# Constraint Satisfaction: Algorithms, Complexity Results, and Applications 

by<br>Francesco Lupia

University of Calabria<br>DIMES - Dipartimento di Ingegneria Informatica, Modellistica, Elettronica e Sistemistica PhD Program in Systems and Computer Engineering

(c) Francesco Lupia 2015

All Rights Reserved

To my family

## ACKNOWLEDGEMENTS

First, I want to thank my advisors Francesco Scarcello and Gianluigi Greco. I believe they are the best advisors I could ever desire. They complete me in many aspects. Francesco has been an endless source of inspiration for me. His advice goes well beyond the research and academic context to cover many aspects of my life. Perhaps, the greatest thing that I can think of is that he constantly push my abilities to think and reason about things in a meticulous way. He is always willing to take time out of his busy days to help with important problems and all the unexpected issues that come up. Thanks to Francesco I met Gianluigi. He is truly another source of countless motivations. He constantly impresses me at how good he is in everything he does and for his ability to be in involved in many different things at the same time. Francesco and Gianlugi are constantly teaching me to take a simple idea, to find generalizations and implications and to push it to the limits. But this is only the beginning.

I thank Walter Lucia for sparking my initial interest for research and for all the interactions that we had in the last 10 years. I am also indebted with Antonella Guzzo and Luigi Pontieri for their priceless help in the last three years. Now let me thank my office colleagues Valeria Fionda and Carmine Dodaro with whom during my "PhD days" I had a huge number of valuable discussions. A special thank you goes to Luca Ghionna for his valuable feedback on some aspects of this dissertation. I especially thank Domenico Saccà who worked behind the scenes to support me
with the fundings and I must say thanks to Nicola Leone for allowing me to stay and feel part of the department of Mathematics and Computer Science while I was working with Gianluigi. Moreover I want to thank all the friends from the DIMES and Mathematics and Computer Science departments. In the last three years, it was not all about research and work. I especially thank Massimo and Alessandro for being always present. And for the same reason Francesco deserves a big thank you and our friend Antonio. I also special thank my other two cousins Pasquale and Salvatore. Finally, I am who I am thanks to my mother, my father, my sister Valentina and my grandmother Maria and grandfather Pasquale. I wish I will have the opportunity to return their love and support.

## TABLE OF CONTENTS

DEDICATION ..... ii
ACKNOWLEDGEMENTS ..... iii
LIST OF FIGURES ..... viii
LIST OF TABLES ..... x
CHAPTER

1. Introduction ..... 1
1.1 Outline of the dissertation ..... 4
2. Setup and Preliminaries ..... 6
2.1 Constraint satisfaction problems ..... 6
2.2 Solving CSPs ..... 9
2.2.1 Backtracking search ..... 10
2.3 Structural properties of CSPs ..... 12
2.4 Structural decomposition methods ..... 13
2.4.1 Tree decompositions ..... 13
2.4.2 (Generalized) Hypertree decompositions ..... 14
2.4.3 Tree projections ..... 17
2.5 Counting problems ..... 20
2.5.1 The Langford problem ..... 22
2.5.2 Configuration problems: Renault example ..... 24
2.6 Preliminaries on Complexity Theory ..... 25
I CSP and Databases ..... 27
3. Basic Definitions from Database Theory ..... 28
3.1 Relational structures ..... 28
3.2 Relational databases ..... 29
3.3 Conjunctive queries ..... 30
3.4 Relational algebra. ..... 31
3.5 Hypergraphs and structural restrictions ..... 31
3.6 SQL queries ..... 32
4. A Weighted Structural Decomposition Technique ..... 34
4.1 Introduction ..... 34
4.2 Monotone Greedy Tree projections ..... 35
4.3 Normal form ..... 36
4.4 An algorithm for computing greedy tree projections in normal form ..... 38
4.4.1 Speeding-up computation through greedy coverings ..... 40
4.5 Evaluation functions ..... 41
4.6 Non-monotic valuation functions ..... 44
4.7 Experimental results ..... 46
5. An Hybrid Approach for Counting Solutions ..... 49
5.1 Introduction to Hybrid Tractability ..... 49
5.2 An algorithm for counting answers ..... 52
5.3 Implementation issues and System Architecture ..... 55
5.3.1 Hacking PostgreSQL ..... 58
5.4 Some experiments ..... 59
5.4.1 Query 3 (Q3) on TPC-DS ..... 61
5.4.2 Query 4 (Q4) on TPC-DS ..... 63
II CSP and Game Theory ..... 65
6. An Introduction to Game Theory ..... 66
6.1 Coalitional games ..... 67
6.2 Allocation games ..... 70
6.3 Solution concepts ..... 72
6.3.1 The Core ..... 73
6.3.2 The Nucleoulus ..... 76
6.3.3 The Kernel ..... 77
6.3.4 The Bargaining set ..... 78
6.4 The Shapley value ..... 79
6.4.1 The Banzhaf value ..... 81
6.4.2 Back to the allocation games ..... 82
6.5 A motivating example: The Italian Research Assessment Program (VQR) ..... 83
6.5.1 Division rules ..... 84
6.5.2 Desiderata for division rules ..... 86
6.5.3 Marginal contribution ..... 88
6.5.4 A simple scenario ..... 90
6.5.5 Using the Shapley value as a division rule ..... 92
6.5.6 Discussion ..... 93
7. Structural Tractability of Shapley and Banzhaf Values in Allocation Games ..... 97
7.1 Introduction ..... 97
7.2 Intractability of computation ..... 99
7.3 Characterizations of the Shapley value ..... 102
7.3.1 A closer look at marginal contributions ..... 102
7.4 Islands of tractability ..... 107
7.4.1 Bounded sharing ..... 107
7.4.2 Bounded treewidth ..... 110
7.5 CSP encoding (for the Banzhaf value) ..... 111
7.6 From the Banzhaf value to the Shapley value ..... 117
7.7 Summary ..... 119
III CSP and Mining ..... 121
8. Process Mining ..... 122
8.1 An overview of process discovery ..... 123
8.2 Bottom-up vs top-down design methods ..... 127
8.3 Causal nets and logs ..... 129
8.4 Dependency graphs and process mining: basic results ..... 135
9. Precedence Constraints: Mining Problems and Complexity ..... 146
9.1 Introduction ..... 146
9.2 Syntax and Semantics ..... 146
9.3 Precedence Constraints ..... 151
9.4 Complexity Analysis ..... 159
9.5 An exact solution approach for computing process models ..... 169
9.5.1 CSP encoding for precedence constraints ..... 170
9.5.2 Structural optimization ..... 174
9.6 Classes of Tractable Precedence Constraints and Algorithms ..... 175
9.6.1 Precedence constraints without negated paths ..... 177
9.6.2 The case of negated path constraints ..... 187
9.7 Experimental evaluation ..... 193
9.7.1 Case study: a product-recall process ..... 196
9.7.2 Comparative analysis on benchmark data ..... 206
9.7.3 Further tests on synthesized data ..... 211
9.8 Summary ..... 215
10. An Application Scenario (Beyond Process Mining): Urban Congestion ..... 219
10.1 Introduction ..... 219
10.2 Project goals ..... 220
10.3 Urban log mining ..... 221
IV Conclusions ..... 225
11. Conclusion and Future Research ..... 226
11.1 Contributions ..... 226
11.1.1 Databases ..... 226
11.1.2 Game theory ..... 227
11.1.3 Process mining ..... 228
11.2 Future research directions ..... 230

## LIST OF FIGURES

## Figure

2.1 A hypergraph $\mathcal{H}$ and a tree decomposition of $\mathcal{H}$ of width 3 ..... 14
2.2 A hypergraph $\mathcal{H}$ and a hypertree decomposition of $\mathcal{H}$ of width 2 ..... 16
2.3 A tree projection $\mathcal{H}_{a}$ for the pair of hypergraphs $\left(\mathcal{H}_{1}, \mathcal{H}_{2}\right)$. ..... 18
2.4 Langford problem ..... 22
2.5 A portion of hypertree decomposition of width 4 for the Renault megane configu- ration problem ..... 25
4.1 Weighted candidates graph in example of Section 4.6 ..... 44
5.1 The hypergraph $\mathcal{H}\left(Q_{2}\right)$ and a hypertree decomposition of $\mathcal{H}\left(Q_{2}\right)$ of width 2 ..... 50
5.2 The speedup of PostgreSQL 9.3.5 compared to 9.0.2 ..... 61
5.3 Execution time for query Q3 ..... 62
5.4 Number of solutions against different values for the attribute inv_quantity_on_hand ..... 63
5.5 Execution time for query Q4 ..... 64
6.1 Allocation scenario $\mathcal{A}_{0}$ in Example 6.3. ..... 70
6.2 Running example ..... 90
6.3 A clooser look at strategical behaviours. ..... 94
6.4 Co-autorships at University of Calabria. Researchers against the number of co- authored products ..... 95
6.5 The number of components at University of Calabria ..... 95
6.6 An example component ..... 96
7.1 Illustration in the proof of Theorem 7.1: Construction of the scenario $\mathcal{A}_{\alpha}$ based on $\mathcal{A}$. ..... 100
7.2 Decomposition in Example 7.12- the label of the root modified as in the proof of Theorem 7.13 is on the top. ..... 112
7.3 CSP encoding in Section 7.5. ..... 114
8.1 Process models in the running example. ..... 124
8.2 A process model involving a cycle, with an example unfolding. ..... 139
9.1 Example construction in the proof of Theorem 9.4. ..... 149
9.2 Computation of a possibly extended causal net in Theorem 9.6. ..... 152
9.3 Tractability frontiers. A set $S \subseteq\{\rightarrow, \rightsquigarrow, \nrightarrow, \not \nsim\}$ above (resp., below) the frontier means that the corresponding problem is NP-hard (resp., feasible in polynomial time) on $\mathcal{C}[\mathrm{S}]$. ..... 160
9.4 Example reduction in the proof of Lemma 9.14. ..... 162
9.5 Algorithm PCtoCSP. ..... 171
9.6 Algorithm DG-DiscoveryToCSP. ..... 173
9.7 Precedence graphs for the examples in Section 9.6. ..... 179
9.8 Algorithm COMPUTE-CN (on $\mathcal{C}[\{\rightarrow, \rightsquigarrow, \nrightarrow\}])$. ..... 182
9.9 Algorithm COMPUTE-CN (on $\mathcal{C}[\{\nsim\}\}]$ ) ..... 186
9.10 Causal net of the ProductRecall process. ..... 196
9.11 F-measure scores obtained by COMPUTE-CN when varying the percentages of posi-tive edge/path constraints and of negative edge/path constraints, for four differentfamilies of $\log$ samples, corresponding to $3 \%$ (a) and $6 \%$ (b) of traces in the $\log L$,respectively.201
9.12 Computation time spent by algorithm COMPUTE-CN with different amounts of traces and constraints' percentages in input. A base-2 logarithmic scale is used for the vertical axis in both figures, as well as for the horizontal axis in the left-hand figure.
9.13 Results on synthesized $\log$ data, with different degrees of parallelism (AP): accuracy of discovered process models (in terms of F-measure w.r.t. real activity dependencies), and rates of unsatisfied (negative path) constraints. Both measures are reported for different amounts of a-priori constraints of all types (expressed as percentages w.r.t. the sizes of their respective population).
10.1 A screenshot of the plug-in CNMining implementing the algorithms presented in the paper. The causal net is drawn by exploiting the Flex interface available in ProM, where bindings are displayed inside tooltips activated by hovering the mouse over the corresponding node.223

## LIST OF TABLES

## Table

2.1 The provinces of Calabria and its colored graph ..... 8
2.2 Known solutions for the Langford problem. A dash (-) means that there is no solution for the given value of $n$ while (?) stands for unknown. ..... 24
4.1 The time (in seconds) required to solve a number of instances generated from real- world, academic, random and pattern problems. A dash (-) means timeout while $(+)$ denotes overflow ..... 47
4.2 Number of correctly decomposed instances and number of wins in finding the min- imal width. ..... 47
9.1 Process discovery algorithms used in the experiments: legend of symbols. ..... 194
9.2 F-measure scores obtained by algorithm COMPUTE-CN and its competitors, for different percentages of traces of process ProductRecall (Figure 9.10), without (left) and with (right) a-priori knowledge on parallelism relationships. For both cases, maximal scores on each trace percentage are written in bold. ..... 199
9.3 Benchmark logs: structural characteristics and statistics (see also Weerdt et al. [2012]). Each $\log$ consists of 300 (not necessarily distinct) traces. ..... 206
9.4 Average conformance measures obtained, on benchmark logs, by different discov- ery methods-including the one proposed in this thesis (Here). For each row, the best average score(excluding ground-truth models) is underlined, while the re- sults that were recognized as significantly better than the average of the outcomes over the various methods (for the same metrics and setting) are in bold. ..... 208
9.5 Statistics on the "critical" sub-log extracted for each of the benchmark logs in Table 9.7.2. ..... 210
9.6 Results on "critical" samples without and with background knowledge. For each row, the best average scoreis underlined, while the results that were recognized as significantly better than the average of the outcomes over the various methods (for the same metrics and setting) are in bold. ..... 211
11.1 Summary of constraint support. ..... 229

ABSTRACT<br>Constraint Satisfaction: Algorithms, Complexity Results, and Applications<br>by<br>Francesco Lupia

A fundamental problem in the field of Artificial Intelligence and related disciplines, in particular Database theory, is the constraint satisfaction problem (or CSP) which comes as a unifying framework to express a wide spectrum of computational problems. Examples include graph colorability, planning, and database queries. The goal is either to find one solution, to enumerate all solutions, or counting them. As a very general problem, it comes with no surprise that in most settings CSPs are hard to solve. Indeed considerable effort has been invested by the scientific community to shed light on the computational issues of this problem, with the objective of identifying easy instances (also called islands of tractability) and exploiting the knowledge derived from their solution to help solving the harder ones.

My thesis investigates the role that structural properties play in the computational aspects of CSPs, describes algorithms to exploit such properties, and provides a number of specific tools to solve efficiently problems arising in database theory, game theory, and process mining.

## CHAPTER 1

## Introduction

Constraint satisfaction is a very general framework to express a wide range of computational problems. Examples include graph colorability, planning, satisfiability, scheduling, and database queries. Being such a general framework, we well expect computational issues to come up. Indeed, a great effort has been invested by the scientific community to study the complexity of this general problem with the objective to identify classes of easy instances known as "islands of tractability". It is well known that the structure of the constraint satisfaction problem (short: CSP), which is represented by a constraint (hyper)graph, plays a significant role in the study of its computational complexity. Indeed, while such problems are intractable in general, whenever the input has some kind of structures (for example, the constraint graph is a tree), we may compute a solution efficiently, e.g., via dynamic programming. Moreover, by precisely understanding the structure of easy problems, we may also be able to approximate hard problems by manipulating their structure in a suitable way. Unlike most approaches in the CSP research community that focus on the problem of computing just one solution to CSPs, we are mainly interested here in finding all solutions or in counting the number of solutions efficiently (i.e., without actually computing them). Database queries are an important domain of
application of our techniques. Indeed, even structurally simple queries may have an exponential number of solutions w.r.t. the input size, so that even state-of-the-art solvers are not helpful. We believe that structural decomposition methods, besides their theoretical interest, may find useful applications in solving real-world instances of these problems.

The thesis also presents a further line of research where CSP techniques are employed to model and solve Process Mining problems. Such problems arise in the study of workflow management systems that aim at supporting the execution of business processes in various contexts. In particular, to support the (re)design and the optimization of existing processes, it is desirable to identify the best strategies to schedule the execution of any given set of activities, by looking at the choices made in the past and registered in execution logs.

In summary, the thesis exploits the general CSP framework as a weapon to analyze and to attack problems arising in different areas of research. The main results, of both theoretical and practical interest, are listed in the next section.

### 1.1 Contributions

- Database theory. In this area of research, we use CSP structural techniques for the problem of counting solutions of SQL queries involving "COUNT" aggregates, which occur very often in practice. First, we propose a novel structural decomposition notion that extends the framework of greedy tree projections by adding weights. Then we design, analyze, and give a polynomial time implementation of an algorithm for computing a restricted variant of such decompositions, called monotone greedy tree projections. Then, we identify a restricted class of weighting functions for which the computation of an optimal weighted
decomposition is tractable. Furthermore, we propose optimization algorithms, which exploit the novel structural framework, that can be used to boost the performances of relational DBMS.
- Game theory. In this field we focused on coalitional games, which model situations where players can obtain higher worths by collaborating with each other, rather than by acting in isolation. We consider resource allocation problems and in particular the notion of Shapley value and Banzhaf value, which are two well-known solution concepts to provide fair worths distribution. For this class of games, is known that computing either solution concept is an intractable problem, formally \#P-hard. Motivated by these bad news, we identified large islands of tractability for both solution concepts, by exploring the structure of the interaction graphs among the players. We showed that tractability holds in scenarios where each good is owned by at most two agents (independently of their interactions). Moreover, we study allocation games where interaction graphs have bounded treewidth: Our main result is that both the Shapley value and the Banzhaf value can be computed in polynomial time for this class of games. A key ingredient for the polynomial-time tractability of the problem is a CSP encoding allowing us to exploit the CSP machinery for counting solutions. Note that our restrictions capture scenarios of practical interest, such as the allocation problem arising in the Italian Research Assessment program promoted by ANVUR.
- Process mining. The third direction of research is devoted to the study of workflow management systems that aim at supporting the execution of business processes in various contexts. We design and implement a "hybrid" approach to process discovery. A mining method is conceived which can take into account
a wide variety of constraints over the causal dependencies that are available to the analyst. This is particularly useful for circumventing problems emerging when $\log$ completeness does not hold. We formulate algorithms in terms of reasoning problems over precedence constraints, and we analyze the computational complexity of the proposed setting, by taking into account various qualitative properties regarding the kinds of constraints being allowed, and by tracing the tractability frontier w.r.t. them. In particular, we show that for the classes of constraints that are not covered by our algorithms, an efficient solution algorithm is unlikely to exist at all, because process discovery turns out to be NP-hard over them.


### 1.2 Outline of the dissertation

The rest of the dissertation is organized as follows. In Chapter 2, we discuss the setting of constraint satisfaction problems, including structural decomposition methods, counting problems and preliminaries on computational complexity. In Chapter 3, we briefly review some basic concepts from database theory. In Chapter 4, we present the novel notion of weighted monotone greedy tree projection. We analyze the complexity of computing greedy tree projections in different settings, and we give an algorithm that can be exploited for all problems that can be solved efficiently on acyclic and quasi-acyclic instances. In Chapter 5 we take a first step towards designing an Hybrid optimizer for counting problems. In Chapter 6 we review basic concepts from game theory, and we discuss a real-world application modeled as a coalitional game. In Chapter 7 we identify structural restrictions for the class of allocation games that guarantee the tractability of computing the Shapley value. We also strengthen some hardness result about this problem. In Chapter 8 we intro-
duce the framework of process mining and we discuss process-discovery methods. In Chapter 9, we analyze the complexity of computing business process models under precedence constraints imposed over the topology of the mined model. We chart the frontier of tractability for this problem. We present a CSP encoding for this problem and heuristic algorithms that however solve exactly the problem over some restricted classes. In Chapter 10 we describe a real-world scenario that can be modeled as a mining problem, and hence can benefit of our approaches. Finally, in Chapter 11 we conclude by discussing further research directions.

## CHAPTER 2

## Setup and Preliminaries

In this chapter we review the framework of Constraint satisfaction problems (CSPs) and we explore the structure of these problems from a complexity-theoretic perspective exploiting the concept of decomposition. In particular, we will first present two relevant decomposition methods for (hyper)graph based structures: the treewidth, which is the most powerful decomposition method on graphs, and the hypertree width, which is its natural counter-part over hypergraphs. Then we will show that both methods are specializations of a more general decomposition scheme called tree projection, which we will describe in Section 2.4.3. Finally we will introduce counting problems in Section 2.5. Meanwhile, we need to define the basic concepts that we will use throughout the thesis.

### 2.1 Constraint satisfaction problems

A fundamental problem in the field of Artificial Intelligence and related disciplines, in particular Database theory, is the constraint satisfaction problem (or CSP) which provides a unifying framework to express a wide spectrum of computational problems, such as graph colorability, model checking, planning, satisfiability, scheduling, theorem proving and database queries. Crudely, the goal of a CSP is to assign one value to each variable so that the constraints are all satisfied. While this problem
appears very general and indeed suitable to describe many problems, it is not surprising that the price to be paid for this generality is its high computational complexity. Formally, a CSP instance is a triple $\mathcal{I}=\langle\operatorname{Var}, U, \mathbf{C}\rangle$, where Var is a finite set of variables, $U$ is a finite domain of values, and $\mathbf{C}=\left\{C_{1}, C_{2}, \ldots, C_{q}\right\}$ is a finite set of constraints. Each constraint $C_{v}$, for $1 \leq v \leq q$, is a pair ( $S_{v}, r_{v}$ ), where $S_{v} \subseteq \operatorname{Var}$ is a set of variables called the constraint scope, and $r_{v}$ is a constraint relation, i.e., a set of substitutions $\theta: S_{v} \rightarrow U$ indicating the allowed combinations of simultaneous values for the variables in $S_{v}$. For each variable $X \in \operatorname{Var}$, we denote by dom $(X)$ its domain, i.e., the set of all elements $u \in U$ for which some constraint relation contains a substitution $\theta$ with $\theta(X)=u$. In the following, a substitution from a set of variables $V \subseteq \operatorname{Var}$ to $U$ is transparently viewed as the set of pairs of the form $X / u$, where $\theta(X)=u$ is the value to which $X \in V$ is mapped. A substitution $\theta$ satisfies a constraint $C_{v}$ if its restriction to $S_{v}$, i.e., the set of all pairs $X / u \in \theta$ such that $X \in S_{v}$, occurs in $r_{v}$. A solution to $\mathcal{I}$ is a substitution $\theta: \operatorname{Var} \mapsto U$ that satisfies all constraints. The set of all solutions is denoted by $\Theta(\mathcal{I})$. If $\mathcal{W}$ is a set of variables, then $\Theta(\mathcal{I}, \mathcal{W})$ denotes the set of all solutions in $\Theta(\mathcal{I})$ restricted to the variables in $\mathcal{W}$. Variables outside $\mathcal{W}$ can be viewed as auxiliary ones. That is, they are used for internal encoding activities, and they are not required in the output. The problem is either to find all solutions, one solution, or to verify that a certain substitution $\theta$ is a solution [Dechter and Pearl, 1987]. While the last problem is easy, the first two are NP-hard and in particular the first can be much harder than the second, as we will discuss later in this chapter.

Example 2.1. As an example, consider the map coloring problem (i.e., graph colorability) reported in Figure 2.1. The goal is to assign k-colors to the provinces of Calabria such that no two adjacent provinces have the same color. We can formulate


Table 2.1: The provinces of Calabria and its colored graph
this problem as CSP where we define the variables to be the provinces: $\{C S, C Z$, $K R, V V$ and $R C\}$, the domain of each variable is the set $\{$ red, green, blue $\}$ and the constraints impose that two adjacent provinces must have different colors. A solution to this problem is a combination of simultaneous values for the variables that satisfies all constraints. For the example above, there are many possible solutions, such as $\{C S=$ red, $C Z=$ green, $K R=$ blue, $V V=$ red, $R C=$ blue $\}$.

Observe that, the problem described above is a binary constraint satisfaction problem over a finite domain, since all constraints are imposed between pairs of variables (i.e., the arity of these constraints is just 2). In addition, it is worth noting that any n-ary CSP can be converted to a binary CSP by means of a suitable transformation. Also note that we described the problem in terms of graph and this is not by chance. In fact, with each CSP instance $\mathcal{I}$, it is naturally to associate a constraint (hyper)graph $G(\mathcal{I})$ whose nodes are the variables and where there is a (hyper)edge between any pair of variables appearing within the same constraint scope. Indeed, the constraint graph describes the structure induced by the constraints on the variables and it plays a crucial role in our understanding of the complexity of CSPs. We will come back to this in Section 2.3.

For the sake of completeness we show also an example of n-ary CSP.

Example 2.2. Consider the Boolean Satisfiability Problem (SAT, for short), that is, given a Boolean formula $\phi$ in conjunctive normal form (CNF), is there an assignment of 0 s and 1 s to the variables in input such that the formula returns the value 1 . SAT is a constraint satisfaction problem where the CNF formula $\phi$ corresponds to a CSP instance $\mathcal{I}$ whose variables are the same variables of $\phi$, whose domain is $U=\{0,1\}$ and where each constraint $C_{i} \in \mathbf{C}$ corresponds to a clause $c_{i} \in \phi$. For example, the clause $c_{x y z}=(x \vee \neg y \vee \neg z)$ corresponds to a constraint $\left(S_{x y z}, r_{x y z}\right)$ where $r_{x y z}$ is the ternary relation $\{(0,0,0),(0,0,1),(0,1,0),(1,0,0),(1,0,1),(1,1,0),(1,1,1)\}$ on $\{0,1\}$.

We conclude this section by observing that CSPs can be expressed via Constraint programming, a declarative programming paradigm where users are just in charge of specifying the problem at hand in terms of a constraint satisfaction problem, instead of formalizing the steps needed for its solution. Constraint programming systems exploit, indeed, general search mechanisms (such as backtracking) enriched with powerful speed-up methods (such as constraint propagation techniques) in order to find a solution such that meets declarative specification. We will detail how such mechanisms work in the next section. Finally, it is worth noting that a constraint language is tractable if we can solve the associated CSP in polynomial time.

### 2.2 Solving CSPs

Now that we introduced the basic terminology, we go back to the main issue of CSP and we point out some interesting properties of this problem. When dealing with general CSPs, the major drawback is represented by their high complexity that often calls for heuristic methods to achieve a solution in reasonable time. Therefore,
much effort has been spent by the scientific community in finding an efficient solution for the constraint satisfaction problem. Perhaps, the most important direction is the identification of tractable classes of CSPs and taking advantage of such classes. In fact, we can use the knowledge about solving tractable CSP classes to help the solution of harder problems. The approaches to identify tractable CSP classes can be divided into two main categories [Gottlob et al., 2000]:

- Tractable structural restrictions. This includes all tractable classes of CSPs that are identified only by analyzing the structure of the constraint scopes.
- Tractability due to restricted constraint relations. This includes all classes that are tractable as a consequence of particular properties of the constraint relations.

In this dissertation, we will focus on tractable structural restrictions. But before we start describing the tools that help us address this issue, let us discuss what happens in the general case. It turns out that the most commonly used algorithm for solving CSPs is the backtracking search.

### 2.2.1 Backtracking search

Backtracking search is as simple as systematic. Basically it consists of the following steps. Given a CSP instance $\mathcal{I}=\langle\operatorname{Var}, U, \mathbf{C}\rangle$, it starts in a state where no assignment has been done for the variables in Var. Then it traverses the variables in a predetermined order, it picks a variable, it assigns a value and if the assigment is consistent (i.e., it satisfies all the constraints), it repeats. Eventually, if no consistent assignment can be found for the next variable then the algoritm backtracks the search tree until it finds a variable that can have its value set consistently to something else. Backing all the way up the tree to the root, and finding no more values means that there is no solution. A backtracking algorithm that computes one
solution is given below.

```
Algorithm 1 backtrack_search
Input: \(\mathcal{I}=\langle\operatorname{Var}, U, \mathbf{C}\rangle\) a CSP instance
Output: \(T R U E\) if there exists a solution \(F A L S E\) otherwise.
    \((\mathcal{F}, \mathcal{A}):=(\emptyset, \emptyset) ;\)
    while there are unassigned var do
        \(X_{i}:=\operatorname{pick}(\) Var \() ;\)
        for all \(v \in U\left(X_{i}\right)\) do
            \(\mathcal{A}:=\mathcal{A} \cup\left\{X_{i}=v\right\} ;\)
            if \(\left\{X_{i}=v\right\} \in\) Forbid then
                continue;
            end if
            \(\mathcal{F}:=\mathcal{F} \cup\left\{X_{i}=v\right\} ;\)
            \(\mathcal{A}:=\mathcal{A} \cup\left\{X_{i}=v\right\} ;\)
            if \(\exists c \in \mathbf{C}\) s.t. \(c\) is not satisfied then
                \(\mathcal{A}:=\mathcal{A} \backslash\left\{X_{i}=v\right\} ;\)
            else
                goto 2;
            end if
        end for
        unroll previous choice; \(\quad\) also clear \(\mathcal{F}\)
        if \(\mathcal{A}=\emptyset\) then
            return FALSE
        end if
    end while
    return TRUE
```

When we use this algorithm for finding one solution, it is natural to ask the following questions.

- How do we chose the next variable and its value?
- Can we prune the search space, and thereby search less?
- Can we learn from our mistakes, i.e., bad variable selections?

For instance, the problem of selecting a variable and assigning a value plays a significant role in the execution time. In fact, the easy approach that consists in picking the next variable and next value from a static list is not always the most efficient approach. Indeed is not easy to make the right choice and thus several heuristic proposal have been offered by the AI community. For example we can choose the variable with the smallest number of remaining values in its domain.

Another approach could be the one that chooses the variable that is part of the most remaining unsatisfied constraints. Moreover, once we choose a variable, we select from its domain the value that rules out the fewest choices for neighboring variables. It turns out that Algorithm 1 is not very efficient or intelligent since it does not use all the information available from the constraint set and the assigned variables. In fact, we can use the value assigned to available and the contraints in order to probe the search space, that is we restrict the available future values for the other variables. If we empty the domain of another variable, that is we are unable to assign a consistent value to another variable, then we know that we can backtrack without exploring the remaining subtree below the current variable.

We conclude this section, by pointing out again that there are many situations (such as database query answering), in which we are interested not only in finding one solution but we require all solutions. In this setting, it easy to see that things become computationally much harder since even the simplest constraint satisfaction problem instance may have an exponential number of satisfying solutions. Therefore, it is sensible to propose algorithms that generate solutions with polynomial delay. For the remainder of this thesis, we will consider both these settings.

### 2.3 Structural properties of CSPs

In this section, we change our perspective by focusing on the structure of CSPs with the objective to identify structural properties that lead to tractable classes of CSPs. Indeed, taking advantage of the understading of such classes often increases our abilities in solving harder problems. As we already noted, many problems arising in different areas of computer science can be described via CSPs. Often, such problems are intractable in general, but whenever the input has a desired structure
(for example, the constraint graph is a tree), we can compute solution efficiently via dynamic programming. In graph-based problems, the intractability is frequently due to cyclicity. A natural question to ask is therefore whether the degree of cyclicity affects the complexity of the problem. The answer to this question is "yes": the more the input is tree-like the easier it becomes. In particular, it was recognized in both AI and database theory community that the most important property in the context of CSPs (and conjunctive database queries) is acyclicity [Gottlob et al., 2000].Therefore, multiple effort has been made in finding a way to measure the degree of cyclicity of graphs (and hypergraphs) with the objective to identify nearly-acyclic instances resulting in the definition of a broad set of structural decomposition methods. Decomposition techniques work as follows: First we associate to each instance problem $\mathcal{I}$ the graph $G$ (or a hypergraph $\mathcal{H}$ ). Then, clusters of vertices are arranged in some tree-shape in order to obtain an acyclic data structure. Finally we can formalize the concept of degree of cyclicity using the width of the decomposition (of the instance problem $\mathcal{I}$ ) that corresponds to the size of the largest cluster of vertices occurring in the decomposition. The treewidth, which we will review more technically in the next section, is currently the best way to measure the degree of cyclicity of a graph since it subsumes all the other measures.

### 2.4 Structural decomposition methods

### 2.4.1 Tree decompositions

A tree decomposition [Robertson and Seymour, 1984] of a graph $G$ is a pair $\langle T, \chi\rangle$, where $T=(N, E)$ is a tree, and $\chi$ is a labeling function assigning to each vertex $v \in N$ a set of vertices $\chi(v) \subseteq \operatorname{nodes}(G)$, such that the following conditions are satisfied:


Figure 2.1: A hypergraph $\mathcal{H}$ and a tree decomposition of $\mathcal{H}$ of width 3

1. for each node $Y \in \operatorname{nodes}(G)$, there exists $p \in N$ such that $Y \in \chi(p)$;
2. for each edge $\{X, Y\} \in \operatorname{edges}(G)$, there exists $p \in N$ such that $\{X, Y\} \subseteq \chi(p)$;
3. for each node $Y \in \operatorname{nodes}(G)$, the set $\{p \in N \mid Y \in \chi(p)\}$ induces a (connected) subtree of $T$.

The width of $\langle T, \chi\rangle$ is the number $\max _{p \in N}(|\chi(p)|-1)$.
In particular, the treewidth of $\mathcal{G}$ is the minimum width over all the tree decompositions that can be constructed. Deciding whether a given graph has treewidth bounded by a fixed natural number $k$ is known to be feasible in linear time [Bodlaender and Fomin, 1996]. Moreover, the treewidth of a hypergraph $\mathcal{H}$ is 1 if $\mathcal{H}$ is acyclic, and it is equal to the treewidth of its primal graph otherwise. As an illustrative example, consider the Figure 2.1

### 2.4.2 (Generalized) Hypertree decompositions

Unfortunately, there exist a number of real-world problems where very complex relationships between objects may arise, leading to loss of information whenever these relationships are represented as graphs. In particular, while binary relationships can naturally be illustrated with simple graphs, general relationships are best described
by hypergraphs. Since hypergraphs generalize graphs, we can define the degree of cyclicity of a hypergraph and we can carry out investigations on its properties.

A hypergraph $\mathcal{H}$ is a pair $(V, H)$, where $V$ is a finite set of nodes and $H$ is a set of hyperedges such that, for each $h \in H, h \subseteq V$. In the following, we denote $V$ and $H$ by $\operatorname{nodes}(\mathcal{H})$ and edges $(\mathcal{H})$, respectively.

A hypergraph $\mathcal{H}$ is acyclic (more precisely, $\alpha$-acyclic [Fagin, 1983]) if, and only if, it has a join tree [Bernstein and Goodman, 1981], i.e., a tree whose vertices are the hyperedges of $\mathcal{H}$ such that, whenever a node $X \in V$ occurs in two hyperedges $h_{1}$ and $h_{2}$ of $\mathcal{H}$, then $h_{1}$ and $h_{2}$ are connected in $J$, and $X$ occurs in each vertex on the unique path linking $h_{1}$ and $h_{2}$. In words, the set of vertices in which $X$ occurs induces a connected subtree of $J T$. Similar to the notion of tree decomposition, there exists the (generalized) hypertree decomposition over hypergraphs which, in turn, leads to the associated notion of hypertree-width.

A hypertree for a hypergraph $\mathcal{H}$ is a triple $\langle T, \chi, \lambda\rangle$, where $T=(N, E)$ is a rooted tree, and $\chi$ and $\lambda$ are labeling functions that associate each vertex $v \in N$ with two sets $\chi(v) \subseteq \operatorname{nodes}(\mathcal{H})$ and $\lambda(v) \subseteq \operatorname{edges}(\mathcal{H})$.

The width of a hypertree is the size of its largest $\lambda$, i.e., $\max _{p \in T}(|\lambda(p)|)$. We denote the set of vertices $N$ of $T$ by vertices $(T)$, and the root of $T$ by $\operatorname{root}(T)$. Moreover, for any $p \in N, T_{p}$ denotes the subtree of $T$ rooted at $p$. If $T^{\prime}$ is a subtree of $T$, we define $\chi\left(T^{\prime}\right)=\bigcup_{v \in \text { vertices }\left(T^{\prime}\right)} \chi(v)$.

A generalized hyertree decomposition of a hypergraph $\mathcal{H}$ is a hypertree $H D=$ $\langle T, \chi, \lambda\rangle$ such that: (1) for each hyperedge $h \in \operatorname{edges}(\mathcal{H})$, there exists $p \in \operatorname{vertices}(T)$ such that $\operatorname{vars}(h) \in \chi(p) ;(2)$ for each node $Y \in \operatorname{nodes}(\mathcal{H})$, the set $\{p \in \operatorname{vertices}(T) \mid$ $Y \in \chi(p)\}$ induces a (connected) subtree of $T$; (3) and finally, for each node $p \in$ $\operatorname{vertices}(T), \chi(p) \subseteq \operatorname{nodes}(\lambda(p))$


Figure 2.2: A hypergraph $\mathcal{H}$ and a hypertree decomposition of $\mathcal{H}$ of width 2

In particular, the hypertree width of $\mathcal{H}$ is the minimum width over all the hypertree decompositions that can be constructed.

The hypertree decomposition $H D$ is a complete decomposition if for each hyperedge $h \in \operatorname{edges}(\mathcal{H})$, there exists $p \in \operatorname{vertices}(T)$ such that $\operatorname{vars}(h) \in \chi(p)$ and $h \in \lambda(p)$. Moreover, a generalized hypertree decomposition is a hypertree decomposition if it satisfies the following special condition: for each node $p \in$ $\operatorname{vertices}(T), \operatorname{nodes}(\lambda(p)) \cap \chi\left(T_{p}\right) \subseteq \chi(p)$.

Observe that, deciding whether a given hypergraph has hypertree width bounded by a fixed natural number $k$ is feasible in polynomial time (and, actually it is highlyparallelizable) [Gottlob et al., 2002]. Moreover, it is known that for any hypergraph $\mathcal{H}$ the generalized hypertree width $\operatorname{ghw}(\mathcal{H})$ is always smaller than or equal to the hypertree width $h w(\mathcal{H})$, and more precisely, $g h w(\mathcal{H}) \leq h w(\mathcal{H}) \leq g h w(\mathcal{H})+1[$ Adler et al., 2007].

Deciding whether a given hypergraph has generalized hypertree width bounded by a fixed natural number $k$ is NP-complete [Gottlob et al., 2009].

Figure 5.1 shows an example of hypertree decomposition.

### 2.4.3 Tree projections

As noted before, these methods have in common the idea of transforming a given cyclic hypergraph (graph) into an acyclic one, by organizing its vertices (or its edges) into a polynomial number of clusters and by arranging these clusters into a treeshape structure - a so-called decomposition tree. Given the decomposition tree, we can then evaluate the original problem with a cost that is exponential in the width of the decomposition (i.e., the cardinality of the largest cluster) but polynomial in the combined size of the input and the output. Consequently, great effort has been spent by the scientific community in finding the best method to identify nearly acyclic graphs (hypergraphs) resulting in the definition of various methods that regardless of the technical differences, are just a special case of a general framework called tree projections. The result is the following formal definition for this general framework. For a given pair of hypergraphs $\left(\mathcal{H}_{1}, \mathcal{H}_{2}\right)$, a tree projection of $\mathcal{H}_{1}$ w.r.t $\mathcal{H}_{2}$ is an acyclic hypergraph $\mathcal{H}_{a}$ such that each hyperedge of $\mathcal{H}_{1}$ is contained in some hyperedge of $\mathcal{H}_{a}$, that is in turn contained in some hyperedge of $\mathcal{H}_{2} . \mathcal{H}_{2}$ is called the resource hypergraph and it is arbitrary. Hence, whenever we compute $\mathcal{H}_{2}$ with some specific method from the hypergraph $\mathcal{H}_{1}$, we obtain, as special cases, some known purely structural decomposition methods. In fact, observe that the decomposition method based on the submodular width [Dániel Marx, 2010], which is not purely structural, is the only one among all the structural decomposition methods that does not fit this framework.

Let $\mathcal{H}_{1}, \mathcal{H}_{2}$ be two hypergraphs. We assume w.l.o.g. that $\operatorname{nodes}\left(\mathcal{H}_{1}\right)=\operatorname{nodes}\left(\mathcal{H}_{2}\right)$. Given $\left(\mathcal{H}_{1}, \mathcal{H}_{2}\right)$, we say that $\mathcal{H}_{1} \leq \mathcal{H}_{2}$, iff each hyperedge of $\mathcal{H}_{1}$ is contained in at least one hyperedge of $\mathcal{H}_{2}$. Let $\mathcal{H}_{1} \leq \mathcal{H}_{2}$; then, a tree projection of $\mathcal{H}_{1}$ w.r.t $\mathcal{H}_{2}$ is an acyclic hypergraph $\mathcal{H}_{a}$ such that $\mathcal{H}_{1} \leq \mathcal{H}_{a} \leq \mathcal{H}_{2}$. We say that $\mathcal{H}_{a}$ is a tree projection


Figure 2.3: A tree projection $\mathcal{H}_{a}$ for the pair of hypergraphs $\left(\mathcal{H}_{1}, \mathcal{H}_{2}\right)$.
for the pair $\left(\mathcal{H}_{1}, \mathcal{H}_{2}\right)$ if it exists.
In fact, by slightly abusing notation, we use the term tree projection to refer to a join tree of such a sandwich acyclic hypergraph or, equivalently, a width-1 hypertree decomposition of $\mathcal{H}_{a}$. More precisely, we say that a tree projection of $\mathcal{H}_{1}$ w.r.t $\mathcal{H}_{2}$ is a triple $H D=\langle T, \chi, \lambda\rangle$ where $T=(N, E)$ is a rooted tree; $\chi$ and $\lambda$ are labeling functions that associate each vertex $v \in N$ with two sets of vertices $\chi(v) \subseteq \operatorname{nodes}\left(\mathcal{H}_{1}\right)$ and $\lambda(v) \in \operatorname{edges}\left(\mathcal{H}_{2}\right)$; and such that the following conditions hold: (1) for each hyperedge $h \in \operatorname{edges}\left(\mathcal{H}_{1}\right)$, there exists $v_{h} \in N$ with $h \subseteq \chi\left(v_{h}\right) ;(2)$ for each node $v \in \operatorname{nodes}\left(\mathcal{H}_{1}\right)$, the set $\left\{v_{h} \in N \mid v \in \chi\left(v_{h}\right)\right\}$ induces a (connected) subtree of $T$;

Throughout the dissertation we assume that $\mathcal{I}$ is a given CSP instance, and we shall seek to compute its solutions (possibly restricted to a desired set of output variables) by combining the solutions of suitable sets of subproblems, available as additional distinguished constraints called resource views. Intuitively, each original constraint of $\mathcal{I}$ is associated with a new distinguished hyperedge in the resource hypergraph. Importantly, note that a resource view in order to be a legal subproblem must be not more restrictive than the full problem. Moreover, any resource view has to be at least as restrictive as the original constraint associated with it.

Example 2.3. Consider the hypergraph $\mathcal{H}_{1}$ on the left of Figure 2.3 together with other two hypergraphs that are discussed next. $\mathcal{H}_{1}$ represents the CSP instance $\mathcal{I}$ whose variables are given by $\operatorname{nodes}(\mathcal{H})$ and where each constraint $C_{i} \in \mathbf{C}$ corresponds to a hyperedge appearing in edges $(\mathcal{H})$. For example, the constraint set $\mathbf{C}$ is given by

$$
\begin{aligned}
\mathbf{C}:= & \{\langle A, B, C\rangle,\langle A, F\rangle,\langle C, D\rangle,\langle D, E, F\rangle, \\
& \langle E, F, G\rangle,\langle G, H, I\rangle,\langle I, J\rangle,\langle J, K\rangle\} .
\end{aligned}
$$

To solve $\mathcal{I}$, assume that a set of resource views is available, playing the role of additional distinguished constraints. The set of variables of each resource view is a hyperedge in the hypergraph $\mathcal{H}_{2}$. In the middle between $\mathcal{H}_{1}$ and $\mathcal{H}_{2}$, Figure 2.3 it is reported the hypergraph $\mathcal{H}_{a}$ which covers $\mathcal{H}_{1}$, and which is in its turn covered by $\mathcal{H}_{2}$ e.g., $\{C, D\} \subseteq\{A, B, C, D\} \subseteq\{A, B, C, D, H\}$. Since $\mathcal{H}_{a}$ is in addition acyclic, $\mathcal{H}_{a}$ is a tree projection of $\mathcal{H}_{1}$ w.r.t. $\mathcal{H}_{2}$.

Note that all the (known) structural decomposition methods can be recast as special cases of tree projections, since they just differ in how they define the set of views to be built for evaluating the CSP instance. For example, consider the generalized hypertree decomposition approach. Given a hypergraph $\mathcal{H}$ and a natural number $k>0$, let $\mathcal{H}^{k}$ denote the hypergraph over the same set of nodes as $\mathcal{H}$, and whose set of hyperedges is given by all possible unions of $k$ edges in $\mathcal{H}$, i.e., $\operatorname{edges}\left(\mathcal{H}^{k}\right)=\left\{h_{1} \cup h_{2} \cup \cdots \cup h_{k} \mid\left\{h_{1}, h_{2}, \ldots, h_{k}\right\} \subseteq \operatorname{edges}(\mathcal{H})\right\}$. Then, it is well known and easy to see that $\mathcal{H}$ has generalized hypertree width at most $k$ if, and only if, there is a tree projection for $\left(\mathcal{H}, \mathcal{H}^{k}\right)$.

Similarly, for tree decompositions, let $\mathcal{H}^{t k}$ be the hypergraph over the same set of nodes as $\mathcal{H}$, and whose set of hyperedges is given by all possible clusters $B \subseteq$ nodes $(\mathcal{H})$ of nodes such that $|B| \leq k+1$. Then, $\mathcal{H}$ has treewidth at most $k$ if, and
only if, there is a tree projection for $\left(\mathcal{H}, \mathcal{H}^{t k}\right)$.

### 2.5 Counting problems

So far, we discussed how to solve CSPs exploiting its structural properties. However, in real-world applications like database query answering, often we do not concentrate on finding one solution or just counting the number of solutions, possibly focusing on a restricted set of output variables (as we will discuss in Chapter 5). For example, the "select" part of an SQL query allows users to specify a set of output variables, so that query answering amounts at enumerating all solutions projected over those variables, rather than just deciding whether there is any. In [Atserias et al., 2013], it has been shown that the enumeration problem where all variables are output variables (i.e., "SELECT *" queries where no variable is projected out) can be solved in polynomial time on a class $\mathbf{C}$ of queries if, and only if, the number of solutions is always polynomial, and that this is the case, if and only if, the queries in $\mathbf{C}$ have bounded fractional edge cover number. Similar tight worst-case bounds for conjunctive queries with arbitrary sets of output variables have been derived in [Gottlob et al., 2012]. Note however that even easy instances may have an exponential number of solutions. Therefore, for enumeration problems it is sensible to propose algorithms computing solutions with polynomial delay, or even with linear delay [Bagan et al., 2007]. Again, acyclicity has been shown to be a key for tractability in this setting [Yannakakis, 1981; Koch, 2006; Kimelfeld and Sagiv, 2006; Bagan et al., 2007], and generalizing these results to decomposition methods is an active area of research (see, e.g., [Bulatov et al., 2012; Greco and Scarcello, 2013a]). Closely related to the enumeration problem is the problem of counting the solutions of constraint satisfaction problems (see, e.g, [Bulatov et al., 2013; Bulatov, 2013; Gomes
et al., 2007; Pesant, 2005]). This problem is equivalent to counting the number of query answers in the related field of database theory, which occurs in the presence of SQL queries specifying "COUNT" aggregates. Note that these queries are very often at the basis of decision support systems examining large volumes of data in order to get insights on critical business questions. In this thesis we analyze the well known TPC-DS, which is a decision support benchmark, ${ }^{1}$ consisting of business-oriented complex queries (many of them involve counting) over which structural decomposition methods can exhibit their effectiveness. Indeed, recently structural tractability results have been generalized to counting problems, as summarized below.

Theorem 2.4 (cf. [Pichler and Skritek, 2013; Greco and Scarcello, 2014b]). Counting the number of substitutions in $\Theta(\mathcal{I}, \mathcal{W})$ is feasible in polynomial time, on classes of CSP instances $\mathcal{I}$ such that the treewidth of $G(\mathcal{I})$ is bounded by a constant, and the size of the domain of each variable not in $\mathcal{W}$ is bounded by some constant, too.

Note that, differently from the case of the standard decision and computation problems, the result is established under the additional condition that auxiliary variables have a bounded domain. If the condition is not met, then \#P-complete instances can be exhibited Pichler and Skritek [2013].

Example 2.5. As a tiny example, consider the following counting problem. We want to count the number of PhD students that attended some Summer School about Constraint Satisfaction. We can express this problem as a CSP: Consider the CSP instance $\mathcal{I}$ whose variables are the set $\{$ Name, Af filiation, School, Year, Subject, Location $\}$, whose domain is $U=\{\ldots\}$ and where the constraint set $\mathbf{C}$ is as follows.

$$
\begin{gathered}
\{\text { student }\langle\text { Name, Affiliation }\rangle \text {, attendend }\langle\text { Name, School }\rangle \text {, } \\
\text { school }\langle\text { School, Year, Subject, Location }\rangle, \text { equal }\langle\text { Subject, " CS" }\rangle\} \text {. }
\end{gathered}
$$

[^0]The intended meaning is that we want to count over all solutions of this CSP the number of distinct occurrences of any student for which there exists some school in some year that she attended and that was about Constraint Satisfaction. Note that, we have just one output variable that is the student's name. All other variables are just auxiliary (i.e, existential variables in the logic-based syntax). The objective is to compute just the number of (desired) solutions. In fact, in counting problems we may get wrong numbers if such variables are not properly considered. In this example, we could get as a result the number 8 that may come from the one student attending 3 summer school on constraint satisfaction, 3 on Game Theory and 2 on Mining, or by 8 distinct students on constraint satisfaction. The counting problem is quite different from the enumeration problem where even if the complexity may change, we still have all the solutions and thus we can get the right number by just projecting out the useless variables.

### 2.5.1 The Langford problem



Figure 2.4: Langford problem

Langford's problem goes back to the Scottish mathematician C. Dudley Langford who once observed his young boy stacking colored blocks. While he was watching, he observed that his son piled six colored blocks in a such a way that there was one block
between the blue pair, two blocks between the red pair and three blocks between the green pair. This is exemplified on the left of Figure 2.5.1. Then he added a pair of yellow blocks and he came up with a solution like the one depicted in the right of Figure 2.5.1. It turns out that, these solutions are unique for 3 and 4 pairs of colored blocks. One could always reverse the order and cycling colors but this is irrelevant. Indeed, the number of distinct solutions is just one. Then, Langford tried to find a solution for 5 and 6 pairs of blocks but he did not succeed. Later he was able to find a solution for 7 and 8 pairs but again he was unable to find a solution for 9 and 10 pairs. Therefore, it was natural for him to investigate more and he came up with the following question. What values of $n$ admit a solution? [Davies, 1959] gave an answer to that question proving that for $5,6,9$ and 10 there not exists a solution. In particular he proved that it is possible to find a solution only if $n$ is a multiple of four or one less than a multiple of four. The study of this problem is fascinating in that for small numbers of blocks, we are able to find a solution even without using a computer. Indeed, we have done that for $n=3$ and $n=4$ in Figure 2.5.1. Therefore, we wonder how many solutions exist for higher values of $n$. While finding one solution for any number of blocks is quite easy, counting all solutions becomes pretty hard. Just to give an idea, the exact number of distinct solutions for $n \geq 28$ is not known as of today. Moreover, $n=27$ was determined only in 2015. Perhaps, we want to look at the numbers in Table 2.5.1 in order to realize how difficult are these kinds of problems.

| $n$ | \# of solutions |
| :--- | :--- |
| 1 | - |
| 2 | - |
| 3 | 1 |
| 4 | 1 |
| 5 | - |
| 6 | - |
| 7 | 26 |
| 8 | 150 |
| 9 | - |
| 10 | - |
| 11 | 17,792 |
| 12 | 108,144 |
| 13 | - |
| 14 | - |
| 15 | $39,809,640$ |
| 16 | $326,721,800$ |
| 17 | - |
| 18 | - |
| 19 | $256,814,891,280$ |
| 20 | $2,636,337,861,200$ |
| 21 | - |
| 21 | - |
| 23 | $3,799,455,942,515,488$ |
| 24 | $46,845,158,056,515,936$ |
| 25 | - |
| 26 | - |
| 27 | $111,683,606,778,027,803,456$ |
| 28 | $?$ |
| 29 | - |

Table 2.2:
Known solutions for the Langford problem. A dash (-) means that there is no solution for the given value of $n$ while (?) stands for unknown.

### 2.5.2 Configuration problems: Renault example

As an interesting real-world example of counting problem, we describe the Renault Megane configuration problem [Amilhastre et al., 2002] used in CSP competitions as a benchmark for CSP solvers. The problem consists of 108 variables that encodes informations about the Renault megane such as the type of engine, country, options like air cooling, etc. with domains ranging from 2 to 43 and with a total of 858 constraints which can be compressed to 149 constraints. Furthermore, many of them are non binary indeed arity of constraint scopes is up to 10 . Finally, the constraint


Figure 2.5:
A portion of hypertree decomposition of width 4 for the Renault megane configuration problem
relations comprise about 200000 tuples. We analyzed this problem by applying our decomposition tools and despite the big numbers in play, we discovered that the hypertree-width associated with it is just 4 as we shall see in more detail in Chapter 4. Moreover, we observed that it is possible to find one solution for this problem very quickly. However, if our objective is to find all the solutions, the problem becomes much harder to solve. In Figure 2.5 we show a portion of the decomposition tree for this problem. The total number of solutions known for this problem is $2.84 * 10^{12}$.

We conclude this chapter by giving some preliminaries on complexity theory.

### 2.6 Preliminaries on Complexity Theory

We recall some notions of complexity theory which we will use later in this thesis, by referring the reader to the book by [Garey and Johnson, 1979] for more on this subject.

Decision problems are maps from strings (encoding the input instance over a fixed alphabet, e.g., the binary alphabet $\{0,1\}$ ) to the set $\{$ "yes","no" $\}$. We are often interested in computations carried out by non-deterministic Turing machines. We recall that these are Turing machines that, at some points of the computation, may not have one single next action to perform, but a choice between several possible next actions. A non-deterministic Turing machine answers a decision problem if, on any input $x,(i)$ there is at least one sequence of choices leading to halt in an accepting state if $x$ is a "yes" instance (such a sequence is called accepting computation path); and (ii) all possible sequences of choices lead to a rejecting state if $x$ is a "no" instance. The class of decision problems that can be solved by non-deterministic Turing machines in polynomial time is denoted by NP.

A decision problem $A_{1}$ is polynomially reducible to a decision problem $A_{2}$, if there is a polynomial-time computable function $h$ (called reduction) such that, for every $x$, $h(x)$ is defined and $x$ is a "yes" instance of $A_{1}$ if and only if $h(x)$ is a "yes" instance of $A_{2}$. A decision problem $A$ is NP-hard if every problem in NP is polynomially reducible to $A$; if $A$ is NP-hard and belongs to NP, then $A$ is said to be NP-complete. Thus, problems that are complete for NP are the most difficult problems in NP. In particular, it is unlikely that an NP-complete problem can be solved in polynomial time because it would entail $\mathrm{P}=\mathrm{NP}$.

## Part I

## CSP and Databases

## CHAPTER 3

## Basic Definitions from Database Theory

In this chapter we will discuss the connection between constraint satisfaction problems and database theory. Various important problems in the field of database theory, such as the problem of evaluating Boolean Conjunctive Queries over a relational database, or the problem of conjunctive query containment are equivalent to solving constraint satisfaction problems as pointed out by several authors, e.g., [Gyssens et al., 1994; Kolaitis and Vardi, 2000].

We will start with the required background of database theory by examining relational structures, relational databases, the concept of conjunctive query and SQL queries. The notions defined in this chapter will be intensely used in the next chapters.

### 3.1 Relational structures

It is well known that constraint satisfaction problems can be recast as homomorphism problems for relational structures. Let $\mathcal{U}$ and $\mathcal{X}$ be disjoint infinite sets that we call the universe of constants and the universe of variables, respectively. A (relational) vocabulary $\tau$ (also called schema in the field of databases) consists of a name $r$ and a finite set of relation symbols of specified finite arities. A relational structure $\mathcal{A}$ over $\tau$ (short: $\tau$-structure) is an instance of that (relational) schema with
a specific domain. More formally, it consists of a universe $A \subseteq \mathcal{U} \cup \mathcal{X}$ and, for each relation symbol $r$ in $\tau$, of a relation $r^{\mathcal{A}} \subseteq A^{\rho}$, where $\rho$ is the arity of $r$. For instance, directed and undirected graphs are relational structures where the vocabulary has a single binary relation $E$, which for digraphs corresponds to the edge relation, while for graphs includes the pairs $(x, y)$ and $(y, x)$ for each edge $\{x, y\}$. In the same way as for graphs, we can define homomorphisms of structures. Let $\mathcal{A}$ and $\mathcal{B}$ be two $\tau$-structures with universes $A$ and $B$, respectively. A homomorphism from $\mathcal{A}$ to $\mathcal{B}$ is a mapping $h: A \mapsto B$ such that $h(c)=c$ for each constant $c$ in $A \cap \mathcal{U}$, and such that, for each relation symbol $r$ in $\tau$ and for each tuple $\left\langle a_{1}, \ldots, a_{\rho}\right\rangle \in r^{\mathcal{A}}$, it holds that $\left\langle h\left(a_{1}\right), \ldots, h\left(a_{\rho}\right)\right\rangle \in r^{\mathcal{B}}$. For any mapping $h$ (not necessarily a homomorphism), $h\left(\left\langle a_{1}, \ldots, a_{\rho}\right\rangle\right)$ is used, as usual, as a shorthand for $\left\langle h\left(a_{1}\right), \ldots, h\left(a_{\rho}\right)\right\rangle$.

A $\tau$-structure $\mathcal{A}$ is a substructure of a $\tau$-structure $\mathcal{B}$ if $A \subseteq B$ and $r^{\mathcal{A}} \subseteq r^{\mathcal{B}}$, for each relation symbol $r$ in $\tau$.

### 3.2 Relational databases

Let $\tau$ be a given vocabulary. A database instance (or, simply, a database) D over $D \subseteq \mathcal{U}$ is a finite $\tau$-structure D whose universe is the set $D$ of constants. For each relation symbol $r$ in $\tau, r^{\mathrm{D}}$ is a relation instance (or, simply, relation) of D .

A database schema DS is a finite set $\left\langle R_{1}, \ldots, R_{m}\right\rangle$ ) of relation schemas.
The elements of relations are called tuples. In this thesis, we will adopt the logical representation of a database [Ullman, 1989; Abiteboul et al., 1995], where a tuple $\left\langle a_{1}, \ldots, a_{\rho}\right\rangle$ of values from $D$ belonging to the $\rho$-ary relation (over symbol) $r$ is identified with the ground atom $r\left(a_{1}, \ldots, a_{\rho}\right)$. Accordingly, a database D can be viewed as a set of ground atoms.

### 3.3 Conjunctive queries

Among all the classes of relational database queries, the most thoroughly analyzed and probably the most important one is the class of Conjunctive queries.

A Conjunctive query $Q$ is a first order formula $\exists \bar{X} \Phi$ where $\Phi=r_{1}\left(\mathbf{u}_{\mathbf{1}}\right) \wedge \ldots \wedge r_{m}\left(\mathbf{u}_{\mathbf{m}}\right)$ is a conjunction of atoms (relation names), $r_{1}, \ldots, r_{m}$ (with $m>0$ ) are relation symbols, $\mathbf{u}_{\mathbf{1}}, \ldots, \mathbf{u}_{\mathbf{m}}$ are lists of terms (i.e., variables or constants), and $\bar{X}=X_{1}, \ldots, X_{n}$ is the list of quantified variables of $\Phi$. We say that the query $Q$ is simple if every atom is defined over a distinct relation symbol. The conjunction $\Phi$ is denoted by form $(Q)$, while the sets of all atoms in $Q$ is denoted by $\operatorname{atoms}(Q)$. For any set $A$ of atoms, $\operatorname{vars}(A)$ is the set of all variables in $A$, and $\operatorname{vars}(Q)$ is used for short in place of $\operatorname{vars}(\operatorname{atoms}(Q))$. Moreover, we define $\operatorname{free}(Q)=\operatorname{vars}(Q) \backslash\left\{X_{1}, \ldots, X_{n}\right\}$ as the set of the free variables in $Q$.

For a database D over $D, Q^{\text {D }}$ denotes the set of all substitutions $\theta: \operatorname{vars}(Q) \mapsto D$ such that for each $i \in\{1, \ldots, m\}, \theta^{\prime}\left(r_{\alpha_{i}}\left(\mathbf{u}_{\mathbf{i}}\right)\right) \in \mathrm{D}$, where $\theta^{\prime}(t)=\theta(t)$ if $t \in \operatorname{vars}(Q)$ and $\theta^{\prime}(t)=t$ otherwise (i.e., if the term $t$ is a constant).

Note that the conjunction form $(Q)$ can be viewed as a relational structure $\mathcal{Q}$, whose vocabulary $\tau_{Q}$ and universe $U_{Q}$ are the set of relation symbols and the set of terms occurring in its atoms, respectively. For each symbol $r_{i} \in \tau_{Q}$, the relation $r_{i}^{\mathcal{Q}}$ contains a tuple of terms $\mathbf{u}$, for any atom of the form $r_{i}(\mathbf{u}) \in \operatorname{atoms}(Q)$ defined over $r_{i}$. In the special case of simple queries, every relation $r_{i}^{\mathcal{Q}}$ of $\mathcal{Q}$ contains just one tuple of terms. According to this view, elements in $Q^{\mathrm{D}}$ are in a one-to-one correspondence with homomorphisms from $\mathcal{Q}$ to $\mathrm{D}_{Q}$, where the latter is the (maximal) substructure of D over the (sub)vocabulary $\tau_{Q}$. Hereinafter, for the sake of presentation, we freely use interchangeably queries and databases with their relational structures, e.g., we
may use $Q$ and D in place of $\mathcal{Q}$ and $\mathrm{D}_{Q}$. Moreover, $\|Q\|$ and $\|\mathrm{D}\|$ denote the sizes of the underlying relation structures, according to their standard encoding (see, e.g., [Grohe, 2007]).

According to the above notation, $\pi_{\text {free }(Q)}\left(Q^{\mathrm{D}}\right)$ is the set of answers of the conjunctive query $Q$ on the database D (over the output variables free $(Q)$ ). We denote by count $(Q, \mathrm{D})$ the problem of computing the cardinality of this set, i.e., the number of (distinct) answers of $Q$ on D.

### 3.4 Relational algebra.

For any set $W \subseteq \operatorname{vars}(Q)$ of variables and set $S$ of substitutions, we denote by $\pi_{W}(S)$ the set of the restrictions of the substitutions in $S$ over the variables in $W$. If $\theta$ is a substitution with domain $W$, then we denote by $\sigma_{\theta}(S)$ the set $\left\{\theta^{\prime} \in S \mid \pi_{W}\left(\left\{\theta^{\prime}\right\}\right)=\{\theta\}\right)$. If $S_{1}$ and $S_{2}$ are sets of substitutions with domains $W_{1}$ and $W_{2}$, respectively, then we denote by $S_{1} \bowtie S_{2}$ the set of all substitutions $\theta$ over $W_{1} \cup W_{2}$ such that $\pi_{W_{1}}(\{\theta\}) \subseteq S_{1}$ and $\pi_{W_{2}}(\{\theta\}) \subseteq S_{2}$. We use $S_{1} \ltimes S_{2}$ as a shorthand for $\pi_{W_{1}}\left(S_{1} \bowtie S_{2}\right)$.

### 3.5 Hypergraphs and structural restrictions

There is a very natural way to associate a hypergraph $\mathcal{H}_{\mathcal{V}}=(N, H)$ with any set $\mathcal{V}$ of atoms: the set $N$ of nodes consists of all variables occurring in $\mathcal{V}$; for each atom in $\mathcal{V}$, the set $H$ of hyperedges contains a hyperedge including all its variables; and no other hyperedge is in $H$. For a query $Q$, the hypergraph associated with atoms $(Q)$ is denoted by $\mathcal{H}_{Q}$.

Let $Q$ be a conjunctive query, and let $\mathcal{V}$ be a set of atoms each one defined over a specific relation symbol not occurring in $Q$. These atoms play the role of additional resources that can be used to answer (counting problems on) $Q$, so that
we are abstracting here all purely structural decomposition methods proposed in the literature, which just differ in how they define and build the set of such available resources. Following the framework in [Greco and Scarcello, 2010], we say that $\mathcal{V}$ is a view set for $Q$ if, for each atom $q \in \operatorname{atoms}(Q), \mathcal{V}$ contains an atom $w_{q}$ with the same list of variables as $q$ (but with a different relation symbol). Each atom in $\mathcal{V}$ is called a view; in particular, atoms of the form $w_{q}$ are called query views. The set of all query views will be denoted by views $(Q)$.

Let D be a database whose vocabulary includes all relation symbols in $Q$ and $\mathcal{V}$. In a structural decomposition method, query views are initialized with the same tuples as their associated query atoms while, for any other view $w \in \mathcal{V}, w^{\text {D }}$ is initialized by including solutions of some subquery over the variables in $w$. Using such views for evaluating $Q$ is possible whenever D is a legal database (on $\mathcal{V}$ w.r.t. $Q$ ), i.e., (i) $w_{q}^{\mathrm{D}} \subseteq q^{\mathrm{D}}$ holds, for each query view $w_{q} \in \operatorname{views}(Q)$; and (ii) $w^{\mathrm{D}} \supseteq \pi_{v a r s(w)}\left(Q^{\mathrm{D}}\right)$, for each view $w \in \mathcal{V}$. Intuitively, all original "constraints" are there, and views are not more restrictive than the original query.

In this thesis, we look for "structural" restrictions to $Q$ guaranteeing that the problem count $(Q, \mathrm{D})$ can be efficiently solved by exploiting the views in $\mathcal{V}$.

### 3.6 SQL queries

Structured Query Language ( $S Q L$ ) is a declarative language used to define, query and update data in a relational database. In SQL, which is without doubts the most widely accepted database programming language, the user is in charge to specify what he wants rather than how to do it. Indeed, the complex problem of generating an efficient query plan is left to the query optimizer of the DBMS, which will decide how to process and retrieve the data. Not only SQL is interesting due to its wide use,
but also because it is well known that SQL queries without nested statements and inequality comparisons (i.e., Select/Project/Join) are equivalent in expressive power to conjunctive queries. Therefore it is of great interest to propose efficient solutions for SQL query execution plan.

Select/Project/Join queries are generated using the following SQL statement.

$$
\begin{aligned}
& \text { SELECT } R_{1}\left(\mathbf{u}_{\mathbf{1}}\right), \ldots, R_{m}\left(\mathbf{u}_{\mathbf{m}}\right) \\
& \text { FROM } R_{1}, \ldots, R_{m} \\
& \text { WHERE } \quad R_{i}\left(\mathbf{u}_{\mathbf{k}}\right)=R_{j}\left(\mathbf{u}_{\mathbf{z}}\right)
\end{aligned}
$$

We next focus on this kind of SQL queries and in additional functionality that includes aggregating data (e.g., counting). We will discuss this in more detail in Chapter 5.

## CHAPTER 4

## A Weighted Structural Decomposition Technique

### 4.1 Introduction

In this chapter we move a further step towards the understanding of the structural tractability of constraint satisfaction problems by focusing on a general framework based on the notion of tree projection, which generalizes all known decomposition methods. Indeed we have seen in Chapter 2 that the tree projections unify all purely structural decomposition methods proposed in the literature. However, it is well known that, computing general tree projections is intractable, formally NPhard [Gottlob et al., 2009]. This could in principle result in a limited applicability of the framework. Motivated by the bad news, in [Greco and Scarcello, 2010] a relevant tractable class of tree projections is identified, known as the class of greedy tree projections. In this chapter we design and implement an algorithm for a restricted monotone variant, called monotone greedy tree projection. Up to now, our algorithm, known as GreedyTreeProjection ${ }_{D M}$ (in short $G T P_{D M}$ ) is the only polynomial-time algorithm for constructing such decompositions. We note however, that in many applications, finding a purely structural decomposition is not the best we can do. For instance, when dealing with CSP evaluation we can think of minimizing the number of vertices of the decomposition tree having the largest width $w$. Equivalently, in the
related query answering problem we would like to exploit data information such as the size of relations, the selectivity of attributes, indexes, etc. In our framework, we address this issue by extending $G T P_{D M}$ with the novel notion of weighted (monotone) greedy tree projections, in order to combine the purely structural decomposition approach with quantitative aspects. Precisely we equip the decomposition method with an evaluation function that can be used to model many practical contexts where we can exploit additional knowledge on the problem at hand. Therefore we aim at computing greedy tree projections having the smallest weights. To this end, it is worth nothing that, computing a minimal weighted monotone greedy tree projection may be in general harder than computing a (purely structural) one. In this chapter, we still show a polynomial time algorithm for the weighted setting based on a restricted class of evaluation functions. Our results, find application in all those problems that can be solved efficiently on acyclic and quasi-acyclic instances, and in particular, it can be exploited immediately for solving CSPs where constraints are represented as finite relations encoding allowed tuples of values. Finally, we discuss implementation issues and a preliminary testing activity is performed, giving evidence of competitive results compared to the state-of-the-art decomposition methods known in the literature. This is the first step that needs to be done in order to emphasize the practical usefulness of structural approaches to CSP evaluation and query answering. We will return to this in more detail in the next chapter.

### 4.2 Monotone Greedy Tree projections

Most structural decomposition methods can be characterized through hypergraph games that are variations of the Robber and Cops graph game that characterizes the notion of treewidth. In particular, decomposition trees somehow correspond to mono-
tone winning strategies, where the escape space of the robber on the hypergraph is shrunk monotonically by the cops. In fact, unlike the treewidth case, there are hypergraphs where monotone strategies do not exist, while the robber can be captured by means of more complex non-monotone strategies. However, these powerful strategies do not correspond in general to valid decompositions.

In [Greco and Scarcello, 2010] it is given a general way to exploit the power of nonmonotone strategies, by allowing a disciplined form of non-monotonicity, that induces valid decomposition trees. Moreover, it is shown that deciding the existence of a (non-monotone) greedy winning strategy (and compute one, if any) is tractable. In fact, greedy strategies can be computed in polynomial time. To establish the result, a useful technical property is that greedy strategies can only involve a polynomial number of configurations (a configuration is a pair $(h, C)$, where $h$ is the hyperedge controlled by the cop, and $C$ is a $[h]$-component where the Robber stands).

Thus, based on them (even on non-monotone ones) it is possible to construct, again in polynomial time, tree projections, which are called greedy. We refer the reader to the work done by [Greco and Scarcello, 2010] for more on this subject. In this thesis, we give an implementation of greedy tree projections based on monotone strategies, a restricted variant that is easier to deal with for our purposes, which include the extension to the optimal weighted framework. Hereinafter, we write simply greedy tree projections for monotone greedy tree projections.

### 4.3 Normal form

To formalize our results, let us introduce some additional definitions and notations, which we will use later. Given two hypergraphs $\mathcal{H}_{1}$ and $\mathcal{H}_{2}$, in general many greedy tree projections with useless redundancies may exist. Hence, we would like
to have a suitable notion of minimality that allow us to identify the most desirable greedy tree projection. In fact, for several structural decomposition methods, normal forms have been defined to restrict the search space of decomposition trees, without loosing any useful decomposition. We make this more precise in the following.

Let $H D=\langle T, \chi, \lambda\rangle$ and $H D=\left\langle T^{\prime}, \chi^{\prime}, \lambda^{\prime}\right\rangle$ be two greedy tree projections. We write $H D \preceq H D^{\prime}$ if for each $h \in H \backslash H^{\prime}$, there is a set $h^{\prime} \in H^{\prime} \backslash H$ with $h \subseteq h^{\prime}$, where $H=\{\chi(p) \mid p \in \operatorname{vertices}(T)\}$ and $H^{\prime}=\left\{\chi^{\prime}\left(p^{\prime}\right) \mid p^{\prime} \in \operatorname{vertices}\left(T^{\prime}\right)\right\}$

Definition 4.1. Let $H D=\langle T, \chi, \lambda\rangle$ be a greedy tree projection of $\mathcal{H}_{1}$ w.r.t. $\mathcal{H}_{2}$. We say that $H D$ is in normal form if the following conditions hold:

1. for each vertex $p \in \operatorname{vertices}(T), \chi\left(T_{p}\right) \cap \lambda(p)=\chi(p)$.
2. for each tree projection $H D^{\prime}$ of $\mathcal{H}_{1}$ w.r.t. $\mathcal{H}_{2}$ that satisfies (4), $H D \preceq H D^{\prime}$ holds.

Hereinafter, we denote by $\operatorname{nfTP}\left(\mathcal{H}_{1}, \mathcal{H}_{2}\right)$ the set of all greedy tree projections in normal form of $\mathcal{H}_{1}$ w.r.t. $\mathcal{H}_{2}$.

Definition 4.2. Assume that a hypergraph $\mathcal{H}$ is given. Let $V, W$, and $\{X, Y\}$ be sets of nodes. Then, $X$ is said $[V]$-adjacent (in $\mathcal{H}$ ) to $Y$ if there exists a hyperedge $h \in \operatorname{edges}(\mathcal{H})$ such that $\{X, Y\} \subseteq(h-V)$. A $[V]$-path from $X$ to $Y$ is a sequence $X=X_{0}, \ldots, X_{\ell}=Y$ of nodes such that $X_{i}$ is [V]-adjacent to $X_{i+1}$, for each $i \in$ [0...८-1]. We say that $W$ is [ $V$ ]-connected if $\forall X, Y \in W$ there is a [ $V$ ]-path from $X$ to $Y$. A $[V]$-component $($ of $\mathcal{H})$ is a maximal $[V]$-connected non-empty set of nodes $W \subseteq(\operatorname{nodes}(\mathcal{H})-V)$.

For any [V]-component $C$, let $\operatorname{edges}(C)=\{h \in \operatorname{edges}(\mathcal{H}) \mid h \cap C \neq \emptyset\}$, and for a set of hyperedges $H \subseteq \operatorname{edges}(\mathcal{H})$, let $\operatorname{nodes}(H)$ denote the set of nodes occurring in $H$, that is $\operatorname{nodes}(H)=\bigcup_{h \in H} h$. For any component $C$ of $\mathcal{H}$, we denote by $\operatorname{Fr}(C, \mathcal{H})$ the frontier of $C($ in $\mathcal{H})$, i.e., the set nodes $($ edges $(C))$. Moreover, $\partial(C, \mathcal{H})$
denote the border of $C$ (in $\mathcal{H}$ ), i.e., the set $\operatorname{Fr}(C, \mathcal{H}) \backslash C$. Note that $C_{1} \subseteq C_{2}$ entails $\operatorname{Fr}\left(C_{1}, \mathcal{H}\right) \subseteq \operatorname{Fr}\left(C_{2}, \mathcal{H}\right)$. We write simply $\operatorname{Fr}(C)$ or $\partial(C)$, whenever $\mathcal{H}$ is clear from the context.

With all the technical notions in place, we are now ready to present $G T P_{D M}$.

### 4.4 An algorithm for computing greedy tree projections in normal form

As we already said in the previous sections, $G T P_{D M}$ is a polynomial time algorithm for computing greedy tree projections in normal form. It is worth noting that our algorithm finds application in all those problems that can be solved efficiently on acyclic and quasi-acyclic instances, and it can be exploited immediately for solving constraint satisfaction problems where constraints are represented as finite relations encoding allowed tuples of values. $G T P_{D M}$ receives as input a constraint hypergraph $\mathcal{H}_{1}$ and resource hypergraph $\mathcal{H}_{2}$ and output a greedy tree projection of $\mathcal{H}_{1}$ w.r.t. $\mathcal{H}_{2}$ if it exists. It starts by building a directed bipartite graph $C G$, called Candidates Graph, that stores all the informations that we need to compute any greedy tree projection in normal form. The nodes are partitioned in two sets which we call Component and Resource nodes. Component nodes have the form ( $S, C$ ) where $S$ is some hyperedge contained in $\mathcal{H}_{2}$ and $C$ is the component which we have to decompose using available Resource nodes. In particular, the root node is a special Component node $\left\{\left(\emptyset\right.\right.$, nodes $\left.\left.\left(\mathcal{H}_{1}\right)\right)\right\}$ which represents the entire problem. Resource nodes have the form $\left(R, C^{\prime}\right)$ where $R$ is some hyperedge contained in $\mathcal{H}_{2}$ and $C^{\prime}$ is the set of $[R]$-component that have to be decomposed. The node $\left(R, C^{\prime}\right)$ has a number of outgoing arcs pointing to all the nodes of the form $(S, C)$ for which it is a candidate solution, that is, for which $\operatorname{Fr}(C) \subset \operatorname{vars}(R)$ and $\operatorname{vars}(R) \cap \partial(C) \neq \emptyset$ holds. Moreover, $\left(R, C^{\prime}\right)$ has a number of incoming arcs for each $[R]$-component that is included in $C^{\prime}$,
as each of these nodes represents a subproblem of $\left(R, C^{\prime}\right)$. If a node $C \in C$ omponent nodes has no candidate solutions, i.e., if $\operatorname{outgoing}(C)=\emptyset$, then it is not solvable. We immediately exploit this information by removing all of its outgoing arcs, for which it was a subproblem. On the other hand, if it has some candidates, whenever all of them have been completely evaluated, it can propagate this information through its outgoing arcs. Once all nodes have been processed, the information encoded in the candidates graph $C G$ is enough to compute every greedy tree projection of $\mathcal{H}_{1}$ w.r.t $\mathcal{H}_{2}$ in normal form, if any. One of these greedy tree projections is eventually selected through the recursive procedure SelectGreedyTreeProj.

```
Algorithm 2 GreedyTreeProjection \({ }_{\text {DM }}\)
Input: \(\left\langle\mathcal{H}_{1}, \mathcal{H}_{2}\right\rangle\) two hypergraphs
Output: \(H D=\langle T, \chi, \lambda\rangle\) a hypertree of \(\mathcal{H}_{1}\) w.r.t. \(\mathcal{H}_{2}\)
    \(N_{\text {Root }} \leftarrow\left\{\left(\emptyset, \operatorname{nodes}\left(\mathcal{H}_{1}\right)\right)\right\} \quad \triangleright\) Root component
    \(\operatorname{RESOLVE}\left(\left(N_{\text {Root }}\right)\right)\)
    if \(\operatorname{incoming}\left(N_{\text {Root }}\right) \geq 1\) then
        \(E \leftarrow\{\emptyset\} \quad \triangleright\) The tree of decomposition is empty, initially
        Select \(p \in \operatorname{incoming}\left(N_{\text {Root }}\right)\)
        add the edge \(\left\{p, N_{\text {Root }}\right\}\) to \(E\)
        SelectGreedyTreeProj \((p)\)
        Output \(H D\)
    else
        Outputfailure
    end if
    procedure SelectGreedyTreeProj \((p)\)
        for all \(q \in \operatorname{incoming}(p)\) do
            Select \(p^{\prime} \in \operatorname{incoming}(q)\)
            add the edge \(\left\{p^{\prime}, q\right\}\) to \(E\)
            SelectGreedyTreeProj ( \(p^{\prime}\) )
        end for
    end procedure
```

```
Algorithm 3 Resolve \(\left(N_{C}\right)\)
    CandidateEdges \(\leftarrow\left\{h_{R} \in \operatorname{edges}\left(H_{2}\right) \mid h_{R} \cap \operatorname{Fr}\left(N_{C}\right) \neq \emptyset\right\} ;\)
    for all cover \(\in \operatorname{covers}\left(\operatorname{Fr}\left(N_{C}\right)\right.\), CandidateEdges \()\) do
        if \(\operatorname{vars}(\) cover \(\left.) \cap \partial\left(N_{C}\right)\right) \neq \emptyset\) then
            \(N_{R} \leftarrow\left(\right.\) cover, vars \((\) cover \(\left.) \cap\left(\partial\left(N_{C}\right) \cup \operatorname{Fr}\left(N_{C}\right)\right)\right)\)
            \(N_{\text {Rcomps }} \leftarrow\) findComponents \(\left(\partial\left(N_{c}\right)\right.\), vars(cover \(\left.)\right)\)
            for all \(N_{C^{\prime}} \in N_{\text {Rcomps }}\) do
                    RESOLVE \(\left(N_{C^{\prime}}\right)\)
                    if incoming \(\left(N_{C^{\prime}}\right)>0\) then
                        add an edge from \(N_{C^{\prime}}\) to \(N_{R}\)
                    else
                        break
                        end if
            end for
            if incoming \(\left(N_{R}\right)=\left|N_{\text {Rcomps }}\right|\) then
                add an edge from \(N_{R}\) to \(N_{C}\)
                return
            else continue
            end if
        else continue
        end if
    end for
```


### 4.4.1 Speeding-up computation through greedy coverings

In line 1 of Algorithm 3, rather than trying to decompose the component node using all the available resources, we can restrict the search space by selecting a subset of edges $\mathcal{H}_{2}$ such that for each $R \in \mathcal{H}_{2}$ we have that $\operatorname{Fr}(C) \subset \operatorname{vars}(R)$. In addition to this, maintaining appropriate data structures (such as a list of "good" resources sorted by the size of $\operatorname{vars}(R) \cap \partial(C))$ allows us to quickly decompose the considered component node.

We now examine the complexity of computing pure greedy tree projections in more detail.

Theorem 4.3. Given a constraint hypergraphs $\mathcal{H}_{1}$ and a resource hypergraph $\mathcal{H}_{2}$, Algoritm 2 runs in time $O\left(n m^{2} \log m\right)$ using dynamic programming, where $n$ is the number of nodes of $\mathcal{H}_{1}$ and $m$ is the number of edges of $\mathcal{H}_{2}$.

Proof. Let $\mathcal{H}_{1}, \mathcal{H}_{2}$ be the constraint and the resource hypergraphs respectively. Let
$n=\operatorname{nodes}\left(\mathcal{H}_{1}\right)$ and $m=\operatorname{edges}\left(\mathcal{H}_{2}\right)$. Observe that, for each edge $R \in \mathcal{H}_{2}$ there are at most $O(n)[R]$-components. Thus, since we store the component nodes, the number of recursive calls to Algorithm 3 is bounded by $O(n m)$. What it remains to analyze is the number of loops in line 2 of Algorithm 3. Observe that, for any component node $C$ the auxiliary procedure covers computes a subset of edges $\mathcal{H}_{2}$ such that for each $R \in \operatorname{covers}\left(\mathcal{H}_{2}\right)$ we have that $\operatorname{Fr}(C) \subset \operatorname{vars}(R)$. This number is clearly bounded by $m$. Eventually, to speed up the computation these edges are sorted by the size of $\operatorname{vars}(R) \cap \partial(C)$. Therefore, an upper bound for the overall complexity is given by $O\left(n m^{2} \log m\right)$. Importantly, note that the polynomial runtime is guaranteed under the assumption that already visited components are stored and reused, otherwise, the same subtree could be constructed multiple times which would result in an exponential runtime.

### 4.5 Evaluation functions

In this section, we define the notion of weighted greedy tree projection, in order to combine the power of pure greedy tree projections with quantitative approaches. More precisely, we equip our restricted tree projections with evaluation functions that can be used to model situations where we have further information at hand on the given problem, besides its hypergraph representation. For instance, consider concrete contexts like database query answering where we can use statistics such as the size of relations, the selectivity of attributes, indexes, etc., to improve the runtimes of join algorithms. In fact, all commercial database systems exploit such statistics (and do not care about structural properties). In our implementation, users can specify an evaluation function $f$ and the algorithm will compute, if it exists, a $f$-minimal greedy tree projection.

Now we describe evaluation functions in detail. Let $\mathbb{D}$ be a given domain of values, and assume that $\preceq$ is a total order defined on it. An evaluation function $f$ over $\mathbb{D}$ associates an element $f(V, h, \mathcal{S}) \in \mathbb{D}$ with each $V, h \subseteq \mathcal{V}$ and multiset $\mathcal{S}$ of elements in $\mathbb{D}$. The value of a greedy tree projection $H D=\langle T, \chi, \lambda\rangle$, with $p=\operatorname{root}(T)$, under $f$ is inductively defined as:

$$
f(H D)=f\left(\chi(p), \lambda(p),\left\{f\left(\left\langle T_{s}, \chi, \lambda\right\rangle\right) \mid s \text { is a child of } p\right\}\right) .
$$

A greedy tree projection $H D \in \operatorname{nfP}\left(\mathcal{H}_{1}, \mathcal{H}_{2}\right)$ is $f$-minimal if $f(H D) \preceq f\left(H D^{\prime}\right)$, for each $H D^{\prime} \in \operatorname{nfTP}\left(\mathcal{H}_{1}, \mathcal{H}_{2}\right)$. Moreover, we say that an evaluation function is valuation-monotonic if for each $V, h, \mathcal{S}, \mathcal{S}^{\prime}, \mathcal{S}^{\prime}=\mathcal{S} \backslash\{v\} \cup\left\{v^{\prime}\right\}$ with $v \preceq v^{\prime}$ implies $f(V, h, \mathcal{S}) \preceq f\left(V, h, \mathcal{S}^{\prime}\right)$.

Theorem 4.4. For any valuation-monotonic function $f$, a $f$-minimal greedy tree projection can be computed in polynomial time.

Proof. Consider the algorithm WeightedGreedyTreeProjection ${ }_{D M}$ shown below, which returns a $f$-minimal greedy tree projection of a hypergraph $\mathcal{H}_{1}$ w.r.t. a hypergraph $\mathcal{H}_{2}$. We note that, this algorithm is derived from Algorithm 2 with a simple change in line 2 where we call Algorithm 5 in place of Algorithm 3. Algorithm 5 is in its turn a variant of Algorithm 3 with the addition of lines 15, 16 and 17. This variant does the same operations performed in Algorithm 2 and in addition for every successful resource node $r$, it set weight $(r) \leftarrow f(r)$. In particular, for each component node $c$, we keep only the child having the minimum weight between all children. The decomposition tree is then computed very efficiently using a branch-and-bound technique to avoid exploring solutions that are certainly not optimal. Once the weighted candidates graph is build, all the informations that we need are in place to compute the $f$-minimal greedy tree projection.

```
Algorithm 4 WeightedGreedyTreeProjection \({ }_{\text {DM }}\)
Input: \(\left\langle\mathcal{H}_{1}, \mathcal{H}_{2}\right\rangle\) two hypergraphs
Output: \(H D=\langle T, \chi, \lambda\rangle\) a hypertree of \(\mathcal{H}_{1}\) w.r.t. \(\mathcal{H}_{2}\)
    \(N_{\text {Root }} \leftarrow\left\{\left(\emptyset, \operatorname{nodes}\left(\mathcal{H}_{1}\right)\right)\right\} \quad \triangleright\) Root component
    \(\operatorname{RESOLVE}_{\mathrm{W}}\left(\left(N_{\text {Root }}\right)\right)\)
    if incoming \(\left(N_{\text {Root }}\right) \geq 1\) then
        \(E \leftarrow\{\emptyset\} \quad \triangleright\) The tree of decomposition is empty, initially
        Select a minimum weighted \(p \in \operatorname{incoming}\left(N_{\text {Root }}\right)\)
        add the edge \(\left\{p, N_{\text {Root }}\right\}\) to \(E\)
        SelectWGreedyTreeProj \((p)\)
        Output \(H D\)
    else
        Outputfailure
    end if
    procedure \(\operatorname{SelectWGreedyTreeProj}(p)\)
        for all \(q \in \operatorname{incoming}(p)\) do
            Select a minimum weighted \(p^{\prime} \in \operatorname{incoming}(q)\)
            add the edge \(\left\{p^{\prime}, q\right\}\) to \(E\)
            SelectWGreedyTreeProj \(\left(p^{\prime}\right)\)
        end for
    end procedure
```

```
Algorithm 5 RESOLVE \(_{\mathrm{W}}\left(N_{C}\right)\)
    CandidateEdges \(\leftarrow\left\{h_{R} \in \operatorname{edges}\left(H_{2}\right) \mid h_{R} \cap \operatorname{Fr}\left(N_{C}\right) \neq \emptyset\right\} ;\)
    for all cover \(\in \operatorname{covers}\left(\operatorname{Fr}\left(N_{C}\right)\right.\), CandidateEdges \()\) do
        if \(\operatorname{vars}(\) cover \(\left.) \cap \partial\left(N_{C}\right)\right) \neq \emptyset\) then
            \(N_{R} \leftarrow\left(\right.\) cover, vars \((\) cover \(\left.) \cap\left(\partial\left(N_{C}\right) \cup \operatorname{Fr}\left(N_{C}\right)\right)\right)\)
            \(N_{\text {Rcomps }} \leftarrow\) findComponents \(\left(\partial\left(N_{c}\right)\right.\), vars(cover \(\left.)\right)\)
            for all \(N_{C^{\prime}} \in N_{R c o m p s}\) do
                    RESOLVE \(\mathrm{W}_{\mathrm{W}}\left(N_{C^{\prime}}\right)\)
                    if incoming \(\left(N_{C^{\prime}}\right)>0\) then
                        add an edge from \(N_{C^{\prime}}\) to \(N_{R}\)
                    else
                        break
                end if
            end for
            if incoming \(\left(N_{R}\right)=\left|N_{\text {Rcomps }}\right|\) then
                    weight \(\left(N_{R}\right) \leftarrow f\left(N_{R}\right)\)
                    if weight \(\left(N_{C}\right)<\) weight \(\left(N_{R}\right)\) then
                    weight \(\left(N_{C}\right) \leftarrow\) weight \(\left(N_{R}\right)\)
                    add an edge from \(N_{R}\) to \(N_{C}\)
                    end if
            else continue
            end if
        else continue
        end if
    end for
```

The running time of WeightedGreedyTreeProjection ${ }_{\mathrm{DM}}$ is $O\left(n m^{2} \log m\right)$ from The-
orem 4.3, plus the time needed to compute $f(r)$ for each resource node $r$. Because we require the function $f$ to be polynomial-time computable, it follows that Algorithm 4 can find a $f$-minimal greedy tree projection, if it exists, in polynomial time.

### 4.6 Non-monotic valuation functions



Figure 4.1: Weighted candidates graph in example of Section 4.6

There are many practical situations where we required to process large datasets. Example includes social networks, linked data and also biological motifs. To deal with the complexity of processing such large data, it is usually exploited parallel computing. However, it is known that parallelization performs bad on trees that are unbalanced. Therefore in this setting, we aim at computing the most balanced decomposition. We formalize this idea by defining the evaluation function as follows.

$$
\psi(p)=(h(p), i(p)) \quad \text { where }
$$

$$
\begin{aligned}
& h(p)=1+\max _{s \in \operatorname{children}(p)} h(s), \\
& i(p)=\max \left(\max _{s \in \operatorname{children}(p)} i(s),\left(\max _{s \in \operatorname{children}(p)} h(s)-\min _{s \in \operatorname{children}(p)} h(s)\right)\right)
\end{aligned}
$$

Let $h(p)$ be the height and $i(p)$ be the imbalance of the tree $T$ rooted at $p$ and consider the illustration reported in Figure 4.1 to help the intuition. Observe the candidates graph and note that we represent resource nodes with a circle while component nodes are represented with a rectangle. Our objective is to single out the decomposition having the smallest height $h(R)$. From the illustration above we see that the children of $R$ are $C_{1}$ and $C_{2}$. The weight of $R$ therefore depends on the costs (and hence on the domain values) of $C_{1}$ and $C_{2}$. Consider $C_{1}$ and assume that $\psi(R)$ is valuation-monotonic. Therefore since $R_{i}$ and $R_{j}$ have the same imbalance value, i.e., 0 , at this level they are interchangeable and thus we choose $R_{i}$ to solve $C_{1}$. After selecting $R_{i}$ as a solution node for the subproblem node $C_{1}$ we move to $C_{2}$ and we select $R_{k}$ for the same reason. Thus we can compute $\psi(R)$ with $R_{i}$ and $R_{k}$ being the children of $R$ and we get as result the domain value $(11,9)$ and hence the weight $(R)=9$. Now assume to select $R_{j}$ in place of $R_{i}$ and $R_{z}$ in place of $R_{k}$. Again at this level, from a local perspective it does not change anything since the imbalance is always 0 . However, if we compute again $\psi(R)$ we get a domain value of $(301,0)$ and hence the weight $(R)=0$, which is smaller than before. It turns out that if we assume $\psi(R)$ to be valuation-monotonic, we fail at computing most balanced decomposition. We can easily fix this by assuming $\psi(R)$ to be non-monotonic. In this case it easy to see that it can explore all the possible domain values without getting stuck in local minima.

### 4.7 Experimental results

In this section we present the performances of our implementation compared with those of the approaches listed in Table 4.7 which represent the state-of-the-art algorithms for computing a minimal (purely structural) hypertree decomposition. In particular, opt-k-decomp is the first polynomial time algorithm for constructing hypertree decomposition of minimal width less than or equal to a given constant $k$. det-k-decomp is a backtracking-based algorithm for computing hypertree decompositions of width smaller or equal than $k$. det- $k$-decomp is the best algorithm known for this task up to date. We observe that these methods compute a hypertree decomposition of the given CSP instance, therefore we set our general tool in hypertree decomposition mode. Our experiments evaluate the methods on four sets of instances involving many problems from industry and academic research. This is detailed below.

Problem instance selection. To compare these algorithms, we tried to identify a set of instances that could cover the largest number of application domains. Therefore we identified 4 categories: real-world problems, academic, random and pattern. We performed a random selection of instances in each category. Finally 26 instances were considered in our benchmark.

Hardware and software configuration. Experiments have been performed on a dedicated machine, equipped with an Intel Core i7-3770k 3.5 GHz processor with 12 GB (DDR3 1600 MHz ) of RAM, and running Linux Debian Jessie. Our algorithm is implemented in Java and the code was executed on the JDK 1.8.0_05-b13. Moreover, we used a C++ implementation of opt-k-decomp and the source code was compiled with GCC 4.9.3. Regarding det-k-decomp, we downloaded the binary executable
available from here ${ }^{1}$. Default settings were used for the various methods. We set a timeout of 300 seconds to the total execution time of the methods. For each instance we report the average computation time (based on three runs), the number of instances correctly decomposed and the number of wins in finding the minimal width. Of course if a method times out we exclude the related running time from the means. Results are reported in Table 4.7 and4.7.

|  | Instance |  |  |  |  | opt-k- | comp | det-k-d | comp |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | constraints | variables | width | time | width | time | width | time |
| $e^{20^{20}}$ | instExtCW-m1c-ogd-vg16-20 | 36 | 320 | 16 | 0 | - | - | 16 | 0 |
|  | instExtCW-m1c-ogd-vg16-29 | 35 | 304 | 16 | 0 | - | - | 16 | 0 |
|  | instExtCW-m1c-words-vg16-18 | 34 | 288 | 16 | 0 | - | - | 16 | 0 |
|  | instExtCW-m1c-words-vg16-20 | 36 | 320 | 16 | 0 | - | - | 16 | 0 |
|  | instExtCW-m1c-lex-vg15-17 | 32 | 255 | 15 | 0 | - | - | 15 | 0 |
|  | instExtCW-m1c-lex-vg16-18 | 34 | 288 | 16 | 0 | - | - | 16 | 0 |
|  | renault-mod-16-ext | 149 | 111 | 3 | 0 | $+$ | + | 3 | 0 |
|  | renault-mod-32-ext | 154 | 111 | 5 | 0 | $+$ | + | 5 | 0 |
|  | renault-mod-48-ext | 149 | 108 | 4 | 0 | $+$ | + | 4 | 0 |
|  | instExtlangford-2-8 | 128 | 16 | 8 | 20 | $+$ | + | 8 | 0 |
|  | magicSquare-3-glb | 9 | 9 | 1 | 0 | 1 | 1 | 1 | 0 |
|  | magicSquare-20-glb | 400 | 43 | 1 | 0 | 1 | 0 | 1 | 0 |
|  | perfect-square-packing-20-20 | 953 | 800 | 3 | 0 | $+$ | + | 3 | 0 |
|  | perfect-square-packing-38-38 | 3518 | 2888 | 3 | 6 | $+$ | $+$ | 3 | 10 |
|  | schurr-lemma-100-9-mod | 2450 | 100 | - | - | $+$ | + | - | - |
|  | schurr-lemma-66-4 | 1056 | 66 | - | - | $+$ | + | - | - |
|  | schurr-lemma-50-9-mod | 600 | 50 | 31 | 4 | $+$ | $+$ | 30 | 0 |
|  | schurr-lemma-23-3 | 210 | 30 | 9 | 5 | $+$ | $+$ | 8 | 0 |
| $e^{\gamma^{20}}$ | tightness-rand-2-40-8-753-100-80-ext | 753 | 40 | 37 | 212 | + | + | 31 | 3 |
|  |  | 753 | 40 | 41 | 69 | $+$ | + | 31 | 3 |
|  | ModelRB-frb59-26-5-mgd-ext | 539 | 59 | - | - | $+$ | $+$ | 38 | 0 |
|  | ModelRB-frb59-26-3-mgd-ext | 547 | 59 | - | - | + | + | 37 | 6 |
|  | ModelB-rand-27-27-351-163-63021-ext | 351 | 27 | 22 | 30 | - | - | 19 | 4 |
| $e^{\sigma^{x^{x^{2}}}}$ | graph-coloring-geom-40-6-ext | 40 | 78 | 3 | 0 | - | - | 3 | 0 |
|  | graph-coloring-queens-5-5-5-ext | 160 | 25 | 13 | 2 | - | - | 10 | 40 |
|  | graph-valiente-val49-50 | 1128 | 128 | 3 | 2 | + | $+$ | 3 | 0 |

Table 4.1:
The time (in seconds) required to solve a number of instances generated from real-world, academic, random and pattern problems. A dash ( - ) means timeout while ( + ) denotes overflow

| Algorithm | \# of decomposed instances | \# number of wins |
| :--- | ---: | ---: |
| Here | 22 | 15 |
| opt-k-decomp | 2 | 2 |
| det-k-decomp | 27 | 23 |

Table 4.2:
Number of correctly decomposed instances and number of wins in finding the minimal width.

From the tables above we note that det-k-decomp, as expected, shows the best overall performances. However GreedyTreeProjection ${ }_{D M}$ closely follows on almost all

[^1]instances. opt-k-decomp is the worst one because it frequently times out on difficult instances. We observe that our implementation is not optimized for speed. In fact, we note that our tool is more general than those tools that implement opt-k-decomp and det-k-decomp and in order to reduce development times we were forced to reuse as much code as possible to implement all variants and that inevitably results in performance limitations.

## CHAPTER 5

## An Hybrid Approach for Counting Solutions

### 5.1 Introduction to Hybrid Tractability

We have seen in Chapter 2 that the problem of counting the solutions of constraint satisfaction problems is equivalent to counting the number of query answers in the related field of database theory. This problem occurs in the presence of SQL queries specifying "COUNT" aggregates. Note that these queries are very often at the basis of decision support systems examining large volumes of data in order to get insights on critical business questions. Dealing with queries processing large amounts of data has been a crucial algorithmic challenge in relational databases. Indeed, the definition and implementation of advanced optimization strategies for SQL queries are topics widely discussed in the literature, but at the same time extremely demanding in order to design efficient RDBMS [Codd, 1970]. The optimizers are asked to calculate the best query plan, using a strategy for evaluating a query over a relational database that provides the best performance compared to all possible strategies. This process, however, forces traditional optimizers to find a good compromise between two conflicting factors [Chaudhuri, 1998]. The first factor is related to the cost of evaluating a query that can be exponential in the length of the query itself [Libkin, 2004]. The second factor is related to the complexity of the calculation


Figure 5.1: The hypergraph $\mathcal{H}\left(Q_{2}\right)$ and a hypertree decomposition of $\mathcal{H}\left(Q_{2}\right)$ of width 2
of the optimal query execution plan; optimizers are forced to limit the search space to query plans with simple structures, thus avoiding a costly search. The choice of the most promising plan, rather than optimal, is determined by a cost model based on quantitative information on the collected data such as the size of database relations and attributes selectivity. For these reasons, traditional optimizers are based on quantitative methods. While generally very efficient, these optimizers can not ensure a "bound" on the response time of a query; Furthermore, the execution plan produced by them is often approximate and far away from the ideal plan. To cope with the limitations of traditional optimizers and calculate an optimal query plan that guarantees a polynomial bound on response time, it has been proposed in the community of database theory, a completely different approach. The latter is based on the use of the structural properties of the query, which, as we already noted, are usually represented by the hypergraph associated with the query.

By exploiting these properties, one can respond to a wide class of queries efficiently: this class is known as the class of acyclic queries (or more generally, of bounded hypertree width). As an example, consider the following query $Q_{2}$.

```
SELECT n_name,
    sum(l_extendedprice*(1-1_discount)) AS revenue
FROM customer, orders, lineitem,
    supplier, nation, region
WHERE c_custkey = o_custkey
    and l_orderkey = o_orderkey
    and l_suppkey = s_suppkey
    and c_nationkey = s_nationkey
    and s_nationkey = n_nationkey
    and n_regionkey = r_regionkey
    and r_name = '[REGION]'
    and o_orderdate >= date '[DATE]'
    and o_orderdate < date '[DATE]' + interval '1' year
GROUP BY n_name ORDER BY revenue desc;
```

Observe its associated hypergraph $\mathcal{H}\left(Q_{2}\right)$ that is depicted on the top of Figure 5.1 and the corresponding hypertree decomposition of $\mathcal{H}\left(Q_{2}\right)$ on the bottom.
$Q_{2}$ can be now answered using a two-step plan. First, one computes for each cluster of the decomposition tree, the join of its nodes. Then, starting from the leaves, one moves bottom-up towards the root, by performing upward semi-joins of intermediate results.

Despite their very nice computational properties, structural decomposition methods have not had any serious impact on the design of commercial database optimizers, and the interest for these techniques mainly remained at a theoretical level. This is mainly because decomposition methods flatten in the query hypergraph all the quantitative aspects of data, do not take care of the output of the variables, and do not support aggregate operators and nested queries. Therefore, since we are dealing with a setting where the query and the database are both given as input (neither is constant), it is meaningful to consider decomposition techniques that are able to exploit structural properties of the query in combination with properties of the given data, such as functional dependencies or any other feature that may simplify the evaluation. Intuitively, such features may induce different structural properties that
are not identified by the "worst-possible database" perspective of purely structural methods. The motivating idea is to end up with algorithms that can be practically applied over a wide range of real world settings where, in particular, such methods do not suffice alone.
[Pichler and Skritek, 2013] have recently proposed an algorithm for counting answers to acyclic conjunctive queries, whose scaling is in the worst case exponential w.r.t. the maximum number of tuples over the database relations, denoted by $m$ hereinafter. In fact, they also shown that counting query answers remains \#P-hard over this structurally simple class. Moreover, they pointed out that their algorithm can be extended to queries that are not necessarily acyclic, but have bounded hypertree width. [Greco and Scarcello 2014], moved a first step towards proposing a hybrid decomposition method for counting problems. In the next section we analyze such an extension. Finally we give an implementation and we discuss the issues related to this algorithm, we also present preliminary experimental results.

Before moving to the technical details of the proposed algorithm, note that for the sake of presentation, we mostly focus on the database setting and the query answering problem. However, we remark that all results can immediately be applied to all problems that can be recast as constraint satisfaction problems (as we shall see with a tiny example).

### 5.2 An algorithm for counting answers

We recall that a hypertree decomposition [Gottlob et al., 2002] of a conjunctive query $Q$ is a generalized hypertree decomposition enjoying an additional property, often called descendant condition. In particular, a hypertree for a query $Q$ is a triple $\langle T, \chi, \lambda\rangle$, where $T$ is a rooted tree, and $\chi$ and $\lambda$ are labeling functions associating
each vertex $p$ of $T$ with two sets $\chi(p) \subseteq \operatorname{vars}(Q)$ and $\lambda(p) \subseteq \operatorname{atoms}(Q)$. The set of the vertices of $T$ is denoted by vertices $(T)$, whereas its root is denoted by $\operatorname{root}(T)$. Moreover, for any $p \in N$, we denote by $T_{p}$ the subtree of $T$ rooted at $p$ (hence, $\left.T=T_{\text {root }(T)}\right)$, and by $\chi\left(T_{p}\right)$ the set of all variables occurring in the $\chi$ labeling of $T_{p}$.

A hypertree decomposition of $Q$ is a hypertree $H D=\langle T, \chi, \lambda\rangle$ for $Q$ such that: (1) for each $q \in \operatorname{atoms}(Q)$, there exists $p \in \operatorname{vertices}(T)$ such that $\operatorname{vars}(q) \subseteq \chi(p)$; (2) for each $X \in \operatorname{vars}(Q)$, the set $\{p \in \operatorname{vertices}(T) \mid X \in \chi(p)\}$ induces a (connected) subtree of $T$; (3) for each $p \in \operatorname{vertices}(T), \chi(p) \subseteq \operatorname{vars}(\lambda(p))$; and (4) for each $p \in$ $\operatorname{vertices}(T)$, $\operatorname{vars}(\lambda(p)) \cap \chi\left(T_{p}\right) \subseteq \chi(p)$ (descendant condition). The decomposition $H D$ is said complete if for each atom $q \in \operatorname{atoms}(Q)$, there exists $p \in \operatorname{vertices}(T)$ such that $q \in \lambda(p)$.

This notion is a true generalization of acyclicity, as acyclic queries are precisely those having hypertree width 1 . Importantly, for any fixed natural number $k \geq 1$, deciding whether a query has hypertree width at most $k$ is in LOGCFL, and thus it is a tractable and highly parallelizable problem [Gottlob et al., 2002].

Finally for each vertex $p$ of $T$ we define $r_{p}=\pi_{\chi(p)}\left(\bowtie_{q \in \lambda(p)} q^{\mathrm{D}}\right)$, and $\delta(p)$ the number of its children, which we assume to be ordered in some way; If $R$ and $R^{\prime}$ are sets of sets of substitutions, then we define an ad hoc semi join as $R \ltimes R^{\prime}=\left\{S \ltimes S^{\prime} \mid\right.$ $\left.S \in R, S^{\prime} \in R^{\prime}, S \ltimes S^{\prime} \neq \emptyset\right\} ;$

```
Algorithm 6 counting-answers
Input: \(\left\langle\mathcal{H}_{1}, \mathcal{H}_{2}\right\rangle\) two hypergraphs
Output: \(H D=\langle T, \chi, \lambda\rangle\) a hypertree of \(\mathcal{H}_{1}\) w.r.t. \(\mathcal{H}_{2}\)
    for all \(p \in \operatorname{vertices}(T)\) do \(\quad \triangleright(\mathbf{1}) \longrightarrow\) Initialization
        \(R_{p}^{0} \leftarrow\left\{\sigma_{\theta}\left(r_{p}\right) \mid \theta \in \pi_{\text {free }(Q)}\left(r_{p}\right)\right\}\)
    end for
    \(B U:=\{ \} ; \quad \triangleright(2) \longrightarrow\) Bottom-Up
    while \(B U \neq \operatorname{vertices}(T)\) do
        Let \(p \notin B U\) whose children \(q_{1}, \ldots, q_{\delta(p)}\) are all in \(B U\)
        \(B U:=B U \cup\{p\}\)
        for \(\alpha:=1\) to \(\delta(p)\) do
            Let \(q\) be the \(\alpha^{\text {th }}\) child of \(p\)
                with its computed \(R_{q}:=R_{q}^{\delta(q)}\) and \(c_{q}=c_{q}^{\delta(q)}\)
            \(R_{p}^{\alpha}:=R_{p}^{\alpha-1} \ltimes R_{q}\);
            \(c^{\alpha}(S):=\sum_{(*)} c_{p}^{\alpha-1}\left(S_{p}\right) \times c_{q}\left(S_{q}\right)\), for each \(S \in R_{p}^{\alpha}\)
                    \((*): S_{p} \in R_{p}^{\alpha-1}, S_{q} \in R_{q} \mid S=S_{p} \ltimes S_{q} ;\)
        end for
    end while
    return \(\sum_{S \in R_{r}^{\delta(r)}} c_{r}^{\delta(r)}(S)\), where \(r=\operatorname{root}(T) ; \quad \triangleright(\mathbf{3}) \longrightarrow\) Finalization
```

With the above notation in place, consider Algorithm 6, which takes as input a complete hypertree decomposition $H D=\langle T, \chi, \lambda\rangle$ of a query $Q$, and a database D . The algorithm works as follows: In the initialization step, for each vertex $p$, the set $r_{p}$ (of the partial solutions over $p$ ) is partitioned according to the profiles given by the configurations of values allowed for the free variables in $\chi(p)$. Let $R_{p}^{0}$ be the resulting partition. In the subsequent steps, $p$ will be associated with a set $R_{p}^{\alpha}$ where $\alpha \in\{0,1, \ldots, \delta(p)\}$, with $\delta(p)$ being the number of children of $p$. In particular, each of these sets consists of a set of substitutions that are manipulated via an ad-hoc semi join " $\ltimes$ " operator such that: if $R$ and $R^{\prime}$ are sets of sets of substitutions (hereinafter called \#-relations), then $R \ltimes R^{\prime}=\left\{S \ltimes S^{\prime} \mid S \in R, S^{\prime} \in R^{\prime}, S \ltimes S^{\prime} \neq \emptyset\right\}$. In the bottom-up evaluation, upward semi joins are performed, and $R_{p}^{\alpha}$ is the \#-relation associated with $p$ after that the first $\alpha$ children (according to some fixed ordering) have been processed. Each set $S \in R_{p}^{\alpha}$ is associated with a value $c_{p}^{\alpha}(S)$ initialized to 1 and updated bottom-up. At the end, the value computed in the finalization phase is returned as the desired number of solutions. The correctness of this algorithm follows
from the arguments in [Pichler and Skritek, 2013]. In fact, because the running time of the algorithm in [Pichler and Skritek, 2013] is in general exponential w.r.t. m, [Greco and Scarcello, 2014] introduced the following notion to keep under control such a blow-up.

Definition 5.1. Let $b \geq 0$ be a natural number and let $H D=\langle T, \chi, \lambda\rangle$ be a hypertree (not necessarily decomposition) of a query $Q$. We say that a database D is $b$-unbound for $Q$ w.r.t. $H D$ if there is a vertex $v$ of $T$ and a substitution $\theta \in \pi_{\text {free }(Q)}\left(r_{v}\right)$ such that

$$
\left|\sigma_{\theta}\left(r_{v}\right)\right|>b, \quad \text { where } r_{v}=\pi_{\chi(v)}\left(\bowtie_{q \in \lambda(v)} q^{\mathrm{D}}\right)
$$

The minimum $b$ for which D is not $b$-unbound is the boundness of D for $Q$ w.r.t. $H D$, and it is shortly denoted by bound (D, HD).

Intuitively, the boundness value bound $(\mathrm{D}, H D)$ for $Q$ provides an estimate on the size of the information that is required to flow along the given hypertree decomposition $H D$ in order to answer a counting problem over D. Finally they proved the following theorem.

Theorem 5.2. Let $H D=\langle T, \chi, \lambda\rangle$ be a width-k hypertree decomposition of $Q$, and let D be a database such that bound $(\mathrm{D}, H D) \leq h$. Then, Algorithm 6 runs in time $O\left(|\operatorname{vertices}(T)| \times m^{2 \times k} \times 4^{h}\right)$.

### 5.3 Implementation issues and System Architecture

In the last section, we discussed Algorithm 6 which can especially be exploited to find all solutions to a counting problem by considering the whole set of variables as output ones. However, it should be noted this high level description is of great interest mainly from a theoretic perspective. Indeed a verbatim implementation will
not perform very well in practical settings where the domain of each output variable even if bounded can still contain a very high number of values. This is especially true when counting CSP solutions for classes of instances having not only a large domain but also a huge number of constraints. In fact, the algorithm requires to create a "view" for each possible tuple thus generating a terrible overhead which soon vanish all the nice theoretical computationally properties. Therefore, in complex real-world settings, this observation once more strengthen the choice of backtracking approaches based on arc-consistency and propagations over the structural ones. Indeed it is not by chance that all the CSP and SAT solvers to date completely discard decomposition methods. In this dissertation we show how to efficiently exploit our tools both standlone or to help reduce the search space of backtracking approaches.

We observe that when counting all solutions, at low level we do not really need to create a "subview" for each tuple of the relation. Indeed, it suffices to virtually augment the relation and hence each tuple with a new attribute encoding the number of extensions of that tuple. This value is initially set to 1 for each tuple and is then updated through an ad-hoc semi join operator in the bottom-up process. We give details of this operation in the next section. Intuitively our idea is to exploit the acyclicity and our semi join operator to efficiently compute the counting task. Indeed, we focused in the database setting by designing an HybridOptmizer that works on top of PostgreSQL (and possibly of any other DBMS) and is in charge of computing on optimal query plan for queries specifying the "count" aggregate. The architecture is depicted in Figure ??. Given an SQL query, the Parser preliminary verifies its correctness. To this aim, we exploit the syntactical analyzer J-SQLParser ${ }^{1}$ that has been further enhanced to check whether the query complies with the un-

[^2]derlying database schema. The resulting parse tree is then processed by the Query Rewriter, which prepares all data structures needed for structural analysis and optimization. In particular, the module modifies the parse tree by associating each nested SQL statement with a distinguished node storing all information to evaluate it independently of the others, in particular by propagating join constraints on the basis of the scope of attributes and relations. The HybridOptimizer module is the main responsible for the structural optimization. Its implementation relies on the decomposition technique described in Chapter 4 exploiting an evaluation function based on the cardinality of relations and the selectivity of attributes as presented in [Scarcello et al., 2007].

In a nutshell, the optimizer works as follow. Starting from the leaves, it navigates the rewritten parse tree and calls on each node (i.e., subquery) the Subplan Handler, which is responsible for computing a suitable optimized plan for it. The handler preliminary invokes the Hypergraph Manager, which builds the hypergraph associated with the current input node. The Hypergraph Manager collects also information about statistics of the relations and attributes appearing in the node, by invoking the underlying system catalog. The computed hypergraph and its associated statistics are then passed to the decomposition module that computes a hypertree decomposition, whose width is fixed by the user at configuration time. Note that, according to the q-decomposition strategy, described in [Scarcello et al., 2007] we force all output variables (i.e., selected items, attributes appearing in group by or in order by clauses) to appear in the root of the decomposition. Furthermore, we also request that the attributes appearing in some non-join condition are covered together by at least a cluster of the computed decomposition.

Finally, the decomposition is translated in an actual query plan by the View

Builder, which rephrases it in terms of a tree of SQL views, possibly materialized. During this phase, all non-join constraints are assigned to the node of the decomposition, for being later evaluated. The computed optimized plan is then returned to the SubPlan Hadler and finally to the HybridOpitmizer module, which arranges the computed plan node as an element of the current query plan tree. Once all nodes have been process, the final plan tree gives us the query rewrite. The key to implementing this query rewrite idea is to develop a new ad-hoc physical semi-join operator and to integrate it within the PostgreSQL database management system.

### 5.3.1 Hacking PostgreSQL

In this section we detail the implementation of the ad-hoc Semi Join operator. We start by sorting the two tables and scan them looking for a match. When a match is found, we stop moving on the outer table, we mark a pointer to the last tuple visited, set a counter, and continue scanning inner tuples until they have a match with the marked tuple. During this scan, we properly update the counter at each match. When no more matches occur on the inner, we return the marked tuple and the counter, and start again moving on the outer table. Each outer tuple matching the same set of inner tuples of the last marked one is immediately returned (without a rescan of the inner table) with the last computed counter. Note that we implemented the Semi Join operator inside Postgresql 9.0.2. That is because our work is based on the preliminary work published in ??. Thus to reduce development times we decided to extend the implementation that was already available with the new ideas. However, to stress the effectiveness of our approach, in the next section we compare our hybrid 9.0 .2 with the more recent 9.3 .5 regular version. We are pretty sure that if we can get good results with an older version of PostgreSQL, than we can expect to be even more fast with the latest releases.

### 5.4 Some experiments

This section reports the results of the testing activity that we conducted in order to evaluate the effectiveness of hybrid 9.02 compared to the regular 9.3.5. Before we discuss the results, let us describe the methodology used.

Dataset. We evaluate the performances of on the TPC-DS benchmarks. TPCDS is the de-facto industry standard benchmark for measuring the performance of modern decision support systems including, but not limited to, Big Data systems. It models analytical workloads (reporting, data mining, iterative OLAP and so on), including SQL queries of various complexity, i.e., performing aggregations (GROUP BY), various joins on large amounts of data. This benchmark effectively extends and deprecates TPC-H benchmark, improving it in multiple ways to make it broadly representative of modern decision support systems. Firstly, it makes the schema more complex (e.g. more tables), and uses less uniform distributions of the data (which makes cardinality estimations way more difficult). It also increases the number of query templates from 22 to 99 , and uses modern features like CTEs, window functions and grouping sets. As recommended by TPC-DS specification we decided to use 16GB of data, specified as a parameter to the $D S G E N$ tool used to create the CSV file. After loading the CSV into the database, the resulting size was 80GB (this is probably due to overhead, indexes, etc.). System configuration. We ran the experiments on a Core i7 Machine equipped with 12 GB of RAM and running Linux Debian Jessie. We installed from source the three variants of the PostgreSQL DBMS: version 9.3.5, version 9.0.2 and hybrid 9.0.2 (i.e., with structural optimization). We used GCC 4.9.3 to compile the code. Moreover, we use the default configuration for all the versions. Before we give the results, it is convenient to show one of the query
templates involved in our testing.

```
select
    cc_call_center_id as Call_Center,
    cc_name as Call_Center_Name,
    cc_manager as Manager,
    sum(cr_net_loss) as Returns_Loss
from
    call_center,
    catalog_returns,
    date_dim,
    customer,
    customer_address,
    customer_demographics,
    household_demographics
where
            cr_call_center_sk = cc_call_center_sk
and cr_returned_date_sk = d_date_sk
and cr_returning_customer_sk= c_customer_sk
and cd_demo_sk = c_current_cdemo_sk
and hd_demo_sk = c_current_hdemo_sk
and ca_address_sk = c_current_addr_sk
and d_year = [YEAR]
and d_moy = [MONTH]
and ( (cd_marital_status = 'M' and cd_education_status = 'Unknown')
and hd_buy_potential like '[BUY_POTENTIAL]%'
and ca_gmt_offset = [GMT]
group by cc_call_center_id,cc_name,cc_manager,cd_marital_status,cd_education_status
order by sum(cr_net_loss) desc;
```

In particular, this query template consists of a join between 6 tables and aggregations. This template is of average complexity compared to others. Moreover not all the templates are compatible with PostgreSQL. Indeed we are able to execute 41 queries out of 99. The first set of experiments compares the performance of the regular 9.0.2 version with that of 9.3.5. This preliminary test will help us to better judge the effectiveness of our approach. Figure 5.4 shows the average computation times of the two versions over 41 queries. The speedup of PostgreSQL 9.3.5 compared to the regular version 9.0.2 is about 3 x . This is good news for us, because things get interesting when we turn on our structural optimization. We recall that our hybrid approach works only for counting queries, although it can easily be extended also to other aggregates. Therefore, we restrict the attention to queries involving only the count aggregate.


Figure 5.2: The speedup of PostgreSQL 9.3.5 compared to 9.0.2

### 5.4.1 Query 3 (Q3) on TPC-DS

```
select count(i_item_id)
from item, inventory, store_sales
where inv_item_sk = i_item_sk and inv_quantity_on_hand >= 100
and inv_quantity_on_hand <= 200 and ss_item_sk = i_item_sk
```

Query Q3 is a simple acyclic query consisting of three joins which asks to count all sold item_id from the store channel with an inventory quantity on-hand between 100 and 200. Note that we are not counting distinct types of item but we are counting all the occurrences (i.e., answers). At this point, the hybrid optimizer performs the set of operations that we discussed in the previous sections and yields to the following query rewrite.

```
SELECT sum(VT1P.OUTER_COUNT)
FROM (
    SELECT VT1.i_item_id,VT1.i_item_sk,VT1.OUTER_COUNT AS OUTER_COUNT,1 AS INNER_COUNT
    FROM
        (SELECT item.i_item_sk AS i_item_sk, item.i_item_id AS i_item_id, 1 AS OUTER_COUNT
            FROM store_sales, item
            WHERE item.i_item_sk=store_sales.ss_item_sk OFFSET 0
            ) VT1 SEMI JOIN inventory ON VT1.i_item_sk=inventory.inv_item_sk
            and inventory.inv_quantity_on_hand>=100
```



Figure 5.3: Execution time for query Q3

```
    and inventory.inv_quantity_on_hand<=200 OFFSET 0
) VT1P
```

We observe that in the query rewrite, we exploit the ad-hoc SEMI JOIN operator in addition to the nested views in which we store the counters which we need to compute the solutions. The comparison of the execution times of the three variants is shown in Figure 5.4.1.

In order to assess, in a deeper and more systematic manner, the capability of our approach to exploit both structural and statistics information, the tests were carried out on the same scenario while using different values for the attribute inv_quantity_on_hand.

We observe that hybrid 9.0.2 has the lowest runtime followed closely by regular 9.3.5. In particular a noticeably boost seems to be achieved when the number of computed solutions increases Figure 5.4.1. The reason is that the regular optimizer produces large intermediate results, as shown by the "explain analyze".

```
QUERY PLAN
Aggregate (cost=2732592.31..2732592.32 rows=1 width=17)
(actual time=302591.508..302591.509 rows=1 loops=1)
    -> Hash Join (cost=214476.16..618195.75 rows=845758624 width=17)
        (actual time=11291.097..204287.431 rows=449493233 loops=1)
        Hash Cond: (inventory.inv_item_sk = item.i_item_sk)
```



Figure 5.4: Number of solutions against different values for the attribute inv_quantity_on_hand

```
-> Seq Scan on inventory (cost=0.00..239309.00 rows=58725 width=4)
    (actual time=2.414..7619.955 rows=2241403 loops=1)
    Filter: ((inv_quantity_on_hand >= 0) AND (inv_quantity_on_hand <= 200))
-> Hash (cost=160187.11..160187.11 rows=2880404 width=25)
    (actual time=11282.395..11282.395 rows=2880404 loops=1)
    Buckets: 2048 Batches: 256 Memory Usage: 999kB
    -> Hash Join (cost=1742.00..160187.11 rows=2880404 width=25)
                    (actual time=112.532..8689.848 rows=2880404 loops=1)
            Hash Cond: (store_sales.ss_item_sk = item.i_item_sk)
            -> Seq Scan on store_sales (cost=0.00..88240.04 rows=2880404 width=4)
                            (actual time=1.782..4777.135 rows=2880404 loops=1)
            -> Hash (cost=1411.00..1411.00 rows=18000 width=21)
                            (actual time=110.520..110.520 rows=18000 loops=1)
                Buckets: 2048 Batches: 2 Memory Usage: 529kB
                -> Seq Scan on item (cost=0.00..1411.00 rows=18000 width=21)
                    (actual time=1.232..89.643 rows=18000 loops=1)
```

Total runtime: 302591.836 ms

### 5.4.2 Query 4 (Q4) on TPC-DS

Query (Q4) is an example of knowledge extraction in the TPC-DS decision support workload. It asks for all the distinct goods from sales comprising both a good of category "Women" and a good of category "Jewelry".

```
select count(distinct i3.i_item_id)
from item i1, store_sales s1, item i2, store_sales s2, item i3, store_sales s3
where s1.ss_item_sk = i1.i_item_sk and s2.ss_item_sk = i2.i_item_sk and
    s1.ss_ticket_number = s2.ss_ticket_number and i3.i_item_sk = s3.ss_item_sk
    and s1.ss_ticket_number = s3.ss_ticket_number and i1.i_category='Jewelry'
    and i2.i_category='Women'
```



Figure 5.5: Execution time for query Q4

This query is an acyclic query involving 6 joins. The total number of distinct answers for this query is only 8952 . Figure 5.4 .2 shows the performance the three optimizers for this query. Interestingly, when counting distinct solutions, even if the number of solutions is small, the advantage of hybrid 9.0.2 over regular 9.3.5 is emphasized. To know why this happens, we take again a closer look at the cost of the query plan generated by the 9.3 .5 optimizer reported in Figure ?? and we see that the bottleneck for the 9.3.5 is given by the cost to generate the large intermediate results (i.e., item join store_sales). Thus acyclicity plays a crucial role in this setting which goes well beyond all the gains obtained from the sophisticated quantitative machinery that traditional DBMS often perform.

## Part II

## CSP and Game Theory

## CHAPTER 6

## An Introduction to Game Theory

Game theory is one of the most important field in economy and mathematics. It is concerned with analyzing the strategic interactions that occur in systems of intelligent agents, both in noncooperative (where agents are pursuing their individual goals) and in cooperative contexts (where they jointly pursue some common goal). Noncooperative game theory, model those situations where players make individual choices since they are unable to cooperate with other players (or without additional benefits), and thus can apply even in situations where the interests of different agents are not in conflict. In this context, a prominent role is assumed by the concept of Nash equilibrium. This concept, that is certainly the most influential solution concept to date in the field of game theory, model in fact the rational behavior of players whom, having available a set of possible strategies with associated utility functions (that establish the payoffs that each player receives), reach a stable condition where each of them has selected an optimal strategy (with respect to the action of any other player) and thus has no incentive to deviate from it. Again, even if she knew what strategies the other agents were following, she would not want to change her strategy. Deciding whether a game admits a Nash equilibrium and eventually computing it efficiently, are problems of great interest to the scientific community. In
the case of games with mixed strategies, i.e., regulated by probability distributions, the existence of an equilibrium is guaranteed by the famous Nash's theorem, and recently the computational complexity of computing the Nash equilibrium has been fully characterized. However, in this thesis our studies are devoted to cooperative games, which have been introduced by the AI community to provide a solid mathematical framework to study scenarios where agents can obtain higher worths by collaborating with each other rather than by acting in isolation (see, e.g., [Nisan et al., 2007; Osborne and Rubinstein, 1994]). However, in many real-world application, not all the outcomes are equally fair (or rational), and as a consequence we need suitable mechanisms for coalition formation such that no group of agents would earn more by leaving the coalition for another potential coalition. In other words, we require coalition formation mechanisms to enforce the agreements between agents. In this chapter we proceed as follows. First, we formally define the most widely used model of coalitional games and we discuss a special class of coalitional games known as allocation games. Then we study coalition formation and we analyze how the payoff obtained by the coalition can be divided fairly among the agents of that coalition. This is addressed by discussing the notion of solution concepts; we describe the most important solution concepts used in cooperative game theory by focusing on the concept of Shapley value. We conclude the chapter with a real-world example of allocation game.

### 6.1 Coalitional games

Coalitional (i.e., cooperative) games provide a rich mathematical framework to analyze interactions between intelligent agents. Before we start studying these games, it is helpful to formalize a summary of relevant concepts as follows. In abstract
terms, a coalitional game $\mathcal{G}$ is a tuple $\langle N, v\rangle$, where $N$ is a set of agents and $v$ is a function associating each coalition $C \subseteq N$ with the worth that agents in $C$ can guarantee to themselves. In this thesis we restrict our attention to coalition games with transferable utilities. In the Transferable Utility (TU) setting, coalition worths can be freely distributed within a coalition, while in the non-Transferable Utility setting (NTU) coalitions are allowed to distribute worths only in some specified configurations, called consequences [Osborne and Rubinstein, 1994]. In fact, the crucial problem is to single out the most desirable distributions (of the worth associated with the grand-coalition $N$ ), usually called solution concepts, which can be perceived as fair and stable. Therefore, Coalitional games can be formalized as tuples $\mathcal{G}=\langle N, v\rangle$ where each coalition $C \subseteq N$ is associated with a real value $v(C)$ meant to encode the worth that agents in $C$ obtain by collaborating with each other. The function $v$ is supermodular (resp., submodular) if $v(R \cup T)+v(R \cap T) \geq v(R)+v(T)$, (resp., $v(R \cup T)+v(R \cap T) \leq v(R)+v(T))$ holds for each pair of coalitions $R, T \subseteq N$. After this brief formal description, we recall that in a coalition game, we are not interested in finding how agents within a coalition act, that is we are not concerned with the detail of their agreements, rather we look only at the payoff taken by the coalition. Thus we would like to answer these two basic questions:

1. Among all possible coalitions, which coalition will be formed?
2. How that coalition should distribute its worth among its members ?

Often, the answer to (1) is the "grand coalition", i.e, the coalition that includes all the agents in $N$. However, in many situations this answer depends also on the choice that has been made by taking into consideration (2).

Example 6.1. As an example of a coalitional game, consider the voting game. We have a set of agents $N$ and for each coalition $C \subseteq N, C$ is a winning coalition (i.e.,
group of agents who can come together to "win"), if it is sufficient to pass a bill into a law. Coalitions that are not willing to pass the law are called losing coalitions. For a winning coalition $C$ we have $v(C)=1$, and otherwise $v(C)=0$. Therefore, every "single" agent coalition is losing while the grand coalition is winning.

Example 6.2. As another example, consider that there is a research project and you can choose with which colleague you will work. Then a number of questions may come to mind: for instance, do I want someone who is easy to work with or do I prefer a talented colleague, etc.? Of course, every colleague will have the same reasoning. Therefore, a condition that must be satisfied in order to form a coalition is that all the agents inside that coalition are willing to stay. Since no one has an interest to leave, we can say that the coalition that has formed is stable. It follows from this discussion that stability is a necessary condition. A natural question is whether stability is also a sufficient condition. We will answer this question in Section 6.4.

In coalitional game theory, an interesting property often assumed for the function $v$ is the notion of superadditivity. Formally $v$ is superadditive if $v(S \cup T) \geq v(S)+$ $v(T), \forall S, T \subset N, S \cap T=\emptyset$. In words, this property says that cooperation leads to higher utilities for each player than the individual action. This is true if coalitions can work without interfering with one another. Note that, every game can be transformed to its superadditive form with the intuition that in the worst case agents can agree to not cooperate in order to obtain greater payoffs. Note that superadditivity has an important implication for how agents will form coalitions. In fact, the grand coalition will always form since its utility is higher than the utility of any other coalition. Moreover, if coalitions can never affect one another, either positively or negatively then we have another form of game called additive games. Formally, a game is additive if $v(S \cup T)=v(S)+v(T), \forall S, T \subset N, S \cap T=\emptyset$. In the rest of this


Figure 6.1: Allocation scenario $\mathcal{A}_{0}$ in Example 6.3.
chapter and in later chapters, when we refer to the function $v$, we will be making the assumption that it is superadditive.

In the next section we restrict the attention to a class of coalitional games called allocation games.

### 6.2 Allocation games

Among the various classes of coalitional games, we consider the class of allocation games, which is a setting to analyze fair division problems where monetary compensations are allowed and utilities are quasi-linear [Moulin, 1992]. Allocation games naturally arise in various application domains, ranging from house allocation to room assignment-rent division, to (cooperative) scheduling and task allocation, to protocols for wireless communication networks, and to queuing problems (see, e.g., [Moulin, 1992; Maniquet, 2003; Mishra and Rangarajan, 2007; ?] and the references therein). In this thesis, we analyze the setting where, we are given an allocation scenario $\mathcal{A}$ comprising a set of goods and a set of agents, and each agent is to be assigned at most one good she is interested in. Each good $g$ has a value $\operatorname{val}(g) \in \mathbb{R}$ and the worth $v_{\mathcal{A}}(C)$ associated with any coalition $C \subseteq N$ is the maximum overall value that can be obtained over the assignments to agents in $C$ only, also called allocations, hereinafter.

Example 6.3. As an illustrative example, consider the allocation scenario $\mathcal{A}_{0}$ that is reported in Figure 6.1, by using an intuitive graphical notation. We have a set $\left\{g_{1}, g_{2}, g_{3}, g_{4}\right\}$ of goods that have to be allocated to three agents. Each edge connects an agent to a good she is interested in. Edges in bold identify an optimal allocation, i.e., a feasible allocation whose sum of values of the allocated goods is the maximum possible one. The value of this allocation is $\operatorname{val}\left(g_{1}\right)+\operatorname{val}\left(g_{2}\right)+\operatorname{val}\left(g_{3}\right)=3+2+1=6$.

For each $C \subset\{1,2,3\}$ with $C \neq \emptyset$, an optimal allocation restricted to the agents in $C$ is also reported. Then, the associated coalitional game is $\mathcal{G}_{\mathcal{A}_{0}}=\left\langle\{1,2,3\}, v_{\mathcal{A}_{0}}\right\rangle$, where $v_{\mathcal{A}_{0}}(\{1,2,3\})=6, v_{\mathcal{A}_{0}}(\{1,2\})=5, v_{\mathcal{A}_{0}}(\{1,3\})=v_{\mathcal{A}_{0}}(\{2,3\})=4, v_{\mathcal{A}_{0}}(\{1\})=$ $v_{\mathcal{A}_{0}}(\{2\})=3$, and $v_{\mathcal{A}_{0}}(\{3\})=1$.

Allocation Scenario. Assume that a set $\mathbb{G}$ of goods have to be allocated to a set $N=\{1, \ldots, n\}$ of agents. Each good $g \in \mathbb{G}$ is associated with a real value $\operatorname{val}(g) \in \mathbb{R}$, and each agent $i \in N$ can receive at most one good taken from her set of interest $\Omega(i) \subseteq \mathbb{G}$. The tuple $\mathcal{A}=\langle N, \mathbb{G}, \Omega$, val $\rangle$, with $\Omega: N \rightarrow 2^{\mathbb{G}}$ and val : $\mathbb{G} \rightarrow \mathbb{R}$, is an allocation scenario. Moreover, goods are indivisible and unshareable. Hence, an allocation for $\mathcal{A}$ is a function $\pi: N \rightarrow \mathbb{G} \cup\{\emptyset\}$ such that: (1) for each agent $i \in N, \pi(i) \neq \emptyset$ implies $\pi(i) \in \Omega(i)$; and (2) for each pair $i, i^{\prime} \in N$ with $i \neq i^{\prime}$, $\pi(i) \cap \pi\left(i^{\prime}\right)=\emptyset$ holds. We denote by $\operatorname{img}(\pi)$ the set of all goods in the image of $\pi$, i.e., $\operatorname{img}(\pi)=\{\pi(i) \mid i \in N \wedge \pi(i) \neq \emptyset\}$.

By slightly abusing of notation, if $S \subseteq \mathbb{G}$ is a set of goods, then $\operatorname{val}(S)$ denotes the sum of their values. Moreover, if $\pi$ is an allocation, then $\operatorname{val}(\pi)$ denotes the value of $\operatorname{img}(\pi)$. An allocation $\pi$ is optimal (w.r.t. $\mathcal{A})$ if $\operatorname{val}(\pi) \geq \operatorname{val}\left(\pi^{\prime}\right)$ holds, for each allocation $\pi^{\prime}$. The value associated with any optimal allocation w.r.t. $\mathcal{A}$ is hereinafter denoted as $\operatorname{opt}(\mathcal{A})$.

Let $\mathcal{A}=\langle N, \mathbb{G}, \Omega$, val $\rangle$ be an allocation scenario and let $C \subseteq N$ be a set of agents.

The restriction of $\mathcal{A}$ to $C$ is the sub-scenario $\mathcal{A}[C]=\left\langle C, \mathbb{G}, \Omega_{C}\right.$, val $\rangle$ where $\Omega_{C}$ is the restriction of $\Omega$ over $C$. The allocation game induced by $\mathcal{A}$ is the tuple $\mathcal{G}_{\mathcal{A}}=\left\langle N, v_{\mathcal{A}}\right\rangle$, where $v_{\mathcal{A}}: 2^{N} \rightarrow \mathbb{R}$ is such that $v_{\mathcal{A}}(C)=\operatorname{opt}(\mathcal{A}[C])$, for each $C \subseteq N$. The value of an empty set of goods is 0 . Then, the definition trivializes for $C=\emptyset$, with $v_{\mathcal{A}}(\emptyset)=0$. Moreover, note that $v_{\mathcal{A}}(C) \geq 0$ holds, for each $C \subseteq N$, since the allocation where no agent receives some good is a feasible one.

We are now ready to analyze the most important solution concepts used in coalitional game theory. Among all the solution concepts we will focus on the Shapley value, concluding the chapter with an application of this solution concept to a real allocation problem arising in a research assessment program in Italy described by [Greco and Scarcello, 2013b].

### 6.3 Solution concepts

As we said, a fundamental problem for a coalitional game $\mathcal{G}=\langle N, v\rangle$ is to single out the most desirable outcomes, usually called solution concepts, in terms of appropriate notions of worth distributions, i.e., of payoff vectors of the form $\left(x_{1}, \ldots, x_{|N|}\right) \in \mathbb{R}^{|N|}$ where $x_{i}+\cdots+x_{|N|}$ equals the worth associated with the whole set $N$ of agents. A payoff vector $x$ is called an imputation if it is individual rational i.e., for all players $i \in N$ we have $x_{i} \geq v(i)$, and efficient i.e., $\sum_{i \in N} x_{i}=v(N)$. Efficiency in this context means that the whole available worth is distributed while individual rationality indicates that if $x_{i} \geq v(i)$ the player $i$ opts for participating since he would never do better going on his own. In the previous sections, we have seen that in superadditive games, we can always expect the grand coalition to form and thus the real issue is just that of the division of the whole available worth among the players. Many solution concepts have been considered in the coalitional game
theory and thoroughly studied through the years, and among all, the most known and widely-accepted are the core, kernel, bargaining set, nucleolus and the Shapley value. Below, we recall the notions of core, kernel, bargaining set, and nucleolus and we briefly discuss each of them.

### 6.3.1 The Core

Throughout the chapter we have emphasized the concept of stability several times. Stability leads to the question of whether the agents would be always willing to form the grand coalition given any payment profile. The answer to this question is given by the concept of core [Edgeworth, 1881; Gillies, 1959]. The core of a coalitional game is the set of all imputations $x$ to which there not exists an objection. Formally, it is defined as follows.

Definition 6.4. Let $\mathcal{G}=\langle N, v\rangle$ be a coalitional game, and let $x$ be an imputation. An imputation $x$ is in the core of $\mathcal{G}$, if and only if

$$
\forall S \subseteq N, \sum_{i \in S} x_{i} \geq v(S)
$$

That is, the sum of the payoffs given to any group of agents must be always greater or equal to the sum of the payoffs that these agents are able to obtain independently. Consequently, we say that an imputation $x$ in the core is "stable" because there is no coalition whose members will receive a higher payoff than in $x$ by leaving the grand-coalition.

Example 6.5. As an example, consider let $\mathcal{G}=\langle N, v\rangle$ be a coalitional game, with the caracheristic function $v$ given by

$$
\begin{array}{ll}
v(\{a\})=0 & v(\{a, b\})=20 \\
v(\{b\})=0 & v(\{a, c\})=30 \\
v(\{c\})=0 & v(\{b, b\})=40
\end{array}
$$

The set of imputations for this game are the triples $\left(x_{a}, x_{b}, x_{c}\right)$ such that $x_{a}+x_{b}+$ $x_{c}=45$ and $x_{a} \geq 0, x_{b} \geq 0, x_{c} \geq 0$. For instance, if we consider the imputation $x$ such that: $x_{a}=5, x_{b}=15, x_{c}=25$ it is easily seen that $x$ is in the core of $\mathcal{G}$.

Given that the core provides a concept of stability, it is a very attractive solution concept. However, there are games for which the core is empty.

Example 6.6. Let $\mathcal{G}=\langle N, v\rangle$ be a coalitional game, with $N=\{a, b, c\}, v(\{a\})=$ $v(\{b\})=v(\{c\})=0, v(\{a, b\})=20, v(\{a, c\})=30, v(\{b, c\})=40, \operatorname{andv}(\{a, b, c\})=$ 42. Consider the imputation $x$ such that: $x_{a}=4, x_{b}=14$, and $x_{c}=24$. Since $v(\{b, c\})=40>38=x(\{b, c\})$ such imputation $x$ is not in the core of $\mathcal{G}$. In particular, we can show that for this game the core is empty. In fact, observe that $v(\{a, b, c\})=42$ and consider the coalitions $S_{1}=\{a, b\}, S_{2}=\{a, c\}$ and $S_{3}=\{b, c\}$. Recall from the definition that an imputation $x$ must satisfy the following inequalities:

$$
\begin{aligned}
& x_{a}+x_{b} \geq 20 \\
& x_{a}+x_{c} \geq 30 \\
& x_{b}+x_{c} \geq 40
\end{aligned}
$$

It follows that $\left(x_{a}+x_{b}\right)+\left(x_{a}+x_{c}\right)+\left(x_{b}+x_{c}\right) \geq 20+30+40$, that is $x_{a}+x_{b}+x_{c} \geq 45$. Thus, the grand coalition should get a payoff of 45 instead of 42 in order to satisfy the agents. Consequently the grand coalition will never form for this game.
$\epsilon$-core. While the core can be empty, sometimes we may also be happy with a set of imputations that are approximately stable. That is, we relax the notion of core and this leads to the definition of $\epsilon$-core.

Definition 6.7. Let $\mathcal{G}=\langle N, v\rangle$ be a coalitional game, and let $x$ be an imputation. An imputation $x$ is in the $\epsilon$-core of $\mathcal{G}$, if and only if

$$
\forall S \subseteq N, \sum_{i \in S} x_{i} \geq v(S)-\epsilon
$$

Example 6.8. Let $\mathcal{G}=\langle N, v\rangle$ and $N=\{a, b, c\}$. Let $v(S)=1$ if $|S|>1, v(S)=0$ otherwise. Then it is easy to see that $(1 / 3,1 / 3,1 / 3)$ is an imputation and that is in $1 / 3$-core. Therefore $1 / 3$-core is non-empty while $\epsilon$-core is empty for any $\epsilon<1 / 3$.

The interpretation of the $\epsilon$-core is that for any coalition there is at most an $\epsilon$ cost to deviate from the grand coalition. As a result, if the deficit $v(S)-x(S)$ is less than $\epsilon$, the imputation $x$ is still stable. Of course, we are interested in outcomes that minimize the $\epsilon$.

Least core. From the last example we note that even the $\epsilon$-core may be empty for some $\epsilon$. However, there always exists some $\epsilon$ that is large enough to ensure that the $\epsilon$-core is non-empty. Then, the natural problem is to find the smallest $\epsilon$ for which the $\epsilon$-core is non-empty. This leads to the definition of the concept of the least core.

Definition 6.9. Let $\mathcal{G}=\langle N, v\rangle$ be a coalitional game, and let $x$ be an imputation. An imputation $x$ is in the $\epsilon$-core of $\mathcal{G}$, if and only if $x$ is the solution the following linear program.

$$
\begin{array}{ll}
\operatorname{minimize} & \epsilon \\
\text { subject to } & \sum_{i \in S} x_{i} \geq v(S)-\epsilon \quad \forall S \subseteq N
\end{array}
$$

Intuitively, the least core consists of all payoffs that minimize the incentive of all coalitions to deviate. Thus, it can be considered a refinement of the core. Observe also that if $\epsilon$ is large enough, the least core is always non-empty but it may contain
more than one imputation. Therefore, we would like to identify the most stable point in the least core. We pursue this objective, with the definition of the nucleolus.

### 6.3.2 The Nucleoulus

Given an imputation $x$, the least core minimize the maximum inequity $\epsilon(S, x)=$ $v(S)-x(S)$ (also called deficit or excess). This quantity is a measure of the dissatisfaction of $S$ with the proposed imputation $x$. The nucleolus of a game $\mathcal{G}$ not only minimize the worst inequity but given this, it also minimizes the second largest inequity. Intuitively, we look first for the coalition with the maximum degree of dissatisfaction then we try to make it smaller if possible by adjusting $x$. Once the largest dissatisfaction has been made as small as possible, then we address the second largest dissatisfaction and we try to adjust $x$ to lower it as much as possible. This is formalized as follows.

Definition 6.10. Let $\mathcal{G}=\langle N, v\rangle$ be a coalitional game, and let $x$ be an imputation. An imputation $x$ is in the nucleoulus of $\mathcal{G}$, if and only if $x$ is the solution to series of optimization programs $O_{1}, O_{2}, \ldots, O_{|N|}$ defined as follows.

$$
\begin{gathered}
\left(O_{1}\right)= \begin{cases}\text { minimize } \epsilon \\
\text { subject to } & \sum_{i \in S} x_{i} \geq v(S)-\epsilon\end{cases} \\
\left(O_{i}\right)= \begin{cases}\text { minimize } \epsilon \\
\text { subject to } & \sum_{i \in S} x_{i} \geq v(S)-\epsilon \\
\vdots \\
\sum_{i \in S} x_{i} \geq v(S)-\epsilon_{i-1} & \forall S \subseteq S_{i-1} S_{i-2} \\
\sum_{i \in S} x_{i} \geq v(S)-\epsilon & \forall S \subseteq 2^{N} S_{i-1}\end{cases}
\end{gathered}
$$

Theorem 6.11. For any coalitional game $\mathcal{G}(N, v)$, the nucleolus of $\mathcal{G}$ always exists and is unique.

### 6.3.3 The Kernel

The kernel is a solution concept introduced by [Davis and Maschler]. Let us start by recalling its formal definition. For any pair of players $i$ and $j$ of a coalitional game $\mathcal{G}=\langle N, v\rangle$, we denote by $I_{i, j}$ the set of all coalitions containing player $i$ but not player $j$. Recall that $\epsilon(S, x)=v(S)-x(S)$. Define the surplus $s_{i, j}(x)$ of player $i$ against player $j$ at an imputation $x$ as the value si, $j(x)=\max _{S \in I_{i, j}} e(S, x)=$ $\max _{S \in I_{i, j}}(v(S)-x(S))$.

Definition 6.12. Let $\mathcal{G}=\langle N, v\rangle$ be a coalitional game. The kernel of $\mathcal{G}$ is the set of all imputations $x$ such that $s_{i, j}(x)>s_{j, i}(x) \Rightarrow x_{j}=v(j), \forall i, j \in N, i \neq j$.

Intuitively, the surplus of player $i$ against $j$ at $x$ is the highest payoff that player $i$ can gain (or the minimal amount $i$ can lose, if it is a negative value) without the cooperation of $j$, by assuming to form coalitions with other players that are satisfied at $x$; thus, $s_{i, j}(x)$ is the weight of a possible threat of $i$ against $j$. In particular, player $i$ has more bargaining power than $j$ at $x$ if $s_{i, j}(x)>s j, i(x)$; however, player $j$ is immune to such threat whenever $x_{j}=v(j)$, since in this case $j$ can obtain $v(j)$ even by operating on her own. We say that player $i$ outweighs player $j$ at $x$ if $s_{i, j}(x)>s_{j, i}(x)$ and $x_{j}>v(j)$. Therefore, the kernel is the set of all imputations where no player outweighs another one.

Example 6.13. Let $\mathcal{G}=\langle N, v\rangle$ be a coalitional game and let $N=\{a, b, c\}, v(\{a\})=$ $v(\{b\})=v(\{c\})=0, v(\{a, b\})=20, v(\{a, c\})=30, v(\{b, c\})=40$, and $v(\{a, b, c\})=$ 42. It is easily verified that the imputation $x$ such that $x_{a}=4, x_{b}=14$, and $x_{c}=24$ is in the kernel of $\mathcal{G}$. Indeed, we note first that every player in $N$ receives in $x$ a payoff
strictly greater than what she is able to obtain by acting on her own. For this reason, in order for $x$ to belong to the kernel of $\mathcal{G}$ it must be the case that $s_{i, j}(x)<s_{j, i}(x)$, for all distinct players $i$ and $j$. By the definition of the worth function, the maximum excess that a coalition $S$ including $i$ and excluding $j$ can achieve is obtained by the coalition $S \in I_{i, j}$ such that $|S|=2$. By this, $s_{i, j}(x)=s_{j, i}(x)=2$ for all pairs of different players $i, j$. Thus, the imputation $x$ is in the kernel of $\mathcal{G}$.

### 6.3.4 The Bargaining set

The concept of bargaining set was defined by [Aumann and Maschler, 1964]. As usual, we recall its formal definition and then we illustrate the concept with an example.

Definition 6.14. Let $\mathcal{G}=\langle N, v\rangle$ be a coaliitonal game, and $x$ be an imputation. Let $S \subseteq N$ be a coalition, and $y$ be an $S$-feasible payoff vector (i.e., $y(S)=v(S)$ ). The pair $(y, S)$ is an objection of player $i$ against player $j$ to $x$ if $i \in S, j \in S$, and $y_{k}>x_{k} \forall k \in S$. A counterobjection to the objection $(y, S)$ of $i$ against $j$ to $x$ is a pair $(z, T)$ where $j \in T, i \in T$, and $z$ is a $T$-feasible payoff vector such that $z_{k} \geq x_{k}$ $\forall k \in T \backslash S$ and $z_{k} \geq y_{k} \forall k \in T \cap S$. If there does not exist any counterobjection to $(y, S)$, we say that $(y, S)$ is a justified objection.

The bargaining set of a calitional game $\mathcal{G}$ is the set of all imputations $x$ to which there is no justified objection.

Example 6.15. Let $\mathcal{G}=\langle N, v\rangle$ be a coalitional game with $N=\{a, b, c\}, v(\{a\})=$ $v(\{b\})=v(\{c\})=0, v(\{a, b\})=20, v(\{a, c\})=30, v(\{b, c\})=40$, and $v(\{a, b, c\})=$ 42. Consider the imputation $x$ such that $x_{a}=8, x_{b}=10$, and $x_{c}=24$. An objection of player $c$ against player $a$ to $x$ is $((12,28),\{b, c\})$. Player $a$ can counterobject to this objection using $((8,12),\{a, b\})$. Another objection of player $c$ against player $a$
to $x$ is $((14,26), b, c)$. In this case, player a cannot counterobject. The reason is that coalition $\{a, b\}$ receives a payoff 20 and this is not sufficient for player a to counterobject since she needs at least 8 for herself and at least 14 for player $b$, in order to respond to the proposal of player $c$. Therefore the imputation $x$ does not belong to $B(G)$. The intuitive reason is that player a receives too much, according to this profile. Consider now the imputation $x^{\prime}$ such that $x_{a}^{\prime}=4, x_{b}=14$, and $x_{c}=24$. We focus on the objections of player a against player $c$. We note that, in order to object, player $a$ has to form the coalition $S=\{a, b\}$. The excess $e(S, x)$ of $S$ at $x$ is 2, hence players $a$ and $b$ have the possibility to distribute among themselves a payoff of 2 to make the objection. But player $c$ can always counterobject to player $a$ because she can form the coalition $T=\{b, c\}$ whose excess at $x$ is 2 and hence she can always match the proposal made to player $b$ by player $a$ in order to object. A similar argumentation holds for every objection of every player against any other. Thus $x$ belongs to the bargaining set of $\mathcal{G}$. $\triangleleft$

Observe that the core is the set of all imputations that do not admit an objection. Hence core $\subseteq$ bargaining set. To conclude our analsys on the most important solution concepts in game theory, we have to consider the Shapley value (see, e.g., [Shapley, 1953]) to which we will dedicate the next section.

### 6.4 The Shapley value

The core defined a way to form coalition that are stable. However it does not make fairness considerations. In fact, outcomes in the core may be really unfair. For instance consider the following example.

Example 6.16. Let $\mathcal{G}=\langle N, v\rangle$ be a coalitional game, with $N=a, b, v(a)=v(b)=$ $5, v(a, b)=20$. Consider the imputation $x$ such that: $x_{a}=20, x_{b}=5$. We can easily
check that $x$ is in in the core of $\mathcal{G}$, in fact both players cannot benefit by deviating from the grand coalition. However, the outcome is really unfair for $b$ since she is perfectly symmetric with $a$.

Consequently, we answer the question posed in Section 6.1 negatively. Therefore, we would like to find a mechanism that allow us to divide the total payoff among the members of the grand coalition in a fair way. It turns out that the answer to this problem is given by The Shapley value [Shapley, 1953]. Crudely, the Shapley value of each player $i$ is given by the average marginal contribution of $i$ over all possible coalitions (s)he may participate in. Thus, while the core formalize the concept that the coalition has to be stable, the Shapley value provides a unique way to distribute coalitional payoff among agents in such a way that satisfy various fairness criteria. Let us start by axiomatizing the concept of fairness in our context.

The Shapley Axioms for the characteristic function $v$.

1. $\sum_{i \in N} v_{i}(G)=v(N)$ (efficiency);
2. For any $v$, if $i$ and $j$ are such that $v(S \cup i)=v(S \cup j)$ for every coalition $S$ not containing $i$ and $j$, then $\phi_{i}(\mathcal{G})=\phi_{j}(\mathcal{G})$ (simmetry);
3. For any $v$, if $i$ is a dummy player then $\phi_{i}(\mathcal{G})=v(i)$ (dummy player);
4. If $u$ and $v$ are characteristic functions, then $\phi(u+v)=\phi(u)+\phi(v)$ (additivity);
5. If $v$ is supermodular, the Shapley value belongs to the core of the game (stability);
6. If $G^{\prime}=\left(N, v^{\prime}\right)$ is a game such that $v^{\prime}(C) \geq v(C), \forall C \subseteq N$, then $\phi_{i}\left(G^{\prime}\right) \geq$ $\phi_{i}(G), \forall i \in N$ (monotonicity).

Theorem 6.17. Given a coalitional game $\mathcal{G}=\langle N, v\rangle$, there is a unique imputation that satisfies the Shapley axioms and it is called the Shapley value.

As we already stated, the Shapley value is based on the value that each player is able to add to the possible coalitions, i.e., on its marginal contribution. Suppose we form the grand coalition by entering the players into this coalition one at a time. As each player enters the coalition, he receives the amount by which his entry increases the value of the coalition he enters. The amount a player receives by this scheme depends on the order in which the players are entered. The Shapley value is just the average payoff to the players if the players are entered in completely random order.

Definition 6.18. It turns out that the payoff associated with each agent $i \in N$ is given by

$$
\phi_{i}(\mathcal{G})=\sum_{C \subseteq N \backslash\{i\}} \frac{|C|!(n-|C|-1)!}{n!}(v(C \cup\{i\})-v(C)) \text {, where } v(C \cup\{i\})-v(C)
$$

is the marginal contribution of $i$ to the coalition $C \cup\{i\}$.

Example 6.19. Consider again the allocation game $\mathcal{G}_{\mathcal{A}_{0}}$ introduced in Example 6.3. For agent 1, we have:

$$
\begin{aligned}
\phi_{1}\left(\mathcal{G}_{\mathcal{A}_{0}}\right)= & \frac{2!0!}{3!}\left(v_{\mathcal{A}_{0}}(\{2,3\} \cup\{1\})-v_{\mathcal{A}_{0}}(\{2,3\})\right)+ \\
& \frac{1!1!}{3!}\left(v_{\mathcal{A}_{0}}(\{2\} \cup\{1\})-v_{\mathcal{A}_{0}}(\{2\})\right)+ \\
& \frac{1!1!}{3!}\left(v_{\mathcal{A}_{0}}(\{3\} \cup\{1\})-v_{\mathcal{A}_{0}}(\{3\})\right)+ \\
& \frac{0!2!}{3!}\left(v_{\mathcal{A}_{0}}(\emptyset \cup\{1\})-v_{\mathcal{A}_{0}}(\emptyset)\right)=\frac{15}{6} .
\end{aligned}
$$

Similarly, we can derive $\phi_{2}\left(\mathcal{G}_{\mathcal{A}_{0}}\right)=\frac{15}{6}$ and $\phi_{3}\left(\mathcal{G}_{\mathcal{A}_{0}}\right)=\frac{6}{6}$.

### 6.4.1 The Banzhaf value

Another solution concept which is similar in spirit to the Shapley value is the Banzhaf value (or index) (see, e.g., [Banzhaf, 1965]) for which the payoff of $i$ is

$$
\beta_{i}(\mathcal{G})=\frac{1}{2^{n-1}} \sum_{C \subseteq N \backslash\{i\}}(v(C \cup\{i\})-v(C)) .
$$

In the next chapter we will analyze in detail the relationship between these two solution concepts.

### 6.4.2 Back to the allocation games

In the context of allocations games (and when monetary transfers are possible), the prototypical solution concepts considered in the literature are the Shapley value and the Banzhaf value, since they lead to outcomes enjoying desirable properties such as equal treatment of equals, envy-freeness, and a stronger property called individual-optimality guaranteeing that, for every agent, her/his utility is the maximum possible one over any alternative optimal allocation [?]. However, it is well known that, in general, computing such values is intractable, formally \#P-complete. This is a serious obstruction to their applicability in allocation scenarios involving many agents, and it motivates the design of approximation algorithms and the identification of subclasses of practical interest where exact computation can be carried out efficiently. In Chapter 7 we focus precisely on the latter approach. But before we do so, let us recall some basic properties about allocations games.

Basic Properties. We discuss below some properties pointed out by [?] and [Moulin, 1992]. Let $\mathcal{A}=\langle N, \mathbb{G}, \Omega$, val $\rangle$ be an allocation scenario, and let $C$ and $C^{\prime}$ be two coalitions with $C^{\prime} \subseteq C \subseteq N$. Then,
(allocation) monotonicity: $v_{\mathcal{A}}(C) \geq v_{\mathcal{A}}\left(C^{\prime}\right)$. Moreover, if $\pi$ is an optimal allocation for $\mathcal{A}[C]$, then there is an optimal allocation $\pi^{\prime}$ for $\mathcal{A}\left[C^{\prime}\right]$ such that $\operatorname{img}\left(\pi^{\prime}\right) \subseteq \operatorname{img}(\pi) ;$
submodularity: $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C) \leq v_{\mathcal{A}}\left(C^{\prime} \cup\{i\}\right)-v_{\mathcal{A}}\left(C^{\prime}\right)$, for each $i \in N \backslash C$.
Furthermore, it has be shown that, for these games, the Shapley value grants to every coalition at least its marginal contribution to the worth of the grand-coalition. Indeed, the Shapley value of $\mathcal{G}_{\mathcal{A}}$ is in the core of the game $\overline{\mathcal{G}}_{\mathcal{A}}=\left\langle N, \bar{v}_{\mathcal{A}}\right\rangle$, where $\bar{v}_{\mathcal{A}}(C)=v_{\mathcal{A}}(N)-v_{\mathcal{A}}(N \backslash C)$, for each coalition $C \subseteq N .{ }^{1}$ More formally, the following

[^3]holds.

Theorem 6.20. For each scenario $\mathcal{A}=\langle N, \mathbb{G}, \Omega$, val $\rangle$, we have that: $\sum_{i \in N} \phi_{i}\left(\mathcal{G}_{\mathcal{A}}\right)=$ $v_{\mathcal{A}}(N)$ and $\sum_{i \in C} \phi_{i}\left(\mathcal{G}_{\mathcal{A}}\right) \geq v_{\mathcal{A}}(N)-v_{\mathcal{A}}(N \backslash C)$, for each $C \subseteq N$.

Example 6.21. Note that $\phi_{1}\left(\mathcal{G}_{\mathcal{A}_{0}}\right)+\phi_{2}\left(\mathcal{G}_{\mathcal{A}_{0}}\right)+\phi_{3}\left(\mathcal{G}_{\mathcal{A}_{0}}\right)=\frac{36}{6}=v_{\mathcal{A}_{0}}(\{1,2,3\})$. Moreover, for each $C \subseteq N$, observe that $\sum_{i \in C} \phi_{i}\left(\mathcal{G}_{\mathcal{A}_{0}}\right) \geq 6-v_{\mathcal{A}_{0}}(N \backslash C)$. For instance, $\phi_{1}\left(\mathcal{G}_{\mathcal{A}_{0}}\right)+\phi_{3}\left(\mathcal{G}_{\mathcal{A}_{0}}\right)=\frac{21}{6} \geq 6-v_{\mathcal{A}_{0}}(\{2\})=6-\frac{15}{6}$.

### 6.5 A motivating example: The Italian Research Assessment Program (VQR)

The Italian Research Assessment Program (VQR in short) is aimed at evaluating the results of the Italian scientific research in the period 2011-2014 carried out by public universities, private universities, and other research institutes. The outcome of the evaluation will be used by the National Agency for the Evaluation of Universities and Research Institutes (ANVUR) to proportionally distribute the funds allocated by the Ministry to universities and research institutes in order to support their research activities in the next years (until the subsequent evaluation process). In addition to the institutions, the VQR will also evaluate departments of universities and similar internal structures of other research institutes. At this point, it is easy to see that every research structure $R$ is in competition with all other research structures, as every structure $R$ is then interested in selecting and submitting its best research products (e.g., publications). The program is articulated as follows:

1. Authors affiliated to $R$ are asked to self-evaluate their products, according to some evaluation criteria defined by groups of experts chosen by ANVUR. We suppose that every researcher is able to associate to each product a quality score.
2. Based on the self-evaluations being collected, $R$ is in charge of selecting the best products to submit to ANVUR (at most) products for each one of its authors. Note that a product cannot be allocated to two researchers.
3. ANVUR carries out its assessment on the quality of products submitted and the quality profile (VQR score) of $R$ is determined by the sum of the "true" scores associated with the products (i.e., those resulting by ANVUR actual evaluation). Eventually, $R$ will receive funds in the next years proportionally to this score. Moreover, by using such products scores, ANVUR evaluates all substructures, too (e.g., all departments, if $R$ is a University).

It is then natural for the structure $R$ to exploit the VQR scores of its substructures to divide funds among them. In particular, the outcomes could be in principle used to evaluate single researchers, too. In this chapter we will show that finding a good (fair) division rule for scores/funds distribution is not an easy task. Moreover, we pinpoint that, without and agreement on a fair division rule, the self-evaluation performed at point (1) is unfeasible, and the research structure may miss the optimal possible submission to ANVUR, hence losing funds. Before we point out how we can apply techniques from game theory to this case study, we present the simple division rules that are believed to be applied by the ANVUR and we show that they are unfair for substructures.

### 6.5.1 Division rules

Let $R$ be a research structure, and let $\mathcal{R}$ be the set of researchers affiliated with $R$. For each researcher $r \in R$, let $\operatorname{products}(r, R)$ (or just $\operatorname{products}(r)$, if $R$ is understood from the context) be the set of the research products of $r$ in the given period 20112014. An allocation for a set of researchers $S \subseteq R$ is a function $\psi$ mapping each
researcher $r \in S$ to a set of publications $\psi(r) \subseteq \operatorname{products}(r)$ with $|\psi(r)| \leq 3$ and with $\psi(r) \cap \psi\left(r^{\prime}\right)=\emptyset$, for each $r^{\prime} \in \mathcal{S} \backslash\{r\}$. An allocation for $R$ is an allocation for all researchers $\mathcal{R}$, while an allocation for a substructure $S$ of $R$ is an allocation for the researchers affiliated to $S$. In the VQR program, every research structure $R$ has to submit to ANVUR for its evaluation a set of products $\mathcal{P}_{\psi}$ such that $\mathcal{P}_{\psi}=\bigcup_{r \in \mathcal{R}} \psi(r)$, for some allocation $\psi$ for all researchers $\mathcal{R}$ affiliated to $R$. Then, for each $p \in \mathcal{P}_{\psi}$, ANVUR calculates a quality score $\operatorname{score} e_{\mathrm{VQR}}(p)$, so that $R$ will receive funds proportionally to its overall score $\operatorname{score}_{\mathrm{VQR}}(\psi)=\sum_{p \in \mathcal{P}_{\psi}} \operatorname{score}_{\mathrm{VQR}}(p)$.

While the first aim of the VQR program is to evaluate the various Italian research structures, it is known that the obtained information will be used to evaluate substructures, too (e.g., University departments). Thus, following the same principle of binding funds to VQR scores used for the main structure, it is natural to exploit such scores for money distribution inside every research structure. It is therefore of utmost importance the way VQR assigns scores to substructures. Nevertheless, as already mentioned, up to date there is no official information about such an algorithm. We argue that the score of any structure $R$ should fairly be distributed over its substructures (and possibly over individuals), in such a way to reflect their actual contribution to the result achieved by the structure $R$. Formally, we need a suitable division rule.

Definition 6.22. A division rule $\gamma$ for $R$ is a real-valued function that, given a researcher $r \in \mathcal{R}$ and an allocation $\psi$ for $R$, returns its score $\gamma_{\psi}(r, R) \geq 0$ with respect to the allocation $\psi$. By slightly abusing notation, for any substructure $\mathcal{S} \subseteq \mathcal{R}$, we denote by $\gamma_{\psi}(\mathcal{S}, R)$ the value $\sum_{r \in \mathcal{S}} \gamma_{\psi}(r, R)$. Whenever $R$ is understood from the context, we just write $\gamma_{\psi}(r)$ and $\gamma_{\psi}(\mathcal{S})$, in place of $\gamma_{\psi}(r, R)$ and $\gamma_{\psi}(S, R)$, respectively.

Hereafter, we assume that $\mathcal{S}_{1}, \ldots, \mathcal{S}_{n}$ are the substructures of $R$. They exhaus-
tively cover the researchers of $R$, i.e.,

$$
\left\{\begin{array}{l}
\bigcup_{i=1}^{n} \mathcal{S}_{i}=\mathcal{R} \\
\mathcal{S}_{i} \cap \mathcal{S}_{j}=\emptyset, \forall i \neq j
\end{array}\right.
$$

Note that such a division rule may naturally be used for evaluating individuals, besides substructures. In fact, we shall discuss all fairness properties with reference to individual researchers, because all issues about individuals immediately extend to the substructures they belong to.

### 6.5.2 Desiderata for division rules

In the absence of an official division rule, most researchers believe that the score of any substructure $S$ will be based on the naive proj rule where, for any researcher $r, \operatorname{proj}_{\psi}(r)$ is the sum of the VQR scores of the products allocated to $r$ in $\psi$, i.e.,

$$
\operatorname{proj}_{\psi}(r)=\sum_{p \in \psi(r)} \operatorname{score}_{\mathrm{VQR}}(p)
$$

This rule satisfies a very basic requirement for every division rule, which we state below.

Definition 6.23 (Budget-balance). A division rule $\gamma$ must precisely distribute the VQR score of $R$ over all its members, i.e., $\sum_{r \in \mathcal{R}} \gamma_{\psi}(r)=\operatorname{score}_{\mathrm{VQR}}(R)$. Clearly, because the substructures of $R$ define a partition of its researchers, this implies that $\gamma$ completely distribute the VQR score of $R$ over all its substructures and does not distribute more than that, i.e., $\sum_{i=1}^{n} \gamma_{\psi}\left(\mathcal{S}_{i}\right)=\operatorname{score}_{\mathrm{VQR}}(R)$.

However, this rule is hardly perceivable as a fair one. Indeed, proj might lead to scenarios where some researcher $r$ (and in turn her/his substructure) has reasonable arguments against her/his structure because of possible alternative allocations where
$r$ would get higher scores. Thus, proj fails to satisfy a basic and intuitive requirement of fairness, which we formalize as follows.

Definition 6.24 (Fairness). Assume that $\psi$ and $\bar{\psi}$ are two allocations of the same research products, i.e., $\mathcal{P}_{\psi}=\mathcal{P}_{\bar{\psi}}$. A division rule $\gamma$ must be indifferent w.r.t. the optimal allocation being selected, i.e., for each researcher $r \in \mathcal{R}, \gamma_{\psi}(r)=\gamma_{\bar{\psi}}(r)$. In fact, this implies that, for each substructure $\mathcal{S}_{i}, \gamma_{\psi}\left(\mathcal{S}_{i}\right)=\gamma_{\bar{\psi}}\left(\mathcal{S}_{i}\right)$ also holds.

The above criterion is definitely desirable. However, a closer look reveals that it is not enough, as there is a trivial way to circumvent this kind of fairness: just consider the rule trivial assigning the overall VQR score to one fixed researcher, say $r_{1}$ (independently of its actual contribution), i.e.,

$$
\operatorname{trivial}_{\psi}(r)= \begin{cases}\sum_{p \in \mathcal{P}_{\psi}} \operatorname{score}_{\mathrm{VQR}}(p), & \text { if } r=r_{1} \\ 0 & \text { if } r \in \mathcal{R} \backslash\left\{r_{1}\right\}\end{cases}
$$

A problem with this rule is that it is not symmetric, in that a researcher $r_{2}$ that has co-authored precisely the same set of products as $r_{1}$ would be treated differently, just because of her name. Avoiding these cases and guaranteing an equal treatment of the equals is another very basic requirement, formalized as follows.

Let $\pi: \mathcal{R} \mapsto \mathcal{R}$ be a permutation of the researchers in $\mathcal{R}$. Let $R_{\pi}$ be the research structure over the researchers in $\mathcal{R}$ where, for each $r \in \mathcal{R}, \operatorname{products}\left(r, R_{\pi}\right)=$ $\operatorname{products}(\pi(r), R)$. Moreover, if $\psi$ is an allocation for $R$, then let $\psi_{\pi}$ be the allocation such that, for each $r \in \mathcal{R}, \psi_{\pi}\left(r, R_{\psi}\right)=\psi(\pi(r), R)$. Thus, $R_{\pi}$ and $\psi_{\pi}$ are derived by applying the permutation $\pi$, whose role is just to rename the researchers in $\mathcal{R}$. With these notions in place, we can now state the following property, which is in fact not enjoyed by trivial.

Definition 6.25 (Impartiality). Let $\pi$ be an arbitrary permutation over $\mathcal{R}$. A
division rule $\gamma$ must be such that, for each $r \in \mathcal{R}$ and each allocation $\psi, \gamma_{\psi_{\pi}}\left(r, R_{\pi}\right)=$ $\gamma_{\psi}(\pi(r), R)$ holds.

Yet again this is not enough. To see that, consider the very impartial rule uniform, where the overall ANVUR score is distributed uniformly over the various researchers, i.e.,

$$
\operatorname{uniform}_{\psi}(r)=\frac{\text { score }_{\mathrm{VQR}}(R)}{|\mathcal{R}|} .
$$

Clearly, uniform is unsatisfying because it does not capture our intuition that a division rule should reflect the actual contribution of each researcher to the overall evaluation of the structure. In the rest of this section, we will elaborate on this issue by using the notion of marginal contribution, which fits well our intuition of actual contributions of individual or groups to the performance of a given structure.

### 6.5.3 Marginal contribution

Let $R$ be a research structure and assume that $\psi$ is the allocation selected by the structure $R$, so that the set of products $\mathcal{P}_{\psi}$ have been submitted and evaluated by ANVUR. Let $\mathcal{S} \subseteq \mathcal{R}$ be any set of researchers. An allocation $\psi_{\mathcal{S}}$ for $\mathcal{S}$ is $\psi-$ legal if $\psi_{\mathcal{S}}(r) \subseteq \mathcal{P}_{\psi}$, for each $r \in \mathcal{S}$. That is, any legal allocation only considers for the assignment the products already evaluated by ANVUR. The allocation $\psi_{\mathcal{S}}$ is $\psi$-optimal if there is no $\psi$-legal $\psi_{\mathcal{S}}^{\prime}$ such that $\sum_{r \in \mathcal{S}} \sum_{p \in \psi_{\mathcal{S}}^{\prime}(r)} \operatorname{score} e_{\mathrm{VQR}}(p)>$ $\sum_{r \in \mathcal{S}} \sum_{p \in \psi_{\mathcal{S}}(r)}$ score $_{\mathrm{VQR}}(p)$.

Given the above notions, we can equip $S$ with the following score, which is meant to assess the overall VQR score that researchers in $S$ would have achieved if the research structure had been constituted by them only (i.e., without caring about their co-authors outside $S$ )

$$
\operatorname{best}_{\psi}(\mathcal{S})=\sum_{r \in \mathcal{S}} \sum_{p \in \psi_{\mathcal{S}}(r)} \text { score }_{\mathrm{VQR}}(p),
$$

where $\psi_{\mathcal{S}}$ is any $\psi$-optimal allocation.
Note that in the extreme case where $\mathcal{S}=\mathcal{R}$ we just obtain $\operatorname{best}_{\psi}(\mathcal{R})=\operatorname{score}_{\mathrm{VQR}}(p) .$. That is, when all researchers of $R$ are considered, best $_{\psi}$ precisely coincides with the overall VQR score.

We can now formalize the notion of marginal contribution.

Definition 6.26. Let $\psi$ be an allocation for $R$. Given two sets of researchers $\mathcal{S}_{1}, \mathcal{S}_{2} \subseteq$ $\mathcal{R}$ with $\mathcal{S}_{1} \subseteq \mathcal{S}_{2}$, the marginal contribution of $\mathcal{S}_{1}$ to $\mathcal{S}_{2}$ (in $R$ and $\psi$ ) is the value:

$$
\operatorname{marg}_{\psi}\left(\mathcal{S}_{1}, \mathcal{S}_{2}\right)=\operatorname{best}_{\psi}\left(\mathcal{S}_{2}\right)-\operatorname{best}_{\psi}\left(\mathcal{S}_{2} \backslash \mathcal{S}_{1}\right)
$$

Intuitively, the marginal contribution $\operatorname{marg}_{\psi}\left(S_{1}, S_{2}\right)$ quantifies the loss of VQR score for the group of researchers $\mathcal{S}_{2}$ (e.g., a substructure) if the groups of researchers $\mathcal{S}_{1}$ were not part of it. In particular, $\operatorname{marg}_{\psi}(\{r\}, \mathcal{R})$ measures the loss for the whole structure $R$, if the single researcher $r$ were not part of it.

We are now ready to define the last fairness property, which takes care of the actual contribution of individuals and groups.

Definition 6.27 (marginality). A division rule $\gamma$ must be such that, for each group of researchers $\mathcal{S} \subseteq \mathcal{R}$ and each allocation $\psi, \gamma_{\psi}(\mathcal{S}) \geq \operatorname{marg}_{\psi}(\mathcal{S}, \mathcal{R})$. Therefore, every group is granted at least its marginal contribution to the performance of the structure $R$.

In particular, the above property entails that groups without interactions with other researchers, e.g., departments without collaborations with other departments of the same university, get precisely the total scores of the products assigned to them according to $\psi$. We now describe a practical example with the aim to summarize and clarify the division rules and the issues discussed so far.


Figure 6.2: Running example

### 6.5.4 A simple scenario

Consider the simple scenario depicted in Figure 6.2 and let us exploit an intuitive graphical notation based on a weighted bipartite graph, whose vertices are the researchers and the products, and whose weights are the VQR scores of the products. Assume that there are just three researchers, $r_{1}, r_{2}$ and $r_{3}$, affiliated to $R$. Moreover, consider the allocation $\psi$ such that $\psi\left(r_{1}\right)=\left\{p_{1}, p_{2}, p_{3}\right\}, \psi\left(r_{2}\right)=$ $\left\{p_{4}, p_{5}, p_{6}\right\}$ and $\psi\left(r_{3}\right)=\left\{p_{7}, p_{8}, p_{9}\right\}$. Thus, we have that $\mathcal{P}_{\psi}=\left\{p_{1}, \ldots, p_{9}\right\}$ is the set of publications submitted for the evaluation. In particular, we also assume that products $\left(r_{1}\right) \cap \mathcal{P}_{\psi}=\left\{p_{1}, p_{2}, p_{3}, p_{4}\right\}, \operatorname{products}\left(r_{2}\right) \cap \mathcal{P}_{\psi}=\left\{p_{3}, p_{4}, p_{5}, p_{6}\right\}$, and $\operatorname{products}\left(r_{3}\right) \cap \mathcal{P}_{\psi}=\left\{p_{3}, p_{6}, p_{7}, p_{8}\right\}$. Thus, $p_{3}$ is co-authored by $r_{1}, r_{2}$ e $r_{3}$, while $p_{4}$ is co-authored by $r_{1}$ and $r_{2}$. Then, according to the VQR score assigned to each product (i.e., the weight of the corresponding vertex in Figure 6.2), we get that $\operatorname{products}_{\mathrm{VQR}}(\psi)=66$.

Suppose now that we want to use the rule proj to assign the scores. Recall that proj is defined as follows.

$$
\operatorname{proj}_{\psi}(r)=\sum_{p \in \psi(r)} \operatorname{score}_{\mathrm{VQR}}(p)
$$

Then, according to $\operatorname{proj}$ the researchers would get $\operatorname{proj}_{\psi}\left(r_{1}\right)=20, \operatorname{proj}_{\psi}\left(r_{2}\right)=24$
and $\operatorname{proj}_{\psi}\left(r_{3}\right)=22$. This is unfair for the researcher $r_{1}$ since even if she and $r_{2}$ have co-authored $p_{3}$ and $p_{4}$, they get a different payoff. Therefore, as we said, $r_{1}$ could complain with her structure if $\psi$ is selected.

Now, suppose were using the rule trivial,

$$
\operatorname{trivial}_{\psi}(r)= \begin{cases}\sum_{p \in \mathcal{P}_{\psi}} \operatorname{score}_{\mathrm{VQR}}(p), & \text { if } r=r_{1} \\ 0 & \text { if } r \in \mathcal{R} \backslash\left\{r_{1}\right\}\end{cases}
$$

it is easy to check that $\operatorname{score}\left(r_{1}\right)=66$ and $\operatorname{score}\left(r_{i}\right)=0 \forall r_{i} \in \mathcal{R} \backslash\left\{r_{1}\right\}$. Clearly trivial is unsatisfying.

Finally, if we consider

$$
\operatorname{uniform}_{\psi}(r)=\frac{\operatorname{score}_{\mathrm{VQR}}(R)}{|\mathcal{R}|}
$$

we have that each researcher gets the score $\frac{20+24+22}{3}=22$, that is uniform, but it does not capture our intuition that a division rule should reflect the actual contribution of each researcher to the overall evaluation of the structure.

We can show that the concept of marginal contribution captures this intuition. We proceed as follows. Initially, we compute the best scores that all the possible groups of researches could get if they were the only one in the research structure.

- $\operatorname{best}_{\psi}\left(\left\{r_{1}\right\}\right)=24, \operatorname{best}_{\psi}\left(\left\{r_{2}\right\}\right)=24$ and $\operatorname{best}_{\psi}\left(\left\{r_{3}\right\}\right)=22$;
- $\operatorname{best}_{\psi}\left(\left\{r_{1}, r_{2}\right\}\right)=44$, $\operatorname{best}_{\psi}\left(\left\{r_{1}, r_{3}\right\}\right)=46$ and $\operatorname{best}_{\psi}\left(\left\{r_{2}, r_{3}\right\}\right)=46$;
- $\operatorname{best}_{\psi}\left(\left\{r_{1}, r_{2}, r_{3}\right\}\right)=66$.

Then, by recalling the formula of marginal contributions,

$$
\operatorname{marg}_{\psi}\left(\mathcal{S}_{1}, \mathcal{S}_{2}\right)=\operatorname{best}_{\psi}\left(\mathcal{S}_{2}\right)-\operatorname{best}_{\psi}\left(\mathcal{S}_{2} \backslash \mathcal{S}_{1}\right)
$$

we compute them, first with respect to the whole set $\mathcal{R}$ of researchers in $R$, and then with respect to all the possible pairs of researchers $\left\{r_{i}, r_{j}\right\}$ :

- $\operatorname{marg}_{\psi}\left(\left\{r_{1}\right\}, \mathcal{R}\right)=66-46=20$;
- $\operatorname{marg}_{\psi}\left(\left\{r_{2}\right\}, \mathcal{R}\right)=66-46=20$;
- $\operatorname{marg}_{\psi}\left(\left\{r_{3}\right\}, \mathcal{R}\right)=66-44=22$;
- $\operatorname{marg}_{\psi}\left(\left\{r_{1}, r_{2}\right\}, \mathcal{R}\right)=66-22=44$;
- $\operatorname{marg}_{\psi}\left(\left\{r_{1}, r_{3}\right\}, \mathcal{R}\right)=\operatorname{marg}_{\psi}\left(\left\{r_{2}, r_{3}\right\}, \mathcal{R}\right)=66-24=42$;
- $\operatorname{marg}_{\psi}\left(\left\{r_{1}\right\},\left\{r_{1}, r_{2}\right\}\right)=\operatorname{marg}_{\psi}\left(\left\{r_{2}\right\},\left\{r_{1}, r_{2}\right\}\right)=44-24=20$;
- $\operatorname{marg}_{\psi}\left(\left\{r_{1}\right\},\left\{r_{1}, r_{3}\right\}\right)=\operatorname{marg}_{\psi}\left(\left\{r_{2}\right\},\left\{r_{2}, r_{3}\right\}\right)=46-22=24$;
- $\operatorname{marg}_{\psi}\left(\left\{r_{3}\right\},\left\{r_{1}, r_{3}\right\}\right)=\operatorname{marg}_{\psi}\left(\left\{r_{3}\right\},\left\{r_{2}, r_{3}\right\}\right)=46-24=22$.

From the computed marginal contributions, we can see that researchers $r_{1}$ and $r_{2}$ are completely interchangeable.

### 6.5.5 Using the Shapley value as a division rule

In this section we show that a division rule enjoying all the above fairness properties actually exists and it is the Shapley value. Indeed, it is natural to formalize the considered application as a coalitional game where the agents are the researchers, and the worth of any coalition is the best result that the researchers in that coalition may achieve if acting in isolation. Formally, for a research structure $R$ and an allocation $\psi$, define the coalitional game $\mathcal{G}_{\psi}=\left(\mathcal{R}\right.$, best $\left._{\psi}\right)$. Then, we propose the following division rule, which is easily seen to be precisely the Shapley value of $\mathcal{G}_{\psi}$ :

$$
\gamma_{\psi}^{*}(r)=\sum_{\mathcal{S} \subseteq \mathcal{R}} \frac{(|\mathcal{R}|-|\mathcal{S}|)!(|\mathcal{S}|-1)!}{|\mathcal{R}|!} \operatorname{mar}_{\psi}(r, \mathcal{S}), \quad \forall r \in \mathcal{R}
$$

This rule, satisfies all the nice fairness properties discussed so far as the following theorem states:

Theorem 6.28 (cf. [Greco and Scarcello, 2013b]). Let $\psi$ be any allocation for a given structure $R$. Then, $\gamma_{\psi}^{*}$ is a division rule satisfying the budget-balance, the independence of the product allocation, the impartiality and the marginality fairness properties.

### 6.5.6 Discussion

So far, we focused on a setting where an allocation is selected by the research structure $R$ and we are in charge of specifying a division rule $\gamma_{\psi}^{*}$ that enjoys all the fairness properties discussed earlier. However, fairness depends also on the preliminary selection of the products $\mathcal{P}_{\psi}$. Indeed we cannot guarantee fairness if the given allocation is unfair itself. In the first phase of the evaluation process, the structure $R$ (i.e., the department) is in charge of evaluating the scientific production of its researchers. However, a detailed evaluation of the whole production is not an easy task, indeed the simplest solution for $R$ is to ask its researchers to perform a selfevaluation of their products. Therefore, the goal of $R$ will be achieved if authors correctly/truthfully self-evaluate their products. But what happens if this assumption is removed? It is natural at all to remove this assumption? While it is clear that the main goal for $R$ is to maximize the total value of the products submitted to ANVUR, whenever the same product has different co-authors some strategic issues come into play, and co-authors personal interest may induce them to cheat on the quality of their products. To see why, recall that ANVUR did not specify a division rule. Therefore, researchers consider proj as the division rule. Thus if proj is selected, and if (as we think) it is then used to reward $R$ and hence researchers, co-authors may be competitors and might want to act strategically to improve their own score. It is clear that these kinds of manipulations, should be prevented by a good division
rule. Indeed this issues is addressed by a branch of game theory called Mechanism design. Note that, proj is definetely not fair. Thus, under the assumption that this rule will be used, researchers act strategically in order to have allocated the best products.


Figure 6.3: A clooser look at strategical behaviours.

This is exemplified in Figure 6.3. For example, researcher $r_{1}$ could manipulate the allocation process by being not truthful in order to get the best publications for herself and thus forcing $r_{2}$ to get the bad ones. The result is that the global optimum for $R$ is missed (i.e., no efficiency) and moreover there is no fairness at all. Note that this strategic behavior is likely to happen as soon as two researchers from different departments co-authored some work, and each of them is interested in providing as much as contribution as possible to her/his department. For instance, in Figure 6.4 reports data about co-authorship at University of Calabria while in Figure 6.5 we illustrate the number of research groups with their size.

From these pictures we see that a very particular interactions between researchers emerges. Indeed 30 components out of 14 consist of only two researchers. Moreover, for the largest component which comprise 27 researchers we discovered that treewidth is really low Figure 6.6. We will see in the next chapter that these kinds of information


Figure 6.4:
Co-autorships at University of Calabria. Researchers against the number of co-authored products

## Components at Univ. of Calabria



Figure 6.5: The number of components at University of Calabria
are of great interest (specifically, to the complexity of computing the Shapley value in this setting). We conclude the section by referring the reader to [Greco and Scarcello, 2013b] for more on this subject.


Figure 6.6: An example component

## CHAPTER 7

## Structural Tractability of Shapley and Banzhaf Values in Allocation Games

### 7.1 Introduction

It is often required to reward agents according to their actual contribution to the total worth of their organization. A notable case is when a set of indivisible goods (or tasks, or whatever) should be allocated to a given set of agents, and a monetary compensation is possible. Coalitional games, and in particular allocation games, properly model such applications, and the Shapley value and the Banzhaf value are two well-known solution concepts that provide fair and hence acceptable agent worths. Unfortunately, it turns out that their computation is a \#P-hard problem, for which only quite trivial tractable classes are known.

In this chapter we identify large islands of tractability for both solution concepts, by exploiting the structure of the interactions among agents. The main result is that for any class of games having interaction graphs of bounded treewidth, both the Shapley value and the Banzhaf value are computable in polynomial time. Tractability also holds for games where each good is owned by at most two agents (independently of agents' interactions). These results allow us to deal with games having hundreds of agents in real-world applications. Moreover, we believe that the underlying technical ingredients can be exploited for different kinds of coalitional games, too.

For a better understanding of the problem, we first strengthen the known hardness results to the case of goods with one possible value only. Then, we look for islands of tractability of allocations problems. To this end, we provide a characterization of the marginal contribution of an agent to any coalition in terms of certain properties of good allocations, which are not required to be optimal ones. Such a technical tool allows us to point out the tractability of allocation games where every good is shared (or claimed for) by two agents at most.

Our main result, also based on the tool discussed in Chapter 5 and on further ingredients exploiting constraint satisfaction techniques, is a polynomial-time algorithm for the computation of the Shapley value and the Banzhaf value in allocation games where agent interactions have a tree-like structure - formally, have bounded treewidth [Robertson and Seymour, 1984]. These games capture scenarios of practical interest; for instance, in Chapter 6 we analyzed the data of a concrete instantiation for the setting described by [Greco and Scarcello, 2014b]. The setting refers to an allocation problem arising in the Italian Research Assessment program. The instantiation refers to the data of the University of Calabria by discovering that the treewidth of the associated graph, consisting of more than 500 nodes, is just 9 . Moreover, the main result and the technical tools used to get it have their own theoretical interest, and the analysis of the complexity of reasoning problems related to coalitional games over classes of instances having some useful structural property is an active topic of research in artificial intelligence. In fact, even if we focus on the Shapley and the Banzhaf values, the methodology we propose may have a wider spectrum of applicability, as its salient ideas can be reused in other contexts.

### 7.2 Intractability of computation

Computing the Shapley value is a problem that has been shown to be \#P-complete on different classes of games (see, e.g., [Deng and Papadimitriou, 1994; Nagamochi et al., 1997; Bachrach and Rosenschein, 2009; Aziz and de Keijzer, 2014]), including the allocation games [?]. In particular, hardness has been shown to hold even on instances whose goods have three possible values. Below, we improve the result by showing that there is no advantage in focusing on scenarios where all goods have the same value. To this end, we first prove the \#P-hardness of computing the Banzhaf value.

Theorem 7.1. Computing the Banzhaf value is \#P-hard on allocation games (under Turing reductions), even for scenarios $\mathcal{A}=\langle N, \mathbb{G}, \Omega$, val $\rangle$ such that $\mid\{\operatorname{val}(g) \mid g \in$ $\mathbb{G}\} \mid=1$.

Proof Sketch. Let $(S \cup I, E)$ be a bipartite graph, hence with $S \cap I=\emptyset$ and $E \subseteq S \times I$. Computing the number of subsets $C \subseteq S$ of vertices to which all vertices in $I$ can be matched is \#P-hard [Colbourn et al., 1995].

Based on $(S \cup I, E)$, let us build the allocation scenario $\mathcal{A}=\langle S \cup\{|S|+1\}, I, \Omega$, val $\rangle$ where nodes in $S$ (resp., $I$ ) are transparently viewed as the agents (resp., goods), where $\operatorname{val}(g)=1$ for each $g \in I$, and where $\Omega(i)=\{g \in I \mid\{i, g\} \in E\}$ while $\Omega(|S|+1)=I$. Consider then the allocation game $\mathcal{G}_{\mathcal{A}}=\left\langle N, v_{\mathcal{A}}\right\rangle$ with $N=S \cup$ $\{|S|+1\}$, and the Banzhaf value $\beta_{|S|+1}\left(\mathcal{G}_{\mathcal{A}}\right)$. Observe that, for any given coalition $C \subseteq S=N \backslash\{|S|+1\}, v(C \cup\{|S|+1\})-v(C)=0$ if, and only if, $C \subseteq S$ is a set of vertices to which all vertices in $I$ can be matched. Eventually, $\beta_{|S|+1}\left(\mathcal{G}_{\mathcal{A}}\right) \times 2^{|S|}$ is the number of subsets $C \subseteq S$ for which some vertex in $I$ cannot be matched, and $2^{|S|}-\beta_{|S|+1}\left(\mathcal{G}_{\mathcal{A}}\right) \times 2^{|S|}$ is the desired number, which can be computed in polynomial


Figure 7.1: Illustration in the proof of Theorem 7.1: Construction of the scenario $\mathcal{A}_{\alpha}$ based on $\mathcal{A}$. time once the Banzhaf value $\beta_{|S|+1}\left(\mathcal{G}_{\mathcal{A}}\right)$ is known.

Now, we show that the Banzhaf value of allocation games can be computed in polynomial time based on the knowledge of the Shapley value, so that this latter concept turns out to be \#P-hard too. This property was already known to hold over (certain) simple games [Aziz et al., 2009]. For its proof, we shall exploit the fact that, for each agent $i \in N$, the Shapley value can be rewritten as follows:

$$
\begin{equation*}
\phi_{i}\left(\mathcal{G}_{\mathcal{A}}\right)=\sum_{h=0}^{n-1} \frac{h!(n-h-1)!}{n!} \beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right), \tag{7.1}
\end{equation*}
$$

where, for each $h \in\{0, \ldots, n-1\}$,

$$
\beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)=\sum_{C \subseteq N \backslash\{i\},|C|=h}(v(C \cup\{i\})-v(C)) .
$$

Theorem 7.2. Computing the Shapley value is \#P-hard on allocation games (under Turing reductions), even for scenarios $\mathcal{A}=\langle N, \mathbb{G}, \Omega$, val $\rangle$ such that $\mid\{\operatorname{val}(g) \mid g \in$ $\mathbb{G}\} \mid=1$.

Proof Sketch. Assume that $|\{\operatorname{val}(g) \mid g \in \mathbb{G}\}|=1$; in particular, w.l.o.g., assume that $\operatorname{val}(g)=1$ holds, for each $g \in \mathbb{G}$. Let $i$ be an agent in $N$, and for each $\alpha \in\{1, \ldots, n\}$, consider the allocation scenario $\mathcal{A}_{\alpha}=\langle\{1, \ldots, n+\alpha\}, \mathbb{G} \cup$ $\left.\left\{g_{1}, \ldots, g_{\alpha}\right\}, \Omega_{\alpha}, \operatorname{val}_{\alpha}\right\rangle$ defined as follows:

- on the agents in $N \backslash\{i\}, \Omega_{\alpha}$ coincides with $\Omega$; moreover, $\Omega_{\alpha}(n+j)=\left\{g_{j}, g_{j+1}\right\}$, for each $j \in\{1, \ldots, \alpha-1\} ; \Omega_{\alpha}(i)=\Omega(i) \cup\left\{g_{1}\right\} ;$ and $\Omega_{\alpha}(n+\alpha)=\left\{g_{\alpha}\right\} ;$
- on the goods in $\mathbb{G}, \operatorname{val}_{\alpha}$ coincides with val; moreover, $\operatorname{val}_{\alpha}\left(g_{j}\right)=1$ for each $j \in\{1, \ldots, \alpha\}$.

Let $C \subseteq N \cup\{n+1, \ldots,(n+\alpha-1)\}$, and consider the illustration in Figure 7.1 to help the intuition. Observe that if $\{i, n+1, \ldots,(n+\alpha-1)\} \nsubseteq C$, then $v_{\mathcal{A}_{\alpha}}(C \cup$ $\{n+\alpha\})-v_{\mathcal{A}_{\alpha}}(C)=1$. Otherwise, i.e., if $\{i, n+1, \ldots,(n+\alpha-1)\} \subseteq C$, then $v_{\mathcal{A}_{\alpha}}(C \cup\{n+\alpha\})-v_{\mathcal{A}_{\alpha}}(C)=v_{\mathcal{A}}\left(C^{\prime} \cup\{i\}\right)-v_{\mathcal{A}}\left(C^{\prime}\right)$, where $C^{\prime}=C \backslash\{i, n+1, \ldots, n+\alpha\}$. Therefore, for each $h \in\{0, \ldots,(n+\alpha-1)\}$, we have that:

$$
\beta_{n+\alpha}\left(\mathcal{G}_{\mathcal{A}_{\alpha}}, h\right)= \begin{cases}\kappa_{n+\alpha}(h)+\beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h-\alpha\right) & \text { if } h \geq \alpha \\ \kappa_{n+\alpha}(h) & \text { if } h<\alpha\end{cases}
$$

where $\kappa_{n+\alpha}(h)$ is the number of coalitions $C$ such that $|C|=h$ and such that $\{i, n+1, \ldots,(n+\alpha-1)\} \nsubseteq C$.

Note that $\kappa_{n+\alpha}(h)$ can be computed in polynomial time.
By using the above expression for $\beta_{n+\alpha}\left(\mathcal{G}_{\mathcal{A}_{\alpha}}, h\right)$ in Equation (7.1), we can derive that:

$$
\phi_{n+\alpha}\left(\mathcal{G}_{\mathcal{A}_{\alpha}}\right)=\sum_{h=0}^{n+\alpha-1} \frac{h!(n+\alpha-h-1)!}{(n+\alpha)!} \kappa_{n+\alpha}(h)+\sum_{h=0}^{n-1} \frac{(h+\alpha)!(n-h-1)!}{(n+\alpha)!} \beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right) .
$$

Eventually, by instantiating $\alpha$ with each value in $\{1, \ldots, n\}$, we get a system of $n$ linear equations over the variables $\beta_{i}\left(\mathcal{G}_{\mathcal{A}}, 0\right), \ldots, \beta_{i}\left(\mathcal{G}_{\mathcal{A}}, n-1\right)$. By multiplying by $(n+\alpha)$ ! all the terms of the corresponding equation, we get a system where $(h+$ $\alpha)!(n-h-1)$ ! is the coefficient of any term of the form $\beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$. Given the form of these coefficients, the equations are easily seen to be linear independent-see [Aziz and de Keijzer, 2014; Aziz et al., 2009].

Therefore, the values of $\beta_{i}\left(\mathcal{G}_{\mathcal{A}}, 0\right), \ldots, \beta_{i}\left(\mathcal{G}_{\mathcal{A}}, n-1\right)$ for which the system admits a solution are univocally determined. Moreover, they might be computed in polynomial time (e.g., by Bareiss's implementation of Gaussian elimination [Bareiss, 1968]), if we were able to compute in polynomial time $\phi_{n+\alpha}\left(\mathcal{G}_{\mathcal{A}_{\alpha}}\right)$, for each $\alpha \in\{1, \ldots, n\}$. To conclude, just notice that the Banzhaf value of agent $i$ in $\mathcal{G}_{\mathcal{A}}$ is precisely given by $\beta_{i}\left(\mathcal{G}_{\mathcal{A}}\right)=\frac{1}{2^{n-1}} \sum_{h=0}^{n-1} \beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$.

From these results, it turns out that acting on the values of goods does not help very much in identifying tractable classes of instances. So, we next consider different kinds of restrictions based on the "interactions" that emerge among agents.

### 7.3 Characterizations of the Shapley value

Throughout the section, assume that an allocation scenario $\mathcal{A}=\langle N, \mathbb{G}, \Omega$, val $\rangle$ is given. Let $\left\{w_{1}, \ldots, w_{m}\right\}=\{\operatorname{val}(g) \mid g \in \mathbb{G}\} \cup\{0\}$ be the set of all values associated with goods in $\mathbb{G}$ (plus the null value 0 , if not present), and assume that $w_{1}>w_{2}>$ $\cdots>w_{m}$. W.l.o.g, assume also that $w_{m}=0$.

### 7.3.1 A closer look at marginal contributions

In this section, we elaborate and discuss useful characterizations of the Shapley value of allocation games.

We start with a simple reformulation. Let $i \in N$ be an agent, let $h \in\{0, \ldots, n-1\}$, let $\ell \in\{1, \ldots, m\}$, and let us denote by $\# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{A}, h\right)$ the number of coalitions $C$ such that $|C|=h$ and $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C) \geq w_{\ell}$. Then, the coefficients $\beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$ in the expression for the Shapely value illustrated in Equation (7.1) can be rewritten as follows.

Theorem 7.3. For each agent $i \in N$ and $h \in\{0, \ldots, n-1\}$,

$$
\beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)=w_{1} \times \# \operatorname{col}_{1}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)+\sum_{\ell=2}^{m} w_{\ell} \times\left(\# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)-\# \operatorname{col}_{\ell-1}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)\right)
$$

Proof. Recall that, for each $i \in N$ and $h \in\{0, \ldots, n-1\}$,

$$
\beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)=\sum_{C \subseteq N \backslash\{i\},|C|=h}(v(C \cup\{i\})-v(C)) .
$$

Consider an arbitrary coalition $C \subseteq N \backslash\{i\}$ with $|C|=h$. Let $\bar{\pi}$ be an optimal allocation for $\mathcal{A}[C \cup\{i\}]$, and observe that because of the monotonicity of allocation games, there is an optimal allocation $\bar{\pi}_{-i}$ for $\mathcal{A}[C]$ such that $\operatorname{img}\left(\bar{\pi}_{-i}\right) \subseteq \operatorname{img}(\bar{\pi})$. If the two images coincide, then we clearly have that $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C)=0$. Otherwise, $\operatorname{img}(\bar{\pi}) \backslash \operatorname{img}\left(\bar{\pi}_{-i}\right)$ is a singleton $\{g\}$, and $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C)=\operatorname{val}(\bar{\pi})-\operatorname{val}\left(\bar{\pi}_{-i}\right)=$ $\operatorname{val}(g)$. In both cases, the marginal contribution is a value $w_{\ell}$ with $\ell \in\{1, \ldots, m\}$. Then, we can rewrite the expression for the coefficients as $\beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)=\sum_{\ell=1}^{m} w_{\ell} \times$ $M_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$, where $M_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$ is the number of coalitions $C$ with $|C|=h$ and such that $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C)=w_{\ell}$. By definition of $\# \operatorname{col}_{l}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$, we get $M_{1}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)=$ $\# \operatorname{col}_{1}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$ and, $\forall \ell \geq 2, M_{l}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)=\# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)-\# \operatorname{col}_{\ell-1}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$.

Hence, counting the number of coalitions to which an agent $i$ provides some marginal contribution is deeply related to the computation of the Shapley (and the Banzhaf) value of allocation games. We now further explore this specific task.

For each "level" $\ell \in\{1, \ldots, m\}$ over the possible values, an agent $i \in N$ is said to be dependent at level $\ell$ (short: $\ell$-dependent) if for each $g \in \Omega(i)$ with $\operatorname{val}(g) \geq w_{\ell}$, there is an agent $j \in N \backslash\{i\}$ such that $g \in \Omega(j)$. In particular, note that any agent having goods with value at least $w_{\ell}$ and which are not shared with any other agent is not dependent at level $\ell$; in fact, all her marginal contributions are at least $w_{\ell}$, independently of the coalition to be considered. Let $G_{\ell}=\left(N_{\ell}, E_{\ell}\right)$ be the undirected
graph where $N_{\ell}$ is the set of all $\ell$-dependent agents and where $\{i, j\} \in E_{\ell}$ if, and only if, there is a good $g \in \Omega(i) \cap \Omega(j)$ with $\operatorname{val}(g) \geq w_{\ell}$. Then, a coalition $R \subseteq N_{\ell}$ of agents is called a component at level $\ell$ (short: $\ell$-component) if the subgraph of $G_{\ell}$ induced by the nodes in $R$ is connected. As a special case, if there is no good $g \in \Omega(i)$ with $\operatorname{val}(g) \geq w_{\ell}$, then $\{i\}$ is an $\ell$-component.

Example 7.4. Consider the scenario $\mathcal{A}_{0}$ reported in Figure 6.1. We have $w_{1}=$ $\operatorname{val}\left(g_{1}\right)$. Moreover, $\{1,2\}$ and $\{3\}$ are the only subset-maximal components at level 1. Indeed, $g_{1} \in \Omega(1) \cap \Omega(2)$, and there is no good in $\Omega(3)$ with value $w_{1}$. $\triangleleft$

The concept of $\ell$-component plays a crucial role to characterize useful properties of marginal contributions.

Lemma 7.5. Let $R \cup\{i\}$ be an $\ell$-component, with $R \subseteq N_{\ell} \backslash\{i\}$. Then, the following statements are equivalent:
(1) $v_{\mathcal{A}}(R \cup\{i\})-v_{\mathcal{A}}(R) \geq w_{\ell}$;
(2) there is an optimal allocation $\bar{\pi}$ for $\mathcal{A}[R \cup\{i\}]$ such that $\operatorname{val}(\bar{\pi}) \geq w_{\ell}$, for each $j \in R \cup\{i\}$.

Proof Sketch. (1) $\Rightarrow$ (2) Assume that (2) does not hold. So, we are guaranteed about the existence of an optimal allocation $\bar{\pi}$ for $\mathcal{A}[R \cup\{i\}]$ and of an agent $j^{\prime} \in R \cup\{i\}$ such that $\operatorname{val}\left(\bar{\pi}\left(j^{\prime}\right)\right)<w_{\ell}$. Consider the following two possible cases.

First, assume that $\operatorname{val}(\bar{\pi}(i))<w_{\ell}$. Since the restriction of $\bar{\pi}$ over the agents in $R$ is a feasible allocation for $\mathcal{A}[R]$, then we immediately get that $v_{\mathcal{A}}(R) \geq \operatorname{val}(\bar{\pi})-$ $\operatorname{val}(\bar{\pi}(i))>\operatorname{val}(\bar{\pi})-w_{\ell}$, and hence $v_{\mathcal{A}}(R \cup\{i\})-v_{\mathcal{A}}(R)<w_{\ell}$.

Second, assume that $\operatorname{val}(\bar{\pi}(i)) \geq w_{\ell}$. We start by observing that, due to the optimality of $\bar{\pi}$, for each agent $j^{\prime} \in R \cup\{i\}$ with $\operatorname{val}\left(\bar{\pi}\left(j^{\prime}\right)\right)<w_{\ell},\left\{g \mid g \in \Omega\left(j^{\prime}\right) \wedge\right.$ $\left.\operatorname{val}(g) \geq w_{\ell}\right\} \subseteq \operatorname{img}(\bar{\pi})$. That is, goods that might be in principle allocated to an
agent $j^{\prime} \in R \cup\{i\}$ with $\operatorname{val}\left(\bar{\pi}\left(j^{\prime}\right)\right)<w_{\ell}$ and having value at least $w_{\ell}$ are actually allocated to some different agent in $R \cup\{i\}$. Given that $R \cup\{i\}$ is an $\ell$-component (and that, in particular, each agent is $\ell$-dependent), we are guaranteed about the existence of a succession $i=j_{1}^{\prime}, j_{2}^{\prime}, \ldots, j_{h}^{\prime}$ such that $\bar{\pi}\left(j_{x}^{\prime}\right) \cap \Omega\left(j_{x+1}^{\prime}\right) \neq \emptyset$, for each $x \in\{1, \ldots, h-1\}$; and $\operatorname{val}\left(\pi\left(j_{h}^{\prime}\right)\right)<w_{\ell}$. Consider then the function $\bar{\pi}_{-i}: R \rightarrow \mathbb{G} \cup\{\emptyset\}$ with $\bar{\pi}_{-i}\left(j_{x+1}^{\prime}\right)=\bar{\pi}\left(j_{x}^{\prime}\right)$, for each $x \in\{1, \ldots, h-1\}$; and $\bar{\pi}_{-i}\left(j^{\prime \prime}\right)=\bar{\pi}\left(j^{\prime \prime}\right)$, for each $j^{\prime \prime} \in R \backslash\left\{j_{2}^{\prime}, \ldots, j_{h}^{\prime}\right\}$. Then, $\bar{\pi}_{-i}$ is an allocation for $\mathcal{A}[R]$ and we have that $\operatorname{val}\left(\bar{\pi}_{-i}\right)=$ $\operatorname{val}(\bar{\pi})-\operatorname{val}\left(\bar{\pi}\left(j_{h}^{\prime}\right)\right)$. Hence, $v_{\mathcal{A}}(R) \geq \operatorname{val}\left(\bar{\pi}_{-i}\right)=\operatorname{val}(\bar{\pi})-\operatorname{val}\left(\bar{\pi}\left(j_{h}^{\prime}\right)\right)>\operatorname{val}(\bar{\pi})-w_{\ell}$. That is, $v_{\mathcal{A}}(R \cup\{i\})-v_{\mathcal{A}}(R)<w_{\ell}$.

In both cases, we have concluded that (1) does not hold.
$(2) \Rightarrow(1)$ Let $\bar{\pi}$ be an optimal allocation for $\mathcal{A}[R \cup\{i\}]$ such that $\operatorname{val}(\bar{\pi}) \geq w_{\ell}$, for each $j \in R \cup\{i\}$. Because of the allocation monotonicity property, there is an optimal allocation $\bar{\pi}_{-i}$ for $\mathcal{A}[R]$ such that $\operatorname{img}\left(\bar{\pi}_{-i}\right) \subseteq \operatorname{img}(\bar{\pi})$. Hence, for each $j \in R, \operatorname{val}\left(\bar{\pi}_{-i}(j)\right) \geq w_{\ell}$. Now, observe that $v_{\mathcal{A}}(R \cup\{i\})=\operatorname{val}(\operatorname{img}(\bar{\pi}))$ and $v_{\mathcal{A}}(R)=$ $\operatorname{val}\left(\operatorname{img}\left(\bar{\pi}_{-i}\right)\right)$. So, $v_{\mathcal{A}}(R \cup\{i\})-v_{\mathcal{A}}(R)$ coincides with the value of one of the goods in $\operatorname{img}(\bar{\pi})$, and $v_{\mathcal{A}}(R \cup\{i\})-v_{\mathcal{A}}(R) \geq w_{\ell}$.

Example 7.6. In our running example, $\{1,2\}$ is a component at level 1 , and $g_{1}$ is the only good having value (at least) $w_{1}$. So, there is no allocation $\bar{\pi}$ for $\mathcal{A}_{0}[\{1,2\}]$ satisfying condition (2) in Lemma 7.5 for $\ell=1$, and hence $v_{\mathcal{A}_{0}}(\{1,2\})-v_{\mathcal{A}_{0}}(\{2\})<$ $w_{1}$. Eventually, by moving to the second level, the allocation for $\mathcal{A}_{0}[\{1,2\}]$ depicted in Figure 6.1 witnesses, again by Lemma 7.5, that $v_{\mathcal{A}_{0}}(\{1,2\})-v_{\mathcal{A}_{0}}(\{2\}) \geq w_{2} . \triangleleft$

We now generalize Lemma 7.5 to arbitrary coalitions and allocations that are not required to be optimal. To this end, if $C \subseteq N$ is a coalition and $i \notin C$ is an agent, then we denote by $\mathbf{p}_{\ell}^{i}(C)$ the $\ell$-part of $C$ w.r.t. $i$. This is defined as the emptyset
if $i \notin N_{\ell}$; otherwise, $\mathbf{p}_{\ell}^{i}(C)$ is the subset-maximal (in fact, unique) $\ell$-component $R \subseteq C \cup\{i\}$ with $i \in R$.

Theorem 7.7. Let $C \subseteq N$ be a coalition and let $i \in N \backslash C$ be an agent. Then, the following statements are equivalent:
(1) $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C) \geq w_{\ell}$;
(2) there is an allocation $\bar{\pi}$ for $\mathcal{A}\left[\mathbf{p}_{\ell}^{i}(C)\right]$ such that $\operatorname{val}(\bar{\pi}(j)) \geq w_{\ell}$, for each $j \in$ $\mathbf{p}_{\ell}^{i}(C)$.

Proof. Let us address preliminary the case where $i \notin N_{\ell}$. In this case, $\mathbf{p}_{\ell}^{i}(C)=\emptyset$ and (2) trivially holds. Moreover, since agent $i$ is not $\ell$-dependent, its marginal contribution to any coalition is at least $w_{\ell}$. So, (1) holds trivially, too. In the remaining, consider the case where $i \in N_{\ell}$, so that $i \in \mathbf{p}_{\ell}^{i}(C)$. Let $R=\mathbf{p}_{\ell}^{i}(C) \backslash\{i\}$ and let $S=C \backslash R$.
$(1) \Rightarrow(2)$ By submodularity, we know that $v_{\mathcal{A}}(R \cup\{i\})-v_{\mathcal{A}}(R) \geq v_{\mathcal{A}}(R \cup\{i\} \cup S)-$ $v_{\mathcal{A}}(R \cup S)$ holds. Hence, we get that $v_{\mathcal{A}}(R \cup\{i\})-v_{\mathcal{A}}(R) \geq w_{\ell}$. Then, by Lemma 7.5, we derive that there is an allocation $\bar{\pi}$ for $\mathcal{A}[R \cup\{i\}]$ such that $\operatorname{val}(\bar{\pi}(j)) \geq w_{\ell}$, for each $j \in R \cup\{i\}$.
$(2) \Rightarrow(1)$ Assume that $\pi$ is an allocation for $\mathcal{A}[R \cup\{i\}]$ such that $\operatorname{val}(\pi(j)) \geq w_{\ell}$, for each $j \in R \cup\{i\}$. We first claim that this property must hold for every optimal allocation, too. Assume by contradiction that there is an optimal allocation $\bar{\pi}$ for $\mathcal{A}[R \cup\{i\}]$ with $\operatorname{val}\left(\bar{\pi}\left(j^{\prime}\right)\right)<w_{\ell}$ for some agent $j^{\prime} \in R \cup\{i\}$. Because $\operatorname{img}(\bar{\pi}) \neq$ $\operatorname{img}(\pi)$, there exists a succession of agents $j^{\prime}=j_{1}^{\prime}, j_{2}^{\prime}, \ldots, j_{h}^{\prime}$ with the assignment to the last agent (possibly, just $j^{\prime}$ ) such that $\pi\left(j_{h}^{\prime}\right) \notin \operatorname{img}(\bar{\pi})$ and $\pi\left(j_{x}^{\prime}\right)=\bar{\pi}\left(j_{x+1}^{\prime}\right)$, for each $x \in\{1, \ldots, h-1\}$. Then, consider the function $\bar{\pi}^{*}: R \cup\{i\} \rightarrow \mathbb{G} \cup\{\emptyset\}$ such that $\bar{\pi}^{*}\left(j_{x}^{\prime}\right)=$ $\pi\left(j_{x}^{\prime}\right)$, for each $x \in\{1, \ldots, h\} ;$ and $\bar{\pi}^{*}\left(j^{\prime \prime}\right)=\bar{\pi}\left(j^{\prime \prime}\right)$, for each $j^{\prime \prime} \in(R \cup\{i\}) \backslash\left\{j_{1}^{\prime}, \ldots, j_{h}^{\prime}\right\}$.

Note that $\bar{\pi}^{*}$ is an allocation for $\mathcal{A}[R \cup\{i\}]$ and $\operatorname{val}\left(\bar{\pi}^{*}\right)>\operatorname{val}(\bar{\pi})$, because the only difference is that the new image contains the good $\pi\left(j_{h}^{\prime}\right)$, whose value is at least $w_{\ell}$, instead of the good $\bar{\pi}\left(j_{1}^{\prime}\right)$ for which $\operatorname{val}\left(\bar{\pi}\left(j_{1}^{\prime}\right)\right)<w_{\ell}$ holds. However, this contradicts the optimality of $\bar{\pi}$.

So far, we have shown that there is an optimal allocation $\bar{\pi}$ for $\mathcal{A}[R \cup\{i\}]$ such that $\operatorname{val}(\bar{\pi}(j)) \geq w_{\ell}$, for each $j \in R \cup\{i\}$. By Lemma 7.5, this entails that $v_{\mathcal{A}}(R \cup\{i\})-$ $v_{\mathcal{A}}(R) \geq w_{\ell}$ holds. Moreover, by monotonicity, there is an optimal allocation $\bar{\pi}_{-i}$ for $\mathcal{A}[R]$ such that $\operatorname{img}\left(\bar{\pi}_{-i}\right) \subseteq \operatorname{img}(\bar{\pi})$. Hence, for each $j \in R, \operatorname{val}\left(\bar{\pi}_{-i}(j)\right) \geq w_{\ell}$. Given these properties of the optimal allocations for $\mathcal{A}[R \cup\{i\}]$ and $\mathcal{A}[R]$, and given the definition of $S$, we conclude that $\operatorname{opt}(\mathcal{A}[R \cup\{i\} \cup S])=\operatorname{opt}(\mathcal{A}[R \cup\{i\}])+\operatorname{opt}(\mathcal{A}[S])$ and $\operatorname{opt}(\mathcal{A}[R \cup S])=\operatorname{opt}(\mathcal{A}[R])+\operatorname{opt}(\mathcal{A}[S])$. So, $v_{\mathcal{A}}(R \cup\{i\} \cup S)-v_{\mathcal{A}}(R \cup S)=$ $v_{\mathcal{A}}(R \cup\{i\})-v_{\mathcal{A}}(R) \geq w_{\ell}$.

Example 7.8. By continuing with Example 7.6, note that $\{1,2\}=\mathbf{p}_{\ell}^{1}(\{2,3\})$ holds, for $\ell \in\{1,2\}$. Therefore, the allocation for $\mathcal{A}_{0}[\{1,2\}]$ depicted in Figure 6.1 witnesses, by Theorem 7.7, that $v_{\mathcal{A}_{0}}(\{1,2,3\})-v_{\mathcal{A}_{0}}(\{2,3\}) \geq w_{2} . \triangleleft$

### 7.4 Islands of tractability

### 7.4.1 Bounded sharing

Our analysis intensively uses Theorem 7.7. The first outcome is an island of tractability based on the notion of bounded sharing. Formally, for a given level $\ell$, define the sharing degree of an allocation scenario $\mathcal{A}$, denoted by $s d_{\ell}(\mathcal{A})$, as the maximum, over all goods $g$ with $\operatorname{val}(g) \geq w_{\ell}$, of $|\{j \in N \mid g \in \Omega(j)\}|$. Intuitively, it measures the maximum number of agents competing for the same good (with value at least $w_{\ell}$ ).

Lemma 7.9. Assume that $s d_{\ell}(\mathcal{A}) \leq 2$, for a value level $\ell$. Let $i$ be an agent in $N$,
and let $h \in\{0, \ldots, n-1\}$. Then,

$$
\# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)=\frac{(n-1)!}{(n-1-h)!h!}-\mathcal{X}, \text { where }
$$

- $\mathcal{X}=0$, if $h<\left|\mathbf{p}_{\ell}^{i}(N \backslash\{i\})\right|-1$; or if $i$ is not $\ell$-dependent, or the subgraph of $G_{\ell}$ induced by the nodes in $\mathbf{p}_{\ell}^{i}(N \backslash\{i\})$ contains a cycle, or there are two agents $j$ and $j^{\prime}$ in $\mathbf{p}_{\ell}^{i}(N \backslash\{i\})$ with $\left|\Omega(j) \cap \Omega\left(j^{\prime}\right) \cap\left\{g \mid \operatorname{val}(g) \geq w_{\ell}\right\}\right|>1$.
- $\mathcal{X}=\frac{\left(n-\left|\mathbf{p}_{\ell}^{i}(N \backslash\{i\})\right|\right)!}{(n-h-1)!\left(h+1-\left|\mathbf{p}_{\ell}^{2}(N \backslash\{i\})\right|\right)!}$, otherwise.

Proof Sketch. Let $C \subseteq N$ be any coalition with $i \notin C$. We claim that: (1) $v_{\mathcal{A}}(C \cup$ $\{i\})-v_{\mathcal{A}}(C) \geq w_{\ell}$ holds if, and only if, (2) either $i$ is not $\ell$-dependent; or the subgraph of $G_{\ell}$ induced by the nodes in $\mathbf{p}_{\ell}^{i}(N \backslash\{i\})$ contains a cycle; or there are agents $j$ and $j^{\prime}$ in $\mathbf{p}_{\ell}^{i}(N \backslash\{i\})$ with $\left|\Omega(j) \cap \Omega\left(j^{\prime}\right) \cap\left\{g \mid \operatorname{val}(g) \geq w_{\ell}\right\}\right|>1$; or $\mathbf{p}_{\ell}^{i}(N \backslash\{i\}) \neq \mathbf{p}_{\ell}^{i}(C)$.
$(1) \Rightarrow(2)$ Assume that (2) does not hold. So, the subgraph of $G_{\ell}$ induced by the nodes in $\mathbf{p}_{\ell}^{i}(N \backslash\{i\})$ is acyclic and $\mathbf{p}_{\ell}^{i}(N \backslash\{i\})=\mathbf{p}_{\ell}^{i}(C)$ holds, with $i \in \mathbf{p}_{\ell}^{i}(C)$. Moreover, for each pair $j, j^{\prime}$ of agents in $\mathbf{p}_{\ell}^{i}(C)$, we have $\mid \Omega(j) \cap \Omega\left(j^{\prime}\right) \cap\{g \mid \operatorname{val}(g) \geq$ $\left.w_{\ell}\right\} \mid \leq 1$. Therefore, by definition of $G_{\ell}$ and given that $s d_{\ell}(\mathcal{A}) \leq 2$, the goods in $\bigcup_{j \in \mathbf{p}_{\ell}^{i}(C)} \Omega(j)$ having value at least $w_{\ell}$ have a one-to-one correspondence with the edges of the subgraph of $G_{\ell}$ induced by the nodes in $\mathbf{p}_{\ell}^{i}(C)$, which is a tree. Hence, the number of these goods is $\left|\mathbf{p}_{\ell}^{i}(C)\right|-1$. Therefore, there is no allocation $\pi$ such that $\operatorname{val}(\pi(j)) \geq w_{\ell}$, for each $j \in \mathbf{p}_{\ell}^{i}(C)$. By Theorem 7.7, we conclude that (1) does not hold.
$(2) \Rightarrow(1)$ If $i$ is not $\ell$-dependent, then it is immediate that $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C) \geq w_{\ell}$. So assume that $i$ is $\ell$-dependent.

Assume that there are two agents $j^{\prime}$ and $j^{\prime \prime}$ in $\mathbf{p}_{\ell}^{i}(C)$ such that $\Omega\left(j^{\prime}\right) \cap \Omega\left(j^{\prime \prime}\right) \cap\{g \mid$ $\left.\operatorname{val}(g) \geq w_{\ell}\right\} \supseteq\left\{g^{\prime}, g^{\prime \prime}\right\}$ with $g^{\prime} \neq g^{\prime \prime}$. Consider a spanning tree $T$ of the subgraph of $G_{\ell}$ induced by the nodes in $\mathbf{p}_{\ell}^{i}(C)$, and let us root it at agent $j^{\prime}$. Consider then the
allocation $\bar{\pi}$ built as follows: $\bar{\pi}\left(j^{\prime}\right)=g^{\prime}, \bar{\pi}\left(j^{\prime \prime}\right)=g^{\prime \prime}$; and for each agent $j_{c} \neq j^{\prime \prime}$ whose parent in $T$ is $j_{p}$, let $\bar{\pi}\left(j_{c}\right)$ be any arbitrary good in $\Omega\left(j_{c}\right) \cap \Omega\left(j_{p}\right) \cap\{g \mid \operatorname{val}(g) \geq$ $\left.w_{\ell}\right\}$. Since $s d_{\ell}(\mathcal{A}) \leq 2, \bar{\pi}$ is well-defined. Moreover, $\bar{\pi}$ satisfies condition (2) in Theorem 7.7, and we can conclude that $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C) \geq w_{\ell}$.

In the rest, assume that for each pair $j, j^{\prime}$ of agents in $\mathbf{p}_{\ell}^{i}(N \backslash\{i\}), \mid \Omega(j) \cap \Omega\left(j^{\prime}\right) \cap\{g \mid$ $\left.\operatorname{val}(g) \geq w_{\ell}\right\} \mid \leq 1$. Consider the case where $\mathbf{p}_{\ell}^{i}(N \backslash\{i\}) \neq \mathbf{p}_{\ell}^{i}(C)$, and let $j$ be an agent in $\mathbf{p}_{\ell}^{i}(C)$ adjacent to an agent $j^{\prime} \in \mathbf{p}_{\ell}^{i}(N \backslash\{i\}) \backslash \mathbf{p}_{\ell}^{i}(C)$. Consider the spanning tree $T$ of the subgraph of $G_{\ell}$ induced by the nodes in $\mathbf{p}_{\ell}^{i}(C)$, and let us root it at agent $j$. Consider then the allocation $\bar{\pi}$ built as follows: $\bar{\pi}(j)=g$ with $\{g\}=\Omega(j) \cap \Omega\left(j^{\prime}\right) \cap\left\{g \mid \operatorname{val}(g) \geq w_{\ell}\right\}$; and for each agent $j_{c}$ whose parent in $T$ is $j_{p}$, let $\bar{\pi}\left(j_{c}\right)=g_{c}$ with $\left\{g_{c}\right\}=\Omega\left(j_{c}\right) \cap \Omega\left(j_{p}\right) \cap\left\{g \mid \operatorname{val}(g) \geq w_{\ell}\right\}$. Since $s d_{\ell}(\mathcal{A}) \leq 2, \bar{\pi}$ is well-defined. Moreover, $\bar{\pi}$ satisfies condition (2) in Theorem 7.7, and we conclude that $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C) \geq w_{\ell}$.

Finally, assume that $\mathbf{p}_{\ell}^{i}(N \backslash\{i\})=\mathbf{p}_{\ell}^{i}(C)$ and that for each pair $j, j^{\prime}$ of agents in $\mathbf{p}_{\ell}^{i}(N \backslash\{i\}),\left|\Omega(j) \cap \Omega\left(j^{\prime}\right) \cap\left\{g \mid \operatorname{val}(g) \geq w_{\ell}\right\}\right| \leq 1$. Assume that the subgraph of $G_{\ell}$ induced by the nodes in $\mathbf{p}_{\ell}^{i}(N \backslash\{i\})$ contains a cycle, say $j_{1}, \ldots, j_{k}, j_{1}$. Consider the spanning tree $T$ of the subgraph of $G_{\ell}$ induced by the nodes in $\mathbf{p}_{\ell}^{i}(C)$. Assume, w.l.o.g., that the edge $\left(j_{k-1}, j_{1}\right)$ is not in $T$ and let us root $T$ at $j_{1}$. Consider the allocation $\bar{\pi}$ such that: $\bar{\pi}\left(j_{1}\right)=g$ with $\{g\}=\Omega\left(j_{1}\right) \cap \Omega\left(j_{k-1}\right) \cap\left\{g \mid \operatorname{val}(g) \geq w_{\ell}\right\}$; and for each agent $j_{c}$ whose parent in $T$ is $j_{p}$, let $\bar{\pi}\left(j_{c}\right)=g_{c}$ with $\left\{g_{c}\right\}=\Omega\left(j_{c}\right) \cap \Omega\left(j_{p}\right) \cap\{g \mid$ $\left.\operatorname{val}(g) \geq w_{\ell}\right\}$. As above, $\bar{\pi}$ is well-defined and we have that $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C) \geq w_{\ell}$.

In the light of the above relationship, observe that if $i$ is not $\ell$-dependent, or the subgraph of $G_{\ell}$ induced by the nodes in $\mathbf{p}_{\ell}^{i}(N \backslash\{i\})$ contains a cycle, or there are agents $j$ and $j^{\prime}$ in $\mathbf{p}_{\ell}^{i}(N \backslash\{i\})$ with $\left|\Omega(j) \cap \Omega\left(j^{\prime}\right) \cap\left\{g \mid \operatorname{val}(g) \geq w_{\ell}\right\}\right|>1$, or $h+1<\left|\mathbf{p}_{\ell}^{i}(N \backslash\{i\})\right|$ (so that $\mathbf{p}_{\ell}^{i}(N \backslash\{i\}) \neq \mathbf{p}_{\ell}^{i}(C)$ ), then any coalition $C \subseteq N \backslash\{i\}$
with $|C|=h$ is such that $v_{\mathcal{A}}(C \cup\{i\})-v_{\mathcal{A}}(C) \geq w_{\ell}$. Indeed, these conditions are independent of $C$. Therefore, $\# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$ just consists of the overall number of coalitions $C$ of cardinality $h$ that can be built out of $n-1$ agents. This is given by $\frac{(n-1)!}{(n-1-h)!h!}$.

It remains to consider what happens if the conditions above that are independent of $h$ do not hold, and the cardinality $h$ of the considered coalitions is such that $h>\left|\mathbf{p}_{\ell}^{i}(N \backslash\{i\})\right|-1$. In this case, for all coalitions $C$ with $|C|=h$ obtained by extending $\mathbf{p}_{\ell}^{i}(N \backslash\{i\}) \backslash\{i\}$ with further $h-\left(\left|\mathbf{p}_{\ell}^{i}(N \backslash\{i\})\right|-1\right)$ agents, $v_{\mathcal{A}}(C \cup\{i\})-$ $v_{\mathcal{A}}(C)<w_{\ell}$ holds. Therefore, $\# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$ is obtained by subtracting the number of such coalitions from the whole number of cardinality- $h$ coalitions without agent $i$.

Because of this closed form characterization and of Theorem 7.3, the following is immediately established.

Theorem 7.10. The Shapley value (and the Banzhaf value) of allocation games $\mathcal{G}_{\mathcal{A}}$ can be computed in polynomial time on scenarios $\mathcal{A}=\langle N, G, \Omega$, val $\rangle$ such that $s d_{\ell}(\mathcal{A}) \leq 2$, for each value level $\ell$.

Assessing whether tractability still holds with higher values of sharing degree is left as an open problem. In fact, one may recall that computing the number of perfect matching in a graph (cf. proof of Theorem 7.1) is \#P-hard even on graph whose nodes have degree 3 [Dagum and Luby, 1992]. Hence, a negative answer to the question would be hardly surprising.

### 7.4.2 Bounded treewidth

In this subsection, we analyze the complexity of the Shapley value over allocation games where the interactions among agents have a tree-like structure. We use the
technical tools provided in Section 7.3, by combining them with CSP techniques that are of interest in their own.

For any scenario $\mathcal{A}=\langle N, \mathbb{G}, \Omega$, val $\rangle$, its interaction graph is the undirected graph $G(\mathcal{A})=(N, E)$ such that $\{i, j\} \in E$ if, and only if, there is a good $g \in \Omega(i) \cap \Omega(j)$.

Recall that a tree decomposition of a graph $G=(N, E)$ is a pair $\langle T, \chi\rangle$, where $T=(V, F)$ is a tree, and $\chi$ is a function assigning to each vertex $p \in V$ a set of nodes $\chi(p) \subseteq N$, such that the following conditions are satisfied: (1) $\forall b \in N, \exists p \in V$ such that $b \in \chi(p) ;(2) \forall\{b, d\} \in E, \exists p \in V$ such that $\{b, d\} \subseteq \chi(p) ;(3) \forall b \in N$, the set $\{p \in V \mid b \in \chi(p)\}$ induces a connected subtree of $T$. The width of $\langle T, \chi\rangle$ is $\max _{p \in V}|\chi(p)-1|$, and the treewidth of $G$ (short: $t w(G)$ ) is the minimum width over all its tree decompositions (see, e.g., [Robertson and Seymour, 1984]).

Example 7.11. Consider again the allocation scenario $\mathcal{A}_{0}$ that is reported in Figure 6.1. Then, $G\left(\mathcal{A}_{0}\right)$ is a clique over the nodes in $\{1,2,3\}$. The treewidth of $G\left(\mathcal{A}_{0}\right)$ is 2, as it is witnessed by the tree decomposition $\mathrm{TD}_{0}=\left\langle T_{0}, \chi_{0}\right\rangle$ where $T_{0}$ consists of one vertex $v$ only, and $\chi_{0}(v)=\{1,2,3\}$.

### 7.5 CSP encoding (for the Banzhaf value)

Recall that, given a CSP instance $\mathcal{I}$, deciding whether there is a solution (and compute one, if any) is generally NP-hard, but it is known to be feasible in polynomial time on classes of CSP instances $\mathcal{I}$ whose associated graphs have treewidth bounded by some given constant [Gottlob et al., 2013]. In particular, we already pointed out (in Chapter 2) that these kinds of structural tractability results have been generalized to counting problems too Theorem 2.4.

In order to show the tractability of the Shapley value over allocation games having associated graphs with bounded treewidth, we shall encode the computation of the


Figure 7.2:
Decomposition in Example 7.12-the label of the root modified as in the proof of Theorem 7.13 is on the top.
coefficients \# $\operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$-see Theorem 7.3-in terms of a counting problem over a suitably defined CSP instance and we shall then make use of Theorem 2.4.

We recall that the challenge to be faced is to use a constant number of possible values for the auxiliary variables. For instance, the natural encoding where some variable $X_{j}$ (associated with an agent $j \in N$ ) can take as values the goods in $\Omega(j)$ is not useful here. In fact, we propose an encoding that uses both the given allocation scenario $\mathcal{A}$ and a tree decomposition TD $=\langle T, \chi\rangle$ of $G(\mathcal{A})$. The idea is that each good is associated with some distinguished vertex of $T$, while suitable variables in the labels of the tree encode the roadmaps to reach such goods. In particular, their domain just contains the needed road signs (five values are enough). This is detailed below.

We start by building a tree decomposition with certain desirable properties. Let $\left\langle T^{\prime}, \chi^{\prime}\right\rangle$ be a tree decomposition of $G(\mathcal{A})$ whose width is $k>0$. Note that, for each good $g \in \mathbb{G}$, we are guaranteed about the existence of a vertex $v_{g}^{\prime}$ in $T^{\prime}$ such that $\chi\left(v_{g}^{\prime}\right) \supseteq\{j \mid g \in \Omega(j)\}$. Indeed, the agents in $\{j \mid g \in \Omega(j)\}$ form a clique in $G(\mathcal{A})$.

In a pre-processing step, we modify $\left\langle T^{\prime}, \chi^{\prime}\right\rangle$ by adding a fresh vertex $v_{g}$ as a child
of $v_{g}^{\prime}$, whose label is $\chi\left(v_{g}\right)=\chi\left(v_{g}^{\prime}\right) \cap\{j \mid g \in \Omega(j)\}$. By iterating over all goods, we get the desired tree where each good $g$ is associated with a distinguished vertex (in fact, leaf) $v_{g}$ labeled by the agents to whom $g$ can be allocated. Of course, the transformation is feasible in polynomial time. Eventually, we further transform the decomposition by making it binary: For each vertex $v$ with children $v_{1}, \ldots, v_{n}$, we can create a novel vertex $\bar{v}$ as a child of $v$ and with its label, by subsequently appending under it all these children but $v_{1}$. Let $\mathrm{TD}=\langle T, \chi\rangle$ be the resulting tree decomposition, having the same width as $\left\langle T^{\prime}, \chi^{\prime}\right\rangle$.

Example 7.12. Figure 7.2 illustrates a width-2 tree decomposition $\overline{\mathrm{TD}}_{0}$ of $G\left(\mathcal{A}_{0}\right)$, by evidencing the vertices that are univocally associated with the goods in $\left\{g_{1}, g_{2}, g_{3}, g_{4}\right\}$. Moreover, note that the decomposition is defined over a binary tree.

The input to our encoding is the allocation scenario $\mathcal{A}$, the agent $i \in N$, the natural number $\ell$, and the tree decomposition $\mathrm{TD}=\langle T, \chi\rangle$. Note that, for the moment, we do not consider the size $h$. Then, we define the encoding $\xi$ such that $\xi(\mathcal{A}, i, \ell, \mathrm{TD})$ is the CSP instance $\langle\operatorname{Var}, U, \mathbf{C}\rangle$, where

- $\operatorname{Var}=\bigcup_{j \in N}\left\{X_{j}, Y_{j}\right\} \cup\left\{X_{j}^{v} \mid v\right.$ is in $\left.T \wedge j \in \chi(v)\right\} ;$
- $U=\{0,1, \odot, \swarrow, \searrow, \uparrow\} ;$
and where $\mathbf{C}$ is defined as follows, with constraint relations being reported in tabular form in Figure 7.3:

1. For each agent $j \in N$ and vertex $v$ in $T$ with $j \in \chi(v)$, there is a constraint $\left(S_{v, j}, r_{v, j}\right)$ with $S_{v, j}=\left\{X_{j}^{v}, X_{j}\right\} ;$
2. For each good $g$ with $\operatorname{val}(g) \geq w_{\ell}$ and each $j \in \chi\left(v_{g}\right)$, there is a constraint $\left(S_{g, j}, r_{g, j}\right)$ such that $S_{g, j}=\left\{X_{j}^{v_{g}}\right\} ;$
3. For each agent $j \in N$, there is a constraint $\left(S_{j}, r_{j}\right)$ such that $S_{j}=\left\{Y_{j}, X_{j}\right\}$;

| $X_{j}^{v}$ | $X_{j}$ |  |
| :---: | :---: | :---: |
| 0 | 0 |  |
| $\odot$ | 1 | if $v$ is a vertex of the form $v_{g}$, for some good $g$ with val $(g) \geq w_{\ell}$ |
| $\swarrow$ | 1 | if $v$ is not a leaf, $v_{1}$ is its left child, and $j \in \chi\left(v_{1}\right)$ |
| $\searrow$ | 1 | if $v$ is not a leaf, $v_{2}$ is its right child, and $j \in \chi\left(v_{2}\right)$ |
| $\uparrow$ | 1 | if $v$ is not the root, $p$ is its parent, and $j \in \chi(p)$ |


| $X_{j}^{\text {vg }}$ |  |  | if $j \neq i$if $j \neq i$ | $Y_{j}$ | $X_{j}$ | $X^{\prime}{ }^{\prime}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | $Y_{j}$ | $X_{j}$ |  | 1 | 1 | ${ }^{\prime}$ |
| 0 | 1 | 1 |  | 0 | 0 | 0 |
| $\odot$ | 1 | 0 |  | 0 | 0 | 1 |
| $\uparrow$ | 0 | 0 |  | 1 | 0 | 0 |
|  |  |  |  | 1 | 1 | 0 |


| $X_{j}^{v}$ | $X_{j}^{v_{1}}$ | $X_{j}^{v_{2}}$ |
| :---: | :---: | :---: |
| 0 | 0 | 0 |
| $\odot$ | $\uparrow$ | $\uparrow$ |
| $\uparrow$ | $\uparrow$ | $\uparrow$ |
| $\swarrow$ | $\swarrow$ | $\uparrow$ |
| $\swarrow$ | $\searrow$ | $\uparrow$ |
| $\swarrow$ | $\odot$ | $\uparrow$ |
| $\searrow$ | $\uparrow$ | $\swarrow$ |
| $\searrow$ | $\uparrow$ | $\searrow$ |
| $\searrow$ | $\uparrow$ | $\odot$ |


| $X_{j}^{v g}$ | $X_{j^{\prime}}^{v g}$ |
| :---: | :---: |
| $\forall u \in\{0, \uparrow\}$ | $\odot$ |
| $\forall u \in\{0, \uparrow\}$ | $\odot$ |
| $\forall u \in\{0, \uparrow\}$ | $\odot$ |
| $\forall u \in\{0, \uparrow\}$ | $\odot$ |
| $\odot$ | $\forall u \in\{0, \uparrow\}$ |
| $\odot$ | $\forall u \in\{0, \uparrow\}$ |
| $\odot$ | $\forall u \in\{0, \uparrow\}$ |
| $\odot$ | $\forall u \in\{0, \uparrow\}$ |
| $\forall u \in\{0, \uparrow\}$ | $\forall u \in\{0, \uparrow\}$ |

Figure 7.3: CSP encoding in Section 7.5.
4. For each pair of agents $j \in N$ and $j^{\prime} \in N$ that are adjacent in $G_{\ell}$, there is a constraint $\left(S_{j, j^{\prime}}, r_{j, j^{\prime}}\right)$ such that $S_{j, j^{\prime}}=\left\{Y_{j}, X_{j}, X_{j^{\prime}}\right\}$;
5. For each non-leaf vertex $v$ whose left (resp., right) child is $v_{1}$ (resp., $v_{2}$ ), and for each $j \in \chi(v)$, there is a constraint $\left(S_{v, j}^{\prime}, r_{v, j}^{\prime}\right)$ such that $S_{v, j}^{\prime}=\left\{X_{j}^{v}, X_{j}^{v_{1}}, X_{j}^{v_{2}}\right\}$;
6. For each good $g$ with $\operatorname{val}(g) \geq w_{\ell}$ and for each pair $j, j^{\prime} \in \chi\left(v_{g}\right)$ with $g \in$ $\Omega(j) \cap \Omega\left(j^{\prime}\right)$, there is a constraint $\left(S_{g, j, j^{\prime}}, r_{g, j, j^{\prime}}\right)$ such that $S_{g, j, j^{\prime}}=\left\{X_{j}^{v_{g}}, X_{j^{\prime}}^{v_{g}}\right\}$;
7. No further constraint is in $\mathbf{C}$.

Theorem 7.13. The following properties hold:
(a) $\xi(\mathcal{A}, i, \ell, \mathrm{TD})$ can be built in polynomial time;
(b) the domain of each variable in $\xi(\mathcal{A}, i, \ell, \mathrm{TD})$ consists of at most 5 distinct elements;
(c) $t w(G(\xi(\mathcal{A}, i, \ell, \mathrm{TD}))) \leq 5 \times(t w(G(\mathcal{A}))+1)$;
(d) if $\theta$ is a solution to $\xi(\mathcal{A}, i, \ell, \mathrm{TD})$, then $R_{\theta}=\left\{j \mid \theta\left(Y_{j}\right)=1 \wedge j \neq i\right\}$ is such that $v_{\mathcal{A}}\left(R_{\theta} \cup\{i\}\right)-v_{\mathcal{A}}\left(R_{\theta}\right) \geq w_{\ell} ;$
(e) if $R \subseteq N \backslash\{i\}$ is such that $v_{\mathcal{A}}(R \cup\{i\})-v_{\mathcal{A}}(R) \geq w_{\ell}$, then there is a solution $\theta$ to $\xi(\mathcal{A}, i, \ell, \mathrm{TD})$ with $R=R_{\theta}$;
(f) $\sum_{h=0}^{n-1} \# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)=\left|\Theta\left(\xi(\mathcal{A}, i, \ell, \mathrm{TD}),\left\{Y_{1}, \ldots, Y_{n}\right\}\right)\right|$.

Proof Sketch. Property (a) and Property (b) are immediate.
Concerning Property (c) note that, if $\mathrm{TD}=\langle T, \chi\rangle$, then the tuple $\left\langle T, \chi_{\xi}\right\rangle$ such that for each vertex $v$ in $T, \chi_{\xi}(v)=\bigcup_{j \in \chi(v)}\left\{X_{j}, Y_{j}, X_{j}^{v}\right\} \cup\left\{X_{j}^{v_{1}}, X_{j}^{v_{2}} \mid v\right.$ is not a leaf $\}$ is a tree decomposition of $G(\xi(\mathcal{A}, i, \ell, \mathrm{TD}))$. Note that the decomposition does not depend on $i$ and $\ell$. As an example, the modified label associated with the root node of the tree decomposition of the graph $G\left(\mathcal{A}_{0}\right)$ in Example 7.11 is shown in Figure 7.2.

Concerning Property (d), assume that $\theta$ is a solution to $\xi(\mathcal{A}, i, \ell, \mathrm{TD})$ and let $\bar{R}_{\theta}$ be the set $\left\{j \mid \theta\left(X_{j}\right)=1 \wedge j \neq i\right\}$.

Let $j$ be any agent in $\bar{R}_{\theta} \cup\{i\}$. First, we claim that there is a vertex $v^{*}$ in $T$ such that $\theta\left(X_{j}^{v^{*}}\right)=\odot$. By contradiction, assume there is no such vertex. Let $v$ be the vertex in TD that is the closest to the root with $j \in \chi(v)$. Because of the constraint $\left(S_{v, j}, r_{v, j}\right)$ of type 1, we have that $\theta\left(X_{j}^{v}\right) \in\{\swarrow, \searrow\}$. If $v$ is a non-leaf vertex, then because of constraint $\left(S_{v, j}^{\prime}, r_{v, j}^{\prime}\right)$ of type 5 , we have that $\theta\left(X_{j}^{w}\right) \in\{\swarrow, \searrow\}$ holds, with $w \in\left\{v_{1}, v_{2}\right\}$ being one of its two children in $T$. In particular, given the constraint of type 1 , for the child $w$ it must be the case that $j \in \chi(w)$ holds. Therefore, we can apply the argument again on $w$, and so top-down from $w$ we can eventually reach a leaf $\bar{v}$ such that $\theta(\bar{v}) \in\{\swarrow, \searrow\}$. But, this is impossible by the constraint $\left(S_{\bar{v}, j}, r_{\bar{v}, j}\right)$ of type 1. So, we know that for each $j \in \bar{R}_{\theta} \cup\{i\}$, there is a vertex $v^{*}$ in $T$ such that $\theta\left(X_{j}^{v^{*}}\right)=\odot$. Moreover, for each vertex $w$ in the path connecting $v$ and $v^{*}$, $\theta(w) \in\{\swarrow, \searrow\}$. Therefore, because of constraints of type 5 , for each vertex $u$ with $j \in \chi(u)$ and not occurring in this path, it is the case that $\theta(u)=\uparrow$.

Hence, for each $j \in \bar{R}_{\theta} \cup\{i\}$, there is precisely one vertex $v^{*}$ such that $\theta\left(X_{j}^{v^{*}}\right)=\odot$.

Because of the constraints of type 1 and 2 , it holds that $v^{*}=v_{g}$, for some good $g \in \Omega(j)$ with $\operatorname{val}(g) \geq w_{\ell}$. Moreover, because of the constraints of type 6 , there is no other agent $j^{\prime}$ such that $\theta\left(X_{j^{\prime}}^{v^{*}}\right)=\odot$. In the light of these properties, the function $\pi: N \rightarrow \mathbb{G}$ such that $\pi(j)=g$, for each $j \in \bar{R}_{\theta} \cup\{i\}$ with $\theta\left(X_{j}^{v_{g}}\right)=\odot$, and $\pi(\bar{j})=\emptyset$, for each other agent $\bar{j}$, is well-defined and is an allocation. In particular, for each agent $j \in \bar{R}_{\theta} \cup\{i\}, \operatorname{val}(\pi(j)) \geq w_{\ell}$.

Let now $R_{\theta}=\left\{j \mid \theta\left(Y_{j}\right)=1 \wedge j \neq i\right\}$. Observe that, because of the constraints of type 3 , it holds that $\bar{R}_{\theta} \subseteq R_{\theta}$. In particular, $\theta\left(X_{i}\right)=1$. Moreover, note that because of the constraints of type 4 , whenever $\theta\left(X_{j^{\prime}}\right)=1$ and $\theta\left(Y_{j}\right)=1$ with $j$ and $j^{\prime}$ being adjacent in $G_{\ell}$, then $\theta\left(X_{j}\right)=1$ holds, too. Hence, $\bar{R}_{\theta} \cup\{i\} \supseteq \mathbf{p}_{\ell}^{i}\left(R_{\theta}\right)$. It follows that we can apply Theorem 7.7 on the coalition $R_{\theta}$, and we conclude that $v_{\mathcal{A}}\left(R_{\theta} \cup\{i\}\right)-v_{\mathcal{A}}\left(R_{\theta}\right) \geq w_{\ell}$.

Consider now Property (e). If $R \subseteq N \backslash\{i\}$ is a coalition with $v_{\mathcal{A}}(R \cup\{i\})-v_{\mathcal{A}}(R) \geq$ $w_{\ell}$, then by Theorem 7.7 there is an allocation $\pi$ such that $\operatorname{val}(\pi(j)) \geq w_{\ell}$, for each $j \in \mathbf{p}_{\ell}^{i}(R)$. Consider the substitution $\theta$ such that: $\theta\left(Y_{j}\right)=1$ iff $j \in R \cup\{i\} ; \theta\left(X_{j}\right)=1$ iff $j \in \mathbf{p}_{\ell}^{i}(R) ; \theta\left(X_{j}^{v_{g}}\right)=\odot$ iff $\pi(j)=g ; \theta\left(X_{j}^{v}\right)=0$ iff $j \notin R ; \theta\left(X_{j}^{v}\right)=\uparrow$ iff $\theta\left(X_{j}^{v_{g}}\right)=\odot$ holds for a vertex $v_{g}$ that is not in the subtree of $T$ rooted at $v ; \theta\left(X_{j}^{v}\right)=\swarrow$ (resp., $\left.\theta\left(X_{j}^{v}\right)=\searrow\right)$ if $\theta\left(X_{j}^{v_{g}}\right)=\odot$ for a vertex $v_{g}$ that occurs in the subtree rooted at the left (resp., right) child of $v$. By inspecting the constraints, it can be checked that $\theta$ is in fact a solution to $\xi(\mathcal{A}, i, \ell, \mathrm{TD})$.

Finally, Property (f) derives by Property (d), by Property (e), by the fact that $\theta\left(Y_{i}\right)=\theta\left(X_{i}\right)=1$ holds in any solution, and by the definition of $\# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$.

By combining Theorem 7.13 and Theorem 2.4, the following can be easilly established.

Corollary 7.14. The Banzhaf value of allocation games $\mathcal{G}_{\mathcal{A}}$ can be computed in polynomial time on scenarios $\mathcal{A}=\langle N, G, \Omega$, val $\rangle$ such that $\operatorname{tw}(G(\mathcal{A}))$ is bounded by a constant.

Proof. Recall that the Banzhaf value of $\mathcal{G}_{\mathcal{A}}$ can be written as $\beta_{i}\left(\mathcal{G}_{\mathcal{A}}\right)=\frac{1}{2^{n-1}} \sum_{h=0}^{n-1} \beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$. Given the expressions for $\beta_{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$ in Theorem 7.3, to compute the Banzhaf value in polynomial time, we need a method to efficiently evaluate $\sum_{h=0}^{n-1} \# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$, for each agent $i$ and level $\ell$. By Property $(f)$ in Theorem 7.13, this is possible on scenarios $\mathcal{A}$ where $t w(G(\mathcal{A}))$ is bounded by some constant $k$. Indeed, a width- $k$ tree decomposition can be computed in linear time.

### 7.6 From the Banzhaf value to the Shapley value

The encoding $\xi$ discussed so far does not take $h$ as a parameter. In fact, it just provides us a way to compute the value $\sum_{h=0}^{n-1} \# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$ and, hence, the Banzhaf value. In order to compute the contribution $\# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{\mathcal{A}}, h\right)$ for each cardinality of the coalitions (and, hence, the Shapley value by Theorem 7.3), we need a way to filter, out of all possible solutions, those $\theta$ such that $\left|R_{\theta}\right|=h$. This is not immediate (by preserving the structural properties and the bound on the domains), so that a careful construction is in order.

Let $\mathcal{I}$ be the instance $\langle\operatorname{Var}, U, \mathbf{C}\rangle$ with $\left\{Y_{1}, \ldots, Y_{n}\right\} \subseteq \operatorname{Var}$, let $\mathrm{TD}^{\prime}=\left\langle T^{\prime}, \chi^{\prime}\right\rangle$ be a binary tree decomposition of $G(\mathcal{I})$, and let $h \in\{0, \ldots, n-1\}$. Then, consider the modified CSP instance $\zeta\left(\mathcal{I},\left\{Y_{1}, \ldots, Y_{n}\right\}, h, \mathrm{TD}^{\prime}\right)=\left\langle\operatorname{Var}^{\prime}, U^{\prime}, \mathbf{C}^{\prime}\right\rangle$ such that:

- $\operatorname{Var}^{\prime}=\operatorname{Var} \cup\left\{W_{v} \mid v\right.$ is in $\left.T^{\prime}\right\} ;$
- $U^{\prime}=U \cup\{0, \ldots, h\} ;$
- $\mathbf{C}^{\prime}=\mathbf{C} \cup\left\{\left(S_{v}, r_{v}\right) \mid v\right.$ is in $\left.T^{\prime}\right\}$.

In particular, for each non-leaf vertex $v$ in $T^{\prime}$ with children $v_{1}$ and $v_{2}$, we have $S_{v}=\chi^{\prime}(v) \cup\left\{W_{v}, W_{v_{1}}, W_{v_{2}}\right\}$. Moreover, $r_{v}$ contains all possible substitutions $\theta$ over the variables in $S_{v}$ such that $\theta\left(W_{v}\right), \theta\left(W_{v_{1}}\right), \theta\left(W_{v_{2}}\right) \in\{0, \ldots, h\}$ and $\theta\left(W_{v}\right)-\mid\left\{Y_{j} \in\right.$ $\left.S_{v} \mid v=\operatorname{cr}(j) \wedge \theta\left(Y_{j}\right)=1\right\} \mid=\theta\left(W_{v_{1}}\right)+\theta\left(W_{v_{2}}\right)$, where $\operatorname{cr}(j)$ is the vertex $v^{*}$ that is the closest to the root and such that $Y_{j} \in \chi^{\prime}\left(v^{*}\right)$. Additionally, if $v$ is the root of $T^{\prime}$, then we require that $\theta\left(W_{v}\right)=h+1$ holds.

Instead, if $v$ is a leaf, then $S_{v}=\chi^{\prime}(v) \cup\left\{W_{v}\right\}$, and $r_{v}$ contains all possible substitutions $\theta$ over $S_{v}$ such that $\theta\left(W_{v}\right)=\left|\left\{Y_{j} \in S_{v} \mid v=\operatorname{cr}(j) \wedge \theta\left(Y_{j}\right)=1\right\}\right|$.

Theorem 7.15. Let $k$ be a fixed natural number, and assume that $\mathrm{TD}^{\prime}$ has treewidth $k$. Then, the following properties hold:
(a) $\zeta\left(\mathcal{I},\left\{Y_{1}, \ldots, Y_{n}\right\}, h, \mathrm{TD}^{\prime}\right)$ can be built in polynomial time;
(b) $\operatorname{tw}\left(G\left(\zeta\left(\mathcal{I},\left\{Y_{1}, \ldots, Y_{n}\right\}, h, \mathrm{TD}^{\prime}\right)\right)\right) \leq k+3$;
(c) if $\theta_{h}$ is a solution to $\zeta\left(\mathcal{I},\left\{Y_{1}, \ldots, Y_{n}\right\}, h, \mathrm{TD}^{\prime}\right)$, then there is exactly one solution $\theta$ to $\mathcal{I}$ with $\left|\left\{Y_{j} \mid \theta\left(Y_{j}\right)=1\right\}\right|=h$ and $\theta_{h} \supseteq \theta$.
(d) if $\theta$ is a solution to $\mathcal{I}$ such that $\left|\left\{Y_{j} \mid \theta\left(Y_{j}\right)=1\right\}\right|=h$, then there is exactly one solution $\theta_{h}$ to $\zeta\left(\mathcal{I},\left\{Y_{1}, \ldots, Y_{n}\right\}, h, \mathrm{TD}^{\prime}\right)$ with $\theta_{h} \supseteq \theta$;

Proof Sketch. Property (a) and Property (b) are immediate, by inspection of the construction. For Property (c), just observe that if $\theta_{h}$ is a solution to $\zeta\left(\mathcal{I},\left\{Y_{1}, \ldots, Y_{n}\right\}, h, \mathrm{TD}^{\prime}\right)$, then $\left|\left\{Y_{j} \mid \theta_{h}\left(Y_{j}\right)=1\right\}\right|=h$ and its restriction to the variables in $\operatorname{Var}$, say $\theta$, is a solution to $\mathcal{I}$. Finally, concerning Property (d), note that if $\theta$ is a solution to $\mathcal{I}$ with $\left|\left\{Y_{j} \mid \theta\left(Y_{j}\right)=1\right\}\right|=h$, then we can build a solution $\theta_{h} \supseteq \theta$ to $\zeta\left(\mathcal{I},\left\{Y_{1}, \ldots, Y_{n}\right\}, h, \mathrm{TD}\right)$ in a way that, for each variable $W_{v} \in\left(\operatorname{Var}^{\prime} \backslash \operatorname{Var}\right), \theta_{h}\left(W_{v}\right)$ is the number of variables $Y_{j}$ such that $\theta\left(Y_{j}\right)=1$ and $\operatorname{cr}(j)$ occurs in the subtree of $T$ rooted at $v$. In particular, note that there are no further assignments for these variables leading to
a solution.

By combining the above result with Theorem 7.13, the following can be established.

Theorem 7.16. The Shapley value of allocation games can be computed in polynomial time on all allocation scenarios whose interaction graphs have bounded treewidth.

Proof. Consider such a class of allocation scenarios $\mathcal{A}$ with $\operatorname{tw}(G(\mathcal{A})) \leq k$, for some fixed natural number $k$. Then, a width- $k$ tree decomposition TD of $G(\mathcal{A})$ and the encoding $\mathcal{I}=\xi(\mathcal{A}, i, \ell, \mathrm{TD})$ can be computed in polynomial time. Let $\mathrm{TD}^{\prime}$ be a tree decomposition of $G(\mathcal{I})$ whose width is bounded by $5 \times(k+1)$ (cf. Theorem 7.13). For each $h \in\{0, \ldots, n-1\}$, let $\mathcal{I}^{\prime}=\zeta\left(\mathcal{I},\left\{Y_{1}, \ldots, Y_{n}\right\}, h+1, \mathrm{TD}^{\prime}\right)$. Because of Theorem 7.15 and again Theorem 7.13, we know that $\left|\Theta\left(\mathcal{I}^{\prime},\left\{Y_{1}, \ldots, Y_{n}\right\}\right)\right|$ coincides with the value $\# \operatorname{col}_{\ell}^{i}\left(\mathcal{G}_{A}, h\right)$. We then get the Shapley value by using Theorem 7.3, and Equation (7.1).

Unfortunately, tractability does not follow here by just using Theorem 2.4 on $\mathcal{I}^{\prime}$ and $\left\{Y_{1}, \ldots, Y_{n}\right\}$, since the auxiliary variables of the form $W_{v}$ do not have a bounded domain. However, because of Property (c) and Property (d), we can add such variables to the output variables without altering the number of solutions, because $\left|\Theta\left(\mathcal{I}^{\prime},\left\{Y_{1}, \ldots, Y_{n}\right\}\right)\right|=\mid \Theta\left(\mathcal{I}^{\prime},\left\{Y_{1}, \ldots, Y_{n}\right\} \cup(\right.$ Var $\backslash$ Var $) \mid$ holds. Thus, we actually apply Theorem 2.4 on $\mathcal{I}^{\prime}$ with output variables $\left\{Y_{1}, \ldots, Y_{n}\right\} \cup\left(V a r^{\prime} \backslash V a r\right)$, and we still conclude that $\left|\Theta\left(\mathcal{I}^{\prime},\left\{Y_{1}, \ldots, Y_{n}\right\}\right)\right|$ can be computed in polynomial time.

### 7.7 Summary

We have studied the problem of computing the Shapley value (and the Banzhaf value) of allocation games, which are coalitional games implicitly (and succinctly) specified in terms of an underlying allocation scenario. We have shown that the
problem is \#P-complete, even in stringent settings. Motivated by this bad news, we identified islands of tractability by focusing either on scenarios with sharing degree at most 2 or such that the interactions among agents have a tree-like structure. This way, real world applications with useful structural properties can efficiently be dealt with, even if many agents and goods are involved. Moreover, the technical tools used to get the results may have a wider spectrum of applicability, beyond allocation problems.

A variant of the proposed framework considers scenarios where agents must necessarily get some good. In this case, it makes sense to have goods with negative values, too. In fact, we remark that our algorithms can be extended to manage these cases as well, by just considering suitable negative levels. Our work leaves open the technical question of whether tractability still holds over scenarios with sharing degree bounded by some constant greater than 2 (e.g., $s d_{\ell}(\mathcal{A}) \geq 3$ ). Moreover, it might stimulate further research to analyze the complexity of other solution concepts over allocation games, such as the nucleolus [Schmeidler, 1969].

## Part III

## CSP and Mining

## CHAPTER 8

## Process Mining

Process mining aims to discover, monitor and improve real processes by extracting knowledge from the event logs that are made available by today's information systems van der Aalst [2011]. A prominent process mining task is process discovery, whose goal is to facilitate the (re-)design phase and the implementation of complex process models. Indeed, process discovery algorithms are devoted to automatically deriving a model that can explain all the episodes recorded in an event $\log$ related to the execution of the activities of an underlying process. An approach to accomplish this task is presented which can benefit from the background knowledge that, in many cases, is available to the analysts taking care of the process (re-)design. The approach is based on encoding the information gathered from the log and the (possibly) given background knowledge in terms of precedence constraints, i.e., of constraints over the topology of the resulting process models. Mining algorithms are eventually formulated in terms of reasoning problems over precedence constraints, and the computational complexity of such problems is thoroughly analyzed by tracing their tractability frontier. Solution algorithms are proposed and their properties analyzed. These algorithms have been implemented in a prototype system, and results of a thorough experimental activity are discussed.

### 8.1 An overview of process discovery

Process discovery has emerged as a powerful approach to support the analysis and the design of complex processes. It consists of analyzing a set of traces registering the sequence of tasks performed along several enactments of a transactional system, in order to build a process model that can explain all the episodes recorded over them. Eventually, the "mined" model can be used to design a detailed process model suited to be supported in a workflow management system, or to shed light on the actual process behavior for optimization purposes (for instance, by singling out deviations between the intended conceptual model and the behavior that is actually registered).

In abstract terms, discovery algorithms carry out two different sub-tasks. First, they analyze the traces registering the sequences of activity executions over each enactment of the process, by mining the causal dependencies that are likely to hold among them. In particular, they present these dependencies in form of dependency graphs, that is, directed graphs whose nodes one-to-one correspond with the activities and where, intuitively, edges ingoing into an activity $a$ are meant to model that the possibility of executing $a$ is conditioned by the execution of the activities from which these edges originate. Second, they enrich dependency graphs with advanced facets of process enactments (such as synchronization constructs, branching constructs, duplicate activities, and invisible activities, just to name a few) and return process models formalized in expressive modeling languages. These concepts are next exemplified.

Example 8.1. Consider the set $\mathcal{A}=\{a, b, c, d, e\}$ of activities, and the structures depicted in Figure 8.1. We have three dependency graphs, $\mathcal{G}_{0}, \mathcal{G}_{1}$, and $\mathcal{G}_{2}$, which are meant to encode the causal relationships that hold over the activities in $\mathcal{A}$. Moreover,


Figure 8.1: Process models in the running example.
these graphs are adorned with an intuitive notation that expresses routing constructs over them. For instance, in $\mathcal{G}_{1}$, the flow of execution is split after a over two branches that are to be executed in parallel, and which are eventually synchronized by the activity d-see Section 8.3 for the formalization of the notation and of its associated semantics.

Assume now that the two traces abcde and acbde over $\mathcal{A}$ are given as input to a process discovery algorithm. Then, the graph $\mathcal{G}_{0}$ will be hardly returned as output, as it does not model the flow associated with the trace acbde, where $b$ is executed after c.

Instead, $\mathcal{G}_{1}$ and $\mathcal{G}_{2}$ are good candidates for a dependency graph supporting the two traces. The crucial observation here is that $b$ is executed before $c$ in one trace, while $b$ is executed after $c$ in the other. This likely witnesses that the two activities are not related by any causal dependency, and that they are executed in "parallel", i.e., over different branches of the process. Moreover, one might argue that the graph $\mathcal{G}_{1}$,
together with the associated routing constraints, is more appropriate than $\mathcal{G}_{2}$ for being returned as output. Indeed, there is heuristic evidence in the two given traces that $b$ and $d$ are not parallel, since $b$ seems to be a pre-requisite for the execution of $d$. $\triangleleft$

Several process models can, in general, be associated with a log of traces given as input. Ideally, one might select the most appropriate one among them under the assumption that the $\log$ is complete, i.e., that it registers all the possible traces for the underlying process. Indeed, in this case, if an activity always precedes another, then we can safely discard those models where such activities can be executed over different branches. For instance, if we assume that the traces abcde and acbde of Example 8.1 are the only possible ones, then we can safely conclude that $b$ and $d$ are not parallel activities, hence discarding the graph $\mathcal{G}_{2}$ in Figure 8.1.

Log completeness has received considerable attention in the literature, and it is a crucial assumption under which a number of discovery methods can be proven to be correct (see, e.g., van der Aalst et al. [2004, 2002]). In fact, process discovery is often carried out via heuristics approaches, for which the quality of the resulting models grows with the fraction of the traces given at hand w.r.t. all possible traces for the process. It follows that the quality can be rather poor in those cases where logs are far from being complete due, for instance, to the following reasons:

Temporal bias: First, two parallel activities may always appear in the same relative order, simply because one of them always finishes after the other, due to a different duration of them or of some of their predecessor activities. For instance, even for a process that conforms with the schema $\mathcal{G}_{2}$, we may find no trace where $d$ occurs before $b$, just because $c$ is time consuming, so that $b$ is always completed before $d$. Combinatorial explosion: Second, log completeness might not hold just because the process has not been enacted a sufficient number of times. Indeed, the number
of possible traces grows exponentially w.r.t. the number of activities that can be executed in parallel. For instance, for a process with $n$ branches each one involving $m$ activities, we have more than $n!^{m}$ possible traces. Therefore, for real processes even with just 2 branches and about 20 activities, the number of combinations immediately leads to more than one billion enactments, which in many application domains are definitively more than the enactments actually registered.

Access restrictions: Finally, even when logs are in principle complete, analysts might have a restricted access to them due to a number of reasons, ranging from privacy reasons, to specific policies adopted within the enterprise, and to issues related to the technologic infrastructure. In particular, these latter issues often occur when the process is not (fully) automatized within a unique environment, so that logs are not immediately available and they need to be built by integrating the raw data produced by different information platforms (which can be difficult, timeconsuming, or even just impossible if one of these platforms does not support logging mechanisms). In fact, it frequently happens in real-world applications that only a few representative traces are available (or can be reconstructed) for analysis purposes (for instance, related to a given time range or to some specific use cases of the overall process).

Because of the issues illustrated above, process discovery techniques are still at an early stage of adoption within enterprises. Indeed, analysts are likely to prefer traditional "top-down" design approaches, where models are eventually built by refining and formalizing a number of desiderata and specifications reflecting the prior knowledge they possess about the process to be automatized. For example, by a-priori knowing that $b$ and $d$ are parallel activities, the analyst can immediately discard $\mathcal{G}_{1}$ in Example 8.1, even though no trace is given where $d$ actually occurred before $b$.

### 8.2 Bottom-up vs top-down design methods

Several process discovery approaches have already been proposed in the literature, differing in the mining capabilities, in their goals, in the adopted mining methods, and in the kinds of modeling language they support. For example, processes are often intuitively represented via pure directed graphs Agrawal et al. [1998]; Weijters and van der Aalst [2001]; Weijters and van der Aalst [2003]; Weijters et al. [2006]; Greco et al. [2006]; Chen and Yun [2003], while more expressive representations are used in other proposals, ranging from expression tree models Schimm [2003] and block-structured workflow models Herbst and Karagiannis [2000, 2003]; Hammori et al. [2006], to special classes of Petri nets van der Aalst and van Hee [2002]; van der Aalst et al. [2002, 2004]; de Medeiros et al. [2004]; Medeiros et al. [2007]. Moreover, moving from the observation that extracting a single process model may lead to over-generalized models mixing different usage scenarios, classical discovery algorithms have been often combined with methods for clustering log traces, so that a set of process models can be returned as output. This improves the precision of the underlying algorithms by capturing-from an abstract perspective - constructs that are beyond the expressiveness of the given modeling languages (cf. Greco et al. [2006]).

Despite the technical differences that emerge from the non-exhaustive list of proposals discussed above, it must be pointed out that all of them share the idea of mining process models by gathering statistics from the data and by processing them via heuristics. In particular, these "bottom-up" discovery techniques are not capable to take into account prior knowledge on the underlying process. As a result, over logs that are not complete, mined models may well violate conceptual specifications
and domain-constraints (recall that, in Example 8.1, $\mathcal{G}_{1}$ would be the most probable outcome of a discovery algorithm, but the analyst might a-priori know that $b$ and $d$ are parallel activities), hence turning out to be useless in real-life applications.

Recently, the need of defining process discovery algorithms that can take advantage of prior knowledge has been pointed out by [Goedertier et al., 2009]. In particular, they presented the AGNEs (Artificial Generation of Negative Events) technique, which is a mining algorithm founding on an Inductive Logic Programming (ILP) classification-oriented learner. The algorithm is articulated in four steps: first, temporal constraints are extracted from the input log, in order to capture local dependencies, non-local dependencies, and parallelism relationships. Second, the input log and the temporal constraints are used to generate negative examples, i.e., for each prefix of any trace, negative events are generated stating which activities are not allowed to be executed later on in that trace. Third, by using input log traces (as positive examples) together with the artificial negative events, a logic program is induced to predict whether any given activity is allowed to occur at a given position of a given sequence. Finally, the logic program is translated into a Petri net. An important peculiarity of AGNEs is that the first of the above four phases allows experts to define a "top-down" partial specification of the process, in terms of background knowledge that will be subsequently used to guide the mining algorithm. In particular, it is possible to state that two activities are parallel (resp., not parallel), and that one precedes/succeeds (resp., does not precede/succeed) the other.

Another process discovery method (partially) taking into account domain knowledge has been proposed by [van der Werf et al., 2009]. As a reference model, again Petri nets are considered-moreover, the user can specify that the mined model must belong to some desired subclass of Petri nets (for instance, avoiding self loops or with
a fixed upper bound on the number of places/transitions). Starting with the most liberal (and overgeneralized) net for a given log, with as many transitions as the process activities and with no place, a more refined process model is obtained by iteratively adding a new place for restricting the allowed behavior. Each place is chosen greedily, by solving a system of integer linear inequalities, asking for a place with a minimal (resp., maximal) number of incoming edges (resp., outgoing edges). To curb the growth of the mined model, the search of the places is guided by the causal dependencies derived from the log (by using the metrics of [van der Aalst et al., 2004]). Moreover, in its implementation in the ProM framework [van Dongen et al., 2005], users are allowed to enforce finer grain constraints, by manually modifying the basic activity dependencies extracted from the log, prior to deriving a novel (refined) workflow model. By this way, the learning process can benefit from "top-down" specifications made available as domain knowledge expressed in terms of edge constraints and of constraints enforcing parallelism between pairs of activities.

### 8.3 Causal nets and logs

Several languages have been proposed in the literature which are tailored to the design and the analysis of processes, such as Petri nets [van der Aalst, 1998] or event driven process chains EPCs [van der Aalst et al., 2002]. Given the focus of the thesis, it is convenient to adopt a language that is closer to the needs of process mining applications. Accordingly, our choice is to consider the language of causal nets [van der Aalst et al., 2011], providing a favorable representational bias for such applications while having the same expressiveness as Petri nets.

Let us hereinafter assume that $\mathcal{A}$ is a given alphabet of symbols, univocally identifying the activities of some underlying process. The set $\mathcal{A}$ contains two distinguished
activities $a_{\perp}$ and $a_{\top}$, called the starting and the terminating activity, respectively.
A dependency graph (over $\mathcal{A})$ is a directed graph $\mathcal{G}=(V, E)$ whose nodes are the activities in the set $V \subseteq \mathcal{A}$, with $V \supseteq\left\{a_{\perp}, a_{\top}\right\}$, and whose edges in $E \subseteq V \times V$ encode the causal relationships that hold over them. In particular, for each activity $a \in V \backslash\left\{a_{\perp}, a_{\top}\right\}$, it must be the case that $a$ occurs in some path from $a_{\perp}$ to $a_{\top}$. Moreover, $a_{\perp}$ and $a_{\top}$ have no ingoing and outgoing edges, respectively.

Example 8.2. Consider again the three graphs $\mathcal{G}_{0}, \mathcal{G}_{1}$, and $\mathcal{G}_{2}$ depicted in Figure 8.1. It is immediate to check that they are dependency graphs over the set $\mathcal{A}=\{a, b, c, d, e\}$ of activities. In particular, a (resp., e) is the starting (resp., terminating) activity. $\triangleleft$

A causal net (over $\mathcal{A}$ ) is a tuple $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ where $\mathcal{G}=(V, E)$ is a dependency graph and where $\mathcal{I}$ and $\mathcal{O}$ are two functions such that:

- $\mathcal{I}$ maps each activity $a \in V$ to the set $\mathcal{I}(a)$ of the input bindings for $a$. For any $a \in V \backslash\left\{a_{\perp}\right\}$, an input binding $i b \in \mathcal{I}(a)$ is a non-empty set of edges such that $i b \subseteq\{(x, a) \mid(x, a) \in E\}$, while the empty set $\emptyset$ is the only input binding for $a_{\perp}$, i.e., $\mathcal{I}\left(a_{\perp}\right)=\{\emptyset\}$. For each $a \in V, \bigcup_{i b \in \mathcal{I}(a)} i b=\{(x, a) \mid(x, a) \in E\}$ must hold.
- $\mathcal{O}$ maps each activity $a \in V$ to the set $\mathcal{O}(a)$ of the output bindings for $a$. For any $a \in V \backslash\left\{a_{\mathrm{T}}\right\}$, an output binding $o b \in \mathcal{O}(a)$ is a non-empty set of edges such that $o b \subseteq\{(a, y) \mid(a, y) \in E\}$, while the empty set $\emptyset$ is the only output binding for $a_{\mathrm{T}}$, i.e., $\mathcal{O}\left(a_{\mathrm{T}}\right)=\{\emptyset\}$. For each $a \in V, \bigcup_{o b \in \mathcal{O}(a)} o b=\{(a, y) \mid(a, y) \in E\}$ must hold.

Intuitively, input bindings are meant to encode the pre-conditions for the execution of an activity, while output bindings are meant to encode the effects of this execution.

Example 8.3. Consider the causal net $\mathcal{C}_{0}=\left\langle\mathcal{G}_{0}, \mathcal{I}_{0}, \mathcal{O}_{0}\right\rangle$, where $\mathcal{G}_{0}=(\{a, b, c, d$, $\left.e\}, E_{0}\right)$ is the graph in Figure 8.1, and where $\mathcal{I}_{0}(z)=\left\{i b_{z}\right\}$ and $\mathcal{O}_{0}(z)=\left\{o b_{z}\right\}$ are such that $i b_{z}=\left\{(x, z) \mid(x, z) \in E_{0}\right\}$ and $o b_{z}=\left\{(z, y) \mid(z, y) \in E_{0}\right\}$, for each $z \in\{a, b, c, d, e\}$. The causal net models that the execution of $z$ can start as soon as its predecessor activity in $\mathcal{G}_{0}$ is completed-of course, this is immaterial for the starting activity $a$, which is such that $i b_{a}=\emptyset$. In fact, after its execution and if $z \neq e$, we have that $z$ enables the execution of its unique successor in $\mathcal{G}_{0}$.

Similarly, consider the causal net $\mathcal{C}_{1}=\left\langle\mathcal{G}_{1}, \mathcal{I}_{1}, \mathcal{O}_{1}\right\rangle$, where $\mathcal{G}_{1}=\left(\{a, b, c, d, e\}, E_{1}\right)$ is the graph in Figure 8.1, and where $\mathcal{I}_{1}(z)=\left\{i b_{z}^{\prime}\right\}$ and $\mathcal{O}_{1}(z)=\left\{o b_{z}^{\prime}\right\}$, for each activity $z$, are such that: $i b_{a}^{\prime}=\emptyset, i b_{b}^{\prime}=\{(a, b)\}, i b_{c}^{\prime}=\{(a, c)\}, i b_{d}^{\prime}=\{(b, d),(c, d)\}$, $i b_{e}^{\prime}=\{(d, e)\}, o b_{a}^{\prime}=\{(a, b),(a, c)\}, o b_{b}^{\prime}=\{(b, d)\}, o b_{c}^{\prime}=\{(c, d)\}, o b_{d}^{\prime}=\{(d, e)\}$, and $o b_{e}^{\prime}=\emptyset$. Note that, after its execution, the activity $a$ enables both $b$ and $c$ (in parallel), and that the flow is synchronized by $d$, which can be executed only once $b$ and c are both completed.

Note also that the symbols adorning the edges of the graphs in Figure 8.1 precisely correspond to the input and the output bindings. Formally, if $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ is a causal net, then each input binding $i b \in \mathcal{I}(z)$ (resp., output binding ob $\in \mathcal{O}(z)$ ) is represented by marking the edges in ib (resp., in ob) and by linking all such markings with a line. For instance, in the graph $\mathcal{G}_{1}$, the edges $(a, b)$ and $(a, c)$ are marked, and these markings are linked together, hence meaning that the output binding $\{(a, b),(a, c)\}$ occurs in $\mathcal{O}(a)$.

For completeness, we point out that any given activity can be associated with more than one input/output binding, in general. As an example, consider the causal net $\mathcal{C}_{2}=\left\langle\mathcal{G}_{2}, \mathcal{I}_{2}, \mathcal{O}_{2}\right\rangle$ associated with the graph $\mathcal{G}_{2}$ shown in Figure 8.1. Then, we may note, for instance, that $\mathcal{O}_{3}(a)=\{\{(a, b)\},\{(a, c)\},\{(a, b),(a, c)\}\}$ holds. Indeed, the
edges $(a, b)$ and $(a, c)$ are linked together as in the case of $\mathcal{G}_{1}$, but we additionally have the singleton markings now. Thus, a can either activate b, or $c$, or even both of them.

Formally, the semantics of causal nets is next given in terms of which instantiations are "globally valid". This is different from the operational semantics of designoriented modeling languages, such as the token-game semantics of Petri nets.

A binding activity for the causal net $\mathcal{C}$ is a tuple $\langle a, i b, o b\rangle$, where $a \in V$ is an activity and such that $i b \in \mathcal{I}(a)$ and $o b \in \mathcal{O}(a)$ hold. A sequence $\sigma$ of binding activities $\left\langle a_{1}=a_{\perp}, i b_{1}, o b_{1}\right\rangle, \ldots,\left\langle a_{n}=a_{\top}, i b_{n}, o b_{n}\right\rangle$ is called a binding sequence. The state $S_{j}^{\sigma}$ of $\mathcal{C}$ at the $j$-th step of $\sigma$ is defined inductively as the multi-set ${ }^{1}$ of edges such that $S_{0}^{\sigma}=\emptyset$, and $S_{j}^{\sigma}=S_{j-1}^{\sigma} \cup o b_{j} \backslash i b_{j}$, for each $j \in\{1, \ldots, n\}$. The sequence $\sigma$ is valid w.r.t. $\mathcal{C}$ if $S_{n}^{\sigma}=\emptyset$ and $i b_{j} \subseteq S_{j-1}^{\sigma}$, for each $j \in\{1, \ldots, n\}$.

Example 8.4. Consider the net $\mathcal{C}_{1}$ discussed in Example 8.3, and the binding sequence $\sigma=\sigma_{1} \sigma_{2} \ldots \sigma_{5}$ such that: $\sigma_{1}=\langle a,\{ \},\{(a, b),(a, c)\}\rangle, \sigma_{2}=\langle b,\{(a, b)\},\{(b, d)\}\rangle, \sigma_{3}=$ $\langle c,\{(a, c)\},\{(c, d)\}\rangle, \sigma_{4}=\langle d,\{(b, d),(c, d)\},\{(d, e)\}\rangle$, and $\sigma_{5}=\langle e,\{(d, e)\},\{ \}\rangle$. Note that $S_{0}^{\sigma}=\emptyset, S_{1}^{\sigma}=\{(a, b),(a, c)\}, S_{2}^{\sigma}=\{(a, c),(b, d)\}, S_{3}^{\sigma}=\{(b, d),(c, d)\}, S_{4}^{\sigma}=$ $\{(d, e)\}$, and $S_{5}^{\sigma}=\emptyset$. In fact, $\sigma$ is valid w.r.t. $\mathcal{C}_{1}$.

Transactional systems store partial information about binding sequences, by tracing the events related to the execution of the various activities. Formally, a trace $t$ (over $\mathcal{A}$ ) has the form $t[1] t[2] \ldots t[n]$, with $t[i] \in \mathcal{A}$ being an activity, for each $i \in\{1, \ldots, n\}$, and with $n$ being the length of $t$. W.l.o.g., we shall hereinafter assume that $t[1]=a_{\perp}, t[n]=a_{\text {T }}$, and that for each $i \in\{2, \ldots, n-1\}, t[i] \cap\left\{a_{\perp}, a_{\top}\right\}=\emptyset$. Indeed, we may possibly view $a_{\perp}$ and $a_{\top}$ as two virtual activities, which we add to $t$ in order to satisfy this requirement. The length of $t$ is also denoted as $\operatorname{len}(t)$. A

[^4]multi-set $L$ of traces is hereinafter just called a $l o g$, and the set of all the activities occurring over the traces in $L$ is denoted by $\mathcal{A}(L)$.

Note that, as commonly done in the literature, we are considering here an abstract view of a log, by focusing on what order the various activities were executed (specifically, completed) and by getting rid of all information about (i) timings (e.g., starting times and durations) and about (ii) the data involved in them.

Concerning the first assumption, note that our process modeling language is not capable to support temporal information. In fact, very few mining approaches have been proposed that are able to discover timed models, like stochastic Petri nets (see, e.g, Anastasiou et al. [2011]; Hu et al. [2011]). Therefore, in our setting, information about timings in the log can be helpful to a limited extent only. For instance, we can immediately conclude that two activities are parallel if one starts after the other, but before the completion of it. Moreover, for two activities $a$ and $b$ that are not parallel, in order to assess how likely $a$ is a pre-requisite for the execution of $b$, in addition of their relative positions (see Section 9.6), we can consider the time elapsed between the completion of $a$ and the starting of $b$. In the paper, however, we will not expand on these standard heuristics (conceived for models that do not support temporal information), by referring the interested reader, e.g., to the work by [Wen et al., 2009]. Here, we stress instead that it is an interesting avenue for further research to extend the features of causal nets by directly incorporating timing information and to explore how our techniques can be modified as to deal with the resulting model.

Concerning the second assumption, we point out that even the adaptation of basic process discovery algorithms to multi-dimensional settings, where modeling languages have again to be extended (in this latter case in order to deal with data about activity executions), has been only partially explored in the literature (see,
e.g., Greco et al. [2007]). In fact, this is still largely an open research issue, so that exporting our results to such richer settings is outside the scope of the paper, while constituting another interesting avenue for further research.

Now that we have clarified the assumptions underlying our setting, we can proceed to formalize when a log can be considered as the result of the enactments of a given process model. To this end, we say that a causal net $\mathcal{C}$ supports a trace $t$ if there is a binding sequence $\sigma$ that is valid w.r.t. $\mathcal{C}$ and where the $j$-th binding activity $\left\langle a_{j}, i b_{j}, o b_{j}\right\rangle$ of $\sigma$, for each $j \in\{1, \ldots$, len $(t)\}$, is such that $a_{j}=t[j]$. Moreover, we say that the causal net $\mathcal{C}$ supports a $\log L$, denoted by $\mathcal{C} \vdash L$, if $\mathcal{C}$ supports each trace $t \in L$.

Example 8.5. The causal net $\mathcal{C}_{1}$ discussed in Example 8.3 supports the trace abcde, as it is witnessed by the binding sequence $\sigma$ illustrated in Example 8.4. Moreover, it can be checked that $\mathcal{C}_{1}$ also supports the trace acbde, and that no further trace is supported.

Consider instead the causal net $\mathcal{C}_{2}$, again discussed in Example 8.3. Recall that a can activate either $b$, or $c$, or both of them. Moreover, by looking at the dependency graph $\mathcal{G}_{2}$ depicted in Figure 8.1, observe that $b$ and d are parallel activities. Thus, the traces that $\mathcal{C}_{2}$ supports are abe, acde, abcde, acbde, and acdbe. $\triangleleft$

We leave the section by pointing out that a useful extension of causal nets consists in allowing input and output bindings to be multi-sets, rather than just sets. With this extended model, for instance, an activity $x$ can activate two different instances of an activity $y$, for which $(x, y)$ is an edge in the underlying dependency graph. ${ }^{2}$ A causal net enriched with this capability will be hereinafter called an extended causal

[^5]net.

### 8.4 Dependency graphs and process mining: basic results

In process mining, a $\log L$ is given and the goal is to derive a process model supporting its traces. We next show that the main task for achieving this goal is essentially the discovery of the underlying dependency graph. In fact, based on this property, we can contextually show that the whole semantics of causal nets can be recast in simple graph-theoretic terms, which is convenient for our subsequent elaborations.

Definition 8.6. Let $L$ be a log. A dependency graph $\mathcal{G}$ acyclically supports $L$, denoted by $\mathcal{G} \vdash_{a} L$, if for each trace $t \in L$, there is a subgraph $\mathcal{G}_{t}$ of $\mathcal{G}$ such that:

- $\mathcal{G}_{t}$ is an acyclic dependency graph over $\mathcal{A}(\{t\})$, and
- $t$ is a topologic sort of $\mathcal{G}_{t}$, i.e., for each edge $(t[i], t[j])$ in $\mathcal{G}_{t}$, we have that $i<j$.

Example 8.7. Consider again the graph $\mathcal{G}_{2}$ in Figure 8.1, and note that $\mathcal{G}_{2} \vdash_{a}\{a b e$, acde, abcde, acbde, acdbe\}. For instance, given the trace abe, the subgraph of $\mathcal{G}_{2}$ induced over the activities $\{a, b, e\}$ is an acyclic dependency graph and abe is a topologic sort of it. Moreover, for any $\log L$ including a trace not in $\{a b e, a c d e, a b c d e, a c b d e, a c d b e\}$, we can check that $\mathcal{G}_{2} \vdash_{a} L$ does not hold.

At this point, it is interesting to observe that the set of traces supported by the causal net $\mathcal{C}_{2}$ (see Example 8.5) precisely coincides with the set of traces acyclically supported by the dependency graph $\mathcal{G}_{2}$ on top of which $\mathcal{C}_{2}$ was built. This is not by chance, and is intimately related to the fact that none of the given traces contains repetitions of the same activity. We next formalize this property.

First, we recall that a $\log L$ is linear if there is no trace in $L$ containing repetitions
of the same activity, i.e., for each $t \in L$ and for each $i, j \in\{1, \ldots$, len $(t)\}$ with $i \neq j$, $t[i] \neq t[j]$ holds. Then, we derive the following result.

Theorem 8.8. Let $L$ be a linear log, and let $\mathcal{G}$ be a dependency graph. Then, $\mathcal{G} \vdash_{a} L$ if, and only if, there is a causal net $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ such that $\mathcal{C} \vdash L$.

Proof. (if part). Let $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ be a causal net such that $\mathcal{C} \vdash L$. Let $t$ be any trace in $L$. Let $\mathcal{G}_{t}=\left(V_{t}, E_{t}\right)$ be the subgraph of $\mathcal{G}=(V, E)$ such that $V_{t}=$ $\{t[1], \ldots, t[\operatorname{len}(t)]\}$ and $E_{t}=\{(t[i], t[j]) \in E \mid 1 \leq i<j \leq \operatorname{len}(t)\}$. Hence, $\mathcal{G}_{t}$ is the subgraph of $\mathcal{G}$ induced over the activities occurring in $t$, where we keep those edges of $E$ that conform with the ordering of the activities in $t$. Since $L$ is linear, $\mathcal{G}_{t}$ is acyclic and $t$ is a topologic sort of it. According to Definition 8.6, to conclude, it is then sufficient to prove that $\mathcal{G}_{t}$ is a dependency graph. In fact, $a_{\perp}$ and $a_{\top}$ have no ingoing and outgoing edges, respectively, by construction of $\mathcal{G}_{t}$. The final requirement to be checked is now that, for each activity $a \in V \backslash\left\{a_{\perp}, a_{\top}\right\}$, it must be the case that $a$ occurs in some path from $a_{\perp}$ to $a_{\top}$. In particular, since $\mathcal{G}_{t}$ is acyclic, we can equivalently check that each activity $a \in V \backslash\left\{a_{\top}\right\}$ (resp., $a \in V \backslash\left\{a_{\perp}\right\}$ ) has at least one outgoing (resp., ingoing) edge.

Let $\sigma$ be a valid binding sequence such that $\sigma_{j}=\left\langle t[j], i b_{j}, o b_{j}\right\rangle$ is the $j$-th binding activity of $\sigma$, for each $j \in\{1, \ldots$, len $(t)\}$. Note that $\sigma$ exists, since $\mathcal{C} \vdash L$. Consider first an activity $a \in V \backslash\left\{a_{\mathrm{T}}\right\}$, i.e, $a=t[j]$ where $j \in\{1, \ldots$, len $(t)-1\}$. As $\sigma$ is a binding sequence, $o b_{j} \in \mathcal{O}(t[j])$ holds and hence we have that $o b_{j} \neq \emptyset$. In particular, there is an edge $(t[j], y) \in o b_{j}$ such that $(t[j], y) \in S_{j}^{\sigma}=S_{j-1}^{\sigma} \cup o b_{j} \backslash i b_{j}$. Moreover, again because $\sigma$ is valid, $S_{\operatorname{len}(t)}^{\sigma}=\emptyset$ holds, and hence there is an index $i \in\{j+1, \ldots, l e n(t)\}$ such that $(t[j], y) \in i b_{i}$. It follows that $y=t[i]$ and thus $(t[j], t[i])$ belongs to $E$ (hence, to $E_{t}$ since $j<i$ ). So, each activity $a \in V \backslash\left\{a_{T}\right\}$ has at least one outgoing edge.

We conclude by claiming that, for each $j \in\{2, \ldots, \operatorname{len}(t)\}$, there is an index $i \in$ $\{1, \ldots, j-1\}$ such that $(t[i], t[j])$ is in $E$ (hence, in $E_{t}$ ). In order to prove the claim, let again $\sigma$ be a valid binding sequence such that $\sigma_{j}=\left\langle t[j], i b_{j}, o b_{j}\right\rangle$ is the $j$-th binding activity of $\sigma$, for each $j \in\{1, \ldots$, len $(t)\}$. By definition of binding sequence, we know that $S_{j}^{\sigma} \subseteq S_{j-1}^{\sigma} \cup o b_{j}$ holds, for each $j \in\{1, \ldots, \operatorname{len}(t)\}$. In particular, since $\sigma$ is valid, $S_{0}^{\sigma}=\emptyset$ holds, and hence $S_{j}^{\sigma} \subseteq \bigcup_{h=1}^{j} o b_{h}$, for each $j \in\{1, \ldots$, len $(t)\}$. Recall now from the definition of causal net that $o b_{j} \in \mathcal{O}(t[j])$ implies $o b_{j} \subseteq\{(t[j], y) \mid(t[j], y) \in E\}$. Therefore, we conclude that:

$$
\begin{equation*}
S_{j}^{\sigma} \subseteq \bigcup_{h=1}^{j}\{(t[h], y) \mid(t[h], y) \in E\}, \text { for each } j \in\{1, \ldots, \text { len }(t)\} \tag{8.1}
\end{equation*}
$$

Let us exploit again the fact that $\sigma$ is valid, in order to derive that $i b_{j} \subseteq S_{j-1}^{\sigma}$ holds, for each $j \in\{1, \ldots$, len $(t)\}$. Let now $j$ be an index in the set $\{2, \ldots$, len $(t)\}$. Observe that, by definition of causal net, $i b_{j} \in \mathcal{I}(t[j])$ implies that $i b_{j} \subseteq\{(x, t[j]) \mid$ $(x, t[j]) \in E\}$, with $i b_{j}$ being, in particular, non empty. Therefore, by looking again at Equation 8.1 above and recalling that $i b_{j} \subseteq S_{j-1}^{\sigma}$, we can eventually conclude that an index $i \in\{1, \ldots, j-1\}$ exists such that $(t[i], t[j])$ occurs in $E$.
(only-if part). Assume that, for each trace $t \in L$, there is a subgraph $\mathcal{G}_{t}$ of $\mathcal{G}=(V, E)$ such that $\mathcal{G}_{t}$ is an acyclic dependency graph and $t$ is a topologic sort of $\mathcal{G}_{t}$. For each trace $t$ in $L$ and for each index $j \in\{1, \ldots$, len $(t)\}$, define the input binding $i b_{t, j}=\{(x, t[j]) \in E \mid x \in\{t[1], \ldots, t[j-1]\}\}$ and the output binding $o b_{t, j}=$ $\{(t[j], y) \in E \mid y \in\{t[j+1], \ldots, t[\operatorname{len}(t)]\}\}$. Note that, for each $j \in\{2, \ldots, \operatorname{len}(t)\}$, $i b_{t, j} \neq \emptyset$ holds. Indeed, as $\mathcal{G}_{t}$ is a dependency graph, for each $j \in\{2, \ldots$, len $(t)\}, t[j]$ has at least one ingoing edge $(x, t[j])$ in $\mathcal{G}_{t}$. Moreover, since $t$ is a topologic sort of $\mathcal{G}_{t}$, we actually have that $x \in\{t[1], \ldots, t[j-1]\}$. Similarly, it can be seen that, for each $j \in\{1, \ldots$, len $(t)-1\}, \quad o b_{t, j} \neq \emptyset$ holds. Finally, it is immediate to check that
$i b_{t, 1}=o b_{t, l e n(t)}=\emptyset$.
Now, consider the functions $\mathcal{I}$ and $\mathcal{O}$ such that, for each $a \in V$,

- $\mathcal{I}(a)=\bigcup_{t \in L, j \mid t[j]=a} i b_{t, j} \cup I_{a}$, where $I_{a}=\{(x, a) \mid(x, a) \in E\} ;$
- $\mathcal{O}(a)=\bigcup_{t \in L, j \mid t[j]=a} o b_{t, j} \cup O_{a}$, where $O_{a}=\{(a, y) \mid(a, y) \in E\}$.

Since $\mathcal{G}$ is a dependency graph and given the definition of $I_{a}$ and $O_{a}$, for each $a \in V$, it is immediate to check that $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ is a causal net. It remains to show that $\mathcal{C} \vdash L$, i.e., that for each trace $t$ in $L$, there is a valid binding sequence $\sigma$ such that $\left\langle t[j], i b_{j}, o b_{j}\right\rangle$ is the $j$-th binding activity of $\sigma$, for each $j \in\{1, \ldots$, len $(t)\}$. Let $t[1] \ldots t[n]$ be a trace, and consider the sequence $\sigma=\left\langle t[1], i b_{t, 1}, o b_{t, 1}\right\rangle, \ldots,\left\langle t[n], i b_{t, n}, o b_{t, n}\right\rangle$. We have to show that $S_{n}^{\sigma}=\emptyset$ and $i b_{t, j} \subseteq S_{j-1}^{\sigma}$, for each $j \in\{1, \ldots, n\}$.

To prove the result, we first claim that, for each $j \in\{0,1, \ldots, n\}, S_{j}^{\sigma}=\bigcup_{i=1}^{j}\left\{\left(t[i], t\left[i^{\prime}\right]\right) \mid\right.$ $\left.\left(t[i], t\left[i^{\prime}\right]\right) \in E, i^{\prime}>j\right\}$. In fact, the base case holds because $S_{0}^{\sigma}=\emptyset$, by definition of the state of a causal net. Now, assume that the property holds up to the index $h<j$. We have to show that $S_{h+1}^{\sigma}=\bigcup_{i=1}^{h+1}\left\{\left(t[i], t\left[i^{\prime}\right]\right) \in E \mid i^{\prime}>h+1\right\}$. To this end, recall that $S_{h+1}^{\sigma}=S_{h}^{\sigma} \cup o b_{t, h+1} \backslash i b_{t, h+1}$, so that $S_{h+1}^{\sigma}=\bigcup_{i=1}^{h}\left\{\left(t[i], t\left[i^{\prime}\right]\right) \in E \mid i^{\prime}>h\right\} \cup$ $o b_{t, h+1} \backslash i b_{t, h+1}$ holds, by inductive hypothesis. Eventually, the result derives by the above expression and the fact that $o b_{t, h+1}=\{(t[h+1], y) \in E \mid y \in\{t[h+2], \ldots, t[n]\}\}$ and $i b_{t, h+1}=\{(x, t[h+1]) \in E \mid x \in\{t[1], \ldots, t[h]\}\}$.

Armed with the above property, we can now resume the proof. First, we have to show that $S_{n}^{\sigma}=\emptyset$. In fact, we have that $S_{n}^{\sigma}=\bigcup_{i=1}^{n}\left\{\left(t[i], t\left[i^{\prime}\right]\right) \in E \mid i^{\prime}>n\right\}$, which coincides with the empty set as $n$ is the length of $t$. Second, we have to show that $i b_{t, j} \subseteq S_{j-1}^{\sigma}$, for each $j \in\{1, \ldots, n\}$. To this end, recall that $i b_{t, j}=\{(x, t[j]) \in E \mid$ $x \in\{t[1], \ldots, t[j-1]\}\}$, and eventually just check that $i b_{t, j} \subseteq \bigcup_{i=1}^{j-1}\left\{\left(t[i], t\left[i^{\prime}\right]\right) \in E \mid\right.$ $\left.i^{\prime}>j-1\right\}=S_{j-1}^{\sigma}$.


Figure 8.2: A process model involving a cycle, with an example unfolding.

To deal with arbitrary logs, we next introduce a mechanism for virtually unfolding cycles. For each trace $t$, let $\operatorname{unfold}(t)$ denote the trace obtained from $t$ by substituting the $i$-th occurrence in $t$ of any activity $a$ with the fresh (virtual) activity $a^{\langle i\rangle}$. Moreover, let unfold $(\mathcal{A})$ denote the (infinite) set of all virtual activities that can be built based on $\mathcal{A}$. The starting and terminating activity in $\operatorname{unfold}(\mathcal{A})$ are $a_{\perp}^{\langle 1\rangle}$ and $a_{T}^{\langle 1\rangle}$, respectively.

If $L$ is a $\log$, then we define $\operatorname{unfold}(L)$ as the $\operatorname{linear} \log \{\operatorname{unfold}(t) \mid t \in L\}$. Moreover, if $\overline{\mathcal{G}}=(\bar{V}, \bar{E})$ is a graph where $\bar{V} \subseteq \operatorname{unfold}(\mathcal{A})$, then we define the folding of $\overline{\mathcal{G}}$ as the directed graph fold $(\overline{\mathcal{G}})=(V, E)$ such that $V=\left\{x \in \mathcal{A} \mid \exists x^{\langle i\rangle} \in\right.$ $\operatorname{unfold}(\mathcal{A})$ s.t. $\left.x^{\langle i\rangle} \in \bar{V}\right\}$ and $E=\left\{(x, y) \in \mathcal{A} \times \mathcal{A} \mid \exists x^{\langle i\rangle}, y^{\langle j\rangle} \in \operatorname{unfold}(\mathcal{A})\right.$ s.t. $\left.\left(x^{\langle i\rangle}, y^{\langle j\rangle}\right) \in \bar{E}\right\}$.

Definition 8.9. Let $L$ be a log. A dependency graph $\mathcal{G}$ supports $L$, denoted by $\mathcal{G} \vdash L$,
if for each trace $t \in L$, there is a graph $\overline{\mathcal{G}}_{t}$ such that fold $\left(\overline{\mathcal{G}}_{t}\right)$ is a subgraph of $\mathcal{G}$, $\overline{\mathcal{G}}_{t} \vdash_{a}\{\operatorname{unfold}(t)\}$, and the following two conditions hold:
(1) there is no pair of edges $\left(x^{\langle i\rangle}, y^{\langle j\rangle}\right),\left(x^{\langle i\rangle}, y^{\left\langle j^{\prime}\right\rangle}\right)$ in $\overline{\mathcal{G}}_{t}$ such that $j \neq j^{\prime}$, and
(2) there is no pair of edges $\left(x^{\langle i\rangle}, y^{\langle j\rangle}\right),\left(x^{\left\langle i^{\prime}\right\rangle}, y^{\langle j\rangle}\right)$ in $\overline{\mathcal{G}}_{t}$ such that $i \neq i^{\prime}$.

Example 8.10. Consider the log consisting of the trace $t=$ acdcde, and note that unfold $($ acdcde $)=a^{\langle 1\rangle} c^{\langle 1\rangle} d^{\langle 1\rangle} c^{(2\rangle} d^{(2\rangle} e^{\langle 1\rangle}$. Moreover, check that the graph $\overline{\mathcal{G}}_{t}$ depicted on the right part of Figure 8.2 is such that $\overline{\mathcal{G}}_{t} \vdash_{a}\{$ unfold(acdcde) $\}$. In fact, fold( $\left.\overline{\mathcal{G}}\right)$ is a subgraph of the graph $\mathcal{G}_{3}$ depicted in the left part of the figure, and it is easily seen that conditions (1) and (2) of Definition 8.9 hold on $\overline{\mathcal{G}_{t}}$. Hence, $\mathcal{G}_{3} \vdash\{$ acdcde $\}$ holds.

Consider now the trace $t^{\prime}=a b b e$ and unfold $(a b b e)=a^{\langle 1\rangle} b^{\langle 1\rangle} b^{\langle 2\rangle} e^{\langle 1\rangle}$. Let $\overline{\mathcal{G}}_{t^{\prime}}$ be the graph consisting of the edges $\left(a^{\langle 1\rangle}, b^{\langle 1\rangle}\right),\left(a^{\langle 1\rangle}, b^{\langle 2\rangle}\right),\left(b^{\langle 1\rangle}, e^{\langle 1\rangle}\right)$, and $\left(b^{\langle 2\rangle}, e^{\langle 1\rangle}\right)$. Then, $\overline{\mathcal{G}}_{t^{\prime}} \vdash_{a}\{$ unfold $(a b b e)\}$, and fold $\left(\overline{\mathcal{G}}_{t^{\prime}}\right)$ is a subgraph of $\mathcal{G}_{3}$. However, $\overline{\mathcal{G}}_{t^{\prime}}$ violates conditions (1) and (2) in Definition 8.9. Therefore, $\mathcal{G}_{3}$ does not support $\{a b b e\} . \triangleleft$

Note that whenever the $\log L$ is linear, the above definition reduces to Definition 8.6 and, in particular, conditions (1) and (2) are immaterial. More formally, in this case, $\mathcal{G} \vdash L$ holds if, and only if, $\mathcal{G} \vdash_{a} L$ holds. Hence, for linear logs, we are in the position of applying Theorem 8.8 with ' $\vdash$ ' in place of ' $\vdash_{a}$ '. More generally, the following result is established in order to relate the notion of support over dependency graphs with the notion of support over causal nets. The result is of interest in its own, as it allows to restate the semantics of causal nets in pure graph-theoretic terms. In fact, it will play an important role in our subsequent elaborations.

Theorem 8.11. Let $L$ be a $\log$, and let $\mathcal{G}$ be a dependency graph. Then, $\mathcal{G} \vdash L$ if, and only if, there is a causal net $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ such that $\mathcal{C} \vdash L$.

Proof. (if part). Let $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ be a causal net such that $\mathcal{C} \vdash L$, with $\mathcal{G}=(N, E)$. This means that, for each trace $t$ in $L$, there is a binding sequence $\sigma$ such that $\sigma$ is valid w.r.t. $\mathcal{C}$, and the $j$-th element of $\sigma$ has the form $\left\langle t[j], i b_{j}, o b_{j}\right\rangle$, for each $j \in\{1, \ldots$, len $(t)\}$. Let $S_{j}^{\sigma}$ denote the state of $\mathcal{C}$ at the $j$-th step of $\sigma$, and consider the sequence $\bar{\sigma}$ having the same length $\operatorname{len}(t)$ as $\sigma$ and whose $j$-th element $\left\langle\bar{a}_{j}, \bar{i}_{j}, \overline{o b}_{j}\right\rangle$, for each $j \in\{1, \ldots$, len $(t)\}$, is obtained from $\left\langle t[j], i b_{j}, o b_{j}\right\rangle$ as follows:

- $\bar{a}_{j}$ is the virtual activity $t[j]^{\langle k\rangle}$, where $k$ is the number of occurrences of the activity $t[j]$ in $t[1] \ldots t[j]$, i.e., the number of binding activities defined over $t[j]$ in the first $j$ elements of $\sigma$;
- For each output binding $(t[j], y) \in o b_{j}$, let $\alpha_{y}$ denote the number of binding activities occurring up to the $j$-th step of $\sigma$ where $(t[j], y)$ occurs as an element of the output binding. Note that, since $\sigma$ is a valid sequence w.r.t. $\mathcal{C}$ (and since bindings are sets, i.e., multiple occurrences are not allowed), we are guaranteed about the existence of $\alpha_{y}$ binding activities where these output bindings are consumed, i.e., where they are taken as input. Let $n \operatorname{ext}(j, y)$ be the index of the $\alpha_{y}$-th activity binding of this kind, hence, in particular with $(t[j], y) \in i b_{n e x t}(j, y)$, and note that $\operatorname{next}(j, y)>j$. Then, we define $\overline{o b}_{j}=\left\{\left(\bar{a}_{j}, \bar{a}_{\text {next }(j, y)}\right) \mid(t[j], y) \in\right.$ $\left.o b_{j}\right\}$.
- For each input binding $(x, t[j]) \in i b_{j}$, let $\beta_{y}$ denote the number of binding activities occurring up to the $j$-th step of $\sigma$, where $(x, t[j])$ is taken as input. Note that, since $\sigma$ is a valid sequence w.r.t. $\mathcal{C}$, we are guaranteed about the existence of $\beta_{y}$ binding activities where these input bindings are produced. Let $\operatorname{prev}(j, x)$ be the index of the $\beta_{y}$-th activity binding of this kind, hence, in particular with $(x, t[j]) \in \operatorname{ob}_{\operatorname{prev}(j, x)}$, and note that $\operatorname{prev}(j, x)<j$. Then, we define $\overline{i b}_{j}=\left\{\left(\bar{a}_{\text {prev }(j, x)}, \bar{a}_{j}\right) \mid(x, t[j]) \in i b_{j}\right\}$.

Define now $\overline{\mathcal{G}}_{t}=\left(\bar{V}_{t}, \bar{E}_{t}\right)$ as the graph where $\bar{V}_{t}=\mathcal{A}(\{$ unfold $(t)\})$ and where $\bar{E}_{t}=$ $\bigcup_{j=1}^{l e n(t)}\left(\overline{i b}_{j} \cup \overline{o b}_{j}\right)$. We claim that $\overline{\mathcal{G}}_{t}$ is a dependency graph. Indeed, note first that $\overline{\mathcal{G}}_{t}$ is acyclic because each edge in $\bar{E}_{t}$ has the form $\left(\bar{a}_{i}, \bar{a}_{j}\right)$, with $i<j$. In particular, $t[1]^{\langle 1\rangle}$ and $t[l e n(t)]^{\langle 1\rangle}$ play the role of the starting and the terminating activity, respectively. Then, consider an activity $\bar{a}_{j} \in \bar{V}_{t} \backslash\left\{t[1]^{\langle 1\rangle}\right\}$ (resp., $\bar{a}_{j} \in \bar{V}_{t} \backslash\left\{t[\operatorname{len}(t)]^{\langle 1\rangle}\right\}$ ). Note that $\bar{a}_{j}$ has at least one ingoing (resp., outgoing) edge in $\bar{E}_{t}$, because there is at least an edge of the form $(x, t[j])$ (resp., $(t[j], y)$ ) in $i b_{j}$ (resp., $o b_{j}$ ), by the fact that $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ is a causal net (hence, $i b_{j} \neq \emptyset$ and $o b_{j} \neq \emptyset$ hold) and that $\sigma$ is a binding sequence. Since $\overline{\mathcal{G}}_{t}$ is acyclic, the above properties entail that $\bar{a}_{j}$ occurs in a path from $t[1]^{\langle 1\rangle}$ to $t[\operatorname{len}(t)]^{\langle 1\rangle}$.

Consider now the functions $\overline{\mathcal{I}}_{t}$ and $\overline{\mathcal{O}}_{t}$ such that $\overline{\mathcal{I}}_{t}\left(\bar{a}_{j}\right)=\left\{\overline{i b}_{j}\right\}$ and $\overline{\mathcal{O}}_{t}\left(\bar{a}_{j}\right)=\left\{\overline{o b}_{j}\right\}$, for each $\bar{a}_{j} \in \bar{V}_{t}$. Since $\overline{\mathcal{G}}_{t}$ is a dependency graph and given the construction of its edges, we derive that $\left\langle\overline{\mathcal{G}}_{t}, \overline{\mathcal{I}}_{t}, \overline{\mathcal{O}}_{t}\right\rangle$ is a causal net (over $\left.\mathcal{A}(\{\operatorname{unfold}(t)\})\right)$. Moreover, we claim that $\bar{\sigma}$ is valid w.r.t. $\left\langle\overline{\mathcal{G}}_{t}, \overline{\mathcal{I}}_{t}, \overline{\mathcal{O}}_{t}\right\rangle$. To prove the claim, note first that, given the above construction for $\bar{\sigma}$, if $i=\operatorname{prev}(j, x)$ (resp., $i=\operatorname{next}(j, y)$ ), then $x=t[i]$ (resp., $y=t[j])$ and $j=\operatorname{next}(i, t[j])$ (resp., $j=\operatorname{prev}(i, t[j]))$. Then, consider the state $S_{j}^{\bar{\sigma}}$ of $\left\langle\overline{\mathcal{G}}_{t}, \overline{\mathcal{I}}_{t}, \overline{\mathcal{O}}_{t}\right\rangle$ at the $j$-th step of $\bar{\sigma}$, and observe that the following properties hold.

- For each $j \in\{1, \ldots$, len $(t)\}, \bar{i}_{j} \subseteq S_{j-1}^{\bar{\sigma}}$. Indeed, consider an element $\left(\bar{a}_{\text {prev }(j, x)}, \bar{a}_{j}\right)$ in $\bar{i}_{j}$, let $i=\operatorname{prev}(j, x)$, and recall that $\left(\bar{a}_{\operatorname{prev}(j, x)}, \bar{a}_{j}\right)=\left(\bar{a}_{i}, \bar{a}_{\text {next }(i, t[i])}\right)$. Therefore, $\left(\bar{a}_{\operatorname{prev}(j, x)}, \bar{a}_{j}\right)$ occurs in $\overline{o b}_{i}$ and, hence, in $S_{i}^{\bar{\sigma}}=S_{i-1}^{\bar{\sigma}} \cup \overline{o b}_{i} \backslash \overline{i b}_{i}$. Eventually, as $\bar{a}_{j}$ occurs only at the $j$-th step of $\bar{\sigma}$ and $i<j$, we have that $\left(\bar{a}_{\text {prev }(j, x)}, \bar{a}_{j}\right) \in S_{j-1}^{\bar{\sigma}}$.
- $S_{\text {len }(t)}^{\bar{\sigma}}=\emptyset$. Indeed, assume by contradiction that $(\bar{x}, \bar{y})$ occurs in $S_{\text {len }}^{\bar{\sigma}}(t)$. Then, there is an index $j$ such that $\bar{x}=\bar{a}_{j}$ and $\left(\bar{a}_{j}, \bar{y}\right)$ occurs in $\bar{o} b_{j}$. By construction, we therefore have that $(\bar{x}, \bar{y})=\left(\bar{a}_{j}, \bar{a}_{\text {next }(j, y)}\right)$. However, by letting $i=\operatorname{next}(j, y)$, we can write that $\left(\bar{a}_{j}, \bar{a}_{\text {next }(j, y)}\right)=\left(\bar{a}_{\text {prev }(i, t[j])}, \bar{a}_{i}\right) \in \bar{i}_{i}$. Since $i>j$, it follows
that $(\bar{x}, \bar{y}) \notin S_{i}^{\bar{\sigma}}$, as $\bar{x}$ occurs only at the $j$-th step of $\bar{\sigma}$. We conclude that $(\bar{x}, \bar{y}) \notin S_{\text {len }(t)}^{\bar{\sigma}}$. Contradiction.

By putting together the above results, we have so far shown that $\left\langle\overline{\mathcal{G}}_{t}, \overline{\mathcal{I}}_{t}, \overline{\mathcal{O}}_{t}\right\rangle$ is a causal net with $\left\langle\overline{\mathcal{G}}_{t}, \overline{\mathcal{I}}_{t}, \overline{\mathcal{O}}_{t}\right\rangle \vdash\{\operatorname{unfold}(t)\}$, for each $t$ in $L$. It follows that we can apply Theorem 8.8 on $\left\langle\overline{\mathcal{G}}_{t}, \overline{\mathcal{I}}_{t}, \overline{\mathcal{O}}_{t}\right\rangle$, and we derive that $\overline{\mathcal{G}}_{t} \vdash_{a}\{\operatorname{unfold}(t)\}$ holds, for each $t$ in $L$.

Finally, consider the graph $\overline{\mathcal{G}}_{t}$ and note that conditions (1) and (2) in Definition 8.9 are trivially satisfied, by construction of its edges. Indeed, all edges outgoing from $\bar{a}_{j}$, with $j \in\{1, \ldots, l e n(t)\}$, have the form $\left(\bar{a}_{j}, \bar{a}_{\text {next }(j, y)}\right)$, where $(t[j], y) \in o b_{j}$. In particular, $\left(\bar{a}_{j}, \bar{a}_{\text {next }(j, y)}\right)$ is univocally determined by $y$, so that $\left(\bar{a}_{j}, \bar{a}_{\text {next }(j, y)}\right) \neq\left(\bar{a}_{j}, \bar{a}_{\text {next }\left(j, y^{\prime}\right)}\right)$ implies that $y \neq y^{\prime}$ and hence $\bar{a}_{\text {next }(j, y)}$ and $\bar{a}_{\text {next }\left(j, y^{\prime}\right)}$ are virtual activities built from different true activities. A similar line of reasoning applies to the edges incoming into $\bar{a}_{j}$. Moreover, fold $\left(\overline{\mathcal{G}}_{t}\right)$ is a subgraph of $\mathcal{G}$. Indeed, $\bar{E}_{t}$ is the union of all bindings in $\left\langle\overline{\mathcal{G}}_{t}, \overline{\mathcal{I}}_{t}, \overline{\mathcal{O}}_{t}\right\rangle$, and any element in these bindings is of the form $\left(x^{\langle i\rangle}, y^{\langle j\rangle}\right)$, where $(x, y)$ occurs in a binding of $\mathcal{I}$ or $\mathcal{O}$ and, hence, is an edge in $E$. Therefore, according to Definition 8.9, we have shown that $\mathcal{G} \vdash L$ holds.
(only-if part). Assume that $\mathcal{G} \vdash L$ holds, with $\mathcal{G}=(V, E)$. Thus, for each trace $t \in L$, there is a graph $\overline{\mathcal{G}}_{t}=\left(\bar{V}_{t}, \bar{E}_{t}\right)$ such that fold $\left(\overline{\mathcal{G}}_{t}\right)$ is a subgraph of $\mathcal{G}$, $\overline{\mathcal{G}_{t}} \vdash_{a}\{\operatorname{unfold}(t)\}$, and the following two conditions hold:
(1) there is no pair of edges $\left(x^{\langle i\rangle}, y^{\langle j\rangle}\right),\left(x^{\langle i\rangle}, y^{\left\langle j^{\prime}\right\rangle}\right)$ in $\overline{\mathcal{G}}_{t}$ such that $j \neq j^{\prime}$, and
(2) there is no pair of edges $\left(x^{\langle i\rangle}, y^{\langle j\rangle}\right),\left(x^{\left\langle i^{\prime}\right\rangle}, y^{\langle j\rangle}\right)$ in $\overline{\mathcal{G}}_{t}$ such that $i \neq i^{\prime}$.

By Theorem 8.8 applied on the fact that $\overline{\mathcal{G}}_{t} \vdash_{a}\{$ unfold $(t)\}$ holds, we derive that there is a causal net $\overline{\mathcal{C}}_{t}=\left\langle\overline{\mathcal{G}}_{t}, \overline{\mathcal{I}}_{t}, \overline{\mathcal{O}}_{t}\right\rangle$ such that $\overline{\mathcal{C}}_{t} \vdash\{\operatorname{unfold}(t)\}$ holds. This means that there is binding sequence $\bar{\sigma}$ that is valid w.r.t. $\overline{\mathcal{C}}$ and where the $j$-th binding
activity $\left\langle\bar{a}_{j}, \bar{i}_{j}, \bar{o} \bar{b}_{j}\right\rangle$ of $\bar{\sigma}$, for each $j \in\{1, \ldots$, len $(t)\}$, is such that $\bar{a}_{j}$ is the symbol $t[j]^{\langle k\rangle}$, with $k$ being the number of occurrences of $t[j]$ in $t[1] \ldots t[j]$.

Let $\mathcal{I}_{t}$ (resp., $\mathcal{O}_{t}$ ) be the function such that, for each node $z$ in $\operatorname{fold}\left(\overline{\mathcal{G}}_{t}\right), \mathcal{I}_{t}(z)=$ $\left\{i b_{z} \mid\right.$ exists $j$ s.t. $\left.\overline{i b}_{z} \in \overline{\mathcal{I}}_{t}\left(z^{\langle j\rangle}\right)\right\}\left(\right.$ resp., $\mathcal{O}_{t}(z)=\left\{o b_{z} \mid\right.$ exists $i$ s.t. $\left.\left.\overline{o b}_{z} \in \overline{\mathcal{O}}_{t}\left(z^{\langle i\rangle}\right)\right\}\right)$, where $i b_{z}$ (resp., $o b_{z}$ ) is the binding obtained from $\overline{i b}_{z}$ (resp., $\bar{o} b_{z}$ ) by stripping off the instantiation numbers of the virtual symbols. Similarly, define $\sigma$ as the sequence obtained from $\bar{\sigma}$ by stripping off the instantiation numbers of the virtual symbols. Note that because of the conditions (1) and (2) above, and since bindings are defined over the edges of $\overline{\mathcal{G}}_{t},\left|\bar{b} \bar{b}_{z}\right|=\left|i b_{z}\right|$ and $\left|\overline{o b} b_{z}\right|=\left|o b_{z}\right|$ hold, for each activity $z$. Therefore, if $S_{j}^{\bar{\sigma}}$ is the state of $\overline{\mathcal{C}}_{t}$ at the $j$-th step of $\bar{\sigma}$, then the state $S_{j}^{\sigma}$ of $\mathcal{C}_{t}$ at the $j$-th step of $\sigma$ can be obtained from $S_{j}^{\bar{\sigma}}$ by just stripping off the instantiation numbers of the symbols it contains. Hence, since $\bar{\sigma}$ is valid w.r.t. $\overline{\mathcal{C}}_{t}$, we can conclude that $\sigma$ is valid w.r.t. $\mathcal{C}_{t}=\left\langle\operatorname{fold}\left(\overline{\mathcal{G}}_{t}\right), \mathcal{I}_{t}, \mathcal{O}_{t}\right\rangle$. This witnesses that $\mathcal{C}_{t} \vdash\{t\}$ holds.

Let now $\overline{\mathcal{G}}=(\bar{V}, \bar{E})$ be the graph such that $\bar{V}=\bigcup_{t \in L} \bar{V}_{t}$ and $\bar{E}=\bigcup_{t \in L} \bar{E}_{t}$, and let $\mathcal{I}$ and $\mathcal{O}$ be the functions such that $\mathcal{I}(z)=\bigcup_{t \in L} \mathcal{I}_{t}(z)$ and $\mathcal{O}(z)=\bigcup_{t \in L} \mathcal{O}_{t}(z)$, for each $z$ in $\operatorname{fold}(\overline{\mathcal{G}})$. Note that fold $(\overline{\mathcal{G}})$ is a subgraph of $\mathcal{G}$, and that $\mathcal{C}=\langle$ fold $(\overline{\mathcal{G}}), \mathcal{I}, \mathcal{O}\rangle$ is a causal net such that $\mathcal{C} \vdash L$. Therefore, in the case where $\operatorname{fold}(\overline{\mathcal{G}})=\mathcal{G}$, then we have derived that there is a causal net $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ such that $\mathcal{C} \vdash L$. To conclude the proof, the only remaining case to be analyzed is when $\operatorname{fold}(\overline{\mathcal{G}})=\left(V_{f}, E_{f}\right)$ is a proper subgraph of $\mathcal{G}$. In this case, consider the functions $\mathcal{I}^{\prime}$ and $\mathcal{O}^{\prime}$ such that, for each $z \in V \cap V_{f}, \mathcal{I}^{\prime}(z)=\mathcal{I}(z) \cup I_{z}$ and $\mathcal{O}^{\prime}(z)=\mathcal{O}^{\prime}(z) \cup O_{z}$, and for each $z \in V \backslash V_{f}, I^{\prime}(z)=I_{z}$ and $\mathcal{O}^{\prime}(Z)=O_{z}$, where $I_{z}=\{\{(x, z)\} \mid(x, z) \in E\}$ and $O_{z}=\{\{(z, y)\} \mid(z, y) \in E\}$. As $\mathcal{G}$ is a dependency graph, by construction of $\mathcal{I}^{\prime}$ and $\mathcal{O}^{\prime}$, we trivially have that $\left\langle\mathcal{G}, \mathcal{I}^{\prime}, \mathcal{O}^{\prime}\right\rangle$ is a causal net, which generalizes the behavior of $\mathcal{C}=\langle$ fold $(\overline{\mathcal{G}}), \mathcal{I}, \mathcal{O}\rangle$ over the nodes and the edges in $\mathcal{G}$ that do not occur in fold $(\overline{\mathcal{G}})$.

Since $\mathcal{C} \vdash L$, we conclude that $\left\langle\mathcal{G}, \mathcal{I}^{\prime}, \mathcal{O}^{\prime}\right\rangle \vdash L$. Indeed, if $\sigma$ is a sequence valid w.r.t. $\mathcal{C}$, then it is also valid w.r.t. $\left\langle\mathcal{G}, \mathcal{I}^{\prime}, \mathcal{O}^{\prime}\right\rangle$.

Example 8.12. Consider the causal net $\mathcal{C}_{3}=\left\langle\mathcal{G}_{3}, \mathcal{I}_{3}, \mathcal{O}_{3}\right\rangle$, where $\mathcal{G}_{3}$ is the dependency graph shown in Figure 8.2, and where $\mathcal{I}_{3}$ and $\mathcal{O}_{3}$ are the functions such that: $\mathcal{I}_{3}(a)=$ $\{\emptyset\}, \mathcal{I}_{3}(b)=\{\{(a, b)\}\}, \mathcal{I}_{3}(c)=\{\{(b, c)\},\{d, c\}\}, \mathcal{I}_{3}(d)=\{\{(c, d)\}\}, \mathcal{I}_{3}(e)=$ $\{\{(b, e)\},\{(c, e)\},\{(b, e),(c, e)\}\}, \mathcal{O}_{3}(a)=\left\{\{(a, b),(a, c)\},\{(a, b\},\{(a, c)\}\}, \mathcal{O}_{3}(b)=\right.$ $\{\{(b, d)\}\}, \mathcal{O}_{3}(c)=\{\{(c, d)\}\}, \mathcal{O}_{3}(d)=\{\{(d, e)\}\}$, and $\mathcal{O}_{3}(e)=\{\emptyset\}$. Note that $\mathcal{C}_{3} \vdash\{$ acdcde $\}$ holds. Hence, by the above result, we can conclude that $\mathcal{G}_{3} \vdash\{$ acdcde $\}$ also holds, as we have in fact already observed in Example 8.10.

Note that if conditions (1) and (2) in Definition 8.9 are not guaranteed to hold, then a weaker variant of Theorem 8.11 can be still established.

Theorem 8.13. Let $L$ be a log, and let $\mathcal{G}$ be a dependency graph such that for each trace $t \in L$, there is a graph $\overline{\mathcal{G}}_{t}$ such that fold $\left(\overline{\mathcal{G}}_{t}\right)$ is a subgraph of $\mathcal{G}$ and $\overline{\mathcal{G}}_{t} \vdash_{a}\{\operatorname{unfold}(t)\}$. Then, there is a possibly extended causal net $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ such that $\mathcal{C} \vdash L$.

Proof. In the case where conditions (1) and (2) in Definition 8.9 are not guaranteed to hold, by inspecting the proof of the only-if part of Theorem 8.11, it can be checked that the structure $\mathcal{C}$ built there over $\mathcal{G}$ is still such that $\mathcal{C} \vdash L$ holds. The only difference is that $\mathcal{C}$ might be an extended causal net.

## CHAPTER 9

## Precedence Constraints: Mining Problems and Complexity

### 9.1 Introduction

In this chapter, we propose and analyze a framework to specify additional properties on the process models that can be produced as output by process discovery algorithms. In particular, we formalize the concept of precedence constraints, discuss their application to the problem of mining causal nets, and analyze their complexity.

### 9.2 Syntax and Semantics

A precedence constraint is an assertion aimed at expressing a relationship of precedence among some of the activities in the underlying set $\mathcal{A}$. The language of precedence constraints is next defined in order to smoothly allow the formalization of the kinds of prior knowledge that are usually available to the analyst, such as, parallelism, locality, or exclusivity of activities (cf. Goedertier et al. [2009]).

Definition 9.1. A positive precedence constraint $\pi$ over $\mathcal{A}$ is either

- an expression of the form $S \rightarrow T$, called edge constraint, or
- an expression of the form $S \rightsquigarrow T$, called path constraint,
where $S, T \subseteq \mathcal{A}$, with $|S| \geq 1$ and $|T| \geq 1$, are non-empty sets of activities.
For a positive constraint $\pi, \neg(\pi)$ is a negative precedence constraint.

Precedence constraints are interpreted over directed graphs as follows.
Definition 9.2. Let $\mathcal{G}=(V, E)$ be a directed graph such that $V \subseteq \mathcal{A}$. Then,
(1) $\mathcal{G}$ satisfies an edge constraint $S \rightarrow T$, if there is an edge $(x, y) \in E$ with $x \in S$ and $y \in T$;
(2) $\mathcal{G}$ satisfies a path constraint $S \rightsquigarrow T$, if there is a sequence $x=a_{0}, a_{1}, \ldots, a_{n}=y$, with $n>0$, such that $x \in S, y \in T$ and $\left(a_{i}, a_{i+1}\right) \in E$, for each $i \in\{0, \ldots, n-1\}$;
(3) $\mathcal{G}$ satisfies $\neg(\pi)$, if $\mathcal{G}$ does not satisfy $\pi$.

If $\mathcal{G}$ satisfies each constraint in a set $\Pi$ of precedence constraints, then $\mathcal{G}$ is a model of $\Pi$, denoted by $\mathcal{G} \models \Pi$. The set of all activities occurring in the constraints in $\Pi$ is hereinafter denoted by $\mathcal{A}(\Pi)$.

A foundational task in process mining consists in automatically building a model that can explain the behavior registered in all the traces of some $\log L$ given as input. In this context, precedence constraints can formalize additional requirements that the model discovered from $L$ has to satisfy. This gives rise to the following "Causal Net Mining" (short: CN-Mining) problem.

CN-Mining: Given a set $\mathcal{A}$ of activities, a $\log L$ with $\mathcal{A}(L) \subseteq \mathcal{A}$, and a set $\Pi$ of precedence constraints with $\mathcal{A}(\Pi) \subseteq \mathcal{A}$, compute a possibly extended causal net $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ over $\mathcal{A}$ such that $\mathcal{C} \vdash L$ and $\mathcal{G} \models \Pi$, or check that no net with these properties exists.

Note that, for $\Pi=\emptyset$, the problem above reduces to the standard one considered in the literature. Moreover, observe that we are not considering (for the moment) quality measures for the resulting net. This issue will be explored later on. Finally, note that extended causal nets are allowed as solutions to the problem. In fact,
throughout the chapter, we shall explicitly discuss and show how our results extend to the more stringent setting where the focus is on (standard) causal nets.

Example 9.3. Consider the set $\Pi=\{\neg(\{b\} \rightsquigarrow\{d\}), \neg(\{d\} \rightsquigarrow\{b\})\}$ of precedence constraints. We have two negative path constraints, stating that $b$ and d must be executed over "parallel" branches of the given process.

Consider then the traces abcde and acbde within the setting of Example 8.1, plus the causal nets $\mathcal{C}_{1}$ and $\mathcal{C}_{2}$ discussed in Example 8.3 and depicted in Figure 8.1. Without additional constraints, we have already noticed that $\mathcal{C}_{1}$ and $\mathcal{C}_{2}$ are such that $\mathcal{C}_{1} \vdash\{a b c d e, a c b d e\}$ and $\mathcal{C}_{2} \vdash\{a b c d e, a c b d e\}$. However, the dependency graph $\mathcal{G}_{2}$ (associated with $\mathcal{C}_{2}$ ) is a model of $\Pi$, while $\mathcal{G}_{1}$ is not, as it violates the constraint $\neg(\{b\} \rightsquigarrow\{d\})$. Thus, $\mathcal{C}_{2}$ is a solution to CN-Mining on input $\{a b c d e$, acbde $\}$ and $\Pi$.

As a further remark, note that a solution to the mining problem might involve activities that do not occur in the $\log$ and in the constraints, i.e., in $\mathcal{A} \backslash \mathcal{A}(L) \backslash$ $\mathcal{A}(\Pi)$. As an example, consider the set $\{\neg(\{a\} \rightarrow\{b\}),\{a\} \rightsquigarrow\{b\}\}$ of constraints prescribing the existence of a path from $a$ to $b$, but forbidding the existence of a direct connection. If the log provided as input is empty, then any solution has to be defined over some "fresh" activities (at least one) ensuring the existence of a path from $a$ to $b$ of this kind. However, according to the formulation of CN-Mining, any "fresh" activity (in $\mathcal{A} \backslash \mathcal{A}(L) \backslash \mathcal{A}(\Pi)$ ) has to be explicitly provided as input to the mining problem, too. Therefore, one might wonder whether scenarios exist where all solutions need exponentially many fresh activities (w.r.t. $|\mathcal{A}(L) \cup \mathcal{A}(\Pi)|$ ), so that explicitly listing all of them would artificially blow up the size of the input. We next show that this is not possible, since a "small" solution always exists if the problem admits any solution. Hence, there is no loss of generality, but rather we gain
flexibility, by taking as input the set $\mathcal{A}$ possibly implicitly, i.e., by just specifying the upper bound on the number of activities that is defined in the statement of the result below.

Theorem 9.4. Let $L$ be a log, let $\Pi$ be a set of precedence constraints, and let $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ be a causal net (resp., an extended causal net) with $\mathcal{C} \vdash L$ and $\mathcal{G} \models \Pi$. Then, a causal net (resp., an extended causal net) $\mathcal{C}^{\prime}=\left\langle\mathcal{G}^{\prime}, \mathcal{I}^{\prime}, \mathcal{O}^{\prime}\right\rangle$ exists such that $\mathcal{C}^{\prime} \vdash L, \mathcal{G}^{\prime} \models \Pi$, and $\left|V^{\prime}\right| \leq|\mathcal{A}(L) \cup \mathcal{A}(\Pi)|^{2}+|\mathcal{A}(L) \cup \mathcal{A}(\Pi)|$, with $V^{\prime}$ being the set of nodes of $\mathcal{G}^{\prime}$. Moreover, if $\mathcal{G}$ is acyclic, then $\mathcal{G}^{\prime}$ is acyclic, too.


Figure 9.1: Example construction in the proof of Theorem 9.4.

Proof. Let $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$, with $\mathcal{G}=(V, E)$, be a (resp., an extended) causal net such that $\mathcal{C} \vdash L$ and $\mathcal{G} \models \Pi$. Consider the graph $\mathcal{G}^{\prime}=\left(V^{\prime}, E_{1}^{\prime} \cup E_{2}^{\prime}\right)$ built as follows. The set $V^{\prime}$ consists of all the activities in $\mathcal{A}(L) \cup \mathcal{A}(\Pi)$ plus a fresh activity $a_{x, y}$ for each pair of activities $x, y \in \mathcal{A}(L) \cup \mathcal{A}(\Pi)$ such that there is a path in $\mathcal{G}$ from $x$ to $y$. The set $E_{1}^{\prime}$ consists of all the edges in $E$ defined over the activities in $\mathcal{A}(L) \cup \mathcal{A}(\Pi)$, i.e., $E_{1}^{\prime}=\{(x, y) \in E \mid\{x, y\} \subseteq \mathcal{A}(L) \cup \mathcal{A}(\Pi)\}$. The set $E_{2}^{\prime}$ contains the edges $\left(x, a_{x, y}\right)$ and $\left(a_{x, y}, y\right)$, for each fresh activity $a_{x, y} \in V^{\prime}$ (hence, $\left.a_{x, y} \notin V\right)$, and no further edge is in $E_{2}^{\prime}$.

As an example, consider the graph $\mathcal{G}$ reported on the left of Figure 9.1. The graph is defined over the activities in $\mathcal{A}(L) \cup \mathcal{A}(\Pi)=\left\{a_{\perp}, a_{\top}, x, y\right\}$ plus 6 additional activities, which are depicted as circles. The graph $\mathcal{G}^{\prime}$ that is built based on $\mathcal{G}$ is illustrated on the right part of the same figure. All nodes in $\mathcal{G}^{\prime}$ that do not occur in $\mathcal{G}$ are depicted as black circles. In particular, observe that the node $a_{x, y}$ is responsible for preserving the connectivity that is supported in $\mathcal{G}$ by the nodes not occurring in $\mathcal{A}(L) \cup \mathcal{A}(\Pi)$.

Let us now analyze the properties of $\mathcal{G}^{\prime}$. First, note that the activities $a_{\perp}$ and $a_{\top}$ are in $V^{\prime}$, as in fact they occur in $\mathcal{A}(L)$. Moreover, since $\mathcal{G}$ is a dependency graph, no edge ingoing into $a_{\perp}$ (resp., outgoing from $a_{\top}$ ) occurs in $E_{2}^{\prime}$. Hence, given the construction of the edges in $E_{1}^{\prime}$, we conclude that $a_{\perp}$ and $a_{\top}$ have no ingoing and outgoing edges in $\mathcal{G}^{\prime}$.

Now, we claim that for each pair of activities $x, y \in \mathcal{A}(L) \cup \mathcal{A}(\Pi)$, there is a path from $x$ to $y$ in $\mathcal{G}$ if, and only if, there is a path from $x$ to $y$ in $\mathcal{G}^{\prime}$. Indeed, if there is a path from $x$ to $y$ in $\mathcal{G}$, then the edges $\left(x, a_{x, y}\right)$ and $\left(a_{x, y}, y\right)$ occur in $E_{2}^{\prime}$. Conversely, assume that there is a path $\pi$ from $x$ to $y$ in $\mathcal{G}^{\prime}$, and for the sake of contradiction that there is no path from $x$ to $y$ in $\mathcal{G}$. As all edges of $E$ defined over the activities in $\mathcal{A}(L) \cup \mathcal{A}(\Pi)$ are in $E_{1}^{\prime}$, it must be the case that two edges occur in $E_{2}^{\prime}$ having the form $\left(\bar{x}, a_{\bar{x}, \bar{y}}\right)$ and $\left(a_{\bar{x}, \bar{y}}, \bar{y}\right)$ and such that there is no path from $\bar{x}$ to $\bar{y}$ in $\mathcal{G}$. However, this is impossible given the construction of the edges in $E_{2}^{\prime}$.

In the light of the above property, it follows that if $\mathcal{G}$ is acyclic, then $\mathcal{G}^{\prime}$ is acyclic, too. Indeed, just notice that any cycle in $\mathcal{G}^{\prime}$ must necessarily include a node in $\mathcal{A}(L) \cup \mathcal{A}(\Pi)$. Moreover, we can conclude that each activity $a \in \mathcal{A}(L) \cup \mathcal{A}(\Pi) \backslash$ $\left\{a_{\perp}, a_{\top}\right\}$ is in a path in $\mathcal{G}^{\prime}$ from $a_{\perp}$ to $a_{\top}$. Consider then an activity of the form $a_{x, y}$, which occurs in $V^{\prime} \backslash V$. Since $x$ (resp., $y$ ) either coincides with $a_{\perp}$ (resp., $a_{\top}$ )
or it is reachable from $a_{\perp}$ ( can reach $a_{\top}$ ) in $\mathcal{G}^{\prime}$, because this property holds in fact on $\mathcal{G}$, we also conclude that $a_{x, y}$ is in a path in $\mathcal{G}^{\prime}$ from $a_{\perp}$ to $a_{\top}$. By putting the above observations together, it follows that $\mathcal{G}^{\prime}$ is a dependency graph over the set $V^{\prime}$ of activities. Moreover, $\left|V^{\prime}\right| \leq|\mathcal{A}(L) \cup \mathcal{A}(\Pi)|^{2}+|\mathcal{A}(L) \cup \mathcal{A}(\Pi)|$.

Recall now that $\mathcal{G}^{\prime}$ preserves all the edges defined over the activities in $\mathcal{A}(L)$, and that $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ is a causal net (resp., extended causal net) such that $\mathcal{C} \vdash L$. It follows that there is a (resp., an extended) causal net $\mathcal{C}^{\prime}=\left\langle\mathcal{G}^{\prime}, \mathcal{I}^{\prime}, \mathcal{O}^{\prime}\right\rangle$ such that $\mathcal{C}^{\prime} \vdash L$, where $\mathcal{I}^{\prime}$ and $\mathcal{O}^{\prime}$ just extend $\mathcal{I}$ and $\mathcal{O}$ as to include, for each given activity, a binding defined over the whole set of its ingoing and outgoing edges, respectively. Note that $\mathcal{C}^{\prime}$ supports all the traces in $L$, even without using such fresh bindings. In fact, the construction of $\mathcal{I}^{\prime}$ and $\mathcal{O}^{\prime}$ is just required to ensure that $\mathcal{C}^{\prime}$ is formally a (possibly extended) causal net.

In order to conclude, we have then to show that $\mathcal{G}^{\prime} \models \Pi$. To this end, note that edge constraints and negated edge constraints are satisfied by $\mathcal{G}^{\prime}$, because they are satisfied by $\mathcal{G}$ and since the two graphs coincide over the activities in $\mathcal{A}(\Pi)$. Eventually, recall that, for each pair of activities $x, y \in \mathcal{A}(L) \cup \mathcal{A}(\Pi)$, there is a path from $x$ to $y$ in $\mathcal{G}$ if, and only if, there is a path from $x$ to $y$ in $\mathcal{G}^{\prime}$. Hence, also path constraints and negated path constraints are satisfied by $\mathcal{G}^{\prime}$, because they are satisfied by $\mathcal{G}$. That is, $\mathcal{G}^{\prime} \models \Pi$.

### 9.3 Precedence Constraints

From a conceptual viewpoint, the problem defined above comprises a mining task, i.e., mining a process model supporting a given $\log$, and a reasoning task, i.e., to check whether the model additionally satisfies some precedence constraints. We next show that even the mining task can be declaratively formulated in terms of reasoning

```
Function computeBindings (\mathcal{G},L), with \mathcal{G }=(\mathcal{A},E);
\foralla\in\mathcal{A},\mathcal{I}(a):= Ia}\mathrm{ and }\mathcal{O}(a):=\mp@subsup{O}{a}{}\mathrm{ , where
    Ia}={(x,a)|(x,a)\inE}\mathrm{ and }\mp@subsup{O}{a}{}={(a,y)|(a,y)\inE}
for each trace t in L, and for each }i\in{1,\ldots,len(t)} d
    \mathcal { I } ( t [ i ] ) : = \mathcal { I } ( t [ i ] ) \cup \{ ( t [ j ] , t [ i ] ) \in E \| j < i \} ; ( * t r e a t e d ~ a s ~ m u l t i - s e t s * )
    \mathcal { O } ( t [ i ] ) : = \mathcal { O } ( t [ i ] ) \cup \{ ( t [ i ] , t [ j ] ) \in E \| i < j \} ; ( * t r e a t e d ~ a s ~ m u l t i - s e t s * )
end for
return }\langle\mathcal{G},\mathcal{I},\mathcal{O}\rangle
```

Figure 9.2: Computation of a possibly extended causal net in Theorem 9.6.
about precedence constraints. In other words, the language of precedence constraints is enough expressive as to capture the semantics of causal nets.

Definition 9.5. Let $L$ be a log. For each trace $t \in L$, the set of precedence constraints induced by $t$ is defined as follows:

$$
\begin{aligned}
\pi(t)= & \{\{t[1], \ldots, t[i-1]\} \rightarrow\{t[i]\} \mid 1<i \leq \operatorname{len}(t)\} \cup \\
& \{\{t[i]\} \rightarrow\{t[i+1], \ldots, t[\operatorname{len}(t)]\} \mid 1 \leq i<\operatorname{len}(t)\} .
\end{aligned}
$$

The set of precedence constraints induced by $L$ is defined as $\pi(L)=\bigcup_{t \in L} \pi(t)$.
Intuitively, we have stated that each activity in $t$ can be directly reached by at least one of its predecessors in $t$, and it can directly reach at least one of its successors. Based on this definition, we establish the following crucial result.

Theorem 9.6. Let $\mathcal{G}$ be a dependency graph over a set $\mathcal{A}$ of activities, let $L$ be a log with $\mathcal{A}(L) \subseteq \mathcal{A}$, and let $\Pi$ be a set of precedence constraints with $\mathcal{A}(\Pi) \subseteq \mathcal{A}$. The following statements are equivalent:
(1) The graph $\mathcal{G}$ is a model of $\Pi \cup \pi(L)$.
(2) There is a possibly extended causal net $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ that is a solution to CN Mining on input $\mathcal{A}, L$, and $\Pi$.

Moreover, input and output bindings witnessing that the" $(1) \Rightarrow(2)$ "-part holds can be built via the function reported in Figure 9.2. In particular, if there is no trace in
$L$ containing repetitions of the same activity (for short, we hereinafter say that $L$ is linear), then the function actually builds a (standard) causal net.

Since the proof is rather involved, as a base case we show that the result holds in absence of user-defined constraints and when the focus is on linear logs only. Note that in Chapter 8 we have introduced a number of technical ingredients which will use in this proof.

Theorem 9.7. Let $L$ be a linear log and let $\mathcal{G}$ be a dependency graph. Then, $\mathcal{G} \models$ $\pi(L)$ if, and only if, $\mathcal{G} \vdash_{a} L$.

Proof. (if part). Assume that $\mathcal{G} \vdash_{a} L$, i.e., for each $t \in L$, there is a subgraph $\mathcal{G}_{t}=$ $\left(V_{t}, E_{t}\right)$ of $\mathcal{G}=(V, E)$ such that $\mathcal{G}_{t}$ is an acyclic dependency graph, and $t$ is a topologic sort of $\mathcal{G}_{t}$. Therefore, for each $i \in\{2, \ldots$, len $(t)\}$ (resp., $i \in\{1, \ldots$, len $(t)-1\}$ ), there is a path from $t[1]$ (resp., $t[i])$ to $t[i]$ (resp., $t[\operatorname{len}(t)])$ in $\mathcal{G}_{t}$, by definition of dependency graph. In particular, since $t$ is a topologic sort of $\mathcal{G}_{t}$, we are guaranteed about the existence of an edge in $E_{t}$ (and then in $E$ ) having the form ( $t[j], t[i]$ ) (resp., $\left.\left(t[i], t\left[j^{\prime}\right]\right)\right)$ and such that $j<i$ (resp., $\left.i<j^{\prime}\right)$ holds. Hence, the set $\pi(t)$ of the precedence constraints induced by $t$ are satisfied by $\mathcal{G}$, for each trace $t$ in $L$. That is, $\mathcal{G} \models \pi(L)$.
(only-if part). Assume that $\mathcal{G} \models \pi(L)$, with $\mathcal{G}=(V, E)$. Let $t$ be a trace in $L$, and let $\mathcal{G}_{t}=\left(V_{t}, E_{t}\right)$ be the graph such that $V_{t}=\{t[1], \ldots, t[\operatorname{len}(t)]\}$ and $E_{t}=$ $\{(t[i], t[j]) \in E \mid 1 \leq i<j \leq l e n(t)\}$. Since $L$ is linear, we can note that $\mathcal{G}_{t}$ is acyclic, that $t[1]=a_{\perp}$ has no ingoing edges, and that $t[$ len $(t)]=a_{\top}$ has no outgoing edges. We now claim that each activity $t[i] \in V_{t} \backslash\{t[1]\}$ can be reached from $t[1]$. The above property can be shown by induction on the index $i>1$. In the case where $i=2$, $(t[1], t[2])$ must belong to $E$ (and hence to $E_{t}$ ), in order to satisfy the constraint
$\{t[1]\} \rightarrow\{t[2]\}$ in $\pi(t)$. Assume now that the activities in the set $\{t[2], \ldots, t[i-1]\}$ can be reached from $t[1]$. Then, because of the constraint $\{t[1], \ldots, t[i-1]\} \rightarrow\{t[i]\}$ in $\pi(t)$, we again have that $t[i]$ can be reached from $t[1]$. Similarly, it can be checked that the terminating activity $t[\operatorname{len}(t)]$ can be reached by each activity $t[i] \in V_{t} \backslash\{t[\operatorname{len}(t)]\}$, by using this time the fact that $\{t[i]\} \rightarrow\{t[i+1], \ldots, t[\operatorname{len}(t)]\}$ is in $\pi(t)$. Hence, $\mathcal{G}_{t}$ is an acyclic dependency graph. Moreover, for each edge $(t[i], t[j])$ in $E_{t}$, we have that $i<j$ holds by construction. Thus, $t$ is a topologic sort of $\mathcal{G}_{t}$. As $\mathcal{G}_{t}$ is a subgraph of $\mathcal{G}$, we then have $\mathcal{G} \vdash_{a} L$.

Theorem 9.7 and Theorem 8.8 imply the following corollary, where process mining over linear logs is restated in terms of reasoning about precedence constraints.

Corollary 9.8. Let $\mathcal{G}$ be a dependency graph over a set $\mathcal{A}$ of activities, let $L$ be a linear $\log$ with $\mathcal{A}(L) \subseteq \mathcal{A}$, and let $\Pi$ be a set of precedence constraints with $\mathcal{A}(\Pi) \subseteq \mathcal{A}$. Then, the followings are equivalent: (1) The graph $\mathcal{G}$ is a model of $\Pi \cup \pi(L)$; (2) There is a causal net $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ that is a solution to CN-Mining on input $\mathcal{A}, L$, and $\Pi$.

Example 9.9. Let $\Pi$ be the set of constraints in Example 9.3, and consider the novel set $\Pi^{\prime}=\Pi \cup \pi(\{a b c d e, a c b d e\})$. It can be checked that the dependency graph $\mathcal{G}_{2}$ in Figure 8.2 is a model of $\Pi^{\prime}$. Thus, by Corollary 9.8, we are guaranteed about the existence of a causal net that can be defined on top of $\mathcal{G}_{2}$ and that is a solution to CN-Mining on input \{abcde, acbde\} and $\Pi$. In fact, we already know that the causal net $\mathcal{C}_{2}$ defined in Example 8.3 is a solution. $\triangleleft$

Note that, the above result completely characterizes the cases where the problem CN-Mining admit solutions. Indeed, point (2) above can be equivalently restated as the existence of a possibly extended causal net that is a solution to the mining
problem. This is because of the following fact, which trivially holds since the ability of extended nets to activate multiple instances of the same activity is useless over linear logs.

Fact 9.10. Over linear logs, CN-Mining admits a solution if, and only if, it admits a causal net as a solution.

Moreover, it is useful to remark that the ' $(1) \Rightarrow(2)$ '-part of Corollary 9.8 can be stated constructively. That is, the causal net $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ can be efficiently built given the graph $\mathcal{G}$. The method is just based on inspecting the proofs for the results in Section 8.4, and in fact it coincides with the one algorithmically reported in Figure 9.2.

At this point, the natural question is whether we can extend Theorem 9.7 to deal with arbitrary logs. An answer, which is however only partially positive, is stated below.

Theorem 9.11. Let $L$ be a log and let $\mathcal{G}$ be a dependency graph. Then, $\mathcal{G} \vdash L$ implies that $\mathcal{G} \models \pi(L)$.

Proof. Assume that $\mathcal{G} \vdash L$ holds, with $\mathcal{G}=(V, E)$. By Definition 8.9, for each trace $t \in L$, there is a graph $\overline{\mathcal{G}}_{t}=\left(\bar{V}_{t}, \bar{E}_{t}\right)$ such that, in particular, fold $\left(\overline{\mathcal{G}}_{t}\right)$ is a subgraph of $\mathcal{G}=\left(V_{t}, E_{t}\right)$ and $\overline{\mathcal{G}}_{t} \vdash_{a}\{\operatorname{unfold}(t)\}$. Then, we apply Theorem 9.7 on $\overline{\mathcal{G}}_{t}$, and we conclude that $\overline{\mathcal{G}}_{t} \models \pi(\operatorname{unfold}(t))$. Moreover, we note that fold $\left(\overline{\mathcal{G}}_{t}\right) \models \pi(t)$ also holds. Indeed, for each $i \in\{1, \ldots$, len $(t)\}$, whenever (unfold $(t)[j]$, unfold $(t)[i])$ (resp., $(\operatorname{unfold}(t)[i], \operatorname{unfold}(t)[h]))$ is in $\bar{E}_{t}$, then $(t[j], t[i])$ (resp., $\left.(t[i], t[h])\right)$ is an edge of fold $\left(\overline{\mathcal{G}}_{t}\right)$ with $j<i$ (resp., $h>i$ ). Finally, define $\overline{\mathcal{G}}=\left(\bigcup_{t \in L} \bar{V}_{t}, \bigcup_{t \in L} \bar{E}_{t}\right)$. Then, we claim that fold $(\overline{\mathcal{G}}) \models \pi(L)$. Indeed, the constraints induced by $L$ are positive ones, so that, since $\operatorname{fold}\left(\overline{\mathcal{G}}_{t}\right) \models \pi(t)$, the graph fold $(\overline{\mathcal{G}})$ (of which fold $\left(\overline{\mathcal{G}}_{t}\right)$ is a subgraph) is
still such that $\operatorname{fold}(\overline{\mathcal{G}}) \models \pi(t)$. To conclude, we can eventually observe that fold $(\overline{\mathcal{G}})$ is a subgraph of $\mathcal{G}$, and hence $\mathcal{G} \models \pi(L)$.

This is the best one can hope to do. Indeed, we can see that $\mathcal{G} \models \pi(L)$ does not imply $\mathcal{G} \vdash L$, by just looking again at Example 8.10. There, we have noticed that the graph $\mathcal{G}_{3}$ does not support $\{a b b e\}$. However, it can be checked that $\mathcal{G}_{3} \models \pi(\{a b b e\})$ holds.

Despite the above limitation, the counterpart of Corollary 9.8 for arbitrary logs can still be obtained by moving to possibly extended causal nets-as with Corollary 9.8 , the ' $(1) \Rightarrow(2)^{\prime}$ '-part below is formalized algorithmically in Figure 9.2.

Theorem 9.12. Let $\mathcal{G}$ be a dependency graph over a set $\mathcal{A}$ of activities, let $L$ be a $\log$ with $\mathcal{A}(L) \subseteq \mathcal{A}$, and let $\Pi$ be a set of precedence constraints with $\mathcal{A}(\Pi) \subseteq \mathcal{A}$. The followings are equivalent: (1) The graph $\mathcal{G}$ is a model of $\Pi \cup \pi(L)$; (2) There is a possibly extended causal net $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ that is a solution to CN-Mining on input $\mathcal{A}, L$, and $\Pi$.

Proof. (1) $\Rightarrow(2)$. Assume that $\mathcal{G}$ is a model of $\Pi \cup \pi(L)$, with $\mathcal{G}=(V, E)$. Let $\overline{\mathcal{G}}=(\bar{V}, \bar{E})$ be the graph such that $\bar{V}=\mathcal{A}(\operatorname{unfold}(L))$ and $\bar{E}=\left\{\left(x^{(i\rangle}, y^{\langle j\rangle}\right) \mid(x, y) \in\right.$ $\left.E, x^{\langle i\rangle} \in \bar{V}, y^{\langle j\rangle} \in \bar{V}\right\}$. Since $\mathcal{G}$ is a dependency graph, it is also the case that $\overline{\mathcal{G}}$ is a dependency graph (over $\mathcal{A}(\operatorname{unfold}(L)))$. Moreover, $\overline{\mathcal{G}} \models \pi(\operatorname{unfold}(L))$ holds. Indeed, just note that for each trace $t \in L$ and for each $i \in\{1, \ldots, \operatorname{len}(t)\}$, if $(t[j], t[i])$ (resp., $(t[i], t[h]))$ is in $E$ with $j<i$ (resp., $h>i)$, then $(\operatorname{unfold}(t)[j]$, unfold $(t)[i])$ (resp., $(\operatorname{unfold}(t)[i], \operatorname{unfold}(t)[h]))$ is in $\bar{E}$ by construction. Thus, we can apply Theorem 9.7 in order to conclude that $\overline{\mathcal{G}} \vdash_{a} \operatorname{unfold}(L)$. Eventually, we observe that $\operatorname{fold}(\overline{\mathcal{G}})$ is clearly a subgraph of $\mathcal{G}$. Then, we distinguish two cases. In the case where $\overline{\mathcal{G}}$ satisfies conditions (1) and (2) in Definition 8.9, then we can apply Theorem 8.11
and conclude that a causal net $\mathcal{C}$ can be built over $\mathcal{G}$ such that $\mathcal{C} \vdash L$. Instead, in the case where one of the above conditions does not hold, we can apply Theorem 8.13 and conclude that there is a possibly extended causal net $\mathcal{C}$ built over $\mathcal{G}$ and such that $\mathcal{C} \vdash L$.
$(2) \Rightarrow(1)$. Assume that $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ is a solution to CN-Mining on input $\mathcal{A}, L$, and $\Pi$, and that, for the sake of contradiction, $\mathcal{G}$ is not a model of $\Pi \cup \pi(L)$. The fact that $\mathcal{G}$ is not a model of $\Pi$ is trivially impossible. Thus, assume that $\mathcal{G}$ is not a model of $\pi(L)$, i.e., there is a trace $t \in L$ such that $\mathcal{G}$ does not satisfy the constraints in $\pi(t)$. According to Definition 9.5, there are two possible cases: (a) there is an index $i \in\{2, \ldots$, len $(t)\}$ such that there is no edge in $\mathcal{G}$ from one of the activities in $\{t[1], \ldots, t[i-1]\}$ to $t[i] ;(b)$ there is an index $i \in\{1, \ldots$, len $(t)-1\}$ such that there is no edge in $\mathcal{G}$ from $t[i]$ to one of the activities in $\{t[i+1], \ldots, t[\operatorname{len}(t)]\}$.

Now, as $\mathcal{C}$ is a solution, it must be the case that $\mathcal{C}$ supports the trace $t$, i.e., there is a binding sequence $\sigma$ that is valid w.r.t. $\mathcal{C}$ and where the $j$-th binding activity $\left\langle a_{j}, i b_{j}, o b_{j}\right\rangle$ of $\sigma$, for each $j \in\{1, \ldots, \operatorname{len}(t)\}$, is such that $a_{j}=t[j]$. In particular, the $i$-th activity is $\left\langle t[i], i b_{i}, o b_{i}\right\rangle$. In the case (a) above, the set $i b_{i}$, which is a non-empty set of edges ingoing into $t[i]$, is a subset of the set $\{(x, t[i]) \mid x \notin\{t[1], \ldots, t[i-1]\}\}$, while the state $S_{i-1}^{\sigma}$ of the causal net consists of edges having the form $\left(x^{\prime}, y\right)$, with $x^{\prime} \in\{t[1], \ldots, t[i-1]\}$. It follows that $i b_{i}$ is not contained in $S_{i-1}^{\sigma}$, and $\sigma$ is not valid, which is impossible as $\sigma$ is valid. In the case (b), the set $o b_{i}$, which is a non-empty set of edges outgoing from $t[i]$, is included in the set $\{(t[i], y) \mid y \notin$ $\{t[1+1], \ldots, t[l e n(t)]\}\}$. Thus, $S_{i}^{\sigma}$ necessarily includes an element having the form $(t[i], y)$, with $y \notin\{t[1+1], \ldots, t[\operatorname{len}(t)]\}$. However, as there is no edge in $\mathcal{G}$ from $t[i]$ to one of the activities in $\{t[i+1], \ldots, t[\operatorname{len}(t)]\}$, this element will occur at the last step of the execution of $\sigma$, i.e., $S_{l e n(t)}^{\sigma} \neq \emptyset$, which again is impossible.

By combining the above result with Corollary 9.8, we have established the result in Theorem 9.6. In particular, the constructive " $(1) \Rightarrow(2)$ "-part is seen to hold by inspecting the construction in the proofs and by observing that input and output bindings are constructed there precisely according to the function reported in Figure 9.2.

As an example application, since $\mathcal{G}_{3} \models \pi(\{a b b e\})$ holds, we are guaranteed that a net supporting this trace can be built on top of $\mathcal{G}_{3}$. To be concrete, we can consider the extended causal net $\mathcal{C}_{3}^{\prime}=\left\langle\mathcal{G}_{3}, \mathcal{I}_{3}^{\prime}, \mathcal{O}_{3}^{\prime}\right\rangle$ where $\mathcal{I}_{3}^{\prime}(a)=\{\emptyset\}, \mathcal{I}_{3}^{\prime}(b)=\{\{(a, b)\}\}$, $\mathcal{I}_{3}^{\prime}(e)=\{\{(b, e),(b, e)\}\}, \mathcal{O}_{3}^{\prime}(a)=\{\{(a, b),(a, b)\}\}, \mathcal{O}_{3}^{\prime}(b)=\{\{(b, e)\}\}$, and $\mathcal{O}_{3}^{\prime}(e)=$ $\{\emptyset\}$-the specification of the bindings over the activities $c$ and $d$ is irrelevant here.

Note that Theorem 9.6 reformulates the mining problem and the dynamics of causal nets in purely "static" terms, i.e., in terms of reasoning about the satisfaction of precedence constraints with no reference to the concepts of bindings and of trace/log support. This will be very useful for designing our algorithms and for conducting the complexity analysis. Moreover, the result has an immediate concrete application. Indeed, the reasoning problem can be easily encoded via a "standard" constraints satisfaction problem, CSP for short (see, e.g., Dechter [1992]), so that existing constraint programming platforms can be reused to compute solutions, in the spirit of the works by [De Raedt et al., 2008] and [Nijssen et al., 2009]. A preliminary version of this approach was discussed by [Greco et al., 2012], however it was not designed to deal with causal nets (i.e., with full process models), but its focus was on the discovery of the underlying dependence graphs only.

An advantage of the CSP-based framework is that it can transparently handle arbitrary sets of precedence constraints. The price to be paid, however, is that computing a solution might be very challenging from a computational viewpoint, as we
shall formally illustrate in Section 9.4. In this dissertation, we decided therefore to propose solution algorithms that are specific for some classes of constraints (over which solutions can be efficiently computed), as well as methods that handle the general setting heuristically. In fact, from our experimentation, the heuristics methods emerged to be definitively much faster (orders of magnitude) than the CSP-based method, while being capable to end up with an exact solution in most cases. Accordingly, for those cases where a heuristic solution is acceptable (i.e., where some constraints might be violated), our earlier approach has to be considered as pragmatically superseded by the methods proposed here. However, it is an interesting avenue of further research to define different encodings (possibly, still CSP-based ones).

### 9.4 Complexity Analysis

We now complete the picture by studying the computational complexity of the problem CN-Mining, which is an important step towards developing effective algorithms for its solution. In the analysis that follows, we take into account various qualitative properties regarding the kinds of constraint being allowed, by tracing the tractability frontier w.r.t. them. Formally, let S be a subset of the following set of symbols $\{\rightarrow, \rightsquigarrow, \nrightarrow, \not \nsim\}$. Let $\mathcal{C}[\mathrm{S}]$ denote all the possible constraints that can be built in a way that if $\rightarrow \notin \mathrm{S}$ (resp., $\rightsquigarrow \notin \mathrm{S}, \nrightarrow \notin \mathrm{S}, \nLeftarrow \notin \mathrm{S}$ ), then no edge (resp., path, negated edge, negated path) constraint is in $\mathcal{C}[S])$. Let $\mathrm{CN}-\mathrm{Mining}^{2}[\mathrm{~S}]$ denote the restriction of the problem over any set $\Pi$ of precedence constraints such that $\Pi \subseteq \mathcal{C}[\mathrm{S}]$. And, finally, let CN-Existence[ $[$ ] be the decision version of this problem, where we have just to decide whether a solution exist at all, and we are not asked to compute a solution, if any. Then, our results can be summarized as it is stated next (see also Figure 9.3).


Figure 9.3:
Tractability frontiers. A set $\mathrm{S} \subseteq\{\rightarrow, \rightsquigarrow, \nrightarrow, \not \nsim\}$ above (resp., below) the frontier means that the corresponding problem is NP-hard (resp., feasible in polynomial time) on $\mathcal{C}[\mathrm{S}]$.

Theorem 9.13. If $\mathrm{S} \subseteq\{\rightarrow, \rightsquigarrow, \nrightarrow\}$ or $\mathrm{S} \subseteq\{\nsim\}\}$, then $\left.\mathrm{CN}-\mathrm{Mining}^{[\mathrm{S}}\right]$ is feasible in polynomial time. Otherwise, it is even intractable to check whether there is a solution at all (formally, the problem CN-Existence[S] is NP-complete).

We now provide the proof of the hardness results in Theorem 9.13. Note that these proofs are based on linear logs only. Hence, according to Fact 9.10, for these results it is immaterial whether CN-Mining is defined over extended causal nets or over (standard) causal nets.

For the analysis that follows, we find convenient to introduce a variant of the problem CN-Mining, which we call Acyclic-CN-Mining. Given a set $\mathcal{A}$ of activities, a $\log L$ with $\mathcal{A}(L) \subseteq \mathcal{A}$, and a set $\Pi$ of precedence constraints with $\mathcal{A}(\Pi) \subseteq \mathcal{A}$, the problem asks to compute a possibly extended causal net $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ over $\mathcal{A}$ such that $\mathcal{G}$ is acyclic, $\mathcal{C} \vdash L$, and $\mathcal{G} \models \Pi$, or to check that no net with these properties exists.

We start the proofs of the hardness results stated in Theorem 9.13, by focusing
on edge and path constraints imposed over the problem Acyclic-CN-Existence. Lemma 9.14. Acyclic-CN-Existence [ $\{\rightarrow\}]$ is NP-hard.

Proof. Consider the monotone one-in-three 3SAT problem defined as follows. We are given as input a Boolean formula in conjunctive normal form $\Phi=c_{1} \wedge \ldots \wedge c_{m}$ over the variables $X_{1}, \ldots, X_{n}$, where each clause $c_{j}$, with $j \in\{1, \ldots, m\}$, has the form $\left(t_{j, 1} \vee t_{j, 2} \vee t_{j, 3}\right)$, and $t_{j, 1}, t_{j, 2}$, and $t_{j, 3}$ are three distinct variables. Note that $\Phi$ is always satisfiable by the truth assignment where all variables evaluate true. In fact, the problem asks whether there is a satisfying truth assignment where each clause has exactly one variable evaluating true (and hence two variables evaluating false). This problem is known to be NP-complete Schaefer [1978].

Based on any formula $\Phi$ as above, we build the set $\mathcal{A}(\Phi)$ consisting of the variables in $\Phi$, which are transparently viewed as activities, plus the three distinguished activities $c, a_{\perp}$ and $a_{\top}$, where $a_{\perp}$ and $a_{\top}$ are as usual the starting and the terminating activity, respectively. So, we formally have $\mathcal{A}(\Phi)=\{c\} \cup\left\{X_{1}, \ldots, X_{n}\right\} \cup\left\{a_{\perp}, a_{\top}\right\}$.

Moreover, we build the set $\Pi(\Phi) \subseteq \mathcal{C}[\{\rightarrow\}]$ of edge constraints as follows. For each clause $c_{j}, \Pi(\Phi)$ contains the constraints $\left\{t_{j, 1}, t_{j, 2}, t_{j, 3}\right\} \rightarrow\{c\},\{c\} \rightarrow\left\{t_{j, 1}, t_{j, 2}\right\}$, $\{c\} \rightarrow\left\{t_{j, 1}, t_{j, 3}\right\}$, and $\{c\} \rightarrow\left\{t_{j, 2}, t_{j, 3}\right\}$. No further constraint is in $\Pi(\Phi)$.

We now claim that: There is a satisfying truth assignment to the variables of $\Phi$ such that each clause has exactly one variable evaluating true $\Leftrightarrow$ there is an acyclic dependency graph $\mathcal{G}$ (over $\mathcal{A}(\Phi)$ ) such that $\mathcal{G} \models \Pi(\Phi)$.
$(\Rightarrow)$ Assume that $\sigma$ is a satisfying truth assignment such that each clause has exactly one variable evaluating true. Consider the graph $\mathcal{G}(\Phi, \sigma)=(\mathcal{A}(\Phi), E)$ whose set of edges is defined as follows. For each variable $X_{h}$, with $h \in\{1, \ldots, n\}$, the edges $\left(a_{\perp}, X_{h}\right)$ and $\left(X_{h}, a_{\perp}\right)$ are in $E$. For each clause $c_{j}$ and each variable $t_{j, i}$ evaluating


Figure 9.4: Example reduction in the proof of Lemma 9.14.
true (resp., false) in $\sigma$, the edge $\left(t_{j, i}, c\right)$ (resp., $\left.\left(c, t_{j, i}\right)\right)$ is in $E$. The edges $\left(a_{\perp}, c\right)$ and $\left(c, a_{\mathrm{T}}\right)$ are in $E$, and no further edge is in $E$. As an example, the graph $\mathcal{G}(\Phi, \sigma)$ associated with the formula $\Phi=\left(X_{1} \vee X_{2} \vee X_{3}\right) \wedge\left(X_{3} \vee X_{4} \vee X_{5}\right)$ and the truth assignment $\sigma$, where $X_{1}$ and $X_{4}$ are the only variables evaluating true, is reported in Figure 9.4. In particular, note that the edges whose definition depend on $\sigma$ are depicted in bold.

We first show that $\mathcal{G}(\Phi, \sigma)$ is a dependency graph. Indeed, $a_{\perp}$ and $a_{\text {T }}$ have no ingoing and outgoing edges, respectively. Moreover, for each activity $a \in \mathcal{A}(\Phi) \backslash$ $\left\{a_{\perp}, a_{\top}\right\}$, the edges $\left(a_{\perp}, a\right)$ and $\left(a, a_{\top}\right)$ are in $E$, so that $a$ occurs in a path from $a_{\perp}$ to $a_{\top}$.

Then, we show that $\mathcal{G}(\Phi, \sigma)$ satisfies all the constraints in $\Pi(\Phi)$. Recall that $\sigma$ is a satisfying truth assignment such that each clause has exactly one variable evaluating true. Thus, for each clause $c_{j}$, with $j \in\{1, \ldots, m\}$, by construction there is an edge of the form $\left(t_{j, i}, c\right)$, so that the constraint $\left\{t_{j, 1}, t_{j, 2}, t_{j, 3}\right\} \rightarrow\{c\}$ is satisfied. Moreover, there are also two edges of the form $\left(c, t_{j, i^{\prime}}\right)$ and $\left(c, t_{j, i^{\prime \prime}}\right)$, where $i^{\prime} \neq i$,
$i^{\prime \prime} \neq i, i^{\prime} \neq i^{\prime \prime}$, and $\left\{i^{\prime}, i^{\prime \prime}\right\} \subseteq\{1,2,3\}$. Thus, the constraints $\{c\} \rightarrow\left\{t_{j, 1}, t_{j, 2}\right\}$, $\{c\} \rightarrow\left\{t_{j, 1}, t_{j, 3}\right\}$, and $\{c\} \rightarrow\left\{t_{j, 2}, t_{j, 3}\right\}$ are also satisfied, for each clause $c_{j}$. It follows that all constraints in $\Pi(\Phi)$ are satisfied by $\mathcal{G}(\Phi, \sigma)$, i.e., $\mathcal{G}(\Phi, \sigma) \models \Pi(\Phi)$.

In order to conclude, we now just need to point out that $\mathcal{G}(\Phi, \sigma)$ is acyclic. To this end, assume for the sake of contradiction that a cycle exists in $\mathcal{G}(\Phi, \sigma)$. Note that this cycle is necessarily defined over the variables and the distinguished activity $c$. Therefore, the set $E$ of edges must contain an edge of the form $\left(c, X_{h}\right)$ and an edge of the form $\left(X_{h}, c\right)$. By construction, the existence of the edge $\left(c, X_{h}\right)$ implies that $X_{h}$ is a variable evaluating false in $\sigma$. However, the existence of the edge $\left(X_{h}, c\right)$ implies that $X_{h}$ is a variable evaluating true in $\sigma$. Contradiction.
$(\Leftarrow)$ Assume that $\mathcal{G}=(V, E)$ is an acyclic dependency graph satisfying all the constraints in $\Pi(\Phi)$, and define the truth assignment $\sigma_{\mathcal{G}}$ such that $X_{h}$ evaluates true if, and only if, the edge $\left(X_{h}, c\right)$ occurs in $E$. We first show that $\sigma_{\mathcal{G}}$ is satisfying. Indeed, for each clause $c_{j}$, with $j \in\{1, \ldots, m\}$, consider the associated constraint $\left\{t_{j, 1}, t_{j, 2}, t_{j, 3}\right\} \rightarrow\{c\}$, and note that since $\mathcal{G} \models \Pi(\Phi)$, we are guaranteed about the existence of an edge from one of the variables in $\left\{t_{j, 1}, t_{j, 2}, t_{j, 3}\right\}$ to $c$. Hence, for each clause $c_{j}$, at least one of the variables occurring in $c_{j}$ evaluates true in $\sigma_{\mathcal{G}}$, by definition of this assignment, which is therefore satisfying.

Now, we show that $\sigma_{\mathcal{G}}$ is a such that each clause has exactly one variable evaluating true. Assume, for the sake of contradiction, that a clause $c_{j}$ exists such that two variables, say $t_{j, i^{\prime}}$ and $t_{j, i^{\prime \prime}}$ with $i^{\prime} \neq i^{\prime \prime}$, evaluate true in $\sigma_{\mathcal{G}}$. Thus, the edges $\left(t_{j, i^{\prime}}, c\right)$ and $\left(t_{j, i^{\prime \prime}}, c\right)$ are both in $E$. Consider then the constraint $\{c\} \rightarrow\left\{t_{j, i^{\prime}}, t_{j, i^{\prime \prime}}\right\}$ associated with the clause $c_{j}$, and note that it prescribes that at least one of the edges in $\left\{\left(c, t_{j, i^{\prime}}\right),\left(c, t_{j, i^{\prime \prime}}\right)\right\}$ occurs in $E$. Assume, w.l.o.g., that $\left(c, t_{j, i^{\prime}}\right)$ is in $E$ and observe that we have eventually a cycle over $c$ and $t_{j, i^{\prime}}$. Contradiction.

By Corollary 9.8 and Fact 9.10, the above entails that Acyclic-CN-Mining[\{ $\rightarrow\}$ ] on input $\mathcal{A}(\Phi)$, the empty log, and the set $\Pi(\Phi)$ has a solution if, and only if, there is a satisfying truth assignment to the variables of $\Phi$ such that each clause has exactly one variable evaluating true. As the reduction is feasible in polynomial time, it follows that Acyclic-CN-Existence $[\{\rightarrow\}]$ is NP-hard.

A straightforward adaptation of the above proof, where each edge constraint is replaced by a path constraint over the same sets of activities, can be used to show that the problem remains intractable if we consider path constraints in place of edge constraints. The proof is reported below, for the sake of completeness.

Lemma 9.15. Acyclic-CN-Existence [ $\{\rightsquigarrow\}]$ is NP-hard.

Proof. Consider again the setting in the proof of Lemma 9.14 and, for any formula $\Phi$, define $\Pi^{\prime}(\Phi) \subseteq \mathcal{C}[\{\sim\}]$ as the set of constraints obtained from $\Pi(\Phi)$ by replacing each edge constraint with the analogous path constraint. Therefore, for each clause $c_{j}, \Pi^{\prime}(\Phi)$ contains the constraints $\left\{t_{j, 1}, t_{j, 2}, t_{j, 3}\right\} \rightsquigarrow\{c\},\{c\} \rightsquigarrow\left\{t_{j, 1}, t_{j, 2}\right\},\{c\} \rightsquigarrow$ $\left\{t_{j, 1}, t_{j, 3}\right\}$, and $\{c\} \rightsquigarrow\left\{t_{j, 2}, t_{j, 3}\right\}$. We now claim that: There is a satisfying truth assignment to the variables of $\Phi$ such that each clause has exactly one variable evaluating true $\Leftrightarrow$ there is an acyclic dependency graph $\mathcal{G}$ (over $\mathcal{A}(\Phi)$ ) such that $\mathcal{G} \models \Pi^{\prime}(\Phi)$.
$(\Rightarrow)$ Assume that $\sigma$ is a satisfying truth assignment such that each clause has exactly one variable evaluating true, and consider the graph $\mathcal{G}=(\Phi, \sigma)$ built in the proof of Lemma 9.14. Recall that $\mathcal{G} \models \Pi(\Phi)$. Hence, we trivially have that $\mathcal{G} \models \Pi^{\prime}(\Phi)$, in particular because all path constraints are satisfied by direct connections.
$(\Leftarrow)$ Assume that $\mathcal{G}=(V, E)$ is an acyclic dependency graph satisfying all the constraints in $\Pi^{\prime}(\Phi)$, and define the truth assignment $\sigma_{\mathcal{G}}^{\prime}$ such that $X_{h}$ evaluates true
if, and only if, there is a path from $X_{h}$ to $c$ in $E$. Because of the constraints $\left\{t_{j, 1}, t_{j, 2}, t_{j, 3}\right\} \rightsquigarrow\{c\}$ occurring in $\Pi^{\prime}(\Phi)$ and associated with the clause $c_{j}$, for each $j \in\{1, \ldots, m\}, \sigma_{\mathcal{G}}^{\prime}$ is satisfying. In order to conclude, we need to show that $\sigma_{\mathcal{G}}^{\prime}$ is a such that each clause has exactly one variable evaluating true. Assume, for the sake of contradiction, that a clause $c_{j}$ exists such that two variables, say $t_{j, i^{\prime}}$ and $t_{j, i^{\prime \prime}}$ with $i^{\prime} \neq i^{\prime \prime}$, evaluate true in $\sigma_{\mathcal{G}}$. Thus, there is a path from $t_{j, i^{\prime}}$ to $c$ and a path from $t_{j, i^{\prime \prime}}$ to $c$ are both in $E$. Consider then the constraint $\{c\} \rightsquigarrow\left\{t_{j, i^{\prime}}, t_{j, i^{\prime \prime}}\right\}$ associated with the clause $c_{j}$, and note that it prescribes that there is a path from $c$ to at least one of the nodes in $\left\{t_{j, i^{\prime}}, t_{j, i^{\prime \prime}}\right\}$. Assume, w.l.o.g., that a path from $c$ to $t_{j, i^{\prime}}$ exists. Then, we have a cycle involving the activities $c$ and $t_{j, i^{\prime}}$. Contradiction.

By the above result, Corollary 9.8, Fact 9.10, and the fact that the reduction is feasible in polynomial time, it follows that Acyclic-CN-Existence[\{ $\sim\}]$ is NPhard.

Let us now turn to the case of negated edge constraints, but still focusing on the acyclic variant of the mining problem.

Lemma 9.16. Acyclic-CN-Existence $[\{\nrightarrow\}]$ is NP-hard.
Proof. Let $\Pi=\left\{\left\{b_{i}^{1}, \ldots, b_{i}^{k_{i}}\right\} \rightarrow\left\{a_{i}^{1}, \ldots, a_{i}^{h_{i}}\right\} \mid i \in\{1, \ldots, m\}\right\} \subseteq \mathcal{C}[\{\rightarrow\}]$ be a (nonempty) set of edge constraints, such that $\left\{b_{i}^{1}, \ldots, b_{i}^{k_{i}}\right\} \cap\left\{a_{i}^{1}, \ldots, a_{i}^{h_{i}}\right\}=\emptyset$, for each $i \in\{1, \ldots, m\}$. For each $i \in\{1, \ldots, m\}$, let $c_{i}$ denote a fresh activity not in $\mathcal{A}(\Pi)$. Moreover, let $a_{\perp}$ and $a_{\top}$ be two activities not in $\mathcal{A}(\Pi)$ playing the role of the starting and terminating activity, respectively. Then, consider the $\log L(\Pi)=\left\{t_{1}, \ldots, t_{m}\right\}$ such that $t_{i}=a_{\perp} b_{i}^{1} \ldots b_{i}^{k_{i}} c_{i} a_{i}^{1} \ldots a_{i}^{h_{i}} a_{\top}$, for each $i \in\{1, \ldots, m\}$. Moreover, consider the set $\Pi\urcorner$ of negated edge constraints such that $\Pi\urcorner=\left\{\neg\left(\left\{a_{\perp}\right\} \rightarrow\left\{c_{i}\right\}\right), \neg\left(\left\{c_{i}\right\} \rightarrow\right.\right.$ $\left.\left.\left\{a_{\top}\right\}\right) \mid i \in\{1, \ldots, m\}\right\}$.

We claim that: There is an acyclic dependency graph $\mathcal{G}$ over the set $\mathcal{A}(\Pi) \cup$ $\left\{a_{\perp}, a_{\top}\right\}$ of activities and such that $\mathcal{G} \models \Pi \Leftrightarrow$ there is an acyclic dependency graph $\mathcal{G}^{\prime}$ over the set $\mathcal{A}(\Pi) \cup\left\{a_{\perp}, a_{\top}\right\} \cup\left\{c_{1}, \ldots, c_{m}\right\}$ of activities and such that $\mathcal{G}^{\prime} \models \pi(L(\Pi)) \cup$ $\Pi\urcorner$.
$(\Rightarrow)$ Assume that $\mathcal{G}=(V, E)$ is an acyclic dependency graph over the set $\mathcal{A}(\Pi) \cup$ $\left\{a_{\perp}, a_{\mathrm{T}}\right\}$ of activities and such that $\mathcal{G} \models \Pi$. Consider the graph $\mathcal{G}^{\prime}=\left(V^{\prime}, E \cup E_{1}^{\prime} \cup\right.$ $\left.E_{2}^{\prime}\right)$ where $V^{\prime}=V \cup\left\{c_{1}, \ldots, c_{m}\right\}$ and whose set of edges is built as follows. The set $E_{1}^{\prime}$ contains the edges $\left(a_{\perp}, a\right)$ and $\left(a, a_{\mathrm{T}}\right)$, for each activity $a \in \mathcal{A}(\Pi)$, and no further edge is in $E_{1}^{\prime}$. Moreover, for each $i \in\{1, \ldots, m\}$, and for each edge $(x, y) \in E$ such that $x \in\left\{b_{i}^{1}, \ldots, b_{i}^{k_{i}}\right\}$ and $y \in\left\{a_{i}^{1}, \ldots, a_{i}^{h_{i}}\right\}, E_{2}^{\prime}$ includes the edges $\left(x, c_{i}\right)$ and $\left(c_{i}, y\right)$. No further edge is in $E_{2}^{\prime}$. It is immediately checked that $\mathcal{G}^{\prime}$ is an acyclic dependency graph.

Now, we first claim that $\left.\mathcal{G}^{\prime} \models \Pi\right\urcorner$. Indeed, $\left\{a_{\perp}, a_{\top}\right\} \cap \mathcal{A}(\Pi)=\emptyset$, so that $E_{2}^{\prime}$ does not include any edge of the form $\left(a_{\perp}, c_{i}\right)$ or of the form $\left(c_{i}, a_{\top}\right)$, with $i \in\{1, \ldots, m\}$. Let now $t_{i}$ be a trace in $L(\Pi)$, with $i \in\{1, \ldots, m\}$. Consider the set $\pi\left(t_{i}\right)$ of constraints derived from $t_{i}$ according to Definition 9.5. Because of the edges in $E_{1}^{\prime}$, all constraints in $\pi\left(t_{i}\right)$ are trivially satisfied by $\mathcal{G}^{\prime}$, but the two constraints $\left\{a_{\perp}, b_{i}^{1}, \ldots, b_{i}^{k_{i}}\right\} \rightarrow\left\{c_{i}\right\}$ and $\left\{c_{i}\right\} \rightarrow\left\{a_{i}^{1}, \ldots, a_{i}^{h_{i}}, a_{\mathrm{T}}\right\}$, because there is no edge in $\mathcal{G}^{\prime}$ from $a_{\perp}$ to $c_{i}$, and from $c_{i}$ to $a_{\top}$. Now, recall that $\mathcal{G}$ satisfies $\Pi$ and hence, an edge $(x, y)$ occurs in $E$ such that $x \in\left\{b_{i}^{1}, \ldots, b_{i}^{k_{i}}\right\}$ and $y \in\left\{a_{i}^{1}, \ldots, a_{i}^{h_{i}}\right\}$. By construction of the edges in $E_{2}^{\prime}$, we therefore have that $\left(x, c_{i}\right)$ and $\left(c_{i}, y\right)$ are edges of $\mathcal{G}^{\prime}$, which proves that the two constraints are satisfied, too. Hence, $\mathcal{G}^{\prime} \models \pi(L(\Pi))$.
$(\Leftarrow)$ Assume there is an acyclic dependency graph $\mathcal{G}^{\prime}=\left(V^{\prime}, E^{\prime}\right)$ over the set $\mathcal{A}(\Pi) \cup$ $\left\{a_{\perp}, a_{\mathrm{T}}\right\} \cup\left\{c_{1}, \ldots, c_{m}\right\}$ of activities and such that $\left.\mathcal{G}^{\prime} \models \pi(L(\Pi)) \cup \Pi\right\urcorner$. Let $t_{i}$ be a trace in $L(\Pi)$, with $i \in\{1, \ldots, m\}$, and consider the two constraints $\left\{a_{\perp}, b_{i}^{1}, \ldots, b_{i}^{k_{i}}\right\} \rightarrow\left\{c_{i}\right\}$
and $\left\{c_{i}\right\} \rightarrow\left\{a_{i}^{1}, \ldots, a_{i}^{h_{i}}, a_{\top}\right\}$ in the set $\pi\left(t_{i}\right)$, which are satisfied by $\mathcal{G}^{\prime}$. Since $\left.\mathcal{G}^{\prime} \models \Pi\right\urcorner$, from the above we conclude that $\mathcal{G}^{\prime} \models\left\{\left\{b_{i}^{1}, \ldots, b_{i}^{k_{i}}\right\} \rightarrow\left\{c_{i}\right\},\left\{c_{i}\right\} \rightarrow\left\{a_{i}^{1}, \ldots, a_{i}^{h_{i}}\right\}\right\}$ holds.

Consider now the graph $\mathcal{G}=(V, E)$ where $V=V^{\prime} \backslash\left\{c_{1}, \ldots, c_{m}\right\}$ and where the set of edges is defined as follows. The set $E$ contains the edges $\left(a_{\perp}, a\right)$ and $\left(a, a_{\top}\right)$, for each activity $a \in V$. Moreover, for each $i \in\{1, \ldots, m\}, E$ includes all edges of the form $(x, y)$ such that $\left\{\left(x, c_{i}\right),\left(c_{i}, y\right)\right\} \subseteq E^{\prime}, x \in\left\{b_{i}^{1}, \ldots, b_{i}^{k_{i}}\right\}$, and $y \in\left\{a_{i}^{1}, \ldots, a_{i}^{h_{i}}\right\}$. Note that, in the light of the above observation, at least one edge of this kind is included in $E$, so that $\mathcal{G} \models \Pi$ holds. Note also that $\mathcal{G}$ is a dependency graph. Eventually, to conclude the proof, just note that $\mathcal{G}$ is acyclic, as the existence of a cycle in $\mathcal{G}$ would immediately entail the existence of a cycle in $\mathcal{G}^{\prime}$.

Observe now that the $\log L(\Pi)$ is linear, so that we are in the position of applying Corollary 9.8 and Fact 9.10 on the result proven above. By this way, we conclude that Acyclic-CN-Mining $[\{\rightarrow\}]$ on input $\mathcal{A}(\Pi) \cup\left\{a_{\perp}, a_{\top}\right\}$, the empty log, and the set $\Pi$ has a solution if, and only if, Acyclic-CN-Mining $[\{\nrightarrow\}]$ on input $\mathcal{A}(\Pi) \cup$ $\left\{a_{\perp}, a_{\top}\right\} \cup\left\{c_{1}, \ldots, c_{m}\right\}$, the $\log L(\Pi)$, and the set $\left.\Pi\right\urcorner$ has a solution. Eventually, by inspecting the proof of Lemma 9.14, note that the constraints used to prove the NPhardness of Acyclic-CN-Mining $[\{\rightarrow\}]$ are precisely of the form considered here for the set $\Pi$, and that the result is established even for logs that are empty. Therefore, we have reduced Acyclic-CN-Mining $[\{\nrightarrow\}]$ to Acyclic-CN-Mining $[\{\rightarrow\}]$, so that the problem is hence shown to be NP-hard, too.

To complete the picture, we now move to the case of arbitrary process models, i.e., of models that are not required to be acyclic. In this context, the picture is easily completed, as negated path constraints can be used to enforce acyclicity.

Theorem 9.17. The problems CN-Mining $[\{\rightarrow, \nrightarrow\}\}], \mathrm{CN}-\mathrm{Mining}[\{\sim, \nrightarrow \psi\}]$, and CN$\operatorname{Mining}[\{\nrightarrow, \nrightarrow \psi\}]$ are NP-hard.

Proof. Let $\Pi$ be a set of constraints in $\mathcal{C}[\{\rightarrow\}]$ (resp., $\mathcal{C}[\{\rightsquigarrow\}], \mathcal{C}[\{\nrightarrow\}])$. Based on $\Pi$, we build the set $\Pi^{\prime}$ of constraints including all the constraints in $\Pi$, plus the novel constraint $\neg(\{a\} \nLeftarrow \rightarrow\{a\})$, for each activity $a$ taken from the underlying set of symbols $\mathcal{A}$. Of course, $\Pi^{\prime}$ belongs to $\mathcal{C}[\{\rightarrow, \not \not \nsim\}]$ (resp., $\left.\mathcal{C}[\{\rightsquigarrow, \not \not \nsim\}], \mathcal{C}[\{\nrightarrow, \not \not \nsim\}\}\right]$ ). In particular, the novel constraints just enforce that the resulting process model is acyclic. Then, Acyclic-CN-Mining $[\{\rightarrow\}]$ (resp., Acyclic-CN-Mining $[\{\rightsquigarrow\}]$, Acyclic-CN-Mining $[\{\nrightarrow\}])$ has a solution on input the set $\mathcal{A}$ of activities, a log $L$, and the set $\Pi$ of constraints if, and only if, CN-Mining $[\{\rightarrow, \nrightarrow\}\}$ (resp., CN$\operatorname{Mining}[\{\leadsto, \not \not \nsim\}], \operatorname{CN}-\operatorname{Mining}[\{\nrightarrow, \nsim\}\}])$ has a solution on input $\mathcal{A}, L$, and $\Pi^{\prime}$. The result therefore follows from Lemma 9.14, Lemma 9.15, and Lemma 9.16.

We now show that the decision version belongs to the complexity class NP. Combined with the NP-hardness results, this entails the corresponding NP-completeness results.

Theorem 9.18. CN-Existence $[\mathrm{S}]$ is in NP , for each set $\mathrm{S} \subseteq\{\rightarrow, \rightsquigarrow, \nrightarrow, \nprec \rightarrow\}$.

Proof. We need to decide about the existence of an extended causal net $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ over $\mathcal{A}$ such that $\mathcal{C} \vdash L$ and $\mathcal{G} \models \Pi$, where $L$ and $\Pi$ are the log and the set of constraints, respectively, provided as input. Thus, we can build a non-deterministic Turing machine that guesses a graph $\mathcal{G}=(V, E)$, with $V \subseteq \mathcal{A}$, and checks that $\mathcal{G} \models \Pi$. Moreover, for each trace $t \in L$, the machine guesses a sequence $\sigma_{t}$ of binding activities $\left\langle t[1], i b_{1}, o b_{1}\right\rangle, \ldots,\left\langle t[l e n(t)], i b_{n}, o b_{n}\right\rangle$ and checks that $\sigma_{t}$ is valid w.r.t. $\mathcal{C}$. By Theorem 9.4, we can assume w.l.o.g. that the size of $\mathcal{G}$ is polynomially bounded. Thus, the overall size of the structures guessed by the machine is polynomially bounded.

Moreover, the operations performed on them are feasible in polynomial time. Hence, the problem belongs to NP. For completeness, note that if one restricts the problem to causal nets only, then the problem is still feasible in NP, as the machine has just to additionally check that $\mathcal{I}$ and $\mathcal{O}$ are sets rather than multi-sets as in the extended model.

### 9.5 An exact solution approach for computing process models

The complexity analysis we have conducted evidenced that polynomial time algorithms are unlikely to exists for the problem of computing models (and minimal models) of sets of precedence constraints. This bad news calls for the design of sophisticated exact solution approaches that perform well in practice, and of heuristic methods that further speed up the computation These complementary solution approaches will be explored in the following two sections.

In this section, we start by considering the issue of designing exact solution approaches. To this end, we propose to encode precedence constraints in terms of "standard" constraints satisfaction problems (short: CSPs) and to reuse existing constraint programming platforms to compute models of them. Indeed, such platforms have been developed to solve NP-hard problems declaratively specified in terms of CSPs, and embody sophisticated solution algorithms (see, e.g., Apt . [2003]) allowing them to scale over large datasets, as their application in data mining contexts have already demonstrated De Raedt et al. [2008]; Nijssen et al. [2009].

Various kinds of constraints are supported by constraint programming systems in the literature. To our ends, we just need to consider two kinds of constraints defined over binary domains. Let $X_{1}, \ldots, X_{n}$ be $n$ variables in $\operatorname{Var}$, let $U\left(X_{i}\right)=\{0,1\}$, and let $w_{1}, \ldots, w_{n}, \gamma$ be $n+1$ real numbers. Then,
(1) A summation constraint is an expression of the form $\sum_{i=1}^{n} w_{i} \times X_{i} \geq \gamma$. The constraint is satisfied by an assignment $\theta$ if $\sum_{i=1}^{n} w_{i} \times \theta\left(X_{i}\right) \geq \gamma$ holds.
(2) A reified (summation) constraint is an expression of the form $\sum_{i=1}^{n} w_{i} \times X_{i} \geq$ $\theta \leftrightarrow X$, where $X$ is a variables such that $U(X)=\{0,1\}$. The constraint is satisfied by an assignment $\theta$ if $\sum_{i=1}^{n} w_{i} \times \theta\left(X_{i}\right) \geq \gamma$ holds if, and only, if $\theta(X)=1$.

In some cases, among all the possible solutions, we are interested in the one optimizing some given criterium. This gives rise to a constraint satisfaction optimization problem (short: CSOP) instance, that is, a quadruple ( $\operatorname{Var}, U, \mathcal{C}, f)$ where ( $\operatorname{Var}, U, \mathcal{C}$ ) is the underlying CSP instance, and where $f$ is a valuation function assigning a real number $f(\theta)$ to each solution $\theta$ to (Var $, U, \mathcal{C})$. A solution $\theta$ to the CSOP instance is then a solution to the CSP instance such that $f(\theta) \leq f\left(\theta^{\prime}\right)$, for each other CSP solution $\theta^{\prime}$.

In several constraint programming platforms, the function $f$ can be specified in terms of a linear minimization constraint. Let $X_{1}, \ldots, X_{n}$ be $n$ variables in Var, let $U\left(X_{i}\right)=\{0,1\}$, and let $w_{1}, \ldots, w_{n}$ be $n$ real numbers. Then, a linear minimization constraint is an expression of the form $\sum_{i=1}^{n} w_{i} \times X_{i}$. Under this function, the value associated with a CSP solution $\theta$ is $f(\theta)=\sum_{i=1}^{n} w_{i} \times \theta\left(X_{i}\right)$.

### 9.5.1 CSP encoding for precedence constraints

Let $\mathcal{A}=\left\{a_{1}, \ldots, a_{n}\right\}$ be any set of activities, and let $\Pi$ be any set of precedence constraints such that $\mathcal{A}(\Pi) \subseteq \mathcal{A}$. Figure 9.5 illustrates an algorithm, named PC тoCSP, to encode $\mathcal{A}$ and $\Pi$ into a CSP instance ( $\operatorname{Var}, U, \mathcal{C}$ ). The instance defined by the algorithm is such that, for each pair $a_{i}$ and $a_{j}$ of activities taken from $\mathcal{A}$, Var contains an "edge" variable, denoted as $e\left[a_{i}, a_{j}\right]$, plus $n+n^{2}$ "path" variables,

```
Input: A set \mathcal{A}={\mp@subsup{a}{1}{},\ldots,\mp@subsup{a}{n}{}}\mathrm{ of activities,}
    a set \Pi of precedence constraints with \mathcal{A (\Pi) \subseteq\mathcal{A}}\mathrm{ , and}
    the TYPE }\in{\mathrm{ ARBITRARY, ACYCLIC } of the problem;
Output: A CSP instance (Var, U,\mathcal{C});
    let Var = {e[ai, aj],p[\mp@subsup{a}{i}{},\mp@subsup{a}{j}{}\mp@subsup{]}{}{\ell},p[\mp@subsup{a}{i}{},\mp@subsup{a}{k}{},\mp@subsup{a}{j}{}\mp@subsup{]}{}{\ell}|\mp@subsup{a}{i}{},\mp@subsup{a}{j}{},\mp@subsup{a}{k}{}\in\mathcal{A},\ell\in{1,\ldots,n}};
    let U(X)={0,1} for each X \inVar;
    let }\mathcal{C}=\emptyset\mathrm{ ;
    for each edge constraint S->a}\mp@subsup{a}{j}{}\mathrm{ in }\Pi\mathrm{ do
        \mathcal { C } : = \mathcal { C } \cup \{ \sum _ { a _ { i } \in S } e [ a _ { i } , a _ { j } ] \geq 1 \} ;
    for each path constraint S\rightsquigarrow\mp@subsup{a}{j}{}\mathrm{ in }\Pi\mathrm{ do}
        C}:=\mathcal{C}\cup{\mp@subsup{\sum}{\mp@subsup{a}{i}{}\inS}{}p[\mp@subsup{a}{i}{},\mp@subsup{a}{j}{}\mp@subsup{]}{}{n}\geq1
    for each negated edge constraint }\negS->\mp@subsup{a}{j}{}\mathrm{ in }\Pi\mathrm{ do
            \mathcal { C } : = \mathcal { C } \cup \{ e [ a _ { i } , a _ { j } ] = 0 \| a _ { i } \in S \} ;
    for each negated path constraint }\negS\rightsquigarrow\mp@subsup{a}{j}{}\mathrm{ in }\Pi\mathrm{ do
        \mathcal { C } : = \mathcal { C } \cup \{ p [ a _ { i } , a _ { j } ] ^ { n } = 0 \| a _ { i } \in S \} ;
    if TYPE=ACYCLIC then
        \mathcal { C } : = \mathcal { C } \cup \{ p [ a _ { i } , a _ { j } ] ^ { n } + p [ a _ { j } , a _ { i } ] ^ { n } \leq 1 \| \{ a _ { i } , a _ { j } \} \subseteq \mathcal { A } \} ;
    for each pair of distinct activities ai and a}\mp@subsup{a}{j}{}\mathrm{ do
        \mathcal { C } : = \mathcal { C } \cup \{ e [ a _ { i } , a _ { j } ] \geq 1 \leftrightarrow p [ a _ { i } , a _ { j } ] ^ { 1 } \} ;
        \mathcal{C}:=\mathcal{C}\cup{e[\mp@subsup{a}{i}{},\mp@subsup{a}{k}{}]+p[\mp@subsup{a}{k}{},\mp@subsup{a}{j}{}\mp@subsup{]}{}{\ell-1}\geq2\leftrightarrowp[\mp@subsup{a}{i}{},\mp@subsup{a}{k}{},\mp@subsup{a}{j}{}\mp@subsup{]}{}{\ell}|\mp@subsup{a}{k}{}\in\mathcal{A},\ell\in{2,\ldots,n}};
        \mathcal { C } : = \mathcal { C } \cup \{ \sum _ { k } p [ a _ { i } , a _ { k } , a _ { j } ] ^ { \ell } \geq 1 \leftrightarrow p [ a _ { i } , a _ { j } ] ^ { \ell } \| \ell \in \{ 2 , \ldots , n \} \} ;
    for each activity }\mp@subsup{a}{i}{}\mathrm{ , with }i\not\in{1,n} d
        \mathcal{C}:=\mathcal{C}\cup{p[\mp@subsup{a}{1}{},\mp@subsup{a}{i}{}\mp@subsup{]}{}{n}+p[\mp@subsup{a}{i}{},\mp@subsup{a}{n}{}\mp@subsup{]}{}{n}\geq2};
        \mathcal { C } : = \mathcal { C } \cup \{ \sum _ { a _ { k } \in \mathcal { A } } p [ a _ { k } , a _ { 1 } ] ^ { n } + \sum _ { a _ { k } \in \mathcal { A } } ^ { - } p [ a _ { 1 } , a _ { k } ] ^ { n } \leq 0 \} ;
    return(Var,U,\mathcal{C});
```

Figure 9.5: Algorithm PCToCSP.
denoted as $p\left[a_{i}, a_{j}\right]^{\ell}$ and $p\left[a_{i}, a_{k}, a_{j}\right]^{\ell}$, with $a_{k} \in \mathcal{A}$ and $\ell \in\{1, \ldots, n\}$. All variables are defined over the binary domain $\{0,1\}$.

To help the intuition, we anticipate that $e\left[a_{i}, a_{j}\right]$ is meant to encode the existence of an edge from $a_{i}$ to $a_{j}$, while $p\left[a_{i}, a_{j}\right]^{\ell}$ is meant to encode the existence of a path of length at most $\ell$.

Steps 4, 5, 6, and 7 in the algorithm in Figure 9.5 encode in a straightforward manner edge, path, negated edge, and negated path constraints of $\Pi$, respectively. For example, for each edge constraint $S \rightarrow T$ occurring in $\Pi$, the constraint $\sum_{a_{i} \in S} \sum_{a_{j} \in T} e\left[a_{i}, a_{j}\right] \geq 1$ is added to $\mathcal{C}$ stating that at least one of the edge variables from any activity in $S$ to any activity in $T$ must be mapped to 1 .

Step 8 is responsible of enforcing acyclicity if the additional input parameter TYPE is ACYCLIC. Indeed the constraint states that for each pair of distinct activities $a_{i}$ and $a_{j}$, at most one path variable in $\left\{p\left[a_{i}, a_{j}\right]^{n}, p\left[a_{j}, a_{i}\right]^{n}\right\}$ can be set to 1 .

Step 9, is responsible of providing the semantics to the variable of the form $p\left[a_{i}, a_{j}\right]^{n}$. Indeed, it enforces that $p\left[a_{i}, a_{j}\right]^{n}$ is mapped to 1 if, and only if, $e\left[a_{i}, a_{j}\right]$ is mapped to 1 or there is an intermediate activity $a_{k}$ with $e\left[a_{i}, a_{k}\right]+p\left[a_{k}, a_{j}\right]^{n-1} \geq 1$.

Note that, the definition is recursive, and is explicitly unfolded in the encoding as the maximum length of any path is $n$. The base case occurs for variables of the form $p\left[a_{i}, a_{j}\right]^{1}$, whose value coincides with that of the corresponding edge variables. More generally, the definition exploits the auxiliary variables having the form $p\left[a_{i}, a_{k}, a_{j}\right]^{\ell}$, encoding the existence of an edge from $a_{i}$ to $a_{k}$ and of a path of length $\ell-1$ at most from $a_{k}$ to $a_{j}$.

Finally, step 10 and 11 enforce the conditions required to hold over dependency graphs, where $a_{1}$ and $a_{n}$ play the role of the starting and terminating activity, respectively. In particular, each activity $a_{i}$ must occur in a path connecting $a_{1}$ and $a_{n}$, and $a_{1}$ (resp., $a_{n}$ ) does not have ingoing (resp., outgoing) edges.

To analyze the algorithm, for any solution $\theta$ to the $\operatorname{CSP}(\operatorname{Var}, U, \mathcal{C})$ computed by РСтоСSP (on input $\mathcal{A}, \Pi$ and Type), let $\mathcal{G}_{\theta}=(\mathcal{A}, E)$ be the graph such that $\left(a_{i}, a_{j}\right)$ is in $E$ if, and only if, $\theta\left(e\left[a_{i}, a_{j}\right]\right)=1$. With this notation, the crucial properties of the encodings are stated below.

Theorem 9.19. Let $\theta$ be a solution to (Var, $U, \mathcal{C}$ ). Then, $\mathcal{G}_{\theta}$ is a dependency graph such that $\mathcal{G}_{\theta} \models \Pi$. Moreover, if TYPE=ACYCLIC, then $\mathcal{G}$ is acyclic.

Theorem 9.20. Let $\mathcal{G}$ be a dependency graph (resp., acyclic dependency graph) such that $\mathcal{G} \models \Pi$. If TYPE=ARBITRARY (resp., ACYCLIC), then there is a solution $\theta$ to $($ Var $, U, \mathcal{C})$ such that $\mathcal{G}=\mathcal{G}_{\theta}$.

By combining Theorem 9.19, Theorem 9.20 and Corollary 9.8 (with the function computeBindings discussed in the previous section), we get that PCToCSP is an effective method to solve CN-Mining and Acyclic-CN-Mining over linear logs. In order to deal with arbitrary logs, we need and extension of the algorithm illustrated in Figure 9.5. This extension is the algorithm DGToCSP illustrated in Figure 9.6.

```
Input: A 酋 L, a set \Pi of precedence constraints over \mathcal{A}}(L)={\mp@subsup{a}{1}{},\ldots,\mp@subsup{a}{n}{}}\mathrm{ , and
    the TYPE of the problem;
Output: A CSP instance (Var, U,\mathcal{C});
    let (Var
    let (Var}2,\mp@subsup{U}{2}{},\mp@subsup{\mathcal{C}}{2}{})\mathrm{ be the output of РСтоСSP(П,TYPE);
    let Var= Var }\cup\cup\mp@subsup{Var}{2}{}\cup{eProjected [a, , a j] | a i, 和 \in\mathcal{A}}
    let U(X)={0,1} for each X G Var;
    let \mathcal{C}=\mp@subsup{\mathcal{C}}{1}{}\cup\mp@subsup{\mathcal{C}}{2}{};
    for each pair of distinct activities }\mp@subsup{a}{i}{}\mathrm{ and }\mp@subsup{a}{j}{}\mathrm{ do
        \mathcal { C } : = \mathcal { C } \cup \{ e [ a _ { i } , a _ { j } ] \geq 1 \leftrightarrow ~ e P r o j e c t e d ~ [ a , ~ , ~ a ~ a ~ ] ~ ] ; ;
        \mathcal { C } : = \mathcal { C } \cup \{ \sum _ { h , k } e [ a _ { i } \langle h \rangle , a _ { j } \langle k \rangle ] \geq 1 \leftrightarrow e P P r o j e c t e d ~ [ ~ a ~ i , ~ a j ] \} ;
    return (Var, U,\mathcal{C});
```

Figure 9.6: Algorithm DG-DiscoveryToCSP.
Corollary 9.21. Let $L$ be a $\log$ where no trace contains multiple occurrences of the same activity. Let $\mathcal{G}$ be a graph (resp., acyclic graph) over $\mathcal{A}(L)$, and $\Pi$ be a set of precedence constraints over $\mathcal{A}(L)$. Then, the followings are equivalent:
(1) $\mathcal{G}$ is a solution to CN-Mining (resp., Acyclic-CN-Mining) on input $L$ and $\Pi$.
(2) $\mathcal{G}=\mathcal{G}_{\theta}$ where $\theta$ is a solution to the CSP computed by PCTOCSP on input $\pi(L) \cup \Pi$ with TYPE=ARBITRARY (resp., ACYCLIC).

The algorithm receives as input a set $\mathcal{A}$ of activities, a $\log L$ with $\mathcal{A}(L) \subset \mathcal{A}$, a set $\Pi$ of precedence constraints with $\mathcal{A}(\Pi) \subset \mathcal{A}$, and a TYPE.

First, it computes the CSP instances $\left(\operatorname{Var}_{1}, U_{1}, \mathcal{C}_{1}\right)$ and ( $\operatorname{Var}_{2}, U_{2}, \mathcal{C}_{2}$ ) produced by PCtoCSP on input $(\mathcal{A}(\operatorname{unfold}(L)), \Pi(u n f o l d(L))$, $\operatorname{ACYCLIC})$ and $(\mathcal{A}, \Pi$, TYPE $)$, respectively. Then, a novel CSP instance $(V A R, U, \mathcal{C})$ is defined as the union of these two instances, and the set $\mathcal{C}$ of constraints is eventually enlarged by including some constraints that server to relate the variables in $V a r_{1}$ with the variables in $\operatorname{Var}_{2}$. In particular, observe that the CSP instance $\left(\operatorname{Var}_{1}, U_{1}, \mathcal{C}_{1}\right)$ is build given as input the set $\mathcal{A}(\operatorname{unfold}(L))$ of fresh virtual activities, and recall from section 8.4 that any activity in $\mathcal{A}(\operatorname{unfold}(L))$ has the form where $a$ is an activity in $\mathcal{A}$ and $k$ is the number of its occurrences in some trace of the $\log$ (up to some given position). Then the constraints added in step 6 guarantee that the variable $e\left[a_{i}, a_{j}\right] \in \operatorname{Var}_{2}$ is mapped to 1 , for each pair of activities $a_{i}$ and $a_{j}$ if, and only if, there is at least one variable
of the form $e\left[a_{i}\langle h\rangle, a_{j}\langle k\rangle\right] \in \operatorname{Var}_{1}$ being mapped to 1 . In words, variables in $\operatorname{Var}_{2}$ are defined as the "projection" of the corresponding variables in $V a r_{1}$, i.e., by just stripping off the subscripts denoting the number of occurrences.

Proposition 9.22. Let $L$ be a $\log$, let $\mathcal{G}$ be a graph over $\mathcal{A}(L)$, and let $\Pi$ be a set of precedence constraints over $\mathcal{A}$. Then, the followings are equivalent:
(1) $\mathcal{G}$ is a solution to CN-Mining (resp., Acyclic-CN-Mining) on input $L$ and $\Pi$.
(2) $\mathcal{G}=\mathcal{G}_{\theta}$ where $\theta$ is a solution to the CSP computed by DG-DiscoveryToCSP on input $L, \Pi$, and TYPE=ARBITRARY (resp., ACYCLIC).

Proof. Note that the constraints in step 6 guarantees that $\mathcal{G}_{\theta}$ is the folding of the graph $\overline{\mathcal{G}}_{\theta}$, for each solution $\theta$. Thus, the result follows from the correctness of algorithm PCтоCSP (cf. Theorem 9.19 and Theorem 9.20) guaranteeing that $\overline{\mathcal{G}}_{\theta} \models \pi(\bar{L})$, and by the application of Corollary 9.8.

### 9.5.2 Structural optimization

Several process models (in principle, up to exponentially many) might be built as solutions to CN-Mining and Acyclic-CN-Mining given a log and a set of precedence constraints provided as input. Returning an arbitrary one is usually not enough in practical applications, and we might rather want to formalize an objective function over the candidate solutions, as compute the best one over them. In order to formalize an optimality criterium, we next propose to focus on those models whose dependencies are maximal w.r.t. some (reflexive and transitive) order Let $\mathcal{A}$ denote, as usual, a given set of activities. Moreover, assume that a weighting function $w$ : $\mathcal{A} \times \mathcal{A} \longmapsto \mathbb{R}$ is given, which associates a value $w(e) \in \mathbb{R}$ with each element $e \in \mathcal{A} \times \mathcal{A}$, Then, for any pair of directed graphs $\mathcal{G}=(V, E)$ and $\mathcal{G}^{\prime}=\left(V^{\prime}, E^{\prime}\right)$ over $\mathcal{A}$, we write
$\mathcal{G} \sqsubseteq_{w} \mathcal{G}^{\prime}$ if $\sum_{e \in E} w(e) \leq \sum_{e^{\prime} \in E^{\prime}} w\left(e^{\prime}\right)$. Hence, $\mathcal{G} \sqsubseteq_{w} \mathcal{G}^{\prime}$ means that the overall weight associated with $\mathcal{G}$ is at most the overall weight associated with $\mathcal{G}^{\prime}$, A solution $\mathcal{C}=\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ to CN-Mining (resp., AcYCLIC-CN-Mining) is $\sqsubseteq_{w}$-maximal if, for each other solution $\mathcal{C}^{\prime}=\left\langle\mathcal{G}^{\prime}, \mathcal{I}^{\prime}, \mathcal{O}^{\prime}\right\rangle$ to CN-Mining (resp., Acyclic-CN-Mining), it holds that $\mathcal{G}^{\prime} \sqsubseteq_{w} \mathcal{G}$. In words, by considering $\sqsubseteq_{w}$-maximal solutions, we look for dependency graphs having associated the maximum total weight.

Example 9.23. As a very simple weighting function, one might consider the constant function $1: \mathcal{A} \times \mathcal{A} \longmapsto\{-1\}$ assigning a unitary negative weight to each element in $\mathcal{A} \times \mathcal{A}$. Then, $\sqsubseteq_{1}$-maximal solutions are those with the minimum possible number of edges. Consider, for instance, the graphs $\mathcal{G}_{2}$ and $\mathcal{G}_{3}$ depicted in Figure 8.1. It is easily seen that $\mathcal{G}_{3} \sqsubseteq_{1} \mathcal{G}_{2}$, because $\mathcal{G}_{2}$ consists of 5 edges, while $\mathcal{G}_{3}$ consists of 6 edges. As a further example, consider the graph $\mathcal{G}_{1}$ depicted in Figure 8.1, which consists of 5 edges. Therefore, $\mathcal{G}_{1} \sqsubseteq_{1} \mathcal{G}_{2}$ and $\mathcal{G}_{2} \sqsubseteq_{1} \mathcal{G}_{1}$ hold.

By using linear maximization constraints on top of the rewriting in Figure 9.6, we may compute $\sqsubseteq_{w}$-maximal solutions. In particular, if we consider solutions $\theta$ to the $\operatorname{CSOP}\left(\operatorname{Var}, U, \mathcal{C}, \sum_{a_{i}, a_{j}} w\left(\left(a_{i}, a_{j}\right)\right) \times e\left[a_{i}, a_{j}\right]\right)$, where $(\operatorname{Var}, U, \mathcal{C})$ is the CSP computed by DG-DiscoveryToCSP, then solutions in point (2) of Theorem 9.6 are $\sqsubseteq_{w}$-maximal ones. In fact, this is a simple consequence of the semantics of linear maximization constraints. Eventually, we stress that as a useful weighting function, we might want to consider here the causal score defined in Section 9.6.

### 9.6 Classes of Tractable Precedence Constraints and Algorithms

In this section, we define two algorithms for efficiently solving CN-Mining over the classes $\mathcal{C}[\{\rightarrow, \rightsquigarrow, \nrightarrow\}]$ and $\mathcal{C}[\{\nsim\}\}]$, respectively. The algorithms have been conceived as to make them still applicable over larger classes of constraints, by providing heuris-
tic solution methods in these (NP-hard) cases. The efficiency of the methods and their efficacy as heuristics will be eventually assessed in Section 9.7.

Hereinafter, we assume that a set $\mathcal{A}$ of activities is given, together with a $\log L$ such that $\mathcal{A}(L) \subseteq \mathcal{A}$ and with a set $\Pi$ of precedence constraints such that $\mathcal{A}(\Pi) \subseteq \mathcal{A}$. Accordingly, to simplify the notation, we shall omit to indicate $\mathcal{A}, L$, and $\Pi$, unless we want to explicitly point out some dependency on some of them. Moreover, we shall denote by $\Pi^{\rightarrow}$ and $\Pi^{\leadsto}$ (resp., $\Pi^{\nrightarrow}$ and $\Pi^{\curvearrowright>}$ ) the sets of all positive (resp., negated) edge and path constraints in $\Pi$, respectively.

Both algorithms are based on a succession of graph manipulations, i.e., insertions and deletions of edges, starting with an initial (dependency) graph built from the log. In order to facilitate reasoning about such graph manipulations, for any directed graph $\mathcal{G}=(V, E)$ and for any set $E^{\prime} \subseteq V \times V$ of edges, we define $\mathcal{G} \oplus E^{\prime}$ as the graph $\left(V, E \cup E^{\prime}\right)$, and $\mathcal{G} \ominus E^{\prime}$ as the graph $\left(V, E \backslash E^{\prime}\right)$. Moreover, we observe that, while performing these operations, it is practically relevant to (resp., exclude) only those edges that very likely (resp., hardly) witness the existence of true causal relationships. In order to provide a formal measure of the 'quality' of an edge, we consider here the notion of causal score inspired by the works of [Weijters and van der Aalst, 2001], [Weijters and van der Aalst, 2003], and [Weijters et al., 2006].

Let $\delta$ be a real number with $0<\delta<1$. Then, the causal score (w.r.t. $\delta$ ) is defined as the function $\mathbf{c s}_{\delta}: \mathcal{A} \times \mathcal{A} \mapsto \mathbb{R}$ such that $\mathbf{c s}_{\delta}\left(a_{i}, a_{j}\right)=D\left(a_{i}, a_{j}\right) / \mid\left\{t \in L \mid a_{i}=\right.$ $t[k]$, for some index $k\} \mid$, and where:

$$
D\left(a_{i}, a_{j}\right)=\sum_{t \in L \mid a_{i} \text { preceedes } a_{j} \text { in } t} \delta^{\mathbf{c s}^{+}\left(t, a_{i}, a_{j}\right)}-\sum_{t \in L \mid a_{j} \text { preceedes } a_{i} \text { in } t} \delta^{\mathbf{c s}^{-}\left(t, a_{i}, a_{j}\right)}
$$

with

$$
\text { - } \mathbf{c s}^{+}\left(t, a_{i}, a_{j}\right)=\min \left\{k-h-1 \mid a_{i}=t[h] \wedge a_{j}=t[k] \wedge h<k\right\} \text { being the minimum }
$$

number of symbols between an occurrence of $a_{i}$ and a subsequent occurrence of $a_{j}$, and

- $\mathbf{c s}^{-}\left(t, a_{i}, a_{j}\right)=\max \left\{k-h-1 \mid a_{i}=t[k] \wedge a_{j}=t[h] \wedge h<k\right\}$ being the minimum number of symbols between an occurrence of $a_{j}$ and a subsequent occurrence of $a_{i}$.

To illustrate the above definition, note that, for each trace $t[1] \ldots t[n], D\left(a_{i}, a_{j}\right)$ is incremented by a term $\delta^{k-h-1}$ if $a_{i}$ occurs $k-h$ positions before an occurrence of $a_{j}$ (and this is the minimum distance between the symbols in the given order), and decremented by the same term (in absolute value) if $a_{i}$ occurs $k-h$ positions after an occurrence of $a_{j}$ (again, with this distance being the minimum one). Moreover, the positive and the negative terms exponentially decrease at the growing of the distance between $a_{i}$ and $a_{j}$ in the traces. Note that $-1 \leq \mathbf{c s}_{\delta}\left(a_{i}, a_{j}\right) \leq 1$ holds, since $0<\delta<1$.

### 9.6.1 Precedence constraints without negated paths

We start by illustrating an algorithm to solve CN-Mining over $\mathcal{C}[\{\rightarrow, \rightsquigarrow, \nrightarrow\}]$. The algorithm is designed so that, over linear logs (cf. Theorem 9.6), a causal net is always returned, whenever a solution in fact exists. Instead, over arbitrary logs, the algorithm might well return an extended causal net.

Conceptually the algorithm is organized in three main phases:

- First, a dependency graph is built by guaranteeing that each trace in the given log can be supported by properly enriching the graph with suitable bindings. While doing so, we avoid to make use of edges that are forbidden because of negated edge constraints. If this is not possible, then (we shall show that) no solution exists at all.
- Second, we incrementally add to the dependency graph edges that are not forbidden, until all positive edge and path constraints are satisfied. Note that, in principle, all edges that are not forbidden might be added to the dependency graph (and if some constraint is still violated, then we are again guaranteed that no solution exists). Therefore, in order to minimize the size of the resulting model (and the probability of introducing spurious dependencies), edges are selected in the algorithm based on their causal scores.
- Finally, the dependency graph resulting from the above manipulations is equipped with bindings guaranteeing that all the traces can be supported. This leads to a causal net that is returned as output.

These steps are now detailed in the rest of the section.

## Precedence Graphs

The starting point of the algorithm is the construction of the precedence graph $\mathrm{PG}(L, \Pi)$ over $\mathcal{A}(L)$. This graph is built by avoiding the inclusion of the edges that are forbidden because of negated edge constraints in $\Pi$, i.e., the edges in $\mathrm{FE}(\Pi)=$ $\left\{(x, y) \mid x \in S, y \in T, \neg(S \rightarrow T) \in \Pi^{\nrightarrow}\right\}$. We start with an exemplification.

Example 9.24. Consider a $\log L$ including only the trace abcde, and assume that $\Pi=\{\neg(\{c\} \rightarrow\{d\})\}$. The precedence graph $\mathrm{PG}(L, \Pi)$ is illustrated in the left part of Figure 9.7. Intuitively, each activity $x \in\{a, b, c, d, e\}$ has an edge incoming (resp., outgoing) to any activity that precedes (resp., follows) $x$ in some trace. However, we avoid the edges that are forbidden according to the constraints. In particular, the graph does not contain the edge $(c, d)$, and the connectivity of $d$ is guaranteed via the edge $(b, d)$, i.e., $d$ is reached by the node closest to it in the trace abcde and for which no violation in $\Pi^{\nrightarrow}$ occurs. Moreover, c can reach the final activity via the edge (c,e).


Figure 9.7: Precedence graphs for the examples in Section 9.6.
$\triangleleft$

Formally, $\mathrm{PG}(L, \Pi)=(V, E)$ is defined as the directed graph where $V=\mathcal{A}(L)$ and where the set $E$ of its edges is built as follows. For each trace $t \in L$ and for each $i \in\{2, \ldots$, len $(t)\}$ (resp., $i \in\{1, \ldots$, len $(t)-1\}$ ), if there is an index $j \in\{1, \ldots, i-1\}$ (resp., $h \in\{i+1, \ldots$, len $(t)\})$ such that $(t[j], t[i]) \notin \mathrm{FE}(\Pi)$ (resp., $(t[i], t[h]) \notin \mathrm{FE}(\Pi))$, then $E$ contains the edge $\left(t\left[j^{*}\right], t[i]\right)$ (resp., $\left(t[i], t\left[h^{*}\right]\right)$ ) witnessing that the property holds (i.e., $\left(t\left[j^{*}\right], t[i]\right) \notin \mathrm{FE}(\Pi)$ (resp., $\left.\left(t[i], t\left[h^{*}\right]\right) \notin \mathrm{FE}(\Pi)\right)$ ) and having the highest causal score. ${ }^{1}$ No further edge is in $E$.

Lemma 9.25. Assume that $\operatorname{PG}(L, \Pi) \not \vDash \pi(L)$ holds. Then, there is no dependency graph $\mathcal{G}$ such that $\mathcal{G} \models \pi(L) \cup \Pi$. Otherwise, i.e., if $\operatorname{PG}(L, \Pi) \models \pi(L)$, then $\operatorname{PG}(L, \Pi)$ is a dependency graph over $\mathcal{A}(L)$ such that $\operatorname{PG}(L, \Pi) \models \Pi^{\nrightarrow}$.

Proof. Assume that $\operatorname{PG}(L, \Pi) \not \vDash \pi(L)$. Then, there is a trace $t \in L$ such that $\mathrm{PG}(L, \Pi) \not \models \pi(t)$. By Definition 9.5, there are two possible cases. First, there might be an index $i \in\{2, \ldots$, len $(t)\}$ such that $\{t[1], \ldots, t[i-1]\} \rightarrow\{t[i]\}$ is not satisfied by

[^6]$\mathrm{PG}(L, \Pi)$. By construction of $\mathrm{PG}(L, \Pi)$, this means that $(t[j], t[i]) \in \mathrm{FE}(\Pi)$ holds, for each $j \in\{1, \ldots, i-1\}$. Consider then a dependency graph $\mathcal{G}$ such that $\mathcal{G} \models \pi(L)$. Note that $\mathcal{G}$ must include an edge $\left(t\left[j^{*}\right], t[i]\right)$, with $j^{*} \in\{1, \ldots, i-1\}$, in order to satisfy the above constraint. However, $\left(t\left[j^{*}\right], t[i]\right)$ is in $\mathrm{FE}(\Pi)$, and there is a negated edge constraint $\neg(S \rightarrow T) \in \Pi$ such that $t\left[j^{*}\right] \in S$ and $t[i] \in T$. So, if $\mathcal{G} \models \pi(L)$, then $\mathcal{G} \not \vDash \Pi$.

Assume now that $\operatorname{PG}(L, \Pi) \models \pi(L)$. We first observe that $\operatorname{PG}(L, \Pi)$ contains the two activities $a_{\perp}$ and $a_{\top}$, playing the role of the starting and terminating activity. In particular, $a_{\perp}$ and $a_{\top}$ have no ingoing and outgoing edges, respectively. Consider then any other activity $a$ occurring in $\operatorname{PG}(L, \Pi)$, and note that there is a trace $t \in L$ and an index $i \in\{2, \ldots$, len $(t)-1\}$ such that $t[i]=a$. Since $\operatorname{PG}(L, \Pi) \models \pi(L)$, we are then guaranteed about the existence of two edges having the form $(t[j], t[i])$ and $(t[i], t[h])$, with $j \in\{1, \ldots, i-1\}$ and $h \in\{i+1, \ldots$, len $(t)\}$. By structural induction on the index $i$, it then follows that $t[i]$ occurs in a path connecting $t[1]=a_{\perp}$ to $t[\operatorname{len}(t)]=a_{\top}$. Hence, $\mathrm{PG}(L, \Pi)$ satisfies all the conditions for being a dependency graph over $\mathcal{A}(L)$. Then, in order to conclude the proof, we need to show that $\operatorname{PG}(L, \Pi) \models \Pi^{\nrightarrow}$. Indeed, assume by contradiction that a negated edge constraint $\neg(S \rightarrow T)$ exists in $\Pi$ such that $x \in S, y \in T$, and the edge $(x, y)$ is in $\mathrm{PG}(L, \Pi)$. Hence, $(x, y) \in \mathrm{FE}(\Pi)$, which is impossible as all edges of $\operatorname{PG}(L, \Pi)$ do not belong to $\mathrm{FE}(\Pi)$, by construction.

## Positive Precedence Constraints

According to the above result, precedence graphs can be used as a preliminary representation of inter-activity dependencies. However, these graphs do not guarantee that positive constraints are satisfied. Therefore, we define a method for identifying the edges needed to satisfy positive constraints in $\Pi$.

Example 9.26. Consider the precedence graph $\mathrm{PG}(L, \Pi)$ discussed in Example 9.24
and assume that $\Pi$ also contains the constraint $\{d\} \rightsquigarrow\{b\}$. Note that $\operatorname{PG}(L, \Pi)$ does not satisfy the positive constraint $\{d\} \rightsquigarrow\{b\}$. Thus, we have to update the graph by including a path starting from $d$ and terminating into $b$ and where the edge $(c, d)$ does not occur in it. Of course, in this case we can just simply add an edge from d to $b$, as it is shown in the central part of Figure 9.7. $\triangleleft$

Let $\mathcal{G}=(V, E)$ be a dependency graph with $V \supseteq \mathcal{A}(\Pi)$. For each pair of nodes $x, y \in V$, define the weight of $(x, y)$ for $\mathcal{G}$ w.r.t. $\delta$, as the real number $w_{\delta}(\mathcal{G}, x, y)$ such that ${ }^{2}$ :
$w_{\delta}(\mathcal{G}, x, y)= \begin{cases}0 & \text { if }(x, y) \in E \\ +\infty & \text { if } x=a_{\top} ; \text { or } y=a_{\perp} ; \text { or } x \in S, y \in T, \text { and } \neg(S \rightarrow T) \in \Pi \\ 2-\mathbf{c s}_{\delta}(x, y) & \text { in the remaining cases }\end{cases}$
If $a_{1}, \ldots, a_{h}$ is a sequence of nodes forming a path in $\mathcal{G}$, with $h \geq 2$, then we define its weight $w_{\delta}\left(\mathcal{G}, a_{1}, \ldots, a_{h}\right)$ as the value $\sum_{i=1}^{h-1} w_{\delta}\left(\mathcal{G}, a_{i}, a_{i+1}\right)$. Note that $w_{\delta}\left(\mathcal{G}, a_{1}, \ldots, a_{h}\right) \geq$ 0 , because $\mathbf{c s}_{\delta}(x, y) \leq 1$ holds, for each pair $x, y$. Moreover, for each path constraint $S \rightsquigarrow T$ (resp., edge constraint $S \rightarrow T$ ) in $\Pi$, we define $\operatorname{BestPath}_{\delta}(\mathcal{G}, S \rightsquigarrow T)=$ $a_{1}, \ldots, a_{h}$ (resp. BestEdge ${ }_{\delta}(\mathcal{G}, S \rightarrow T)=a_{1}, a_{2}$ ) as the minimum-weight path (resp., minimum-weight edge) such that $a_{1} \in S$ and $a_{h} \in T$. Note that, if there is a precedence constraint $S \rightsquigarrow T$ (resp., $S \rightarrow T$ ) such that the weight of $\operatorname{BestPath}_{\delta}(\mathcal{G}, S \rightsquigarrow T)$ (resp., $\left.\operatorname{BestEdge}_{\delta}(\mathcal{G}, S \rightarrow T)\right)$ is $+\infty$, then all the paths (resp., edges) connecting any activity in $S$ to any activity in $T$ must include an edge that cannot occur in a model of $\Pi^{\nrightarrow}$. Hence, the following is immediately established.

Lemma 9.27. Let $\mathcal{G}$ be a graph such that $\mathcal{G} \models \Pi^{\nrightarrow}$. If there is a path constraint $S \rightsquigarrow T$ (resp., edge constraint $S \rightarrow T)$ such that the weight of $\operatorname{BestPath}_{\delta}(\mathcal{G}, S \rightsquigarrow T)$ (resp., $\operatorname{BestEdge}_{\delta}(\mathcal{G}, S \rightarrow T)$ ) is $+\infty$, then CN-Mining has no solution (on $\mathcal{A}, L$,

[^7]```
Input: A set \mathcal{A of activities, with }\mp@subsup{a}{\perp}{}(\mp@subsup{a}{\top}{})\mathrm{ being the starting (terminating) one,}
            a log}L\mathrm{ with }\mathcal{A}(L)\subseteq\mathcal{A}\mathrm{ , and
            a set }\Pi\in\mathcal{C}[{->,\rightsquigarrow,\not->}}] of precedence constraints with \mathcal{A}(\Pi)\subseteq\mathcal{A}
Parameters: Real numbers 1>\delta>0 and 1\geq\tau\geq0;
Output: A triple }\langle\mathcal{G},\mathcal{I},\mathcal{O}\rangle\mathrm{ , or 'no';
    \Pi:=\Pi\cup{{\mp@subsup{a}{\perp}{}}\rightsquigarrow{a},{a}\rightsquigarrow{\mp@subsup{a}{\top}{}}|a\in\mathcal{A}\\mathcal{A}(L)};
    if PG (L,\Pi) }\vDash=\pi(L)\mathrm{ then return 'no' (and HALT);
    let }\mp@subsup{\mathcal{G}}{1}{}=(\mathcal{A},\mp@subsup{E}{1}{})\mathrm{ be the graph where E E consists of the edges in PG (L, ח), k:=1;
    \mathcal{G}}:==\mp@subsup{\mathcal{G}}{1}{}\ominus{(x,y)|(x,y)\mathrm{ is supported by (at least) }\tau\times|L|\mathrm{ traces }; (*see Footnote 1*)
    \mathcal{F}}:={S->T\in\Pi|\mathcal{G}k\not|=S->T}
    while F}\not=\emptyset\emptyset\mathrm{ do
        let S->T be in \mathcal{F}}\mathrm{ , and let BestEdge }\mp@subsup{\delta}{\delta}{}(\mp@subsup{\mathcal{G}}{k}{},S,T)=\mp@subsup{a}{1}{},\mp@subsup{a}{2}{}\mathrm{ ;
        if w
        \mathcal{G}
        \mathcal{F}}:={S->T\in\Pi|\mathcal{G}k\not|=S->T}
    end while
    \mathcal{F}}:={S->T\in\Pi|\mp@subsup{\mathcal{G}}{k}{}\not\vDash=S\rightsquigarrowT}
    while }\mathcal{F}\not=\emptyset\mathrm{ do
        let S\rightsquigarrowT be in \mathcal{F}}\mathrm{ , and let BestPath}\mp@subsup{\mathcal{B}}{\delta}{}(\mp@subsup{\mathcal{G}}{k}{},S\rightsquigarrowT)=\mp@subsup{a}{1}{},\ldots,\mp@subsup{a}{h}{}
        if }\mp@subsup{w}{\delta}{}(\mp@subsup{\mathcal{G}}{k}{},\mp@subsup{a}{1}{},\ldots,\mp@subsup{a}{h}{})\geq+\infty then return 'no'(and HALT)
        \mathcal{G}
        \mathcal { F } : = \{ S \rightsquigarrow T \in \Pi \| \mathcal { G } k \not \vDash \vDash S \rightsquigarrow T \} ;
    end while
    return computeBindings(\mp@subsup{\mathcal{G}}{k}{},L); (* see Figure 9.2*)
```

Figure 9.8: Algorithm COMPUTE-CN (on $\mathcal{C}[\{\rightarrow, \rightsquigarrow, \nrightarrow\}])$.
and $\Pi$ ).

If the hypothesis in the above lemma does not hold, in order to satisfy a path constraint $S \rightsquigarrow T$ (resp., edge constraint $S \rightarrow T$ ) we can just update the graph $\mathcal{G}$ as to include $\operatorname{BestPath}_{\delta}(\mathcal{G}, S \rightsquigarrow T)\left(\right.$ resp., $\left.\operatorname{BestEdge}_{\delta}(\mathcal{G}, S \rightarrow T)\right)$ as a path (resp., an edge).

## Putting Things Together

Now that we have discussed all the salient ingredients, we can illustrate the algorithm compute- CN shown in Figure 9.8.

The algorithm starts in step 1 by adding to $\Pi$ a set of path constraints stating that each activity $a \in \mathcal{A} \backslash \mathcal{A}(L)$, i.e., not occurring in the log, has still to occur in a path from the starting activity to the terminating one. Step 2 is responsible for checking whether the precedence graph "supports" the log. Note that we directly
check the satisfaction of the constraints induced by the log. Indeed, if this graph does not satisfy the condition, then we report that no solution exists at all.

In step 3 , a graph $\mathcal{G}_{1}$ is initially built over the nodes in $\mathcal{A}$ and the edges in $\mathrm{PG}(L, \Pi)$. In the subsequent steps, $\mathcal{G}_{k}$ denotes the graph obtained from $\mathcal{G}_{1}$ after having performed $k-1$ manipulations on it. In fact, the process starts with step 4 , which is a heuristic step that removes any edge that is not supported by a sufficient number of traces, i.e., by $\tau \times|L|$ traces, where $\tau$ is a threshold received as an additional input parameter. In fact, it can be checked that for $\tau=0$, this is immaterial and the graph remains unchanged. Then, step 5-11 (resp., 12-18) are responsible for adding a number of edges to the precedence graph, as to satisfy all edge constraints (resp., path constraints), according to the strategy described in Section 9.6.1. A failure in this step leads the algorithm to exit while reporting that no solution exists at all.

Finally, a (possibly extended) causal net $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ is built and returned as output. This latter step is carried out by the function in Figure 9.2 that just implements the strategy of including an input (resp., output) binding for each trace $t$, and of defining this binding with all predecessor (resp., successor) activities in $t$. Moreover, in order to formally guarantee that $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ is a (possibly extended) causal net, we add the binding $I_{a}$ (resp., $O_{a}$ ) to $\mathcal{I}(a)$ (resp., $\mathcal{O}(a)$ ), for each activity $a$ in $\mathcal{G}$, being defined as the union of all incoming (resp., outgoing) edges.

The correctness of the whole algorithm is stated next. Note that the result explicitly differentiates the case when $L$ is a linear $\log$ from the case when $L$ is not linear. In the former case we are guaranteed that the output $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ returned by the algorithm is a causal net, whereas in the latter case, $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ might well be an extended causal net. In both cases, however, $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle \vdash L$ and $\mathcal{G} \models \Pi$ hold.

Theorem 9.28. The following properties hold on COMPUTE-CN, receiving as input
$\mathcal{A}, L, \Pi \in \mathcal{C}[\{\rightarrow, \rightsquigarrow, \nrightarrow\}]$, the parameter $\delta$, and for $\tau=0$ :

- if it returns ' $n o$ ', then there is no solution;
- if it returns $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ and $L$ is (resp., is not) a linear log, then $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ is a (resp., a possibly extended) causal net such that $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle \vdash L$ and $\mathcal{G} \models \Pi$.

Proof. By Lemma 9.25 and Theorem 9.6, if the algorithm returns 'no' at step 2, then we are guaranteed that there is no solution. Assume that we are not in this case. Then, by Lemma 9.25, we know that $\mathrm{PG}(L, \Pi)$ is a dependency graph over $\mathcal{A}(L)$ with $\mathrm{PG}(L, \Pi) \models \pi(L)$ and $\mathrm{PG}(L, \Pi) \models \Pi^{\nrightarrow}$. So, $\mathcal{G}_{1}$ is such that $\mathcal{G}_{1} \models \pi(L)$ and $\mathcal{G}_{1} \models \Pi^{\nrightarrow}$.

Consider now the steps $5-18$. Note that whenever the current graph $\mathcal{G}_{k}$ is updated (in steps 9 and 16), we are guaranteed that any edge ( $x, y$ ) that is inserted does not violate any negated edge constraint, for otherwise the weight of BestEdge ${ }_{\delta}(\mathcal{G}, S \rightarrow T)$ or $\operatorname{BestPath}_{\delta}(\mathcal{G}, S \rightsquigarrow T)$ would be $+\infty$. Moreover, in the remaining steps, no edge is added. Therefore, $\mathcal{G}_{k} \models \Pi^{\ngtr}$ holds, for each $k \geq 1$. In fact, steps 6-18 try to enforce the satisfaction of the edge and the path constraints. By Lemma 9.27 and since $\mathcal{G}_{k} \models \Pi^{\nrightarrow}$ holds, for each $k \geq 1$, we derive that if the algorithm returns 'no', then we are guaranteed that no solution exists at all. Again, let us assume that this is not the case, in order to complete the analysis. So, we have now reached step 19, where we are guaranteed that the current graph $\mathcal{G}_{k}$ is such that $\mathcal{G}_{k} \models \Pi$. Moreover, as we have just added edges (with no edge incoming into the starting activity or outgoing from the terminating one) and initially $\mathcal{G}_{1} \models \pi(L)$ holds, then $\mathcal{G}_{k} \models \pi(L)$ holds, too. In particular, the subgraph of $\mathcal{G}_{k}$ induced over the nodes in $\mathcal{A}(L)$ is a dependency graph. Moreover, because of the constraints added in step 1, every other activity in $a \in \mathcal{A} \backslash \mathcal{A}(L)$ is also in a path from the starting to the terminating activity. Hence, $\mathcal{G}_{k}$ is a dependency graph.

Finally, step 19 invokes a function that equips $\mathcal{G}_{k}$ with the sets $\mathcal{I}$ and $\mathcal{O}$. The
reader may check that the function coincides with the constructive implementation stated in Theorem 9.6. Therefore, the tuple $\mathcal{C}=\left\langle\mathcal{G}_{k}, \mathcal{I}, \mathcal{O}\right\rangle$ returned as output is such that $\mathcal{C} \vdash L$ and $\mathcal{G}_{k} \models \Pi \backslash \Pi \nsim$. Hence, if $\Pi \in \mathcal{C}[\{\rightarrow, \rightsquigarrow, \nrightarrow\}]$, then we have actually computed a solution, which is a (standard) causal net when $L$ is linear.

Implementation issues and computation time analysis Let $n_{t}=|L|$ and $n_{a}=|\mathcal{A}|$ be the number of traces and activities in input, respectively, and let $l_{t}$ be the maximal trace length. Let $n^{\rightarrow}$ (resp., $n^{\nrightarrow}, n^{\rightsquigarrow}$ ) be the number of constraints of type $\mathcal{C}[\{\rightarrow$ \}] (resp., $\mathcal{C}[\{\nrightarrow\}], \mathcal{C}[\{\rightsquigarrow\}])$ given as input. Moreover, let $n_{c}$ be the total number of input constraints (independently of the type), and let $k$ be the maximum number of elements in either side of them all, i.e., the maximum size of all their associated sets $S$ and $T$.

In the implementation, constraints are indexed with a number in $\left\{1, \ldots, n_{c}\right\}$, and two arrays of lists of activity identifiers are used to keep trace of the activities appearing in the left and right sides, respectively, of each constraint. Causal scores, forbidden edges, and dependency graphs are all represented via $n_{a} \times n_{a}$ matrices. The initialization of these structures and steps 1-5 can be done in $O\left(n_{t} \times l_{t}^{2}+n_{c} \times k+n_{a}^{2}\right)$ time, where in particular the leftmost term corresponds to computing the causal score matrix and $\operatorname{PG}(L, \Pi)$ plus assessing whether this latter satisfies $\pi(L)$.

The cost of the first loop (steps 6-11) is $O\left(n^{\rightarrow} \times k^{2}\right)$. This accounts for checking the satisfaction of all positive edge constraints, and for the cost of computing the "best edge" for each of them. Since no edge is removed, each constraint is considered at most in one iteration, and will remain satisfied in the subsequent ones. The second loop (steps 13-18) resembles the first one, except for the focus on computing "best paths" rather than "best edges". For each constraint, the task can be carried out via

```
Input: A set \mathcal{A of activities, with }\mp@subsup{a}{\perp}{}(\mp@subsup{a}{\top}{})\mathrm{ being the starting (terminating) one,}
    a }\operatorname{log}L\mathrm{ with }\mathcal{A}(L)\subseteq\mathcal{A}\mathrm{ , and
    a set }\Pi\in\mathcal{C}[{\not~}}]\mathrm{ of precedence constraints with }\mathcal{A}(\Pi)\subseteq\mathcal{A}\mathrm{ ;
Parameters: Real numbers 1>\delta>0, and 1\geq\tau\geq0;
Output: A triple }\langle\mathcal{G},\mathcal{I},\mathcal{O}\rangle\mathrm{ , or 'no';
    \mathcal{G}}:=\textrm{PG}(L,\emptyset)
    \mathcal{G}}:==\mp@subsup{\mathcal{G}}{1}{}\ominus{(x,y)|(x,y) is supported by (at least) \tau\times|L| traces}, k:=1, removed :=\emptyset;
    while FakeEdge(\mp@subsup{\mathcal{G}}{k}{})\not=\emptyset\mathrm{ do}
```



```
        let z,w be any pair of nodes in }\mp@subsup{\mathcal{G}}{k}{}\mathrm{ such that
            (c1) z\in\operatorname{succ}(\mp@subsup{\mathcal{G}}{k}{},\mp@subsup{x}{}{*}),w\in\operatorname{pred}(\mp@subsup{\mathcal{G}}{k}{},\mp@subsup{y}{}{*}),{(\mp@subsup{x}{}{*},z),(w,\mp@subsup{y}{}{*})}\cap\mathrm{ removed }=\emptyset;
            (c2) }\forall\mp@subsup{z}{}{\prime},\mp@subsup{w}{}{\prime}\mathrm{ satisfying (c1), 枚}(\mp@subsup{x}{}{*},z)+\mp@subsup{\mathbf{cs}}{\delta}{}(w,\mp@subsup{y}{}{*})\geq\mp@subsup{\mathbf{cs}}{\delta}{}(\mp@subsup{x}{}{*},\mp@subsup{z}{}{\prime})+\mp@subsup{\mathbf{cs}}{\delta}{}(\mp@subsup{w}{}{\prime},\mp@subsup{y}{}{*})
        \mathcal{G}
    end while
    if \mathcal{G}
    return computeBindings(\mp@subsup{\mathcal{G}}{k}{},L);
```

Figure 9.9: Algorithm COMPUTE-CN (on $\mathcal{C}[\{\nsim \sim\}]$ ).
the classical Djikstra algorithm (possibly provided with an artificial "super"-source node, linked to all activities in the lefthand set of the constraint), with a cost $O\left(n_{a}^{2}\right)$. Since $k \leq n_{a}$ and we have $n^{\leadsto}$ constraints, the overall cost is $O\left(n^{\rightsquigarrow} \times n_{a}^{2}\right)$.

Finally, the input and output bindings of each activity can be stored in two dictionaries, whose entries are multi-sets of activities, with each activity identifying the other vertex of an associated incoming/outgoing edge. Such multi-sets (which reduce to sets in the case of pure causal nets) are simply stored as vectors, encoding edges' multiplicity. Regarding such vectors as strings of length $n_{a}$ (i.e., the sequence of occurrence counts, one per edge), these dictionaries can be implemented as tries (i.e., prefix trees), which can be built in $O\left(n_{t} \times l_{t} \times n_{a}\right)$. This cost accounts for (i) scanning each input trace $s$ in both directions (forward and backward) with an index, say $i$, while incrementally building the multi-set of all activities in its first (resp., last) $i$ positions, and for (ii) generating the resulting multi-set and adding it to the input (resp., output) bindings of $s[i]$. In total, we get $O\left(n_{t} \times l_{t} \times \max \left(l_{t}, n_{a}\right)+\right.$ $\left.n_{a}^{2} \times\left(1+n^{\rightsquigarrow}\right)+n^{\rightarrow} \times k^{2}+n^{\nrightarrow} \times k\right)$.

### 9.6.2 The case of negated path constraints

We now present an algorithm to solve CN-Mining on the class $\mathcal{C}[\{\psi\}$,$] , which is$ illustrated in Figure 9.9. Most of the ingredients discussed so far will still play a role, but the approach is substantially different. The algorithm starts again with the construction of the precedence graph, but this time without taking care of negated edge constraints-as usual, note that edges with low causal scores are removed. That is, we start with the graph $\operatorname{PG}(L, \emptyset)$, whose main properties are stated below and are easily seen to hold by inspecting the proof of Lemma 9.25.

Lemma 9.29. $\mathrm{PG}(L, \emptyset)$ is a dependency graph over $\mathcal{A}(L)$ and $\mathrm{PG}(L, \emptyset) \models \pi(L)$.
The algorithm subsequently breaks any path that witnesses a violation of the negated path constraints in $\Pi$. Formally, if $\mathcal{G}=(V, E)$ is a dependency graph over $\mathcal{A}$, then the set of all fake edges is defined as FakeEdge $(\mathcal{G})=\{(x, y) \in E \mid$ $(x, y)$ occurs in a path from an activity in $S$ to an activity in $T$, with $\neg(S \rightsquigarrow T) \in$ $\Pi\} \backslash\left\{(x, y) \in E \mid x=a_{\perp}\right.$ or $\left.y=a_{\top}\right\}$. Note that edges outgoing from the starting activity and incoming to the terminating one are treated differently. Intuitively, we would like to remove all fake edges from the graph. However, while doing so we might miss the ability of supporting the $\log L$, so that we might need to repair the connectivity.

Let $\mathcal{G}=(V, E)$ be a graph over $\mathcal{A}$, and let $y$ be an activity in $V$. The set $\operatorname{pred}(\mathcal{G}, y)$ of the causal predecessors of $y$ in $L$ is defined as the set of all the activities $w \in V$ such that there is a path from $a_{\perp}$ to $w$ in $\mathcal{G} \backslash\left\{\left(x^{\prime}, y\right) \mid\left(x^{\prime}, y\right) \in E\right\}$ or $a_{\perp}=w$, and for each trace $t \in L$ where $y$ occurs, i.e., $y=t[i]$ for some $i \in\{1, \ldots$, len $(t)\}$, then $w$ also occurs in $t$ before $y$, i.e., $w=t[j]$ where $j<i$ holds. Symmetrically, let $x$ be an activity in $V$. The set $\operatorname{succ}(\mathcal{G}, x)$ of the causal successors of $x$ in $L$ is the set of all the activities
$z \in V$ such that there is a path from $z$ to $a_{\top}$ in $\mathcal{G} \backslash\left\{\left(x, y^{\prime}\right) \mid\left(x, y^{\prime}\right) \in E\right\}$ or $a_{\top}=z$, and for each trace $t \in L$ where $x$ occurs, i.e., $x=t[i]$ for some $i \in\{1, \ldots$, len $(t)\}$, then $z$ also occurs in $t$ after $x$, i.e., $z=t[j]$ where $j>i$ holds. Note that the following is immediate.

Lemma 9.30. Let $\mathcal{G}=(V, E)$ be a dependency graph such that $\mathcal{G} \models \pi(L)$, and let $(x, y)$ be in $E$, hence with $x \neq a_{\top}$ and $y \neq a_{\perp}$. Then, $a_{\perp} \in \operatorname{pred}(\mathcal{G}, y)$ and $a_{\top} \in \operatorname{succ}(\mathcal{G}, x)$.

Proof. Note first that, since $\mathcal{G}$ is a dependency graph, it must be the case that $x \neq a_{\top}$ and $y \neq a_{\perp}$. Moreover, it can be also established that $w \cap\left\{y, a_{\top}\right\}=\emptyset$ and $z \cap\left\{x, a_{\perp}\right\}=\emptyset$, by definition of causal predecessor and successor. Therefore, the graph $\mathcal{G}^{\prime}=\mathcal{G} \oplus\{(w, y),(x, z)\} \ominus\{(x, y)\}$ is such that the starting activity $a_{\perp}$ and the terminating activity $a_{\top}$ have no ingoing and outgoing edges, respectively. Consider now an activity $a \in V \backslash\left\{a_{\perp}, a_{\top}\right\}$. Note that, since $\mathcal{G}$ is a dependency graph, either (i) there is a path in $\mathcal{G} \ominus\{(x, y)\}$ from $a_{\perp}$ to $a$, or (ii) the edge $(x, y)$ occurs in each path in $\mathcal{G}$ from $a_{\perp}$ to $a$. In the case (i), we immediately conclude that there is a path from $a_{\perp}$ to $a$ in $\mathcal{G}^{\prime}$, too. Hence, let us focus on case (ii). Recall that in $\mathcal{G}^{\prime}$ we have the edge $(w, y)$ where $w \neq x$. In particular, there is a path from $a_{\perp}$ to $w$ in $\mathcal{G} \backslash\left\{\left(x^{\prime}, y\right) \mid\left(x^{\prime}, y\right) \in E\right\}$, or $w=a_{\perp}$. Hence, we have derived that there is a path from $a_{\perp}$ to $a$ in $\mathcal{G}^{\prime}$, too. By symmetric arguments, we derive also that there is a path from $a$ to $a_{\mathrm{T}}$ in $\mathcal{G}^{\prime}$. Hence, $\mathcal{G}^{\prime}$ is a dependency graph.

In order to conclude the proof, we have now to show that $\mathcal{G}^{\prime} \models \pi(L)$, too. Indeed, assume by contradiction that a trace $t$ exists in $L$ such that $\mathcal{G}^{\prime}$ does not model $\pi(t)$. Given the differences between $\mathcal{G}$ and $\mathcal{G}^{\prime}$ and since the constraints induced by the traces are only positive ones, it must be the case that the removal of the edge $(x, y)$ is the source of the violation. Formally, we can be in one of the following two
scenarios:
(1) It holds that $y=t[i]$ and $x=t[j]$, with $j<i$, and $\mathcal{G}^{\prime}$ does not satisfy the constraint $\{t[1], \ldots, t[i-1]\} \rightarrow\{y\}$, which is instead satisfied by $\mathcal{G}$ precisely because of the edge $(x, y)$. However, we recall that the edge $(w, y)$ occurs in $\mathcal{G}^{\prime}$ and that $w$ occurs before $y$ in any trace where $y$ occurs. That is, there is an index $j^{\prime} \in$ $\{1, \ldots, i-1\}$ such that $w=t\left[j^{\prime}\right]$ where $j^{\prime}<i$ holds. Hence, $\mathcal{G}^{\prime}$ satisfies the constraint. Contradiction.
(2) It holds that $x=t[i]$ and $y=t[j]$, with $i<j$, and $\mathcal{G}^{\prime}$ does not satisfy the constraint $\{x\} \rightarrow\{t[i+1], \ldots, t[l e n(t)]\}$, which is instead satisfied by $\mathcal{G}$ precisely because of the edge $(x, y)$. However, we recall that the edge $(x, z)$ occurs in $\overline{\mathcal{G}}^{\prime}$ and that $z$ occurs after $x$ in any trace where $x$ occurs. That is, there is an index $j^{\prime} \in\{i+1, \ldots$, len $(t)\}$ such that $z=t\left[j^{\prime}\right]$ where $j^{\prime}>i$ holds. Hence, $\mathcal{G}^{\prime}$ satisfies the constraint. Contradiction.

Therefore, $\mathcal{G}^{\prime} \models \pi(L)$ holds.

Lemma 9.31. Let $\mathcal{G}=(V, E)$ be a dependency graph such that $\mathcal{G} \models \pi(L)$, let $(x, y)$ be in $E$, and let $w \neq x$ and $z \neq y$ be in $\operatorname{pred}(\mathcal{G}, y)$ and $\operatorname{succ}(\mathcal{G}, x)$, respectively. Then, $\mathcal{G}^{\prime}=\mathcal{G} \oplus\{(w, y),(x, z)\} \ominus\{(x, y)\}$ is a dependency graph such that $\mathcal{G}^{\prime} \models \pi(L)$.

Example 9.32. Consider the trace abcde and the constraints $\neg(\{c\} \rightsquigarrow\{d\})$ and $\neg(\{b\} \rightsquigarrow\{d\})$. In this setting, for the dependency graph $\operatorname{PG}(\{a b c d e\}, \emptyset)$, the edges $(c, d)$ and $(b, c)$ are fake ones. The left part of Figure 9.7 evidences how the graph has to be updated when removing the edge $(c, d)$-in fact, this coincides with the graph built when the edge $(c, d)$ is forbidden, as we already discussed in Example 9.24. In the resulting graph, $(b, c)$ is no longer fake. However, the graph does not still satisfy the constraints, and $(b, d)$ is a fake edge. On the right part of Figure 9.7, a further
update is reported accommodating the deletion of $(b, d)$.
The specific strategy adopted to select causal predecessors and causal successors is formalized in the steps $3-7$. Eventually we return the causal net built on top of $\mathcal{G}_{k}$ via the function computeBindings. The correctness of the whole approach is shown below. Note that, similarly to Theorem 9.28 , the result explicitly differentiates the case when $L$ is a linear $\log$ from the case when $L$ is not linear. As usual, in the former case, $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ is a causal net, while in the latter case it might be an extended causal net.

Theorem 9.33. The following properties hold on COMPUTE-CN, receiving as input $\mathcal{A}, L, \Pi \in \mathcal{C}[\{\nsim\}]$, the parameter $\delta$, and for $\tau=0$ :

- if it returns 'no', then there is no solution;
- if it returns $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ and $L$ is (resp., is not) a linear $\log$, then $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle$ is a (resp., possibly extended) causal net such that $\langle\mathcal{G}, \mathcal{I}, \mathcal{O}\rangle \vdash L$ and $\mathcal{G} \models \Pi$.

Proof. Because of Lemma 9.29, we know that $\mathcal{G}_{1}$ is a dependency graph over $\mathcal{A}(L)$ and $\mathcal{G}_{1} \models \pi(L)$. Consider then all the update operations performed in the steps 3-7. Because of Lemma 9.30, these operations are well-defined: for each fake edge detected, we can always find a predecessor and a successor (coinciding, at most, with the starting and terminating activities, respectively) that allow for removing it. Moreover, by Lemma 9.31, the graph $\mathcal{G}_{k}$ is still a dependency graph with $\mathcal{G}_{k} \models \pi(L)$.

Let us now focus on step 8. Note that when all fake edges are removed, the only edges that remain and that can violate a negated path constraints have the form $\left(a_{\perp}, a\right),\left(a, a_{\top}\right)$, or $\left(a_{\perp}, a_{\top}\right)$. However, if there is a constraint preventing the existence of a path from $a_{\perp}$ to $a$, or from $a$ to $a_{\top}$, or from $a_{\perp}$ to $a_{\top}$, then there can be no solution because of the definition of dependency graph. Hence, it is correct
that the algorithm halts there．Otherwise，for the analysis of the last step，recall that the function computeBindings is a constructive implementation of the proof of Theorem 9．6，in order to build a possibly extended causal net from the given graph $\mathcal{G}_{k}$ ．In particular，whenever $L$ is linear，we end up with a causal net．

Implementation issues and computation time analysis In addition to the notation used to analyze the costs of the algorithm of Figure 9．8，let us denote by $n^{\not 九 \iota}$ the number of constraints of type $\mathcal{C}[\{\nsim \psi\}]$ that are taken as input，and by $m^{\nrightarrow \iota}$ the number of（distinct）paths prohibited by them，i．e．，$m^{\nprec}=\mid\{(x, y) \in \mathcal{A} \times \mathcal{A} \mid \exists \neg(S \nLeftarrow$ $T)$ such that $x \in S$ and $y \in T\} \mid$ ．Note that typically $m^{\nrightarrow}$ is much smaller than $k^{2}$ ， where $k$ still denotes the maximum number of elements appearing in either side of any constraint．

The data structures are the same as those discussed for the algorithm of Figure 9．8． In order to speed up the calculation of $\operatorname{pred}\left(\mathcal{G}_{k}, x\right)$ and $\operatorname{succ}\left(\mathcal{G}_{k}, x\right)$ ，for each activity $x$ ，we precompute a relaxed version of both sets，denoted by $\operatorname{pred}(x)$ and $\operatorname{succ}(x)$ ， only accounting for the ordering of activities in the log traces．More precisely，for any activity $y$ ，we have $y \in \operatorname{succ}(x)$（resp．，$y \in \operatorname{pred}(x)$ ）if，and only if，for each trace $t \in L$ where $x$ occurs，then $y$ also occurs in $t$ after（resp．，before）$x$ ．All such sets of（potential）successors and predecessors are stored as boolean vectors，one for each activity．

Initializing the data structures and performing the first two steps in the algorithm takes $O\left(n_{t} \times l_{t}^{2}+n^{\nrightarrow} \times k+n_{a}^{2}\right)$ time．

The loop spanning over steps 3－7 can be iterated $m \nprec \sim$ times at most，seeing as each iteration removes at least one of the fake edges and one of the paths violating some constraint－which are $n_{a}^{2}$ and $m^{\not 九 九}$ ，at most，respectively（with $m^{\nrightarrow \iota} \leq k^{2} \leq n_{a}^{2}$ ）．

The computation of all fake edges, at the beginning of each iteration, can be accomplished in $O\left(n^{\nsucc} \times n_{a}^{2}\right)$ time as follows. For each constraint that is still unsatisfied, two symmetric multiple-source visits of $\mathcal{G}_{k}$ are carried out, starting from all the activities in its left and right sides, respectively; in particular, in the latter case, edges are considered as they were reversed. By reckoning all edges traversed in both directions as fake edges, we compute that with the minimum causal score- $\left(x^{*}, y^{*}\right)$ in the figure.

We can then compute the transitive closure $\mathcal{G}_{k}^{+}$of the current dependency graph via a matrix-multiplication method, in $O\left(n_{a}^{\omega}\right)$ time, and materialize it into a matrix. In our current implementation, based on the famous Strassen's method, it is $\omega=$ 2.8074. Anyway, answering path queries against $\mathcal{G}_{k}^{+}$, in $O\left(n_{a}\right)$ time we can find the nodes $z$ and $w$ mentioned in step 5 . To this end, we can simply select the activity in $\operatorname{succ}\left(x^{*}\right)\left(\right.$ resp., in $\left.\operatorname{pred}\left(y^{*}\right)\right)$ with the maximal score $\mathbf{c s}_{\delta}\left(x^{*}, z\right)$ (resp., $\left.\mathbf{c s}_{\delta}\left(w, y^{*}\right)\right)$, among those satisfying all involved path constraints-involving the existence of a path from $z$ to $a_{\top}$ (resp., from $a_{\perp}$ to $w$ ), which is a required property of causal successors (resp., predecessors). Therefore, the overall cost of the loop in Figure 9.9


A cost of $O\left(n_{a}\right)$ is enough for accomplishing all remaining (edge update) operations in the loop. The same result holds for step 8, where we just need to check whether all constraints are marked as satisfied. As explained previously, a $O\left(n_{t} \times l_{t} \times n_{a}\right)$ cost is needed for computing all bindings (while storing them in a concise form).

In conclusion, we get a total cost of $O\left(n_{t} \times l_{t} \times \max \left(l_{t}, n_{a}\right)+\left(n_{a}^{\omega}+n_{a}^{2} \times n^{\not 九 \dagger}\right) \times m^{\nrightarrow \dagger}\right)$.

### 9.7 Experimental evaluation

The algorithms proposed in this chapter have been implemented as a plug-in for the well-known process mining suite ProM van Dongen et al. [2005]. The plugin receives as input a log file in the (standard) MXML or XES format ${ }^{3}$ plus the constraints that users can define by using an intuitive XML-based specification language. As output, it produces a (possibly extended) causal net that, according to the philosophy of ProM, is made available and can be re-used in the suite for subsequent elaborations. In fact, ProM does not natively support bindings defined as multisets, and therefore the mined models are made available by flattening all multisets into standard sets. This means that, when logs are not linear (cf. Theorem 9.28 and Theorem 9.33), some loss of information might occur when re-using in ProM the models mined with our plug-in. However, the true mined models are always exported into a file encoding dependencies and bindings (again) via an XML-based specification language. The plug-in together with documentation on its usage, plus all datasets illustrated in this section are publicly available at http://staff.icar.cnr.it/wfmining/http/public_html/cnmining/.

Internally, the plug-in combines the computation schemes described in Figure 9.8 and Figure 9.9. In particular, the latter scheme is adopted when user-defined constraints belong to the class $\mathcal{C}[\{\nsim\}\}$, whereas the former has been slightly generalized, ${ }^{4}$ in order to provide a heuristic solution approach in all remaining cases (while still being an exact solution approach over the class $\mathcal{C}[\{\rightarrow, \nrightarrow, \rightsquigarrow\}])$. The generalization consists of a post-processing procedure applied to the discovered causal net returned by step 19 in Figure 9.8. The procedure removes "useless" edges with the intended goal

[^8]| Symbol | Meaning |
| :--- | :--- |
| ILP | The ILP-based mining algorithm defined in van der Werf et al. $[2009]$ |
| AGNEs | The $A G N E s$ mining algorithm in Goedertier et al. $[2009]$ |
| $\alpha$ | The $\alpha$ mining algorithm defined in van der Aalst et al. $[2004]$ |
| HM | The Heuristics mining algorithm defined in Weijters et al. $[2006]$ |
| GM | The Genetics mining algorithm defined in Medeiros et al. $[2007]$ |
| Here | The COMPUTE-CN algorithm defined in Figure 9.8 and Figure 9.9 |

Table 9.1: Process discovery algorithms used in the experiments: legend of symbols.
of returning a more compact model, by increasing the chances of satisfying negated path constraints. Moreover, note that the removal of edges is consistently reflected by suitably updating the associated bindings where they occur.

Formally, the procedure iteratively removes the edge with the lowest causal score over all the edges $(x, y)$ satisfying the following three conditions:
(i) none of the input (resp., output) bindings associated with $y$ (resp., $x$ ) coincides with $\{(x, y)\}$, i.e., the edge does not appear as a singleton binding;
(ii) the removal of $(x, y)$ will not violate any user-defined precedence constraint;
(iii) $\frac{\mathbf{c s}_{\delta}(x, y)}{\min \left\{\mathbf{c s}_{\delta}\left(x, y^{*}\right), \mathbf{c s}_{\delta}\left(x^{*}, y\right)\right\}}<\tau_{r 2 b}$, where $\left(x, y^{*}\right)$ (resp., $\left(x^{*}, y\right)$ ) is the outgoing edge of $x$ (resp., the incoming edge of $y$ ) with the highest causal score. Notice that $\tau_{r 2 b}$ is used here as a relative (lower) threshold to check the strength of any edge $(x, y)$, relatively to those of the best $y$ 's predecessor and of the best $x$ 's successor-it is similar to the "relative to best" threshold proposed in Weijters et al. [2006].

In the tests described in the remainder of the chapter, we fixed the values $\delta=0.9$, $\tau=0.05$, and $\tau_{r 2 b}=0.75$. These values were chosen pragmatically based on a series of tests conducted on a wide range of synthesized data.

The performances of our implementation have been compared with those of the approaches listed in Table 9.7. Note that we considered the approaches proposed by [van der Werf et al., 2009] (ILP) and [Goedertier et al., 2009] (AGNEs), which we have already discussed in Section 8.2 and which can in fact incorporate a-priori knowl-
edge on the existence of activity dependencies. Moreover, we considered some classical discovery approaches, very popular in the Process Mining community: "Heuristics Miner" Weijters et al. [2006] (HM), "Genetic Miner" Medeiros et al. [2007] (GM), and the algorithm $\alpha$ van der Aalst et al. [2004]. The performances of these methods are discussed here just to delineate baselines suitable for assessing the gain that can be obtained by empowering mining methods with the capability to exploit knowledge on the real structure of the process under analysis. In fact, it will come with no surprise that these "traditional" methods are outperformed in presence of background knowledge (that they are not able to exploit). Hence, it is not sensible to consider a wider range of methods of this kind in the analysis that follows.

In order to test the competitors, we exploited their implementations available in the release 6.3 of the ProM framework van Dongen et al. [2005]. In fact, for AGNEs, ILP, and GM the implementations in the version 5.2 were used. Indeed, AGNEs is only available up to this earlier version, while the implementation of ILP in the latest release lacks of the ability of improving the models discovered by expressing parallelism and direct dependency relationships (which is the feature we are interested in). Concerning GM, we just obtained better performances with the version 5.2. Moreover, since an earlier implementation of $\alpha$ in the version 5.2 of ProM also allows users to express relationships of parallelism and direct dependencies, we have included this method too in all the tests performed in presence of background knowledge. Default settings were used for the various methods. For GM, we took the best model derived from an initial population of 100 models by means of 100 iterations (of genetic operations).

In the rest of the section, we shall illustrate results of experimental activity conducted over our method and its competitors. In particular, in Section 9.7.1 we


Figure 9.10: Causal net of the ProductRecall process.
consider an archetypical application scenario, in Section 9.7.2 we illustrate results on benchmark logs, and in Section 9.7.3 we present results on synthesized logs. A usage scenario (not directly related to process mining) is discussed in Section ??. Experiments have been performed on a dedicated machine, equipped with an Intel core i7-3770k 3.5 GHz processor with 12 GB (DDR3 1600 MHz ) of RAM, and running Windows 8 Professional.

### 9.7.1 Case study: a product-recall process

Let us consider the product recall process defined by [Wynn et al., 2009], in accordance to the guidelines established by several public institutions (e.g., in Australia, New Zealand, USA, and EU). This is an archetypical application scenario, and its discussion is meant here to assess the importance of having background knowledge when an incomplete sample of the possible traces is given at hand.

## Testbed Description

The process concerns the main activities that must be performed by a recall sponsor (usually the manufacturer of a suspect product), in response to a recall incident, which can be possibly triggered by consumer complaints, supplier notifications, or failed quality tests. Specifically, the reported problem has to be investigated and a comprehensive risk analysis must be done (macro-activity PROLOGUE-see Wynn et al. [2009] for details on the sequence of activities comprised in it), in order to decide whether the product should be recalled or not. The model associated with the activities occurring in the former case is reported in Figure 9.10: after starting the recall procedure, a case can proceed along a number of concurrent threads, including the following tasks: (i) stopping the distribution of the product, (ii) identifying remedies, (iii) arranging the disposal of items already distributed, (iv) keeping records for subsequent monitoring and analysis purposes, and (v) notifying third parties about the recall. In addition, depending on the kinds of product and of defect involved, it can be necessary to halt the production of the product and to destroy/modify other products that might have been contaminated. Once these recall actions have been completed, a sequence of finalizing activities must be performed (macro-activity EPILOGUE), ranging from monitoring the effectiveness of the process, to implementing changes to prevent similar problems in the future, to preparing reports for regulatory authorities and/or other third parties. Notice that, as pointed out by [Wynn et al., 2009], the need of handling product recall operations, while taking care of traceability and notification issues, arises in a wide variety of real applications.

## Evaluation Setting

In order to assess the quality of findings, we contrast the set $D_{\text {out }}$ of causal dependencies discovered by the mining methods towards the set $D_{i n}$ of the real dependencies existing in the a-priori known process model, by resorting to the classical $F$-measure metric, defined as $\frac{2 \times P \times R}{P+R}$, where $P$ (standing for precision) is the fraction of the dependencies in the mined model that do exist in the real model, i.e., $P=\left|D_{\text {out }} \cap D_{\text {in }}\right| /\left|D_{\text {out }}\right|$, whereas $R$ (standing for Recall) is the fraction of real dependencies captured by the mined model, i.e., $R=\left|D_{\text {out }} \cap D_{\text {in }}\right| /\left|D_{\text {in }}\right|$.

## Test with variable amounts of log traces

First, we generated a log (of 10000 traces) that covers a significant number of all the possible behaviors and over which the various mining methods are able to perfectly reconstruct the process model for this case study. This step was just meant to check that all methods can be in principle effective to face the given scenario. Then, we progressively filtered the log, by therefore increasing its level of incompleteness. We stopped the process with a resulting $\log L$ of about 250 traces (of which about 200 are distinct ones), over which HM (the simplest discovery method) is still able to perfectly reconstruct the underlying model based on them. Experiments have been conducted over logs that are obtained from $L$ by taking $x \%$ of its traces, for $x \in\{10,20, \ldots, 100\}$. In particular, for each $x$, we built 10 samples. Each sample has been built by (uniformly at random) removing one trace, until a sub-log of the desired size of $\frac{x}{100} \times|L|$ is obtained. Experiments are performed on each of them, so that average values (for any given value of $x$ ) will be discussed in our analysis. Table 9.7.1 summarizes our findings.

Let us first consider the columns reporting results for the scenario where no con-

| trace <br> $\%$ | without constraints |  |  |  |  | HM | $\alpha$ | GM | ILP | AGNEs |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Here | $\alpha$ | ILP | AGNEs | Here |  |  |  |  |  |  |
| 10 | 0.657 | 0.733 | 0.553 | 0.848 | 0.674 | $\mathbf{0 . 9 2 9}$ | 0.772 | 0.848 | 0.667 | $\mathbf{0 . 9 5 6}$ |
| 20 | 0.869 | 0.871 | 0.578 | 0.924 | 0.727 | $\mathbf{0 . 9 8 7}$ | 0.899 | 0.924 | 0.736 | $\mathbf{0 . 9 9 2}$ |
| 30 | 0.901 | 0.925 | 0.603 | 0.914 | 0.677 | $\mathbf{0 . 9 8 4}$ | 0.950 | 0.914 | 0.720 | $\mathbf{1 . 0 0 0}$ |
| 40 | 0.927 | 0.944 | 0.658 | 0.914 | 0.720 | $\mathbf{0 . 9 9 2}$ | 0.983 | 0.914 | 0.730 | $\mathbf{1 . 0 0 0}$ |
| 50 | 0.964 | 0.969 | 0.591 | 0.904 | 0.748 | $\mathbf{1 . 0 0 0}$ | 0.968 | 0.904 | 0.727 | $\mathbf{1 . 0 0 0}$ |
| 60 | 0.974 | 0.968 | 0.628 | 0.903 | 0.745 | $\mathbf{1 . 0 0 0}$ | 0.970 | 0.903 | 0.774 | $\mathbf{1 . 0 0 0}$ |
| 70 | 0.979 | 0.990 | 0.584 | 0.882 | 0.774 | $\mathbf{1 . 0 0 0}$ | 0.990 | 0.882 | 0.763 | $\mathbf{1 . 0 0 0}$ |
| 80 | 0.982 | 0.993 | 0.629 | 0.893 | 0.763 | $\mathbf{1 . 0 0 0}$ | 0.993 | 0.893 | 0.779 | $\mathbf{1 . 0 0 0}$ |
| 90 | $\mathbf{1 . 0 0 0}$ | $\mathbf{1 . 0 0 0}$ | 0.652 | 0.882 | 0.763 | $\mathbf{1 . 0 0 0}$ | $\mathbf{1 . 0 0 0}$ | 0.882 | 0.779 | $\mathbf{1 . 0 0 0}$ |
| 100 | $\mathbf{1 . 0 0 0}$ | $\mathbf{1 . 0 0 0}$ | 0.684 | 0.882 | 0.763 | $\mathbf{1 . 0 0 0}$ | $\mathbf{1 . 0 0 0}$ | 0.882 | 0.782 | $\mathbf{1 . 0 0 0}$ |
| Avg | 0.911 | 0.939 | 0.616 | 0.897 | 0.735 | $\mathbf{0 . 9 8 9}$ | 0.952 | 0.897 | 0.752 | $\mathbf{0 . 9 9 5}$ |

Table 9.2:
F-measure scores obtained by algorithm COMPUTE-CN and its competitors, for different percentages of traces of process ProductRecall (Figure 9.10), without (left) and with (right) a-priori knowledge on parallelism relationships. For both cases, maximal scores on each trace percentage are written in bold.
straints are provided as input. It is easily seen that the performances of all methods are rather poor against very small log samples, and tend to improve when augmenting the number of input traces. Such an effect is evident in the case of the classical methods HM and $\alpha$, which get full accuracy when provided with at least $90 \%$ of the log traces. Interestingly, our approach managed to reconstruct the real set of task dependencies already with $50 \%$ samples of the log, hence outperforming the competitors even without being provided as input with additional background knowledge. This confirms the validity of the processing scheme of Figure 9.8 and Figure 9.9 and the efficacy of the post-processing method illustrated in this section, where different heuristics (mainly considered in isolation in earlier approaches) are synergically integrated and exploited within an original computation scheme. In particular, note that in most of the classical approaches, a dependency from an activity $x$ to an activity $y$ can be added to the model only when $x$ is a direct predecessor of $y$ in some trace and the pair $(x, y)$ is associated with a high causal score value. Our algorithms, instead, are more liberal because the definition of precedence graphs (see Section 9.6.1) is
entirely based on causal scores, so that a dependency from $x$ to $y$ can be added even when $x$ never occurs as a direct predecessor of $y$. On the other hand, recall that dependencies are then pruned based on their support in the log with the threshold $\tau$ and by using the "relative to best" threshold proposed in Weijters et al. [2006]. The results in Table 9.7.1 evidence the efficacy of the combined usage of these ingredients. In fact, opposed to these good news related to our method, it emerges that low quality scores (with no evident gain w.r.t. the given percentage of traces) are achieved when using ILP and AGNEs, which are to be seen as our more direct competitors.

The four rightmost columns in Table 9.7.1 report the results obtained by providing domain knowledge to the discovery methods. To this end, we assumed that the analyst knows that the activity Complete_optional_actions is parallel with each of the activities in \{Identify_remedies, Keep_records, Stop_distribution, Arrange_disposal, Notify_third_parties\}. Note that this knowledge can be incorporated in the two competitors AGNEs, and ILP, as well as in the $\alpha$ algorithm, so that we are not exploiting for the moment the richer expressiveness of our framework. By looking at the table, it emerges that our method is able to fully exploit these constraints, by getting impressive accuracy results even on very small samples. Instead, the benefits of the background knowledge are lower for $\alpha$ and especially for AGNEs. Moreover, note that ILP is unable to exploit at all the given knowledge. In fact, in our experiments, we have noticed that the method is quite often not capable to benefit from knowledge about parallelism, while significant improvements can be obtained in presence of knowledge about relationships of precedences, as we shall see later.


Figure 9.11:
F-measure scores obtained by COMPUTE-CN when varying the percentages of positive edge/path constraints and of negative edge/path constraints, for four different families of $\log$ samples, corresponding to $3 \%$ (a) and $6 \%$ (b) of traces in the $\log L$, respectively.

## Varying the quantity and type of background knowledge

In order to assess, in a deeper and more systematic manner, the capability of our approach to exploit a-priori knowledge, further tests were carried out on the same scenario while using different amounts of precedence constraints. The focus is to understand the impact of each of the types of precedence constraints on the quality of the mining results. Therefore, only very small portions of the log were used, containing $3 \%$ and $6 \%$ of the traces of the $\log L$ (of about 250 traces). For each trace percentage $x \%$ (with $x \in\{3,6\}$ ), four samples were generated (as discussed in Section 9.7.1) and, as usual, average results are discussed.

Different inter-activity dependencies were extracted directly from the known model of the process, shown in Figure 9.10. Regarding this model as a dependency graph, four relations over its activities can be defined and computed trivially, each encoding some basic kind of (singleton body/head) precedence constraints: edges and paths (i.e., pairs of activities, where the second one depends on the first either directly or indirectly, respectively), and the associated complementary (w.r.t. all possible activ-
ity pairs) relations of negative edges and negative paths. These sets are denoted in the figures discussed below by E, P, NE, and NP, respectively. And, they consists of 19, 46, 150, and 123 different constraints, respectively. Then, in order to automatically generate different sets of precedence constraints, while controlling the relative amount of each type of them, a sample of elements is extracted randomly from each of the core relations discussed above, by using some given percentage value for each of them. In fact, five samples were generated and results were averaged over them.

Figure 9.11 reports the average F-measure scores obtained by our solution approach against the different percentages of log traces, while using one of the above described kinds of pairwise constraints per time. For each kind of constraints, different amounts were considered, ranging from 0 to $50 \%$ of its whole population. It clearly emerges that the use of the background knowledge improves the performances of the algorithm. Indeed, higher F-measure scores are achieved when increasing the amount of whichever kind of constraints. However, a noticeably boost seems to be given to the accuracy when increasing the percentage of negative (edge or path) constraints, no matter of how many traces are taken as input. Conversely, lower improvements are obtained when only positive constraints are used, especially when these are expressed on paths. Intuitively, this is due to the fact that negative constraints are able to reduce the number of spurious flows of execution, which negatively impact on the quality of the resulting process model. This finding was confirmed over all our tests. Hence, as a practical guideline, the users of our plug-in are encouraged to introduce as much as possible negated constraints in the specification of the mining problem.

Note that Figure 9.11 evidences that the level of improvement is neatly higher when working on smaller log samples (hardly capturing all actual process behaviors).

Moreover, note that even in absence of constraints and with very few traces at hand (3\%), our method is capable of achieving good performances in reconstructing the underlying model. This is in line with our findings discussed in Section 9.7.1.

## Rate of unsatisfied constraints

So far our algorithm has been tested in scenarios where precedence constraints of different kinds are not mixed together. Therefore, according to the results discussed in Section 9.6, it provides an exact solution in these cases, i.e., no constraint can be violated by the resulting process model. Hence, in order to study the efficacy of the method as a heuristic, we performed an additional series of experiments with heterogeneous combinations of precedence constraints, mixing up negative path constraints with other kinds of constraints. To this end, we applied our algorithm to the same log samples built as in the previous subsection, while providing it with variable amounts of negative path constraints, and fixing the percentage of any other kind of constraints to $25 \%$. Note that, given the adaptation discussed at the beginning of this section, the algorithm returns a process model that still satisfies all positive edge constraints, path constraints, and negated edge constraints. However, the satisfaction of negated path constraints is just greedily enforced by the post-processing phase, and it is therefore not guaranteed. In fact, we computed the rate of negated path constraints left unsatisfied by our approach, in correspondence of different amounts of distinct traces and of negated path constraints provided as input (expressed as percentages w.r.t. the size of their respective populations). In the worst case, just $0.8 \%$ of all negated path constraints given as input were not satisfied, in the average. This rate shrinked of about a half when using either $10 \%$ samples or $50 \%$ samples of negative paths. A similar trend was observed for trace samples with a $6 \%$ of distinct process traces. The amount of violated constraints became negligible when bigger


Figure 9.12:
Computation time spent by algorithm COMPUTE-CN with different amounts of traces and constraints' percentages in input. A base-2 logarithmic scale is used for the vertical axis in both figures, as well as for the horizontal axis in the left-hand figure.
samples of $\log$ traces (we tested both $12 \%$ or $24 \%$ ), no matter how many negated paths are passed to the algorithm.

## Running time

As we have pointed out in the Introduction, it often happens in real process mining applications that only a "small" set of (representative) traces are available for analysis purposes. For this reason, scalability plays generally a secondary role in the design of process mining algorithms, and many works just omit at all this kind of analysis. Nonetheless, we feel that it is relevant to complete the picture by providing the reader with a general idea of the scaling of our algorithms, hence complementing the theoretical analysis carried out in Section 9.6. To this end, we performed a series of tests on logs of different sizes. Logs were generated again with a variable level of completeness, by extracting different amounts of distinct traces out of the original log. First of all, we randomly extracted a sample of (distinct) traces out of the complete log, to produce different logs with size $s \in\{1000,4000,16000,64000\}$. In particular, when traces have to be added to obtain a log bigger than the original
one, we simply duplicated the original traces in a balanced way (i.e., all traces are replicated approximatively the same number of times). ${ }^{5}$

Figure 9.12 reports the computation time (measured in milliseconds) when varying the amounts of traces and constraints in input. In particular, on the left, it shows results grouped according to the different kinds of constraints we have considered (i.e., E, P, NE, NP) and averaged over the samples built when considering $p \%$, with $p \in\{0,10,20,30,40,50\}$, of constraints derived from the population of the given kind (within the same setting as in Section 9.7.1).

On the right, results are reported where, for each particular percentage $p \%$ of constraints, all the sets of different constraints built for $p$ have been merged together. For each value of $p$, five different heterogeneous sets of constraints were generated, containing $p \%$ constraints for each constraint type, extracted at random from their respective populations. Four trials were performed over different logs built as discussed above, and as usual average measures are reported. Notably, whatever percentage is chosen, the computation time scales basically linearly with respect to the number of input traces. On the other hand, a positive correlation seems to exist as well between the overall running time and the quantity of constraints taken as input-if ignoring the case of 1000 -sized logs, where the times measured are too low to safely infer any general significant trend of behavior. Anyway, the impact of constraints on times is negligible if compared with that of the log size.

Finally, we notice that the computation time of algorithm COMPUTE-CN is comparable to that of the standard process mining methods $\alpha$, ILP, and HM, and neatly lower than those spent by GM, AGNEs. Indeed, the latter method took about 600 times longer than COMPUTE-CN to compute a model, in the average, while the

[^9]| log <br> name | distinct <br> activities | pairs of <br> activities | control-flow <br> constructs | distinct <br> traces | log <br> size |
| :--- | :---: | :---: | :---: | :---: | :---: |
| parallel5 | 10 | 10 | - | 107 | 3000 |
| a10skip | 12 | 1 | skip | 6 | 2665 |
| a12 | 14 | 2 | - | 5 | 2492 |
| a5 | 7 | 1 | loop | 13 | 2189 |
| a6nfc | 8 | 1 | nf-choice | 3 | 2040 |
| a7 | 9 | 4 | - | 14 | 2023 |
| a8 | 10 | 1 | - | 4 | 1803 |
| choice | 12 | 0 | - | 16 | 2400 |
| driversLicense | 9 | 0 | - | 2 | 2100 |
| herbstFig3p4 | 12 | 3 | loop | 32 | 5637 |
| herbstFig6p18 | 7 | 0 | loop | 153 | 9844 |
| herbstFig6p36 | 12 | 0 | nf-choice | 2 | 3000 |
| herbstFig6p37 | 16 | 36 | - | 132 | 4800 |
| herbstFig6p41 | 16 | 4 | - | 12 | 3600 |
| herbstFig6p45 | 8 | 4 | - | 12 | 2400 |
| l2l | 6 | 0 | loop | 10 | 4932 |
| l2lOptional | 6 | 0 | loop, skip | 9 | 2622 |
| l2lSkip | 6 | 0 | loop | 8 | 4554 |

Table 9.3:
Benchmark logs: structural characteristics and statistics (see also Weerdt et al. [2012]). Each log consists of 300 (not necessarily distinct) traces.
computation time of GM was about 4500 times that of COMPUTE-CN when using the default population size of 1000 (the ratio only decreased to $600: 1$ with 100 -model populations).

### 9.7.2 Comparative analysis on benchmark data

In order to assess the capability of our approach to discover a good-quality process model in a wider range of settings, we performed a series of tests on some benchmark logs, while measuring the accuracy of each discovered model by way of several "logconformance" metrics, very popular in the fields of Process Mining and Business Process Analysis. Differently from the pure (edge-oriented) F-measure employed in the previous section, to contrast a discovered model to the true (a-priori known) one, these metrics allow for evaluating how much the behaviors registered in a given log comply with those allowed by the model under analysis.

## Testbed: logs and conformance metrics

Our experimental activities were carried out over some of the benchmark logs provided with the ProM framework van Dongen et al. [2005], which have been widely used in the literature (see, e.g., Medeiros et al. [2007]; Goedertier et al. [2009];

De Weerdt et al. [2011]; Weerdt et al. [2012]) in order to evaluate process mining approaches - this benchmark is also congenial to our ends, since the underlying "true" models are available, and can be used both for comparison purposes and for automatically extracting background knowledge. In particular, owing to our special interest toward incompleteness issues, we focused on those logs exhibiting the highest degree of non-determinism and concurrency. The logs were generated from models including a number of routing constructs/patterns, such as (non free) choices, skips, and loops-the reader interested in expanding on these concepts is referred to the above literature. Table 9.7.2 summarizes their features, by reporting in particular the number of distinct activities composing the process, the number of pairs of activities belonging to different parallel branches ("parallel || pairs"), the presence of special constructs, the number of traces, and the log size measured as the number of episodes stored in it, i.e., as the sum of the occurrences of the activities over all traces. Note that six logs derive from process models that contains loops, and in fact they are not linear logs as they contain traces with multiple occurrences of the same activity.

In order to assess the capability of a process model to accurately capture the behavior recorded in a given log, several alternative conformance metrics have been proposed in the literature. In our analysis, given the focus of our approach, we consider precision metrics and recall metrics. Precision metrics attempt to estimate the amount of the "extra" (unseen and likely unwanted) behavior allowed by the model, with respect to that actually registered in the log, whereas recall metrics try to evaluate how much of the behavior recorded in a $\log$ is really captured by the model. All these metrics range over the real interval $[0,1]$ and have been defined in the literature for Petri-net models. In order to use them with a model represented

| Metric | True | AGNEs | $\alpha$ | GM | HM | ILP | Here |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Fitness Rozinat and van der Aalst [2008] | 1.000 | 0.995 | 0.988 | $\mathbf{1 . 0 0 0}$ | 0.994 | $\mathbf{1 . 0 0 0}$ | 0.994 |
| Alignment Based Fitness Adriansyah et al. [2011] | 1.000 | $\mathbf{0 . 9 6 8}$ | 0.830 | $\mathbf{0 . 9 7 8}$ | $\mathbf{0 . 9 5 8}$ | 0.889 | $\underline{\mathbf{0 . 9 9 5}}$ |
| Behavioral Recall Goedertier et al. [2009] | 0.997 | $\mathbf{0 . 9 9 2}$ | $\mathbf{0 . 9 7 8}$ | $\mathbf{0 . 9 8 9}$ | 0.846 | $\mathbf{1 . 0 0 0}$ | $\mathbf{0 . 9 8 9}$ |
| Proper Completion Rozinat and van der Aalst [2008] | 1.000 | 0.938 | 0.851 | $\mathbf{0 . 9 9 9}$ | $\mathbf{0 . 9 3 3}$ | $\underline{\mathbf{1 . 0 0 0}}$ | 0.931 |
| Adv. Behav. Appropriateness Rozinat and van der Aalst [2008] | 0.814 | $\mathbf{0 . 8 2 3}$ | $\mathbf{0 . 8 5 4}$ | 0.802 | 0.783 | $\underline{\mathbf{0 . 8 5 6}}$ | 0.809 |
| Alignment Based Precision van der Aalst et al. $[2012]$ | 0.921 | $\mathbf{0 . 9 2 6}$ | $\mathbf{0 . 9 4 8}$ | 0.874 | 0.871 | $\mathbf{0 . 9 1 2}$ | $\mathbf{0 . 9 3 9}$ |
| Behavioral Specificity Goedertier et al. [2009] | 1.000 | $\mathbf{0 . 9 9 2}$ | 0.978 | $\mathbf{0 . 9 9 7}$ | 0.990 | $\mathbf{1 . 0 0 0}$ | 0.986 |
| (Wtd) Behavioral Precision De Weerdt et al. $[2011]$ | 0.915 | $\mathbf{0 . 8 8 4}$ | 0.882 | $\mathbf{0 . 8 9 0}$ | 0.748 | $\underline{\mathbf{0 . 8 9 1}}$ | 0.887 |
| Negative Event Precision Goedertier et al. [2009] | 0.970 | 0.927 | 0.927 | $\underline{\mathbf{0 . 9 5 3}}$ | $\mathbf{0 . 9 4 3}$ | 0.934 | $\mathbf{0 . 9 3 7}$ |

Table 9.4:
Average conformance measures obtained, on benchmark logs, by different discovery methods-including the one proposed in this thesis (Here). For each row, the best average score (excluding ground-truth models) is underlined, while the results that were recognized as significantly better than the average of the outcomes over the various methods (for the same metrics and setting) are in bold.
in another language (for instance, a causal net), we preliminary translated the given model into a Petri net, with the help of suitable conversion plug-ins available in the ProM framework. The actual computation of the metrics was carried out by taking advantage of the CoBeFra tool, recently proposed van den Broucke et al. [2013] as a practical support to conformance analysis. In fact, we noticed that slight different measures can be obtained over different runs of the tool (but deviations are very limited), so that hereinafter we refer to measures averaged over three runs. A summary of all the considered metrics is reported in the first column of Table 9.7.2.

## Results

For each of the conformance metrics described above, Table 9.7.2 reports the average value obtained by applying each of the methods to each of the logs in Table 9.7.2. As a term of comparison, column True reports the conformance results obtained for the process models that were actually used to generate the logs (which are all available). Since many methods got very similar results over several metrics, a statistical testing procedure was carried out, in order to check whether their behaviors are really different. To this end, for each of the considered metrics, a paired one-tail Student's test was applied to compare the outcomes of each method with the average of those obtained by all the methods. Notably, for each of the metrics, we did
not find significantly enough differences, apart from a small number of cases, which are emphasized in bold in the table - in almost all cases we could not reject (with a $95 \%$ level of confidence) the null hypothesis that a method behaves identically to the "average" one, rather than improving the latter. In fact, since all methods achieve good performances and differences are not statistically relevant, it is not informative to devote further space to analytically illustrate these performances over each of the logs. Accordingly, we just focus below on illustrating Table 9.7.2, where a synthectic picture is depicted by averaging the results over the whole family of benchmark logs.
"Traditional" approaches are able to reconstruct the originating models with high levels of precision and recall (this is hardly surprising given that these methods have been designed precisely to face the challenges in such kinds of benchmarks logs). However, AGNEs, ILP, and our method also exhibit good performances over these logs, and in some cases they perform better than the competitors even though no background knowledge is available. In any case, we observe that our evaluation procedure tends to disadvantage a model that was not originally built as a Petri net (as in our case), because the conversion plug-ins often produce results with hidden/duplicated transitions, which are considered as a source of extra-log behaviors by certain metrics. ${ }^{6}$

As a further comment, we stress here that, in terms of computation time, our approach is in line with the most efficient process discovery techniques, and definitely faster than both AGNEs and GM. The average values and standard deviations computed over running times (in milliseconds) are: $21851.2 \pm 12334.6$ for AGNEs, $341012.2 \pm$ 385997.7 for GM, $95.6 \pm 73.4$ for HM, $203.1 \pm 53.4$ for ILP, and $138.2 \pm 85.3$ for our approach.

[^10]| original <br> benchmark log | (\%) of traces in <br> the sub-log |
| :--- | :---: |
| herbstFig3p4 | $41 \%$ |
| herbstFig6p37 | $4 \%$ |
| herbstFig6p41 | $23 \%$ |
| herbstig6p45 | $18 \%$ |
| parallel5 | $3 \%$ |

Table 9.5:
Statistics on the "critical" sub-log extracted for each of the benchmark logs in Table 9.7.2.

## Tests on (critical) sublogs

In order to test all discovery methods in a more challenging setting, we extracted a critical sample out of each log, defined as an incomplete collection of traces, covering a limited portion of the whole variety of the behaviors in the original log. An empirical iterative procedure was devised to this purpose, where increasing amounts of traces are randomly removed from a given benchmark log, until we register a loss of $15 \%$ in the F-measure of the models discovered with algorithm HM. In order to further emphasize the effect of $\log$ incompleteness, we focused our attention on a subset of the logs in Table 9.7.2, which allowed all the tested discovery methods to perfectly rediscover the associated model (i.e., to achieve a maximal value of 1 over both F-measure and all of the recall metrics). The number of traces of each of these sub-logs is reported in the Table 9.7.2.

Given the incompleteness of the resulting logs, experiments were conducted also in the presence of background knowledge. In this case, for each critical sub-log and for each of the tested methods, we considered the background knowledge consisting of the three constraints extracted from the given true model that allow us to obtain the best performances w.r.t. the F-measure. Note that here we exploit the full expressiveness of our approach. In fact, it emerged that the best results are obtained in presence of negative constraints only.

Results for these tests are reported in Table 9.7.2, whose precision and recall scores

|  | with constraints |  |  |  |  |  |  |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Metric | AGNEs | $\alpha$ | GM |  |  |  | HM | ILP | Here | $\alpha$ | AGNEs |
| ILP | Here |  |  |  |  |  |  |  |  |  |  |
| Fitness | 0.900 | 0.852 | 0.950 | 0.882 | 0.941 | $\underline{0.951}$ | 0.865 | 0.950 | 0.969 | $\mathbf{1 . 0 0 0}$ |  |
| Behavioral Recall | 0.903 | 0.792 | $\underline{0.948}$ | 0.881 | 0.945 | 0.939 | 0.806 | 0.955 | 0.972 | $\mathbf{1 . 0 0 0}$ |  |
| Alignment Based Fitness | 0.858 | 0.744 | $\underline{0.948}$ | 0.834 | $\mathbf{0 . 9 4 5}$ | $\underline{0.948}$ | 0.799 | 0.919 | 0.572 | $\underline{1.000}$ |  |
| Proper Completion | 0.081 | 0.102 | 0.523 | 0.081 | 0.459 | $\underline{\mathbf{0 . 5 4 3}}$ | 0.340 | 0.400 | 0.655 | $\underline{\mathbf{1 . 0 0 0}}$ |  |
| Adv. Behav. Appropriateness | 0.625 | 0.711 | 0.901 | 0.580 | 0.828 | $\underline{0.910}$ | 0.788 | 0.579 | 0.880 | $\underline{0.928}$ |  |
| Alignment Based Precision | 0.952 | 0.957 | $\underline{0.966}$ | 0.956 | 0.963 | 0.965 | 0.967 | 0.966 | $\underline{0.973}$ | 0.972 |  |
| Behavioral Specificity | 0.903 | 0.792 | 0.948 | 0.881 | $\mathbf{0 . 9 4 5}$ | $\underline{0.949}$ | 0.806 | 0.955 | 0.972 | $\mathbf{1 . 0 0 0}$ |  |
| (Wtd) Behavioral Precision | 0.634 | 0.701 | 0.794 | 0.615 | $\underline{\mathbf{0 . 7 9 4}}$ | 0.776 | 0.764 | 0.843 | 0.788 | $\mathbf{0 . 9 6 6}$ |  |
| Negative Event Precision | 0.678 | 0.746 | 0.827 | 0.658 | $\underline{0.833}$ | 0.822 | 0.815 | 0.872 | 0.816 | $\underline{\mathbf{0 . 9 7 6}}$ |  |
| F-measure | 0.741 | 0.730 | $\underline{0.855}$ | 0.735 | 0.831 | 0.852 | 0.806 | 0.941 | 0.915 | $\underline{\mathbf{1 . 0 0 0}}$ |  |

Table 9.6:
Results on "critical" samples without and with background knowledge. For each row, the best average score is underlined, while the results that were recognized as significantly better than the average of the outcomes over the various methods (for the same metrics and setting) are in bold.
were computed by comparing each model discovered (from the sub-logs described above) with the respective (entire) benchmark log. Moreover, the F-measure values are also reported (see Section 9.7.1). As above, a one-tail Student's test was carried out to assess the significance of the performance differences. The results obtained in absence of background knowledge further confirm the validity of our approach and, in particular, the effectiveness of its underlying (causal-score driven) heuristics for pruning useless/unlikely edges. The advantage of using additional background knowledge clearly emerges. In particular, our method perfectly reconstructs the true models of all benchmark logs (see the lowermost row), and achieves full recall without incurring into overgeneralization.

### 9.7.3 Further tests on synthesized data

Another series of tests was conducted on different synthesized logs, in order to assess the effectiveness of the proposed approach against logs following different data distributions, as far as concerns the structure of the processes producing them.

## Generation of process models

The data used in this experimentation were produced with the help of a $\log$ generator (based upon the one described in Burattin and Sperduti [2011]), which
allows for both constructing a process model randomly, and generating a log of random execution traces out of it. The synthesis of the model is carried out via an incremental block-oriented procedure, where an initially empty model is extended iteratively by adding a new subprocess, until a given number of elementary activities is obtained. Each subprocess can be either an elementary activity, a sequence of activities, or a more complex control-flow block. The whole procedure is governed by a combination of gaussian and multinomial probability distributions, and it can be controlled by way of a number of parameters, which are listed below:

- Task Number (TN): the number of elementary activities to be generated;
- Branching Factor $(B F)$ : the number of branches in any branching structure;
- Singleton Probability $(S P)$ : probability that any current subprocess will be instantiated with just one activity;
- Fork Probability $(F P)$ : probability that the subprocess is a "fork" structure (i.e., a number of subprocesses that can be executed in parallel or that are mutually exclusive), given that it is not an elementary activity;
- AND Probability (AP): probability that the branches of any fork structure are in parallel, rather than in mutual exclusion.

In order to prevent models from having an excessively unbalanced shape, one further parameter can be set, named the Maximum Nesting Depth (MND), stating the maximum level of nesting of any activity with respect to all the control-flow blocks it is enclosed within-precisely, whenever the nesting level of a subprocess equals $M N D$, the subprocess is forced to take the form of a single activity (i.e., it cannot give rise to a complex control-flow block).


Figure 9.13:
Results on synthesized log data, with different degrees of parallelism (AP): accuracy of discovered process models (in terms of F-measure w.r.t. real activity dependencies), and rates of unsatisfied (negative path) constraints. Both measures are reported for different amounts of a-priori constraints of all types (expressed as percentages w.r.t. the sizes of their respective population).

In our experiments, we fixed $S P=0.1$ and $M N D=4$, and simulated three different levels of concurrency, hinged on different values of the parameter $A P \in$ $\{0.0,0.5,1.0\}$. In order to consider a wide enough variety of process schemas, the remaining parameters were chosen in a random manner. Specifically, we picked up $T N$ from a truncated normal distribution ranging between 10 and 30 , with location 25 and scale 8 . Instead, $B F$ and $F P$ were extracted from two uniform distributions ranging over the sets $\{2,4,6\}$ and $\{10,20, \ldots, 80\}$, respectively. Overall, we generated 25 different schemas, and produced 4 logs for each of them, by simulating 1000 random enactments of the schema for each of them (as in Burattin and Sperduti [2011]). In this way, 100 different logs were eventually obtained, each one consisting of 1000 traces.

## Results

Figure 9.13 reports the results of tests performed by applying our approach to the logs described above, while providing it with heterogenous sets of precedence
constraints, and hence expressing partial information on real activity relationships. As in Section 9.7.1, these constraints were randomly extracted from the model which generated the log, according to different levels of coverage (w.r.t. the entire population of constraints that can be derived from the model), expressed in terms of a percentage $p \in\{0,10,20, \ldots, 50\}$. In more details, for each $\log L$ and each value of $p$, we randomly generated 5 sets of constraints, by sampling $p \%$ pair from each of the binary relations (of the form $\mathrm{E}, \mathrm{P}, \mathrm{NE}, \mathrm{NP}$ ) extracted from the a-priori known model of $L$. Note that the binary relations E, P, NE, and NP vary depending on the specific model from which they have been extracted. As discussed above, we have considered 25 different schemas. Their sizes, averaged over the 25 different scenarios, are of 24 , 140, 440, and 324 elements, respectively. The maximum (resp., minimum) sizes over the scenarios are $32,251,752$, and 533 (resp., 14, 48, 130, and 94), respectively.

Each value was computed by averaging all the results obtained when using (i) any $\log$ generated with the given value of $A P$, and (ii) any constraint set covering $p \%$ of each (positive/negative direct/indirect) precedence relation. Clearly, higher levels of concurrency (i.e., higher values of $A P$ ) tend to lead to lower accuracy results, as far as concerns the recognition of real activity dependencies. Providing increasing amounts of constraints definitely helps to improve the average accuracy of the resulting models.

As a further effectiveness indicator (connected with the capability of our approach to really satisfy the constraints it was provided with), we finally present, in the right side of Figure 9.13, the rate of unsatisfied negative path constraints, i.e., the percentage of constraints of this kind that were taken as input in some run of the algorithm, but were not fulfilled by the causal net discovered in the same run-recall that all other constraints are necessarily satisfied. Again, these results are reported for different values of both $p$ and $A P$. It is good news that the rates of unsatisfied
negative path constraints are always relatively small, especially in the case of low concurrency levels. In fact, all the models that were induced from concurrency-free logs do not violate any constraint at all. Higher violation rates affect the models extracted from choice-free logs $(A P=1)$, i.e., logs produced by a model where all the branches going out of a fork node are always executed concurrently. However, even in such an extreme case, the rate is quite low, reaching at most $15 \%$, precisely when $p \%=10 \%$. Therefore, the absolute number of constraints violated is still very small, and only represents a $0.15 \%$ of all possible constraints associated with the underlying models. Moreover, note that increasing the amount of input constraints seems to help our approach reduce violation rates. This can be explained by observing that, as the availability of more background knowledge makes each discovered model more similar to the true one, it is more likely that the former will eventually satisfy a higher amount of all the a-priori constraints associated with the latter.

### 9.8 Summary

Current research is rather active in proposing process mining techniques supporting increasingly expressive modeling languages. However, most of the approaches proposed in the literature mine the causal dependencies that hold over the activities, while completely ignoring the prior knowledge that in many cases is available to the analyst. This dissertation moves a systematic step to fill the gap, by proposing and analyzing a constraint-based framework for process mining, where background knowledge can be encoded in the form of precedence constraints. The computational complexity of the framework has been studied, and the whole approach has been implemented in a prototype system, which has been tested on different log data.

Note that the use of constraints is currently gaining attention in a different con-
text, yet still related to process management, namely in the context of developing declarative approaches to support business automation. In fact, traditional modeling languages, such as Petri nets or causal nets, are procedural ones for they explicitly represent all the allowed behavior of the process, according to a "closed world" assumption. Opposed to this approach, recent research focused on developing declarative models where any possible enactment is allowed unless a constraint (expressed in some suitable formal logic) is known to hold and explicitly forbids it (see, e.g., van der Aalst et al. [2009]; Sadiq et al. [2005]; Reichert et al. [2009]). Following a number of earlier attempts to use logic-based languages for the specification of business processes (see, e.g., Bonner [1999]; Senkul et al. [2002]; Dourish et al. [1996]; Joeris [2000]; Wainer and de Lima Bezerra [2003]; Lu et al. [2006]; Attie et al. [1996]), Declare Pesic et al. [2009] is nowadays the most solid platform adopting a declarative perspective for process modeling, where the semantics of each constraint is provided in terms of an associated Linear Temporal Logic (short: LTL) formula, whose (finite) models are precisely the set of all the allowed traces Pesic et al. [2007, 2010].

Interestingly, the problem of automatically inferring a process model from a given log available at hand has been considered within these declarative frameworks for process management, too. In particular, techniques specifically designed to infer Declare constraints have been presented by [Maggi et al., 2011] and [Maggi et al., 2012] and subsequently enhanced by [Maggi et al., 2013], in order to discover dataaware models. Another approach has been proposed by [Di Ciccio and Mecella, 2012], where constraints are eventually expressed in terms of regular expressions rather than in terms of LTL formulas. By sharing the spirit of the above proposals, but focusing on a slightly different problem, [Chesani et al., 2009] proposed a methodology to analyze a log whose traces are labeled as compliant or non-compliant, and whose goal
is to learn a classification model defined as a set of rules/constraints expressed in the SCIFF Alberti et al. [2008] language (eventually mapped in the Declare notation). Other approaches to building rule-based classification models have been proposed, for instance, by [Bellodi et al., 2010] and [Ferreira and Ferreira, 2006].

By looking at the above body of literature, it clearly emerges that the use of constraints in these works is completely different from ours. Indeed, in these declarative frameworks, constraints are used to specify which traces are allowed, while in our approach they help define the set of those possible process models that are of interest to the analyst. In fact, our perspective is the traditional one where the analyst wants to end up with a procedural model, and where constraints are used to prune the search space of all the models that could be induced from a given log. Accordingly, we adopted a language to specify topological constraints on the process model, rather than a language (such as LTL) tailored to define constraints on traces.

Even though the two approaches are completely orthogonal, we stress however that the research reported in this thesis might have also an impact in the context of developing mining approaches for declarative process models. Indeed, in this setting, the idea of incorporating a-priori knowledge in the mining phase has been largely unexplored, and constitutes a promising avenue of further research. Currently, analysts are provided with the rules induced via mining methods and might refine them to incorporate their knowledge in a post-processing phase. Instead, it would be interesting to define approaches where analysts can a-priori formalize their knowledge and where rule/constraints adhering with it are automatically inferred from the available $\log$.

Finally, we recall from Section 8.3 that logs have been viewed in the chapter as multi-set of traces, by disregarding the data involved over the activities. Therefore,
it is natural to look for extensions of the proposed mining techniques that could deal with context information (about, e.g., parameters and functional features of the activities) and process ontologies. Moreover, we observe that we have implicitly considered throughout the chapter a setting such that, whenever an activity is executed, then the corresponding event is registered in the log. In some cases, however, we might have to deal with hidden activities, i.e., activities that exist in the process model but that are not registered in the log. This is very often the case for activities corresponding to routing constructs, which therefore give rise to logs that are "incomplete" in a sense that is completely orthogonal to the one we have considered in this thesis. Indeed, algorithms tailored to deal with hidden activities have been also proposed in the literature, such as the $\alpha \#$ algorithm Wen et al. [2010], which however assumes (as usual) that the log registers all the possible traces for the underlying process. In our experiments, we have considered $\alpha \#$ too, and we observed that it is outperformed by the standard $\alpha$ algorithm - this is not surprising given that our experimental setting does not consider models with hidden activities (where $\alpha \#$ might be more effective). In fact, extending our approach with techniques inspired by $\alpha \#$ and conducting experimentations over process models with hidden activities are further interesting avenues of research.

## CHAPTER 10

## An Application Scenario (Beyond Process Mining): Urban Congestion

### 10.1 Introduction

So far, we have analyzed our algorithms within standard process mining settings, but it is not difficult to envisage that they have wider applicability. In this chapter, we complete the picture by discussing a different kind of application motivated by the TETRis research project, ${ }^{1}$ funded by the Italian Ministry of Education, University and Research and within the general frame of reference "Internet of Things" supporting Smart City/Smart Territory, in which the acquisition of data by objects can be applied to large territorial areas by the widespread availability of communication networks [European Commission, 2012], [http://www.internet- of-thingsresearch.eu/documents.htm]. In more detail, the activities described in this chapter are related to the identification of scenarios for the realization of innovative services aimed at an intelligent management of a urban and suburban territory.

Within these activities innovative solutions and technology platforms have been identified in order to enable a new way of working for the business entities of the area of interest (municipalities, provinces, regions, universities, etc.) and for citizens and operators involved. Infact, its extended and enhanced the wealth of information

[^11]acquired on the territory through the use of sensors and devices interconnected by local and remote communications systems that support high added value services to improve the quality of the territory itself in terms of livability and sustainability through the involvement of citizens who become the main tutors of the territory, the so-called "social sensors", for the detection of critical situations related to the context in which they live.

### 10.2 Project goals

TETRis project main objective is to create high added value services within Smart City and Smart Territory context [Kanter and Litow, 2009], [Komninos et al, 2011] also extending, wherever possible, the functions of TETRA communication system and acting along three main axes: The evolution and the opening of the application fields of TETRA communication system in order to define new information services for operators, exploiting new models and open source tools for the interconnection of TETRA with other networks and the identification of new type of devices obtained through TETRA system interoperability with existing sensors and sensor networks; The modeling and prototyping of an Open Source framework that allows to define a Smart Objects cooperation model and a Smart Objects management within the related Smart Environments; The identification of scenarios and application models in the perspective of Smart City/Smart territory services through the definition of Smart Environments and high added value Open Source services applied to territory monitoring, emergency management, urban and suburban mobility and services to citizens.

The specific objectives of TETRis project can be read as follows: Bring economic and social benefits to the community through more targeted and effective actions by

Public Administration and Public Security operators in different application scenarios such as emergency management, environmental protection, mobility and services to citizens, with the contribution of the same citizens through the sharing of information and the use of innovative tools for social networking; Extend the pervasiveness and effectiveness of public administration services, instrumental bodies, local police, health operators, transport companies in the reference areas; Improving the quality of life and the sense of safety of citizens through the spreading of safe and reliable technology infrastructures "always on".

### 10.3 Urban log mining

Within this project, the goal of our application is to discover knowledge about congestion dynamics in a urban transportation network. The analysis was performed on data extracted from a large collection of raw mobility records, concerning the circulation of buses (basically, over non-preferential roads/streets) across a network of two neighboring Italian cities (Cosenza and Rende) during year 2012. Each of these records registers the fact that a certain bus (following a certain route) occupied a given geographical position (identified via GPS coordinates) at a certain time.

The analysis of congestion dynamics was carried out by preliminary converting such raw data in the form of a process log, as explained next. First, we transformed the original data into a collection of bus moves, each indicating the transfer of a bus from one bus stop to the subsequent one. For any move $m$, let $\operatorname{pos}_{\text {from }}(m)$ and $\operatorname{pos}_{t o}(m)$ denote its initial and final positions (i.e., bus stops), respectively, let time $(m)$ be the time when the move was completed, and let duration $(m)$ be the travel time of $m$. The set of all the pairs $\left\langle\operatorname{pos}_{f r o m}(m), \operatorname{pos}_{t o}(m)\right\rangle$ over all the moves $m$ constitutes the alphabet of the symbols for the log we are going to construct. Each
pair is also called a route segment. For each bus, the sequence of route segments associated with the movements of the bus in a given day constitutes one trace in a base version of the log.

The base version is subsequently filtered by removing from each trace the segments that are not associated with an abnormally long travel time. In particular, a move $m$ is considered as an anomalous event if $\operatorname{duration}(m)$ is greater than 1.3 times the average duration of the moves over the same segment. Note that while removing the moves that are not anomalous, it might happen that two consecutive moves $m_{1}$ and $m_{2}$ in a trace result to be completely unrelated, in that either $\operatorname{pos}_{\text {from }}\left(m_{2}\right) \neq$ $\operatorname{pos}_{t o}\left(m_{1}\right)$, or time $\left(m_{2}\right)-\operatorname{duration}\left(m_{2}\right)-\operatorname{time}\left(m_{1}\right)>\Delta_{t}$, with $\Delta_{t}$ being a threshold fixed to 15 minutes in our experiments. In these cases, the trace is broken in two sub-traces, one ending with $m_{1}$ and the other starting with $m_{2}$, and the analysis is repeated until a fixed point is reached. Eventually, two additional dummy activities are placed at the beginning and at the end of each trace.

Starting with a setting with about 300000 moves and 3000 segments, by exploiting the procedure described above, we obtained a non-linear log over 1743 activities and 2883 traces. Note that the log is focused on the segments over which abnormal travel times have been registered (which witness some traffic congestion) and, in particular, each trace is meant to report the propagation of the congestion across the transportation network. In fact, by analyzing the resulting log with a process mining method, it is now easily seen that one can derive a (conceptual) process model representing how congestions "flow" over the network. In fact, the approach is close in the spirit to the work by [Liu et al., 2011]. The main difference is that in our setting the resulting flow model is not limited to a tree as in Liu et al. [2011] (which basically covers only branching constructs modeling scenarios where a congestion


Figure 10.1:
A screenshot of the plug-in CNMining implementing the algorithms presented in the paper. The causal net is drawn by exploiting the Flex interface available in ProM, where bindings are displayed inside tooltips activated by hovering the mouse over the corresponding node.
is split over two streets), but it can be an arbitrary process model with different routing constructs (in particular, synchronizing constructs modeling scenarios where congestions with different origins are merged together).

Figure 10.1 shows (portions of) one causal net discovered from the events recorded during April and May 2012, within a specific geographical area with diameter of about 8 Km - in this case, we have 108 activities and 52 traces. Each node in the model of Figure 10.1 is labeled with the names of the two bus stops it identifies. According to a process-oriented interpretation of anomaly traces, an edge in the model indicates that a delay in one route segment (the origin of the edge) is likely to propagate to another segment (the ending point). Therefore, in order to be adherent to this semantics, there cannot be an edge in the model from a segment to a subsequent adjacent segment along the same (oriented) route. Similarly, there cannot be an edge from a segment to a non-adjacent one. For these reasons, it comes with no surprise that standard approaches to process discovery might be not appropriate in
this application scenario, as they are not able to properly accommodate this kind of background knowledge.

Instead, after they were equipped with this knowledge, our algorithms produced models that truly reflect the semantics of the congestions. Moreover, discovered models emerged to highlight the problems inherent to the network. Actually, in this case, there is no "true" congestion model to be used to quantify the quality of the findings and to conduct experiments on the measures reported in Table 9.7.2. On the other hand, the setting is congenial to stress the scalability of the algorithms w.r.t. the number of the activities (which is large differently from classical applications of process mining). Indeed, recall from Section 9.6 that a factor in the scaling is $n_{t} \times$ $n_{a} \times l_{t}$, where $n_{t}$ is the number of traces, $n_{a}$ the number of activities, and $l_{t}$ the trace length. While in general $l_{t}$ might linearly grow with $n_{a},{ }^{2}$ the scaling we registered (when varying the diameter of the area) was actually linear w.r.t. $n_{a}$, too. The reason is that even over the whole original log, $90 \%$ of the traces have length 6 at most, so that $l_{t}$ acts as a fixed constant. This comes with no surprise, as we are considering causal correlations over anomalies.

[^12]
## Part IV

## Conclusions

## CHAPTER 11

## Conclusion and Future Research

### 11.1 Contributions

In this dissertation we studied Constraint satisfaction as a general framework that allows us to describe many computational problems arising in various areas of research, such as database theory, game theory and process mining. We look for classes of instances where such problems are tractable (islands of tractability), as well as for heuristics that are able to deal with real-world instances, possibly exploiting suitable properties of the instance structure. We summarize next the main contributions discussed in the dissertation.

### 11.1.1 Databases

In this area of research, we studied the problem of counting solutions in SQL queries specifying "COUNT" aggregates. We proposed a novel structural decomposition notion that extends the framework of greedy tree projections by adding weights. An experimental evaluation of the implementation is conducted, where the algorithm is contrasted with the available algorithms for computing hypertree decompositions (because there are no further algorithms dealing with the general framework). We also designed optimization algorithms which exploit our novel structural notion to identify classes of queries that are structurally tractable. We integrated our hybrid
optimizer inside a well-known relational DBMS and we compared the performances of our approach with those of the traditional optimizer. The results evidenced the effectiveness of our tool.

### 11.1.2 Game theory

We study coalitional game theory, which provides a solid mathematical framework to analyze scenarios where agents con obtain higher payoffs by collaborating with each other rather than by acting in isolation. In particular, we consider the class of allocation games, which has been originally defined by [Moulin, 1992] as a way to analyze fair division problems where monetary compensations are allowed and utilities are quasi-linear. In these contexts (and when monetary transfers are possible), the prototypical solution concept for them is the Shapley value, whose computation is known to be intractable, formally \#P-complete. Motivated by this bad news, the dissertation investigates structural requirements that can be used to isolate islands of tractability. We study parameters of the underlying allocation scenarios which allow us to single out islands of tractability. In more detail, we show that computing the Shapley value, as well as the related concept of Banzhaf value, is feasible in polynomial time when each good is of interest for at most two agents, or when interactions among agents have a tree-like structure. We observe that these games capture scenarios of practical interest. For instance, we analyzed the data of one concrete instantiation for the setting described by [Greco and Scarcello, 2014b]. The setting refers to an allocation problem arising in the Italian Research Assessment program. The instantiation refers to the data of the University of Calabria, by discovering that the treewidth of the associated graph, consisting of more than 500 nodes, is 9 only. Moreover, the result has its own theoretical interest, as the analysis of the complexity of reasoning problems (even specifically related to coalitional games) over structures
of bounded treewidth is an active topic of research in artificial intelligence. Note that the tractability result has been established by encoding the problem of computing the Shapley value of allocation games $\mathcal{G}_{A}$ in terms of counting problems over suitably defined CSP instances [Dechter, 2003], and by using known tractability results for CSP instances having bounded treewidth. Indeed, based on the notion of treewidth, islands of tractability for counting problems have been recently singled out [Pichler and Skritek, 2013]. Our results evidence that subtle issues come into play when trying to reuse such recent results. In fact, while being focused on the Shapley value, the methodology we propose has a wider spectrum of applicability, and its salient ideas might be profitably reused in other contexts. Part of this work appeared as: Gianluigi Greco, Francesco Lupia and Francesco Scarcello Structural Tractability of Shapley and Banzhaf Values in Allocation Games in the Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-15).

### 11.1.3 Process mining

While studying process mining applications, it emerged from our analysis that topdown design methods and bottom-up process discovery techniques have been almost separate worlds, so far. Indeed, while the use of background knowledge for improving the quality of results has already been considered in a number of traditional data mining tasks (on relational data and on sequence data) including pattern mining and clustering (see, e.g., [Guns et al., 2011; de Amo and Furtado, 2007]), designing counter-parts of such techniques in the context of process discovery is still largely an open issue. In this thesis, we move a further step to synergically integrate top-down design methods and bottom-up process discovery techniques. Indeed, we propose a "hybrid" approach to process discovery, where a mining method is conceived which

|  | Traditional | Goedertier et al. [2009] | van der Werf et al. [2009] | Here |
| :--- | :---: | :---: | :---: | :---: |
| edge constraints | $\square$ | $\square$ | $\square$ | $\square$ |
| path constraints | $\square$ | $\square$ | $\square$ | $\square$ |
| constraints for parallelism | $\square$ | $\square$ | $\square$ | $\square$ |
| constraints over sets | $\square$ | $\square$ | $\square$ | $\square$ |

Table 11.1: Summary of constraint support.
can take into account a wide variety of constraints over the causal dependencies that are possibly available to the analyst. This is particularly useful for circumventing the problems emerging when log completeness does not hold. In more detail,
(1) We propose a formal framework to specify additional properties on the process models that can be produced as output by process discovery algorithms. The framework is based on defining a set of precedence constraints over the activities, and it supports the kind of prior knowledge that is usually available to the analyst. Table 11.1.3 summarizes the features of our framework, by comparing them with those supported (to some extent) by earlier approaches in the literature. Note that "traditional" process discovery methods do not provide any support to deal with prior knowledge.
(2) In the light of our formulation, process discovery can be conceptually viewed as a mining task (i.e., building all possible models for a given input log) followed by a reasoning task (i.e., filtering out those models that do not satisfy the precedence constraints defined by the analyst). However, exponentially many process models might be built in general as a result of the mining phase, hence a literal implementation of such a two-phase approach is unfeasible. In fact, we identify relevant classes of constraints where the two tasks can be addressed synergically, and we propose efficient algorithms to deal with them.
(3) We analyze the computational complexity of the proposed setting, by taking into account various qualitative properties regarding the kinds of constraints being al-
lowed, and by tracing the tractability frontier w.r.t. them. In particular, we show that for the classes of constraints that are not covered by the algorithms discussed in the point (2) above, an efficient solution algorithm is unlikely to exist at all, because process discovery turns out to be NP-hard over them.
(4) All the algorithms discussed in the thesis have been implemented and integrated in a prototype system, which is made available as a plug-in for the well-known process mining suite ProM van Dongen et al. [2005]. In particular, to face the above intractability results, our efficient solution algorithms (originally conceived for special cases only) are generalized by making them applicable as heuristic solution approaches for arbitrary classes of constraints. Case studies are illustrated. Results for the experimental activity we have conducted in order to validate the effectiveness of the proposed approach are also reported.

Part of this work has appeared as: Gianluigi Greco, Antonella Guzzo, Francesco Lupia and Luigi Pontieri. Process Discovery under Precedence Constraints. ACM Transactions on Knowledge Discovery from Data, 9(4): 32 (2015).

### 11.2 Future research directions

The arguments discussed in this dissertation give us natural directions for future research. In general, our idea is that is possible to develop new hybrid solvers that can effectively use structural properties to address the problem of computing all solutions to CSPs efficiently. In more detail, the most important direction for future research and applications to the database setting, is the one that integrates structural approaches with new developments in join processing, see, e.g., [Ngo et al., 2014].

In the field of game theory, and in particular in coalitional games, a future research of direction consists in extending our structural tractability results to other solutions
concept such as the nucleolus.
Finally in process mining, an important direction for future research consists in extending the framework of precedence constraints in the setting of block-structured process models.

## Bibliography

A. Atserias, M. Grohe, and D. Marx. Size bounds and query plans for relational joins. SIAM Journal on Computing, 42(4):1737-1767, 2013.
G. Bagan, A. Durand, and E. Grandjean. On acyclic conjunctive queries and constant delay enumeration. In Proc. of $C S L^{\prime} 07$, pp. 208-222, 2007.
A. Bulatov, V. Dalmau, M. Grohe, and D. Marx. Enumerating Homomorphisms. Journal of Computer and System Sciences, 78(2): 638-650, 2012.

Isolde Adler. Tree-related widths of graphs and hypergraphs. SIAM J. Discret. Math., 22(1):102-123, February 2008.

Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. Journal of the ACM, 54(1), Article 1, 2007.
G. Greco and F. Scarcello. Structural tractability of enumerating CSP solutions. Constraints, 18(1):8-74, 2013.

Hung Q Ngo, Christopher Ré, and Atri Rudra. Skew strikes back: New developments in the theory of join algorithms. SIGMOD Rec., 42(4):5-16, February 2014.

Francesco Scarcello, Gianluigi Greco, and Nicola Leone. Weighted hypertree decompositions and optimal query plans. J. Comput. Syst. Sci., 73(3):475-506, 2007.
R. Dechter and J. Pearl. Network-based heuristics for constraint-satisfaction problems. Artif. Intell., 34(1):1-38, December 1987.

Georg Gottlob, Nicola Leone, and Francesco Scarcello. A comparison of structural CSP decomposition methods. Artificial Intelligence, 124(2):243-282, 2000.
G. Greco and F. Scarcello. Fair division rules for funds distribution: The case of the italian research assessment program (vqr 2004-2010). Intelligenza artificiale, 7 (1):45-56, 2013.
G. Greco and F. Scarcello. The power of tree projections: Local consistency, greedy algorithms, and larger islands of tractability. In Proceedings of the Twenty-ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS '10, pages 327-338, New York, NY, USA, 2010. ACM.

Tractable hypergraph properties for constraint satisfaction and conjunctive queries. In Proceedings of the Forty-second ACM Symposium on Theory of Computing, STOC '10, pages 735-744, New York, NY, USA, 2010. ACM.
H.L. Bodlaender and F.V. Fomin. A Linear-Time Algorithm for Finding Tree Decompositions of Small Treewidth. SIAM Journal on Computing, 25(6), pp. 1305-1317, 1996.

Georg Gottlob, Nicola Leone, and Francesco Scarcello. Hypertree decompositions and tractable queries. Journal of Computer and System Sciences, 64(3):579-627, 2002.

Isolde Adler, Georg Gottlob, and Martin Grohe. Hypertree width and related hypergraph invariants. European Journal of Combinatorics, 28(8):2167-2181, 2007.

EuroComb 05 - Combinatorics, Graph Theory and ApplicationsEuroComb 05 Combinatorics, Graph Theory and Applications.

Georg Gottlob, Zoltán Miklós, and Thomas Schwentick. Generalized hypertree decompositions: Np-hardness and tractable variants. J. ACM, 56(6):30:1-30:32, September 2009.
P.A. Bernstein and N. Goodman. The power of natural semijoins. SIAM Journal on Computing, 10(4), pp. 751-771, 1981.
R. Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. Journal of the ACM, 30(3):514-550, 1983.

Leonid Libkin. Elements Of Finite Model Theory (Texts in Theoretical Computer Science. An Eatcs Series). SpringerVerlag, 2004.

Surajit Chaudhuri. An overview of query optimization in relational systems. In Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, PODS '98, pages 34-43, New York, NY, USA, 1998. ACM.
E. F. Codd. A relational model of data for large shared data banks. Commun. ACM, 13(6):377-387, June 1970.
S. Abiteboul, R. Hull, and V. Vianu. Foundations of Databases. Addison-Wesley, 1995.
J. D. Ullman. Principles of Database and Knowledge Base Systems. Computer Science Press, 1989.

Phokion G. Kolaitis and Moshe Y. Vardi. Conjunctive-query containment and constraint satisfaction. J. Comput. Syst. Sci., 61(2):302-332, October 2000.

Marc Gyssens, Peter G. Jeavons, and David A. Cohen. Decomposing constraint satisfaction problems using database techniques. Artif. Intell., 66(1):57-89, March 1994.
H. Aziz and B. de Keijzer. Shapley meets shapley. In Proc. of STACS'14, pp. 99-111.
H. Aziz, O. Lachish, M. Paterson, and R. Savani. Power indices in spanning connectivity games. In Proc. of AAIM'09, pp. 55-67.
Y. Bachrach and Y. S. Rosenschein. Power in threshold network flow games. pages 106-132, 2009.
J.F. Banzhaf. Weighted Voting Doesn't Work: A Mathematical Analysis. Rutgers Law Rev., 19:317-343, 1965.
E. H. Bareiss. Sylvester's Identity and Multistep Integer-Preserving Gaussian Elimination. Mathematics of Computation, 22:565-565, 1968.
C. J. Colbourn, J.S. Provan, and D. Vertigan. The complexity of computing the tutte polynomial on transversal matroids. Combinatorica, 15(1):1-10, 1995.
P. Dagum and M. Luby. Approximating the permanent of graphs with large factors. Theor. Comput. Sci., 102(2):283-305, 1992.

Xiaotie Deng and Christos H. Papadimitriou. On the complexity of cooperative solution concepts. Mathematics of Operations Research, 19:257-266, May 1994.
G. Gottlob, G. Greco, and F. Scarcello. Treewidth and hypertree width. In Tractability: Practical Approaches to Hard Problems. Cambridge University Press, 2013.
G. Gottlob, S. Tien Lee, G. Valiant, and P. Valiant. Size and Treewidth Bounds for Conjunctive Queries. Journal of the ACM, 59(3), 2012.
G. Greco and F. Scarcello. Counting solutions to conjunctive queries: structural and hybrid tractability. In Proc. of PODS'14, pp. 132-143.
G. Greco and F. Scarcello. Mechanisms for fair allocation problems: No-punishment payment rules in verifiable settings. Journal of Artificial Intelligence Research, 49:403-449, 2014.
F. Maniquet. A characterization of the Shapley value in queueing problems. Journal of Economic Theory, 109(1):90-103, 2003.
D. Mishra and B. Rangarajan. Cost sharing in a job scheduling problem. Social Choice and Welfare, 29(3):369-382, 2007.
H. Moulin. An application of the Shapley value to fair division with money. Econometrica, 60(6):1331-49, 1992.
H. Nagamochi, D.-Z. Zeng, N. Kabutoya, and T. Ibaraki. Complexity of the minimum base game on matroids. Mathematics of Operations Research, 22:146-164, 1997.
N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. Algorithmic Game Theory. Cambridge University Press, Cambridge, UK, 2007.
M. J. Osborne and A. Rubinstein. A Course in Game Theory. The MIT Press, Cambridge, MA, USA, 1994.
R. Pichler and S. Skritek. Tractable counting of the answers to conjunctive queries. Journal of Computer and System Sciences, 79(6):984-1001, 2013.
G. Pesant. Counting solutions of CSPs: a structural approach. In Proc. of IJCAI'05, pp. 260-265, 2005.
N. Robertson and P.D. Seymour. Graph minors iii: Planar tree-width. Journal of Combinatorial Theory, Series B, 36(1):49-64, 1984.
D. Schmeidler. The nucleolus of a characteristic function game. SIAM Journal of Applied Mathematics, 17(6):1163-1170, 1969.
L. S. Shapley. A value for n-person games. Contributions to the theory of games, 2:307-317, 1953.

Arya Adriansyah, Boudewijn F. van Dongen, and Wil M. P. van der Aalst. 2011. Conformance Checking Using Cost-Based Fitness Analysis. In Proc. of EDOC'11. 55-64.

Rakesh Agrawal, Dimitrios Gunopulos, Frank Leymann. 1998. Mining process models from workflow logs. In Proc. of EDBT'98. 469-483.

Marco Alberti, Federico Chesani, Marco Gavanelli, Evelina Lamma, Paola Mello, and Paolo Torroni. 2008. Verifiable agent interaction in abductive logic programming: The SCIFF framework. ACM Transactions on Computational Logic 9 (2008), 29:129:43. Issue 4.

Nikolas Anastasiou, Tzu-Ching Horng, and William Knottenbelt. 2011. Deriving generalised stochastic Petri net performance models from high-precision location tracking data. In Proc. of Value Tools'11. 91-100.

Krzysztof Apt. 2003. Principles of Constraint Programming. Cambridge University Press, New York, NY, USA, 2003.

Paul C. Attie, Munindar P. Singh, E. Allen Emerson, Amit P. Sheth, and Marek Rusinkiewicz. 1996. Scheduling workflows by enforcing intertask dependencies. Distributed Systems Engineering 3, (1996), 222-238.

Elena Bellodi, Fabrizio Riguzzi, and Evelina Lamma. 2010. Probabilistic Declarative Process Mining. In Proc. of KSEM'10. 292-303.

Anthony J. Bonner. 1999. Workflow, Transactions, and Datalog. In Proc. of PODS'99. 294-305.
A.A. Bulatov, M. Dyer, L.A. Goldberg, M. Jerrum, and C. Mcquillan. The expressibility of functions on the boolean domain, with applications to counting CSPs. Journal of the ACM, 60(5), Article 32, 2013.
A.A. Bulatov. The complexity of the counting constraint satisfaction problem. Journal of the ACM, 60(5), Article 34, 2013.

Andrea Burattin and Alessandro Sperduti. 2011. PLG: A Framework for the Generation of Business Process Models and Their Execution Logs. In BPM Workshops. 214-219.

Kevin C. W. Chen, David Y. Y. Yun. 2003. Discovering Process Models from Execution History by Graph Matching. In Proc. of IDEAL'03. 887-892.

Federico Chesani, Evelina Lamma, Paola Mello, Marco Montali, Fabrizio Riguzzi, and Sergio Storari. 2009. Exploiting Inductive Logic Programming Techniques for Declarative Process Mining. Transactions on Petri Nets and Other Models of Concurrency 2 (2009), 278-295.

Sandra de Amo and Daniel A. Furtado. 2007. First-order temporal pattern mining with regular expression constraints. Data EJ Knowledge Engineering 62 (2007), 401-420. Issue 3.

Ana Karla Alves de Medeiros, Boudewijn F. van Dongen, Wil M. P. van der Aalst, and A. J. M. M. Weijters. 2004. Process Mining: Extending the $\alpha$-algorithm to Mine Short Loops. Technical Report. University of Technology, Eindhoven.

Ana Karla Alves de Medeiros, A.J.M.M. Weijters, and Wil M. P. van der Aalst. 2007. Genetic process mining: an experimental evaluation. Data Mining and Knowledge Discovery 14 (2007), 245-304.

Luc De Raedt, Tias Guns, and Siegfried Nijssen. 2008. Constraint programming for itemset mining. In Proc. of $K D D^{\prime} 08.204-212$.

Jochen De Weerdt, Manu De Backer, Jan Vanthienen and Bart Baesens. 2011. A robust F-measure for evaluating discovered process models. In Proc. of CIDM'11. 148-155.

Jochen De Weerdt, Manu De Backer, Jan Vanthienen, and Bart Baesens. 2012. A multi-dimensional quality assessment of state-of-the-art process discovery algorithms using real-life event logs. Information Systems 37, 7 (2012), 654-676.

Rina Dechter. 1992. Constraint networks. Encyclopedia of Artificial Intelligence (1992), 276-285.

Claudio Di Ciccio and Massimo Mecella. 2012. Mining Constraints for Artful Processes. In Proc. of BIS'12. 11-23.

Paul Dourish, Jim Holmes, Allan MacLean, Pernille Marqvardsen, and Alex Zbyslaw. 1996. Freeflow: Mediating Between Representation and Action in Workflow Systems. In Proc of CSCW'96. 190-198.

Hugo M. Ferreira and Diogo R. Ferreira. 2006. An Integrated Life Cycle for Workflow Management Based on Learning and Planning. International Journal on Cooperative Information Systems 15, 4 (2006), 485-505.

Michael R. Garey and David S. Johnson. 1979. Computers and Intractability. A Guide to the Theory of NP-completeness. Freeman and Comp., NY, USA.

Stijn Goedertier, David Martens, Jan Vanthienen, and Bart Baesens. 2009. Robust Process Discovery with Artificial Negative Events. Journal of Machine Learning Research 10 (2009), 1305-1340.

Gianluigi Greco, Antonella Guzzo, and Luigi Pontieri. 2007. An InformationTheoretic Framework for Process Structure and Data Mining. International Journal of Data Warehousing and Mining 3, 4 (2007), 99-119.

Gianluigi Greco, Antonella Guzzo, and Luigi Pontieri. 2012. Process Discovery via Precedence Constraints. In Proc. of ECAI'12. 366-371.

Gianluigi Greco, Antonella Guzzo, Luigi Pontieri, and Domenico Saccà. 2006. Discovering Expressive Process Models by Clustering Log Traces. IEEE Transactions on Knowledge and Data Engineering 18, 8 (2006), 1010-1027.
C.P. Gomes, W.J. Van Hoeve, A. Sabharwal, and B. Selman. Counting CSP solutions using generalized XOR constraints. In Proc. of AAAI'07, pp. 204-209, 2007.

Tias Guns, Siegfried Nijssen, and Luc De Raedt. 2011. Itemset mining: A constraint programming perspective. Artificial Intelligence 175 (2011), 1951-1983.

Markus Hammoria, Joachim Herbsta, and Niko Kleinerb. 2006. Interactive workflow mining - requirements, concepts and implementation. Data \& Knowledge Engineering 56, 1 (2006), 41-63.

Joachim Herbst and Dimitris Karagiannis. 2000. Integrating Machine Learning and Workflow Management to Support Acquisition and Adaptation of Workflow Models. Journal of Intelligent Systems in Accounting, Finance and Management 9 (2000), 67-92.

Joachim Herbst and Dimitris Karagiannis. 2003. Workflow mining with InWoLvE. Computers in Industry. Special Issue: Process/Workflow Mining 53, 3 (2003), 245264.

Haiyang Hu, Jianen Xie, and Hua Hu. 2011. A Novel Approach for Mining Stochastic Process Model from Workflow Logs. Journal of Computational Information Systems 7, 9 (2011), 3113-3126.
C. Koch. Processing queries on tree-structured data efficiently. In Proc. of PODS'06, pp. 213-224, 2006.
B. Kimelfeld and Y. Sagiv. Incrementally computing ordered answers of acyclic conjunctive formulas. In Proc. of NGITS'06, pp. 33-38, 2006.

Gregor Joeris. 2000. Decentralized and Flexible Workflow Enactment Based on Task Coordination Agents. In Proc. of AOIS'00+CAiSE'00. 41-62.

Wei Liu, Yu Zheng, Sanjay Chawla, Jing Yua and Xie Xing. 2011. Discovering spatio-temporal causal interactions in traffic data streams. In Proc. of $K D D^{\prime} 11$. 1010-1018.

Ruopeng Lu, Shazia Sadiq, Vineet Padmanabhan, and Guido Governatori. 2006. Using a Temporal Constraint Network for Business Process Execution. In Proc. of $A D C^{\prime} 06.157-166$.

Fabrizio Maria Maggi, R. P. Jagadeesh Chandra Bose, and Wil M. P. van der Aalst. 2012. Efficient Discovery of Understandable Declarative Process Models from Event Logs. In Proc. of CAiSE'12. 270-285.

Fabrizio Maria Maggi, Marlon Dumas, Luciano Garcia-Banuelos, and Marco Montali.
2013. Discovering Data-Aware Declarative Process Models from Event Logs. In Proc. of BPM'13. 81-96.

Fabrizio Maria Maggi, Arjan J. Mooij, and Wil M. P. van der Aalst. 2011. Userguided discovery of declarative process models. In Proc. of CIDM'11. 192-199.

Siegfried Nijssen, Tias Guns, and Luc De Raedt. 2009. Correlated itemset mining in ROC space: a constraint programming approach. In Proc. of KDD'09. 647-656.

Maja Pesic, Dragan Bosnacki, and Wil M. P. van der Aalst. 2010. Enacting Declarative Languages Using LTL: Avoiding Errors and Improving Performance. In Proc. of SPIN'10. 146-161.

Maja Pesic, Helen Schonenberg, and Wil M. P. van der Aalst. 2009. DECLARE Demo: A Constraint-based Workflow Management System. In Proc. of BPM'09 (Demos).

Maja Pesic, M. H. Schonenberg, Natalia Sidorova, and Wil M. P. van der Aalst. 2007. Constraint-Based Workflow Models: Change Made Easy. In Proc. of OTM'07. 7794.

Manfred Reichert, Stefanie Rinderle-Ma, and Peter Dadam. 2009. Transactions on Petri Nets and Other Models of Concurrency II. Springer-Verlag, Berlin, Heidelberg, Chapter Flexibility in Process-Aware Information Systems, 115-135.

Anna Rozinat and Wil M. P. van der Aalst. 2008. Conformance checking of processes based on monitoring real behavior. Information Systems 33, 1 (2008), 64-95.

Shazia Wasim Sadiq, Maria E. Orlowska, and Wasim Sadiq. 2005. Specification and validation of process constraints for flexible workflows. Information Systems 30, 5 (2005), 349-378.

Thomas J. Schaefer. 1978. The complexity of satisfiability problems. In Proce. of STOC’78. 216-226.

Guido Schimm. 2003. Mining Most Specific Workflow Models from Event-Based Data. In Proc. of $B P M^{\prime} 03.25-40$.

Pinar Senkul, Michael Kifer, and Ismail H. Toroslu. 2002. A logical Framework for Scheduling Workflows Under Resource Allocation Constraints. In Proc. of VLDB'02. 694-702.

Seppe K.L.M. van den Broucke, Jochen De Weerdt, Jan Vanthienen, and Bart Baesens. 2013. A comprehensive benchmarking framework (CoBeFra) for conformance analysis between procedural process models and event logs in ProM. In Proc. of CIDM'13. 254-261.

Wil M. P. van der Aalst. 1998. The Application of Petri Nets to Worflow Management. Journal of Circuits, Systems, and Computers 8, 1 (1998), 21-66.

Wil M. P. van der Aalst. 2011. Process Mining: Discovery, Conformance and Enhancement of Business Processes (1st ed.). Springer Publishing Company, Incorporated.

Wil M. P. van der Aalst. 2012. Decomposing Process Mining Problems Using Passages. In Proc. of ICATPN'12. 72-91.

Wil M. P. van der Aalst, Arya Adriansyah, and Boudewijn F. van Dongen. Replaying history on process models for conformance checking and performance analysis. Wiley Interdisc. Rew.: Data Mining and Knowledge Discovery 2, 2 (2012), 182192.

Wil M. P. van der Aalst, Arya Adriansyah, and Boudewijn F. van Dongen. 2011. Causal nets: a modeling language tailored towards process discovery. In Proc. of CONCUR'11. 28-42.

Wil M. P. van der Aalst, Jörg Desel, and Ekkart Kindler. 2002. On the Semantics of EPCs: A Vicious Circle. In Proc. of EPK'02. 71-80.

Wil M. P. van der Aalst, Maja Pesic, and Helen Schonenberg. 2009. Declarative workflows: Balancing between flexibility and support. Computer Science - R $8 D$ 23, 2 (2009), 99-113.

Wil M. P. van der Aalst, Boudewijn F. van Dongen, Joachim Herbst, Laura Maruster, Guido Schimm, and A.J.M.M. Weijters. 2003. Workflow Mining: A Survey of Issues and Approaches. Data $\mathcal{E}^{3}$ Knowledge Engineering 47, 2 (2003), 237-267.

Wil M. P. van der Aalst and Kees M. van Hee. 2002. Workflow Management: Models, Methods, and Systems. MIT Press.

Wil M. P. van der Aalst, A.J.M.M. Weijters, and Laura Maruster. 2004. Workflow Mining: Discovering Process Models from Event Logs. IEEE Transactions on Knowledge and Data Engineering 16, 9 (2004), 1128-1142.

Jan Martijn van der Werf, Boudewijn F. van Dongen, Cor A. J. Hurkens, and Alexander Serebrenik. 2009. Process Discovery using Integer Linear Programming. Fundamenta Informaticae 94 (2009), 387-412. Issue 3-4.

Boudewijn F. van Dongen, Ana Karla Alves de Medeiros, H.M.W. Verbeek, A.J.M.M. Weijters, and Wil M. P. van der Aalst. 2005. The ProM Framework: A New Era in Process Mining Tool Support. In Proc. of ATPN'05, 1105-1116.

Jacques Wainer and Fabio de Lima Bezerra. 2003. Constraint-Based Flexible Workflows. In Proc. of CRIWG'03, 151-158.
A.J.M.M. Weijters and Wil M. P. van der Aalst. 2001. Process Mining: Discovering Workflow Models from Event-Based Data. In Proc. of BNAIC'01. 283-290.
A.J.M.M. Weijters, Wil M. P. van der Aalst, and Ana Karla Alves de Medeiros. 2006. Process Mining with the HeuristicsMiner Algorithm. Technical Report. Eindhoven University of Technology, Eindhoven.
A.J.M.M. Weijters and Wil M. P. van der Aalst. 2003. Rediscovering workflow models from event-based data using Little Thumb. Integrated Computer-Aided Engineering 10, 2 (2003), 151-162.

Lijie Wen, Jianmin Wang, Wil M. P. van der Aalst, Biqing Huang, and Jiaguang Sun. 2009. A novel approach for process mining based on event types. J. on Intelligent Information Systems 32, 2 (2009), 163-190.

Lijie Wen, Jianmin Wang, Wil M. P. van der Aalst, Biqing Huang, and Jiaguang Sun. 2010. A novel approach for process mining based on event types. Data $\mathcal{B}$ Knowledge Engineering 69, 10 (2010), 999-1021.

Moe Thandar Wynn, Chun Ouyang, Arthur H.M. ter Hofstede, and Colin J. Fidge. 2009. Workflow Support for Product Recall Coordination. Technical Report. BPMcenter.org.
M. Yannakakis. Algorithms for acyclic database schemes. In Proc. of VLDB'81, pp. 82-94, 1981.


[^0]:    ${ }^{1}$ http://www.tpc.org/tpcds/default.asp

[^1]:    ${ }^{1}$ http://www.dbai.tuwien.ac.at/proj/hypertree/downloads.html

[^2]:    ${ }^{1}$ Available at http://jsqlparser.sourceforge.net/

[^3]:    ${ }^{1}$ It even holds that $\phi_{i}\left(\mathcal{G}_{\mathcal{A}}\right)=\phi_{i}\left(\overline{\mathcal{G}}_{\mathcal{A}}\right)$, for each $i \in N$.

[^4]:    ${ }^{1}$ Hereinafter, set operations are transparently applied to multi-sets with the usual intended meaning.

[^5]:    ${ }^{2}$ Alternatively, we might think that input and output bindings are sets as usual, but that $x$ activates two hidden activities, say $h_{1}$ and $h_{2}$, which both activate $y$ in their turn. While hidden activities play a role in the enactments, they are not registered in the log. In fact, in process discovery applications, hidden activities are frequently used to enrich the basic expressivity of the process modeling languages.

[^6]:    ${ }^{1}$ We will say that the edges $\left(t\left[j^{*}\right], t[i]\right)$ and $\left(t[i], t\left[h^{*}\right]\right)$ are supported by $t$.

[^7]:    ${ }^{2}$ Here, $+\infty$ stands for any large enough positive real number, e.g., $+\infty>|\mathcal{A}|^{2} \times \max _{(x, y) \in E} \mathbf{c s}_{\delta}(x, y)$.

[^8]:    ${ }^{3}$ See http://www.processmining.org for details on the mining suite and on its usage.
    ${ }^{4}$ We also implemented a generalization of the algorithm in Figure 9.9, but results of experimentation evidenced that its efficacy as a heuristic was not satisfying.

[^9]:    ${ }^{5}$ Duplicated traces might be processed efficiently by "weighting" each trace with the number of its occurrences in the log. We avoid this trick, as we want to test the algorithms at the varying of the log size.

[^10]:    ${ }^{6}$ In fact, since even the original models are not available as Petri nets, they get lower scores, for some metrics, than those obtained for methods returning true Petri nets.

[^11]:    ${ }^{1}$ PON Project 01_00451.

[^12]:    ${ }^{2}$ The quadratic scaling w.r.t. the number of activities is intrinsic to the computation of the causal scores, and hence characterizes most of the process discovery algorithms in the literature. In general, it might be avoided by considering less accurate heuristics (fixing the size of a sliding window over the trace), or decomposition approaches as the one recently proposed by [van der Aalst, 2012].

