to study the possible use of local consistency techniques [BMR97b] for a better approximation of the semiring values to be associated to each value combination in order to obtain better bounds for the search. We will also study the relationship between semiring-based arc-consistency and classical arc-consistency, since in some cases by using classical arc-consistency one loses some information but gains in speedup.

We also plan to study the possible use of an approach similar to that of the s-semantics in logic programming [BGLM94] to give a more compact formal characterization of the fix-point semantics of SCLPs.

Moreover, we plan to collaborate with the LOCO research group at INRIA-Roquencourt, which has implemented a version of our framework on top of the existing constraint language clp(fd), to reach a reasonable trade-off between efficiency, semantic elegance of the language, and user needs.

SCLP for OR problems. Using structures more complex than semirings, or by considering trees instead of paths, we could also give semantics using OR algorithms to a greater class of SCLP programs. We plan to investigate these structures in the future.

In particular, we plan to use SCLP programs to model and solve best tree problems. We also intend to apply our methodology to specify and solve planning problems, since in many cases such problems can be cast as best tree problems over the graph describing all possible moves of the planning scenario.

Soft concurrent constraints. We plan to extend the concurrent constraint (cc) framework ([Sar89]); the possibility to handle together soft constraints and concurrency will give new strength to the language and will also open new applicability fields for this concurrent framework. As an example, consider the work described in Chapter 8 where we show how security protocols could be represented using soft concurrent constraints.

Soft constraint databases. The benefits given by the introduction of fuzziness and costs to constraint could also be used in the Constraint Databases framework [GKR95, KG94]. In this framework the queries and the database objects are represented as constraints; adding levels of preference (or cost, or probability) could add to the framework a way to discriminate between several answers for a given query.

In particular, the probability semiring could be used to represent data that have an intrinsic uncertainty to be (or not to be) in the real problem. A possible application of this idea is represented by the geographic data bases: in this framework each object could have a different probability to be in a specific position.

Soft web query. The declarative fashion of constraints together with the levels of preference given by the semiring structure could be used to represent (and execute) query on web documents. In fact, the levels of the semiring could be used to represent the levels of satisfiability of the query, and also the structural properties of the document.

Tools for soft constraints programming. Finally, a more practical application could be the insertion of soft constraints in some programming environment, like, for example, the Constraint Handling Rules formalism [Fru95].

Bibliography

[AB93] A. Aggoun and N. Beldiceanu. Overview of the chip compiler system. In A. Colmerauer and F. Benhamou, editors, *Constraint Logic Programming: Selected Research*. MIT Press, 1993.

[All43] M.F.C. Allais. *A La Recherche d'une Discipline Économique*. 1943.

[AN96] M. Abadi and R.M. Needham. Prudent engineering practice for cryptographic protocols. *TOSE*, 22(1):6–15, 1996.

[Apt98] K.R. Apt. The essence of constraint propagation. *CWI Quarterly*, 11(2&3):215–248, 1998.

[Apt99a] K.R. Apt. The essence of Constraint Propagation. *Theoretical Computer Science*, 221(1-2):179–210, 1999.

[Apt99b] K.R. Apt. The Rough Guide to Constraint Propagation (corrected version). In *Proc. of the 5th International Conference on Principles and Practice of Constraint Programming (CP'99*, number 1713 in LNCS, pages 1–23. Springer-Verlag, 1999. invited lecture.

[B9̈4] E. Börger. *Specification and Validation Methods*. Oxford University Press, 1994.

[BB72] U. Bertelè and F. Brioschi. *Nonserial Dynamic Programming*. Academic Press, 1972.

[BB00] G. Bella and S. Bistarelli. Scsps for modelling attacks to security protocols. In *Proc. CP2000 workshop on soft Constraints*, 2000.

[BB01] G. Bella and S. Bistarelli. Sof constraints for security protocol analysis: Confidentiality. In *Proc. of the 3nd International Workshop on Practical Aspects of Declarative Languages (PADL'01)*, LNCS. Springer-Verlag, 2001.

[BBD+95] C. Beierle, E. Börger, I. Durdanivich, U. Glässer, and E. Riccobene. Refining abstract state machine specifications of the steam boiler control to well documented executable code. In J. R. Abrial, E. Börger, and H.Langmaack, editors, *Formal Methods for Industrial Applications: The Steam-Boiler Case Study*, number 1165 in LNCS, pages 52–78. Springer-Verlag, 1995.

[BCGR00a] S. Bistarelli, P. Codognet, Y. Georget, and F. Rossi. Abstracting soft constraints. In *Proc. Cyprus ERCIM/COMPULOG workshop*, number 1865 in LNAI. Springer-Verlag, 2000.

[BCGR00b] S. Bistarelli, P. Codognet, Y. Georget, and F. Rossi. Labeling and partial local consistency for soft constraint programming. In *Proc. of the 2nd International Workshop on Practical Aspects of Declarative Languages (PADL'00)*, number 1753 in LNCS. Springer-Verlag, 2000.

[BCP92] J. Bellone, A. Chamard, and C. Pradelles. Plane - an evolutive planning system for aircraft production. In *Proc. 1st Int. Conf. on Practical Applications of Prolog (PAP92)*, 1992.

[BCR00] S. Bistarelli, P. Codognet, and F. Rossi. An abstraction framework for soft constraint. In *Proc. SARA2000 Symposium on Abstraction, Reformulation and Approximation*, number 1864 in LNAI. Springer-Verlag, 2000.

[BD62] R.E. Bellman and S.E. Dreyfus. *Applied Dynamic Programming*. Princeton University Press, 1962.

[Bel57] R.E. Bellman. *Dynamic Programming*. Princeton University Press, 1957.

[Bel99] G. Bella. Modelling security protocols based on smart cards. In M. Blum and C.H. Lee, editors, *Proc. of the International Workshop on Cryptographic Techniques & E-Commerce (CrypTEC'99)*, pages 139–146. City University of Hong Kong, 1999.

[Bes94] C. Bessière. Arc-consistency and arc-consistency again. *Artificial Intelligence*, 1994.

[BFBW92] A. Borning, B. Freeman-Benson, and M. Wilson. Constraint hierachies. *Lisp and Symbolic Computation*, 5(3):223–270, 1992.

[BFM$^+$96] S. Bistarelli, H. Fargier, U. Montanari, F. Rossi, T. Schiex, and G. Verfaillie. Semiring-based CSPs and valued CSPs: Basic properties and comparison. In *Over-Constrained Systems*, number 1106 in LNCS. Springer-Verlag, 1996.

[BFM$^+$99] S. Bistarelli, H. Fargier, U. Montanari, F. Rossi, T. Schiex, and G. Verfaillie. Semiring-based CSPs and Valued CSPs: Frameworks, properties, and comparison. *CONSTRAINTS: An international journal. Kluwer*, 4(3), 1999.

[BGLM94] A. Bossi, M. Gabrielli, G. Levi, and M. Martelli. The s-semantics approach: Theory and applications. *The Journal of Logic Programming*, 19,20:149–197, 1994.

[BGR00] S. Bistarelli, R. Gennari, and F. Rossi. Constraint propagation for soft constraint satisfaction problems: Generalization and termination conditions. In *Proc. Sixth Int. Conf. on Principles and Practice of Constraint Programming (CP2000)*, LNCS. Springer-Verlag, 2000.

[Bis94] S. Bistarelli. Programmazione con vincoli pesati e ottimizzazione. Master's thesis, Dipartimento di Informatica, Università di Pisa, Italy, 1994. In Italian.

[BKC94] G. Baues, P. Kay, and P. Charlier. Constraint based resource allocation for airline crew management. In *Proc. ATTIS'94*, 1994.

[BL65] G. Birkhoff and S. Mac Lane. *A Survey of Modern Algebra*. The Macmillan Company, 1965.

[BMMW89] A. Borning, M. Maher, A. Martindale, and M. Wilson. Constraint hierarchies and logic programming. In G. Levi and M. Martelli, editors, *Proc. 6th International Conference on Logic Programming*, pages 149–164. MIT Press, 1989.

[BMPT00] G. Bella, F. Massacci, L.C. Paulson, and P. Tramontano. Formal Verification of Cardholder Registration in SET. In *Proc. of European Symposium on Research in Computer Security (ESORICS 2000)*, LNCS. Springer-Verlag, 2000. In press.

[BMR95] S. Bistarelli, U. Montanari, and F. Rossi. Constraint Solving over Semirings. In *Proc. IJCAI95*, San Francisco, CA, USA, 1995. Morgan Kaufman.

[BMR97a] S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based Constraint Logic Programming. In *Proc. IJCAI97*, San Francisco, CA, USA, 1997. Morgan Kaufman.

[BMR97b] S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based Constraint Solving and Optimization. *Journal of the ACM*, 44(2):201–236, Mar 1997.

[BMR99] S. Bistarelli, U. Montanari, and F. Rossi. Sclp semantics for (multi-criteria) shortest path problems. In *CP-AI-OR'99 - Workshop on Integration of AI and OR technique in Constraint Programming for Combinatorial Optimization Problems*, 1999.

[BMR00a] S. Bistarelli, U. Montanari, and F. Rossi. Higher order semiring-based constraints. In *Proc. CP2000 workshop on soft Constraints*, 2000.

[BMR00b] S. Bistarelli, U. Montanari, and F. Rossi. Soft constraint logic programming and generalized shortest path problems. *Journal of Heuristics*, 2000.

[BMR01] S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based Constraint Logic Programming: Syntax and Semantics. *ACM Transactions on Programming Languages and System (TOPLAS)*, 2001. To Appear.

[Bor81] A. Borning. The programming language aspects of thinglab, a constraint oriented simulation laboratory. *ACM Transaction on Programming Languages and Systems*, 3(4):353–387, 1981.

[BP98] G. Bella and L.C. Paulson. Kerberos version iv: Inductive analysis of the secrecy goals. In J.-J Quisquater, Y. Desware, C. Meadows, and D. Gollmann, editors, *ESORICS98*, volume 1485 of *LNCS*, pages 361–375. SV, 1998.

[BR97a] G. Bella and E. Riccobene. Formal analysis of the kerberos authentication system. *JUCS*, 3(12):1337–1381, 1997.

[BR97b] S. Bistarelli and E. Riccobene. An operational model for the SCLP language. In *Proc. ILPS97 workshop on Tools and Environments for (Constraint) Logic Programming*, 1997.

[BR98] S. Bistarelli and F. Rossi. About Arc-Cconsistency in Semiring-based Constraint Problems. In *Presented at the Fifth International Symposium on Artificial Intelligence and Mathematics*, 1998.

[CC77] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proc. Fourth ACM Symp. Principles of Programming Languages*, pages 238–252, 1977.

[CC79] P. Cousot and R. Cousot. Systematic design of program analyis. In *Proc. Sixth ACM Symp. Principles of Programming Languages*, pages 269–282, 1979.

[CC92] P. Cousot and R. Cousot. Abstract interpretation and application to logic programming. *The Journal of Logic Programming*, 13(2,3), 1992.

[CCTY99] A.H.W. Chun, S.H.C. Chan, F.M.F. Tsang, and D.W.M. Yeung. Stand allocation with constraint technologies at Chek Lap Kok international airport. In *Proc. PACLP99*, 1999.

[CD96] P. Codognet and D. Diaz. Compiling constraints in `clp(fd)`. *The Journal of Logic Programming*, 27(3), 1996.

[CHV98] P. Chan, K. Heus, and G. Veil. Nurse scheduling with global constraints in CHIP: Gymnaste. In *Proc. PACT98*, 1998.

[Col96] C. Collignon. Gestion optimisee de ressources humaines pour l'audiovisuel. In *Proc. CHIP users' club*, 1996.

[Cou77] P. Cousot. Asynchronous iterative methods for solving a fixed point system of monotone equations in a complete lattice. Technical Report R.R.88, Institut National Polytechnique de Grenoble, 1977.

[CP97] K.P. Chow and M. Perrett. Airport counter allocation using constraint logic programming. In *Proc. PACT97*, 1997.

[DD88] R. Dechter and A. Dechter. Ief maintenance in dynamic constraint networks. In *Proc. AAAI88*, pages 37–42, 1988.

[DDP90] R. Dechter, A. Dechter, and J. Pearl. Optimization in constraint networks. In R.M. Oliver and J.Q. Smith, editors, *Influence Diagrams, Belief Nets and Decision Analysis*, pages 411–425. John Wiley and Sons Ltd., 1990.

[DFP93] D. Dubois, H. Fargier, and H. Prade. The calculus of fuzzy restrictions as a basis for flexible constraint satisfaction. In *Proc. IEEE International Conference on Fuzzy Systems*, pages 1131–1136. IEEE, 1993.

[DJ96] A. Dubos and A. Du Jeu. Application EPPER planification des ageents roulaants. In *Proc. CHIP users' club*, 1996.

[DP88a] R. Dechter and J. Pearl. Network-Based Heuristics for Constraint-Satisfaction Problems. In Kanal and Kumar, editors, *Search in Artificial Intelligence*, pages 370–425. Springer-Verlag, 1988.

[DP88b] R. Dechter and J. Pearl. Tree-clustering schemes for constraint processing. In *Proc. AAAI88*, pages 150–154, 1988.

[DP90] B.A. Davey and H.A. Priestley. *Introduction to lattices and order*. Cambridge University Press, 1990.

[Dre69] S.E. Dreyfus. An appraisal of some shortest-paths algorithms. *Operation Research*, 17:395–412, 1969.

[DS91] M. Dincbas and H. Simonis. Apache - a constraint based, automated stand allocation system. In *Proc. of Advanced Software technology in Air Transport (ASTAIR'91)*, 1991.

[Ell93] T. Ellman. Synthesis of abstraction hierarchies for constraint satisfaction by clustering approximately equivalent objects. In *Proc. of International Conference on Machine Learning*, pages 104–111, 1993.

[Emd86] M.H. Van Emden. Quantitative deduction and its fixpoint theory. *The Journal of Logic Programming*, 3(1), 1986.

[FB97] T. Fruhwirth and P. Brisset. Optimal planning of digital cordless telecommunication systems. In *Proc. PACT97*, 1997.

[FGS95] M. Fromherz, V. Gupta, and V. Saraswat. Model-based computing: constructing constraint-based software for electro-mechanical systems. In *Proc. Conf. on Practical Applications of Constraint Technology (PACT95)*, 1995.

[Fik70] R.E. Fikes. REF-ARF: A system for solving problems stated as procedures. *Artificial Intelligence*, 1:27–120, 1970.

[Fit81] M. Fitting. Bilattices and the semantics of logic programming. *The Journal of Logic Programming*, 11(1–2):91–116, 1981.

[FL93] H. Fargier and J. Lang. Uncertainty in constraint satisfaction problems: a probabilistic approach. In *Proc. European Conference on Symbolic and Qualitative Approaches to Reasoning and Uncertainty (ECSQARU)*, number 747 in LNCS, pages 97–104. Springer-Verlag, 1993.

[FLMM97] F. Focacci, E. Lamma, P. Mello, and M. Milano. Constraint logic programming for the crew rostering problem. In *Proc. PACT97*, 1997.

[Flo62] R.W. Floyd. Algorithm 97: Shortest path. *Communication of ACM*, 5:345, 1962.

[Fre78] E.C. Freuder. Synthesizing constraint expressions. *Communication of the ACM*, 21(11), 1978.

[Fre88] E.C. Freuder. Backtrack-free and backtrack-bounded search. In Kanal and Kumar, editors, *Search in Artificial Intelligence*, pages 343–369. Springer-Verlag, 1988.

[Fre91] E.C. Freuder. Eliminating interchangeable values in constraint satisfaction subproblems. In P. Marcotte and S. Nguyen, editors, *Proc. AAAI-91*, 1991.

[Fru95] T. Fruhwirth. Constraint handling rules. In A. Podelski, editor, *Constraint Programming: Basics and Trends*, volume 910 of *LNCS*. Springer-Verlag, 1995.

[FS95] E.C. Freuder and D. Sabin. Interchangeability supports abstraction and reformulation for constraint satisfaction. In *Proc. of Symposium on Abstraction, Reformulation and Approximation (SARA'95)*, 1995. http://www.cs.unh.edu/csprojects/csp/csp.html.

[FW92] E.C. Freuder and R.J. Wallace. Partial constraint satisfaction. *AI Journal*, 58, 1992.

[GC98] Y. Georget and P. Codognet. Compiling semiring-based constraints with clp(fd,s). In *Proc. CP98*, number 1520 in LNCS. Springer-Verlag, 1998.

[GCM99] C. Gervet, Y. Caseau, and D. Montaut. On refining ill-defined constraint problems: A case study in iterative prototyping. In *Proc. PACLP99*, 1999.

[GDL92] R. Giacobazzi, S.K. Debray, and G. Levi. A Generalized Semantics for Constraint Logic Programs. In *Proc. FGCS92*, pages 581–591. ICOT, 1992.

[Geo99] Y. Georget. *Extensions de la Programmation par Contraintes*. PhD thesis, Ecole Polytecnique, Paris, 1999.

[GJC94] M. Gyssens, P.G. Jeavons, and D.A. Cohen. Decomposing Constraint Satisfaction Problems Using Database Techniques. *Artificial Intelligence*, 66(1):57–89, 1994.

[GKR95] P.C. Kanellakis G.M. Kuper and P.Z. Revesz. Constraint query languages. *Journal of Computer and System Science*, 1995.

[GMM81] S. Gnesi, A. Martelli, and U. Montanari. Dynamic programming as graph searching: an algebraic approach. *Journal of the ACM*, 28(5):737–751, 1981.

[GR97] F. Glaisner and L.-M. Richard. FORWARD-C: A refinery scheduling system. In *Proc. PACT97*, 1997.

[Gur95] Y. Gurevich. Evolving algebras 1993: Lipari guide. In E. Börger, editor, *Specification and Validation Methods*. Oxford University Press, 1995.

[HDT92] P. Van Hentenryck, Y. Deville, and C.M. Teng. A generic arc-consistency algorithm and its specializations. *Artificial Intelligence*, 57:291–321, 1992.

[Hen89] P. Van Hentenryck. *Constraint Satisfaction in Logic Programming*. MIT Press, 1989.

[JL86] J. Jaffar and J.L. Lassez. Constraint logic programming. Technical report, IBM T.J. Watson Research Center, Yorktown Heights, 1986.

[JL87] J. Jaffar and J.L. Lassez. Constraint logic programming. In *Proc. POPL*. ACM, 1987.

[JM94] J. Jaffar and M.J. Maher. Constraint logic programming: A survey. *The Journal of Logic Programming*, 19 and 20, 1994.

[KG94] P.C. Kanellakis and D.Q. Goldin. Constraint programming and database query languages. In *Proc. of the Int. Symp. on Theoretical Aspects of Computer Software*, number 789 in LNCS, pages 66–77, 1994.

[KG98] F. Kokkoras and S. Gregory. D-WMS: Distributed workforce management using CLP. In *Proc. PACT98*, 1998.

[Kum92] V. Kumar. Algorithms for constraint satisfaction problems: a survey. *AI Magazine*, 13(1):32–44, 1992.

[Lau78] J-L. Lauriere. A language and a program for stating and solving combinatorial problems. *Artifical Intelligence Journal*, 10:29–127, 1978.

[Lel88] W. Leler. *Constraint Programming Languages, their specification and generation*. Addison-Wesley, 1988.

[Llo87] J.W. Lloyd. *Foundations of Logic Programming*. Springer-Verlag, 1987.

[Low95] G. Lowe. An attack on the needham-schroeder public-key authentication protocol. *Information Processing Letters*, 56(3):131–133, 1995.

[Low96] G. Lowe. Some new attacks upon security protocols. In *CSFW96*, pages 139–146. IEEEP, 1996.

[LR97] G. Lowe and B. Roscoe. Using csp to detect errors in the tmn protocol. *TOSE*, 3(10), 1997.

[Mac77] A.K. Mackworth. Consistency in networks of relations. *Artificial Intelligence*, 8(1):99–118, 1977.

[Mac92] A.K. Mackworth. Constraint satisfaction. In S.C. Shapiro, editor, *Encyclopedia of AI (second edition)*, pages 285–293. John Wiley & Sons, 1992.

[MF85] A.K. Mackworth and E.C. Freuder. The complexity of some polynomial network consistency algorithms for constraint satisfaction problems. *Artificial Intelligence*, 25:65–74, 1985.

[MH86] R. Mohr and T.C. Henderson. Arc and path consistency revisited. *Artificial Intelligence*, 28(2):225–233, 1986.

[MM72] A. Martelli and U. Montanari. Nonserial dynamic programming: on the optimal strategy of variable elimination for the rectangular lattice. *Journal of Mathematical Analysis and Application*, 40:226–242, 1972.

[Mon71] U. Montanari. Fundamental properties and applications to picture processing, information sciences. Technical Report 71-1, Dept. of Computer Science, Carnegie Mellon University, jan 1971.

[Mon74] U. Montanari. Networks of constraints: Fundamental properties and applications to picture processing. *Information Science*, 7:95–132, 1974. Also Technical Report, Carnegie Mellon University, 1971.

[Mor92] P. Morris. On the density of solutions in equilibrium points for the queens problem. In *Proc. AAAI-92*, pages 428–433, 1992.

[Mou88] H. Moulin. *Axioms for Cooperative Decision Making*. Cambridge University Press, 1988.

[MPS97] B. Mobasher, D. Pigozzi, and G. Slutzki. Multi-valued logic programming semantics: An algebraic approach. *Theoretical Computer Science*, 171:77–109, jan 1997.

[MR86] U. Montanari and F. Rossi. An efficient algorithm for the solution of hierachical networks of constraints. In *Proc. International Workshop on Graph Grammars and their application to Computer Science*, number 291 in LNCS, pages 440–457. Springer-Verlag, 1986.

[MR91] U. Montanari and F. Rossi. Constraint relaxation may be perfect. *Artificial Intelligence Journal*, 48:143–170, 1991.

[MR98] W. Maximuck and M. Reviakin. A prototype of train schedule module. In *Proc. PACT98*, 1998.

[MS98] K. Marriott and P. Stuckey. *Programming with Constraints*. MIT Press, 1998.

[Nat77] National Bureau of Standards. *Data Encryption Standard*, Jan 1977. Federal Information Processing Standards Publications, FIPS Pub. 46.

[NS78] R.M. Needham and M.D. Schroeder. Using encryption for authentication in large networks of computers. *CACM*, 21(12):993–999, 1978.

[Par06] V. Pareto. *Manual of Political Economy*. 1906.

[Pau98] L.C. Paulson. The inductive approach to verifying cryptographic protocols. *JCS*, 6:85–128, 1998.

[Per91] M. Perrett. Using constraint logic programming techniques in container port planning. *ICL Technical Journal*, pages 537–545, 1991.

[PJ99] L. Purvis and P. Jeavons. Constraint tractability theory and its application to the product development process for a constraint-based scheduler. In *Proc. PACLP99*, 1999.

[PS88] S. Pallottino and M.G. Scutellà. Shortest path algorithms in transportation models: classical and innovative aspects. In P. Marcotte and S. Nguyen, editors, *Equilibrium and Advanced Transportation Modelling*, pages 245–281, 1988.

[Pug94] J-F. Puget. A c++ implementation of clp. In *Proc. of SPICIS 94*, nov 1994.

[RHZ76] A. Rosenfeld, R.A. Hummel, and S.W. Zucker. Scene labelling by relaxation operations. *IEEE Transactions on Systems, Man, and Cybernetics*, 6(6):420–433, 1976.

[Ros95] F. Rossi. Redundant hidden variables in finite domain constraint problems. In M. Meyer, editor, *Constraint Processing*, number 923 in LNCS. Springer-Verlag, 1995.

[RSA76] R. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *CACM*, 21(2):120–126, 1976.

[Rut94] Z. Ruttkay. Fuzzy constraint satisfaction. In *Proc. 3rd IEEE International Conference on Fuzzy Systems*, pages 1263–1268, 1994.

[Rut98] Z. Ruttkay. Constraint satisfaction a survey. *CWI Quarterly*, 11(2–3):163–214, 1998.

[Sar89] V.A. Saraswat. *Concurrent constraint programming language*. PhD thesis, Carnegie-Mellon University, 1989.

[SC95] H. Simonis and T. Cornelissens. Proceedings 1st int. conference on principles and practice of constraint programming (cp95). In *Proc. PACLP99*, number 976 in LNCS. Springer-Verlag, 1995.

[SC98] H. Simonis and P. Charlier. Cobra - a system for train crew scheduling. In *Proc. DIMACS workshop on constraint programming and large scale combinatorial optimization*, 1998.

[Sch92] T. Schiex. Possibilistic constraint satisfaction problems, or "how to handle soft constraints?". In *Proc. 8th Conf. of Uncertainty in AI*, pages 269–275, 1992.

[Sco82] D.S. Scott. Domains for denotational semantics. In *Proc. ICALP*. Springer-Verlag, 1982.

[SdGS97] G. Verfaillie S. de Givry and T. Schiex. Bounding the optimum of constraint optimization problems. In G. Smolka, editor, *Proc. of the Third International Conference on Principles and Practice of Constraint Programming CP97*, number 1330 in LNCS, pages 405–419. Springer-Verlag, oct 1997.

[SF97] D. Sabin and E.C. Freuder. Understanding and improving the mac algorithm. In *Proc. CP97*, number 1330 in LNCS, pages 167–181. Springer-Verlag, 1997.

[SFV95] T. Schiex, H. Fargier, and G. Verfaille. Valued Constraint Satisfaction Problems: Hard and Easy Problems. In *Proc. IJCAI95*, pages 631–637, San Francisco, CA, USA, 1995. Morgan Kaufmann.

[SG91] R. Sosic and J. Gu. 3000000 queens in less than one minute. *SIGART Bulletin*, 2(3):22–24, 1991.

[She91] P. Shenoy. Valuation-based systems for discrete optimization. In L.N. Kanal P.P. Bonissone, M. Henrion and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*. Elsevier Science, 1991.

[She94] P. Shenoy. Valuation-based systems: A framework for managing uncertainty in expert systems. In L.Zadeh and J.Kacprzyk, editors, *Fuzzy Logic for the Management of Uncertainty*, pages 83–104. Wiley Prof. Computing, 1994.

[SM92] R. Schrag and D. Miranker. An evaluation of domain reduction: Abstraction for unstructured CSPs. In *Proc. of the Symposium on Abstraction, Reformulation, and Approximation*, pages 126–133, 1992.

[Smy78] M.B. Smyth. Power domains. *Journal of Computer and System Sciences*, 16(1):23–36, feb 1978.

[SS80] G.J. Sussman and G.L. Steele. CONSTRAINTS - a language for expressing almost-hierarchical descriptions. *Artificial Intelligence*, 14(1):1–39, 1980.

[Ste80] G.L. Steel. *The definition and implementation of a computer programming language based on Constraints*. PhD thesis, MIT, 1980.

[Sut63] I. Sutherland. Sketchpad: a man-machine graphical communication system. In *Proc. IFIP Spring Joint Computer Conference*, 1963.

[Tar55] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.

[TCR90] C.E. Leiserson T.H. Cormen and R.L. Rivest. *Introduction to Algorithms*, chapter 26. MIT Press, 1990.

[Wal75] D.L. Waltz. Understanding line drawings of scenes with shadows. In Winston, editor, *The Psychology of Computer Vision*. McGraw-Hill, 1975.

[Wal96] M. Wallace. Practical applications of constraint programming. *Constraints: An International Journal*, 1(1 and 2), 1996.

[WG92a] T. Walsh and F. Giunchiglia. Theories of abstraction: A historical perspective. In *Proc. of AAAI-92*, 1992. Also available in LNCS 923, pp. 121-138, 1995.

[WG92b] T. Walsh and F. Giunchiglia. A theory of abstraction. *Artificial Intelligence*, 56(2-3):323–390, 1992.

[Zad65] L.A. Zadeh. Fuzzy sets. *Information and Control*, 8:338–353, 1965.