Constraint satisfaction problems and related logic

Thesis submitted for the degree of
Doctor of Philosophy

Florent MADELAINE

Department of Mathematics
and Computer Science
University of Leicester

March 2003

Abstract: Feder and Vardi have proved that the class captured by a monadic fragment of existential second-order logic, MMSNP, is computationally equivalent (via randomised reductions) to the class of constraint satisfaction problems (CSP) while the latter is strictly included in the former. I introduce a new class of combinatorial problems, the so-called forbidden patterns problems (FP), that correspond exactly to the logic MMSNP and introduce some novel algebraic tools like the recolouring that allow me to construct a normal form. This leads to a constructive characterisation of the borderline of CSP within FP: a given problem in FP is either given as a problem in CSP or we build counter-examples. I relate this result to a recent and independent work by Tardif and Nešetřil which relies heavily on a correspondence between duality and density. I generalise this approach to FP. Finally, I investigate homomorphism problems for unary algebras.
Keywords: logic, combinatorics, computational complexity, descriptive complexity, finite model theory, constraint satisfaction problems (CSP), homomorphism problems, fragment of existential second order logic: monotone monadic syntactic NP without inequalities (MMSNP), Heyting algebra, dichotomy.

Problèmes de satisfactions de contraintes et logique associée

Résumé : Feder et Vardi ont prouvé que la classe capturée par un fragment monadique de la logique du second ordre existentiel, MMSNP, est calculatoirement équivalente (via des réductions probabilistes) à la classe des problèmes de satisfaction de contraintes (CSP), mais que la seconde est strictement incluse dans la première. Je caractérise exactement cette inclusion. J'introduis les problèmes de motifs interdits (FP) qui correspondent exactement à MMSNP et développe des outils algébriques originaux comme le recoloriage qui permettent de définir une forme normale et conduisent à une preuve de nature constructive: soit le problème donné est transformé en un problème de CSP, soit des contre-exemples sont construits. Je contraste par ailleurs ce résultat avec un résultat récent, dû à Tardif et Nešetřil qui utilise une correspondance entre dualité et densité que je généralise par ailleurs à FP. Finalement, je considère les problèmes de contraintes dans le cas de fonctions unaires.
Mots-clés : logique, combinatoire, complexité algorithmique, complexité descriptive, théorie des modèles finis, problèmes de satisfaction de contraintes (CSP), problèmes d'homomorphisme, fragment syntaxique et monotone de la logique du second ordre existentiel monadique sans $\neq$ (MMSNP), algèbre de Heyting, dichotomie.

Foreword

First, I would like to thank my supervisors Iain Stewart and Etienne Grandjean for their continuous support. I am also grateful to Vincent Schmitt and Dietrich Kuske whose doors were always open to my naive questions.
I wish to thank my family, friends and colleagues for their moral support. I wish to thank Claudine Madelaine, who believed in me all these years.
I would like to thank also all the people who passed onto me their love for sciences, without which I would never have attempted this work: among these "fathers in sciences", I recall especially Reinhard Pöschel, Eric Lehman, Monsieur Batman, Monsieur Brisemure and Jacques Madelaine.
Finally, I am very gratefull to Prof. Clemens Lauteman and Prof. Rajeev Raman for their time and the effort they put in reading and commenting this thesis: I greatly enjoyed my defense held in Leicester on the 11th of March 2003.
This work was supported by the EPSRC grant GR/M12933.

mehri shehezadee lei,

Contents

Chapter 1

Introduction

Descriptive complexity theory, a branch of complexity theory, seeks to classify problems, *i.e.* sets of finite structures, as to whether they can be defined using formulae of some specific logic. One of the seminal result in this theory is Fagin's theorem [13] relating definability in second-order logic with the complexity class NP (non-deterministic polynomial-time).

Theorem 1.1 (Fagin)
A problem can be defined in existential-second-order logic if and only if it can be solved in NP. *That is,* NP=*ESO.*

(Note how we equate a logic with the class of problems definable by the sentences of that logic.)

The present work originated in an attempt to find a logical characterisation of a family of combinatorial problem known as *constraint satisfaction problems* (CSP for short). These problems are of great importance in computer science and artificial intelligence and have strong ties with database theory, graph-theory and universal algebra (hence the following keywords are strongly related to constraint satisfaction problens: conjunctive-query containment problem, H-colouring, homomorphism problem, Generalised Satisfiability). In this work, we define CSP in term of existence of homomorphisms between finite structures. The most striking behaviour of these problems with respect to their complexity is that they seem to have a *dichotomy property*: that is, they are either hard (NP-complete), or tractable (P-solvable); and, furthermore there seems to be a relatively simple procedure to decide in which type a given problem falls (which is rather surprising considering Ladner's theorem and Rice's theorem, respectively). There are numerous outstanding supportive results to this conjecture: under others, those by Schaefer [52] in the Boolean case and by Hell and Nešetřil [23] in the case of undirected graph. These results have been recently generalised by Jeavons *et al.* [28, 28–34] using some tools from universal algebra and by Vardi *et al.* in [16, 35, 36] using Datalog, group theory and game theory. The latter have also attempted to give a logical characterisation of constraint satisfaction problems. They investigated numerous fragment of Fagin's ESO, showing that none of them satisfied the dichotomy property, before settling on the fragment of *monotone, monadic, syntactic NP without inequalities* (MMSNP for short). Even though they were unable to show that MMSNP has the dichotomy property, they closely related MMSNP and CSP in a theorem that can be stated as follows.

Theorem 1.2 (Feder and Vardi)
Every problem in CSP is definable by a sentence of MMSNP. Every problem definable by a sentence of MMSNP is computationally *equivalent to a problem in CSP.*

(Note the 'computationally' in the above result: by this, we mean the equivalence relation over problems induced by the existence of *reductions*.)

Furthermore, these authors exhibited examples of problems definable in MMSNP, that are not in CSP: their proof relies essentialy on counting arguments. We gave further examples of such problems in [43]. Our proof is of a different nature: it involves the *construction* of particular families of graphs. In order to give an exact logical characterisation for CSP, we attempted to generalise this approach to *any* problem definable in MMSNP. Instead of working in a logical framework, we introduced a new family of combinatorial problems that corresponds exactly to the logic MMSNP: the so-called *forbidden patterns problems* (FP for short). In this new setting, the above question can be reformulated as follows. Which problems in FP are not in CSP? Furthermore, given a problem in FP, is it decidable whether it is in CSP or not; and, if it is the case then can we give it as a CSP; that is, can we construct its template? Forbidden patterns problems are given by a *representation*, that involves a finite set of *coloured structures*. We introduced the key notion of a *recolouring* between representations. Note that the notions of a representation and a recolouring somehow generalise the notion of a structure and a homomorphism. The concept of a recolouring, together with two notions that were implicitly present in the proof of the above theorem (the notion of a *template of a representation* and of a *Feder-Vardi transformation*) allowed us to give *any* forbidden patterns problem in term of a *normal representation*. Given any normal representation, we are then able to decide (according to a simple criteria) whether the corresponding problem is in CSP or not; and, if it is in CSP then we show how to construct its template. In other words, we were able to answer by the affirmative to the above question. The proof of this result relies on the construction of particular families of structures: the so-called *witness families*. They correspond to winning strategies for Spoiler in the following two player game: Spoiler is opposed to Duplicator; for a given representation, Duplicator's aim is to argue that the corresponding problem is in CSP and to present some structure to Spoiler, claiming that it is the template of the problem. Then, Spoiler has to prove him wrong: either by exhibiting some instance that is a yes-instance of the original problem and a no-instance of the CSP (induced by the template proposed by Duplicator), or vice-versa.

These new notions of a forbidden patterns problem, a representation and a recolouring, go beyond the scope of this characterisation: our result can be related to an outstanding result by Tardif and Nešetřil. These authors have established a very surprising relationship between *duality* and *density*: they proved the existence of a correspondence between *duality pairs* and particular *gap pairs*. The former correspond to particular forbidden patterns problems: these problems are *monochrome* and have *only one* forbidden pattern (which we call *monochrome forbidden pattern problems* – notice the absence of a plural here). The latter correspond to places in the quasi-order over finite structures (induced by the existence of homomorphism) that fail to be dense. Their work leads to a "neat" characterisation of monochrome forbidden pattern problems. Such problems are in CSP if and only if the (core of the) forbidden pattern is a *tree*. However, the construction they provide to compute the template of such problems (whenever possible) uses two rather intricate constructions, namely the *exponential* and the *arrow construction*. First, we briefly relate their work and compare their results with ours. Secondly, we show that the correspondence between duality and density stems from the fact that the above mentionned quasi-order is closely related to a *Heyting algebra* (the approach is not original and follows the line of Tardif and Nešetřil's proof) Finally, the fact that the notion of a representation and a recolouring generalise the notion of a structure and a homomorphism is strengthen further: we show that representations and recolouring can also be related to a suitable *Heyting algebra*. Hence, it follows that the correspondence between duality and density also holds in this more general setting. However, this result is not fully satisfactory and yields a conjecture that we motivate with partial results.

Turning to complexity aspects of forbidden patterns problems, we were also able to give examples of problems in FP that are not in CSP (via our characterisation) that are *complete* for the following standart complexity classes: NL, P and NP. In other words according to some known "fine" complexity results for tractable problems in CSP (see [28]), we strengthen further Feder and Vardi's theorem: in the sense that the class FP seems to behave very much in the same way as CSP with respect to complexity classes. This fact could seem trivial to the reader: MMSNP=FP implies that FP is computationally equivalent to CSP, by Feder and Vardi's theorem. However, in this theorem, the reductions considered were randomized polynomial-time reductions and polynomial-time reductions, respectively: these reductions are too powerful to allow any "fine" complexity results. In order to speed up the proof of such results, we adapted some examples given by Grädel in [21], and introduced further examples according to his elegant characterisation of subclasses of NP, byhand of fragment of ESO. We also briefly reviewed common heuristics and restriction (that are known to lead to tractability for other problems) to forbidden patterns problem.

We interested ourselves as well to a topic only loosely related to the above in [42]: we noticed that while numerous partial results had been proved for constraint satisfaction problems in the case of relational structures, the case of algebras had never been investigated. We proved a result that indicates that it would be at least as hard to obtain a general dichotomy result in this case as in the case of relational structures. We concentrated on an extremely restricted case: the case of unary algebras. We were able to show that even in the presence of only two unary function symbols, the uniform homomorphism problem was NP-complete ('uniform' means that an instance is a pair of algebras and the question is to decide whether there exists an homomorphism of the first algebra into the second). Further, in the case of only one function symbol, we obtained an interesting dichotomy result: the non-uniform problems were either trivial or L-complete (by opposition, 'non-uniform' means that the instance is a single algebra and the question is to decide whether there exists a homomorphism of this algebra into some *fixed* algebra: the latter is called the *template* of the problem). Notice further that this result provides the first known examples of constraint satisfaction problem that are L-complete.

We tried to keep this work as self-contained as possible: however some basic knowledge in complexity theory and descriptive complexity theory is needed. With respect to complexity theory, we refer the reader to the following textbooks: [46] by Papadimitriou or [25] by Hopcroft and Ullman in English and [39] by Lassaigne and Rougemont in French (the latter is also a good introduction to descriptive complexity theory). We refer to [20], Garey and Johnson's guide to NP-completeness. With respect to descriptive complexity theory, we refer to [12], by Ebbinghaus and Flum. We provided some definitions in Appendix A. Furthermore, a little background in algebra would help the reader to understand the motivation of some of our results: with respect to universal algebra, we refer the reader to [44]; and, with respect to category theory, we advise [38]. We provided some definitions in Appendix B.

This work is organised as follows: in Chapter 2, we introduce CSP as the class of (non-uniform) homomorphism problems and comment some known dichotomy results. We also introduce the logic MMSNP and relate Feder and Vardi's theorem in some detail. The remainder of this chapter is devoted to some examples of graph problems definable in MMSNP but not in CSP: this is joint work with Iain Stewart that has been published (see [43]). Chapter 3 is devoted to MMSNP and is concluded by a proof of Feder and Vardi's theorem. In Chapter 4, we introduce the forbidden patterns problems and the related notions of a representation, a recolouring, the template of a representation *etc.* The aim being to derive the notion of a normal representation, which is used to build witness families (our tool to prove that a given forbidden patterns problem can not be in CSP). We finally state our main result: an exact characterisation of the forbidden patterns problems that are not in CSP. Next, in Chapter 5, we relate our main result with Tardif and Nešetřil's results: one of their result concerns the correspondence between duality and density, that we prove in a more general setting of Heyting algebras. The other result is a characterisation of gap pairs for structures, that we skecth and briefly comment. We also show that representations and recolourings can be related to a Heyting algebra . We conclude this chapter with some open problems. In Chapter 6, we present examples of forbidden patterns problems that are not in CSP and complete for some standart complexity subclasses of NP. We also review some standart heuristics that can be applied to forbidden patterns problems in order to gain tractability. Finally, in Chapter 7, we present some results concerning the complexity of homomorphism problems in the case of unary algebras (this is joint work with Iain Stewart: it is available as a preprint [42]; and, has been merged with an independant and contemporary work by Feder and recently submitted as [15]).

We strongly recommend to read Chapter 2 before the remainder of this work. It is not necessary to read Chapter 3 to understand the following Chapters, except for Section 4.1.4 that relates forbidden patterns problems with the logic MMSNP. Chapter 4 must be read before Chapter 5. However, in order to read Chapter 6, only the definition of forbidden patterns problems given in Section 4.1.3 is necessary. Finally, Chapter 7 can be read independantly. Note also that an index that includes the most important notions treated, is available towards the end of this work.

Chapter 2

Homomorphism problems

We introduce briefly the reader to homomorphisms problems and cite some results to motivate the definition of Feder and Vardi's logic MMSNP. We then prove that this logic captures problems that are not homomorphism problems

Constraint satisfaction problems consist of finding assignments of values to variables subject to constraints on the values which can be simultaneously assigned to certain specified subsets of variables. They are of great importance in computer science and artificial intelligence, and have strong links with database theory, combinatorics and universal algebra. For example, the general constraint satisfaction problem is also known as the conjunctive-query containment problem from database theory and the homomorphism problem from combinatorics [4]; and, there is a strong link between the tractability of constraint satisfaction problems and the study of the closure of relations under certain operations in universal algebra [32]. This diversity has meant that the study of these constraint satisfaction problems has progressed on a number of different fronts and according to different motivations.

Our formulation of constraint satisfaction involves the existence of a homomorphism of one finite structure to another, and in some parts of this work we are concerned with the computational complexity of constraint satisfaction problems when the structures involved are restricted. The *general constraint satisfaction problem* has: as its instances pairs of finite structures (A, B) over the same signature; and, as its yes-instances instances (A, B) for which there is a homomorphism of A to B. The general constraint satisfaction problem is trivially in NP and is easily shown to be NP-complete; and it is usual to restrict the problem so that all finite structures come from some specific class of finite structures or, further, so that the second component, the *template*, of any instance is some fixed finite structure. The former problems are called *uniform* constraint satisfaction problems, as the two structures in an instance can be arbitrarily drawn from the given class of structures, whilst the latter problems are called *non-uniform* constraint satisfaction problems, as the second structure in an instance must be a given fixed structure (rather than thinking of instances of non-uniform constraint satisfaction problems as pairs of finite structures (A, T), with T fixed, we simply think of them as finite structures A, with yes-instances those instances A for which there exists a homomorphism to T). The computational complexity of these restricted problems is then studied with the ultimate goal being a classification as to the conditions under which a (uniform or non-uniform) constraint satisfaction problem has a given computational complexity. In this chapter, we shall concentrate on the non-uniform case.

This chapter is organised as follows. In the first section we shall give some basic definitions and results. In Section 2.2 we shall relate briefly the main known results concerning the complexity of non-uniform constraint satisfaction problem: in particular, the so-called *dichotomy results* of Schaefer for boolean problems and of Hell and Nešetřil for the case of undirected graphs. In the final section, we shall outline a logic introduced by Feder and Vardi together with one of their results that states that the class of problems captured by this logic is *computationally equivalent* to the class of non-uniform constraint satisfaction problems. However, we shall prove that various problems over graphs that are expressible in this logic are *not* realisable as non-uniform constraint satisfaction problems.

2.1 Preliminaries

Let σ be a signature with relation symbols only, that is, symbols $R_1, R_2, \ldots, R_s$ with respective arities $r_1 \geq 1, r_2 \geq 1, \ldots, r_s \geq 1$.

Recall that a finite σ-structure A consists of a finite set[1], called the domain of A and denoted by $|A|$, together with an interpretation $R_i^A \subseteq |A|^{r_i}$ for every symbol R_i in σ, $1 \leq i \leq s$. The *size* of A, that is the cardinal of the set $|A|$, is also denoted by $|A|$ (this does not cause confusion).

Let A and B be two σ-structures. We call a *homomorphism* of A to B any mapping $h : |A| \to |B|$ satisfying:

- for any r-ary symbol in σ and for any $\bar{a}$ in $|A|^r$, if $R^A(\bar{a})$ holds then $R^B(h(\bar{b}))$ holds (where $h(\bar{a})$ denotes the r-tuple obtained from $\bar{a}$ via an application of the mapping h component-wise).

If h is a homomorphism of A to B then we write $A \xrightarrow{h} B$; we write $A \overset{h}{\twoheadrightarrow} B$ if h is a surjective homomorphism of A to B; and, we write $A \overset{h}{\hookrightarrow} B$ if h is an injective homomorphism of A to B. If there exists some homomorphism of A to B then we write $A \longrightarrow B$; and, if none exists $A \not\longrightarrow B$.

If $A \overset{h}{\hookrightarrow} B$ then we say that A is a *substructure* of B. If, further, for any r-ary symbol R of σ and any $\bar{a}$ in $|A|^r$, if $R^B(h(\bar{a}))$ holds then $R^A(\bar{a})$ holds, then we say that A is an *induced* substructure of B.

[1] Contrary to usage in finite model theory, we do consider the void structure and the structure with a single element domain.

An *isomorphism* is a bijective homomorphism whose inverse is a homomorphism. When an isomorphism exists between A and B then we say that A and B are *isomorphic* and we write $A \approx B$. Denote by $STRUC(\sigma)$ the class of finite σ-structures.

The *homomorphic image* of A via h, denoted $h(A)$, is the (not necessarily induced) substructure of B such that:

- $|h(A)| := \{b \in |B| | \exists a \in |A| \text{ such that } h(a) = b\}$; and
- for any r-ary symbol R in σ and any $\bar{b}$ in $|h(A)|^r$, $R^{h(A)}(\bar{b})$ holds, if, and only if, there exists some $\bar{a}$ in $|A|^r$ such that $h(\bar{a}) = \bar{b}$ and $R^A(\bar{a})$ holds.

Moreover, it is immediate that the composition of two homomorphisms is a homomorphism and that for any structure A, there exists an identity homomorphism $A \xrightarrow{id_A} A$ (defined by setting $id_A(x) := x$ for any x in $|A|$) such that for any structures B and C and homomorphisms $B \xrightarrow{f} A$ and $A \xrightarrow{g} C$, we have $id_A \circ f = f$ and $g \circ id_A = g$. Furthermore, the composition of homomorphisms being associative, one can speak of the *category of finite σ-structures*. As we shall see later, this category has some interesting properties: in fact, if one considers structures up to homomorphism equivalence then we get a *Heyting Algebra* (*cf.* Chater 5 on page 137).

Let A be a σ-structure. Recall that the *(non-uniform) homomorphism problem* with *template* A, denoted $CSP(A)$, has yes-instances those σ-structures B such that $B \longrightarrow A$. Denote by CSP_σ the class of homomorphism problems having as template a σ-structure and set:

$$CSP := \bigcup_{\sigma \text{ rel sign}} CSP_\sigma.$$

Proposition 2.1 *Let A and B be two σ-structures. $CSP(A) \subseteq CSP(B)$ if, and only if, $A \longrightarrow B$.*

PROOF. If $CSP(A) \subseteq CSP(B)$ then since $A \xrightarrow{id_A} A$, it follows that A belongs to $CSP(A)$. Hence that A is in $CSP(B)$; that is, $A \longrightarrow B$. Conversely, if $A \xrightarrow{h} B$ for some h then for any C in $CSP(A)$; that is, such that $C \xrightarrow{g} A$ for some g; by composition, it follows that $C \xrightarrow{hog} B$; hence that C belongs to $CSP(B)$. □

2.2 Known complexity results

As we mentioned previously, the general constraint satisfaction problem is NP-complete. There are two main ways of restricting this problem in order to obtain tractability. The first way consists in imposing that the first structure of any instance is somehow like a tree, to be precise that it has bounded tree-width, to ensure that the standard resolution algorithms' backtrack is bounded. This approach has been developed by Freuder (*cf.* [18, 19]) but is a direct consequence of a more recent result due to Courcelle (*cf.* [6]). The second approach consists in restricting the second structure of any instance; which often leads to so-called *dichotomy results*; that is, results in which restrictions of the general problem are either NP-complete or decidable in polynomial time. These dichotomy results are best appreciated to the light of Ladner's theorem (cf. [37]); one version of which is as follows.

Theorem 2.2 (Ladner) *If* P $\neq$ NP *then there is a language in* NP *which is neither in* P *nor* NP*-complete.*

Notice that we do not know of any natural problem with such a property (under the assumption that P $\neq$ NP): some problems that resist any classification attempts, such as GRAPH-ISOMORPHISM, are conjectured to be such natural problems.

In practice many problems can be easily specified as constraint satisfaction problems (*e.g.* optimisation problems such as the FREQUENCY ASSIGNMENT problem, *cf.* [11]). For this reason, constraint staisfaction solvers are of real practical importance, which motivates further the study of constraint satisfaction problems in theoretical computer science. Indeed, note that constraint satisfaction problems capture many benchmark problems: in [28], various natural problems are encoded as uniform constraint satisfaction problems (in this work the general constraint satisfaction problem is even referred not without humour as the *great combinatorial problem*). The encodings are in general much more natural and straightforward than reductions to other well-known NP-complete problems.

Next, we shall briefly relate two important dichotomy results; namely the case of undirected graphs (the problem is known also as the *H-colouring* problem) due to Hell and Nešetřil and the case of structures with Boolean domains due to Schaefer (the problem is known as the *Generalised Satisfiability* problem).

2.2.1 H-colouring

The non-uniform constraint satisfaction problem when restricted to undirected graphs is known as the H-colouring problem, where H denotes the template of the problem studied. For example, when H is a triangle, the H-colouring problem is nothing else than 3-COL (the problem that consists of all graphs whose vertices can be coloured with three colours such that no two adjacent vertices are coloured with the same colour). The latter is known to be NP-complete (*cf.* [20]). Hell and Nešetřil proved the following in [23].

Theorem 2.3 (Hell and Nešetřil)
The H-colouring problem is NP*-complete whenever H is not bipartite and can be decided in polynomial time otherwise.*

Their proof makes use of three constructions over graphs that allow one to reduce the question of whether the H-colouring is NP-complete to the question of whether the H'-colouring problem is NP-complete, where H and H' are related via one of these three constructions. They show further that the case when the template is a bipartite graph is tractable; indeed, it can be easily shown that *as a decision problem,* for any bipartite graph B, the B-colouring problem coincides with 2-COL (the problem that consists of all graphs whose vertices can be coloured with two colours such that no two adjacent vertices are coloured with the same colour) which is known to be decidable in polynomial time: as the *core* of a bipartite graph is the graph consisting of a single edge; in other words nothing else than the template of 2-COL , *cf.* Subsection 4.2.1 on page 85. The main part of their proof is indirect and consists in assuming that for some non bipartite graph H, the H-colouring problem is not NP-complete (under the more general assumption that P and NP do not coincide). By properties of the three constructions mentioned above, and the facts that H is not bipartite and can not be a clique (otherwise H would be either bipartite or the H-colouring problem NP-complete), they reduce the H-colouring problem to the H'-colouring problem, where H' can not exist. This part of their proof is fairly technical and involves a case study on the structure of H and its properties to derive some contradicting properties on H'. Notice that no constructive proof is presently known for this result. Furthermore, some unsuccessful attempts have been made to generalise this result to other structures; even the case of directed graph remains open.

2.2.2 Generalised Satisfiability

There exists another type of dichotomy result which is not quite comparable to the former result. Given some fixed domain D of values (that corresponds to the domain of the template) call a set Γ of relations *tractable* if for any structure T with domain D and relations in Γ, the constraint satisfaction problem with template T is decidable in polynomial time. Denote by $CSP(\Gamma)$ the class of non-uniform constraint satisfaction problems whose template T consists of relations from Γ as above. The uniform constraint satisfaction problem where T is drawn from the class $\mathcal{B}$ of Boolean structures is known as GENERALISED-SAT and was studied by Schaefer in [52]. Schaefer proved the following dichotomy result.

Theorem 2.4 (Schaefer). *Let Γ_0 be a subset of $\Gamma_{\mathcal{B}}$, the set of all Boolean finitary relations. If Γ_0 falls within one of the following 6 classes, that is if:*

1. Γ_0 *is* 0*-valid;*
2. Γ_0 *is* 1*-valid;*
3. Γ_0 *is affine;*
4. Γ_0 *is bijunctive;*
5. Γ_0 *is Horn; or*
6. Γ_0 *is anti-Horn,*

then $CSP(\Gamma_0)$ is tractable, otherwise it is NP*-complete.*

These 6 classes have simple characterisations in term of closure properties. For example, a relation is 0-valid if, and only if, it is closed under the Boolean constant operation 0; and, it is Horn if, and only if, it is closed under the binary Boolean operation $\wedge$.

Schaefer's dichotomy result has been generalised to optimisation complexity classes and to counting classes by Creignou *et al.* (*cf.* [7]).

2.2.3 Further selected results

Schaefer was inspired by the work of Post on Boolean functions and relations, work that has been extended in a branch of *universal algebra* known as *clone theory*. Schaefer's approach has been applied by Jeavons *et al.* to larger domains and partial dichotomy results have been obtained (*cf.* [28–34]). For an introduction to this approach see, for example, [41]. Notice that this method leads only to partial results as it relies heavily on what is known about the *clone lattice*. The Boolean clone lattice was completely described by Post in [47]; and, is countable whereas it is known that the clone lattice for larger domains is not (for more on clone theory, see the excellent book in German by Pöschel and Kalužnin [49], a technical report in English by Pöschel [48] or the first chapter of Szendrei's exposition [55]). As a matter of fact, there is presently no description of the clone lattice even for a domain of size 3. However, some progress has been made as regards a conjecture that dichotomy results *à la* Schaefer exists for any finite domain. Recent work by Bulatov, Krokhin and Jeavons involves the use of deep results from universal algebra in [3].

Note that the dichotomy results of the two previous theorems are not comparable. It was proved in [2] that $CSP(\Gamma_{\mathcal{B}})$ is not tractable, where $\Gamma_{\mathcal{B}}$ denotes the set of the edge relations of any finite bipartite graph.

Apart from Jeavons *et al.*, there is another group of researchers that have attempted to develop general methods to classify non-uniform constraint satisfaction problems, namely Feder and Vardi in [16]. In their work, tractable sets of relations fall into two main classes, one being defined in terms of Datalog, the other in terms of group theory. Some of these results have been extended or proved in a more concise way by Kolaitis and Vardi in [35] and [36]. The terminology of uniform and non-uniform constraint satisfaction problem was taken from [35], where the authors proved that many known dichotomy results *uniformise*; that is, can be generalised to the uniform case. We shall explain in more detail what we understand by this in Chapter 7, where homomorphism problems for unary functions are studied. However, for the moment we shall be concerned mainly with a specific result of Feder and Vardi from [16], where they defined the logic MMSNP in an attempt to characterise logically CSP. First, they conjectured the dichotomy of CSP as follows;

Conjecture 2.5 (dichotomy of CSP)
Every problem in CSP is either in P *or* NP*-complete.*

Recall that Syntactic NP (SNP for short) is the fragment of Fagin's existential second order logic (ESO for short)that consists of sentences of the form $\exists\bar{S}\forall\bar{x}\phi$, where ϕ is quantifier-free; that is, second order sentences with a universal first-order part. In order to find some logic for CSP, Feder and Vardi looked for a logic $\mathcal{L}$ that is a restriction of *SNP* (CSP is easily seen to be captured by SNP) and would have the dichotomy property (as SNP itself does not). They investigated 3 types of restrictions on *SNP*: namely *monotonicity, monadicity* and *no inequalities*. That is, imposing that each input predicate occurs with the same *polarity* within a sentence, respectively imposing the second order predicates to be *monadic*, and respectively that no inequality symbol occurs within a sentence. They showed that imposing two of these restrictions is not sufficient by proving the following theorems ($\mathcal{L}$ denotes here the logic obtained from SNP by imposing any two restrictions among the three listed above).

Theorem 2.6 (Feder and Vardi)
Every problem A in NP *has an equivalent (under polynomial-time reductions) problem B in the class of problems expressed by sentences of $\mathcal{L}$.*

(by 'equivalent' we mean that: the problem A reduces to the problem B; and, conversely, the problem B reduces to the problem A.) Therefore as a corollary from Ladner's theorem, it follows that none of these three logics could be adequate to capture exactly *CSP* according to the dichotomy conjecture. They were however unable to extend Ladner's diagonalisation arguments when the three restrictions mentioned above were imposed simultaneously on the logic SNP. They called this fragment of SNP, *Monotone Monadic SNP without inequalities*, which they denoted by MMSNP for short.

EXAMPLE. Consider the signature $\sigma_2 := (E)$, where E is a binary symbol. We can see problems over σ_2 as the realisation of some abstract graph problems via the following encoding *"there exists an edge between two vertices u and v if, and only if, $E(u,v)$ holds or $E(v,u)$ holds"*. In this setting, the well known abstract graph problem 3-COL (that consists of those graphs whose vertices can be coloured with three colours such that every pair of adjacent vertices have been assigned different colours) can be realised over σ_2 as the problem captured by the following sentence of MMSNP.

$$\begin{array}{rl} \exists R \exists G \exists B \forall x \forall y & \neg(B(x) \wedge R(x)) \wedge \neg(B(x) \wedge G(x)) \wedge \neg(R(x) \wedge G(x)) \\ \wedge & \neg(\neg R(x) \wedge \neg G(x) \wedge \neg B(x)) \\ \wedge & \neg(E(x,y) \wedge R(x) \wedge R(y)) \wedge \neg(E(x,y) \wedge G(x) \wedge G(y)) \\ \wedge & \neg(E(x,y) \wedge B(x) \wedge B(y)). \end{array}$$

▲

2.3 Feder and Vardi's MMSNP

In [16] Feder and Vardi attempted to give a logic for CSP: they introduced the logic MMSNP and showed that the set of problems captured by MMSNP is computationally equivalent to CSP. In this section, we introduce briefly this result.

2.3.1 Definition

Monotone Monadic SNP without inequality is a fragment of ESO and consists of the set of formulae of the following form:

$$\exists \bar{M} \forall \bar{x} \bigwedge_i \neg(\alpha_i(\bar{R},\bar{x}) \wedge \beta_i(\bar{M},\bar{x})),$$

where for every *negated conjunct* $\neg(\alpha_i \wedge \beta_i)$:

- the α-*part* α_i consists of a conjunction of *positive* atoms involving relational symbols from σ and variables from $\bar{x}$; and
- the β-*part* (or colouring) β_i consists of a conjunction of atoms or negated atoms involving the monadic existentially-quantified predicates $\bar{M}$ and variables from $\bar{x}$.

Notice that the equality symbol does not occur in Φ. Monotone Monadic SNP without inequality is denoted by MMSNP, for short.

2.3.2 MMSNP is computationally equivalent to CSP

The result we are about to quote has initiated the present work (except for Chapter 7). In the remainder of this work, when we write 'Feder and Vardi's theorem' we understand the following key result.

Theorem 2.7 (Feder and Vardi)
Every problem in CSP is expressible by a sentence of MMSNP. Every problem $\mathbf{P}_\Phi$ expressible by a sentence Φ of MMSNP is equivalent to a problem CSP(T_Φ) in CSP: $\mathbf{P}_\Phi$ reduces to CSP(T_Φ) in polynomial time; and, CSP(T_Φ) reduces to $\mathbf{P}_\Phi$ in randomised polynomial time.

We shall give a proof of the previous theorem in Chapter 3. Feder and Vardi showed that MMSNP captures more than just CSP *i.e.*, that there are problems captured by MMSNP that are not in CSP. They gave two examples of such problems over graphs; the problem consisting of those graphs that are triangle-free; and the problem consisting of those graphs G for which one can colour the elements of $|G|$ black or white such that the coloured graph contains no monochromatic triangle. They gave a sketch of this proof in which they used a counting argument. In the next section, we shall give further examples of such problems, using a different type of proof, involving the construction of families of graphs with special properties.

2.4 MMSNP captures more than CSP

We exhibit some problems over σ_2 that are captured by MMSNP and show that they can not be in CSP (this section is an extended version of [43]).

2.4.1 Some problems expressible by a sentence of MMSNP

The problem TRI-FREE is the problem over σ_2 defined by the following first-order sentence:

$$\begin{aligned}&\forall x(\neg E(x,x))\wedge\\ &\forall x\forall y\forall z(\neg(E(x,y)\vee E(y,x))\vee\neg(E(x,z)\vee E(z,x))\vee\neg(E(y,z)\vee E(z,y))).\end{aligned}$$

Note that the above sentence can be considered to be a realisation of the abstract decision problem consisting of those undirected graphs in which there is no triangle. TRI-FREE is also expressible by a sentence of MMSNP since although the above sentence is not directly a sentence of MMSNP according to our definition, it is logically equivalent to one: it is logically equivalent to the following sentence using the identity $\neg(P\vee Q)\equiv\neg P\wedge\neg Q$

$$\begin{aligned}&\forall x(\neg E(x,x))\wedge\\ &\forall x\forall y\forall z((\neg E(x,y)\wedge\neg E(y,x))\vee(\neg E(x,z)\wedge\neg E(z,x))\vee(\neg E(y,z)\wedge\neg E(z,y))).\end{aligned}$$

Then using the distributivity of $\wedge$ by $\vee$ we obtain the following equivalent sentence

$$\begin{aligned}\forall x\quad &\neg E(x,x)\wedge\forall x\forall y\forall z\\ &(\neg E(x,y)\vee\neg E(x,z)\vee\neg E(y,z))\wedge(\neg E(x,y)\vee\neg E(x,z)\vee\neg E(z,y))\\ &\wedge(\neg E(x,y)\vee\neg E(z,x)\vee\neg E(y,z))\wedge(\neg E(x,y)\vee\neg E(z,x)\vee\neg E(z,y))\\ &\wedge(\neg E(y,x)\vee\neg E(x,z)\vee\neg E(y,z))\wedge(\neg E(y,x)\vee\neg E(x,z)\vee\neg E(z,y))\\ &\wedge(\neg E(y,x)\vee\neg E(z,x)\vee\neg E(y,z))\wedge(\neg E(y,x)\vee\neg E(z,x)\vee\neg E(z,y)).\end{aligned}$$

Finally, using the fact that $\neg P\vee\neg Q\equiv\neg(P\wedge Q)$ and rewriting the sentence in prenex form, we obtain the following equivalent MMSNP sentence

$$\begin{aligned}\Phi_1:=\forall x\forall y\forall z\quad &\neg E(x,x)\wedge\neg\ell_1(x,y,z)\wedge\neg\ell_1(x,z,y)\wedge\neg\ell_2(x,y,z)\wedge\neg\ell_1(z,y,x)\\ &\wedge\neg\ell_1(y,x,z)\wedge\neg\ell_2(y,x,z)\wedge\neg\ell_1(y,z,x)\wedge\neg\ell_1(z,y,x),\end{aligned}$$

where

$$\ell_1(x,y,z)=(E(x,y)\wedge E(x,z)\wedge E(y,z)),$$

and

$$\ell_2(x,y,z)=(E(x,y)\wedge E(z,x)\wedge E(y,z)).$$

The problem NO-MONO-TRI is the problem over σ_2 defined by the following sentence:

$$\exists C(\forall x(\neg E(x,x)) \wedge \forall x\forall y\forall z(((E(x,y) \vee E(y,x)) \wedge (E(x,z) \wedge E(z,x)) \wedge(E(y,z) \vee E(z,y))) \Rightarrow (\neg(C(x) \wedge C(y) \wedge C(z)) \wedge \neg(\neg C(x) \wedge \neg C(y) \wedge \neg C(z))))).$$

Note that the problem NO-MONO-TRI can be considered as a realisation of the abstract decision problem consisting of those undirected graphs for which there exists a 2-colouring of the vertices so that the vertices of every triangle in the graph are not monochromatically coloured. Note that the problem NO-MONO-TRI can also be captured by a sentence of MMSNP. The previous sentence can be rewritten using the same technique as previously, since the polarity of each occurrence of the symbol E is odd. We prefer to work with the previous sentence as it is much more compact. The same shall hold for any further sentence we shall consider in this section.

The problem TRI-FREE-TRI is the problem over σ_2 defined by the following sentence:

$$\exists R\exists W\exists B(\forall x((R(x) \wedge \neg W(x) \wedge \neg B(x)) \vee (\neg R(x) \wedge W(x) \wedge \neg B(x)) \vee(\neg R(x) \wedge \neg W(x) \wedge B(x))) \wedge \forall x\forall y((E(x,y) \vee E(y,x)) \Rightarrow (\neg(R(x) \wedge R(y)) \wedge \neg(W(x) \wedge W(y)) \wedge \neg(B(x) \wedge B(y)))) \wedge \forall x(\neg E(x,x)) \wedge\forall x\forall y\forall z(\neg(E(x,y) \vee E(y,x)) \vee \neg(E(x,z) \vee E(z,x)) \vee \neg(E(y,z) \vee E(z,y)))).$$

Note that the problem TRI-FREE-TRI can be considered as a realisation of the abstract decision problem consisting of those undirected graphs that are tripartite and in which there is no triangle; that is, as a restriction of TRI-FREE to tripartite graphs.

The problem NO-WALK-5 is the problem over σ_2 defined by the following first-order sentence:

$$\forall x(\neg E(x,x)) \wedge \forall x_1 \forall x_2 \forall x_3 \forall x_4 \forall x_5 (\neg((E(x_1,x_2) \vee E(x_2,x_1)) \wedge (E(x_2,x_3) \vee E(x_3,x_2)) \wedge (E(x_3,x_4) \vee E(x_4,x_3)) \wedge (E(x_4,x_5) \vee E(x_5,x_4)) \wedge (E(x_5,x_1) \vee E(x_1,x_5)))).$$

Note that NO-WALK-5 can be considered to be a realisation of the abstract decision problem consisting of those undirected graphs in which there is no closed walk of length 5. The problem NO-WALK-7 is defined similarly. Moreover, consider the problems NO-WALK-5-TRI and NO-WALK-7-TRI respectively, as the restrictions of NO-WALK-5 and NO-WALK-7 respectively, to tripartite graphs, as above.

Our first observation is that, if the template defining a problem in CSP over σ_2 has a self-loop, then the problem must consist of the class of all σ_2-structures. Hence, we may assume that any template has no self-loops as none of the problems we consider in this section are trivial. Our second observation is that the template defining a problem in CSP over σ_2 must be a yes-instance of the problem (as the identity map of the template to the template is a homomorphism).

Lemma 2.8 *Let $G,T \in STRUC(\sigma_2)$. Suppose that, $T \in$ TRI-FREE. Furthermore, suppose that in the undirected graph encoded by G, there is a path of length 3 joining two non-adjacent vertices u and v. Then for any $G \xrightarrow{h} T$, $h(u) \neq h(v)$.*

PROOF. Let u and v be two non-adjacent vertices of G. Suppose further that there is a homomorphism h of G to T such that $h(u) = h(v)$. By definition, there is a path u, w_1, w_2, v in the graph encoded by G. Because T has no self-loops, we must have that $h(u)$, $h(w_1)$ and $h(w_2)$ are pairwise distinct in T and since h is a homomorphism, we have $(E(h(u),h(w_1))$ or $E(h(w_1),h(u)))$ and $(E(h(w_1),h(w_2))$ or $E(h(w_2),h(w_1)))$ and $(E(h(w_2),h(u))$ or $E(h(u),h(w_2)))$ that hold in T; that is, the graph encoded by T has a triangle. Thus $T \notin$ TRI-FREE. This yields a contradiction. □

Suppose that some problem **P** over σ_2 is such that:

- every σ_2-structure in **P** is in TRI-FREE; and

- for every n, **P** contains a structure H_n that encodes a graph with n mutually non-adjacent vertices where there is a path of length 3 joining every pair of such vertices.

Then, by Lemma 2.8 on the preceding page, **P** is not in CSP (any homomorphism of H_n to the template must have an image of size at least n). In the following, we construct such a family of graphs for all the first-order problems that have been introduced in this section.

2.4.2 Construction of H_n.

Define the structure H_n as follows. The domain of H_n consist of the union of the sets:

- $V_n = \{1, 2, \ldots, n\}$;
- $U_n^1 = \{(i, j) : 1 \leq i, j \leq n, i < j\}$; and
- $U_n^2 = \{(i, j) : 1 \leq i, j \leq n, i > j\}$.

E^{H_n} consist of the union of the sets:

- $\{(i, (i, j)) : 1 \leq i, j \leq n, i < j\}$;
- $\{(i, (i, j)) : 1 \leq i, j \leq n, i > j\}$; and
- $\{((i, j), (j, i)) : 1 \leq i, j \leq n, i \neq j\}$.

The graph encoded by H_n can be depicted as in Fig. 2.1 on the next page Note that: the graph encoded by H_n is triangle-free; there is a path of length 3 joining any two distinct vertices of V_n; V_n forms an independent set in this graph; and this graph is tripartite.

Lemma 2.9 *There does not exist a closed walk of length 5 or 7 in the graph encoded by H_n.*

PROOF. Suppose that there exists a closed walk W of length 5 or 7 in the graph encoded by H_n. As this graph is tripartite, W must have at least one vertex, w_1 say, in V_n. Hence, there is $w_2 \in V_n \setminus \{w_1\}$ such that either:

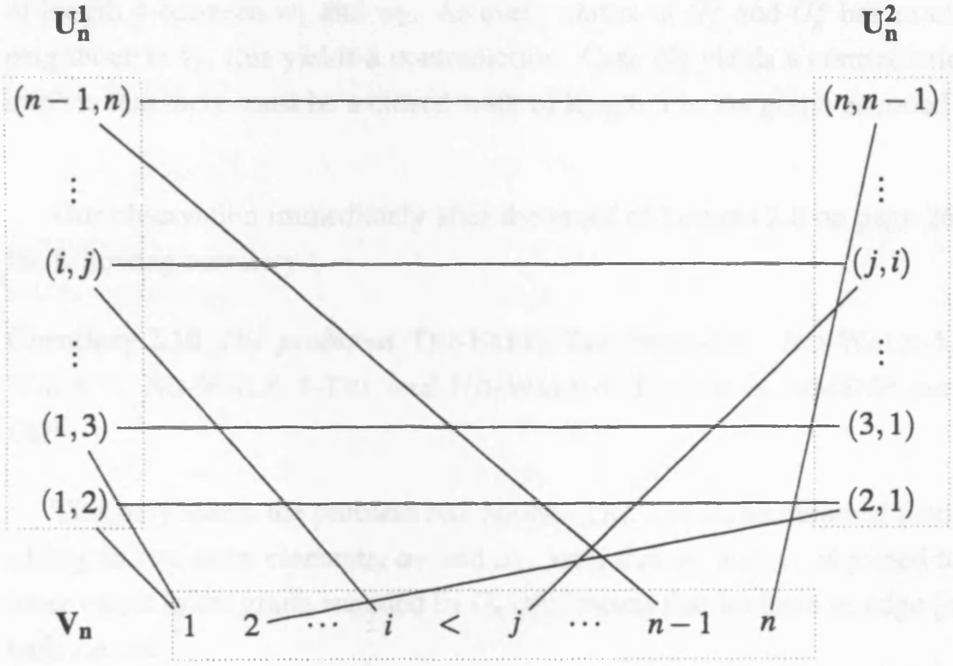

Figure 2.1: The Graph encoded by H_n.

1. w_1, (w_1,w_2), (w_2,w_1) is a sub-walk of W; or

2. w_1, (w_1,w_2), w_1 is a sub-walk of W.

Suppose that the length of W is 5. In case (1), we obtain a contradiction as every vertex of U_n^1 and U_n^2 is joined to exactly one vertex of V_n. In case (2), we also obtain a contradiction as this would imply that the graph encoded by H_n has a triangle. Hence, this graph has no closed walk of length 5.

Suppose that the length of W is 7. In case (1), we must have a closed walk of length 4 between w_1 and w_2. As every vertex of U_n^1 and U_n^2 has exactly one neighbour in V_n, this yields a contradiction. Case (2) yields a contradiction as it implies that there must be a closed walk of length 5 in the graph encoded by H_n.

□

Our observation immediately after the proof of Lemma 2.8 on page 26 yields the following corollary.

Corollary 2.10 *The problems* TRI-FREE, TRI-FREE-TRI, NO-WALK-5, NO--WALK-7, NO-WALK-5-TRI *and* NO-WALK-7-TRI *are in MMSNP but not in CSP.*

This only leaves the problem NO-MONO-TRI. Let G_n be obtained from H_n by adding in two extra elements, $a_\top$ and $a_\bot$, such that $a_\top$ and $a_\bot$ is joined to every other vertex in the graph encoded by G_n (this means that we have an edge $(a_\top, a_\bot)$ too); *i.e.* set

$$E^{G_n} := E^{H_n} \cup \{(a_\bot, w), (a_\top, w) |\text{ such that } w \in |G_n|\} \cup \{(a_\top, a_\bot)\}.$$

Lemma 2.11 *Suppose that u and v are vertices of V_n in the graph encoded by G_n and let T be a σ_2-structure in* NO-MONO-TRI *such that there is a homomorphism h of G_n to T. Then $h(u) \neq h(v)$.*

PROOF. Suppose that $h(u) = h(v)$. By arguing as in Lemma 2.8 on page 26, there are vertices w_1 and w_2 of $G_n \setminus \{a_\top, a_\bot\}$ such that $h(w_1), h(u)$ and $h(w_2)$ are pairwise distinct. Also, both $h(a_\top)$ and $h(a_\bot)$ must be different from the image of any other vertex of G_n. Hence, $E(x,y)$ or $E(y,x)$ holds in T for every distinct pair of elements x and y from the set $\{h(u), h(w_1), h(w_2), h(a_\top), h(a_\bot)\}$ of 5 elements. We obtain a contradiction as this implies that $T \notin$ NO-MONO-TRI, since a structure encoding a clique of size 5 is not in NO-MONO-TRI. □

Lemma 2.12 *For every $n \geq 2$, $G_n \in$* NO-MONO-TRI.

PROOF. Colour the elements $a_\top$ and $a_\bot$ 'black' and the other elements 'white'. This is a valid colouring since the part of G_n coloured 'white' is a copy of the structure H_n, and encodes a graph that is triangle-free. □

By arguing as above, we immediately obtain the following.

Corollary 2.13 NO-MONO-TRI *is in MMSNP but not in CSP.*

Notice that among the problems that are in MMSNP but not in CSP, there are tractable problems (all the problems of Corollary 2.10 on the preceding page are first-order expressible hence in the complexity class L; *i.e.* deterministic logarithmic space) as well as intractable problem (NO-MONO-TRI is NP-complete, *cf.* Chapter 6). We shall provide in Chapter 6 further examples of such problems that are complete for the complexity classes NL, P and NP.

In Chapter 4, we shall take the approach that has been developed in this chapter one step further: we shall completely characterise those problems in MMSNP that are not in CSP where the underlying signature is arbitrary.

Chapter 3

Monotone Monadic SNP without inequalities

We introduce in detail Feder and Vardi's logic MMSNP in order to give a detailed proof of Feder and Vardi's result concerning the computational equivalence of (the problems expressed by) MMSNP with CSP.

In this chapter we give a detailed proof of Feder and Vardi's theorem (quoted in Chapter 2 as Theorem 2.7). The approach is not original as we follow the lines of the proof given in [16]. However Feder and Vardi's proof is rather short and hard to understand. This motivated me to expand on the original proof. The idea of the proof is the following. First, to notice that some sentences of MMSNP of a particular form, the so-called *conform sentences* define homomorphism problems: an example of such a sentence was given in the previous chapter for the problem 3-COL. Secondly, to transform every sentence of MMSNP into a sentence that is "as conform as possible" such that one can associate some canonical constraint satisfaction problem to it (notice that most probably this constraint satisfaction problem has a different signature). In the computational equivalence, one reduction is relatively straightforward, namely from the MMSNP problem to its canonical constraint satisfaction problem. However, this reduction is not onto: there are instances of the constraint satisfaction problem that do not correspond to any instance of the MMSNP problem. The key idea to circumvent this difficulty consists in transforming further the sentence of MMSNP into a *special form* where every negated conjunct is *biconnected*: this shall allow us to define the converse reduction (from the canonical constraint satisfaction problem to the MMSNP problem) as some canonical inversion of the reduction mentioned above (from the MMSNP problem to its canonical constraint satisfaction problem). We can then prove that problems might arise only for instances of the canonical constraint satisfaction problem that have *small cycles* (*i.e.* are of small *girth*: here, "small" is a function of the MMSNP sentence). There are well-known constructions in graph theory that allow one to build graphs for any fixed parameters, in particular, in the case where these parameters are the chromatic number and the girth. Feder and Vardi adapted a construction of Erdös in order to reduce, in randomized polynomial time, a given instance of a constraint satisfaction problem into an equivalent instance of high girth. Hence the reduction from the canonical constraint satisfaction problem to the MMSNP problem consists first in reducing any instance into an equivalent constraint satisfaction problem instance of sufficiently high girth via the generalisation of Erdös' construction; and, secondly, in applying the inverse of the canonical reduction mentioned above. Notice that it remains open whether this reduction can be derandomized.

This chapter is organised as follows. In Section 3.1, we introduce some basic notation, definition and example. In Section 3.2, we define the conform sentences of MMSNP and show that they correspond to problems in CSP. Next, in Section 3.3, we prove Feder and Vardi's theorem for our main example of a MMSNP problem (the problem NO-MONO-TRI) along the lines of Feder and Vardi's proof in order to help the reader to understand the proof in the general case. Then, in Section 3.4, we show how to transform any sentence of MMSNP into an equivalent sentence of the special form. Section 3.5 is devoted to the main part of the proof: namely, the reductions between a MMSNP problem and its canonical constraint satisfaction problem (associated with the special form of this MMSNP sentence). Finally, in Section 3.6, we give the proof of Theorem 2.7.

3.1 Preliminaries

3.1.1 Good sentences

In the following, we show how to rewrite any sentence of MMSNP as a "good sentence": informally, we remove redundant negated conjuncts and we enforce that for every first-order variable occurring in a negated conjunct, a full choice of validity for the monadic predicates is inherent.

Notation Let Φ be a sentence of MMSNP over the signature σ, that is a sentence of the following form

$$\exists \bar{M} \forall \bar{x} \bigwedge_i \neg(\alpha_i(\bar{R},\bar{x}) \wedge \beta_i(\bar{M},\bar{x})).$$

Let $\kappa(\Phi) = (M_1, M_2, \ldots)$ be the signature consisting of the monadic symbols occurring in Φ but not in σ (when this does not cause confusion, we write simply κ instead of $\kappa(\Phi)$). Set $\sigma' = \sigma \dot{\cup} \kappa$. Let $\gamma(\Phi)$ denote the set of negated conjuncts that occur in Φ. Let $\gamma \in \gamma(\Phi)$. Denote by X_γ, the set of first-order variables that occur in the negated conjunct γ.

Let Φ be a sentence of MMSNP. For any negated conjunct $\gamma = \neg(\alpha \wedge \beta)$ in $\gamma(\Phi)$:

(i) if an atom occurs once positively and once negatively in β then discard γ; and

(ii) if an atom occurs more than once in γ then remove all occurrences of this atom in γ but one.

The sentence hence obtained is clearly equivalent to the original.
From now on, we only ever consider sentences for which this transformation has been carried out.

A partial order over negated conjuncts Let X be a set of variables. We define a binary relation $\precsim_{\sigma'}$ over the set of conjunctions of atoms involving relational symbols from some signature σ'. Let δ_1 and δ_2 be two conjunctions of atoms involving relational symbols from σ' and variables from X. Let i be a bijective mapping of X to X. Denote by $i(\delta_1)$ the conjunction obtained by replacing every variable x occurring in δ_1 by its image via i. We set $\delta_1 \precsim_{\sigma'} \delta_2$ whenever there exists a bijective mapping i of X to X such that $i(\delta_1)$ is a subconjunction of δ_2. Clearly, this binary relation is a partial order. If $\delta_1 \precsim_{\sigma'} \delta_2$ then we say that δ_1 *is a subconjunction of* δ_2 *up to a renaming of variables.* If both $\delta_1 \precsim_{\sigma'} \delta_2$ and $\delta_2 \precsim_{\sigma'} \delta_1$ then we write $\delta_1 \sim_{\sigma'} \delta_2$. Note that $\sim_{\sigma'}$ is an equivalence relation.

This partial order induces a partial order over the negated conjuncts of a sentence of MMSNP. Let Φ be a sentence of MMSNP. Let $\gamma_1 = \neg(\alpha_1 \wedge \beta_1)$ and $\gamma_2 = \neg(\alpha_2 \wedge \beta_2)$, in $\gamma(\Phi)$, be two negated conjuncts. If $\alpha_1 \wedge \beta_1 \precsim_{\sigma'} \alpha_2 \wedge \beta_2$ then we write that γ_1 is a *sub-negated-conjunct (up to a renaming of variables)* of γ_2. If γ_1 is *not* a sub-negated-conjunct of γ_2 for any two distinct negated conjuncts γ_1 and γ_2 in $\gamma(\Phi)$ then we write that Φ is *simplified.*

Simplifying a sentence Let Φ be a sentence of MMSNP. Discard all the negated conjuncts γ in $\gamma(\Phi)$ that are not minimal with respect to the partial order defined previously, keeping only one occurrence of a negated conjunct for each equivalence class. Since up to a permutation of the variable names, there is a unique sentence obtained in this way, by an abuse of notation we speak of *the* sentence obtained from Φ by simplification, and we denote it by $\mathbf{Simp}(\Phi)$.

Lemma 3.1 *Let Φ be a sentence of MMSNP. Then* $\mathbf{Simp}(\Phi)$ *is a sentence of MMSNP that is simplified and is equivalent to Φ.*

PROOF. Let Φ be a sentence of MMSNP that is not simplified; *i.e.* there are two distinct negated conjuncts $\gamma_1 = \neg(\alpha_1 \wedge \beta_1)$ and $\gamma_2 = \neg(\alpha_2 \wedge \beta_2)$ in $\gamma(\Phi)$, and there exists a bijective mapping i of X_{γ_1} to X_{γ_2} such that $i(\alpha_1 \wedge \beta_1)$ is a subconjunction of $\alpha_2 \wedge \beta_2$. The sentence Φ is of the form:

$$\exists \bar{M} \forall \bar{x} (\phi \wedge \gamma_1 \wedge \gamma_2).$$

It is equivalent to:

$$\exists \bar{M} \forall \bar{x} \phi \wedge \forall \bar{x} \gamma_1 \wedge \forall \bar{x} \gamma_2.$$

Since i is a bijection, renaming the variables we obtain equivalently:

$$\exists \bar{M} \forall \bar{x} \phi \wedge \forall \bar{x} \neg (i(\alpha_1 \wedge i(\beta_1)) \wedge \forall \bar{x} \gamma_2.$$

The previous sentence is clearly equivalent to the following sentence:

$$\exists \bar{M} \forall \bar{x} \phi \wedge \forall \bar{x} (\neg (i(\alpha_1) \wedge i(\beta_1)) \wedge \neg (\alpha_2 \wedge \beta_2)).$$

We can rewrite it as follows:

$$\exists \bar{M} \forall \bar{x} \phi \wedge \forall \bar{x} \neg ((i(\alpha_1) \wedge i(\beta_1)) \vee (\alpha_2 \wedge \beta_2)).$$

Since $(i(\alpha_1) \wedge i(\beta_1))$ is a subconjunction of $(\alpha_2 \wedge \beta_2)$, we obtain equivalently:

$$\exists \bar{M} \forall \bar{x} \phi \wedge \forall \bar{x} \neg (i(\alpha_1) \wedge i(\beta_1)).$$

Renaming the variables via the inverse of the bijection i, we get:

$$\exists \bar{M} \forall \bar{x} \phi \wedge \forall \bar{x} \gamma_1.$$

The previous sentence is finally equivalent to

$$\exists \bar{M} \forall \bar{x} (\phi \wedge \gamma_1).$$

This sentence is clearly a sentence of MMSNP and is equivalent to the original sentence Φ, and can be obtained from Φ by discarding the negated conjunct γ_2 that is not minimal. Since **Simp**(Φ) is simplified by construction and can be obtained via an iteration of the above basic simplification, the result follows. □

In the following, we shall give some examples of this construction.

EXAMPLE. Recall the sentence Φ_1 of MMSNP that expresses the problem TRI-FREE introduced in Section 2.4.1:

$$\begin{aligned} &\forall x (\neg E(x,x)) \\ &\wedge \forall x \forall y \forall z (\neg (E(x,y) \vee E(y,x)) \vee \neg (E(x,z) \vee E(z,x)) \vee \neg (E(y,z) \vee E(z,y))). \end{aligned}$$

It is not simplified and contains in fact only two types of negated conjuncts apart from $\neg E(x,x)$:

- $\gamma_1 := \neg \ell_1(x,y,z) = \neg (E(x,y) \wedge E(x,z) \wedge E(y,z))$; and
- $\gamma_3 := \neg \ell_2(x,y,z) = \neg (E(x,y) \wedge E(z,x) \wedge E(y,z))$.

For example $\gamma_2 := \neg\ell_1(x,z,y) = \neg(E(x,y) \wedge E(x,z) \wedge E(z,y))$ is equivalent to γ_1: indeed:

$$(E(x,y) \wedge E(x,z) \wedge E(y,z)) \precsim_{\sigma_2} (E(x,y) \wedge E(x,z) \wedge E(z,y))$$

via the permutation (y,z); and

$$(E(x,y) \wedge E(x,z) \wedge E(z,y)) \precsim_{\sigma_2} (E(x,y) \wedge E(x,z) \wedge E(y,z))$$

via the inverse of the permutation (y,z) (that is (y,z) itself). Another example is $\gamma_4 := \neg\ell_1(z,y,x) = \neg(E(x,y) \wedge E(z,x) \wedge E(z,y))$ that is also equivalent to γ_1: indeed,

$$(E(x,y) \wedge E(x,z) \wedge E(y,z)) \precsim_{\sigma_2} (E(x,y) \wedge E(z,x) \wedge E(z,y))$$

via the permutation (x,z,y); and,

$$(E(x,y) \wedge E(z,x) \wedge E(z,y)) \precsim_{\sigma_2} (E(x,y) \wedge E(x,z) \wedge E(y,z))$$

via the permutation (z,x,y).
Hence, we finally have:

$$\mathbf{Simp}(\Phi_1) = \forall x \forall y \forall z \neg E(x,x) \wedge \neg\ell_1(x,y,z) \wedge \neg\ell_2(x,y,z).$$

As a second example, consider the following sentence Φ_2 of MMSNP that expresses the problem NO-MONO-TRI that we introduced in Section 2.4.1 (it is not exactly the sentence given there, but an equivalent sentence rewritten in a similar way as for the case of the problem TRI-FREE).

$$\begin{aligned} \exists C \forall x \forall y \forall z \quad & \neg E(x,x) \wedge \neg(\ell_1(x,y,z) \wedge w(x,y,z)) \wedge \neg(\ell_1(x,z,y) \wedge w(x,y,z)) \\ & \wedge \neg(\ell_2(x,y,z) \wedge w(x,y,z)) \wedge \neg(\ell_1(z,y,x) \wedge w(x,y,z)) \\ & \wedge \neg(\ell_1(y,x,z) \wedge w(x,y,z)) \wedge \neg(\ell_2(y,x,z) \wedge w(x,y,z)) \\ & \wedge \neg(\ell_1(y,z,x) \wedge w(x,y,z)) \wedge \neg(\ell_1(z,y,x) \wedge w(x,y,z)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge b(x,y,z)) \wedge \neg(\ell_1(x,z,y) \wedge b(x,y,z)) \\ & \wedge \neg(\ell_2(x,y,z) \wedge b(x,y,z)) \wedge \neg(\ell_1(z,y,x) \wedge b(x,y,z)) \\ & \wedge \neg(\ell_1(y,x,z) \wedge b(x,y,z)) \wedge \neg(\ell_2(y,x,z) \wedge b(x,y,z)) \\ & \wedge \neg(\ell_1(y,z,x) \wedge b(x,y,z)) \wedge \neg(\ell_1(z,y,x) \wedge b(x,y,z)), \end{aligned}$$

where:

$$w(x,y,z) := C(x) \wedge C(y) \wedge C(z) \text{ and } b(x,y,z) := \neg C(x) \wedge \neg C(y) \wedge \neg C(z).$$

We proceed as in the previous case and we get $\mathbf{Simp}(\Phi_2)$:

$$\begin{aligned}\forall x \forall y \forall z \quad & \neg E(x,x) \wedge \neg(\ell_1(x,y,z) \wedge w(x,y,z)) \wedge \neg(\ell_2(x,y,z) \wedge w(x,y,z)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge b(x,y,z)) \wedge \neg(\ell_2(x,y,z) \wedge b(x,y,z))\end{aligned}$$

▲

Let X be a set of variables. A conjunction β of positive or negative atoms involving the monadic symbols from κ and the variables from X is said to be a *complete colouring of* X *with respect to* κ if for any variable x in X and any predicate M in κ, either $M(x)$ occurs in β or $\neg M(x)$, but not both. Let Φ be a sentence of MMSNP. If β is a complete colouring of X_γ with respect to $\kappa(\Phi)$ for every forbidden conjunct $\gamma := \neg(\alpha \wedge \beta)$ in $\gamma(\Phi)$ then we say that Φ has *complete colourings*.

Let X be a set of variables. Let $\mathcal{K}$ be the set of complete colourings of one variable x in X with respect to κ. We call an equivalence class of $\mathcal{K}$ for $\sim_\kappa$ a κ-*colour*, or simply *colour* when this does not cause confusion.

Enforcing complete colouring on a sentence Let Φ be a sentence of MMSNP. For any negated conjunct $\gamma = \neg(\alpha \wedge \beta)$ in $\gamma(\Phi)$, if the β-part of γ is not a complete colouring of X_γ relative to κ then there exist a variable x in X_γ and a monadic symbol M in κ, such that neither $M(x)$ nor $\neg M(x)$ occur in β. Replace γ by the two following negated conjuncts $\neg(\alpha \wedge \beta \wedge M(x))$ and $\neg(\alpha \wedge \beta \wedge \neg M(x))$. Repeat this process until a fixed point is reached and denote the new sentence by $\mathbf{Comp}(\Phi)$.

Lemma 3.2 $\mathbf{Comp}(\Phi)$ *is a well-defined sentence of MMSNP that has complete colourings and that is equivalent to* Φ.

PROOF. $\mathbf{Comp}(\Phi)$ is well defined since a fixed point must be reached after finitely many steps ($\kappa(\Phi)$ is finite). $\mathbf{Comp}(\Phi)$ is a sentence of MMSNP equivalent to Φ because at each step the sentence obtained is a sentence of MMSNP and is equivalent to the sentence from the previous stage (note that $\exists \bar{M} \forall \bar{x} \gamma$ is logically equivalent to the sentence $\exists \bar{M} \forall \bar{x} \neg(\alpha \wedge \beta \wedge M(x)) \wedge \neg(\alpha \wedge \beta \wedge \neg M(x))$). □

We say that a sentence of MMSNP that is both simplified and has complete colourings is a *good* sentence of MMSNP.

Proposition 3.3 *Let Φ be a sentence of MMSNP. Then* $\mathbf{Simp}(\mathbf{Comp}(\Phi))$ *is a good sentence of MMSNP that is equivalent to Φ. Moreover*

$$\mathbf{Simp}(\mathbf{Comp}(\mathbf{Simp}(\Phi))) = \mathbf{Simp}(\mathbf{Comp}(\Phi)).$$

PROOF. By Lemma 3.2 Φ is equivalent to $\mathbf{Comp}(\Phi)$, which has complete colourings. By Lemma 3.1 $\mathbf{Comp}(\Phi)$ is equivalent to $\mathbf{Simp}(\mathbf{Comp}(\Phi))$, which is simplified. The latter also has complete colourings, since it is obtained by discarding some negated conjuncts from the former. This proves the first assertion.

The second assertion follows from the fact that if a simplification is carried out before completing the colourings, it can still be carried out afterwards. Let $\gamma_1 = \neg(\alpha_1 \wedge \beta_1)$ and $\gamma_2 = \neg(\alpha_2 \wedge \beta_2)$ be two distinct negated conjuncts from $\gamma(\Phi)$ such that $(\alpha_1 \wedge \beta_1) \precsim_{\sigma'} (\alpha_2 \wedge \beta_2)$ via some bijection i of X_{γ_1} to X_{γ_2}: *i.e.* γ_2 does not appear in $\mathbf{Simp}(\Phi)$. Moreover, assume that β_1 is not a full colouring, that is that there exists some variable x in X_{γ_1} and some monadic predicate M in $\kappa(\Phi)$ such that neither $M(x)$ nor $\neg M(x)$ occur in β_1. Then either $M(i(x))$ or $\neg M(i(x))$ or neither of them occur in β_2. Hence in the two first cases, either

$$(\alpha_1 \wedge \beta_1 \wedge M(x)) \precsim_{\sigma'} (\alpha_2 \wedge \beta_2) \text{ or } (\alpha_1 \wedge \beta_1 \wedge \neg M(x)) \precsim_{\sigma'} (\alpha_2 \wedge \beta_2),$$

that is, a completion of the colouring of γ_1 is a sub-negated-conjunct of γ_2 via i. In the third case

$$(\alpha_1 \wedge \beta_1 \wedge M(x)) \precsim_{\sigma'} (\alpha_2 \wedge \beta_2 \wedge M(x))$$
$$\text{and}$$
$$(\alpha_1 \wedge \beta_1 \wedge \neg M(x)) \precsim_{\sigma'} (\alpha_2 \wedge \beta_2 \wedge \neg M(x)),$$

that is, the completions of the colouring of γ_1 in the variable x and the monadic predicate M are respective sub-negated-conjuncts of the completions of γ_2 in the variable $i(x)$ and the monadic predicate M via i. Thus in any case, the completions of the colouring of γ_2 in **Comp**(Φ) do not appear in **Simp**$($**Comp**$(\Phi))$. □

Notice however that

$$\mathbf{Comp}(\mathbf{Simp}(\Phi)) = \mathbf{Simp}(\mathbf{Comp}(\Phi))$$

does not hold in general, since completing a simplified sentence might yield new simplifications. We shall provide an example for this in the following.

EXAMPLE. The sentence **Simp**(Φ_1) is a trivial example of a good sentence as it is a first-order formula.

Consider as another example of a good sentence the sentence **Comp**$($**Simp**$(\Phi_2))$:

$$\begin{aligned}\exists C \quad & \forall x \forall y \forall z \neg(\ell_1(x,y,z) \wedge w(x,y,z)) \wedge \neg(\ell_2(x,y,z) \wedge w(x,y,z)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge b(x,y,z)) \wedge \neg(\ell_2(x,y,z) \wedge b(x,y,z)) \\ & \wedge \neg(E(x,x) \wedge C(x)) \wedge \neg(E(x,x) \wedge \neg C(x)).\end{aligned}$$

Indeed, in this particular case, there is no need for further simplification. However, this shall not be the case for our next example.

Consider the sentence that expresses the problem TRI-FREE-TRI introduced in Section 2.4.1: it can be rewritten as the following equivalent sentence Φ_3 of MMSNP:

$$\begin{aligned}\exists R \exists W \exists B \forall x \forall y \forall z \quad & \neg(R(x) \wedge W(x)) \wedge \neg(R(x) \wedge B(x)) \wedge \neg(W(x) \wedge B(x)) \\ & \wedge \neg E(x,x) \\ & \wedge \neg(E(x,y) \wedge R(x) \wedge R(y)) \wedge \neg(E(y,x) \wedge R(x) \wedge R(y)) \\ & \wedge \neg(E(x,y) \wedge W(x) \wedge W(y)) \wedge \neg(E(y,x) \wedge W(x) \wedge W(y)) \\ & \wedge \neg(E(x,y) \wedge B(x) \wedge B(y)) \wedge \neg(E(y,x) \wedge B(x) \wedge B(y)) \\ & \wedge \neg\ell_1(x,y,z) \wedge \neg\ell_1(x,z,y) \wedge \neg\ell_2(x,y,z) \wedge \neg\ell_1(z,y,x) \\ & \wedge \neg\ell_1(y,x,z) \wedge \neg\ell_2(y,x,z) \wedge \neg\ell_1(y,z,x) \wedge \neg\ell_1(z,y,x)\end{aligned}$$

We want to find a good sentence of MMSNP expressing TRI-FREE-TRI. First simplify the sentence; $\mathbf{Simp}(\Phi_3)$ is

$$\begin{array}{ll}\exists R \exists W \exists B \forall x \forall y \forall z & \neg(R(x) \wedge W(x)) \wedge \neg(R(x) \wedge B(x)) \wedge \neg(W(x) \wedge B(x)) \\ & \wedge \neg E(x,x) \\ & \wedge \neg(E(x,y) \wedge R(x) \wedge R(y)) \wedge \neg(E(x,y) \wedge W(x) \wedge W(y)) \\ & \wedge \neg(E(x,y) \wedge B(x) \wedge B(y)) \wedge \neg \ell_1(x,y,z) \wedge \neg \ell_2(x,y,z)\end{array}$$

Then, complete its colourings and simplify the sentence to obtain the good sentence $\mathbf{Simp}(\mathbf{Comp}(\mathbf{Simp}(\Phi_3)))$ of MMSNP as follows:

$$\begin{array}{ll}\exists R \exists W \exists B \forall x \forall y \forall z & \neg(R(x) \wedge W(x) \wedge \neg B(x)) \wedge \neg(R(x) \wedge W(x) \wedge B(x)) \\ & \wedge \neg(R(x) \wedge \neg W(x) \wedge B(x)) \wedge \neg(\neg R(x) \wedge W(x) \wedge B(x)) \\ & \wedge \neg(E(x,x) \wedge r(x)) \wedge \neg(E(x,x) \wedge w(x)) \wedge \neg(E(x,x) \wedge b(x)) \\ & \wedge \neg(E(x,y) \wedge r(x) \wedge r(y)) \wedge \neg(E(x,y) \wedge w(x) \wedge w(y)) \\ & \wedge \neg(E(x,y) \wedge b(x) \wedge b(y)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge r(x) \wedge w(y) \wedge b(z)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge r(x) \wedge b(y) \wedge w(z)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge w(x) \wedge r(y) \wedge b(z)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge w(x) \wedge b(y) \wedge r(z)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge b(x) \wedge r(y) \wedge w(z)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge b(x) \wedge w(y) \wedge r(z)) \\ & \wedge \neg(\ell_2(x,y,z) \wedge r(x) \wedge w(y) \wedge b(z))\end{array}$$

where:

$$\begin{array}{l} r(x) := R(x) \wedge \neg W(x) \wedge \neg B(x) \\ b(x) := \neg R(x) \wedge W(x) \wedge \neg B(x) \\ w(x) := \neg R(x) \wedge \neg W(x) \wedge B(x) \end{array}$$

We prove that $\mathbf{Comp}(\mathbf{Simp}(\Phi))$ is not necessarily simplified. Consider the case of Φ_3; and, note that, *e.g.*,

$$\neg(R(x) \wedge W(x) \wedge \neg B(x))$$

is a negated conjunct of $\mathbf{Comp}(\mathbf{Simp}(\Phi_3))$ while,

$$\neg(\ell_1(x,y,z) \wedge R(x) \wedge W(x) \wedge \neg B(x))$$

is a sub-negated-conjunct of some negated conjuncts of **Comp**(**Simp**(Φ_3)). ▲

From now on, we shall only consider good sentences of MMSNP.

3.1.2 Structure induced by a negated conjunct

Let Φ be a sentence of MMSNP. Let $\neg(\alpha \wedge \beta) = \gamma \in \gamma(\Phi)$ be a negated conjunct of this sentence.

Denote by G_α the σ-structure induced as follows:

- its universe $|G_\alpha|$ consists of the variables that occur in γ; and
- for every r-ary relation symbol R in σ, define R^{G_α} as follows: for every r-tuple $\bar{x}$ of elements of $|G_\alpha|$, $R(\bar{x})$ holds in G_α if, and only if, it occurs in α.

We call the σ-structure G_α *the structure induced by* α.

Recall that $\sigma' = \sigma \cup \kappa$. In the following we usually denote σ'-structures with a $'$ (as in G') to distinguish them from σ-structures. Let G' be a σ'-structure. Recall that the *reduct* of G' over σ is the σ-structure G that; has the same domain as G'; and, as relation $R^{G'}$ for every relation symbol R in σ. Conversely, we say that G' is an *extension* of G over σ'.

G'_γ is the extension of G_α over σ' defined as follows.

- for any monadic symbol M in κ, define $M^{G'_\gamma}$ as follows: for any x in $|G'_\gamma|$, $M(x)$ holds in G'_γ if, and only if, $M(x)$ occurs as an atom in β.

We call the σ'-structure G'_γ *the structure induced by* γ.

3.1.3 Connected and biconnected structures

We shall be concerned with a generalisation of the graph-theoretic notions of connectivity and biconnectivity for arbitrary relational structures.

Let t be some finite tuple: we denote by $\{t\}$ the set of elements occurring in t.

Let A be a σ-structure and u and v be two elements of $|A|$. If there exist $n > 0$ and n tuples $t_0, t_1, \ldots, t_{n-1}$ of respective arities $r_{i_0}, r_{i_1}, \ldots, r_{i_{n-1}}$ such that:

- $R_{i_0}(t_0), \ldots, R_{i_{n-2}}(t_{n-2})$ and $R_{i_{n-1}}(t_{i_n-1})$ hold in A;
- $u \in \{t_0\}$, $v \in \{t_{n-1}\}$; and
- for any $0 \leq j \leq n-2$, there exists u_j in $|A|$ such that $u_j \in \{t_j\}$ and $u_j \in \{t_{j+1}\}$,

then we say that $t_0, t_1, \ldots, t_{n-1}$ form a *path* of length n from u to v.
The structure A is said to be *connected* if, and only if, for any distinct u and v in $|A|$, there exists a path from u to v.
Let A be a connected σ-structure. A is said to be *1-connected* if, and only if, there exists some u in $|A|$ and a pair (P_0, P_1) of induced substructures of A satisfying

- $|P_0| \cap |P_1| = \{u\}$;
- $|P_0| \cup |P_1| = |A|$;
- $\text{size}(P_i) := \sum_{R \in \sigma} |R^{P_i}| \geq 1$, for $i = 0, 1$; and
- for every r-ary symbol R in σ, if $R^A(t)$ holds then either $R^{P_0}(t)$ holds or $R^{P_1}(t)$ holds, but not both.

We say that (P_0, P_1) forms a *decomposition* of A in the *articulation point u*. If such a decomposition does not exist and that A is connected then A is said to be *biconnected.*
A σ-structure A is said to be *antireflexive* if, and only if, for every r-ary symbol R in σ, for any $t \in |A|^r$ such that $R^A(t)$ holds, all components of t are distinct.
A structure A is said to be *monotuple* if $\sum_{R \in \sigma} |R^A| = 1$ (note that a monotuple connected structure is biconnected).
Let C be a monotuple structure such that $R^C(t)$ holds for some r-ary symbol in σ and some $t \in |C|^r$. If every element of the domain of C is mentioned in the tuple t and that some element u occurs at least twice in the tuple t then we say that C is a *1-cycle*. In other words, the structure C consists of a single tuple, which contains an element u occurring at least twice: we call u an *articulation point* of the 1-cycle C.
Let $n > 1$. Let C be a structure such that,

- there exists n substructures $P_0, \ldots, P_{n-1}$ of C with $|C| = \bigcup_{i=0}^{i=n-1} |P_i|$ such that:

- for any $0 \leq i \leq n-2$, there exist some $x_i \in |C|$ with $|P_i| \cap |P_{i+1}| = \{x_i\}$;
- there exist some x_{n-1} such that $|P_0| \cap |P_{n-1}| = \{x_{n-1}\}$; and
- for any $0 \leq i < j < n$ such that $i+1 \neq j \mod n$, $|P_i| \cap |P_j| = \emptyset$; and

- for any $0 \leq i < n$, P_i is monotuple and there exists some r_i-ary symbol R_i in σ and some $\bar{y} \in |P_i|^{r_i}$ such that $R^{P_i}(\bar{y})$ and $|\bar{y}| = r_i$ and $|P_i| = \{\bar{y}\}$,

We say that C is an *n-cycle*: furthermore, the x_i's are called *articulation points* of C; and, the $R_i(\bar{y})$'s the *tuples of the cycle C*.
Let A be some σ-structure that contains a cycle as a substructure. Define the *girth* of A as the least integer $n \geq 1$ such that there exist an n-cycle C that is a substructure of A. We write $\mathbf{girth}(A) := n$. We extend this definition to *acyclic* σ-structures (structures that do not contain any cycle as a substructure) by setting $\mathbf{girth}(A) := \infty$ for any acyclic structure A.

We shall need the following technical result later in this chapter. The proof of this result can be found in [16]: it is an adaptation from Erdös' construction of graphs of arbitrary girth and chromatic number. This result is used to reduce an instance of a problem in CSP to an instance without any "small" cycles: indeed, the converse transformation we mentioned earlier (from the canonical constraint satisfaction problem back to the MMSNP problem) can be guaranteed to be a reduction for such instances.

Lemma 3.4 *Let $g,d > 0$. For every σ-structure A, there exists a σ-structure B with:*

- $|B| = |A|^{\delta_{g,d}}$ *(where $\delta_{g,d}$ is a function dependent only on g and d);*
- $\mathbf{girth}(B) \geq g$*;*
- $B \longrightarrow A$*; and*
- *for every σ-structure T with $|T| \leq d$,*

$$A \longrightarrow T \textit{ if, and only if, } B \longrightarrow T.$$

Furthermore, B can be constructed from A in randomised polynomial time.

The definition of a 'randomised polynomial time reduction' can be found in Appendix A.

3.2 Conform sentence and CSP

Let Φ be a sentence of MMSNP. We say that Φ is *conform* if, and only if, every negated conjunct $\gamma \in \gamma(\Phi)$ is either of the form

1. $\gamma = \neg(\alpha \wedge \beta)$; and, the structure induced by α is connected, monotuple and antireflexive; or of the form

2. $\gamma = \neg(\beta)$, where $|X_\gamma| = 1$ and β is a complete colouring of X_γ with respect to κ.

EXAMPLE. The following conform sentence can be considered as the realisation of the abstract problem 3-COL.

$$\begin{aligned} \exists M_1 \exists M_2 \forall x \forall y \quad & \neg(E(x,y) \wedge M_1(x) \wedge M_2(x) \wedge M_1(y) \wedge M_2(y)) \\ & \wedge \neg(E(x,y) \wedge \neg M_1(x) \wedge M_2(x) \wedge \neg M_1(y) \wedge M_2(y)) \\ & \wedge \neg(E(x,y) \wedge M_1(x) \wedge \neg M_2(x) \wedge M_1(y) \wedge \neg M_2(y)) \\ & \wedge \neg(\neg M_1(x) \wedge \neg M_2(x)) \end{aligned}$$

The two monadic predicates M_1 and M_2 encode 4 colours, the fourth of which is forbidden by the last negated conjunct. ▲

Lemma 3.5 *Every problem in CSP is expressible by a good sentence of MMSNP. Moreover, every problem expressed by a conform sentence of MMSNP is in CSP.*

PROOF. We start with the first assertion. Let T be a σ-structure. Let κ be the signature that consists of monadic symbols M_i that do not occur in σ, where i ranges from 1 to $|T|$. The following sentence Φ_T defines the problem $CSP(T)$ and belongs to MMSNP:

$$\begin{aligned} \exists \bar{M} \forall \bar{x} \quad & \neg\Big(\bigwedge_{M_i \in \kappa} \neg M_i(x_0)\Big) \wedge \bigwedge_{M_i, M_j \in \kappa (i \neq j)} \neg(M_i(x_0) \wedge M_j(x_0)) \\ & \wedge \bigwedge_{R_k \in \sigma} \Big(R_k(\bar{x}) \Rightarrow \bigvee_{t \in R_k^T} \varphi_{k,t}(\bar{x})\Big) \end{aligned}$$

where: x_0 is a variable of $\bar{x}$, R_k has arity r_k, $t = (t_1, t_2, \ldots, t_{r_k})$ and $\varphi_{k,t}(\bar{x}) := M_{t_1}(x_1) \wedge \ldots \wedge M_{t_{r_k}}(x_{r_k})$.

The existential monadic predicates $\bar{M}$ represent the elements of $|T|$ and the first part of the sentence states that they associate one element of $|T|$ with every element of an input structure A. The last part of the sentence says that this assignment is a homomorphism. This sentence is not necessarily good. By Proposition 3.3 it can be transformed into a good sentence that is logically equivalent.

We now prove the second assertion. Let Φ be a conform sentence of MMSNP. Construct the σ-structure T_Φ defined as follows:

1. $|T_\Phi|$ consists of those κ-colours that are not forbidden by the sentence (that is, that do not correspond to a negated conjunct of type (2) in the definition of a conform sentence): *i.e.* set
$$|T_\Phi| := \{k \ \kappa\text{-colour} | \forall \gamma \in \gamma(\Phi) \gamma \not\sim_\kappa \neg k(x)\}; \text{ and}$$
2. for any r-ary symbol R in σ and any r-tuple $t = (k_{i_1}, k_{i_2}, \ldots, k_{i_r})$ of elements of $|T_\Phi|$, set $R(t)$ to hold, if, and only if, there is no negated conjunct γ in $\gamma(\Phi)$ such that $\gamma \sim_{\sigma'} \gamma_t$, where
$$\gamma_t = \neg(R(x_{i_1}, x_{i_2}, \ldots, x_{i_r}) \wedge k_{i_1}(x_{i_1}) \wedge k_{i_2}(x_{i_2}) \wedge \ldots \wedge k_{i_r}(x_{i_r})).$$

We now prove that T_Φ is a template for the problem expressed by the sentence Φ.

Let A be a σ-structure and A' an extension of A that mentions only colours from $|T_\Phi|$ (that is, the colours allowed by the sentence Φ). We can clearly restrict ourselves to such extensions: indeed, $A \models \Phi$ if, and only if, there exists an extension A' of A to σ', such that
$$A' \models \forall \bar{x} \bigwedge \gamma.$$

This is equivalent to: there exists an extension A' that mentions only colours from $|T_\Phi|$ such that
$$A' \models \forall \bar{x} \bigwedge_{\gamma \text{ not of type (2)}} \gamma.$$

Note that such an extension A' induces a mapping h of $|A|$ to $|T_\Phi|$: map any v in $|A|$ to its κ-colour in A'. Conversely, such a mapping h induces an extension A' of A over σ' as follows: for any v in $|A|$, set $k(v)$ to hold in A', where k is the

κ-colour given by $h(v) = k$.

We show that h is a homomorphism if, and only if,

$$A' \models \forall \bar{x} \bigwedge_{\gamma \text{ not of type (2)}} \gamma.$$

Suppose that h is a homomorphism. Let γ be one of the negated conjuncts of type (1) in $\gamma(\Phi)$. It follows that

$$\gamma \sim_{\sigma'} \gamma_t = \neg(R(x_{i_1}, x_{i_2}, \ldots, x_{i_r}) \wedge k_{i_1}(x_{i_1}) \wedge k_{i_2}(x_{i_2}) \wedge \ldots \wedge k_{i_r}(x_{i_r}))$$

via some bijection i (renaming of the variables). Let $\pi : X_\gamma \to |A'|$ be some assignment. We must have $A' \models \gamma(\bar{x}/\pi(\bar{x}))$: otherwise, we would have

$$A' \models R(\pi \circ i(x_{i_1}), \ldots, \pi \circ i(x_{i_r})) \wedge k_{i_1}(\pi \circ i(x_{i_1})) \wedge \ldots \wedge k_{i_r}(\pi \circ i(x_{i_r})).$$

Hence, there would be a tuple $t = (\pi \circ i(x_{i_1}), \ldots, \pi \circ i(x_{i_r}))$ such that $R^A(t)$ holds and $R(h(t))$ does not hold in T_Φ, where $h(t) = (k_{i_1}, \ldots, k_{i_r})$. A contradiction.

Conversely, assume that

$$A' \models \forall \bar{x} \bigwedge_{\gamma \text{ not of type (2)}} \gamma.$$

Let t be a r-tuple of elements of $|A|$, let $X = \{x_{i_1}, x_{i_2}, \ldots, x_{i_r}\}$ be a set of variables and let $\pi : X \to |A|$ be a mapping given by $x_{i_j} \mapsto t[j], (1 \leq j \leq r)$. If $A \models R(\bar{x}/\pi(\bar{x}))$ then there can not be a negated conjunct of type (1) γ in $\gamma(\Phi)$ such that $\gamma \sim_{\sigma'} \gamma_{h(t)}$: otherwise, we would have

$$A' \not\models \gamma_{h(t)}(\bar{x}/h \circ \pi \circ i^{-1}(\bar{x})),$$

where $i : X_\gamma \to X_{\gamma_{h(t)}}$ is a bijective mapping witnessing that $\gamma \sim_{\sigma'} \gamma_{h(t)}$. Therefore, by construction of T_Φ, we must have $T_\Phi \models R(\bar{x}/h \circ \pi(\bar{x}))$. □

3.3 A problem in CSP computationaly equivalent to No-Mono-Tri

The remaining sections of this chapter will lead to a proof of Feder and Vardi's theorem. This proof might seem rather involved to some readers: so, in the present section we construct a problem in CSP that is computationaly equivalent to No-Mono-Tri (in the line of the forthcoming proof). We have seen in the example in the last paragraph of Section 3.1.1 that the problem No-Mono-Tri is expressed by the following good sentence of MMSNP:

$$\begin{aligned}\exists C \forall x \forall y \forall z \quad & \neg(\ell_1(x,y,z) \wedge w(x,y,z)) \wedge \neg(\ell_2(x,y,z) \wedge w(x,y,z)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge b(x,y,z)) \wedge \neg(\ell_2(x,y,z) \wedge b(x,y,z)) \\ & \wedge \neg(E(x,x) \wedge C(x)) \wedge \neg(E(x,x) \wedge \neg C(x)).\end{aligned}$$

Notice that in this sentence, replacing $\forall x \forall y \forall z$ by $\forall x \forall y \forall z (x \neq y) \wedge (x \neq z) \wedge (y \neq z)$, leads to a sentence that is logically equivalent. Let τ be the signature consisting of three symbols: two ternary symbols R_1 and R_2 and a unary symbol R_3. Consider the following sentence Ψ over τ:

$$\begin{aligned}\exists C \forall x \forall y \forall z \quad & \neg(R_1(x,y,z) \wedge w(x,y,z)) \wedge \neg(R_2(x,y,z) \wedge w(x,y,z)) \\ & \wedge \neg(R_1(x,y,z) \wedge b(x,y,z)) \wedge \neg(R_2(x,y,z) \wedge b(x,y,z)) \\ & \wedge \neg(R_3(x) \wedge C(x)) \wedge \neg(R_3(x) \wedge \neg C(x)).\end{aligned}$$

Let $\mathbf{P}_\Psi$ be the problem expressed by Ψ. We refer the reader to Appendix A for the definition of an interpretation. No-Mono-Tri can be reduced to the problem $\mathbf{P}_\Psi$ via the following interpretation Π of τ in σ of width one:

$$\Pi := (\phi_1, \phi_2, \phi_3)$$

where

$$\phi_1 := x \neq y \wedge x \neq z \wedge y \neq z \wedge \ell_1(x,y,z),$$

$$\phi_2 := x \neq y \wedge x \neq z \wedge y \neq z \wedge \ell_2(x,y,z) \text{ and } \phi_3 := E(x,x).$$

Note that the sentence Ψ is conform, thus by Lemma 3.5 $\mathbf{P}_\Psi$ belongs to CSP and according to the proof of this lemma, the τ-structure T defined as follows can be considered as a template for $\mathbf{P}_\Psi$:

- $|T| := \{b, w\}$;
- $R_1^T := |T|^3 \setminus \{(b,b,b),(w,w,w)\}$;
- $R_2^T := |T|^3 \setminus \{(b,b,b),(w,w,w)\}$; and
- $R_3^T := \emptyset$.

Let $\Pi^{-1} := (\psi)$ be the first-order interpretation of σ in τ of width one, where

$$\begin{aligned}\psi \;:=\; & \exists z(R_1(x,y,z) \vee R_1(x,z,y) \vee R_1(z,x,y) \vee R_2(x,y,z) \vee R_2(y,z,x) \vee R_2(z,x,y)) \\ & \vee (x = y \wedge (R_3(x)).\end{aligned}$$

We work over different signatures in the following: so, when we give a structure, we write its signature as a superscript (as in A^τ).

Fact 3.6 *Let A^τ be an antireflexive τ-structure and let $B^\tau := \Pi(\Pi^{-1}(A^\tau))$. If* $\mathbf{girth}(A^\tau) > 3$ *then $B^\tau \models \Psi$ if, and only if, $A^\tau \models \Psi$.*

PROOF. A^τ is a substructure of B^τ: hence, the direct implication holds (problems in MMSNP are closed under inverse homomorphism).

We now prove the converse implication. Let $A^{\tau'}$ be a valid extension of A^τ with respect to Ψ: that is, $A^{\tau'}$ is a model of the first-order part of Ψ. Consider the extension $B^{\tau'}$ of B^τ induced by the extension $A^{\tau'}$ of A^τ (recall that the structures share the same domain as we consider width one interpretations). Call informally 'new tuples' the tuples of B^τ that were not present in A^τ. We only need to check the validity of the extension over those new tuples: there are different cases to consider.

1. A new tuple belongs to R_1: that is, there exist some a, b and c such that $R_1^{B^\tau}(a,b,c)$ holds and $R_1^{A^\tau}(a,b,c)$ does not hold. Since $R_1^{B^\tau}(a,b,c)$ holds then $a \neq b \wedge a \neq c \wedge b \neq c \wedge \ell_1(a,b,c)$ holds in $\Pi^{-1}(A^\tau)$. In particular, $E(a,b)$ holds in $\Pi^{-1}(A^\tau)$ and $a \neq b$: thus, according to the definition of

Π^{-1}, there exist some d_1 in $|A^\tau|$ such that some tuple t_1 holds in A^τ and involves the elements d_1, a and b. Similarly for $E(b,c)$ and $E(a,c)$, there exist two further elements, say d_2 and d_3 and two tuples t_2 and t_3; where the tuple t_2 involves d_2, b and c; and, the tuple t_3 involves d_3, a and c. We now prove that the tuples t_1, t_2 and t_3 coincide. We must have $d_1 \neq a$ and $d_1 \neq b$ (otherwise t_1 is a 1-cycle contradicting the fact that $\mathbf{girth}(A^\tau) > 3$). Similarly, we must have $d_2 \neq b, d_2 \neq c, d_3 \neq a$ and $d_3 \neq c$. If t_1 is different from t_2 then $d_1 \neq d_2$ (otherwise t_1 and t_2 would form a 2-cycle). Similarly for the tuples t_2 and t_3 and the tuples t_1 and t_3, this implies $d_2 \neq d_3$ and $d_1 \neq d_3$. Thus, if the tuples t_1, t_2 and t_3 were pairwise distinct then we would have a 3-cycle (which can not happen since $\mathbf{girth}(A^\tau) > 3$). So, we proved that t_1, t_2 and t_3 are the same tuple. This enforces $d_1 = c, d_2 = a$ and $d_3 = b$. Hence, we now know that there exists only one tuple in A^τ that involves a, b and c. Since $a \neq b \wedge a \neq c \wedge b \neq c \wedge \ell_1(a,b,c)$ holds in $\Pi^{-1}(A^\tau)$ this tuple can only correspond to a tuple in some relation R^{A^τ} whose interpretation in σ includes the interpretation of R_1 in σ up to a renaming of variables. The relation R_1 is the only one that satisfies this criteria: it follows that $R_1^{A^\tau}(a,b,c)$ holds. This yields a contradiction.

2. A new tuple belongs to R_2. This case is similar to the previous one.

3. A new tuple belongs to R_3: that is, there exists some a such that that $R_3^{B^\tau}(a)$ holds and $R_3^{A^\tau}(a)$ does not hold. Since $R_3^{B^\tau}(a)$ holds then $E^{\Pi^{-1}(A^\tau)}(a,a)$ holds. There can not be any element d such that the first part of ψ is satisfied: this would mean that a tuple that involves d and a repeated twice would occur in $R_1^{A^\tau}$ or $R_2^{A^\tau}$ (recall that A^τ is antireflexive as by assumption $\mathbf{girth}(A^\tau) > 3$). Hence, according to the definition of Π^{-1}, we must have $R_3^{A^\tau}(a)$. This yields a contradiction.

There are no new tuples, thus the converse implication holds. □

Remark. Note that we really proved the following. If $\mathbf{girth}(A^\tau) > 3$ then B^τ and A^τ coincide. However, this shall not be true in the general case.

It follows from this fact and from Lemma 3.4 that $\mathbf{P}_\Psi = CSP(T)$ can be reduced to NO-MONO-TRI via a randomised polynomial time reduction: first, use the randomised reduction from the lemma to get an equivalent structure of girth greater than 3; then, use Π^{-1}. Hence, we can state the following corollary of Feder and Vardi's theorem for the problem NO-MONO-TRI.

Corollary 3.7 *There exists a structure T such that:*

- NO-MONO-TRI *reduces to CSP(T) via qfps; and*
- *CSP(T) reduces to* NO-MONO-TRI *in randomised polynomial-time.*

notation. Let $\bar{x} := x_0, x_1, \ldots, x_{n-1}$. We write $\forall_{\neq}\bar{x}\phi$, as an abbreviation for:

$$\forall \bar{x}(\bigwedge_{0 \leq i < j < n} x_i \neq x_j \phi).$$

Remark. Before we move onto the proof of Feder-Vardi's theorem, we shall make some remarks on the sentence Φ (used previously as the defining sentence of the problem NO-MONO-TRI). Recall that Φ is the following sentence.

$$\begin{aligned} \exists C \forall x \forall y \forall z \quad & \neg(\ell_1(x,y,z) \wedge w(x,y,z)) \wedge \neg(\ell_2(x,y,z) \wedge w(x,y,z)) \\ & \wedge \neg(\ell_1(x,y,z) \wedge b(x,y,z)) \wedge \neg(\ell_2(x,y,z) \wedge b(x,y,z)) \\ & \wedge \neg(E(x,x) \wedge C(x)) \wedge \neg(E(x,x) \wedge \neg C(x)). \end{aligned}$$

Note that the sentence Φ has the following key properties:

1. Φ is a good sentence;
2. The sentence obtained by replacing $\forall x \forall y \forall z$ in Φ by $\forall_{\neq} x, y, z$ is equivalent to Φ; and
3. for any negated conjunct $\gamma = \neg(\alpha \wedge \beta)$ in $\gamma(\Phi)$, the structure induced by α is biconnected and the structure induced by γ is connected.

(1) is necessary to ensure that Lemma 3.5 can be used to prove that the new sentence Ψ expresses a problem that is a CSP. Each symbol in τ corresponds to the α-part of a negated conjunct of Φ. (2) is necessary to ensure that (3) makes sense together with Lemma 3.4 in the proof of Fact 3.6: that is, for structures of girth

greater than 3, Π^{-1} is a reduction from $\mathbf{P}_\Psi$ to NO-MONO-TRI. Indeed, the existence of cycles is derived from the fact that we have biconnected structures (which make sense only if the variables can not be identified). In Section 3.4.1, we shall introduce the notion of a 'collapsed' sentence of MMSNP, that corresponds to (2); and, in Section 3.4.2, the notion of a 'biconnected' sentence of MMSNP, that corresponds to (3).

3.4 Transforming a sentence into a special form

In this Section we transform a sentence of MMSNP into a special form: this special form is used in the proof of Theorem 2.7. There are two main steps: first, we *collapse* the sentence of MMSNP; that is, we transform the original sentence into an equivalent sentence where the sequence of universal first order quantifiers $\forall x \forall y \ldots$ are replaced by the variant $\forall_{\neq} x, y$ whose semantic is "for every choice of distinct elements of the structure x, y,..." (*cf.* previous Section). Secondly, we *split* each negated conjunct of this collapsed sentence into its biconnected components. This transformation is quite trivial in the case of a negated conjunct that has disjoint components; it involves introducing new nullary predicates (basically it corresponds to a transformation of a MMSNP problem into the union of connected MMSNP problems). However, it is slightly more subtle in the case of a 1-connected negated conjunct that is not biconnected and that is split along some *articulation point*; it involves the introduction of a new monadic predicate.

3.4.1 Collapsed sentences

Later on we shall need the notion of a biconnected negated conjunct; this makes sense only if we deal with sentences where the first-order variables within a negated conjunct can not be identified. In other words, we want to restrict ourselves to injective assignment when checking the satisfiability of a sentence.

Let Φ be a sentence of MMSNP. If the sentence obtained by replacing $\forall \bar{x}$ in Φ by $\forall_{\neq} \bar{x}$ is equivalent to Φ then we say that Φ is *collapsed.*

collapsing a sentence Let Φ be a sentence of MMSNP. Let γ be some negated conjunct occurring in Φ and let m be some mapping of X_γ to X_γ. Denote by $m(\gamma)$

the negated conjunct obtained from γ by replacing in γ every first order variable x in X_γ by its image $m(x)$ and removing redundancies. For every negated conjunct γ in Φ and for every mapping $m : X_\gamma \to X_\gamma$, add to the sentence all negated conjuncts $m(\gamma)$ that are not trivially true [1]. Simplify this sentence and denote it by $\mathbf{Coll}(\Phi)$.

Lemma 3.8 *If Φ is a good sentence of MMSNP then* $\mathbf{Coll}(\Phi)$ *is a good sentence of MMSNP that is collapsed and equivalent to Φ.*

PROOF. Notice that if γ had a complete colouring relatively to κ then so has $m(\gamma)$. Thus $\mathbf{Coll}(\Phi)$ is a good sentence of MMSNP.

By construction, for any σ-structure A, if $A \models \mathbf{Coll}(\Phi)$ then $A \models \Phi$ (as $\mathbf{Coll}(\Phi)$ is obtained from Φ by adding negated conjuncts).

Conversely, suppose that $A \models \Phi$. Then, there exists an extension $A^{\sigma'}$ of A to σ' such that for any assignment $\pi : X \to |A|$, $A^{\sigma'} \models \phi(\bar{x}/\pi(\bar{x}))$, where ϕ is the quantifier-free first-order part of the formula Φ. We shall show now for any assignment $\pi : X \to |A|$, $A^{\sigma'} \models \psi(\bar{x}/\pi(\bar{x}))$, where ψ is the quantifier-free first-order part of $\mathbf{Coll}(\Phi)$. Let γ be some negated conjunct in Φ and $m : X_\gamma \to X_\gamma$ be some mapping such that $m(\gamma)$ occurs in ψ. Let $\pi : X \to |A|$. Since $A \models \Phi$, it follows that $A^{\sigma'} \models \gamma(\bar{x}/\pi \circ m(\bar{x}))$. Hence, $A^{\sigma'} \models m(\gamma)(\bar{x}/\pi(\bar{x}))$. Thus, $A \models \mathbf{Coll}(\Phi)$.

It remains to show that the construction yields a collapsed sentence: that is, we show that after this construction, we can restrict ourselves to assignments to the first-order variables that do not identify any two variables occurring in the same negated conjunct. More precisely, $A \models \mathbf{Coll}(\Phi)$ if, and only if, there exists an extension $A^{\sigma'}$ such that for any negated conjunct γ in $\gamma(\mathbf{Coll}(\Phi))$ and for any one-to-one $\pi : X_\gamma \to |A|$, we have $A^{\sigma'} \models \gamma(\bar{x}/\pi(\bar{x}))$. The direct implication is clear. For the converse, we have to show, that this holds for assignments that are non-injective. Let $\pi : X_\gamma \to X_\gamma$ be a non-injective mapping. Denote by $R(\pi) := \{(x,y) \in X_\gamma^2 | \pi(x) = \pi(y)\}$ the equivalence relation associated with π. Take some representatives for each equivalence class. Denote by $\dot{x}$ the representative of x. Then, let $m : X_\gamma \to X_\gamma$ be the mapping defined as follows. $m(x) = \dot{x}$. By assumption, $A^{\sigma'} \models m(\gamma)(\bar{x}/\pi \circ m(\bar{x}))$. Hence, $A^{\sigma'} \models \gamma(\bar{x}/\pi(\bar{x}))$. Thus, the result follows. □

[1] By this we mean a negated conjunct satisfying the condition (*ii*) in the first paragraph of Section 3.1.1.

To illustrate the above construction, consider the following example.

EXAMPLE. Let Ψ be the following sentence of MMSNP:

$$\begin{aligned}\exists C\forall x\forall y\forall z \quad & \neg(\ell_1(x,y,z)\wedge w(x,y,z))\wedge\neg(\ell_2(x,y,z)\wedge w(x,y,z))\\ & \wedge\neg(\ell_1(x,y,z)\wedge b(x,y,z))\wedge\neg(\ell_2(x,y,z)\wedge b(x,y,z)).\end{aligned}$$

Then $\mathbf{Coll}(\Psi)$ is the following sentence:

$$\begin{aligned}\exists C\forall x\forall y\forall z \quad & \neg(\ell_1(x,y,z)\wedge w(x,y,z))\wedge\neg(\ell_2(x,y,z)\wedge w(x,y,z))\\ & \wedge\neg(\ell_1(x,y,z)\wedge b(x,y,z))\wedge\neg(\ell_2(x,y,z)\wedge b(x,y,z))\\ & \wedge\neg(E(x,x)\wedge C(x))\wedge\neg(E(x,x)\wedge\neg C(x)).\end{aligned}$$

Notice that this is a sentence of MMSNP that expresses the problem NO-MONO-TRI; and, moreover that this is the sentence we used previously to find an equivalent problem in CSP. ▲

3.4.2 Biconnected sentences

Let Φ be a sentence of MMSNP. If G_α is connected (respectively biconnected) for every negated conjunct $\gamma = \neg(\alpha\wedge\beta)$ in $\gamma(\Phi)$ then we say that Φ is *connected* (respectively *biconnected*).

We extend the logic MMSNP by allowing existential quantification over nullary predicates and call MMSNP with nullary predicates this new logic; all the notions introduced in this chapter; that is, the notion of a simplified sentence, of a sentence with full colourings etc, are naturally extended.

Lemma 3.9 *Let Φ be a good sentence of MMSNP. Then, there exists a good biconnected sentence Ψ in MMSNP with nullary predicates that is equivalent to Φ.*

The remainder of this section is devoted to the proof of this result.

Let Φ be a good sentence of MMSNP. We shall construct an equivalent sentence Ψ that is good and biconnected. There are different cases to consider. From now on, we denote by Φ the sentence equivalent to the original MMSNP sentence and that has been obtained up to this point of the construction, and we denote by Ψ the (to be shown) logically equivalent new sentence. As long as there exists a negated conjunct γ that is not biconnected, we proceed as follows, depending on γ's form:

1. **disjoint case**: $\gamma = \neg(\delta_0(\bar{x}) \wedge \delta_1(\bar{y}))$ with $\{\bar{x}\}$ and $\{\bar{y}\}$ disjoints.
   We introduce a new existential nullary predicate p (*i.e.* a Boolean variable) and replace γ by $(\delta_0(\bar{x}) \Rightarrow p) \wedge (\delta_1(\bar{y}) \Rightarrow \neg p)$.

   Fact 3.10 *The new sentence is equivalent.*

   PROOF. Let A^σ be a σ-structure. Suppose that $A^\sigma \models \Phi$. Let $\sigma'' := \sigma' \dot{\cup} \{p\}$. Then there exists an extension $A^{\sigma'}$ of A^σ such that $A^{\sigma'} \models \forall\bar{x}\forall\bar{y}\phi$, where ϕ denotes the quantifier-free first-order part of Φ. In particular, $A^{\sigma'} \models \forall\bar{x}\forall\bar{y}\gamma$. Thus it can not be the case that there exist some $\pi : X_\gamma \rightarrow |A|$ such that both $A^{\sigma'} \models \delta_0(\bar{x}/\pi(\bar{x}))$ and $A^{\sigma'} \models \delta_1(\bar{x}/\pi(\bar{x}))$. Extend $A^{\sigma'}$ as follows: if there exist some $\pi : X_\gamma \rightarrow |A|$ such that $A^{\sigma'} \models \delta_0(\bar{x}/\pi(\bar{x}))$ holds then set $p^{A^{\sigma''}} :=$ **true**, otherwise set $p^{A^{\sigma''}} :=$ **false**. Clearly $A^{\sigma''}$ witnesses that $A^\sigma \models \Psi$. Conversely, assume that $A^\sigma \models \Psi$. Then there exist some extension $A^{\sigma''}$ such that $A^{\sigma''} \models \forall\bar{x}\forall\bar{y}\psi$, where ψ denotes the quantifier-free first-order part of Ψ. Let $A^{\sigma'}$ denotes the reduct of $A^{\sigma''}$ to σ'. We finally show that $A^{\sigma'} \models \forall\bar{x}\forall\bar{y}\phi$: w.l.o.g. $p^{A^{\sigma''}} =$ **true** thus for any assignment $\pi : X_\gamma \rightarrow |A|$, we have $A^{\sigma'} \models \neg\delta_1(\bar{y})$ hence $A^{\sigma'} \models \neg\gamma$. □

2. **1-connected case**: $\gamma = \neg(\delta_0(\bar{x}, z) \wedge \delta_1(\bar{y}, z))$, with $\bar{x}$ and $\bar{y}$ disjoints.
   We replace γ by $\delta_0(\bar{x}, z) \Rightarrow M_\gamma(z)) \wedge (\delta_1(\bar{y}, z) \Rightarrow \neg M_\gamma(z))$ and introduce a new existential monadic predicate M_γ.

   Fact 3.11 *The new sentence is equivalent.*

   PROOF. $A \models \Phi$ if, and only if, there exists an extension $A^{\sigma'}$ of A on σ' such that for each negated conjunct γ in $\gamma(\Phi)$ and for every assignment $\pi : X_\gamma \rightarrow |A|$, $A^{\sigma'} \models \gamma(\bar{x}/\pi(\bar{x}))$. Let σ'' be $\sigma' \cup \{M_\gamma\}$. Extend $A^{\sigma'}$ on σ'' as follows. set

   $$M_\gamma^{A^{\sigma''}} := \{z \in |A| \text{ such that } A^{\sigma'} \models \exists\bar{x}\delta_0(\bar{x}, z)\}.$$

Now, let $\pi_0 : X_{\delta_0} \to |A|$. By definition of $A^{\sigma''}$, we have

$$A^{\sigma''} \models \neg(\delta_0(\bar{x}/\pi_0(\bar{x}), z/\pi_0(z)) \wedge \neg M_\gamma(z/\pi_0(z))).$$

Let $\pi_1 : X_{\delta_1} \to |A|$. We must have

$$A^{\sigma''} \models \neg(\delta_1(\bar{x}/\pi_1(\bar{x}), z/\pi_1(z)) \wedge M_\gamma(z/\pi_1(z))),$$

otherwise,

$$A^{\sigma''} \models \delta_1(\bar{x}/\pi_1(\bar{x}), z/\pi_1(z)) \wedge M_\gamma(z/\pi_1(z)).$$

Hence, by definition of $M_\gamma^{\sigma''}$, we would have $A^{\sigma'} \models \exists\bar{x}\delta_0(\bar{x}, z/\pi_1(z))$, that is there exists some $\pi_0 : X_{\delta_0} \to |A|$ with $\pi_0(z) := \pi_1(z)$ such that $A^{\sigma'} \models \delta_0(\bar{x}/\pi(\bar{x}), z/\pi(z))$. Hence, we would have some $\pi : X_\gamma \to |A|$ induced by π_0 and π_1 such that

$$A^{\sigma'} \models \delta_0(\bar{x}/\pi(\bar{x}), z/\pi(z)) \wedge \delta_1(\bar{y}/\pi(\bar{y}), z/\pi(z)),$$

a contradiction. It follows that $A \models \Psi$.

Conversely, $A \models \Psi$ if, and only if, there exists some extension $A^{\sigma''}$ of A over σ'' such that for all negated conjunct γ in $\gamma(\Psi)$, and for all $\pi : X_\gamma \to |A|$, $A^{\sigma''} \models \gamma(\bar{x}/\pi(\bar{x}))$. In particular, for any $\pi : X_\gamma \to |A|$.

$$A^{\sigma''} \models \neg(\delta_0(\bar{x}/\pi(\bar{x}), z/\pi(z)) \wedge \neg M_\gamma(z/\pi(z)))$$

and

$$A^{\sigma''} \models \neg(\delta_1(\bar{x}/\pi(\bar{x}), z/\pi(z)) \wedge M_\gamma(z/\pi(z))).$$

It follows that

$$A^{\sigma''} \models \neg(\delta_0(\bar{x}/\pi(\bar{x}), z/\pi(z)) \wedge \neg M_\gamma(z/\pi(z))) \wedge \neg(\delta_1(\bar{x}/\pi(\bar{x}), z/\pi(z)) \wedge M_\gamma(z/\pi(z))).$$

Let $A^{\sigma'}$ be the reduct of $A^{\sigma''}$ over σ'. Then,

$$A^{\sigma'} \models \neg\delta_0(\bar{x}/\pi(\bar{x}), z/\pi(z)) \text{ if } A^{\sigma''} \models M_\gamma(z/\pi(z)$$

$$\text{and } A^{\sigma'} \models \neg\delta_1(\bar{x}/\pi(\bar{x}), z/\pi(z)) \text{ if } A^{\sigma''} \models \neg M_\gamma(z/\pi(z).$$

It follows that there exists an extension of $A^{\sigma'}$ of A over σ' such that for any $\pi : X_\gamma \to |A|, A^{\sigma'} \models \neg(\delta_0(\bar{x}/\pi(\bar{x}), z/\pi(z)) \wedge \delta_1(\bar{y}/\pi(\bar{y}), z/\pi(z)))$. Hence, $A \models \Phi$. □

Once every negated conjunct is biconnected, we transform the sentence into a good sentence; *i.e.* we complete the colouring and simplify the sentence. This concludes the proof of Lemma 3.9.

Together with Lemma 3.8 this yields the following corollary (since if one assumes Φ to be collapsed in Lemma 3.9 then the sentence Ψ is also collapsed).

Corollary 3.12 *Let Φ be a good sentence of MMSNP. Then there exists a sentence of MMSNP with nullary predicates equivalent to Φ, that is good, collapsed and biconnected (we call this sentence the* special form *of Φ).*

Remark on MMSNP with nullary predicates Notice that a problem defined by a sentence with nullary predicates simply corresponds to a finite union of problems expressed by sentences without nullary predicates. Lemma 3.5 can be generalised to conform sentences of MMSNP with nullary predicates; indeed, we can do a case analysis on the values of these nullary predicates and for each of these cases apply the lemma and construct a template T_i, and make T the disjoint union of these templates. However, we must ensure that the cases are disjoint for the non-uniform CSP problem as well, and that disconnected instances are in the problem if, and only if, there is an homomorphism into a single T_i. Hence, we add a binary symbol R to σ and set $R^T := \bigcup_i |T_i|^2$ and for every instance A we set $R^A := |A|^2$. Note that this can be achieved via qfps from the constraint satisfaction problem to the MMSNP problem and via a polynomial-time reduction from the MMSNP problem to the constraint satisfaction problem.

3.5 Main part of the reduction

The idea of the reduction is as follows: given a problem expressed by a sentence Φ over σ (of the special form given by the previous corollary) we consider the problem over the signature τ, where τ is induced by the α-parts of the negated conjuncts occurring in Φ; one new relational symbol R_α is introduced for every equivalence class of $\alpha(\Phi)$ for $\sim_\sigma$; and, its arity is the number of different variables occurring in α. Now, choose one α in each equivalence class and let

$$\phi_\alpha := \bigwedge_{x_i \neq x_j \in X_\alpha} x_i \neq x_j \wedge \alpha.$$

This provides an interpretation of τ in σ of width one: $\Pi = (\phi_\alpha | R_\alpha \in \tau)$.

Replace every α-part $\alpha(\bar{x})$ of the negated conjuncts in Φ by the corresponding symbol $R_\alpha(\bar{x})$ Denote this sentence by Ψ.
Note that Ψ is conform and that $A^\sigma \models \Phi$ if, and only if $\Pi(A^\sigma) \models \Psi$. However, we are also interested in the reduction from the problem expressed by Ψ to the problem expressed by Φ. Let B^τ be a τ-structure. If $R_\alpha(t)$ holds in B^τ for some tuple of elements t suitable in length then we want $\alpha(t)$ to hold in the structure A^σ obtained from B^τ. In other words, we just reverse the interpretation Π as follows: for every R in σ, let

$$\phi_R(\bar{x}) = \bigvee_{R(\bar{x}) \text{ occurs in } \alpha(\bar{x},\bar{y})} \exists \bar{y} \alpha(\bar{x}, \bar{y})$$

This provides an interpretation of σ in τ of width one: $\Pi^{-1} = (\phi_R | R \in \sigma)$ (note that for simplicity in the above, we did not take into account the fact that we might have to rename variables). We want:
$B^\tau \models \Psi$ if, and only if $\Pi^{-1}(B^\tau) \models \Phi$.
This would clearly hold if:

$$\Pi(\Pi^{-1}(B^\tau)) = B^\tau$$

would hold, but this is not the case in general. This is where the notion of high girth is needed. Indeed, each tuple in a relation in the τ-structures $\Pi(\Pi^{-1}(B^\tau))$ corresponds either:

1. to a monotuple connected substructure of $\Pi^{-1}(B^\tau)$; or

2. to a non-monotuple biconnected substructure of $\Pi^{-1}(B^\tau)$.

So, according to case (2): different tuples in B^τ could give rise to some tuples in $\Pi^{-1}(B^\tau)$; these latter tuples might satisfy some α in $\alpha(\Phi)$; and, it may yield a

tuple in $\Pi(\Pi^{-1}(B^\tau))$ that is not present in B^τ.
Let g_Φ be the maximal number of atoms occurring in an α-part of Φ. If B^τ has girth greater than g_Φ and $R(t)$ holds in $\Pi^{-1}(B^\tau)$ (for some relation symbol R of arity r in σ and some r-tuple t) then t must be induced according to case (1): *i.e.* t must be contained in some tuple t_α in some relation R_α in B^τ.
Hence, we have to enforce the following: if a colouring β_1 is forbidden by a negated conjunct γ_1, whose α-part α_2 is a subconjunction of a strictly larger α-part of some other negated conjunct γ_2 then the constraint given by β_1 is propagated to α_2. In the following, we amend our construction of Ψ to make sure that this is the case.

Construction of Ψ. First, for every negated conjunct $\gamma_1 = \neg(\alpha_1 \wedge \beta_1)$, $\gamma_2 = \neg(\alpha_2 \wedge \beta_2)$ in $\gamma(\Phi)$ and permutation $m : X_{\gamma_1} \to X_{\gamma_2}$ such that $m(\alpha_1)$ is a subconjunction of α_2; we add the following negated conjunct to Φ:

$$\gamma_{1,2} = \neg(\alpha_2 \wedge m(\beta_1)).$$

Secondly, we complete the colouring of this new sentence and denote it by $\tilde{\Phi}$.
Note that $\tilde{\Phi}$ is equivalent to Φ and also that $\tilde{\Phi}$ is not necessarily simplified. However, $\tilde{\Phi}$ has all the other properties that a sentence obtained via Corollary 3.12 would have; it is biconnected and collapsed and has complete colourings. Denote by Ψ the formula obtained from $\tilde{\Phi}$ by replacing every α by the corresponding symbol R_α in τ.
Note that Ψ is conform.

Lemma 3.13 *Let Φ be a sentence of MMSNP with nullary predicates that is of the special form (given by Corollary 3.12). There exist a signature τ, an interpretation Π of width one from τ in σ, an interpretation Π^{-1} of width one from σ in τ and a conform sentence Ψ over τ such that the following holds:*

(i) for any σ-structure A^σ, $A^\sigma \models \Phi$ if, and only if $\Pi(A^\sigma) \models \Psi$; and

(ii) for any τ-structure B^τ of girth greater than g_Φ, $\Pi(\Pi^{-1}(B^\tau)) \models \Psi$ if, and only if $B^\tau \models \Psi$.

PROOF. Let τ be the signature induced by $\tilde{\Phi}$, let Π be the interpretation of width one of τ in σ let Π^{-1} be the corresponding interpretation of σ in τ and let Ψ be defined as previously.

(i) is clear. We now prove (ii). By monotonicity of Ψ and because B^τ can be embedded in $\Pi(\Pi^{-1}(B^\tau))$, clearly $\Pi(\Pi^{-1}(B^\tau)) \models \Psi$ implies $B^\tau \models \Psi$. Now, suppose that $B^\tau \models \Psi$. Then there exists some extension $B^{\tau'}$ of B^τ to $\tau' := \sigma \cup \kappa$ such that for each negated conjunct γ in $\gamma(\Psi)$, and for every assignment $\pi : X_\gamma \to |B^\tau|$ $B^{\tau'} \models \gamma(\bar{x}/\pi(\bar{x}))$ holds. Let $A^{\tau'}$ be the extension of $\Pi(\Pi^{-1}(B^\tau))$ to τ' constructed as follows: the reduct of $A^{\tau'}$ over κ is the same as the reduct of $B^{\tau'}$ over κ. We show that this extension witnesses that $\Pi(\Pi^{-1}(B^\tau)) \models \Psi$.

Note that, we have to check only those tuples that were not present in B^τ. We call informally "new tuples" such tuples. Since B^τ has girth greater than g_Φ, a new tuple must be the projection over some indices of a longer tuple present in B^τ. Indeed, any k tuples t_i in $R_i^{B^\tau}$ of arity r_i give rise to an acyclic substructure of B^τ because B^τ has girth $g_\Phi > k$. Therefore, a new tuple t_1 in some $R_{\alpha_1}^{\Pi(\Pi^{-1}(B^\tau))}$ must be induced by some tuple t_2 in $R_{\alpha_2}^{B^\tau}$, where α_1 and α_2 belong to $\alpha(\Phi)$ and $\alpha_1 \precsim_\sigma \alpha_2$ (recall that Φ is biconnected). Hence, if there exist $\gamma_1 = \neg(R_{\alpha_1} \wedge \beta_1)$ in $\gamma(\Psi)$ and $\pi_1 : X_{\gamma_1} \to |B^\tau|$ such that

$$\Pi(\Pi^{-1}(B^\tau)) \models R_{\alpha_1}(\bar{x}/\pi(\bar{x})).$$

Then, there exist a negated conjunct $\neg(\alpha_2 \wedge \beta_2)$ in $\gamma(\Phi)$, a one-to-one mapping $m : X_{\alpha_1} \to X_{\alpha_2}$ such that $m(\alpha_1)$ is a subconjunction of α_2; and, moreover, there exists $\pi_2 : X_{\alpha_2} \to |B^\tau|$ such that $\pi_2 \circ m = \pi_1$ over X_{α_1} and $B^\tau \models R_{\alpha_2}(\bar{y}/\pi_2(\bar{y}))$. By construction, some negated conjunct obtained from $\gamma_{1,2}$ is present in Ψ; that is, a negated conjunct of the following form:

$$\gamma_{1,2}^{\beta} = \neg(R_{\alpha_2} \wedge m(\beta_1) \wedge \beta).$$

Since $B^\tau \models \Psi$, it follows that for all such β:

$$B^{\tau'} \models \neg(R_{\alpha_2}(\bar{y}/\pi_2(\bar{y}) \wedge m(\beta_1)(\bar{y}/\pi_2(\bar{y})) \wedge \beta(\bar{y}/\pi_2(\bar{y}))).$$

Hence,

$$B^{\tau'} \models \neg m(\beta_1)(\bar{y}/\pi_2(\bar{y})),$$

and it follows that:

$$B^{\tau'} \models \neg\beta_1(\bar{x}/\pi_1(\bar{x})).$$

Therefore,

$$A^{\tau} \models \neg\beta_1(\bar{x}/\pi_1(\bar{x})).$$

Finally, we get:

$$\Pi(\Pi^{-1}(B^{\tau})) \models \Psi.$$

□

3.6 A Proof of Feder and Vardi's theorem

Combining together the results of this chapter, we can now give a proof of Feder and Vardi's theorem.

Theorem 3.14 (Feder and Vardi)
Every problem in CSP is expressible by a sentence of MMSNP. Every problem $\mathbf{P}_\Phi$ expressible by a sentence Φ of MMSNP is equivalent to a problem $CSP(T_\Phi)$ in CSP: $\mathbf{P}_\Phi$ reduces to $CSP(T_\Phi)$ in polynomial time; and, $CSP(T_\Phi)$ reduces to $\mathbf{P}_\Phi$ in randomised polynomial time.

PROOF.CSP is contained in MMSNP by lemma 3.5.

By Corollary 3.12, we can assume that Φ is a sentence of MMSNP with nullary predicates that is good, collapsed and biconnected. Then, it follows from Lemma 3.13 that there exists a conform sentence Ψ (with possibly some nullary predicates) over a signature τ such that: the problem expressed by Φ reduces to that of Ψ via a qfp of width one; and, the problem expressed by Ψ, when restricted to τ-structures of girth greater than g_Φ, reduces to the problem expressed by Φ via a positive first-order interpretation of width one.
It follows from the remark on nullary predicates on the end of Subsection 3.4.2 that the problem $\mathbf{P}_\Psi$ (the problem expressed by Ψ) is computationaly equivalent to a problem $CSP(T_\Psi)$ in CSP: $\mathbf{P}_\Psi$ reduces to $CSP(T_\Psi)$ via a polynomial-time reduction; and, $CSP(T_\Psi)$ reduces to $\mathbf{P}_\Psi$ via a qfp.
It follows from Lemma 3.4 that the problem $CSP(T_\Psi)$ reduces to its restriction

over τ-structures of girth greater than g_Φ in randomised polynomial-time. This restricted constraint satisfaction problem reduces to $\mathbf{P}_\Psi$ via a trivial qfp that shall not decrease the girth; it consists only in dropping one relation symbol (the symbol introduced to enforce that a disconnected instance would map into a single template). Thus, altogether we provided a randomised polynomial-time reduction from $CSP(T_\Psi)$ to $\mathbf{P}_\Phi$. □

In [16], the authors mention the possibilities of using *quasi-random graphs* to *derandomize* the reduction from the constraint satisfaction problem to the problem expressed by a sentence of MMSNP problem. In other words,

Question 3.15 *is it possible to have polynomial-time reductions in Theorem 2.7 in both directions?*

An unsuccessful attempt along this line lead me to the following question:

Question 3.16 *which problems in MMSNP are not in CSP?*

We know that such problems exist (*cf.* Section 2.4). Moreover, if hopefully I could provide some exact characterisation for the latter question, I could possibly answer negatively to a restriction of the former question. Indeed, proving a negative result for any polynomial-time reduction seems to be rather tricky in front of the immmense diversity that such reductions have to offer. However, if we restrict ourselves to some particular meaningful reductions, say first-order projections, we could hopefully prove that some property that ensures that a problem is not a CSP could be conserved by such transformations. As a matter of fact, I have not yet answered even a restriction of the former question. I answered however the latter. This rather innocent looking question has lead me to a proof involving objects and notions which I consider personally as interesting by themselves. I hope to convince the reader in the next chapter, which is fully devoted to this characterisation. If the reader was not yet convinced of the interest of Question 3.16, we hope to eventually convince him in Chapter 5. There, we shall relate in some detail some recent and independent results by Tardiff and Nešetřil (*cf.* [45]), which can be obtained as a corollary of our forthcoming characterisation.

Chapter 4

Forbidden patterns problems

We introduce a new class of combinatorial problems; the class of forbidden patterns problems. These problems correspond exactly to Feder and Vardi's logic MMSNP. We introduce the concept of a representation for such problems and introduce recolourings. We show that representations and recolourings generalise the notions of structures and homomorphisms. Finally we introduce a normal form for forbidden patterns problems and characterise exactly the forbidden patterns problems that are not homomorphism problems. The proof is constructive in the sense that given a representation, we can *compute* its normal form; according to its normal form, either we are in the case of a CSP and we can *compute* a template, or we can *compute* a family of counter-examples, the so-called witness family, that show that the problem is not a CSP.

In this chapter we introduce a class of combinatorial problems, the so-called *forbidden patterns problems*. We shall see that they correspond exactly to the class of problems captured by the logic MMSNP. Thus according to Feder and Vardi's theorem these problems are computationally equivalent to problems in CSP. However, some of these problems are not in CSP as a corollary of our results from Section 2.4. We have seen that if a problem is expressed by a sentence of MMSNP that is conform then it is in CSP. Among the forbidden patterns problems, we would like to be able to completely differentiate those problems that are in CSP from those that are not, that is, answer Question 3.16 in this new setting. A forbidden pattern shall be given by a *representation*, which in some sense generalises the notion of a structure. Formulating a problem captured by a sentence of MMSNP in terms of a representation, that is, a more algebraic formulation, is not just a technical reformulation that we do for its own sake: it shall allow us to develop a notion of a *morphism for representations*, the so-called *recolouring*. Once we have a notion of morphism we can look at the induced notions of *retract* and therefore of *core*, since we deal with finite objects. Another important notion we shall introduce is the notion of a *template for a representation*; a structure induced by some particular forbidden patterns. The homomorphism problem induced by this structure contains the forbidden patterns problem considered. Then we shall adapt techniques used in Section 3.4.2 which together with the notion of a *core of representations* shall lead us to *normal representations*. Generalising the idea of the proof given in Section 2.4, in order to show that a particular forbidden patterns problem is not in CSP, we shall construct particular families of structures, the so-called *witness families*. We are then able to describe a *generic construction* of a witness family for a forbidden patterns problem, whenever it can be given by a *normal connected representation that is not conform*. Thus we partially answer the above question. Later, we are able to answer the question for any representation, by defining the *set of normal (connected) representations* of a representation.

This chapter is organised as follows. In Section 4.1, we introduce the notions of a representation, a recolouring and a forbidden patterns problem. We also show that the forbidden patterns problems correspond exactly to the problems captured by sentences of MMSNP. In Section 4.2 we investigate the notion of a retract. We start by recalling this notion in the case of structures before extending it to coloured structures and finally introducing it for representations. Section 4.3 deals with the notion of a template of a representation. Next in Section 4.4 we provide tools to construct biconnected representations, adapting the technique used in the proof of Feder and Vardi's theorem in the previous chapter. In Section 4.5 we introduce the key notion of a normal representation and give numerous examples. In Section 4.6, we define witness families for a forbidden patterns problem and show that if a problem has such a family then it can not be in CSP. Then, we describe a generic way of constructing witness families for representations. Finally, in Section 4.7, we state our main result, *i.e.* an exact characterisation of those forbidden patterns problems that are not in CSP (provided they are connected). We illustrate this result by numerous examples. The remainder of the chapter is devoted to the general case, *i.e.* to representations that are not necessarily connected.

4.1 Preliminaries

In this section, we start by introducing the notion of a coloured structure and of a homomorphism for coloured structures, the so-called *colour preserving homomorphisms*. Next in Subsection 4.1.2 we provide various examples to illustrate these notions; these examples are rather numerous as we shall need them later to provide examples of representations. In Subsection 4.1.3 we introduce the notion of a *representation* together with a new combinatorial problem, the so called *forbidden patterns problem* associated to a given representation. We introduce the key notion of a *recolouring* between representations and show that it is a morphism for representations and that, moreover, the existence of a recolouring between two given representations implies the inclusion of the problems they define; thus we obtain a result similar to Proposition 2.1. Next we illustrate these newly introduced notions by various examples. Finally, in Section 4.1.4 we provide two technical lemmas showing that the logic MMSNP captures exactly FP, the class of forbidden patterns problem.

4.1.1 Finite coloured structures and colour preserving homomorphisms

Let μ be a finite set. We call the elements of μ *colours*. A finite *μ-coloured σ-structure* consists of a finite σ-structure A, together with a mapping $c_\mu^A : |A| \to \mu$. We write (A, c_μ^A). We say that (A, c_μ^A) is *connected* (respectively *biconnected*) whenever A is connected (respectively biconnected). Let (A, c_μ^A) and (B, c_μ^B) be two μ-coloured σ-structures. A *colour preserving homomorphism* of (A, c_μ^A) to (B, c_μ^B) is a homomorphism $A \xrightarrow{h} B$ that preserves the colourings of A and B, *i.e.* such that $c_\mu^B \circ h = c_\mu^A$, and we write, $(A, c_\mu^A) \xrightarrow{h} (B, c_\mu^B)$. If there exists some mapping h such that $(A, c_\mu^A) \xrightarrow{h} (B, c_\mu^B)$ then we write $(A, c_\mu^A) \longrightarrow (B, c_\mu^B)$. If it is not the case that $(A, c_\mu^A) \longrightarrow (B, c_\mu^B)$ then we write $(A, c_\mu^A) \not\longrightarrow (B, c_\mu^B)$. We shall make use of diagrams to illustrate definitions and proofs in the following. If $(A, c_\mu^A) \xrightarrow{h} (B, c_\mu^B)$, we draw the following [1].

$$\begin{tikzcd} A \arrow[r, "h"] \arrow[d, "c^A_\mu"'] & B \arrow[dl, "c^B_\mu"] \\ \mu & \end{tikzcd}$$

When h is a surjective colour preserving homomorphism, we write

$$(A, c^A_\mu) \overset{h}{\twoheadrightarrow} (B, c^B_\mu).$$

When h is an injective colour preserving homomorphism, we write

$$(A, c^A_\mu) \overset{h}{\hookrightarrow} (B, c^B_\mu).$$

If $(A, c^A_\mu) \overset{h}{\hookrightarrow} (B, c^B_\mu)$ then we say that (A, c^A_μ) is a subcoloured structure of (B, c^B_μ) (Note that it may be the case that (A, c^A_μ) is not an induced subcoloured structure of (B, c^B_μ)).

A *colour preserving isomorphism* is a bijective colour preserving homomorphism whose inverse is a colour preserving homomorphism. If $(A, c^A_\mu) \xrightarrow{h} (B, c^B_\mu)$ and h is a colour preserving isomorphism then we write $(A, c^A_\mu) \approx (B, c^B_\mu)$. We denote by $\mathrm{STRUC}_\mu(\sigma)$ the class of all finite μ-coloured σ-structures. To avoid having to use too heavy a notation, when the set of colours is clear from the context, we shall not specify it, as in (A, c^A). Moreover, we shall speak of a homomorphism of (A, c^A) to (B, c^B) as meaning a colour preserving homomorphism.

Notice moreover that the composition of two colour preserving homomorphisms is itself a colour preserving homomorphism. As for the case of structures, we have an identity homomorphism associated with any coloured structure (A, c^A), induced by the identity map over $|A|$, which we shall denote $id_{(A,c^A)}$. One can therefore speak of the *category of finite μ-coloured σ-structures.*

In the next subsection, we introduce various σ_2-structures and coloured σ_2-structures and discuss the existence of homomorphisms and colour preserving homomorphisms between them: we shall need this later to build further examples of problems captured by sentences of MMSNP.

[1] A μ-colouring c^A_μ of a structure A can be seen as a homomorphism of A to K_μ, the complete structure with domain μ, *cf.* remark on the end of Subsection 4.3.2.

4.1.2 Examples

Some σ_2-structures. Recall that $\sigma_2 := \{E\}$, where E is a binary relation symbol. σ_2-structures can be considered as an encoding of finite directed graphs (possibly with self-loops). Denote by DC_n, $n > 1$, the following σ_2-structure (DC standing for directed cycle).

- $|DC_n| := \{0, 1, \ldots, n-1\}$; and
- for any elements x, y in $|DC_n|$, $E(x,y)$ holds if, and only if, $x + 1 = y \mod n$.

Denote by C_n, $n > 1$, the following σ_2-structure.

- $|C_n| := \{0, 1, \ldots, n-1\}$; and
- for any elements x, y in $|C_n|$, $E(x,y)$ holds if, and only if, $x + 1 = y \mod n$ or $y + 1 = x \mod n$.

Moreover, set C_1 and DC_1 to be the structure with a single element x such that $E(x,x)$ holds. Some of these structures are depicted in Figure 4.1 (the nodes not being labelled for the sake of simplicity). In the case of the structures C_n, we write a double arrow to denote that the relation E^{C_n} is symmetric.

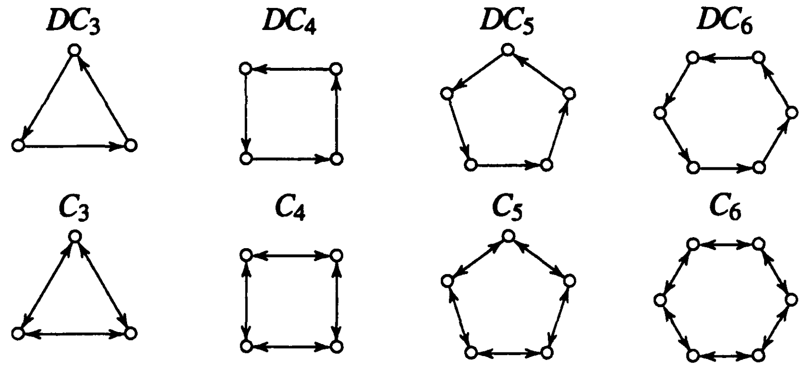

Figure 4.1: Directed Cycles and Cycles

1. Clearly, any graph maps homomorphically into DC_1.

2. For any $n > 0$, there is a natural bijective homomorphism $h_{n,n}$ of DC_n to C_n, where $h_{n,n}$ is the identity mapping over the set $\{0,1,\dots,n-1\}$. However notice that $h_{n,n}$ is not an isomorphism except for $n \leq 2$ since its inverse is not a homomorphism. Thus to sum up, we have the following; for any $n \geq 2$,

$$\begin{array}{rccc} h_{n,n}: & DC_n & \hookrightarrow & C_n \\ & x & \mapsto & x \end{array} \qquad \begin{array}{rccc} h_{n,n}^{-1}: & C_n & \not\rightarrow & DC_n \\ & x & \mapsto & x \end{array}$$

3. Let $n,m > 0$ be such that m divides n. Consider the mapping

$$\begin{array}{rccl} h_{n,m}: & \{0,1,\dots,n-1\} & \rightarrow & \{0,1,\dots,m-1\} \\ & x & \mapsto & x \mod m \end{array}$$

   It is easy to check that $DC_n \xrightarrow{h_{n,m}} DC_m$ and that $C_n \xrightarrow{h_{n,m}} C_m$.

4. If $m < n$ then there is no homomorphism of DC_m to DC_n.

5. Moreover, notice that if $n,m > 1$ are such that $n \neq m$ and n and m are relatively prime then $DC_n \not\rightarrow DC_m$.

6. The case of cycles is different; even cycles are homomorphically equivalent. Let

$$\begin{array}{rccc} f_2: & \{0,1\} & \rightarrow & \{0,1,\dots,2p-1\} \\ & x & \mapsto & x \end{array}$$

   We have $C_{2p} \xrightarrow{h_{2p,2}} C_2$ and $C_2 \xrightarrow{f_2} C_{2p}$. Notice that $h_{2p,2} \circ f_2 = id_2$. However, the two structures are not isomorphic.
   Since any even length cycle $C_{2p}, p > 0$, is homomorphically equivalent to C_2, we have $C_{2p} \longrightarrow C_n$ for any $n > 1$.
   However it is easy to check that odd cycles do not map into even cycles: as for any $q \geq 0$, $C_{2q+1} \not\rightarrow C_2$, it follows from the fact that even cycles are homomorphically equivalent that; for any $p,q \geq 0$, $C_{2q+1} \not\rightarrow C_{2p}$.

7. Let $p \geq 0$. The odd cycle C_{2p+3} maps homomorphically into the odd cycle C_{2p+1}: simply map vertex $2p+1$ of C_{2p+3} to vertex 0 of C_{2p+1}; map vertex $2p+2$ of C_{2p+3} to vertex 1 of C_{2p+1}; and, map any other vertex i of C_{2p+3} to vertex i of C_{2p+1}. Since the composition of two homomorphism is again

a homomorphism, we have proved the following. let $n > m > 1$. If n and m are both odd then $C_m \longrightarrow C_n$.

8. However, it can be easily checked that, if $n > m > 1$ are such that n and m are both odd then $C_m \not\longrightarrow C_n$.

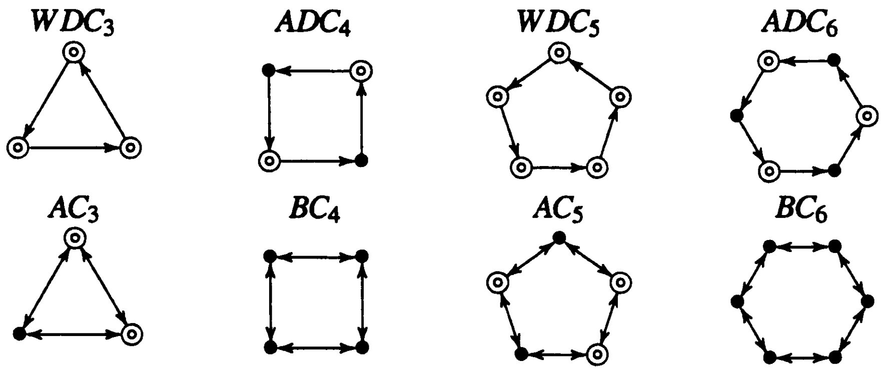

Figure 4.2: some coloured structures

Some 2-coloured σ_2-structures. Let $\mathbf{2} := \{0,1\}$. In our picture, we shall colour an element in white for the colour 0 and in black for the colour 1. Consider the following colourings,

$$w_n^2 : \begin{array}{rcl} |DC_n| & \to & \mathbf{2} \\ x & \mapsto & 0 \end{array} \qquad b_n^2 : \begin{array}{rcl} |DC_n| & \to & \mathbf{2} \\ x & \mapsto & 1 \end{array}$$

$$a_n^2 : \begin{array}{rcl} |DC_n| & \to & \mathbf{2} \\ x & \mapsto & \begin{cases} 0 & \text{if } x \text{ is even} \\ 1 & \text{otherwise.} \end{cases} \end{array}$$

Let $WDC_n := (DC_n, w_n^2)$ and $BDC_n := (DC_n, b_n^2)$ for $n \geq 1$, and for $n > 1$ set $ADC_n := (DC_n, a_n^2)$ (WDC stands for White Directed Cycle, BDC for Black Directed Cycle and ADC for Alternated Directed Cycle). Examples among such structures are depicted in Figure 4.2.

Define similarly, $WC_n := (C_n, w_n^2)$, $BC_n := (C_n, b_n^2)$ and $AC_n := (C_n, a_n^2)$.

1. Let $n,m > 1$ be such that m divides n and $p,q > 1$ such that q divides p. It

is easy to check the following,

$$WDC_n \overset{h_{n,m}}{\longrightarrow} WDC_m \qquad WC_n \overset{h_{n,m}}{\longrightarrow} WC_m$$
$$BDC_n \overset{h_{n,m}}{\longrightarrow} BDC_m \qquad BC_n \overset{h_{n,m}}{\longrightarrow} BC_m$$
$$ADC_{2p} \overset{h_{2p,2q}}{\longrightarrow} ADC_{2q} \qquad AC_{2p} \overset{h_{2p,2q}}{\longrightarrow} AC_{2q}$$

2. However, for any $p,q > 1$, $ADC_{2p+1} \not\longrightarrow ADC_{2q}$ since there is the edge $(n-1,0)$ where both $n-1$ and 0 are coloured white, whereas no such coloured edge occurs in ADC_{2q}. Since no edge of ADC_{2q} can be mapped over the white-white edge of ADC_{2p+1}, if $ADC_{2q} \longrightarrow ADC_{2p+1}$ then this would imply that ADC_{2q} can be mapped homomorphically into a directed path, which is not the case: hence $ADC_{2q} \not\longrightarrow ADC_{2p+1}$.

3. Similarly for any $p,q > 1$, we have $AC_{2p+1} \not\longrightarrow AC_{2q}$.

4. However, $AC_{2p} \longrightarrow AC_{2q+1}$. Indeed, $AC_{2p} \longrightarrow AC_2$.

5. Moreover, clearly there is no homomorphism between any choice of coloured structures of different type among the three types of colouring introduced, white, black or alternated, since the colourings are always incompatible (except for $WDC_2 \approx WC_2 \hookrightarrow AC_{2p+1}$ for $p > 1$).

6. The following gives the relation between the coloured cycles and the directed coloured cycles.

$$WDC_n \overset{h_{n,m}}{\longrightarrow} WC_m$$
$$BDC_n \overset{h_{n,m}}{\longrightarrow} BC_m$$
$$ADC_n \overset{h_{n,m}}{\longrightarrow} AC_m$$

4.1.3 Representations, recolourings and FP

Next, we shall introduce the notion of a representation for a forbidden patterns problem (that shall be defined shortly afterwards). First we shall discuss in some detail the intuition behind these forthcoming definitions. In the case of homomorphism problems, a problem is represented by its template, and for two templates A and B, we have $CSP(A) \subseteq CSP(B)$ if $A \longrightarrow B$ (notice that the converse also holds, *cf.* Proposition 2.1). The notion of a recolouring of one representation to another shall have a similar behaviour:

1. a recolouring shall define a notion of morphism from one representation to another; and

2. if there exists a recolouring of one representation to another then the forbidden patterns problem defined by the first representation is contained in that defined by the second.

A finite σ-*representation* with colours μ is a pair $(\mu, \mathcal{M})$, where μ is a finite set and $\mathcal{M}$ is a finite set of μ-coloured σ-structures. We call the elements of $\mathcal{M}$ the *forbidden patterns* of $(\mu, \mathcal{M})$. Let $REP(\sigma)$ denote the class of finite σ-representations.

EXAMPLE. Let $n \geq 1$ and $p \geq 1$. Consider the following σ_2-representations:

$$\mathfrak{MDC}_n^2 := \{2, \{WDC_n, BDC_n\}\}$$

$$\mathfrak{MC}_n^2 := \{2, \{WC_n, BC_n\}\}$$

$$\mathfrak{ADC}_{2p}^2 - \mathfrak{ME} := \{2, \{ADC_{2p}, WDE, BDE\}\}$$

where WDE, respectively BDE, denotes a single directed edge whose vertices are coloured in white, respectively black (the names of these representations standing for Monochromatic Directed Cycles, Monochromatic Cycles, and, Alternated Directed and Monochromatic Edges, respectively). See Figure 4.3 for some examples. In this picture, each cell in an array stands for a single forbidden pattern (as a forbidden pattern is not necessarily connected), except for the top cell which represents the set of colours.

▲

A coloured structure (A, c_μ^A) in $\text{STRUC}_\mu(\sigma)$ is said to be *valid* with respect to $(\mu, \mathcal{M})$, if, and only if, none of the forbidden patterns maps into (A, c_μ^A) via a colour preserving homomorphism. In other words, for any (M, c_μ^M) in $\mathcal{M}$ and for any mapping h of $|M|$ to $|A|$, either $M \overset{h}{\not\longrightarrow} A$ or $c_\mu^M \circ h \neq c_\mu^A$. When (A, c^A) is not valid with respect to some (M, c_μ^M) in $\mathcal{M}$ via some colour preserving homomor-

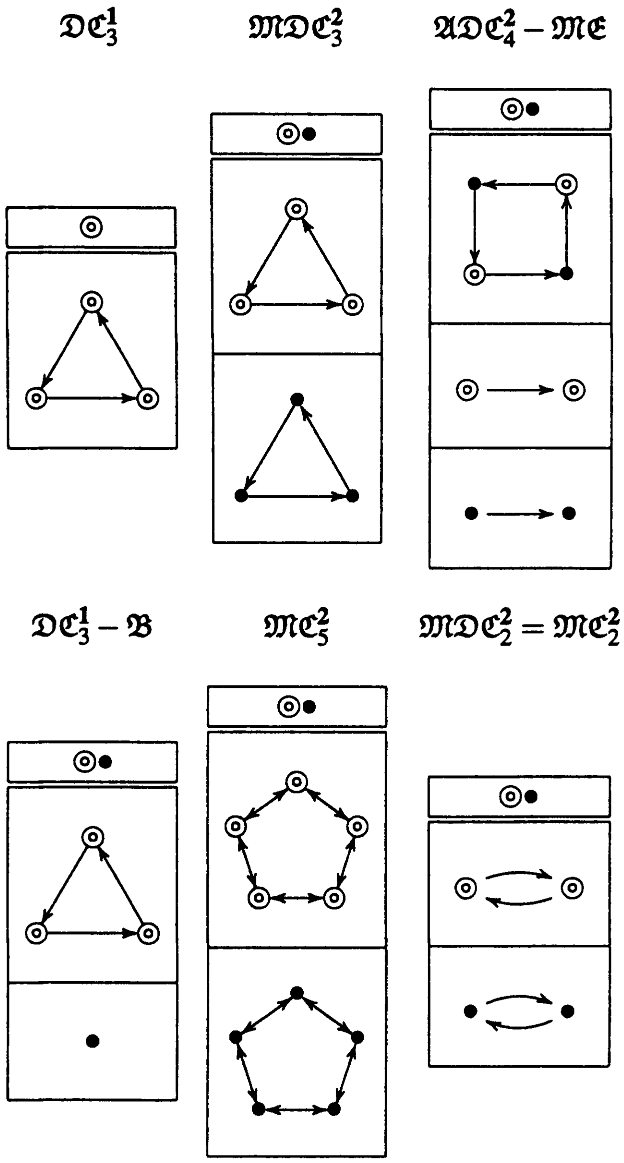

Figure 4.3: some representations for directed graphs

phism h, we shall use the following diagram.

$$\begin{tikzcd} M \arrow[r, "h"] \arrow[d, "\mathcal{M}"', "c_\mu^M"'] & A \arrow[dl, "c_\mu^A"] \\ \mu & \end{tikzcd}$$

EXAMPLE. Consider the representation $\mathfrak{MDC}_3^2$ (see Figure 4.3) and the σ_2-structure DC_5. Consider now this structure together with a colouring that maps every vertex to the colour "white"; that is, the coloured structure WDC_5. WDC_5 is valid with respect to $\mathfrak{MDC}_3^2$ as the forbidden patterns do not map into WDC_5 via a colour preserving homomorphism. ▲

Let $(\mu, \mathcal{M})$ be a σ-representation. Define the *forbidden patterns problem* with representation $(\mu, \mathcal{M})$, denoted $FP(\mu, \mathcal{M})$, to be the problem with yes-instances those σ-structures B such that:

- there exists a mapping c_μ^B such that (B, c_μ^B) is valid for $(\mu, \mathcal{M})$.

Denote by FP_σ the class of forbidden patterns problems given by a σ-representation and set:

$$FP := \bigcup_\sigma FP_\sigma.$$

We now define a notion that is ***absolutely essential*** in the remainder of this work, namely the notion of a ***recolouring between representations***. As we shall see later, the notion of a representation generalises the notion of a template, and the notion of a recolouring generalises the notion of a homomorphism. To grasp the idea behind the following definition, consider the contrapositive of the definition of a homomorphism as given in Section 2.1:

- for any r-ary symbol in σ and for any $\bar{b}$ in $|B|^r$, for any $\bar{a}$ in $|A|^r$ such that $h(\bar{a}) = \bar{b}$, if $R^B(\bar{b})$ does not hold then $R^A(\bar{a})$ does not hold.

That is, informally, the inverse image of a tuple not present in the target structure is not present in the source structure. As we shall see later, a tuple not present in the template of a homomorphism problem corresponds to a forbidden pattern of a special kind. Hence, the intuition behind our definition of a recolouring is that it induces inverse images of forbidden patterns and that the inverse image of something forbidden is forbidden.

- for all forbidden patterns (N, c^N_ν) in $\mathcal{N}$ and all functions c^N_μ of $|N|$ to μ with $c^N_\nu = r \circ c^N_\mu$, the coloured structure (N, c^N_μ) is not valid with respect to the representation $(\mu, \mathcal{M})$.

If r is a recolouring of $(\mu, \mathcal{M})$ to $(\nu, \mathcal{N})$ then we write $(\mu, \mathcal{M}) \xrightarrow{r} (\nu, \mathcal{N})$.

So for any $(N, r \circ c^N_\mu)$ in $\mathcal{N}$, there exists some (M, c^M_μ) in $\mathcal{M}$ with the property that $M \xrightarrow{m} N$ such that the following diagram commutes.

$$\begin{tikzcd}
M \arrow[r, "m"] \arrow[d, swap, "\mathcal{M}\ c^M_\mu"] & N \arrow[d, "\mathcal{N}\ c^N_\nu"] \arrow[dl, swap, "c^N_\mu"] \\
\mu \arrow[r, "r"] & \nu
\end{tikzcd}$$

EXAMPLE. Consider now the following mappings:

$$\begin{array}{rccc@{\qquad}rccc}
id_2: & 2 & \rightarrow & 2 & s_2: & 2 & \rightarrow & 2 \\
 & 0 & \mapsto & 0 & & 0 & \mapsto & 1 \\
 & 1 & \mapsto & 1 & & 1 & \mapsto & 0 \\
c0_2: & 2 & \rightarrow & 2 & c1_2: & 2 & \rightarrow & 2 \\
 & 0 & \mapsto & 0 & & 0 & \mapsto & 1 \\
 & 1 & \mapsto & 0 & & 1 & \mapsto & 1
\end{array}$$

In the following, let $n \geq m > 1$ such that m divides n.

1. We claim that id_2 is a recolouring of $\mathfrak{MDC}^2_n$ to $\mathfrak{MDC}^2_m$. Indeed, the only pre-image of WDC_m via id_2 is WDC_m; and, we have seen previously that $WDC_n \longrightarrow WDC_m$ if $n \geq m > 1$ and m divides n; thus, WDC_m is a valid colouring with respect to $\mathfrak{MDC}^2_n$ The case of the other forbidden pattern BDC_n is similar. Hence, we have shown that:

$$\text{if } n \geq m > 1 \text{ and } m \text{ divides } n \text{ then } \mathfrak{MDC}^2_n \xrightarrow{id_2} \mathfrak{MDC}^2_m.$$

2. By symmetry of the considered representation with respect to its colours, we have:

$$\text{if } n \geq m > 1 \text{ and } m \text{ divides } n \text{ then } \mathfrak{MDC}_n^2 \xrightarrow{s_2} \mathfrak{MDC}_m^2.$$

3. However, notice that id_2 is not a recolouring of $\mathfrak{MDC}_m^2$ to $\mathfrak{MDC}_n^2$, since

$$WDC_m \not\longrightarrow WDC_n \text{ and } BDC_m \not\longrightarrow WDC_n.$$

4. Similarly, s_2 is not a recolouring of $\mathfrak{MDC}_m^2$ to $\mathfrak{MDC}_n^2$.

5. For the two other mappings, $c0_2$ and $c1_2$, one can easily check that they are not recolourings, for example ADC_m is a pre-image of WDC_m via $c0_2$, respectively of BDC_m via $c1_2$, and we have seen that there is no homomorphisms between the alternated coloured cycles and the uniformly coloured cycles in the directed case. These maps are not recolourings of $\mathfrak{MDC}_m^2$ to $\mathfrak{MDC}_n^2$ either.

6. It follows therefore that there is no recolouring of $\mathfrak{MDC}_m^2$ to $\mathfrak{MDC}_n^2$, which we denote by:

$$\text{if } m,n > 1 \text{ and } m \text{ divides } n \text{ then } \mathfrak{MDC}_m^2 \not\longrightarrow \mathfrak{MDC}_n^2.$$

▲

The notion of recolouring we just defined satisfies the properties we required. Indeed, notice that the composition of two recolourings is a recolouring and that we have an identity recolouring associated with any representation $(\mu, \mathcal{M})$ induced by the identity map over μ, which we shall denote $id_{(\mu,\mathcal{M})}$. One can therefore speak of the *category of σ-representations*. This proves (1). As in the case of σ-structures, this category has further interesting properties that shall be investigated in Chapter 5. For (2) consider the following proposition.

Proposition 4.1 *Let* $(\mu, \mathcal{M})$ *and* $(\nu, \mathcal{N})$ *be two σ-representations. If there exists a recolouring* $(\mu, \mathcal{M}) \xrightarrow{r} (\nu, \mathcal{N})$ *then* $FP(\mu, \mathcal{M}) \subseteq FP(\nu, \mathcal{N})$.

PROOF. Let A be a σ-structure. Assume that A is a no-instance of $FP(\nu, \mathcal{N})$. Let

c^A_μ be a colouring of A. $(A, r \circ c^A_\mu)$ can not be valid for $(\nu, \mathcal{N})$. Hence, there exists some forbidden pattern (N, c^N_ν) in $\mathcal{N}$ and some colour preserving homomorphism n such that $(N, c^N_\nu) \xrightarrow{n} (A, r \circ c^A_\mu)$. Since r is a recolouring and $r \circ c^A_\mu \circ n = c^N_\nu$, it follows that there exists some forbidden pattern (M, c^M_μ) in $\mathcal{M}$ and some colour preserving homomorphism m such that $(M, c^M_\mu) \xrightarrow{m} (N, c^A_\mu \circ n)$. Finally, it follows that, $(M, c^M_\mu) \xrightarrow{n \circ m} (A, c^A_\mu)$ (to see this, note that $n \circ m$ is a homomorphism and that it respects colourings). Hence, (A, c^A_μ) is not valid for $(\mu, \mathcal{M})$. Thus, A is a no-instance of $FP(\mu, \mathcal{M})$. □

$$\begin{tikzcd}
M \arrow[r, "m"] \arrow[d, "\mathcal{M}\; c^M_\mu"'] \arrow[drr, bend left, "nom"] & N \arrow[d, "\mathcal{N}\; c^N_\nu"] \arrow[dl, "c^N_\mu"'] \arrow[dr, "n"] & \\
\mu \arrow[r, "r"] & \nu & A \arrow[l, "roc^A_\mu"'] \arrow[ll, bend left, "c^A_\mu"]
\end{tikzcd}$$

The converse does not hold in general; we shall provide a non-trivial counter-example at the end of Section 4.4 and some trivial counter-examples in the following.

Trivial representations. Notice that there are only two representations with colour set $\mu = \emptyset$. Indeed, there is only one structure (up to isomorphism) that can be $\emptyset$-coloured: it is the void structure that has no elements, which we shall denote 0_σ. It can be coloured by the mapping $c^{0_\sigma}_\emptyset$ (considering a mapping $c^{0_\sigma}_S$ of $\emptyset$ to some set S as a special binary relation, $\emptyset = c^{0_\sigma}_S \subseteq |0_\sigma| \times S = \emptyset$). Hence the only two representations with an empty set of colours are $o_\sigma := (\emptyset, \{(0_\sigma, c^{0_\sigma}_\emptyset)\})$ and $\tilde{o}_\sigma := (\emptyset, \emptyset)$. The former represents the trivial problem without any yes-instances and the latter represents the problem with a single yes-instance, namely the void structure 0_σ. However, there are some other representations that define the same problems to those defined by these two trivial representations:

- the representations with a non-void set of colours μ and with a set of forbidden patterns $\mathcal{M}$ consisting only of the coloured structure $(0_\sigma, c^{0_\sigma}_\mu)$; and

- the representations with a non-void set of colours μ and with a set of forbidden patterns $\tilde{\mathcal{M}}$ consisting of the coloured structures $\overline{K_1}^k$ with a single element coloured k, for any colour k in μ.

Clearly we have

$$FP(\mu, \mathcal{M}) = \emptyset = FP(o_\sigma)$$

and

$$FP(\mu, \tilde{\mathcal{M}}) = \{0_\sigma\} = FP(\tilde{o}_\sigma).$$

However, there can not be any mapping of μ to $\emptyset$ as μ is non-void. Hence, there is neither a recolouring of $(\mu, \mathcal{M})$ to o_σ nor of $(\mu, \tilde{\mathcal{M}})$ to $\tilde{o}_\sigma$. This provides some trivial counter-examples for the converse of the last proposition. The first problem is not in CSP as it has no yes-instances, and any CSP problem has at least one yes-instance, its template. Note that the second problem is nothing else than the problem $CSP(0_\sigma)$. Having dealt with these problems, we shall assume in the following that none of the representations we consider define problems equal to $FP(o_\sigma)$ or $FP(\tilde{o}_\sigma)$.

As we have seen earlier, with the notion of a recolouring we have a morphism of representations, thus we can consider the induced notion of monomorphism (respectively epimorphism): it corresponds to the recolourings induced by mappings that are injective (respectively surjective). We use a similar notation for recolourings as we did for homomorphisms and colour preserving homomorphisms. If r is an injective recolouring then we say that r is a *monorecolouring* and we write $(\mu, \mathcal{M}) \xhookrightarrow{r} (\nu, \mathcal{N})$. In this case, by analogy with the case of σ-structures, we say that $(\mu, \mathcal{M})$ is a *subrepresentation* of $(\nu, \mathcal{N})$. Let

$$\mathcal{M}' = \{(M, c_\mu^M) \in \mathrm{STRUC}_\mu(\sigma) | (M, r \circ c_\mu^M) \in \mathcal{N}\}$$

We call the representation $(\mu, \mathcal{M}')$ the *subrepresentation of* $(\nu, \mathcal{N})$ *induced by the recolouring* i (or *induced subrepresentation* of $(\nu, \mathcal{N})$ for short). If r is a surjective recolouring then we say that r is an *epirecolouring*, and we write $(\mu, \mathcal{M}) \xrightarrow{r} (\nu, \mathcal{N})$.

A recolouring that is bijective and whose inverse is a recolouring is called an *isorecolouring*. If $(\mu, \mathcal{M}) \xrightarrow{r} (\nu, \mathcal{N})$ and r is an isorecolouring then we write

$$(\mu, \mathcal{M}) \approx (\nu, \mathcal{N}).$$

Let $(\mu, \mathcal{M})$ be a representation. We say that $(\mu, \mathcal{M})$ is *simple* if, either $|\mathcal{M}| \leq 1$ or for any pair of distinct forbidden patterns (M, c_μ^M) and $(M', c_\mu^{M'})$ in $\mathcal{M}$, we have $(M, c_\mu^M) \not\rightarrow (M', c_\mu^{M'})$.

EXAMPLE. The $\mathfrak{MDC}^2_m$'s from the previous example are easily seen to be simple representations. ▲

As the following result shows, for every representation, there exists a simple representation that is equivalent up to isorecolouring.

Lemma 4.2 *Let $(\mu, \mathcal{M})$ be a representation. There exists a simple representation $(\nu, \mathcal{N})$ such that:*

$$(\mu, \mathcal{M}) \approx (\nu, \mathcal{N}).$$

PROOF. Suppose that $(\mu, \mathcal{M})$ is not simple. Set $\nu := \mu$ and construct $\mathcal{N}$ from $\mathcal{M}$ as follows. Start with $\mathcal{N} = \mathcal{M}$ and as long as there exists a pair of distinct forbidden patterns (M_0, c^{M_0}) and (M_1, c^{M_1}) in $\mathcal{N}$ such that

$$(M_1, c^{M_1}) \twoheadrightarrow (M_0, c^{M_0})$$

remove (M_0, c^{M_0}) from $\mathcal{N}$. This construction terminates eventually as $\mathcal{N}$ is finite and clearly $(\nu, \mathcal{N})$ is simple. The mapping $r : \mu \to \nu$ induced by id_μ (recall that $\nu = \mu$) is a recolouring: for every forbidden pattern (N, c^N) in $\mathcal{N}$, its inverse image via r is (N, c^N) itself and is present in $\mathcal{M}$ by construction of $\mathcal{N}$. The inverse of r is clearly a recolouring as for any forbidden pattern (M_0, c^{M_0}) in $\mathcal{M}$ that is no longer present in $\mathcal{N}$, there exists some (M_1, c^{M_1}) in $\mathcal{M}$ such that

$$(M_1, c^{M_1}) \twoheadrightarrow (M_0, c^{M_0}).$$

If (M_1, c^{M_1}) is not present in $\mathcal{N}$ either then, by construction of $(\nu, \mathcal{N})$, there exists some $n > 1$ and forbidden patterns (M_i, c^{M_i}) in $\mathcal{M}$ $(1 < i \leq n)$ such that

$$(M_n, c^{M_n}) \twoheadrightarrow (M_{n-1}, c^{M_{n-1}}) \twoheadrightarrow \dots \twoheadrightarrow (M_1, c^{M_1}) \twoheadrightarrow (M_0, c^{M_0}),$$

and such that (M_n, c^{M_n}) is in $\mathcal{N}$. Since the composition of colour preserving homomorphisms is a colour preserving homomorphism, it follows that for any (M, c^M) in $\mathcal{M}$, its inverse image induced by the mapping id_μ, that is (M, c^M) itself, is such that there exists some (N, c^N) in $\mathcal{N}$ such that

$$(N, c^N) \longrightarrow (M, c^M).$$

In other words, r^{-1} is a recolouring of $(\nu, \mathcal{N})$ to $(\mu, \mathcal{M})$. Thus we have proved that r is an isorecolouring, hence we have

$$(\mu, \mathcal{M}) \approx (\nu, \mathcal{N}).$$

□

In fact, by analogy to σ-structures, a representation that is not simple would correspond to a structure in which we would list more than once a tuple in some relation.

The previous result together with Proposition 4.1 leads to the following.

Corollary 4.3 *Every forbidden patterns problem can be given by a representation that is simple.*

EXAMPLE. With reference to earlier examples, via similar reasoning to that developed in the case of representations involving directed cycles, we obtain the following for the case of cycles:

if $m, n > 1$ and m divides n *then*

$$\mathfrak{MC}_n^2 \xrightarrow{id_2} \mathfrak{MC}_m^2 \quad \mathfrak{MC}_n^2 \xrightarrow{s_2} \mathfrak{MC}_m^2$$

Moreover, for $p \geq 1$, we have $WC_{2p} \stackrel{h_{2p,2}}{\longrightarrow} WC_2$ and $WC_2 \xrightarrow{f_2} WC_{2p}$ and similarly for the BC. Hence, the following holds:

$$\mathfrak{MC}_{2p}^2 \xrightarrow{id_2} \mathfrak{MC}_2^2 \text{ and } \mathfrak{MC}_{2p}^2 \xrightarrow{id_2^{-1}} \mathfrak{MC}_2^2$$

So, id_2 is an isorecolouring between $\mathfrak{MC}_{2p}^2$ and $\mathfrak{MC}_2^2$, and we write:

$$\mathfrak{MC}_{2p}^2 \approx \mathfrak{MC}_2^2.$$

These previous results might seem a bit odd to the reader who is used to the corresponding notion of isomorphism for σ-structures; in fact, note that a forbidden patterns problem can be given by numerous simple representations, that are equivalent via isorecolourings by replacing any forbidden pattern by another that is homomorphically equivalent to it as in the previous example.

We leave the following as an exercise for the reader:

$$\text{if } p,q \geq 1 \text{ and } q \text{ divides } p \textit{ then}$$
$$\mathfrak{AD}\mathfrak{C}^2_{2p} - \mathfrak{ME} \xrightarrow{id_2} \mathfrak{AD}\mathfrak{C}^2_{2q} - \mathfrak{ME} \text{ and } \mathfrak{AD}\mathfrak{C}^2_{2p} - \mathfrak{ME} \xrightarrow{s_2} \mathfrak{AD}\mathfrak{C}^2_{2q} - \mathfrak{ME}$$

Consider as further examples,

$$\mathfrak{WD}\mathfrak{C}^2_3 - \mathfrak{B} := \{2, \{WDC_3, B\}\}$$

$$\mathfrak{D}\mathfrak{C}^1_3 := \{1, \{WDC_3\}\}$$

where B is the structure consisting of a single element coloured black and $1 = \{0\}$ (and we shall consider 0 to be white as before). These representations are depicted in Figure 4.3. It is easy to check that,

$$\mathfrak{WD}\mathfrak{C}^2_3 - \mathfrak{B} \xrightarrow{id_2} \mathfrak{MD}\mathfrak{C}^2_3$$
$$\mathfrak{WD}\mathfrak{C}^2_3 - \mathfrak{B} \xrightarrow{cO_{2,1}} \mathfrak{D}\mathfrak{C}^1_3 \text{ and } \mathfrak{D}\mathfrak{C}^1_3 \xrightarrow{cO_{1,2}} \mathfrak{WD}\mathfrak{C}^2_3 - \mathfrak{B}$$

where,

$$cO_{2,1}: \begin{array}{ccc} 2 & \to & 1 \\ 0 & \mapsto & 0 \end{array} \quad \text{and} \quad cO_{1,2}: \begin{array}{ccc} 1 & \to & 2 \\ 0 & \mapsto & 0 \end{array}$$

However, $\mathfrak{D}\mathfrak{C}^1_3 \not\approx \mathfrak{WD}\mathfrak{C}^2_3 - \mathfrak{B}$. ▲

Furthermore, we shall make use of the notion of an *image of a representation via a recolouring*. Let $(\mu, \mathcal{M})$ and $(\nu, \mathcal{N})$ be two σ-representations and $(\mu, \mathcal{M}) \xrightarrow{r} (\nu, \mathcal{N})$. Define $r(\mu, \mathcal{M}) := (r(\mu), \{(N, c^N_\nu) | c^N_\nu(|N|) \subseteq r(\mu)\})$, where $r(\mu)$ denotes the image of the set of colours μ via the mapping r.

4.1.4 MMSNP captures exactly FP

We have already seen in Section 3.1.2 how to associate a structure to a negated conjunct of a sentence of MMSNP. The κ-colours of a given sentence of MMSNP correspond to the set of colours of a representation whose forbidden patterns are simply the structures induced by the negated conjuncts (throughout we use the notation established in the previous chapter, *e.g.* κ, σ'...). The following lemma shows that the obtained representation characterises the forbidden patterns problem captured by the given sentence of MMSNP.

Lemma 4.4 *Let Φ be a sentence of the logic MMSNP. There exists a representation $(\mu_\Phi, \mathcal{M}_\Phi)$ such that $FP(\mu_\Phi, \mathcal{M}_\Phi)$ is expressed by Φ.*

PROOF. Let Φ be a sentence of MMSNP. By Lemma 3.2 we can assume w.l.o.g. that it has full colourings. For uniformity, let us fix things so that there is at least one monadic predicate. One way to achieve this is as follows. If Φ is a first-order sentence then simply add an existential monadic predicate M, replace any negated conjunct $\gamma = \neg(\alpha)$ by $\neg(\alpha \wedge \beta)$, where $\beta := \bigwedge_{x \in X_\gamma} M(x)$, and add the negated conjunct $\neg(\neg M(x))$, for some particular bound variable x: the new sentence is clearly equivalent to Φ and has full colourings. Hence assume w.l.o.g. that Φ has full colourings and is not a first-order sentence.

Consider $(\mu_\Phi, \mathcal{M}_\Phi)$ to be the representation defined as follows: set μ_Φ to be the set of κ-colours, where κ is the signature containing the existential monadic predicates of Φ (it can not be void as we ensured beforehand that the sentence is not first order); and set $\mathcal{M}_\Phi$ to be the set of μ_Φ-coloured σ-structures $(G, c^G)_\gamma$ induced by each negated conjunct $\gamma = \neg(\alpha \wedge \beta)$ in $\gamma(\Phi)$ as follows:

- G is the σ-structure induced by α (denoted by G_α in Subsection 3.1.2, recall that it has domain X_γ); and
- for any x in $|G|$, set $c^G(x)$ to be the κ-colour given by β to x.

We claim that $FP(\mu_\Phi, \mathcal{M}_\Phi)$ is expressed by Φ.

Let A be a σ-structure. $A \models \Phi$ if, and only if, there exists an extension A' of A to σ' such that for each $\gamma \in \gamma(\Phi)$ and for any $\pi : X_\gamma \to |A|$,

$$A' \models \gamma(\bar{x}/\pi(\bar{x})).$$

The latter holds if, and only if,

$$A' \not\models \alpha(\bar{x}/\pi(\bar{x})) \text{ or } A' \not\models \beta(\bar{x}/\pi(\bar{x})).$$

That is, in the first case that there exists some r-ary symbol R in σ such that $R(\bar{x})$ occurs in α and $R(\bar{x}/\pi(\bar{x}))$ does not hold in A; in other words, according to the definition of $(G, c^G)_\gamma$, that π is not a homomorphism of G to A. In the second case, there exists some monadic symbol M in κ and some variable x in X_β such that $M(x)$ occurs in β and $M(x)$ does not hold in A', or $\neg M(x)$ occurs in β and $M(x)$ holds in A'. Let $c^A_{\mu_\Phi}$ be the mapping induced by A': it maps each element of A to its κ-colour in the extension A'. Then the second case is equivalent to $c^A_{\mu_\Phi} \circ \pi \neq c^G$. The two cases together are equivalent to π not being a colour preserving homomorphism of $(G, c^G)_\gamma$ to $(A, c^A_{\mu_\Phi})$. Hence we have proved that there exists some valid colouring for A; in other words, that A is a yes-instance of $FP(\mu_\Phi, \mathcal{M}_\Phi)$. For any colouring $c^A_{\mu_\Phi}$, one can derive an extension of A by setting the monadic predicates from κ according to the κ-colours of the elements of A given by $c^A_{\mu_\Phi}$; thus, clearly the converse also holds. □

The following lemma deals with the converse translation; that is, converting a representation into a sentence of MMSNP. One can label each element of μ with an integer written in binary, each such integer inducing a κ-colour, where κ contains one monadic predicate for each place (simply consider the binary expansions to be padded with zeros to the left and for each place set the corresponding monadic predicate in κ negatively for a zero and positively for a one). Hence, each forbidden pattern induces a negated conjunct.

Lemma 4.5 *Let $(\mu, \mathcal{M})$ be a non-trivial representation. There exists a sentence $\Phi_{(\mu,\mathcal{M})}$ of MMSNP such that $FP(\mu, \mathcal{M})$ is expressed by $\Phi_{(\mu,\mathcal{M})}$.*

PROOF. We can assume w.l.o.g. that there exists some $n > 0$ such that $|\mu| = 2^n$. Indeed, if it were not the case, add new colours to μ to reach the nearest power of 2 then add to $\mathcal{M}$ the forbidden patterns consisting of a single vertex coloured by one of the new colours. Clearly this new representation defines an equivalent problem. Let $\kappa := (M_1, M_2, \ldots, M_n)$ be a signature consisting of monadic symbols that do not occur in σ. There are 2^n κ-colours, thus we can identify each element of μ with a κ-colour. Consider $\Phi_{(\mu,\mathcal{M})}$ to be the sentence of MMSNP with: existential monadic predicates, the elements of κ; with universal first-order variables, the union of the universes of the σ-structures of the forbidden patterns in $\mathcal{M}$; and to have a negated conjunct $\gamma_{(G,c^G_\mu)}$, for each forbidden pattern (G, c^G_μ) in $\mathcal{M}$, constructed as follows:

- its α-part contains the atom $R(t)$ whenever $R(t)$ holds in G; and

- its β-part is the conjunction, for each element x of $|G|$, of the κ-colour given by $c_{\mu}^{G}(x)$.

If one applies the constructions used in the previous lemma to derive a representation from this sentence, one obtains a representation that is clearly equivalent to $(\mu, \mathcal{M})$. Thus it follows that $FP(\mu, \mathcal{M})$ is expressed by $\Phi_{(\mu,\mathcal{M})}$. □

Notice that in Lemma 4.5, we have not considered the case of trivial representations. The case of the trivial representations equivalent to $\bar{o}_{\sigma}$ is clear, as we can proceed as in the above proof. The case of the representations equivalent to o_{σ} is different. It does not really correspond to any sentence of MMSNP, as the standard semantics for logics ensures that 0_{σ} is always a yes-instance, unless we extend MMSNP by adding the "sentence" '**False**'. With this convention, from the two previous lemmas one can derive the following.

Corollary 4.6 *MMSNP captures exactly FP.*

EXAMPLE. The problem $FP(\mathfrak{MDC}_3^2)$ is expressed by the following sentence of MMSNP:

$$\exists C \forall x \forall y \forall z \neg(\ell_2(x,y,z) \wedge w(x,y,z)) \wedge \neg(\ell_2(x,y,z) \wedge b(x,y,z)).$$

Recall the abbreviation introduced in Section 2.4.1:

$$\ell_2(x,y,z) = \neg(E(x,y) \wedge E(z,x) \wedge E(y,z))$$

$$w(x,y,z) = C(x) \wedge C(y) \wedge C(z) \text{ and } b(x,y,z) := \neg C(x) \wedge \neg C(y) \wedge \neg C(z).$$

▲

4.2 Retracts

The notion of *retract* allows us to define the notion of a *core*, that is of a minimal retract. We recall this notion for the case of structures, extend it to coloured structures and develop a notion of core with respect to recolouring for representations. Together with the notion of *template of a representation* that shall be introduced in the next section, the notion of a *core of a representation* shall allow us to exhibit a structure that is a no-instance of a given forbidden paterns problem but that can be coloured nonetheless in a way that respects particular forbidden patterns: these structures shall be used later in Section 4.6 to build *witness families*.

4.2.1 Retracts, cores of finite structures and automorphic structures

A *retraction* of a structure A is a triplet (B,i,s), where B is a substructure of A via $B \overset{i}{\hookrightarrow} A$ such that $A \overset{s}{\twoheadrightarrow} B$ and $s \circ i = id_B$; that is, such that the following diagram commutes:

$$\begin{tikzcd} B \arrow[r, hook, "i"] \arrow[dr, "id_B"'] & A \arrow[d, "s"] \\ & B \end{tikzcd}$$

In this case we say that B is a *retract* of A. A structure A is said to be *automorphic*[2] if it has no proper retracts, that is, every retract of A is isomorphic to A. An automorphic retract of A is called a *core* of A.

Proposition 4.7 *Every structure has a unique core (up to isomorphism).*

PROOF. We prove the existence first. Let A be a structure. We prove that A has a core by induction on $|A|$. The base case is clear: if $|A| = 0$ then A is clearly automorphic, hence it has a core, itself. Assume that any structure A with $|A| \leq n$ has a core. Let A be a structure such that $|A| = n+1$. If A is automorphic then we are done. Assume that this is not the case. So there exists a proper retract B of A. Hence $|B| \leq n$ and it follows from the induction hypothesis that B has a core. Since clearly a retract of B is a retract of A, it follows that a core of B is a core of A. Finally A has a core.

We now prove the uniqueness of the core of a structure up to isomorphism. Let A be a structure and B_1 and B_2 be cores of A. That is, there are $B_1 \overset{i_1}{\hookrightarrow} A$, $A \overset{s_1}{\twoheadrightarrow} B_1$ such that $s_1 \circ i_1 = id_{B_1}$ and $B_2 \overset{i_2}{\hookrightarrow} A$, $A \overset{s_2}{\twoheadrightarrow} B_2$ such that $s_2 \circ i_2 = id_{B_2}$. Consider the homomorphic image of B_1 via $s_1 \circ i_2 \circ s_2 \circ i_1$: it is clearly a retract of B_1. Since B_1 is automorphic, it follows that $s_1 \circ i_2$ is surjective and $s_2 \circ i_1$ is injective. One can consider as well the homomorphic image of B_2 via $s_2 \circ i_1 \circ s_1 \circ i_2$ and derive that

[2] We use the terminology proposed in [22].

$s_2 \circ i_1$ is surjective and $s_1 \circ i_2$ is injective. Hence we have proved that B_1 and B_2 are isomorphic, since $s_1 \circ i_2$ and $s_2 \circ i_1$ are isomorphisms. □

Let A be a σ-structure. Denote by $\mathbf{core}(A)$ some representative among the set of cores of A.

EXAMPLE. Any DC_n is automorphic. However, for cycles: for $p \geq 2$, C_{2p} is not automorphic and its core is C_2; and for $p \geq 1$, C_{2p+1} is automorphic. ▲

It follows from Proposition 2.1 that $CSP(\mathbf{core}(A)) = CSP(A)$. Hence in our study of homomorphism problems, we can restrict ourselves to problems whose templates are cores without loss of generality. Notice however that if one is interested in counting the number of homomorphisms, that is, in complexity classes like ♯P as in [7] then this is not necessarily the case; *i.e.* the problem $\sharp CSP(A)$ (the number of homomorphism of a given B to A) is not the same as $\sharp CSP(\mathbf{core}(A))$ in general. Furthermore, Hell and Nešetřil have shown in [24], that deciding whether a graph is a core or not is co-NP-complete.

4.2.2 Retracts, cores of coloured structures and automorphic coloured structures

A *retraction* of a coloured structure (A, c^A) is a triplet $((B, c^B), i, s)$, where (B, c^B) is a subcoloured structure of (A, c^A) via $(B, c^B) \overset{i}{\hookrightarrow} (A, c^A)$, $(A, c^A) \overset{s}{\twoheadrightarrow} (B, c^B)$ satisfying $s \circ i = id_{(B, c^B)}$. In this case, (B, c^B) is called a *retract* of (A, c^A). This property can be summarised by the following diagram.

$$\begin{tikzcd} B \arrow[dr, "id_B"'] & A \arrow[l, two heads, "s"'] \arrow[r, "c^A"] & \mu \\ & B \arrow[u, hook, "i"] \arrow[ur, "c^B"'] & \end{tikzcd}$$

A coloured structure (A, c^A) is said to be *automorphic* if it has no proper retracts. An automorphic retract of (A, c^A) is called a *core* of (A, c^A).

Proposition 4.8 *Every coloured structure has a unique core (up to isomorphism)*

PROOF. Similar to the proof of Proposition 4.7. □

Let $(A, c_\mu^A) \in \mathrm{STRUC}_\mu(\sigma)$. Denote by $\mathbf{core}(A, c_\mu^A)$ some representative among the set of cores of (A, c_μ^A).

EXAMPLE. Notice that if A is automorphic then (A, c^A) is automorphic for any c^A, however the converse is not true: consider for a counter-example the 2-coloured structure consisting of two elements, one coloured black, the other white, connected via an edge to some white element, depicted as follows,

$$\bullet \longrightarrow \circledcirc \longleftarrow \circledcirc$$

As a coloured structure it is automorphic, however, if one consider this structure without its colouring, that is as follows,

$$\circ \longrightarrow \circ \longleftarrow \circ$$

then it is not a core.
Let $n > 0$: WDC_n, BDC_n and ADC_n are automorphic.
Let $p \geq 0$: WC_{2p+1}, BC_{2p+1} and AC_{2p+1} are automorphic.
Let $p \geq 1$: WC_{2p}, BC_{2p} and AC_{2p} are not automorphic and have for respective cores, WC_2, BC_2 and AC_2. ▲

Lemma 4.9 *Let $(\mu, \mathcal{M})$ be a representation. There exists a representation $(\nu, \mathcal{N})$ such that every forbidden pattern (N, c^N) in $\mathcal{N}$ is a coloured core and*

$$(\mu, \mathcal{M}) \approx (\nu, \mathcal{N}).$$

PROOF. Set $\nu := \mu$ and $\mathcal{N} := \{\mathbf{core}(M, c^M) \text{ such that } (M, c^M) \in \mathcal{M}\}$. It follows directly from the definition of a coloured core that the mapping r defined as in the proof of Lemma 4.2 is an isorecolouring. □

The following follows from Lemma 4.2 and from the previous lemma.

Corollary 4.10 *Every forbidden patterns problem can be given by a simple representation $(\mu, \mathcal{M})$ such that every forbidden pattern (M, c^M) in $\mathcal{M}$ is a coloured core.*

In the light of this corollary, from now on, unless otherwise stated we shall only ever consider simple representations such that each forbidden pattern is a coloured core. We now show by way of a back and forth argument that *if two such representations are equivalent up to isorecolouring then they are just the same up to a renaming of the colours.* Hence one obtains a notion of isorecolouring nearer to the intuitive one derived from the notion of isomorphism in the case of such representations.

$$\begin{tikzcd}
 & & M' \arrow[dll, "h'"'] \arrow[dl, "h\circ h'", "\mathcal{M}"'] \arrow[ddl, "c_\mu^{M'}"] \\
N \arrow[r, "h"] \arrow[d, "\mathcal{N}", "c_\nu^N"'] & M \arrow[d, "\mathcal{M}"', "c_\mu^M"] \arrow[dl, "c_\nu^M"'] & \\
\nu \arrow[r, "r"] & \mu &
\end{tikzcd}$$

Let $(\mu, \mathcal{M})$ and $(\nu, \mathcal{N})$ be two simple representations whose forbidden patterns are coloured cores. Let r be some isorecolouring of $(\nu, \mathcal{N})$ to $(\mu, \mathcal{M})$. Let $(M, c_\mu^M) \in \mathcal{M}$. Since r is a recolouring, the inverse image of (M, c_μ^M) via r, that is (M, c_ν^M), where $c_\nu^M = r^{-1} \circ c_\mu^M$, is not valid for $(\nu, \mathcal{N})$. Hence there exists some forbidden pattern (N, c_ν^N) in $\mathcal{N}$ and some colour preserving homomorphism $(N, c_\nu^N) \xrightarrow{h} (M, c_\nu^M)$. Now r^{-1} is a recolouring, thus there exists some forbidden pattern $(M', c_\mu^{M'})$ in $\mathcal{M}$ and some colour preserving homomorphism $(M', c_\mu^{M'}) \xrightarrow{h'} (N, c_\mu^N)$, where $c_\mu^N = (r^{-1})^{-1} \circ c_\nu^N = r \circ c_\nu^N$. Hence by composition with r it follows that $(N, c_\mu^N) \xrightarrow{h} (M, c_\mu^M)$. The composition of h and h' leads therefore to $(M', c_\mu^{M'}) \xrightarrow{h\circ h'} (M, c_\mu^M)$. Now, since $(\mu, \mathcal{M})$ is simple $(M', c_\mu^{M'})$ and (M, c_μ^M) must be the same forbidden pattern. Finally, (M, c_μ^M) and (N, c_μ^N) are homomorphically equivalent via h and h'. Since they are cores by assumption, they must be the same forbidden pattern. This proves our claim that simple representations whose forbidden patterns are coloured cores that are equivalent up to isorecolouring are simply obtained from each other via a permutation of the colours.

4.2.3 Retracts, cores of representations and automorphic representations

From now on and unless otherwise stated we only ever consider simple representations whose forbidden patterns are coloured cores. A *retraction* of a representation $(\mu,\mathcal{M})$ is a triplet $((\nu,\mathcal{N}),i,s)$, where $(\nu,\mathcal{N})$ is a subrepresentation of $(\mu,\mathcal{M})$ via the monorecolouring $(\nu,\mathcal{N}) \overset{i}{\hookrightarrow} (\mu,\mathcal{M})$ and s is an epirecolouring $(\mu,\mathcal{M}) \overset{s}{\twoheadrightarrow} (\nu,\mathcal{N})$ such that $s \circ i = id_{(\nu,\mathcal{N})}$. In this case we say that $(\nu,\mathcal{N})$ is a *retract* of $(\mu,\mathcal{M})$. A representation $(\mu,\mathcal{M})$ is said to be *automorphic* if it has no proper retracts, that is, every retract $(\nu,\mathcal{N})$ of $(\mu,\mathcal{M})$ is such that $(\mu,\mathcal{M}) \approx (\nu,\mathcal{N})$. An automorphic retract of $(\mu,\mathcal{M})$ is called a *core* of $(\mu,\mathcal{M})$.

EXAMPLE. Recall that $\mathfrak{WDC}_3^2 - \mathfrak{B} \overset{c0_{2,1}}{\twoheadrightarrow} \mathfrak{DC}_3^1$, and that $\mathfrak{DC}_3^1 \overset{c0_{1,2}}{\twoheadrightarrow} \mathfrak{WDC}_3^2 - \mathfrak{B}$ but that $\mathfrak{DC}_3^1 \not\approx \mathfrak{WDC}_3^2 - \mathfrak{B}$. Notice further that; $c0_{2,1}$ is an epirecolouring; and that $c0_{1,2}$ is a monorecolouring such that $c0_{2,1} \circ c0_{1,2} = id_{\mathfrak{DC}_3^1}$. In other words $(\mathfrak{DC}_3^1, c0_{1,2}, c0_{2,1})$ is a retraction of the representation $\mathfrak{WDC}_3^2 - \mathfrak{B}$. Furthermore the latter is not automorphic since it has a proper retract, namely $\mathfrak{DC}_3^1$. However $\mathfrak{DC}_3^1$ is automorphic since there can not be any recolouring of it to a trivial representation; indeed, there is no mapping of **1** to $\emptyset$. ▲

Proposition 4.11 *Every representation has a unique core (up to isorecolourings).*

PROOF. The proof is similar to the proof of Proposition 4.7. We prove that $(\mu,\mathcal{M})$ has a core by induction on $|\mu| = n$.

The base case is clear: if $(\mu,\mathcal{M})$ is a representation such that $|\mu| = 0$ then it can not have a proper retract.

Assume that any representation with n colours has a core. Let $(\mu,\mathcal{M})$ be a representation such that $|\mu| = n+1$. If $(\mu,\mathcal{M})$ is automorphic then it has a core: itself. Assume that $(\mu,\mathcal{M})$ is not automorphic. So it has a proper retract $(\nu,\mathcal{N})$. It follows that $|\nu| < n+1$, otherwise i being a bijection we would have $s = i^{-1}$ and i would be an isorecolouring contradicting the fact that $(\nu,\mathcal{N})$ is a proper retract. Since $\nu \leq n$, by the induction hypothesis, it follows that $(\nu,\mathcal{N})$ has a core. Hence by composition, $(\mu,\mathcal{M})$ has a core.

We now prove the uniqueness of the core of a representation up to isorecolouring. Let $(\mu,\mathcal{M})$ be a representation and $(\mu_1,\mathcal{M}_1)$ and $(\mu_2,\mathcal{M}_2)$ be cores of $(\mu,\mathcal{M})$. That is, there are $(\mu_1,\mathcal{M}_1) \overset{i_1}{\hookrightarrow} (\mu,\mathcal{M})$, $(\mu,\mathcal{M}) \overset{s_1}{\twoheadrightarrow} (\mu_1,\mathcal{M}_1)$ such that $s_1 \circ i_1 = id_{(\mu_1,\mathcal{M}_1)}$ and $(\mu_2,\mathcal{M}_2) \overset{i_2}{\hookrightarrow} (\mu,\mathcal{M})$, $(\mu,\mathcal{M}) \overset{s_2}{\twoheadrightarrow} (\mu_2,\mathcal{M}_2)$ such that $s_2 \circ i_2 = id_{(\mu_2,\mathcal{M}_2)}$. Consider the image of $(\mu_1,\mathcal{M}_1)$ via the recolouring $s_1 \circ i_2 \circ s_2 \circ i_1$: call this image $(\mu_1',\mathcal{M}_1')$. We now show that $(\mu_1',\mathcal{M}_1')$ is a retract of $(\mu_1,\mathcal{M}_1)$. Indeed,

$s'_1 := s_1 \circ i_2 \circ s_2 \circ i_1$ is an epirecolouring of $(\mu_1, \mathcal{M}_1)$ to $(\mu'_1, \mathcal{M}'_1)$ by definition of $(\mu'_1, \mathcal{M}'_1)$. Moreover set i'_1 to be simply $id_{(\mu_1, \mathcal{M}_1)}$ restricted to μ'_1. It is clearly a monorecolouring of $(\mu'_1, \mathcal{M}'_1)$ to $(\mu_1, \mathcal{M}_1)$. Thus $((\mu'_1, \mathcal{M}'_1), s'_1, i'_1)$ is a retract of $(\mu_1, \mathcal{M}_1)$. Since $(\mu_1, \mathcal{M}_1)$ is automorphic, it follows that $s_1 \circ i_2 \circ s_2 \circ i_1$ is an isorecolouring. Hence $s_1 \circ i_2$ is surjective and $s_2 \circ i_1$ is injective. One can consider as well the image of $(\mu_2, \mathcal{M}_2)$ via the recolouring $s_2 \circ i_1 \circ s_1 \circ i_2$ and derive that $s_2 \circ i_1$ is surjective and $s_1 \circ i_2$ is injective. Hence we have proved that $(\mu_1, \mathcal{M}_1) \approx (\mu_2, \mathcal{M}_2)$, since $s_1 \circ i_2$ and $s_2 \circ i_1$ are isorecolourings. □

Let $(\mu, \mathcal{M})$ be a σ-representation. Denote by $\mathbf{Core}(\mu, \mathcal{M})$ some representative among the set of cores of $(\mu, \mathcal{M})$ that have the properties of being:

- simple; and
- whose forbidden patterns are all coloured cores.

Note that the above is well-defined according to Lemma 4.2 and Lemma 4.9.

The following corollary follows from Proposition 4.1.

Corollary 4.12 *Let $(\mu, \mathcal{M})$ be a σ-representation. Then,*

$$FP(\mu, \mathcal{M}) = FP(\mathbf{Core}(\mu, \mathcal{M})).$$

EXAMPLE. We show that $\mathfrak{MDC}_n^2$ is an automorphic representation for any $n \geq 2$. Notice first that it is simple and each forbidden pattern is a connected core. In order to check whether it is automorphic, it is enough to check for proper retracts induced by retractions that are simple and whose forbidden patterns are coloured cores (by Corollary 4.10). Let $((\nu, \mathcal{N}), i, s)$ be such a retract. The map s must identify at least two colours, that is, w.l.o.g. $\nu = 1$ and i is the recolouring such that $i(0) = 0$. We claim that this implies that $WDC_n \in \mathcal{N}$. Indeed, since i is a recolouring and we assumed that $i(0) = 0$, we would have a forbidden pattern in $\mathcal{N}$ that would map into WDC_n (the inverse image of the forbidden pattern WDC_n via i). Moreover since $s \circ i = id_{(\nu, \mathcal{N})}$ and s is a recolouring, WDC_n would map to this forbidden pattern (the particular inverse image of this forbidden pattern via s coloured in white only). Hence, it follows from our assumption on the retract $(\nu, \mathcal{N})$ (simple and coloured cores only) that this forbidden pattern is nothing else than WDC_n. But then one inverse image of WDC_n via s would be ADC_n. But the latter is valid with respect to $\mathfrak{MDC}_n^2$. So there is only one case left to check which trivially can not hold; the case of the representation with a void colour set. There is simply no mapping to the void set from any set except the void set himself, so there can not be any epirecolouring of the considered representation to this trivial representation and we are done.

It can be easily checked that the representations $\mathfrak{MC}^2_{2p+1}$ and $\mathfrak{MC}^2_2$ are automorphic. The proof is similar to the previous one. The representations $\mathfrak{MC}^2_{2p}$, for any $p \geq 1$, are examples of representations that are automorphic too. They are all equivalent to $\mathfrak{MC}^2_2$, up to isorecolouring, as we said earlier.

The representations $\mathfrak{ADC}^2_{2p} - \mathfrak{MC}$, for any $p \geq 1$, are automorphic. Indeed, if there were some proper retract $((\nu, \mathcal{N}), i, s)$ then the only case to check is the case when ν contains exactly one colour. So assume w.l.o.g. that $\nu = 1$ and $i(0) = 0$. Then there must be some forbidden pattern (N, c^N) in $\mathcal{N}$ such that there exists some $(N, c^N) \xrightarrow{n} WE$. One possible inverse image of (N, c^N) via s being monochromatic and, say, white (since the only monochromatic white forbidden pattern is WE itself and since we assumed $(\nu, \mathcal{N})$ to be simple and to have only coloured cores as forbidden patterns) it follows that (N, c^N) is the coloured structure WE. One possible inverse image of WE is the structure consisting of a single edge whose origin is coloured white and target is coloured black. However, this coloured structure is clearly valid for $\mathfrak{ADC}^2_{2p} - \mathfrak{MC}$. Thus our claim follows. ▲

4.3 Templates

In this section, we shall introduce the notion of a *template for a representation*; it is a structure associated with some particular forbidden patterns of a given representation, the so-called *conform forbidden patterns*. It is constructed in the same way as the template of a problem that is captured by a sentence of MMSNP that is conform (cf. the proof of Lemma 3.5 in the previous chapter). Hence it is not surprising that a problem given by a conform representation is in CSP (in the light of Lemma 4.4). Furthermore, we shall see that *if the template of a given representation has a valid colouring then this representation has a conform retract*. This leads to an important result: the template of an automorphic representation that is not conform has no valid colouring. However, it can be coloured in such a way that the only forbidden patterns that witness that the colouring is not valid are not conform; in other words, one can colour the template such that it is valid if one considers each of its tuples separately.

This section is organised as follows. First, we shall define precisely the notion of a *conform forbidden pattern* and derive from results of the previous chapter that CSP is a strict subset of FP. Secondly, in Section 4.3.2, we shall define the notion of a *template of a representation* and we shall investigate the relation between the existence of recolourings between two given representations and the existence of homomorphisms between their templates (we hope to make clear to the reader in which sense we consider a recolouring to be a generalisation of a homomorphism). Finally, in Section 4.3.4, we prove the result mentioned above.

4.3.1 CSP is included in FP

A coloured structure (A, c^A) is said to be *antireflexive* whenever A is antireflexive. A coloured structure (A, c^A) is said to be *monotuple* whenever A is monotuple, and *non-sbavate* [3] whenever for each $a \in |A|$, there exists some r-ary relation symbol R in σ and some r-tuple $\bar{a}$ such that $R^A(\bar{a})$ holds and $a \in \{\bar{a}\}$. A representation $(\mu, \mathcal{M})$ is said to be *conform* if every forbidden pattern $(M, c^M_\mu) \in \mathcal{M}$ is monotuple, non-sbavate and antireflexive.

Let $(\mu, \mathcal{M})$ be a conform σ-representation. Then the sentence of MMSNP that expresses $FP(\mu, \mathcal{M})$ given by Lemma 4.5 is clearly conform. Thus by Lemma 3.5, it follows that $FP(\mu, \mathcal{M})$ is in CSP. However we can state more.

Proposition 4.13 *Let $(\mu, \mathcal{M})$ be a conform σ-representation. There exists a σ-structure T such that $CSP(T) = FP(\mu, \mathcal{M})$. Conversely, let T be a σ-structure. There exists a conform σ-representation $(\mu, \mathcal{M})$ such that $CSP(T) = FP(\mu, \mathcal{M})$. Moreover this is a one-to-one correspondence.*

PROOF. For the first part, one could use the argument given above, but we shall implement the construction directly, as this construction shall be used later. Let $(\mu, \mathcal{M})$ be a conform σ-representation. Construct T as follows.

- $|T| := \mu$; and
- for any r-ary relation symbol R in σ and any $\bar{t} = (t_1, t_2, \ldots, t_r) \in |T|^r$, set $R^T(\bar{t})$ to hold if, and only if, there is no forbidden pattern in $\mathcal{M}$ that is

[3] from the italian, literally that does not dribble; when a kid is colouring in outside the lines, italians say that the colours have dribbled.

equivalent to (M, c^M) up to colour preserving isomorphism, where M is the antireflexive, non-sbavate and monotuple structure defined as follows:

- $|M| := \{x_1, x_2, \ldots, x_r\}$; and
- $R^M(x_1, x_2, \ldots, x_r)$ is the only tuple to hold,

and is coloured as,

$$\begin{array}{rccl} c^M: & |M| & \to & \mu \\ & x_i & \mapsto & t_i (1 \leq i \leq r). \end{array}$$

Conversely, let T be a σ-structure. We derive a representation $(\mu, \mathcal{M})$ from T as follows:

- $\mu := |T|$; and
- for any symbol R in σ and any r-tuple $\bar{t} \in |T|^r$ such that $R^T(\bar{t})$ does not hold, add the following antireflexive, non-sbavate and monotuple coloured structure (M, c_μ^M) as a forbidden pattern:
  - $|M| := \{x_1, x_2, \ldots, x_r\}$; and
  - $R^M(x_1, x_2, \ldots, x_r)$ is the only tuple to hold,

  and is coloured as,

$$\begin{array}{rccl} c^M: & |M| & \to & \mu \\ & x_i & \mapsto & t_i (1 \leq i \leq r). \end{array}$$

This clearly establishes a one-to-one correspondence between σ-structures and conform σ-representations. We now prove that $FP(\mu, \mathcal{M}) = CSP(T)$. Let A be a

σ-structure.

$$
\begin{aligned}
& A \in FP(\mu, \mathcal{M}) \\
\iff & \exists c_\mu^A : |A| \to \mu, (A, c_\mu^A) \text{ is valid for } (\mu, \mathcal{M}) \\
\iff & \exists c_\mu^A : |A| \to \mu, \bigwedge_{(M, c_\mu^M) \in \mathcal{M}} (M, c_\mu^M) \not\to (A, c_\mu^A) \\
\iff & \exists c_\mu^A : |A| \to \mu, \bigwedge_{(M, c_\mu^M) \in \mathcal{M}} \forall m : |M| \to |A| (M \not\to A \vee c_\mu^A \circ m \neq c_\mu^M) \\
\iff & \exists c_\mu^A : |A| \to \mu, \\
& \bigwedge_{(M, c_\mu^M) \in \mathcal{M}} \forall m : |M| \to |A| (\neg R^A(m(x_1), m(x_2), \ldots, m(x_r)) \vee c_\mu^A \circ m \neq c_\mu^M) \\
\iff & \exists c_\mu^A : |A| \to \mu, \\
& \bigwedge_{(M, c_\mu^M) \in \mathcal{M}} \forall a_1, a_2, \ldots a_r \in |A| (\neg R^A(a_1, a_2, \ldots, a_r) \vee \exists 1 \leq i \leq r, c_\mu^A(a_i) \neq c_\mu^M(x_i)) \\
\iff & \exists c_\mu^A : |A| \to \mu, \bigwedge_{\neg R^T(\bar{t}), \bar{t} \in |T|^r} \forall \bar{a} \in |A|^r, c_\mu^A(\bar{a}) = \bar{t} \Rightarrow \neg R^A(\bar{a}) \\
\iff & \exists c_\mu^A : |A| \to \mu = |T|, \forall r\text{-ary } R \in \sigma, \forall \bar{a} \in |A|^r, \neg R^T(c_\mu^A(\bar{a})) \Rightarrow \neg R^A(\bar{a}) \\
\iff & \exists c_\mu^A : |A| \to \mu, A \xrightarrow{c_\mu^A} T \\
\iff & A \in CSP(T).
\end{aligned}
$$

□

We then derive the following:

Corollary 4.14 *$CSP \subsetneq FP$.*

PROOF. The inclusion comes from the previous lemma. It is strict since, for example, the problem NO-MONO-TRI was shown in Section 2.4.1 to be expressed by a sentence of MMSNP and not in CSP. Thus, by Lemma 4.4 this provides an example of a forbidden patterns problem that is not in CSP. □

4.3.2 Template of a representation

In fact, one can put aside in a given representation those forbidden patterns that are monotuple, non-sbavate and antireflexive, and for this subrepresentation construct a template. Thus, one can associate to any representation a template that shall somehow measure its conform part. Let $(\mu, \mathcal{M})$ be a σ-representation. Consider its subrepresentation $(\mu, \mathcal{D})$ that corresponds to its conform part, that is, where $\mathcal{D}$ is the following subset of $\mathcal{M}$:

$$\{(M, c_\mu^M) \text{ s.t. } (M, c_\mu^M) \in \mathcal{M} \text{ and is antireflexive, non-sbavate and monotuple}\}.$$

The subrepresentation $(\mu, \mathcal{D})$ of $(\mu, \mathcal{M})$ is conform, hence it follows from Proposition 4.13 that there exists some template T such that

$$CSP(T) = FP(\mu, \mathcal{D}) \supseteq FP(\mu, \mathcal{M}).$$

We call T the *template* of the representation $(\mu, \mathcal{M})$.

We claimed that our notion of recolouring generalises the notion of homomorphism; and the following proposition makes this more precise.

Proposition 4.15 *Let $(\mu, \mathcal{M})$ and $(\nu, \mathcal{N})$ be two σ-representations and let $T_{(\mu,\mathcal{M})}$ and $T_{(\nu,\mathcal{N})}$ be their respective templates. If $(\mu, \mathcal{M}) \xrightarrow{h} (\nu, \mathcal{N})$ and every forbidden pattern in $\mathcal{M}$ is non-sbavate then $T_{(\mu,\mathcal{M})} \xrightarrow{h} T_{(\nu,\mathcal{N})}$.*

When we consider a monotuple antireflexive structure N;

- with domain $\{x_1, x_2, \ldots, x_r\}$; and
- with the tuple $R(x_1, x_2, \ldots, x_r)$, where R is some r-ary symbol from σ.

We shall simply speak of the structure $R(x_1, x_2, \ldots, x_r)$.

PROOF. Let R be a r-ary symbol in σ and $\bar{t} \in \mu^r = |T_{(\mu,\mathcal{M})}|^r$. By construction of $T_{(\nu,\mathcal{N})}$, $R^{T_{(\nu,\mathcal{N})}}(h(\bar{t}))$ does not hold if, and only if, there exists some forbidden pattern that is isomorphic to $(R(x_1, x_2, \ldots, x_r), h \circ c_\mu^N)$ in $\mathcal{N}$, where,

$$\begin{array}{rccl} c_\mu^N: & |N| & \to & \mu \\ & x_i & \mapsto & t_i (1 \leq i \leq r). \end{array}$$

As h is a recolouring, $(R(x_1, x_2, \ldots, x_r), c_\mu^N)$ is not valid for $(\mu, \mathcal{M})$. Hence, there exists some $(M, c_\mu^M) \in \mathcal{M}$ and some f such that,

$$(M, c_\mu^M) \xrightarrow{f} (R(x_1, x_2, \ldots, x_r), c_\mu^N).$$

Since (M, c_μ^M) is non-sbavate and since $(R(x_1,x_2,\dots,x_r), c_\mu^N)$ is antireflexive, it follows that $(R(x_1,x_2,\dots,x_r), c_\mu^N)$ is a subcoloured-structure of (M, c_μ^M). In other words, (M, c_μ^M) is homomorphically equivalent to $(R(x_1,x_2,\dots,x_r), c_\mu^N)$. Moreover, since we consider only representations whose forbidden patterns are coloured cores (*cf.* Section 4.2.3) then we must have $(M, c_\mu^M) \approx (R(x_1,x_2,\dots,x_r), c_\mu^N)$. Finally, by definition of the template it follows that $R^{T_{(\mu,\mathcal{M})}}(\bar{t})$ does not hold. □

The notions of a homomorphism and of a recolouring clearly coincide in the case of conform representations and, furthermore, one can state a weaker form of the converse of the previous proposition; the converse itself being obviously false. Indeed, consider $\mathfrak{MDC}_2^2$ and $\mathfrak{MDC}_3^2$. These representations share the same template, the structure with domain 2 and all possible edges between the elements, as they do not have any antireflexive, non-sbavate and monotuple forbidden patterns. However, there is no recolouring of $\mathfrak{MDC}_2^2$ to $\mathfrak{MDC}_3^2$ (since 2 does not divide 3).

Proposition 4.16 *Let $(\mu, \mathcal{M})$ and $(\nu, \mathcal{N})$ be two σ-representations and let $T_{(\mu,\mathcal{M})}$ and $T_{(\nu,\mathcal{N})}$ be their respective templates. If $T_{(\mu,\mathcal{M})} \xrightarrow{h} T_{(\nu,\mathcal{N})}$ and $(\nu, \mathcal{N})$ is conform then $(\mu, \mathcal{M}) \xrightarrow{h} (\nu, \mathcal{N})$.*

PROOF. Let R be a r-ary relation symbol in σ and $\bar{t} \in \mu^r = |T_{(\mu,\mathcal{M})}|^r$.
Let $(R(x_1,x_2,\dots,x_r), h \circ c_\mu^N) \in \mathcal{N}$, where,

$$\begin{array}{rccl} c_\mu^N: & |N| & \to & \mu \\ & x_i & \mapsto & t_i (1 \le i \le r). \end{array}$$

By definition, $R^{T_{(\nu,\mathcal{N})}}(h(\bar{t}))$ does not hold, hence $R^{T_{(\mu,\mathcal{M})}}(\bar{t})$ does not hold since $T_{(\mu,\mathcal{M})} \xrightarrow{h} T_{(\nu,\mathcal{N})}$. Thus by construction of $T_{(\mu,\mathcal{M})}$, it follows that the forbidden pattern $(R(x_1,x_2,\dots,x_r), c_\mu^N)$ belongs to $\mathcal{M}$. Hence, we have $(\mu, \mathcal{M}) \xrightarrow{h} (\nu, \mathcal{N})$. □

Remark. Notice that we can relax a bit the hypothesis "$(\nu, \mathcal{N})$ is conform" to replace it by "any forbidden pattern of the form $(M, h \circ c_\mu^M)$ in $\mathcal{N}$ is conform".

In the following we prove that any non-conform forbidden pattern of a simple representation with template T is in fact a T-coloured structure. Hence, we can give equivalently a simple representation by giving its template together with its set of non-conform forbidden patterns.

Proposition 4.17 *Let $(\mu, \mathcal{M})$ be a representation with template T. If $(\mu, \mathcal{M})$ is simple then for any non-conform (M, c^M) in $\mathcal{M}$, we have $M \xrightarrow{c^M} T$.*

PROOF. Let $(\mu, \mathcal{M})$ be a simple representation. Let T be its template. Suppose that (M, c^M) in $\mathcal{M}$ is a non-conform forbidden pattern such that $M \stackrel{c^M}{\not\longrightarrow} T$. Let R be some r-ary relation symbol in σ and $\bar{y}$ be some r-tuple in M such that $R^M(\bar{y})$ holds but $R^T(c^M(\bar{y}))$ does not. By definition of the template of a representation, there is some conform forbidden pattern (D, c^D) in $\mathcal{M}$ that is isomorphic to $(R(x_1, \ldots, x_r), \bar{x} \mapsto c^M(\bar{y}))$ via some i. Hence we would have $(D, c^D) \xrightarrow{m \circ i^{-1}} (M, c^M)$, where m is defined by setting $m : \bar{x} \mapsto \bar{y}$ (this is well defined as (D, c^D) is conform and so it must be antireflexive). We obtain a contradiction as we assumed the representation $(\mu, \mathcal{M})$ to be simple. □

In the light of the previous proposition, we can give a simple representation equivalently as a σ-structure T together with a set $\mathcal{M}$ of T-coloured non-conform structures, that is, a set of non-conform coloured structures (M, c^M) such that $M \xrightarrow{c^M} T$. We denote by $(T, \mathcal{M})$ a representation in this new setting. The definition of validity of a coloured structure (A, c^A) becomes the following in this new setting. (A, c^A) is valid w.r.t. $(T, \mathcal{M})$ if, and only if, $A \xrightarrow{c^A} T$ and for any $M \xrightarrow{c^M} T$ in $\mathcal{M}$ and any $M \xrightarrow{m} A$, $c^A \circ m \neq c^M$. That is, A is not valid if $A \not\longrightarrow T$ or for any $c^A : A \longrightarrow T$, there exists some (M, c^M) in $\mathcal{M}$ and some $m : M \longrightarrow A$ such that the following diagram commutes.

$$\begin{tikzcd} M \arrow[r, "m"] \arrow[d, "c^M"', "\mathcal{M}"] & A \arrow[dl, "c^A"] \\ T & \end{tikzcd}$$

Notice that our new notation is compatible with the notation previously used, as one can consider that some μ-coloured structure (M, c^M_μ) is in fact a K_μ-coloured structure where K_μ denotes the *clique* with $|\mu|$-elements; that is the σ-structure with domain μ and such that for any r-ary relation symbol R in σ, $R^{K_\mu} = \mu^r$. Hence instead of $(\mu, \mathcal{M})$, read $(K_\mu, \mathcal{M})$. We chose not to incorporate the template within the definition of a representation for various reasons. First, it would have made the translation between MMSNP and FP harder; secondly, it would have complicated a great deal the definition of recolouring and therefore of the key notions of retracts and so forth, unless we had assumed already many properties of a representation, as being non-sbavate, simple, etc which would have made the above mentioned translation "less" one-to-one.

4.3.3 Canonical representation

Recall that in the previous chapter, we introduced the notion of "collapsed" sentences of MMSNP. We shall do something similar with simple representations. We define a representation $(T, \mathcal{M})$ to be *rigid* whenever the validity of a coloured structure is equivalent to a weaker property, namely if, and only if, the following holds.

- Any $|T|$-coloured structure (A, c^A) is valid for $(T, \mathcal{M})$ if, and only if, $A \xrightarrow{c^A} T$ and for any $M \xrightarrow{c^M} T$ in $\mathcal{M}$ and any $M \xhookrightarrow{m} A$, $c^A \circ m \neq c^M$.

That is, A is not valid if $A \not\longrightarrow T$ or for any $c^A : A \longrightarrow T$, there exists some (M, c^M) in $\mathcal{M}$ and some embedding $m : M \hookrightarrow A$ such that the following diagram commutes.

$$\begin{tikzcd} M \arrow[r, hook, "m"] \arrow[d, "c^M"', "\mathcal{M}" description] & A \arrow[dl, "c^A"] \\ T \end{tikzcd}$$

For any simple representation $(T, \mathcal{M})$, there exists a rigid representation that is equivalent up to isorecolouring, namely $(T, \mathbf{H}\mathcal{M})$, where $\mathbf{H}\mathcal{M}$ denotes the set of homomorphic images of structures from $\mathcal{M}$ that preserve the colouring. That is, for any (M, c^M) in $\mathcal{M}$, and any homomorphic image $h(M)$ of M such that the following diagram commutes,

$$\begin{tikzcd} M \arrow[r, "c^M"] \arrow[d, "h"'] & T \\ h(M) \arrow[ur, "c^{h(M)}"'] & \end{tikzcd}$$

consider the coloured structure $(h(M), c^{h(M)})$ as a new forbidden pattern. Notice that the representation hence obtained is not necessarily simple anymore; however we show easily that it is rigid. Assume that some coloured structure (A, c^A) is not valid with respect to this new representation: some forbidden pattern (M, c^M) maps into (A, c^A) via some colour preserving homomorphism m. By construction, the homomorphic image of (M, c^M) via m is also a forbidden pattern, and it embeds in (A, c^A).

We can ensure furthermore that there is no redundancy by removing from $\mathbf{H}\mathcal{M}$ those structures, a proper substructure of which also occurs in $\mathbf{H}\mathcal{M}$; *i.e.*, we simplify with respect to embedding instead of homomorphism, keeping only one isomorphic copy. Denote by $\mathbf{SH}\mathcal{M}$ the set hence obtained. Notice that it follows that $\mathbf{SH}\mathcal{M}$ contains coloured cores only. Call a representation that is rigid and simple (with respect to embedding) and whose forbidden patterns are coloured cores, a *canonical* representation. We have proved the following.

Proposition 4.18 *Any simple representation is equivalent to a canonical representation, up to isorecolouring.*

4.3.4 Valid colourings of the template and retracts

We show that if the template of the representation of a forbidden patterns problem is a yes-instance of this problem then this representation has a particular retract that is conform. Hence, the problem is in fact in CSP.

Proposition 4.19 *Let $(T, \mathcal{M})$ be a simple representation. If (T, c^T) is valid w.r.t. $(T, \mathcal{M})$ then $(T, \mathcal{M})$ has a conform retract, namely $(c^T(T), \emptyset)$.*

PROOF. Let $(T, \mathcal{M})$ be a simple representation. Assume that (T, c^T) is valid. It follows that $T \xrightarrow{c^T} T$ and that there can not be any non-conform forbidden pattern (M, c^M) in $\mathcal{M}$ such that $(M, c^M) \xrightarrow{m} (T, c^T)$. Hence, there is simply no non-conform forbidden pattern of the form $(M, c^T \circ m)$ in $\mathcal{M}$. It follows by (the remark following) Proposition 4.16 that c^T defines an endorecolouring of $(T, \mathcal{M})$ (a recolouring of $(T, \mathcal{M})$ to $(T, \mathcal{M})$). Consider its image; that is, the representation $(c^T(T), \emptyset)$. Let i be the identity of $c^T(T)$. Then $((c^T(T), \emptyset), c^T, i)$ is a retract of $(T, \mathcal{M})$. □

Notice that this result also holds for canonical representations (we do not really use the fact that the representation is simple but a weaker property possessed by canonical representations, namely that the non-conform forbidden patterns are T-coloured structures).

Theorem 4.20 *Let $(T, \mathcal{M})$ be some non-conform simple automorphic representation. There is no valid colouring for T with respect to $(T, \mathcal{M})$.*

PROOF. If T were to have a valid colouring then it would follow from the previous result that $(T, \mathcal{M})$ would have a conform retract. *i.e.* that it is equivalent to a conform representation, up to isorecolouring, since it has no proper retracts. Therefore, it would follow that it is conform itself. Which yields a contradiction. □

4.4 Feder-Vardi transformation of a representation

The idea of this transformation is directly inspired from that performed on sentences of MMSNP in the previous chapter: it consists in picking any forbidden pattern (S, c^S_μ) that can be decomposed into two components (P_0, c^{P_0}) and (P_1, c^{P_1}) with only one common articulation point x of colour $\chi \in \mu$; replacing χ by two copies χ_0 and χ_1; and making copies of the forbidden patterns accordingly (any vertex that has colour χ takes now either the colour χ_0 or the colour χ_1) except for (S, c^S_μ) which is replaced by two families of forbidden patterns: one family is induced by (P_0, c^{P_0}) and the other by (P_1, c^{P_1}); χ_0 and χ_1 replace the colour χ as above, with the exception of the articulation point x; it has colour χ_0 (respectively χ_1) in the forbidden patterns induced by (P_0, c^{P_0}) (respectively (P_1, c^{P_1})). This transformation leads to a representation that defines the same problem. As in the case of sentence of MMSNP, we would like to apply a sequence of these elementary transformations until there are only biconnected forbidden patterns remaining; but, it is not clear whether this procedure terminates. Indeed, at each step we add a colour and get about twice as many forbidden patterns as before. Notice however that this transformation concerns more the structure of a forbidden pattern than its set of colour: we can simultaneously carry out the transformation over a set of forbidden patterns that share the same structure. This leads us to the notion of a *compact* coloured structure that shall allow us to split simultaneously all forbidden patterns that share the same σ-structure.

We say that a representation $(\mu, \mathcal{M})$ is *connected* (respectively *biconnected*) if every forbidden pattern $(M, c^M) \in \mathcal{M}$ is connected (respectively biconnected).

4.4.1 Compact forbidden patterns and compact representation

We call a pair $(M, c^M_{\wp(\mu)})$ where M is a σ-structure and $c^M_{\wp(\mu)}$ a function of $|M|$ to $\wp(\mu)$ (the powerset of μ) a *compact coloured structure*. Note that in the following we see a compact coloured structure as a set of coloured structures: we see the colour set asociated with a vertex as a choice. A compact coloured structure is only a useful shorthand to prove termination. Bearing this in mind, we can extend the definition of a representation to allow compact coloured structures as forbidden patterns, and call such a representation a *compact representation*. All related

notions (*e.g.*, recolouring, validity of a colouring for a given structure *etc*) extend naturally to compact representations.
Let $(\mu, \mathcal{M})$ be a representation. We can easily transform $(\mu, \mathcal{M})$ into a compact representation; *e.g.* consider the compact representation with:

- colour set μ; and
- replace every forbidden pattern (M, c^M_μ) in $\mathcal{M}$ by $(M, c^M_{\wp(\mu)})$ where for every x in $|M|$: $c^M_{\wp(\mu)}(x) := \{c^M_\mu(x)\}$.

4.4.2 Elementary Feder-Vardi transformations

We defined the notion of a decomposition in Subsection 3.1.3 for σ-structures. This notion extends to compact coloured structures. Let $\mu \neq \emptyset$ and let $(S, c^S_{\wp(\mu)})$ be some compact coloured structure. Suppose that there exist an element x of S and two substructures of S, P_0 and P_1 satisfying the following:

- $|P_0| \cup |P_1| = |S|$;
- $|P_0| \cap |P_1| = \{x\}$;
- for every r-ary relation symbol R in σ and for any $\bar{x}$ in $|S|^r$, if $R^S(\bar{x})$ holds then either $R^{P_0}(\bar{x})$ holds or $R^{P_1}(\bar{x})$ holds but not both; and
- P_0 and P_1 have at least one tuple each.

Let c^{P_0} (respectively c^{P_1}) be the restriction of c^S to P_0 (respectively P_1). We say that the pair $((P_0, c^{P_0}), (P_1, c^{P_1}))$ forms a *decomposition* of (S, c^S_μ) in the *articulation point* x. We denote this by $(P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1})$.
Let $(\mu, \mathcal{M})$ be a compact representation such that $\mathcal{M} = \mathcal{M}' \cup (S, c^S)$ and $(P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1})$ forms a decomposition of (S, c^S). Let $C = c^{P_0}(x) = c^{P_1}(x)$. The colour sets C_0 and C_1 are defined as $\{\chi_i | \chi \in C\}$, for $i = 0, 1$. We assume furthermore that C, C_0 and C_1 are mutually disjoint. Consider the representation with:

- colour set $(\mu \setminus C)\dot{\cup}C_0\dot{\cup}C_1$; and with
- compact forbidden patterns induced from $\mathcal{M}$:
  1. (S, c^S) is replaced by the two compact forbidden patterns induced from the decomposition $(P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1})$ of (S, c^S) so that:
     - in the compact forbidden pattern (P_0, c^{P_0}), $c^{P_0}(x) = C_0$; and
     - in the compact forbidden pattern (P_1, c^{P_1}), $c^{P_1}(x) = C_1$.
  2. every remaining occurrence of a colour $\chi \in C$ in a compact forbidden pattern (including the two previous ones that have replaced (S, c^S)) is replaced by χ_0 and χ_1.

We call this representation the *elementary Feder-Vardi transformation* of $(\mu, \mathcal{M})$ with respect to $(P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1})$.

The following result shows that applying some elementary Feder-Vardi transformation to some representation does not change the problem represented.

Proposition 4.21 *Let $(\mu, \mathcal{M})$ be some compact representation such that*

$$\mathcal{M} = \mathcal{M}' \cup (P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1})$$

and let $(\nu, \mathcal{N})$ be its elementary Feder-Vardi transformation with respect to the compact forbidden pattern $(P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1})$. The following holds:

$$FP(\mu, \mathcal{M}) = FP(\nu, \mathcal{N}).$$

PROOF. Let r be the mapping of ν to μ that

- sends every $\chi_i \in C_i$ onto $\chi \in C$, for $i = 0, 1$; and
- leaves the other colours fixed.

We show that r is a recolouring. By construction, the inverse images of any forbidden pattern in $\mathcal{M}'$ belong to $\mathcal{N}$. So, it remains to check the inverse images of (S, c^S). We may assume without loss of generality that we are checking an inverse image of (S, c^S) whose vertex x takes a colour from C_0. Now consider the induced sub-coloured-structure over P_0: by construction, it is forbidden by the compact forbidden pattern with σ-structure P_0 (remember how we see a compact forbidden pattern as a shorthand for a set of forbidden patterns). This proves that r is a recolouring. So, it follows by Proposition 4.1 that $FP(\mu,\mathcal{M}) \supseteq FP(\nu,\mathcal{N})$.

We now prove the converse inclusion. Let A be some yes-instance of the problem $FP(\mu,\mathcal{M})$. There exists some $|A| \xrightarrow{c_\mu^A} \mu$ such that (A,c_μ^A) is valid with respect to $(\mu,\mathcal{M})$. Now, we construct a valid colouring c_ν^A from c_μ^A as follows. For any $y \in |A|$ such that $c_\mu^A(y) \notin C$, set $c_\nu^A(y) := c_\mu^A(y)$. Suppose now that $c_\mu^A(y) = \chi \in C$. $(P_0,c^{P_0}) \underset{x}{\bowtie} (P_1,c^{P_1})$ belongs to $\mathcal{M}$ and (A,c_μ^A) is valid with respect to $(\mu,\mathcal{M})$: it follows that $(P_0,c^{P_0}) \underset{x}{\bowtie} (P_1,c^{P_1}) \not\rightarrow (A,c_\mu^A)$. Thus, we can not have both $(P_0,c^{P_0}) \xrightarrow{h_0} (A,c_\mu^A)$ and $(P_1,c^{P_1}) \xrightarrow{h_1} (A,c_\mu^A)$, where $h_0(x) = h_1(x) = y$.

- If y is such that $(P_0,c^{P_0}) \xrightarrow{h_0} (A,c_\mu^A)$ with $h_0(x) = y$, we can not also have $(P_1,c^{P_1}) \xrightarrow{h_1} (A,c_\mu^A)$ for some h_1 such that $h_0(x) = h_1(x) = y$. Hence, we can safely set $c_\nu^A(y) := \chi_1$.
- Similarly, if y is such that $(P_1,c^{P_1}) \xrightarrow{h_1} (A,c_\mu^A)$ with $h_1(x) = y$, we can set $c_\nu^A(y) := \chi_0$.
- Otherwise, we set arbitrarily $c_\nu^A(y) := \chi_0$ or $c_\nu^A(y) := \chi_1$.

By definition of the elementary Feder-Vardi transformation (A,c_ν^A) is valid with respect to $(\nu,\mathcal{N})$ and we have proved the converse inclusion, that is, $FP(\mu,\mathcal{M}) \subseteq FP(\nu,\mathcal{N})$. □

4.4.3 Rewriting representations

We prove first that every sequence of elementary Feder-Vardi transformations is finite; and, secondly, that the representations resulting from such sequences are the same (up to isorecolouring).

Termination. Let (S,c^S) be a connected compact coloured structure. Assume that (S,c^S) requires i splittings in order to yield biconnected structures only; that is, (S, c^S) is a structure of the form:

$$(P_0, c^{P_0}) \underset{x_0}{\bowtie} ((P_1, c^{P_1}) \underset{x_1}{\bowtie} (\dots (P_{i-1}, c^{P_{i-1}}) \underset{x_{i-1}}{\bowtie} (P_i, c^{P_i}))\dots)),$$

where (P_j, c^{P_j}) is biconnected, for $j = 0, 1, \dots, i$. We call i the *rank* of the structure (S, c^S).

Let $(\mu, \mathcal{M})$ be a connected compact representation. Let a_i be the number of distinct compact forbidden patterns in $\mathcal{M}$ that have rank i. We associate to the representation $(\mu, \mathcal{M})$ the polynomial $P(X) = \Sigma_i a_i X^i$.

Recall that we want to transform a given connected representation via a sequence of elementary Feder-Vardi transformations until there are biconnected or conform forbidden patterns only.

We show that there can not be an infinite sequence of elementary Feder-Vardi transformations. After each elementary transformation, we get a polynomial that is strictly smaller; if we split according to some compact forbidden pattern of rank $j > 1$ with respect to some decomposition that leaves one forbidden pattern of rank $k < j$ and one of rank $j - k$ then we get the polynomial $P'(X) = \Sigma_i b_i X^i$ where,

$$b_i = \begin{cases} a_j - 1 & \text{, if } i = j \\ a_k + 1 & \text{, if } i = k \\ a_{j-k} + 1 & \text{, if } i = j - k \\ a_i & \text{, otherwise.} \end{cases}$$

So, we have $P' < P$, where $<$ denotes the standard linear order over polynomials (which is well-ordered) and the result follows.

Uniqueness. We prove that the order in which the elementary transformations are carried out over a given representation is not relevant [4], the representations are equivalent up to isorecolouring.

Let $\mathcal{M}$ be a set of compact forbidden patterns. We denote by $[\mathcal{M}]_{\chi \curvearrowright \chi_0 \vee \chi_1}$ the set of compact forbidden patterns obtained from $\mathcal{M}$ by replacing every occurrence of the colour χ by χ_0 and χ_1. We sometimes need to have an exception to such a replacement rule and so we denote by $[\mathcal{M} \cup \{(P_0, c^{P_0})^{x \mapsto \chi_0}\}]_{\chi \curvearrowright \chi_0 \vee \chi_1}$ the fact the

[4]In the terminology of rewriting systems, that is, if we see each elementary transformation as a rewriting rule, then our system would be said to be *locally confluent*.

replacement of χ by χ_0 and χ_1 does not apply to vertex x of (P_0,c^{P_0}) which must take colour χ_0 only. According to this notation, the elementary Feder-Vardi transformation of $(\mu, \mathcal{M}' \cup \{(P_0,c^{P_0}) \underset{x}{\bowtie} (P_1,c^{P_1})\})$ with respect to $(P_0,c^{P_0}) \underset{x}{\bowtie} (P_1,c^{P_1})$, if we assume further that vertex x has colour set $\{\chi\}$ in$(P_0,c^{P_0}) \underset{x}{\bowtie} (P_1,c^{P_1})$, has the following set of compact forbidden patterns:

$$[\mathcal{M} \cup \{(P_0,c^{P_0})^{x \mapsto \chi_0}\}, (P_1,c^{P_1})^{x \mapsto \chi_1}\}]_{\chi \curvearrowleft \chi_0 \vee \chi_1}$$

Let $(\mu, \mathcal{M})$ be a connected compact representation.
Consider first the case of different compact forbidden patterns that could be used for an elementary Feder-Vardi transformation; that is, assume that $\mathcal{M} = \mathcal{M}' \cup \{(S,c^S),(U,c^U)\}$, where: $(S,c^S) = (P_0,c^{P_0}) \underset{x}{\bowtie} (P_1,c^{P_1})$ and $(U,c^U) = (Q_0,c^{Q_0}) \underset{y}{\bowtie} (Q_1,c^{Q_1})$. We assume for simplicity that $c^S(x) = \{\chi_x\}$ and $c^U(y) = \{\chi_y\}$. There are different cases to consider.

(α_1) $\boxed{\chi_x \neq \chi_y}$
It can be easily checked that applying a transformation with respect to $(P_0,c^{P_0}) \underset{x}{\bowtie} (P_1,c^{P_1})$, followed by a transformation with respect to the compact forbidden pattern induced by $(Q_0,c^{Q_0}) \underset{y}{\bowtie} (Q_1,c^{Q_1})$ leads to the same transformation as the other way around (note also that this case is very similar to the case (β_1) which is treated thoroughly underneath).

(α_2) $\boxed{\chi_x = \chi_y = \chi}$
Splitting according to $(P_0,c^{P_0}) \underset{x}{\bowtie} (P_1,c^{P_1})$, we get:

$$[\mathcal{M}' \cup \{(P_0,c^{P_0})^{x \mapsto \chi_0}, (P_1,c^{P_1})^{x \mapsto \chi_1}, (Q_0,c^{Q_0}) \underset{y}{\bowtie} (Q_1,c^{Q_1})\}]_{\chi \curvearrowleft \chi_0 \vee \chi_1}$$

Splitting according to $(Q_0,c^{Q_0}) \underset{y}{\bowtie} (Q_1,c^{Q_1})$, we finally get:

$$\left[\begin{array}{lll} \mathcal{M}' & \cup & \{(P_0,c^{P_0})^{x \mapsto \chi_{00} \vee \chi_{01}}, (P_1,c^{P_1})^{x \mapsto \chi_{10} \vee \chi_{11}}, \\ & & (Q_0,c^{Q_0})^{y \mapsto \chi_{00} \vee \chi_{10}}, (Q_1,c^{Q_1})^{y \mapsto \chi_{01} \vee \chi_{11}}\} \end{array} \right]_{\chi \curvearrowleft \chi_{00} \vee \chi_{01} \vee \chi_{10} \vee \chi_{11}}$$

Splitting first according to $(Q_0,c^{Q_0}) \underset{y}{\bowtie} (Q_1,c^{Q_1})$ and then according to

$(P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1})$, we get:

$$\left[\begin{array}{lll} \mathcal{M}' & \cup & \{(P_0, c^{P_0})^{x \mapsto \chi_{00} \vee \chi_{10}}, (P_1, c^{P_1})^{x \mapsto \chi_{01} \vee \chi_{11}}, \\ & & (Q_0, c^{Q_0})^{y \mapsto \chi_{00} \vee \chi_{01}}, (Q_1, c^{Q_1})^{y \mapsto \chi_{10} \vee \chi_{11}}\} \end{array}\right]_{\chi \curvearrowleft \chi_{00} \vee \chi_{01} \vee \chi_{10} \vee \chi_{11}}$$

Consider the mapping r that sends χ_{ij} to χ_{ji}, for $i = 0, 1$ and $j = 0, 1$, and leaves the other colours invariant: r is clearly an isorecolouring.

Consider now the case of a compact forbidden pattern that admits two different decompositions; that is, $\mathcal{M} = [\mathcal{M}' \cup \{(S, c^S)\}]$, where:

$$(S, c^S) = ((P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1})) \underset{y}{\bowtie} (P_2, c^{P_2}) = (P_0, c^{P_0}) \underset{x}{\bowtie} ((P_1, c^{P_1}) \underset{y}{\bowtie} (P_2, c^{P_2})).$$

We assume for simplicity that $c^S(x) = \{\chi_x\}$ and $c^S(y) = \{\chi_y\}$. There are different cases to consider.

(β_1) $\boxed{\chi_x \neq \chi_y}$
Splitting according to $((P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1})) \underset{y}{\bowtie} (P_2, c^{P_2})$, we get:

$$[\mathcal{M}' \cup \{((P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1}))^{y \mapsto \chi_{y0}}, (P_2, c^{P_2})^{y \mapsto \chi_{y0}}\}]_{\chi_y \curvearrowleft \chi_{y0} \vee \chi_{y1}}$$

Splitting according to $((P_0, c^{P_0}) \underset{x}{\bowtie} (P_1, c^{P_1}))^{y \mapsto \chi_{y0}}$, we finally get the following set $\mathcal{N}$ of compact forbidden patterns;

$$\left[\begin{array}{lll} \mathcal{M}' & \cup & \{(P_0, c^{P_0})^{x \mapsto \chi_{x0}}, \\ & & (P_1, c^{P_1})^{x \mapsto \chi_{x1} \wedge y \mapsto \chi_{y0}}, (P_2, c^{P_2})^{y \mapsto \chi_{y1}}\} \end{array}\right]_{\chi_x \curvearrowleft \chi_{x0} \vee \chi_{x1} \wedge \chi_y \curvearrowleft \chi_{y0} \vee \chi_{y1}}$$

Similarly, if we proceed by first splitting according to $(P_0, c^{P_0}) \underset{x}{\bowtie} ((P_1, c^{P_1}) \underset{y}{\bowtie} (P_2, c^{P_2}))$ and then according to $((P_1, c^{P_1}) \underset{y}{\bowtie} (P_2, c^{P_2}))^{x \mapsto \chi_{x1}}$, we get the following set $\mathcal{N}'$ of compact forbidden patterns;

$$\left[\begin{array}{lll} \mathcal{M}' & \cup & \{(P_0, c^{P_0})^{x \mapsto \chi_{x0}}, \\ & & (P_1, c^{P_1})^{x \mapsto \chi_{x1} \wedge y \mapsto \chi_{y0}}, (P_2, c^{P_2})^{y \mapsto \chi_{y1}}\} \end{array}\right]_{\chi_x \curvearrowleft \chi_{x0} \vee \chi_{x1} \wedge \chi_y \curvearrowleft \chi_{y0} \vee \chi_{y1}}$$

Hence that $\mathcal{N} = \mathcal{N}'$ so as the representations obtained are identical, they are *a fortiori* equivalent up to isorecolouring.

(β_2) $\boxed{\chi_x = \chi_y = \chi}$
We leave the last case since it is very similar to case (α_1).

4.4.4 Definition

We define *the Feder-Vardi transformation* of a given connected representation to be the representation obtained from the iteration of elementary Feder-Vardi transformations until there are only biconnected forbidden patterns remaining (we then expand every compact forbidden pattern into its corresponding set of forbidden patterns).

This definition together with the previous proposition leads to the following corollary.

Corollary 4.22 *Let $(\mu, \mathcal{M})$ be a representation and $(\nu, \mathcal{N})$ its Feder-Vardi transformation. Then $FP(\mu, \mathcal{M}) = FP(\nu, \mathcal{N})$.*

4.4.5 Example

Consider the following σ_2-representation $\mathfrak{P}_2^1 := (\mathbf{1}, \{WOP_2\})$, where WOP_2 is a white directed path of length 2; *i.e.* it consists of a structure OP_2 with three elements $\{x,y,z\}$ such that $E^{OP_2} = \{(x,y),(y,z)\}$, that is coloured white (the only colour). The Feder-Vardi transformation of this representation is the following after simplification; $\Omega = (\mathbf{2}, \{WDE, BDE, BDEW\})$, where WDE, respectively BDE, consists of a single directed edge coloured in white, respectively in black and $BDEW$ consists of a single directed edge with its origin coloured in black and its target coloured in white. Indeed, a new colour has been introduced 'black' and WOP_2 has been split in y yielding two types of forbidden patterns; the first type consists of a single directed edge whose target must be coloured white, and whose origin can be either white (the original colour) or black (the copy of the original colour); and the second type consists of a single directed edge whose origin must be coloured black (the new colour) and whose target can be either white or black. This transformation is depicted on Figure 4.4.

By the previous corollary the two representations define the same problem and since the later is conform it follows from Proposition 4.13 that the problem they

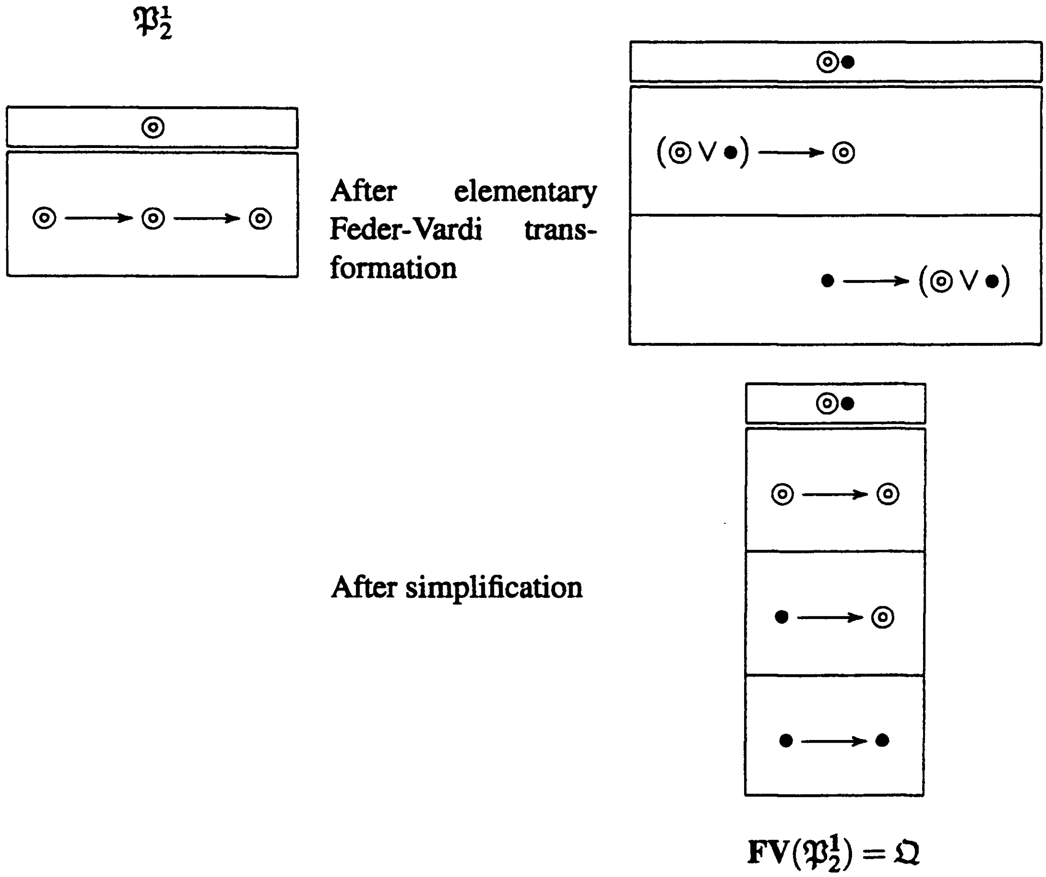

Figure 4.4: example of a Feder-Vardi transformation

define is in CSP. Notice that these representations provide a counter-example to the converse of Proposition 4.1 as there does not exist any recolouring of $\mathfrak{P}_2^1$ to $\mathfrak{O}$. For this, consider the mapping sending white to white, the inverse image of WDE is WDE and there does not exist any colour preserving homomorphism of WOP_2 (the unique forbidden pattern of $\mathfrak{P}_2^1$) to WDE; hence it is not a recolouring. Similarly the mapping sending white to black is not a recolouring. The templates of these representations are depicted on Figure 4.5 (we depicted the templates' element with their corresponding colour, however beware that the template of a representation is a structure and *not* a coloured structure).

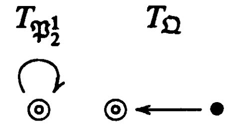

Figure 4.5: Templates of $\mathfrak{P}_2^1$ and $\mathfrak{O}$

4.4.6 Feder Vardi transformation and rigidity

We have seen previously that any simple representation is equivalent up to isorecolouring to a canonical representation. Let $(T,\mathcal{M})$ be some connected canonical representation that is not conform, that is $\mathcal{M} \neq \emptyset$. We claim that the Feder-Vardi transformation of such a representation is rigid.

Suppose there is some (non-conform) forbidden pattern (S,c^S) in $\mathcal{M}$ that admits a decomposition $(P_0,c^{P_0}) \underset{x}{\bowtie} (P_1,c^{P_1})$. Let $(\nu,\mathcal{N})$ be the representation obtained from $(T,\mathcal{M})$ via the elementary Feder-Vardi transformation with respect to $(P_0,c^{P_0}) \underset{x}{\bowtie} (P_1,c^{P_1})$. We have seen in the proof of Proposition 4.21 that there exists a recolouring r of $(\nu,\mathcal{N})$ to $(T,\mathcal{M})$. Furthermore, since $(T,\mathcal{M})$ is canonical it is non-sbavate and by construction so is $(\nu,\mathcal{N})$. Hence if T' is the template of $(\nu,\mathcal{N})$, it follows by Proposition 4.15 that $T' \xrightarrow{r} T$. Let A be some non-valid structure of the problem represented by $(\nu,\mathcal{N})$ (and $(T,\mathcal{M})$ by Proposition 4.21). If $A \xrightarrow{c^A} T'$ then $A \xrightarrow{r \circ c^A} T$. $A \notin FP(T,\mathcal{M})$ and $(T,\mathcal{M})$ rigid implies that there is some $(M, r \circ c^M)$ in $\mathcal{M}$ (recall that r is surjective) such that $(M, r\circ c^M) \xhookrightarrow{m} (A, r \circ c^A)$. Thus, we have $(M,c^M) \xhookrightarrow{m} (A,c^A)$. Now by construction, either (M,c^M) is a forbidden pattern of the new representation or if $(M, r\circ c^M)$ is (S,c^S) then without loss of generality we may assume that some forbidden pattern induced by (P_0, c^{P_0}) is a substructure of (M, c^M). Hence, in any case some forbidden pattern of the $(\nu, \mathcal{N})$ embeds in (A, c^A) by composition. We have therefore proved that $(\nu, \mathcal{N})$ is rigid.

Notice that it is however not necessarily canonical, but it can be altered slightly to obtain a canonical representation; each forbidden pattern can be replaced by its coloured core without affecting the property of being rigid. Furthermore, the set of forbidden patterns can be simplified with respect to embedding without affecting this key property either. Finally, if some forbidden pattern is not properly T'-coloured, simply discard it. We denote by $\mathbf{FV}(T, \mathcal{M})$ the canonical representation hence obtained. Notice that by construction if $(T, \mathcal{M})$ was connected then $\mathbf{FV}(T, \mathcal{M})$ is biconnected.

4.5 Normal representation

In this section we define the *normal form of a connected representation* $(\mu, \mathcal{M})$; essentially, it is an automorphic, biconnected and canonical representation that is equivalent to $(\mu, \mathcal{M})$ (*i.e.* it represents the same problem). Constructing the normal form involves the notions of a core (of a representation) from Section 4.2 and of a Feder-Vardi transformation from Section 4.4.

This section is organised as follows. In Subsection 4.5.1, we define the normal form of a connected representation. In Subsection 4.5.2, we illustrate this notion by computing the normal form of numerous examples.

4.5.1 Definition

Informally the normal form of a canonical connected representation $(T, \mathcal{M})$ is built as follows. First, consider its canonical Feder-Vardi transformation $\mathbf{FV}(T, \mathcal{M})$; recall that it has the following properties:

- each forbidden pattern is biconnected; and
- it is canonical (rigid and simple with respect to embeddings).

Secondly, we want to construct an automorphic representation from $\mathbf{FV}(T, \mathcal{M})$ but keeping the two above properties. So, if $\mathbf{FV}(T, \mathcal{M})$ is automorphic, we are done. Otherwise, we are going to take its core. Recall that the core of a representation is defined up to isorecolouring. So, we consider a particular core to make sure that the key properties listed above are preserved.

Let $\mathfrak{R}$ be a connected representation. Let $(T, \mathcal{M})$ be the canonical representation equivalent to $\mathfrak{R}$ (via some isorecolouring).

1. If $\mathbf{FV}(T, \mathcal{M})$ is automorphic then set $\mathbf{normal}(\mathfrak{R}) := \mathbf{FV}(T, \mathcal{M})$.
2. Otherwise, consider its core $\mathbf{core}(\mathbf{FV}(T, \mathcal{M}))$; that is, $\mathbf{core}(\mathbf{FV}(T, \mathcal{M}))$ is automorphic and there exist some epirecolouring s and some monorecolouring i with $s \circ i = id$ such that $(\mathbf{core}(\mathbf{FV}(T, \mathcal{M})), s, i)$ is a retract of $\mathbf{FV}(T, \mathcal{M})$. Set $\mathbf{normal}(\mathfrak{R})$ to be the subrepresentation of $\mathbf{FV}(T, \mathcal{M})$ *induced* by the monorecolouring i.

We call $\mathbf{normal}(\mathfrak{R})$ the *normal representation* of $\mathfrak{R}$.

The following result shows that the above construction has the properties we required.

Theorem 4.23 *Let $\mathfrak{R}$ be a connected representation.* $\mathbf{normal}(\mathfrak{R})$ *is an automorphic biconnected and canonical representation such that:*

$$FP(\mathfrak{R}) = FP(\mathbf{normal}(\mathfrak{R})).$$

PROOF. We use the same notation as in the above definition. case (1) is clear. We now deal with case (2). Let μ be the colour set of $\mathbf{FV}(T,\mathcal{M})$ and ν that of $\mathbf{core}(\mathbf{FV}(T,\mathcal{M}))$. We show that $\mathbf{core}(\mathbf{FV}(T,\mathcal{M}))$ and $\mathbf{normal}(\mathfrak{R})$ are equivalent up to isorecolouring: More precisely, we show that id_ν is an isorecolouring.

Let (N,c^N_ν) be a forbidden pattern of $\mathbf{normal}(\mathfrak{R})$. Recall that $\mathbf{normal}(\mathfrak{R})$ is an induced subrepresentation of $\mathbf{FV}(T,\mathcal{M})$: that is, by definition, (N,c^N_ν) is a forbidden pattern of $\mathbf{normal}(\mathfrak{R})$ whenever $(N, i\circ c^N_\nu)$ is a forbidden pattern of $\mathbf{FV}(T,\mathcal{M})$. i is a recolouring of $\mathbf{core}(\mathbf{FV}(T,\mathcal{M}))$ to $\mathbf{FV}(T,\mathcal{M})$ implies that the coloured structure (N,c^N_ν) (the inverse image of the forbidden pattern (N,c^N_ν) via id_ν) is not valid for $\mathbf{core}(\mathbf{FV}(T,\mathcal{M}))$. This proves that:

$$\mathbf{core}(\mathbf{FV}(T,\mathcal{M})) \xrightarrow{id_\nu} \mathbf{normal}(\mathfrak{R}).$$

Let (N,c^N_ν) be a forbidden pattern of $\mathbf{core}(\mathbf{FV}(T,\mathcal{M}))$. Recall that $s\circ i = id_\nu$. Since s is a recolouring of $\mathbf{FV}(T,\mathcal{M})$ to $\mathbf{core}(\mathbf{FV}(T,\mathcal{M}))$, it follows that there exists some forbidden pattern (M,c^M_μ) of $\mathbf{FV}(T,\mathcal{M})$ and some homomorphism $(M,c^M_\mu) \xrightarrow{m} (N, i\circ c^N_\nu)$. This means that $c^M_\mu = i\circ c^N_\nu \circ m$ and it follows by definition of an induced subrepresentation that $(M, c^N_\nu\circ m)$ is a forbidden pattern of $\mathbf{normal}(\mathfrak{R})$ such that $(M,c^N_\nu\circ m) \xrightarrow{m} (N,c^N_\nu)$. Thus, (N,c^N_ν) (the inverse image of (N,c^N_ν) via id_ν) is not valid for $\mathbf{normal}(\mathfrak{R})$. We have proved that:

$$\mathbf{normal}(\mathfrak{R}) \xrightarrow{id_\nu} \mathbf{core}(\mathbf{FV}(T,\mathcal{M})).$$

It follows directly from the definition that: $\mathbf{normal}(\mathfrak{R})$ is biconnected; every of its forbidden patterns are coloured cores; and, it is simple with respect to embeddings (any non-conform forbidden pattern is not a substructure of another non-conform forbidden pattern). We show that it is also rigid. Let T' be the template of $\mathbf{normal}(\mathfrak{R})$. $\mathbf{normal}(\mathfrak{R})$ is connected. So if it has some non-sbavate forbidden pattern then it must be a forbidden pattern that consists of a single vertex and no tuple (a forbidden pattern that forbids a colour). But $\mathbf{normal}(\mathfrak{R})$ can not have such a forbidden pattern since it is also automorphic. Hence $\mathbf{normal}(\mathfrak{R})$ is non-sbavate. Thus, by Proposition 4.15, it follows that $T' \xrightarrow{i} T$. Let A be some no-instance of $FP(\mathbf{FV}(T,\mathcal{M}))$. Recall that $\mathbf{FV}(T,\mathcal{M})$ is rigid (it is canonical). If $A \xrightarrow{c^A} T'$ then by composition $A \xrightarrow{i \circ c^A_V} T$. Thus, there exists some non-conform forbidden pattern (M, c^M_μ) of $\mathbf{FV}(T,\mathcal{M})$ such that $(M, c^M_\mu) \xhookrightarrow{m} (A, i \circ c^A_V)$. Thus $c^M_\mu = i \circ c^A_V \circ m$ holds and it follows that $(M, c^A_V \circ m)$ is a forbidden pattern of $\mathbf{normal}(\mathfrak{R})$. Hence, we have proved that $(M, c^A_V \circ m) \xhookrightarrow{m} (A, c^A_V)$. This proves that $\mathbf{normal}(\mathfrak{R})$ is rigid. □

Remark. The construction of the normal form of a given connected representation $\mathfrak{R}$ can be summarised as follows.

$$
\begin{tikzcd}
\mathfrak{R} \arrow[dd, "\approx"', "\text{isorecolouring}"] & \\
& \text{(Subsection 4.3.3)} \\
(T,\mathcal{M}) \arrow[dd, equal, "\text{same problem}"] & \text{canonical representation} \\
& \text{(Subsection 4.4.6)} \\
\mathbf{FV}(T,\mathcal{M}) \arrow[uu, bend left, "\text{epirecolouring}"] \arrow[dd, bend left, "s"] & \text{canonical Feder-Vardi transformation} \\
& \text{(Subsection 4.2.3)} \\
\mathbf{core}(\mathbf{FV}(T,\mathcal{M})) \arrow[uu, bend left, "i"] \arrow[dd, "\approx"', "\text{isorecolouring}"] & \text{core (of the representation)} \\
& \text{(this section)} \\
\mathbf{normal}(\mathfrak{R}) \arrow[uuuu, bend left=60, "i"] & \text{induced subrepresentation}
\end{tikzcd}
$$

4.5.2 Examples

Consider the following representation.

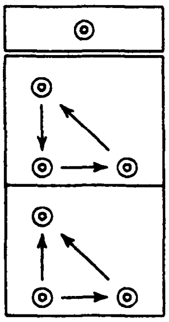

Notice that it corresponds to the problem TRI-FREE that was introduced in Section 2.4. We can make it canonical as in Subsection 4.3.3:

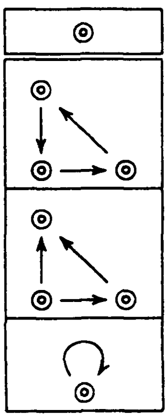

Note that the above representation is also automorphic and biconnected. So, it is the normal form.

Consider as another example, the representation of the problem TRI-FREE-TRI defined in Section 2.4,

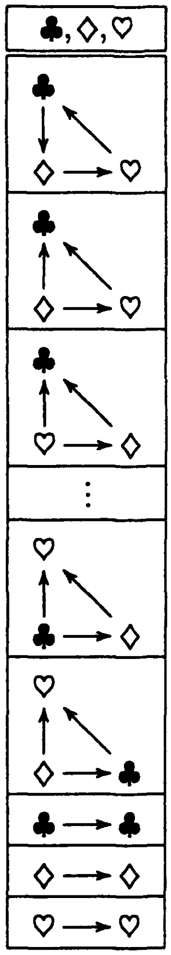

We show that this is already the normal form. Note first that it is rigid and that every forbidden pattern is a biconnected coloured core. It remains to show that it is automorphic. Assume that this was not the case and that $(\mathfrak{R}, s, i)$ is a proper retract of this representation. We may assume w.l.o.g. that $\clubsuit$ is a colour of $\mathfrak{R}$, $i(\clubsuit) = \clubsuit$ and $s(\clubsuit) = s(\heartsuit)$. It follows that:

$$\clubsuit \longrightarrow \clubsuit$$

is a forbidden pattern of $\mathfrak{R}$. However, one of its inverse image via s is the following:

$$\clubsuit \longrightarrow \heartsuit$$

It is valid, so s is not a recolouring. This yields a contradiction.

Consider now, the representation of NO-WALK-5 from the same Section. It has a single colour and as forbidden patterns all possible orientations of the undirected 5-cycle. In particular the following:

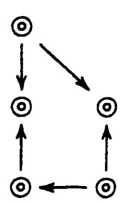

which has the following as

homomorphic image:

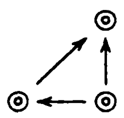

via the homomorphism that identifies element as depicted as follows by double arrows:

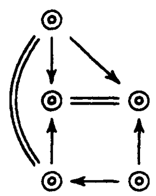

Similarly, we get:

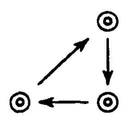

for:

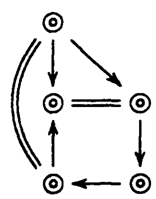

So, making this representation rigid by taking all possible homomorphic images and simplifying with respect to embedding, we get the following representation:

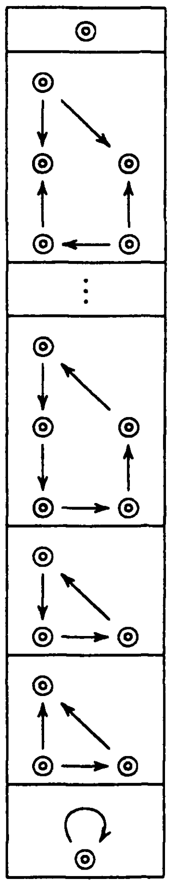

Now, every forbidden pattern being biconnected and the representation being clearly automorphic, the above depicts in fact the normal form of the representation of the problem No-Walk-5.

The restriction No-Walk-5-Tri of the previous problem as defined in Section 2.4 can be depicted as follows:

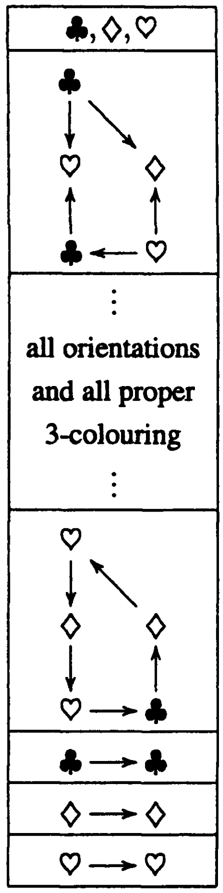

This case is similar to the previous one. Its normal form is as follows:

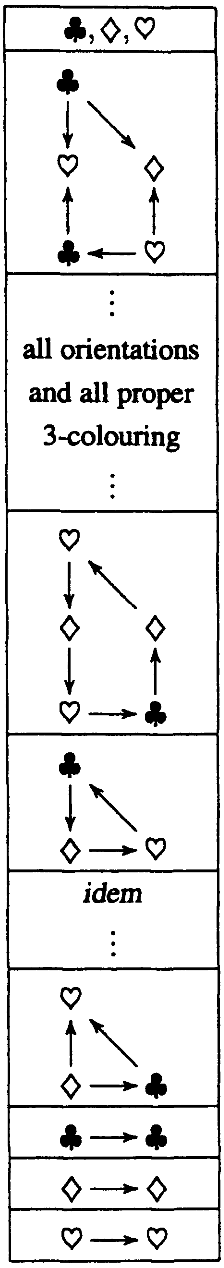

In the above, by 'proper 3-colouring', we mean that the extremities of any edge have different colours.

We shall now compute some of the normal form of the representations we introduced in this Chapter. We leave the cases of $\mathfrak{MDC}^2_3$ and $\mathfrak{MDC}^2_4$ as an exercise for the reader (it is enough to make them canonical as in the previous examples). The normal form of $\mathfrak{MDC}^2_3$ can be depicted as follows:

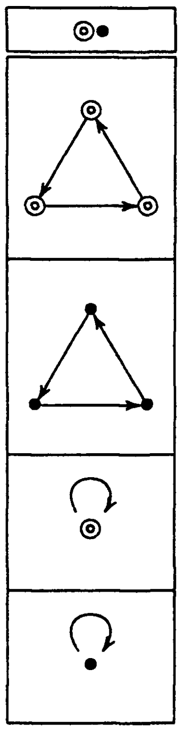

The normal form of $\mathfrak{MDC}^2_4$ can be depicted as follows:

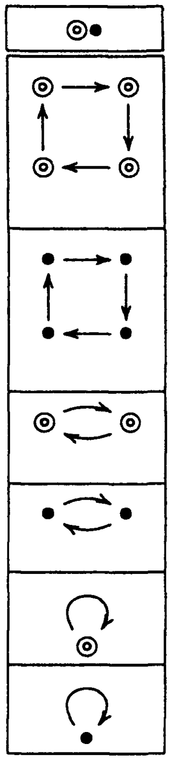

The case of $\mathfrak{MDC}^2_5$ is more interesting; it is the first example of a case where we need to apply the Feder-Vardi transformation. The two colours play a symmetric role, so we may consider only the case of the white forbidden patterns. There are two types of homomorphic image of the directed 5-cycle; the homomorphic images which contain WC_1 (a self-loop), and that which contain both WC_2 and WC_3 but no WC_1.

As WC_1 is also a homomorphic image of WC_5, we may ignore the structures of the first type, as they shall be simplified out later. There are only two structures of the second type (up to isomorphism):

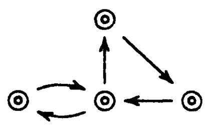

and

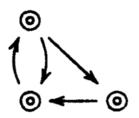

The first structure is not biconnected; hence, during the Feder-Vardi transformation, the colour $\circledcirc$ shall be replaced by two new colours, say, $\diamondsuit$ and $\heartsuit$; and, this structure is replaced by the following two (compact) forbidden patterns (we leave $\circledcirc$ as an abbreviation for $\{\diamondsuit, \heartsuit\}$):

and

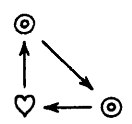

As for the second structure, it can be ignored; it shall be simplified later by one of the two previous forbidden patterns (depending on the choice of the colour). For example,

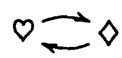

embeds in:

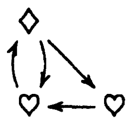

Note that this is a general property of the homomorphic image (S, P^S) of any $(P_0, c^{P_0}) \bowtie_x (P_1, c^{P_1})$ that satisfies that both (P_0, c^{P_0}) and (P_1, c^{P_1}) are substructures of (S, P^S). From now on, we shall ignore such homomorphic images. The case of the black forbidden patterns is symmetric: we denote by $\clubsuit$ and $\spadesuit$ the two copies of the colour $\bullet$ and as above $\bullet$ stands for $\{\clubsuit, \spadesuit\}$. The Feder-Vardi transformation of the canonical representation equivalent to $\mathfrak{MDC}_5^2$ can be depicted as follows:

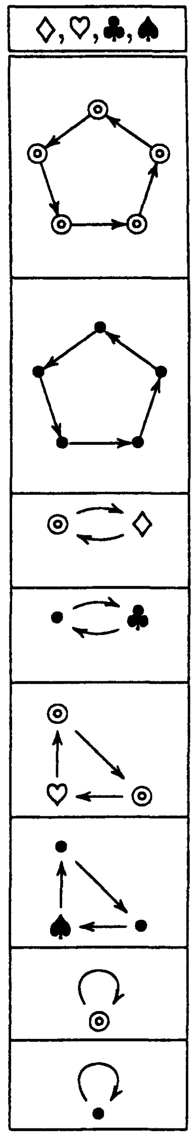

The above depicts the normal form of $\mathfrak{MDC}_5^2$; indeed, it is easy to check that it is automorphic.

Computing the normal form of $\mathfrak{MDC}_n^2$ becomes more tedious as n increases; indeed, there are more possible homomorphic images and in particular more non-biconnected homomorphic images that need to be split. However, notice that any of the biconnected components of the homomorphic image of a directed cycle is non-conform, hence the normal form of any of these representations is not conform.

Consider now the case of representations $\mathfrak{AD}\mathfrak{C}^2_{2p} - \mathfrak{MC}$. The case of $\mathfrak{AD}\mathfrak{C}^2_4 - \mathfrak{MC}$ is easy (no Feder-Vardi transformation is needed) and its normal form can be depicted as follows:

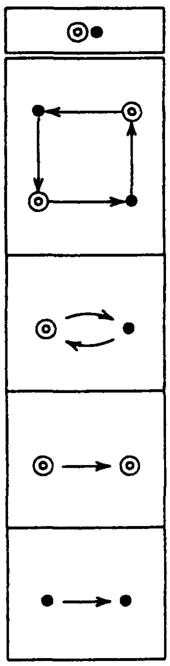

For $p = 3$ and $p = 4$, the normal form is not difficult either and there is no need to split. It becomes more interesting for $p = 5$. As in the case of $\mathfrak{MD}\mathfrak{C}^2_5$, it suffices to consider the following homomorphic images of ADC_{10}:

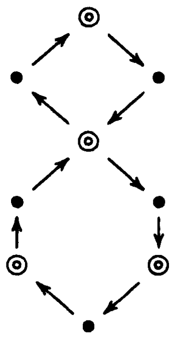

its symmetric:

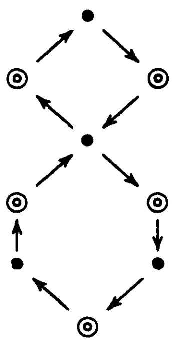

and:

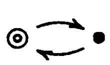

Using the same notation as above, after Feder-Vardi transformation, we finally get (note that some of the compact forbidden patterns represent the same forbidden pattern, so to be completely coherent with the definition of the normal form, we should have listed all possibilities; we beg the reader for some comprehension):

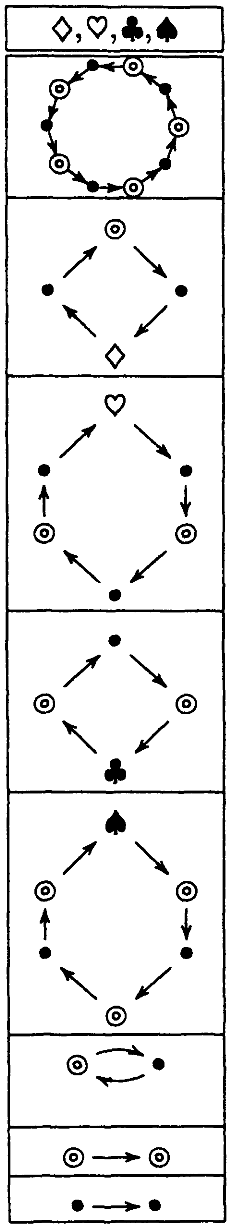

Using a similar argument as in the case of $\mathfrak{MDC}^2_n$, we can prove that: for any $p \geq 1$, the representation $\mathfrak{ADC}^2_{2p}$ has a normal form that is not conform.

It leaves the case of the representations $\mathfrak{MC}^2_n$. We have seen previously that for even n, these representations are all equivalent up to isorecolouring; and, the normal form of $\mathfrak{MC}^2_2$ can be depicted as follows:

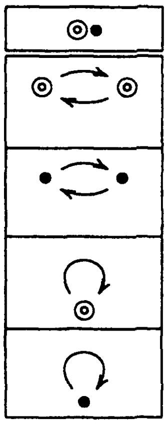

We consider the case of odd n. The same argument as the one used before can be applied to show that the normal form of any representation $\mathfrak{MC}^2_n$ is not conform. We leave as an exercise to the reader that the following depicts the normal form of the representation $\mathfrak{MC}^2_5$.

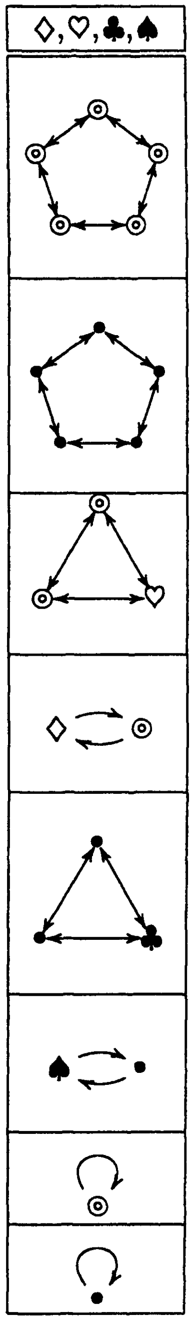

4.6 Witness families

In the first part of the present section, we introduce our main tool to prove that a forbidden patterns problem is not a homomorphism problem, namely a *witness family*. Informally, it can be seen as a particular winning strategy for Spoiler in the following two player game. A representation $\mathfrak{R}$ is given. The first player, *Duplicator*, wants to show that the forbidden patterns problem represented by $\mathfrak{R}$ is in fact a homomorphism problem. The second player, *Spoiler*, wants to prove him wrong. At each round, Duplicator provides some structure B, claiming that the homomorphism problem with template B is the same problem as the forbidden patterns problem represented by $\mathfrak{R}$. Spoiler proves him wrong by giving him either a yes-instance A of $FP(\mathfrak{R})$ such that $A \not\longrightarrow B$ or a no-instance A such that $A \xrightarrow{h} B$. If Spoiler is unable to do this at some round then he has lost the game, otherwise if Spoiler can keep Duplicator going for ever then Spoiler wins. More formally, a witness family for $\mathfrak{R}$ consists of a family of structures $\mathcal{F}$ that are all yes-instances of $FP(\mathfrak{R})$ such that for any fixed σ-structure B (which is a possible candidate for a template if the problem were to be a homomorphism problem) there exists a structure A in $\mathcal{F}$ that witnesses that B can not be such a template. That is, such that either $A \not\longrightarrow B$ or for some $A \xrightarrow{h} B$, $h(A)$ is not in $FP(\mathfrak{R})$.

In the second part of the present section, we only ever consider connected normal representations. If a problem is given by a connected normal representation that is not conform, we shall build a witness family.

The idea behind the construction is as follows. Suppose we have a normal representation $(T, \mathcal{M})$ that is not conform and a structure N that is not valid. Assume further that there exists a colouring c^N for N that is not valid and that (N, c^N) has the following property:

- $N \xrightarrow{c^N} T$ (the colouring c^N is "OK on the edges"); and,
- there exists exactly one forbidden pattern (M, c^M) in $\mathcal{M}$ ((M, c^M) must be a biconnected non-conform forbidden pattern as $(T, \mathcal{M})$ is normal) such that $(M, c^M) \xhookrightarrow{e} (N, c^N)$ and exactly one such embedding e (the colouring is "wrong" but only because of a single occurrence of a biconnected non-conform forbidden pattern).

We can "open-up" this colouring of N: pick some vertex u from this single occurrence of (M, c^M); add a copy v of u; and, from this single occurrence of (M, c^M), pick a tuple t that involves u and replace one occurrence of u in t by v. We call this new structure informally the gadget and u and v its plug-points.

When given some undirected graph G, we can build a large structure S as follows: replace every edge between two vertices x and y of G by a copy of the gadget (identify u with x and v with y).

The structure S is a yes-instance whenever the graph G has a girth higher than the following parameter of the representation $(T, \mathcal{M})$: the size of the largest cycle that is a substructure of any forbidden pattern.

Now, for any candidate B to the role of template for our problem (assume our problem to be in CSP), provided G has a chromatic number higher than the size of this candidate B, any homomorphism of the structure S to B must identify the two plug-points of some copy of the gadget. Hence some homomorphic image of N is a substructure of B and B is a no-instance: therefore, B can not be the template of our problem.

Given Erdös' result on graphs of high girth and high chromatic number, we are therefore able to rule out any B by constructing some witness S from an adequate graph G.

In the examples of this construction described in the following we use the language of graph theory to describe the various structures involved and consider the structures to be graphs even though they should really be directed graphs (all the graphs in the following can be easily transformed into directed graphs in a suitable way). It should be noted that this construction works for problems that correspond to a first-order MMSNP sentence. Consider, for example, the problem TRI-FREE: the structure N in this case is simply a triangle, and opening up N leads to a path of length 3. Call u and v the extremities of this path. Now, if G is a graph of girth g, the structure S obtained by replacing every edge between two vertices x and y by a copy of the path of length 3, identifying u with x and v with y has girth $3g$. So if we consider G to be self-loop-free, that is $g > 1$, S is triangle-free. This construction also works for more complex problems like NO-MONO-TRI: one can consider for the structure N, the 5-clique coloured as follows; 3 vertices coloured in black and the two remaining coloured in white. One can open it to obtain the following gadget: take a 4-clique, add two vertices u and v; connect

u to two elements of the 4-clique and v to the two other elements. Consider the following colouring of this gadget: set u and v to be black; and, for both u and v, one neighbour is black and one neighbour is white. The distance between u and v being 3, any structure S induced by a graph G of girth $g > 1$ is a yes-instance; it can be coloured according to the colouring of the gadget described above; and, any cycle of S that is not a substructure of a copy of the gadget has size strictly greater than 3 (hence, it can not correspond to a forbidden pattern).

For this construction to work we need a structure N that is not valid and for which there is a colouring with a *single* occurrence of a forbidden pattern, or more precisely that can be opened up to yield a structure (the gadget) that has a valid colouring that sends u and v (the plug-points) to the *same colour*. At first I thought that such a property can be achieved by enforcing some condition of minimality on the considered representation. As to whether this is the case remains open, but I was led to the notions of a recolouring and an automorphic representation and consequently to the notion of a normal representation. However the key idea of this construction can be reused. We proved that for any normal representation that is not conform there are non-valid structures N that can be nonetheless coloured in a correct way on the edges; in other words, whose colouring is a homomorphism of N to the template of the considered representation. According to this colouring, the structure N can be opened up, leading to a gadget that is not necessarily a "bipede creature" as above but a many-legged one, a "centipede"... So we can no longer use Erdös' result.

In order to build a large structure, we shall have some set of special vertices corresponding to each type of "leg"(the type of a "leg" being given by the corresponding vertex in N). We can plug copies of the "centipede" in all possible ways between those sets. If the large structures we obtain are always valid then we have a family of witnesses (just like in our examples above) and we are done. If one of the large structure is not valid then we can still colour it via the colouring induced by the colouring of N in such a way that we have a homomorphism into the template of the representation of the considered problem. We can open up this structure and obtain hence some larger structure than the "centipede" we had before, obtaining a new many-legged gadget, let's call it a "millipede" (as a matter of fact it does not have necessarily more legs it is just larger). By carefully choosing the way we open up, we make sure that the large structures obtained from the "millipede" are "sparser". If these large structures obtained from the "millipede" are still not valid then we carry own opening-up: we obtain eventually a family of witnesses.

4.6.1 Definition

We formally define a witness family as follows.

Definition 4.24 *A family of σ-structures $\mathcal{F}$ is said to be a* witness family *for a representation $\mathfrak{R}$ if:*

- $\mathcal{F} \subseteq FP(\mathfrak{R})$*; and*
- *for any σ-structure B, there exists some A in $\mathcal{F}$ such that,*
  - *either $A \notin CSP(B)$; or*
  - *for some $A \xrightarrow{h} B$, $h(A) \notin FP(\mathfrak{R})$.*

The following result is the corner-stone of the proof of our main result.

Lemma 4.25 *If a representation $\mathfrak{R}$ has a witness family then the problem $FP(\mathfrak{R})$ is not a homomorphism problem.*

PROOF. Let $\mathcal{F}$ be a witness family for $\mathfrak{R}$. If $FP(\mathfrak{R})$ were a homomorphism problem with template B then we would have some $A \in FP(\mathfrak{R})$ such that either $A \notin CSP(B)$, or for some $A \xrightarrow{h} B$, $h(A) \notin FP(\mathfrak{R})$, that is either $FP(\mathfrak{R}) \ni A \notin CSP(B)$ or $FP(\mathfrak{R}) \not\ni h(A) \in CSP(B)$, in any case a contradiction. □

We would like to construct a witness family in a generic way for problems given by representations for which we are not able to construct a template; that is, that are not conform. We shall make use of two important features of normal representations that are not conform: first, they are automorphic, therefore by Theorem 4.20 their templates must be no-instances; secondly, every non-conform forbidden pattern is biconnected, by Theorem 4.23.

4.6.2 Opening-up an invalid structure

Let M be a structure and C a cycle such that $C \overset{e}{\hookrightarrow} M$. Let x_0 be some articulation point of C. If C is the 1-cycle $R(\bar{x})$ (with x_0 occurring at least twice in $\bar{x}$) then let G be the structure defined from M and C as follows:

- $|G| := |M| \dot{\cup} \{y_1\}$; and
- G agrees with M everywhere except that the tuple $R(e(\bar{x}))$ is replaced by $R(\bar{y})$, where $\bar{y}$ is obtained from $e(\bar{x})$ by replacing the **first** occurrence of $y_0 := e(x_0)$ by y_1.

If C is a n-cycle ($n > 1$) and $R(\bar{x})$ a tuple from C such that (the articulation point) x_0 occurs in $\bar{x}$ then let G be the structure defined from M and C as follows:

- $|G| := |M| \dot{\cup} \{y_1\}$; and
- G agrees with M everywhere except that the tuple $R(e(\bar{x}))$ is replaced by $R(\bar{y})$, where $\bar{y}$ is obtained from $e(\bar{x})$ by replacing **every** occurrence of $y_0 := e(x_0)$ by y_1.

We call G the *opening* of M with respect to $C, e, R(\bar{x})$ and x_0. We call y_0 and y_1 the *plug-points* of G. Notice that the mapping that sends y_1 to y_0 and fixes the other elements is a homomorphism of G to M.

We extend this definition to coloured structures, setting the colour of the new vertex y_1 to be the same colour as y_0. Figure 4.6 illustrates this construction (notice that in this case there was only one occurrence of y_0 in the tuple $R(e(\bar{x}))$).

EXAMPLE. Let σ_3 be the signature consisting of a single ternary symbol R.

1. Let M be the σ_3-structure with domain $\{a,b,c,d\}$ and let $R^M := \{(a,b,c),(a,d,a)\}$. Consider the 1-cycle $R(x,y,x)$ and let e be the embedding from $R(x,y,x)$ to M defined by $e(x) = a$ and $e(y) = d$. The opening up of M with respect to $R(x,y,x)$ and e in the articulation point x is isomorphic to the structure G with domain $\{a,a',b,c,d\}$ with $R^G = \{(a,b,c),(a',d,a)\}$.
2. Let N be the σ_3-structure with $R^N := \{(a,a,b),(a,b,c),(b,c,d),(a,d,c)\}$ over the domain $\{a,b,c,d\}$. Consider the 3-cycle C with domain $\{x,y,z,t\}$ and $R^C = \{(x,x,y),(x,y,z),(y,z,t)\}$ and let f be the embedding defined by $f(x) = a, f(y) = b$, $f(z) = c$ and $f(t) = d$. The opening up of M with respect to C, e and the tuple $R(x,x,y)$ in

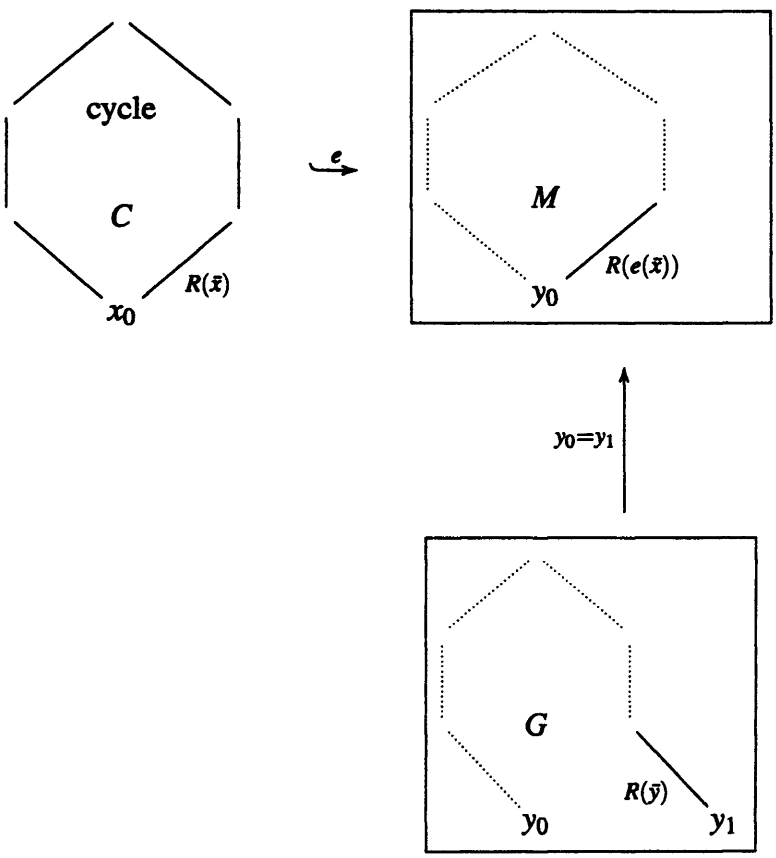

Figure 4.6: Opening a coloured structure

the articulation point x is isomorphic to the structure H with domain $\{a,d',b,c,d\}$ with $R^G = \{(d',d',b),(a,b,c),(b,c,d),(a,d,c)\}$.

▲

In the remainder of this section, let $(T,\mathcal{M})$ be some non-conform normal representation and let (N,c^N) be non-valid with respect to $(T,\mathcal{M})$ such that $N \xrightarrow{c^N} T$.

Since $(T,\mathcal{M})$ is rigid and $N \xrightarrow{c^N} T$, there exists some $(M,c^M) \in \mathcal{M}$ such that $(M,c^M) \xhookrightarrow{e} (N,c^N)$. Since $(T,\mathcal{M})$ is normal, it follows that (M,c^M) is biconnected and therefore that it contains a cycle C. Let $R(\bar{x})$ be a tuple in C and x_0 an articulation point of C with $x_0 \in \{\bar{x}\}$. Let (G,c^G) be the opening of (N,c^N) with respect to $C, e\restriction_C, R(\bar{x})$ and x_0. If (G,c^G) is not valid with respect to $(T,\mathcal{M})$, start this construction over again. Denote by (G,c^G) the valid structure eventually obtained and let

$$\{y_{1,1},y_{1,2},\ldots,y_{1,p_1},y_{2,1},y_{2,2},\ldots,y_{2,p_2},\ldots,y_{q,1},y_{q,2},\ldots,y_{q,p_q}\}$$

be its set of plug-points (the first index giving the type of a plug-point, that is, the $y_{i,-}$'s correspond to the same element of N); in other words $(G,c^G) \xrightarrow{f} (N,c^N)$, where f identifies the plug-points of the same type,

$$\begin{array}{rccl} f: & G & \rightarrow & N \\ & y & \mapsto & \begin{cases} y & \text{, if } y \in |N|, \\ y_{i,1} & \text{, if there is some } 1 \leq j \leq p_i \text{ such that } y = y_{i,j}. \end{cases} \end{array}$$

EXAMPLE.

Refer to Section 4.5.2 for the normal form of the representation $\mathfrak{ADC}^2_{10}$. Its template is as follows.

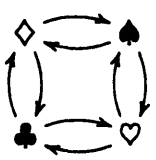

We are going to gradually open it up, considering it as a coloured structure as depicted on the previous figure. Notice that there are many ways of opening up. We highlight the considered forbidden pattern at each stage by using dotted arrows (which shall be seen as a cycle in our case), except for the tuple considered which shall be depicted by a dashed arrow. Moreover we mark the chosen articulation point by enclosing it within a circle. For example, opening up the template according to the following,

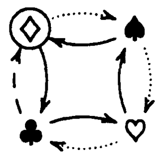

yields the coloured structure.

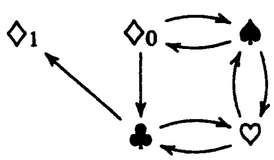

The latter is not valid and we open it up further.

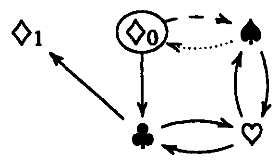

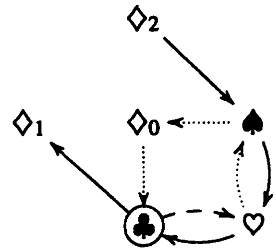

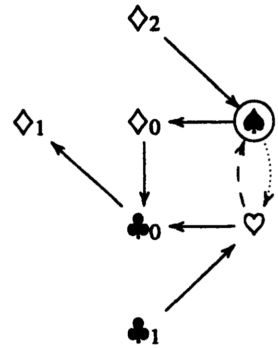

Finally, we obtain the following valid structure.

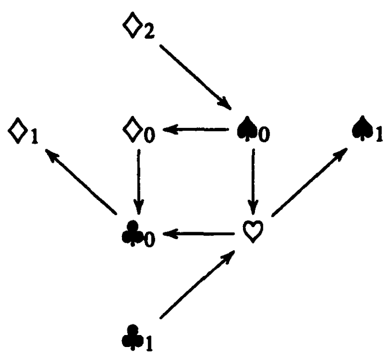

The latter has three types of plug points, that we denoted on the figure by $\diamondsuit_0$, $\diamondsuit_1$, $\diamondsuit_2$, $\clubsuit_0$, $\clubsuit_1$, $\spadesuit_0$ and $\spadesuit_1$.

▲

4.6.3 Constructing a large coloured structure

Let (G, c^G) be a valid coloured structure (informally called the gadget) obtained from some non-valid (N, c^N) as in the previous section. For any $\bar{n} = (n_1, n_2, \ldots, n_q)$ with $n_1 \geq p_1, n_2 \geq p_2, \ldots, n_q \geq p_q$, we build a large coloured structure $(I_{\bar{n}}, c^{I_{\bar{n}}})$ from the gadget as follows. It has a set of *special elements* $|S| \subseteq |I_{\bar{n}}|$ that is partitioned into q pairwise disjoint sets $X_i := \{x_{i,j} | 1 \leq j \leq n_i\}$ $(1 \leq i \leq q)$. For any $1 \leq i \leq q$ and for any choice of p_i elements $x_{i,k_1}, x_{i,k_2}, \ldots, x_{i,k_{p_i}}$ in $|X_i|$ such that $k_1 < k_2 < \ldots < k_{p_i}$, plug in a copy of the gadget (G, c^G), identifying the plug-points of (G, c^G) with the corresponding chosen special vertices; that is, set $x_{i,k_j} := y_{i,j}$ for any $1 \leq i \leq q$ and for any $1 \leq j \leq p_i$.

EXAMPLE. Depicting a large structure with the gadget used in the previous example would not be really helpful as the gadget obtained there is quite complicated. We build therefore an alternative gadget first. Consider for this the structure DC_2. It is clearly a no-instance of the problem $FP(\mathfrak{ADC}^2_{10})$. However, it can be coloured to obtain a valid colouring with respect to the template of the normal representation of $\mathfrak{ADC}^2_{10}$ as follows.

$$\diamondsuit \rightleftarrows \spadesuit$$

According to this colouring, we can open up to obtain the following gadget.

$$\diamondsuit_0 \longrightarrow \spadesuit \longrightarrow \diamondsuit_1$$

It has only two plug-points, denoted by $\diamondsuit_0$ and $\diamondsuit_1$, respectively. The following depicts the "large" coloured structure obtained using this gadget for $n = 3$.

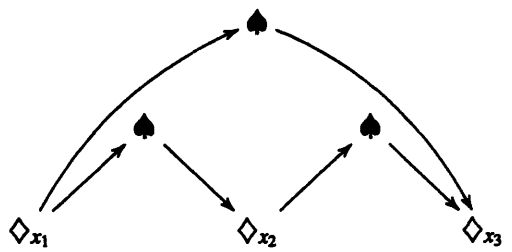

▲

4.6.4 General construction of witness families

By Theorem 4.20, since $(T, \mathcal{M})$ is automorphic and not conform, it follows that its template T is not valid. Consider any homomorphism $c^T : T \longrightarrow T$, e.g. $c^T = id_T$ and set $(N, c^N) := (T, c^T)$. Then we have $N \xrightarrow{c^N} T$ and (N, c^N) not valid with respect to $(T, \mathcal{M})$. Let (G, c^G) be its opening as defined above in Subsection 4.6.2. Let $\mathcal{F}$ be the set of structures $I_{\bar{n}}$ in $STRUC(\sigma)$ for $\bar{n} = n_1, n_2, \ldots, n_q$, with $n_1 \geq p_1, n_2 \geq p_2, \ldots, n_q \geq p_q$ obtained from the gadget (G, c^G) as in Section 4.6.3.

case 1: $\boxed{\mathcal{F} \subseteq FP(T, \mathcal{M})}$
We prove that $\mathcal{F}$ is a witness family with respect to $(T, \mathcal{M})$.
Let B be some σ-structure. We may assume w.l.o.g. that for any A in $\mathcal{F}$, $A \longrightarrow B$. Let $\bar{n} = (n_1, n_2, \ldots, n_q)$, where $n_i > p_i.|B| - |B|$ for any $1 \leq i \leq q$. By assumption, we have $I_{\bar{n}} \xrightarrow{b} B$ for some b. By construction of $I_{\bar{n}}$ there must be a copy of the gadget G in $I_{\bar{n}}$ whose plug-points are all identified by b. Hence $N \xrightarrow{\tilde{b}} B$ for some $\tilde{b}$ induced by b and also $\tilde{b}(N) \notin FP(T, \mathcal{M})$. This proves the claim.

case 2: $\boxed{\mathcal{F} \ni I_{\bar{n}} \notin FP(T, \mathcal{M}), \text{ for some } \bar{n}.}$
Consider the coloured structure $(I_{\bar{n}}, c^{I_{\bar{n}}})$. Notice that the following holds:

- $I_{\bar{n}}$ is a no-instance; and
- $(I_{\bar{n}}, c^{I_{\bar{n}}})$ is not valid but $I_{\bar{n}} \xrightarrow{c^{I_{\bar{n}}}} T$.

We shall repeat the construction, deriving this time a gadget from $(I_{\bar{n}}, c^{I_{\bar{n}}})$. However, we choose with great care the elements at which we open-up: they shall always be special elements of $I_{\bar{n}}$ (as defined in Subsection 4.6.3).

Recall that the only forbidden patterns occurring in $(I_{\bar{n}}, c^{I_{\bar{n}}})$ are biconnected. Moreover, by construction such an occurrence of a biconnected forbidden pattern must involve at least two copies G_1 and G_2 of the gadget. Let x be a special vertex common to G_1, G_2 and to that occurrence. Now, add a new vertex x' and replace every occurrence of the vertex x in every tuple of G_2 by this new vertex x'. Proceed similarly for every occurrence of a forbidden pattern. We call the structure obtained G'. By construction G' is a yes-instance of $FP(T, \mathcal{M})$; consider $c^{G'}$ to be the valid colouring of G' induced by $c^{I_{\bar{n}}}$ and defined as follows. Every vertex occurring in $I_{\bar{n}}$ is coloured according to $c^{I_{\bar{n}}}$ and any new vertex x' takes the same colour as its corresponding vertex x via $c^{I_{\bar{n}}}$. Let $\mathcal{F}'$ be the family of structures obtained from the new gadget G' (the plug-points being the special vertices x at which we opened-up and their copies x' in G'). If $\mathcal{F}' \subseteq FP(T,\mathcal{M})$ then we are back to the first case and we have constructed a witness family. Otherwise, we simply loop back to case 2.

Denote by G^k the gadget used at stage k and by $I^k_{\bar{n}}$ the structures build from G^k.

We claim that we eventually reach case 1. Consider for contradiction the sequence $(u_k)_{k\geq 0}$ defined as follows: u_0 is the minimal distance between any two plug-points of the gadget G (here by distance between two vertices we mean the length of the shortest path between those two vertices); u_k is defined to be the minimal distance between two plug-points of the gadget constructed at stage k. By construction, this sequence is non decreasing; that is, for any $k \geq 0$, we have $u_{k+1} \geq u_k$. Assume further that this sequence is not stationary (we shall prove this shortly). Let d be the size of the largest cycle that embeds into a non-conform forbidden pattern of $\mathcal{M}$. Let $k \geq 0$ such that $u_k > \frac{d}{2}$. By assumption for some $\bar{n}$ the structure $I^k_{\bar{n}}$ is a no-instance of $FP(T,\mathcal{M})$. Consider its canonical colouring $(I^k_{\bar{n}}, c^{I^k_{\bar{n}}})$. This colouring is valid for each copy of the gadget G^k by construction. It follows that some non-conform forbidden pattern must embed in more than one copy of G^k. However, this is not possible: it would imply that this forbidden pattern would contain a cycle of size greater or equal than $2.u_k$, that is strictly greater than d. This yields a contradiction. Therefore we proved the following: if the sequence $(u_k)_{k\geq 0}$ is not stationary then we eventually go out of the loop in case 2; that is, our construction terminates and we eventually obtain a witness family. We now prove that the sequence $(u_k)_{k\geq 0}$ is not stationary.

The sequence $(u_k)_{k\geq 0}$ is not stationary. Assume for contradiction that this is not the case; that is, that for some $k \geq 0$ and for any $k' \geq k$ we have $u_{k'} = u_k$. By construction, in G^{k+1} the distance between two plug-points of the same type (that is, two vertices that correspond to the same special vertex of $I^k_{\bar{n}}$) is greater or equal than $2.u_k$. However, since $u_k = u_{k+1}$, there must be two plug-points x and y at distance u_k in G^{k+1}. These two plug-points x and y must necessarily be incident to the same copy of G^k within G^{k+1}. This leads us to the following definitions.

For any copy G^k_i of G^k in $G^{k'}$ define $P(k,k')$ to be the set of pairs of plug-points of $G^{k'}$ incident to G^k_i that are at distance exactly u_k in G^k_i. For any copy G^k_i of G^k in $G^{k'}$, define $\mathtt{free}(k,k')$ to be the set of plug-points mentioned by the pairs of $P(k,k')$. Furthermore, fix some $x_{k,k'}$ in $\mathtt{free}(k,k')$.

We add another constraint to the construction in case 2: while opening forbidden patterns, for each copy of G^k_i in $G^{k'}$, **never** open-up at $x_{k,k'}$.

Note that the process of opening does not increase the number of new plug-points incident with any copy of G^k, and it does not reduce the distance between any pair of new plug-points incident with any copy of G^k. Hence, for any copy G^k_i of G^k in G^{k+1}, $\mathtt{free}(k,k+1) < \mathtt{free}(k,k)$. It follows that after finitely many steps, say, at step $k' > k$, we must have $u_{k'} > u_k$. This yields a contradiction. So, we have proved that the sequence $(u_k)_{k \geq 0}$ is not stationary.

To summer-up, we have provided a generic construction which allows us to build a witness family for any given non-conform normal representation.

4.7 Characterisation

In this Section, we state our main result, that is the exact characterisation of these forbidden patterns that are not in CSP. We first state this result in the case of connected representation, before illustrating it by some examples. Finally, we extend the result to any representation by generalising the notion of normal representation to disconnected representation; there we introduce the notion of *set of normal representations*.

4.7.1 Main result

The previous results leads to a full characterisation of connected representations with respect to the property of representing a CSP problem.

Theorem 4.26 *(théorème de Louison [5])*

Let $(\mu, \mathcal{M})$ be a connected σ-representation. The following are equivalent.

(i) $\mathbf{normal}(\mu, \mathcal{M})$ *is not conform*

(ii) $FP(\mu, \mathcal{M})$ *is not in CSP*

(iii) There exists a witness family for $(\mu, \mathcal{M})$

PROOF. It follows from the construction of the previous section that $(i) \Rightarrow (iii)$. $(iii) \Rightarrow (ii)$ by Lemma 4.25. Hence, it follows that $(i) \Rightarrow (ii)$. The converse holds since $\neg(i) \Rightarrow \neg(ii)$ by Proposition 4.13. Thus we have proved $(i) \iff (ii)$ and the other equivalences follow. □

4.7.2 Examples

We have seen earlier that numerous representations were normal and not conform, so as a corollary from our main result, we know that they are not in CSP.

Corollary 4.27 *Let* $n \geq 1$. *The forbidden patterns problem represented by* $\mathfrak{MDC}_n^2$ *is not a CSP. The forbidden patterns problem represented by* $\mathfrak{MC}_n^2$ *is not a CSP. Let* $p \geq 0$. *The forbidden patterns problem represented by* $\mathfrak{ADC}_{2p}^2 - \mathfrak{MC}$ *is not a CSP.*

Notice as well that all the problems introduced in Section 2.4 are proved to be not in CSP by hand of the main result, as we computed their normal form in Section 4.5.2 and none of them were conform.

Furthermore, notice that we have given only examples with directed graphs as they are easier examples but the main theorem holds for any signature.

4.7.3 The case of disconnected representation

We can extend the notion of Feder-Vardi transformation of a representation to the disconnected case; that is when a forbidden pattern is not connected. Let $(\mu, \mathcal{M})$ be a representation such that there exists a disconnected forbidden pattern $(F, c_\mu^F) \in \mathcal{M}$, that is F consists of the disjoint union of two structures F_0 and F_1. It is not difficult to see that $FP(\mu, \mathcal{M}) = FP(\mu, \mathcal{M}_0) \cup FP(\mu, \mathcal{M}_1)$ where $\mathcal{M}_i := (\mathcal{M} \setminus \{(F, c_\mu^F)\}) \cup \{(F_1, c_\mu^{F_i})\}$ with $c_\mu^{F_i} := c_\mu^F{}_{\restriction F_i}$.

[5] In the eventuality that the reader might want to refer to this result, please quote it as *le théorème de Louison*, as today I am the proud "republican godfather" of Louison.

So we extend the notion of normal representation to the disconnected case and consider the following recursive definition; the *set of normal representations* of a representation $(\mu, \mathcal{M})$ is

- the set containing $\mathbf{normal}(\mu, \mathcal{M})$ if $(\mu, \mathcal{M})$ is a connected representation; and
- the simplified union of the set of normal representations of $(\mu, \mathcal{M}_0)$ and $(\mu, \mathcal{M}_1)$ if $(\mu, \mathcal{M})$ is as above,

where by simplified union, we mean that we remove a representation whenever there exists a recolouring into another (analogous operation as when we dealt with forbidden patterns). We denote the set of normal representations of a representation $(\mu, \mathcal{M})$ by $\mathbf{Normal}(\mu, \mathcal{M})$.

We can extend our main result to disconnected instances.

Theorem 4.28 *Let $(\mu, \mathcal{M})$ be a σ-representation. The following are equivalent.*

(*i*) *The set* $\mathbf{Normal}(\mu, \mathcal{M})$ *contains a single conform connected representation.*

(*ii*) *$FP(\mu, \mathcal{M})$ is a CSP*

PROOF. The case when $\mathbf{Normal}(\mu, \mathcal{M})$ is a singleton was done previously; so, let $(\mu_0, \mathcal{M}_0), (\mu_1, \mathcal{M}_1) \in \mathbf{Normal}(\mu, \mathcal{M})$. Let T_0 and T_1 be their respective templates. We claim that T_0 is a no-instance of $(\mu_1, \mathcal{M}_1)$. Indeed, if T_0 were accepted then it would induce the existence of a recolouring of $(\mu_0, \mathcal{M}_0)$ to $(\mu_1, \mathcal{M}_1)$ which would contradict the definition of set of normal representations (the proof is very similar to the proof of Proposition 4.3.4). In the case where T_0 is a yes-instance of $(\mu_0, \mathcal{M}_0)$ then the latter is a conform representation and $FP(\mu_0, \mathcal{M}_0)$ a CSP. So assume further that not all the representations among the set of normal representations of $(\mu, \mathcal{M})$ are conform (we shall deal later with this case) and therefore without loss of generality that T_0 is not a yes-instance of $(\mu_0, \mathcal{M}_0)$. Hence, we have a structure that is a no-instance but can be coloured correctly on the edges with respect to the non-conform representation $(\mu_0, \mathcal{M}_0)$. So we can use it to build the gadget for the generic construction that lead to the main result and eventually obtain a witness family. Now if all the representations among the set of normal representations of $(\mu, \mathcal{M})$ are conform then we can see $(\mu, \mathcal{M})$ as the conjunction of CSP of respective templates $T_0, T_1, \ldots, T_n$. Those templates can not map into each other (otherwise this would lead to the existence of a recolouring). If $FP(\mu, \mathcal{M})$ were a CSP then let T be its template. Since $T_i \in FP(\mu, \mathcal{M})$, we would have $T_i \longrightarrow T$ thus the structure S consisting of disjoint copies of the T_i's would satisfy $S \longrightarrow T$ and thus $S \in FP(\mu, \mathcal{M})$. Hence there would be some T_j such that $S \longrightarrow T_j$ and finally we would have $T_i \longrightarrow T_j$ for some $i \neq j$, a contradiction. □

We conclude this chapter with a few remarks. First, notice that the normal form of a representation is quite complicated to compute as the reader may have noticed with the few simple examples provided. In order to implement efficiently an algorithm that would decide whether a forbidden patterns problem is a CSP, some simplifications are needed; representations should be given in a compact form as in Section 4.4. Moreover, notice that we decided to work with coloured structures to simplify the proofs but the same work could be achieved with partially coloured structures. Furthermore on this matter, we enforced the following order when computing the normal representation; first enforcing the representation to be canonical (which involves taking homomorphic images of forbidden patterns, which increases the size of the representation) then applying a canonical Feder-Vardi transformation (which involves adding more colours, thus also increasing the size) and finally taking a particular core (which decreases the size). Notice moreover that the last transformation is the most complicated, as it is clearly NP-hard. Hence, it would be probably more efficient to take the core of the representation as often as possible. Notice however, that since we want an automorphic representation on the end, we must take the core before finishing, as it might be the case that the Feder-Vardi transformation of an automorphic representation is not automorphic. It would be interesting to study in more details the rewriting system associated with the three transformations mentioned above. It is not clear whether it is confluent. In other words, the normal representation might not be definable as the unique rewrite of this system.

Our second remark concerns the gadget used for the construction of witness families. A part of the proof is quite complicated because of the fact that we might deal with a gadget that has many legs of possibly different types. However, for every examples that we investigated on graphs, we were able to build a simple bipede gadget as in the example above in Subsection 4.6.3. If we could prove that such a simple bipede gadget exists for any representation, we could simplify further our proof by using Erdös' theorem.

Finally, notice that recolourings alone do not provide a satisfactory morphism for representation as the converse of Proposition 4.1 does not hold. We shall discuss this issue in more details in the next chapter in Subsection 5.3.2.

In the next chapter, we relate also our main result with some results by Tardif and Nešetřil and we shall investigate the structure of the category of representations.

In Chapter 6, we shall give some examples of complete forbidden patterns problems that are not in CSP for the complexity class NL,P and NP. We shall also investigate some restrictions that lead to tractability of forbidden patterns problems.

Chapter 5

Heyting algebras, density and duality

We relate the results of Tardif and Nešetřil on duality and density to forbidden patterns problems. We introduce the Heyting algebra of cores of structures. We prove Tardif and Nešetřil correspondence between duality and density in the general case of a Heyting algebra. Finally, we show that the cores of representations form a Heyting algebra with respect to recolourings.

In this chapter, we shall investigate in more detail the algebraic properties of the category of σ-structures. We shall see that the cores form a *Heyting algebra*: that is, a distributive lattice with *exponential*. This algebraic machinery allowed Tardif and Nešetřil to relate in a recent work (see [45]) the notion of a *gap* (*i.e.* a place where the order fails to be dense) in this lattice (of cores) with the notion of a *duality pair*. These *duality pairs* correspond in fact to particular forbidden patterns problems that are in CSP. These problems are rather simple: they can be given by a representation with a single colour and only one forbidden pattern. Let us call them *monochrome forbidden pattern problems*. Tardif and Nešetřil have characterised duality pairs (in fact, they characterised gaps, but obtain a characterisation via this correspondence). Our main result from the previous chapter provides an *alternative characterisation of duality pairs*. It is interesting to notice that whereas their characterisation (of duality pairs) is much simpler than ours, our construction for the template (whenever possible) seems much simpler than theirs.

Moreover, we generalise the algebraic machinery mentioned above to representations. Thus, we relate the notion of a *gap* in the category of representations with the notion of a *duality pair* in this category: indeed, we are able to prove the correspondence in the more general setting of *Heyting algebras*.

We defined all the concepts from category theory that we need in Appendix B. For more on category theory, we refer to [38], and for universal algebra and lattice theory, we refer to [44].

5.1 Heyting algebras

In this section, we shall recall the definition and some basic facts about *Heyting algebras*. In a second part we show that the cores form a Heyting algebra.

5.1.1 Definition

A *Heyting Algebra* is a structure over the signature λ_h consisting of three binary function symbols $\wedge, \vee$ and $\Rightarrow$, and of two constant symbols $\mathtt{0}$ and $\mathtt{1}$; this structure is a *lattice* with respect to $\wedge$ and $\vee$ with least element $\mathtt{0}$ and greatest element $\mathtt{1}$, *i.e.* it satisfies the following identities

$$\begin{aligned} x\wedge y &= y\wedge x & x\vee y &= y\vee x \\ x\wedge(y\wedge z) &= (x\wedge y)\wedge z & x\vee(y\vee z) &= (x\vee y)\vee z \\ x\wedge x &= x & x\vee x &= x \\ x\wedge(x\vee y) &= x & x\vee(x\wedge y) &= x \\ x\wedge 0 &= 0 & x\vee 1 &= 1 \end{aligned}$$

We define the partial order $\leq$ that corresponds to this lattice as usual; that is, we set $x \leq y$ if, and only if, $x \wedge y = x$. A further property of these algebras is that, for any x, y and z,

$$z \leq x \Rightarrow y, \text{ if, and only if, } z \wedge x \leq y.$$

5.1.2 The Heyting algebra of cores

The fact that the cores form a Heyting algebra and the existence of the exponential plays an important role in graph theory. It is not quite clear who exactly made this discovery first. It seemed to have been a well known fact in some research groups for a few decades. There is a note about this in case the reader is interested in [50], a survey on Hedetniemi's conjecture, by Norbert Sauer, which we suggest also as it contains further examples of the use of exponentials in graph theory.

Let's consider the quasi-order given by homomorphisms between σ-structures up to homomorphism equivalence: two structures A and B are homomorphically equivalent (denoted by $A \sim B$) whenever $A \longrightarrow B$ and $B \longrightarrow A$. Hence when we factor out $STRUC(\sigma)$ by $\sim$ we obtain a partial order. As representatives for each equivalence class, one can consider cores as we have seen earlier in Proposition 4.7, *i.e.*

$$\frac{\langle STRUC(\sigma), \longrightarrow \rangle}{\sim} \approx \langle CORE(\sigma), \longrightarrow \rangle$$

where $CORE(\sigma)$ denotes the class of cores of σ-structures considered up to isomorphism, that is according to the notation of the previous chapter,

$$CORE(\sigma) := \bigcup_{A \in STRUC(\sigma)} \mathbf{core}(A).$$

In fact, there is a much richer structure than just a partial order. Indeed, one can define the *product* and the *coproduct* of structures with respect to homomorphisms, which lead themselves to the notion of *supremum* and *infimum* for cores. Hence the partial order $\langle CORE(\sigma), \longrightarrow \rangle$ is in fact a *lattice*.

Lemma 5.1 *The category of σ-structures has products and coproducts.*

PROOF. For any given pair of σ-structures (A,B), define the [1] *product* $A \times B$ of A and B as follows.

- $|A \times B| := |A| \times |B|$ (Cartesian product of the two sets); and
- for any r-ary symbol R in σ, and any r-tuple $((a_1,b_1),(a_2,b_2),\ldots,(a_r,b_r))$ of elements of $|A \times B|$, $R((a_1,b_1),(a_2,b_2),\ldots,(a_r,b_r))$ holds in $A \times B$ if, and only if, $R(a_1,a_2,\ldots,a_r)$ holds in A and $R(b_1,b_2,\ldots,b_r)$ holds in B.

We can also define the [1] *coproduct* of A and B denoted by $A+B$ to be simply the structure consisting of the disjoint union of the two structures, that is,

- $|A+B| = A\dot{\cup}B$; and
- for any r-ary symbol R and any r-tuple $(x_1,x_2,\ldots,x_r)$ of elements of $|A+B|$, $R(x_1,x_2,\ldots,x_r)$ holds in $A+B$ if, and only if, $R(x_1,x_2,\ldots,x_r)$ holds either in A or in B.

It is a straightforward exercise to check that these definitions satisfy indeed the defining properties of the product and coproduct; in other words, that for any triple of σ-structures (A,B,C),

- $C \longrightarrow A \times B$ if, and only if, $C \longrightarrow A$ and $C \longrightarrow B$; and
- $A+B \longrightarrow C$ if, and only if, $A \longrightarrow C$ and $B \longrightarrow C$.

□

For any two cores A and B, we set

- $A \wedge B := \mathbf{core}(A \times B)$; and
- $A \vee B := \mathbf{core}(A + B)$.

The following result follows directly from the previous proposition.

Corollary 5.2 *$\langle CORE(\sigma), \wedge, \vee \rangle$ is a lattice.*

Furthermore, this category has *exponentials* (in a lattice, an exponential corresponds to a pseudo-complement; and, in the category of sets, an exponential is simply the set of functions of one set to another).

Lemma 5.3 *The category of σ-structures has exponentials.*

PROOF. For any pair of σ-structures (A, B) we define A^B, as follows.

- $|A^B| := |A|^{|B|}$ (the set of functions of $|B|$ to $|A|$); and
- for any r-ary symbol R and any r functions $f_1, f_2, \ldots, f_r$ of $|B|$ to $|A|$, $R(f_1, f_2, \ldots, f_r)$ holds in A^B if, and only if, for any r-tuple $(b_1, b_2, \ldots, b_r)$ of elements of B, if $R(b_1, b_2, \ldots, b_r)$ holds in B then $R(f_1(b_1), f_2(b_2), \ldots, f(b_r))$ holds in A.

It can be easily checked that A^B satisfies the defining property of the exponential, that is,

$$\text{for any } C \text{ in } STRUC(\sigma), B \times C \longrightarrow A \text{ if, and only if, } C \longrightarrow A^B.$$

□

It follows from the existence of exponentials that the product and the coproduct are distributive with respect to each other: that is, the following *distributive laws* hold.

$$A \times (B + C) \approx (A \times B) + (A \times C) \quad \text{and} \quad A + (B \times C) \approx (A + B) \times (A + C).$$

[1]Note that these notions are defined up to isomorphisms as usual in category theory, in the following we shall feel free to define every categorical notion as such without further warnings.

Moreover, the category of σ-structures has an initial object (a structure that maps into every structure via a single homomorphism) as well as a terminal object (the dual notion; that is, a structure into which every structure maps via a single homomorphism): namely, the structures 0_σ and 1_σ defined as follows,

$$|0_\sigma| := \emptyset, \text{ and for each symbol } R \text{ in } \sigma,\ R^{0_\sigma} := \emptyset;$$
$$|1_\sigma| := \{0\}, \text{ and for each symbol } R \text{ in } \sigma,\ R^{1_\sigma} := \{(0,0,\ldots,0)\}.$$

Hence, together with Corollary 5.2 and Lemma 5.3, this leads to the following result (the notion of a *topos* is defined in Appendix B).

Theorem 5.4 *The category of σ-structures is a* topos.

PROOF.

(i) We prove that $STRUC(\sigma)$ has equalizers. Let B and A be two structures and $B \xrightarrow{f} A$ and $B \xrightarrow{g} A$ be two homomorphisms. Let D be the substructure of B induced by the set:

$$\{x \in |B| \text{ such that } f(x) = g(x)\}$$

and e be the induced embedding $D \overset{e}{\hookrightarrow} B$. By construction, we have $f \circ e = g \circ e$. It remains to show the universality. Let C be a structure and $C \xrightarrow{h} B$ a homomorphism such that $f \circ h = g \circ h$. It follows directly that the image of $|C|$ via h is included in $|D|$. So define $C \overset{h'}{\hookrightarrow} D$ by $h' := e^{-1} \circ h$. Clearly $e \circ h' = h$ and h' is unique.
We have also proved that the category of σ-structures has a terminal object, and that it has products: it follows by Corollary B.1 that $STRUC(\sigma)$ has finite limits.

(ii) Let 2_σ be the disjoint union of two copies of 1_σ. For the subobject classifier, consider the structure 2_σ.

(iii) We have products and exponentials so the category of σ-structures is cartesian closed. □

Notice that 0_σ and 1_σ are cores. Moreover for two cores A and B we set,

$$B \Rightarrow A := \mathbf{core}(A^B).$$

The previous theorem yields the following result.

Corollary 5.5 *$\langle CORE(\sigma), \wedge, \vee, \Rightarrow, 0_\sigma, 1_\sigma\rangle$ is a Heyting algebra.*

Let L be a lattice. Recall that a lattice element a is said to be *(join) prime* if, and only if, for any lattice elements b and c, if $a = b \vee c$ then $a = b$ or $a = c$. In the following, we shall simply write prime for join prime. It can be checked that the prime elements in the lattice of cores are exactly the *connected cores*.

5.2 Duality and density

In this section, we shall investigate the *correspondence between duality and density*. First, we shall define *duality pairs* and relate them to some particular problems; the *monochrome forbidden pattern problems* that are conform. We then derive from the main result of the previous chapter an alternative characterisation of duality pairs, which together with Tardif and Nešetřil's own characterisation, provides a better characterisation of monochrome forbidden pattern problems (as to whether such a problem is in CSP or not). Next, we shall briefly discuss the proof of their result and contrast their better characterisation in the restricted case of monochrome forbidden pattern problems with the superiority of our construction for templates (whenever the problem considered is in CSP) over theirs. Finally, we generalise their proof of the correspondence between duality pairs and gaps in the lattice of cores; we prove such a correspondence for any Heyting algebra.

5.2.1 Duality pairs and monochrome forbidden pattern problems

Let A and B be cores. Notice that the homomorphism problem with template B corresponds to a *principal ideal* in the lattice of cores: namely, the set,

$$\{C \in CORE(\sigma) | C \longrightarrow B\}.$$

Consider now the dual notion for A; that is, the complement of the *principal filter* generated by A: namely, the set,

$$\{C \in CORE(\sigma) | A \not\longrightarrow C\}.$$

This remark leads to the following question: for which structures A and B do these two notions coincide? This yields the following definition. Let A and B be σ-structures. We call (A, B) a *duality pair* if, and only if, the principal ideal generated by $\mathbf{core}(B)$ coincides with the complement of the principal filter generated by $\mathbf{core}(A)$. Notice that the complement of the principal filter generated by $\mathbf{core}(A)$ corresponds to a *monochrome forbidden pattern problem*; that is, a problem with a single colour and a single forbidden pattern (the structure A coloured uniformly with this unique colour). For simplicity, we denote this problem by $FP(A)$ (to be coherent with our notation, we should write $(\mathbf{1}, \{(A, c_{\mathbf{1}}^{A})\})$ instead of A). Such problems correspond to first-order MMSNP sentences with only one negated conjunct and are therefore computationally trivial to solve (within the complexity class L).

Notice that in our settings (A,B) is a duality pair if, and only if, $FP(A) = CSP(B)$. Therefore, the following follows from Theorem 4.28.

Corollary 5.6 *Let A be a structure. There exists a structure B such that (A,B) is a duality pair if, and only if,* $\mathbf{Normal}(\mathbf{1}, \{(A, c_{\mathbf{1}}^{A})\})$ *consists of a single conform representation whose template is homomorphically equivalent to B.*

Another characterisation has been however obtained by Tardif and Nešetřil in [45]; we shall discuss their proof in the next section. In order to state it, we need the following definition. We say that a structure A is a *tree* if, and only if, it is connected and cycle-free (i.e. it has no substructure that is a cycle).

Theorem 5.7 *(Tardif, Nešetřil) Let A be a structure. There exists a structure B such that (A,B) is a duality pair if, and only if,* $\mathbf{core}(A)$ *is a tree.*

One can therefore combine these two results together as follows.

Lemma 5.8 *Let A be a structure.* $\mathbf{Normal}(\mathbf{1}, \{(A, c_{\mathbf{1}}^{A})\})$ *consists of a single conform representation if, and only if,* $\mathbf{core}(A)$ *is a tree.*

This provides therefore a better characterisation for monochrome forbidden pattern problems.

Corollary 5.9 *The problem $FP(A)$ is in CSP if, and only if,* $\mathbf{core}(A)$ *is a tree.*

Notice that in case we would want to prove the above lemma without using Tardif and Nešetřil's characterisation, the indirect implication is clear; if A is a tree then the representation $(\mathbf{1}, \{(A, c_{\mathbf{1}}^{A})\})$ can be broken down by a sequence of elementary Feder-Vardi transformations until there are only conform forbidden patterns remaining (cf. remark in the next subsection). However, the converse implication does not seem to be quite as trivial.

In order to discuss the proof of Tardif and Nešetřil's theorem, we need the following definition. Let A and B be two σ-structures. (A,B) is said to be a *gap pair* if, and only if, $A \longrightarrow B$, $B \not\longrightarrow A$ and for every σ-structure C, if $A \longrightarrow C \longrightarrow B$ then either $A \sim C$ or $C \sim B$. Notice that a gap pair (A,B) simply corresponds to an interval $[\mathbf{core}(A), \mathbf{core}(B)]$ in the lattice of cores that is not *dense*: that is, there is no core C apart from $\mathbf{core}(A)$ and $\mathbf{core}(B)$ such that $\mathbf{core}(A) \leq C \leq \mathbf{core}(B)$. In other words, $\mathbf{core}(B)$ is the *upper cover* of $\mathbf{core}(A)$, which we denote by $\mathbf{core}(A) \prec \mathbf{core}(B)$.

5.2.2 Discussion of Tardif and Nešetřil's proof

Tardif and Nešetřil used the correspondence between gap pairs and duality pairs: as a matter of fact, this correspondence exists because the cores form a Heyting algebra. We shall prove this in the next subsection.

The notion of a duality pair was introduced by Tardif and Nešetřil in an attempt to investigate *good characterisations* of homomorphism problems; that is, to find *obstructing sets*; *e.g.* the set of odd cycles is such an obstructing set in the case of the problem 2-Col. Therefore they looked at the most simple such good characterisation: the case of an obstructing set reduced to a singleton. Hence, the notion of duality pair. It is important to note that since they did not really perceive the problem as a forbidden pattern problem, they did not use colours and did not use a tool like the Feder-Vardi transformation. Their proof relies on the correspondence mentioned earlier: first, gaps are characterised, and therefore so are duality pairs. To characterise gaps, there are two parts: the "positive part" in which they construct what they call *the gap below a tree* and the "negative part" in which they prove that there is no gap below a non-tree.

The first part corresponds, modulo the correspondance, to the construction of a template from the normal form of a conform representation; and, is rather different in its philosophy: Tardif and Nešetřil use a construction called *the arrow construction*. This construction involves the partial order over the subtrees of a given core tree A and the induced notion of a-ideal for some element a of A. For a given core tree A, the arrow construction yields a structure $A^{\downarrow}$ (which is not necessarily a core) such that $\mathbf{core}(A^{\downarrow}) \prec A$. Then, by way of the correspondence between density and duality (*cf.* Lemma 5.11 in the next subsection), Tardif and Nešetřil prove that $(A, (A^{\downarrow})^{A})$ is a duality pair. Hence, for a given core tree A, to construct the template B of the problem $FP(A)$ with their method seems rather difficult (as they point out themselves). Indeed, the arrow construction is already quite intricate and $A^{\downarrow}$ has a size that is exponential in the size of A. Hence, to compose the arrow construction by taking its Ath exponent is doubly exponential! However, our method can be adapted in the case of a tree. Indeed, we do not need to take any homomorphic images of A: a sequence of elementary Feder-Vardi transformations decomposes A into its biconnected components (*i.e.* conform forbidden patterns since A is a tree), and such homomorphic images would be discarded after the canonical Feder-Vardi transformation as they would not be properly coloured according to the new template. Therefore, we obtain a conform representation by applying the canonical Feder-Vardi transformation. Furthermore, we could leave the representation in its compact form. Hence we obtain a description of a structure that is homomorphically equivalent to $(A^{\downarrow})^{A}$, that would be more manageable (we get rid of one exponential that way).

The second part of their proof is quite similar to ours and relies on the same ideas: opening up a non-conform biconnected structure and construct a large structure with this gadget (they take a suitable graph of large girth and high chromatic number, that exists according to a theorem of Erdös, and replace its edges by the gadget). Since they deal with problems of the form $FP(A)$ (where A is a core that is not a tree) they derive a gadget by opening up A (they do not have to deal with the problem of having different colours). Hence, given some B such that $A \not\longrightarrow B$, they produce a structure C such that $A \not\longrightarrow C$ and $C \not\longrightarrow B$ but $C \longrightarrow A$. Thus, the structure $C+B$ is strictly in between A and B, whenever $B \longrightarrow A$. So, for any structure B, (A,B) is not a gap pair.

To conclude on this matter, it seems that combining the two approaches might be quite enriching: the correspondence between duality and density that we extend in the next subsection is a beautiful and useful result (it provides counter examples). However, the approach via representations and computations of a normal form seems to be better when it comes to prove positive results. Indeed, it seems rather hard to picture the exponential of two structures, and this even in simple cases: there are very few general internal descriptions of exponential of graphs known presently ([51]), not to mention the combination of this construction with the intricate arrow construction.

5.2.3 Correspondance between duality and density

In this section we present the correspondence between duality and density that was investigated by Tardif and Nešetřil in [45]. Since we need the same result later for representations, we prove this result in the general setting of Heyting algebra. In the following H denotes such an algebra. Note that the original proof was done in the category of σ-structures rather than in the Heyting algebra of cores (which tends to simplify things a great deal in the proof).

Lemma 5.10 *If (a,b) is a duality pair in H then a is a prime and $(a\wedge b,a)$ is a gap pair.*

PROOF. Assume for contradiction that a is not a prime: that is, there exists some elements a_1 and a_2 such that $a = a_1 \vee a_2$ and $a \neq a_1$ and $a \neq a_2$. It follows that $a \not\leq a_1$ and $a \not\leq a_2$. Since (a,b) is a duality pair, the above yields to the following: $a_1 \leq b$ and $a_2 \leq b$. It follows therefore that $a = a_1 \vee a_2 \leq b$. From $a \leq b$, since (a,b) is a duality pair, we get the following contradiction $a \not\leq a$.

We have $a \wedge b \leq a$. Let c be an element of H such that $a\wedge b \leq c \lneq a$. Since (a,b) is a duality pair and $c \not\geq a$, it follows that $c \leq b$. Hence, we have $c = a\wedge b$. Thus, we have proved that $a \wedge b \prec a$.

□

Lemma 5.11 *If (a,b) is a gap pair in H and b a prime then $(b, b \Rightarrow a)$ is a duality pair.*

PROOF. For any element c of H, we have $a \leq a\vee(b\wedge c) \leq b$. Since $a \prec b$, we have two cases to consider.

1. $a = a\vee(b\wedge c)$: It follows that $b\wedge c \leq a$. Thus, by definition of the exponential it implies that $c \leq b \Rightarrow a$.

2. $b = a\vee(b\wedge c)$: since b is prime and by assumption $a \neq b$, it follows that $b = b\wedge c$ and finally that $b \leq c$.

Thus, we have proved that for any c of H, either $c \leq b \Rightarrow a$ or $b \leq c$: that is, $(b, b \Rightarrow a)$ is a duality pair. □

We now prove that there is a one-to-one correspondence between gap pairs (c,d) where d is a prime and duality pairs.

If we start with a duality pair (a,b) then it follows from Lemma 5.10 that $(a \wedge b, a)$ is a gap pair and a a prime. Hence, it follows from Lemma 5.11 that $(a, a \Rightarrow (a \wedge b))$ is a duality pair. Since (a,b) and $(a, a \Rightarrow (a \wedge b))$ are duality pairs, it follows that $b = a \Rightarrow (a \wedge b)$.

Conversely, let (c,d) be a gap pair with d a prime. Then, by Lemma 5.11, it follows that $(d, d \Rightarrow c)$ is a duality pair. Finally, by Lemma 5.10, it follows that $(d \wedge (d \Rightarrow c), d)$ is a gap pair. We have $c \wedge d = c$. So, in particular, we have $c \leq d \Rightarrow c$ and since $c \leq d$ it follows that $c \leq d \wedge (d \Rightarrow c)$. We also have $c \wedge d \leq c$ hence $d \wedge (c \wedge d) \leq c$. The defining property of the exponential implies that $d \leq (c \wedge d) \Rightarrow c$. But since $(c \wedge d) \Rightarrow c = (c \Rightarrow d) \Rightarrow c$, via the defining property of the exponential we get $d \wedge (c \Rightarrow d) \leq c$. Hence, we get back to the gap pair (c,d) we started with.

5.3 More on representations

We shall first prove that the category of representations is a topos. This yields that normal representations (considered up to iso-recolourings) form a Heyting Algebra. Finally, we discuss the containment problem for forbidden patterns problems.

5.3.1 The topos of representations

In the following, we denote by $REP(\sigma)$ the category of σ-representations: that is, the category whose objects are σ-representations; and, whose morphisms are recolourings. We prove that the category of representations is a topos: indeed, we proved in the previous chapter that a recolouring is a generalisation of a homomorphism; in the same sense, the product, coproduct and exponential of structures can be generalised to representations.

Product of representations. Let $(\mu, \mathcal{M})$ and $(\nu, \mathcal{N})$ be σ-representations. Define $(\mu, \mathcal{M}) \times (\nu, \mathcal{N})$ to be the representation with:

- colours $\mu \times \nu$ (the Cartesian product of the colour set); and

- forbidden patterns

$$\{(F, c^F_{\mu\times\nu}) \in \mathrm{STRUC}_{\mu\times\nu}(\sigma) \mid (F, \pi_\mu \circ c^F_{\mu\times\nu}) \in \mathcal{M}\,\mathbf{or}\,(F, \pi_\nu \circ c^F_{\mu\times\nu}) \in \mathcal{N}\},$$

where π_μ and π_ν are the left and right projections, respectively.

Notice that the "and" of the definition of a product for structures becomes an "or" for representations: intuitively, this is due to the fact that a forbidden pattern is a generalisation of a "no-tuple" in a structure.

Lemma 5.12 *The notion defined above truly is the product in the category $REP(\sigma)$.*

PROOF. Let $(\lambda, \mathcal{L})$ be a representation. Assume that

$$(\lambda, \mathcal{L}) \xrightarrow{r} (\mu, \mathcal{M}) \times (\nu, \mathcal{N}).$$

It follows directly from the above definition and the definition of a recolouring that:

$$(\mu, \mathcal{M}) \times (\nu, \mathcal{N}) \xrightarrow{\pi_\mu} (\mu, \mathcal{M}) \text{ and } (\mu, \mathcal{M}) \times (\nu, \mathcal{N}) \xrightarrow{\pi_\nu} (\nu, \mathcal{N}).$$

Hence, by composition, it follows that:

$$(\lambda, \mathcal{L}) \xrightarrow{\pi_\mu \circ r} (\mu, \mathcal{M}) \text{ and } (\lambda, \mathcal{L}) \xrightarrow{\pi_\nu \circ r} (\nu, \mathcal{N}).$$

Conversely, assume that

$$(\lambda, \mathcal{L}) \xrightarrow{r_\mu} (\mu, \mathcal{M}) \text{ and } (\lambda, \mathcal{L}) \xrightarrow{r_\nu} (\nu, \mathcal{N}).$$

Set $r := (r_\mu, r_\nu)$. Let $(F, r \circ c^F_\lambda)$ be a forbidden pattern of $(\mu, \mathcal{M}) \times (\nu, \mathcal{N})$. We may assume w.l.o.g. that $(F, \pi_\mu \circ r \circ c^F_\lambda) \in \mathcal{M}$. Thus, since $\pi_\mu \circ r = r_\mu$ is a recolouring, it follows that (F, c^F_λ) is not valid for $(\lambda, \mathcal{L})$. So, we have proved that r is a recolouring. □

The following can be easily checked.

representation	template
$(\mu, \mathcal{M})$	M
$(\nu, \mathcal{N})$	N
$(\mu, \mathcal{M}) \times (\nu, \mathcal{N})$	$M \times N$

We discussed in Subsection 4.3.3 an alternative definition of representation, the so-called canonical representation: that is, a representation given as $(M, \mathcal{M})$; where M is a σ-structure M (corresponding to the template of a standard representation); and, where any forbidden pattern $(F, c_\mu^F) \in \mathcal{M}$ satisfies $F \xrightarrow{c_\mu^F} M$ (c_μ^F is a *colouring* in the same sense as in the H-coloring problem). The following is straightforward: for a pair of canonical representations $(M, \mathcal{M})$ and $(N, \mathcal{N})$, the product $(M, \mathcal{M}) \times (N, \mathcal{N})$ is the canonical representation with:

- template $M \times N$; and
- forbidden patterns $F \xrightarrow{c^F} M \times N$, whenever either $F \xrightarrow{\pi_M \circ c^F} M$ belongs to $\mathcal{M}$ or $F \xrightarrow{\pi_N \circ c^F} N$ belongs to $\mathcal{N}$.

Notice that, we can identify a σ-structure M with the canonical representation $(M, \emptyset)$. In that sense, the product of representations generalises the product of structures.

Coproduct of representations Define $(\mu, \mathcal{M}) + (\nu, \mathcal{N})$ to be the representation with:

- colours $\mu \dot{\cup} \nu$ (the disjoint union of the colour sets); and
- forbidden patterns
  1. for every (χ_m, χ_n) in $\mu \times \nu$, the forbidden pattern $(F, c_{\mu\dot{\cup}\nu}^F) \in \mathrm{STRUC}_{\mu\dot{\cup}\nu}(\sigma)$ that consists of two distinct elements x and y and void relations such that $c^F(x) = \chi_m$ and $c^F(y) = \chi_n$;
  2. $\{(F, c_{\mu\dot{\cup}\nu}^F) \in \mathrm{STRUC}_{\mu\dot{\cup}\nu}(\sigma)$ such that $(F, c_{\mu\dot{\cup}\nu}^F) \in \mathcal{M}\}$; **and**
  3. $\{(F, c_{\mu\dot{\cup}\nu}^F) \in \mathrm{STRUC}_{\mu\dot{\cup}\nu}(\sigma)$ such that $(F, c_{\mu\dot{\cup}\nu}^F) \in \mathcal{N}\}$.

Notice that this time the "or" of the definition of a coproduct for structures becomes an "and" for representations.

Lemma 5.13 *The notion defined above is really the coproduct in the category $REP(\sigma)$.*

PROOF. Let $(\lambda, \mathcal{L})$ be a representation. Assume moreover that

$$(\mu, \mathcal{M}) + (\nu, \mathcal{N}) \xrightarrow{r} (\lambda, \mathcal{L}).$$

Since by construction, $(\mu, \mathcal{M})$ and $(\nu, \mathcal{N})$ are subrepresentations of the representation $(\mu, \mathcal{M}) + (\nu, \mathcal{N})$ via the injections ι_μ and ι_ν, that is

$$(\mu, \mathcal{M}) \xrightarrow{\iota_\mu} (\mu, \mathcal{M}) + (\nu, \mathcal{N}) \text{ and } (\nu, \mathcal{N}) \xrightarrow{\iota_\nu} (\mu, \mathcal{M}) + (\nu, \mathcal{N}),$$

by composition it follows that

$$(\mu, \mathcal{M}) \xrightarrow{r \circ \iota_\mu} (\lambda, \mathcal{L}) \text{ and } (\nu, \mathcal{N}) \xrightarrow{r \circ \iota_\nu} (\lambda, \mathcal{L}).$$

Conversely, assume that

$$(\mu, \mathcal{M}) \xrightarrow{r_\mu} (\lambda, \mathcal{L}) \text{ and } (\nu, \mathcal{N}) \xrightarrow{r_\nu} (\lambda, \mathcal{L}).$$

$$\begin{aligned} \text{Set } r : \mu \dot{\cup} \nu &\to \lambda \\ \chi &\mapsto \begin{cases} r_\mu(\chi), & \text{if } \chi \in \mu; \text{ and} \\ r_\nu(\chi), & \text{otherwise.} \end{cases} \end{aligned}$$

We now prove that r is a recolouring. Let $(F, r \circ c^F_{\mu \dot{\cup} \nu}) \in \mathcal{L}$. There are different cases to consider.

1. $c^F_{\mu \dot{\cup} \nu}$ ranges over both μ and ν: that is, there exists some vertex $x \in |F|$ (respectively, $y \in |F|$) and some colour χ_m in μ (respectively, χ_n in ν) such that $c^F_{\mu \dot{\cup} \nu}(x) = \chi_m$ (respectively, $c^F_{\mu \dot{\cup} \nu}(y) = \chi_n$). Hence, $(F, c^F_{\mu \dot{\cup} \nu})$ is not valid for the coproduct (because of the special forbidden patterns consisting of two vertices; one coloured in χ_m; and, the other in χ_n).

2. $c^F_{\mu \dot{\cup} \nu}$ ranges over μ only: we have $r \circ c^F_{\mu \dot{\cup} \nu} = r_\mu \circ c^F_{\mu \dot{\cup} \nu}$, and r_μ being a recolouring it follows that $(F, c^F_{\mu \dot{\cup} \nu})$ is not valid for $(\mu, \mathcal{M})$. Hence there exists

some $(G, c_\mu^G) \in \mathcal{M}$ and some coloured homomorphism g

$$(G, c_\mu^G) \xrightarrow{g} (F, c_{\mu\dot{\cup}\nu}^F).$$

By definition of the coproduct, it follows that (G, c_μ^G) is a forbidden pattern of the coproduct, hence that $(F, c_{\mu\dot{\cup}\nu}^F)$ is not valid for the coproduct.

3. $c_{\mu\dot{\cup}\nu}^F$ ranges over ν only: case similar to the previous one.

□

This construction does not exactly generalise the coproduct of σ-structure. However, if we restrict ourselves to connected and non-sbavate representations then we could amend our construction as follows. Replace the first type of forbidden pattern (those that forbid the simultaneous use of a colour of μ and a colour of ν) by

1′. for any r-ary relation symbol R in σ, for any choice of colours $\chi_1, \chi_2, \dots, \chi_r$ such that there exist $1 \le m, n \le r$ where $m \neq n$, $\chi_m \in \mu$ and $\chi_n \in \nu$, the forbidden pattern $(R(x_1, x_2, \dots, x_r), c_{\mu\dot{\cup}\nu})$, where

$$\begin{aligned} c_{\mu\dot{\cup}\nu} : \{\bar{x}\} &\rightarrow \mu\dot{\cup}\nu \\ x_i &\mapsto \chi_i \end{aligned}$$

Then, the following can be checked.

representation	template
$(\mu, \mathcal{M})$	M
$(\nu, \mathcal{N})$	N
$(\mu, \mathcal{M}) + (\nu, \mathcal{N})$	$M + N$

Exponential of representations Define the representation $(\mu, \mathcal{M})^{(\nu, \mathcal{N})}$ to be the representation with

- colours μ^ν (the set of functions of ν to μ); and

- forbidden patterns all the $(F, c^F_{\mu^\nu}) \in \mathrm{STRUC}_{\mu^\nu}(\sigma)$ such that there exists some $(F, c^F_\mu) \in \mathcal{M}$ and some mapping c^F_ν such that $c^F_\mu = c^F_{\mu^\nu} \otimes c^F_\nu$ and (F, c^F_ν) is valid for $(\nu, \mathcal{N})$, where

$$\begin{aligned} c^F_{\mu^\nu} \otimes c^F_\nu : |F| &\to \mu \\ x &\mapsto (c^F_{\mu^\nu}(x))(c^F_\nu(x)) \end{aligned}$$

The colour set of $(\mu, \mathcal{M})^{(\nu,\mathcal{N})}$ is μ^ν; hence, the colour $(c^F_{\mu^\nu}(x))$ of a vertex x of a forbidden pattern $(F, c^F_{\mu^\nu})$ is some mapping r of ν to μ. Now, if c^F_ν is some colouring of F then $c^F_\nu(x)$ is some colour χ_n of ν. Thus, it makes sense to consider the image of this colour χ_n via the mapping r and $(c^F_{\mu^\nu}(x))(c^F_\nu(x)) = r(\chi_n)$ is indeed some colour χ_m of μ. It makes therefore sense to write $(c^F_{\mu^\nu}(x))(c^F_\nu(x))$ in the above definition.

Lemma 5.14 *The notion defined above really is the exponential in the category REP(σ).*

PROOF. Let $(\lambda, \mathcal{L})$ be a representation. Moreover assume that

$$(\lambda, \mathcal{L}) \times (\nu, \mathcal{N}) \xrightarrow{r} (\mu, \mathcal{M}).$$

Consider the following mapping

$$\begin{aligned} r(\bullet, -) : \lambda &\to \mu^\nu \\ \chi_l &\mapsto \left(\begin{aligned} r(\chi_l, -) : \quad \nu &\to \mu \\ \chi_n &\mapsto r(\chi_l, \chi_n) \end{aligned} \right) \end{aligned}$$

We shall see that it is a recolouring of $(\lambda, \mathcal{L})$ to $(\mu, \mathcal{M})^{(\nu,\mathcal{N})}$. Let $(F, r(\bullet, -) \circ c^F_\lambda)$ be a forbidden pattern of $(\mu, \mathcal{M})^{(\nu,\mathcal{N})}$. By definition of the exponential, there exists some (F, c^F_ν) valid with respect to $(\nu, \mathcal{N})$ such that $(F, c^F_\mu) \in \mathcal{M}$, where $c^F_\mu = (r(\bullet, -) \circ c^F_\lambda) \otimes c^F_\nu = r(c^F_\lambda, c^F_\nu)$. Since (F, c^F_ν) is valid for $(\nu, \mathcal{N})$ and r is a recolouring, it follows from the definition of the product that (F, c^F_λ) is not valid for $(\lambda, \mathcal{L})$.

Conversely, assume that $(\lambda, \mathcal{L}) \xrightarrow{r'} (\mu, \mathcal{M})^{(\nu, \mathcal{N})}$. Consider the following mapping

$$\begin{aligned} (r' \circ \pi_\lambda) \otimes \pi_\nu : \lambda \times \nu &\to \mu \\ (\chi_l, \chi_n) &\mapsto (r'(\chi_l))(\chi_n) \end{aligned}$$

We want to show that $r = (r' \circ \pi_\lambda) \otimes \pi_\nu$ is a recolouring. Let $(F, r \circ c^F_{\lambda\times\nu}) \in \mathcal{M}$. We need to show that $(F, c^F_{\lambda\times\nu})$ is not valid for the product representation $(\lambda, \mathcal{L}) \times (\nu, \mathcal{N})$. There are two cases to consider

1. $(F, \pi_\nu \circ c^F_{\lambda\times\nu})$ is not valid for $(\nu, \mathcal{N})$: by definition of the product, $(F, c^F_{\lambda\times\nu})$ is not valid for $(\lambda, \mathcal{L}) \times (\nu, \mathcal{N})$ and we are done.

2. $(F, \pi_\nu \circ c^F_{\lambda\times\nu})$ is valid for $(\nu, \mathcal{N})$: by definition of the exponential, $(F, r' \circ \pi_\lambda \circ c^F_{\lambda\times\nu})$ is a forbidden pattern of $(\mu, \mathcal{M})^{(\nu, \mathcal{N})}$. Thus, since r' is a recolouring, it follows that $(F, \pi_\lambda \circ c^F_{\lambda\times\nu})$ is not valid for $(\lambda, \mathcal{L})$. Finally, by definition of the product, it follows that $(F, c^F_{\lambda\times\nu})$ is not valid for $(\lambda, \mathcal{L}) \times (\nu, \mathcal{N})$ and we are done.

□

Notice that this construction generalises the exponential of a σ-structure. Indeed, provided that the representation $(\nu, \mathcal{N})$ is simple (or at least canonical), the following can be proved.

representation	template
$(\mu, \mathcal{M})$	M
$(\nu, \mathcal{N})$	N
$(\mu, \mathcal{M})^{(\nu, \mathcal{N})}$	M^N

We have already seen at the end of Section 4.1.3 that the representation $\mathrm{o}_\sigma = (\emptyset, \{(0_\sigma, c^{0_\sigma}_\emptyset)\})$ is an initial object of $REP(\sigma)$. Define further the following representation $\mathrm{1}_\sigma := (\mathbf{1}, \emptyset)$. It is a straightforward exercise to check that it is a terminal object of $REP(\sigma)$.

The following result follows.

Theorem 5.15 *The category of σ-representation is a topos.*

PROOF. The proof is essentially the same as that of Theorem 5.4: for the equalizer of $(\mu, \mathcal{M}) \underset{g}{\overset{f}{\rightrightarrows}} (\nu, \mathcal{N})$, consider the subrepresentation of $(\mu, \mathcal{M})$ induced by the set $\{x \in |B| \text{ such that } f(x) = g(x)\}$; and, for the object classifier, consider the representation $(\mathbf{2}, \emptyset)$. □

Define the relation $\sim$ over $REP(\sigma)$ as follows: $\mathfrak{R}_1 \sim \mathfrak{R}_2$ holds for a pair of representations $\mathfrak{R}_1$ and $\mathfrak{R}_2$ if, and only if, $\mathfrak{R}_1 \twoheadrightarrow \mathfrak{R}_2$ and $\mathfrak{R}_2 \twoheadrightarrow \mathfrak{R}_1$. Clearly, $\sim$ defines an equivalence relation over $REP(\sigma)$. In order to obtain a Heyting algebra, we factor out the quasi-order given by the existence of a recolouring with respect to this equivalence relation. Note that as in the case of structures, cores of representations can be chosen as representatives for each equivalence class: in other words, the following holds.

$$\frac{\langle REP(\sigma), \twoheadrightarrow \rangle}{\sim} \approx \langle COREP(\sigma), \twoheadrightarrow \rangle$$

where $COREP(\sigma)$ denotes the class of cores of σ-representations. Define $\wedge, \vee$ and $\Rightarrow$ for representations as above for structures. It follows that

Corollary 5.16 *$\langle COREP(\sigma), \wedge, \vee, \Rightarrow, 0_\sigma, 1_\sigma \rangle$ is a Heyting algebra.*

Hence, the results from Section 5.2.3 apply to the case of representations; namely, there is also a correspondence between duality and density for representations. However, this result is not fully satisfactory; first, we do not have yet a characterisation of gap pairs in $REP(\sigma)$; and, secondly, note that the Heyting algebra of cores of representations is not as meaningful in our context as the Heyting algebra of cores of structures. Indeed, recall that the converse of Proposition 4.1 does not hold. That is, contrarily to the case of cores of structures where there is an exact correspondence between CSP and cores of structures, in the case of cores of representations, various cores of representations define the same forbidden patterns problem. Hence the real question should concern normal representations and not cores of representations according to the conjecture we motivate in the next subsection.

5.3.2 The containment problem for forbidden patterns problems

A homomorphism problem is given by its template; hence given two homomorphism problems $CSP(A)$ and $CSP(B)$ over the same signature, it is decidable whether $CSP(A) \subseteq CSP(B)$. As a matter of fact, the containment problem for homomorphism problems is nothing else than the uniform homomorphism problem, known to be NP-complete. We would like to extend this result to the more general containment problem for forbidden patterns problems given by their representations. Feder and Vardi proved in [16] that the containment problem for MMSNP is decidable. Hence by our results from Subsection 4.1.4, it follows that the containment problem for forbidden patterns problems is decidable. However, there is no known result about the complexity of the containment problem for MMSNP. Furthermore, even if it were the case, the constructions we use to translate a sentence of MMSNP into a forbidden patterns problem are not meaningful in the context of complexity theory, as the transformation is clearly not polynomial (notice for example, the need for forbidden patterns to be coloured structures, whereas negated conjuncts correspond in general to *partially* coloured structures). The major inconvenience of forbidden patterns problems, by opposition with homomorphism problems, is that the inclusion of two problems does not reduce to the question of the existence of a recolouring: we introduced in Chapter 4 the notion of Feder-Vardi transformation of a representation, which allows one to transform a representation into another representation that represents the same forbidden patterns problem, but that is not necessarily equivalent with respect to recolouring (*cf.* example following Corollary 4.22). In the light of this fact, we could extend our morphisms in the category $REP(\sigma)$. That is, define a morphism between two representations as a finite sequence of recolourings and Feder-Vardi transformations. This yields the following question: does this new category represent faithfully the inclusion relation between forbidden patterns problems? As this question seems still quite hard and because we have at hand a normal form for representations with "good" properties, we can first concentrate on the case of connected normal representations. We shall prove in the remainder of this section some results that support the following conjecture.

Conjecture 5.17 *Let $\mathfrak{R}_1$ and $\mathfrak{R}_2$ be two non-trivial connected representations. $FP(\mathfrak{R}_1) \subseteq FP(\mathfrak{R}_2)$ if, and only if,* $\mathbf{normal}(\mathfrak{R}_1) \longrightarrow \mathbf{normal}(\mathfrak{R}_2)$.

The converse implication holds: we have $FP(\mathfrak{R}_1) = FP(\mathbf{normal}(\mathfrak{R}_1))$ and $FP(\mathfrak{R}_2) = FP(\mathbf{normal}(\mathfrak{R}_2))$, by Theorem 4.23, and by assumption

$$\mathbf{normal}(\mathfrak{R}_1) \longrightarrow \mathbf{normal}(\mathfrak{R}_2);$$

hence, by Proposition 4.1, it follows that $FP(\mathbf{normal}(\mathfrak{R}_1)) \subseteq FP(\mathbf{normal}(\mathfrak{R}_2))$.

We now prove some supportive results with respect to the other implication. Assume that $FP(\mathfrak{R}_1) \subseteq FP(\mathfrak{R}_2)$ and that $\mathbf{normal}(\mathfrak{R}_1)$ is conform, and let T_1 be its template. We have $FP(\mathfrak{R}_1) = CSP(T_1) \ni T_1$. Hence, T_1 is a yes-instance of $FP(\mathfrak{R}_2)$: that is, there exists some r such that $T_1 \xrightarrow{r} T_2$ (where T_2 denotes the template of $\mathbf{normal}(\mathfrak{R}_2)$) such that for any non-conform forbidden pattern $F \xrightarrow{c^F} T_2$ of $\mathbf{normal}(\mathfrak{R}_2)$, we can not have some homomorphism $F \xrightarrow{f} T_1$ and the following commutative diagram

$$\begin{tikzcd} F \arrow[r, "f"] \arrow[dr, "c^F"'] & T_1 \arrow[d, "r"] \\ & T_2 \end{tikzcd}$$

Hence, the remark following Proposition 4.16 implies that:

$$\mathbf{normal}(\mathfrak{R}_1) \xrightarrow{r} \mathbf{normal}(\mathfrak{R}_2).$$

We have just proved that the above conjecture holds when the first representation has a conform normal form.

Proposition 5.18 *Let $\mathfrak{R}_1$ and $\mathfrak{R}_2$ be two connected representation. Furthermore, assume that* $\mathbf{normal}(\mathfrak{R}_1)$ *is conform.*

$$FP(\mathfrak{R}_1) \subseteq FP(\mathfrak{R}_2) \text{ if, and only if, } \mathbf{normal}(\mathfrak{R}_1) \longrightarrow \mathbf{normal}(\mathfrak{R}_2).$$

We shall need the following lemma.

Lemma 5.19 *Let $\mathfrak{R}_1$ and $\mathfrak{R}_2$ be two connected representations. If $\mathfrak{R}_1 \longrightarrow \mathfrak{R}_2$ then* $\mathbf{normal}(\mathfrak{R}_1) \longrightarrow \mathbf{normal}(\mathfrak{R}_2)$.

PROOF. Note that $\mathbf{normal}(\mathfrak{R}_1) \longrightarrow \mathfrak{R}_1$ (*cf.* the remark on the end of Subsection 4.5.1). Hence, if $\mathfrak{R}_1 \longrightarrow \mathfrak{R}_2$ then $\mathbf{normal}(\mathfrak{R}_1) \longrightarrow \mathfrak{R}_2$. So we may assume w.l.o.g. that $\mathfrak{R}_1$ is normal and that $\mathfrak{R}_1 \xrightarrow{r} \mathfrak{R}_2$. Let $\mathfrak{R}_1 = (T, \mathcal{M})$ with $|T| = \mu$ and $\mathfrak{R}_2 = (\mathrm{v}, \mathcal{N})$.

It suffices to check that we can construct a recolouring from r after each elementary Feder-Vardi transformation (the other transformations involved in the computation of the normal form yield representations that are equivalent with respect to recolouring equivalence). For simplicity, we do not consider compact forbidden patterns. This does not change our result, as a compact forbidden pattern $(S, c^S_{\wp(\mathrm{v})})$ stands for a set of forbidden patterns

$$\mathcal{E} := \{(S, c^S_{\mathrm{v})}) \text{ such that for any } x \in |S|,\ c^S_{\mathrm{v}}(x) \in c^S_{\wp(\mathrm{v})}(x)\}$$

and were introduced solely to prove termination: in fact, carrying out an elementary Feder-Vardi transformation with respect to $(S, c^S_{\wp(\mathrm{v})})$ corresponds to carrying out the elementary Feder-Vardi transformations with respect to each forbidden pattern in $\mathcal{E}$ in parallel.

Let $(S, c^S_{\mathrm{v}}) \in \mathcal{N}$ be a non-biconnected forbidden pattern of $\mathfrak{R}_2$ that admits a decomposition $(P_0, c^{P_0}_{\mathrm{v}}) \underset{x}{\bowtie} (P_1, c^{P_1}_{\mathrm{v}})$. Let $\tilde{\mathfrak{R}}_2$ be the elementary Feder-Vardi transformation of $\mathfrak{R}_2$ with respect to the decomposition $(P_0, c^{P_0}_{\mathrm{v}}) \underset{x}{\bowtie} (P_1, c^{P_1}_{\mathrm{v}})$ of (S, c^S_{v}) and let $\chi := c^S_{\mathrm{v}}(x)$.

1. (S, c^S_{v}) is not of the form $(S, r \circ c^S_{\mu})$.
   Consider $\tilde{r}$ to be the mapping that agrees with r for any $\chi' \in \mathrm{v}$ such that $r(\chi') \neq \chi$; and, such that $\tilde{r}(\chi') = \chi_0$, otherwise. Clearly, we have $\mathfrak{R}_1 \xrightarrow{\tilde{r}} \tilde{\mathfrak{R}}_2$.

2. (S, c^S_{v}) is of the form $(S, r \circ c^S_{\mu})$.
   The fact that r is a recolouring and $\mathfrak{R}_1$ is normal implies that any inverse image $(P_0, c^{P_0}_{\mu}) \underset{x}{\bowtie} (P_1, c^{P_1}_{\mu})$ of $(P_0, c^{P_0}_{\mathrm{v}}) \underset{x}{\bowtie} (P_1, c^{P_1}_{\mathrm{v}})$ via r is such that either:

   - for $i \in \{0, 1\}$, the colouring $c^{P_i}_{\mu}$ is not a homomorphism of P_i to T; or
   - for $i \in \{0, 1\}$, there exists some biconnected conform forbidden pattern $(M, c^M_{\mu}) \in \mathcal{M}$ such that $(M, c^M_{\mu}) \xhookrightarrow{m} (P_i, c^{P_i}_{\mu})$.

   Let $\chi' \in \mu$. Let $\mathcal{S}_{\chi'}$ be the set of inverse images $(P_0, c^{P_0}_{\mu}) \underset{x}{\bowtie} (P_1, c^{P_1}_{\mu})$ of

$(P_0, c_\nu^{P_0}) \underset{x}{\bowtie} (P_1, c_\nu^{P_1})$ via r such that $c_\mu^{P_0}(x) = c_\mu^{P_1}(x) = \chi$.
The key property that shall allow us to build a recolouring is some kind of uniformity principle:

Fact 5.20 *There exists some $i \in \{0,1\}$ such that for any $(P_0, c_\mu^{P_0}) \underset{x}{\bowtie} (P_1, c_\mu^{P_1}) \in \mathcal{S}_{\chi'}$, either:*

- *the colouring $c_\mu^{P_i}$ is not a homomorphism of P_i to T; or*
- *there exists some biconnected conform forbidden pattern $(M, c_\mu^M) \in \mathcal{M}$ such that $(M, c_\mu^M) \xhookrightarrow{m} (P_i, c_\mu^{P_i})$.*

We call i an invalid component *of $\mathcal{S}_{\chi'}$.*

To see this fact, note that once the inverse image of the colour of x in the inverse image has been chosen, say χ', the choice of the inverse images for each component is independant. So, if the above did not hold then we could choose a valid colouring for each component and r would not be a recolouring.
Let χ be the colour of x in $(P_0, c_\nu^{P_0}) \underset{x}{\bowtie} (P_1, c_\nu^{P_1})$. We now construct some $\tilde{r}$ from r:

- for any colour $\chi' \in \mu$ such that $r(\chi') \neq \chi$, $\tilde{r}$ agrees with r; and
- otherwise, $\tilde{r}(\chi') := \chi_i$ where i is the invalid component of $\mathcal{S}_{\chi'}$.

By construction, we have $\mathfrak{R}_1 \xrightarrow{\tilde{r}} \tilde{\mathfrak{R}}_2$.

This concludes the proof. □

Consider now the case of monochrome forbidden pattern problems. Let A and B be two σ-structures. Suppose that $FP(A) \subseteq FP(B)$. Since B is a no-instance of $FP(B)$, it follows that B is a no-instance of $FP(A)$; in other words that there exists some homomorphism $A \xrightarrow{h} B$. Hence, that id_1 is a recolouring of the monochrome representation $(\mathbf{1}, (A, c_1^A))$ of the first problem to the monochrome representation of the second problem $(\mathbf{1},(B,c_{\mathbf{1}}^{B}))$ as quite clearly the following diagram commutes.

$$\begin{tikzcd} \mathbf{1} \arrow[r, "id_{\mathbf{1}}"] & \mathbf{1} \\ A \arrow[u, "c_{\mathbf{1}}^{A}"] \arrow[r, "h"'] & B \arrow[u, "c_{\mathbf{1}}^{B}"'] \arrow[ul, "c_{\mathbf{1}}^{B}"'] \end{tikzcd}$$

In the light of Lemma 5.19, it follows that:

$$\mathbf{normal}(\mathbf{1},(A,c_{\mathbf{1}}^{A})) \longrightarrow \mathbf{normal}(\mathbf{1},(B,c_{\mathbf{1}}^{B})).$$

Notice that the above proof extends to the case of monochrome forbidden patterns problems (note the plural). Hence the conjecture holds also in the case of monochrome forbidden patterns problems and we can state the following.

Proposition 5.21 *Let $\mathfrak{N}_1$ and $\mathfrak{N}_2$ be two monochrome forbidden patterns problems.*
$FP(\mathfrak{N}_1) \subseteq FP(\mathfrak{N}_2)$ if, and only if, $\mathbf{normal}(\mathfrak{N}_1) \longrightarrow \mathbf{normal}(\mathfrak{N}_2)$.

I think that one possible approach to the conjecture in the general case would be to use the exponential of a representation. My intuition comes from the fact that the exponential of a representation contains somehow some information about "cleverer" recolourings; these recolourings being adaptive and taking into account the fact that somewhere "local", a structure that defines a forbidden pattern occurs or not.

To conclude this chapter, let us mention the possibility of defining a *hierarchy* of problems. Let $\mathfrak{T}$ be some σ-representation. The *(non-uniform) recolouring problem* with template $\mathfrak{T}$ is the problem that takes as instances σ-representations; and, has yes-instances those σ-representations $\mathfrak{N}$ such that $\mathfrak{N} \longrightarrow \mathfrak{T}$. In the same way that forbidden patterns problems generalise homomorphism problems, one can define problems that generalise the recolouring problems: these problems are given by a *second generation representation* that consists of a representation $\mathfrak{T}$ (the *template*), together with a finite set $\mathcal{F}$ of *forbidden ($\mathfrak{T}$-recoloured) representations* $\mathfrak{F} \xrightarrow{c_{\mathfrak{T}}^{\mathfrak{F}}} \mathfrak{T}$. This problem takes representations as instances and has yes-instances those representations $\mathfrak{A}$ such that there exists a recolouring $\mathfrak{A} \xrightarrow{c_{\mathfrak{T}}^{\mathfrak{A}}} \mathfrak{T}$, such that for every forbidden representation $\mathfrak{F} \xrightarrow{c_{\mathfrak{T}}^{\mathfrak{F}}} \mathfrak{T}$ in $\mathcal{F}$, if $\mathfrak{F} \xrightarrow{c_{\mathfrak{A}}^{\mathfrak{F}}} \mathfrak{A}$ then the following does not commute

$$\begin{tikzcd}
\mathfrak{A} \arrow[r, "c^{\mathfrak{A}}_{\mathfrak{T}}"] & \mathfrak{T} \\
& \mathfrak{F} \arrow[u, swap, "c^{\mathfrak{F}}_{\mathfrak{T}}"] \arrow[ul, "c^{\mathfrak{F}}_{\mathfrak{A}}"]
\end{tikzcd}$$

We could then define a notion of recolouring of second generation and so on.

Chapter 6

On the complexity of forbidden-patterns problems

We show that there are complete forbidden patterns problems for NL, P and NP that are not homomorphism problems. We investigate also some restrictions that ensure the tractability of forbidden patterns problems.

As we have seen in previous chapters, homomorphism problems can be generalised in term of forbidden patterns problems; the latter are the problems that correspond to the logic MMSNP. This logic was introduced by Feder and Vardi in an attempt to capture homomorphism problems. Some extensive investigation has been carried out on the complexity of homomorphism problems in the last decade, the ultimate aim being to prove a dichotomy result. Hence there exists results characterising whether a homomorphism problem is tractable or NP-complete. There are however fewer finer results about the complexity of those problems that are known to be tractable; some tractable problems are known to be in NL (see [28]). In the next chapter, we shall give the first known example of homomorphism problems that are complete for L. In this chapter we show that for forbidden patterns problems that are not homomorphism problems, there are examples that are complete for each main complexity class within NP (except for L). Then, in a second part we investigate some restrictions on the input that can lead to tractability.

Recall that we provided some definitions in Appendix A.

6.1 Examples of complete problems for each class

In order to give complete problems for NL, P and NP, we use first the fact that forbidden-patterns problems correspond to the logic MMSNP to read directly from their defining MMSNP sentence the complexity class to which they belong by hand of Grädel's elegant logical characterisations (see [21]). Then, to prove completeness we simply encode known complete problems using forbidden patterns problems. The present section is by no means an attempt of characterising the complexity of forbidden patterns problems but rather an illustration of what kind of problems can be encoded using forbidden patterns problems.

6.1.1 An NL-complete problem

Let $\sigma_{2,2} := (E_1, E_2)$, where E_1 and E_2 are two binary relation symbols. Consider $\mathfrak{S}$ to be the representation with,

- colour set $\{0,1\}$; and
- forbidden patterns WDC_2^1, WDC_2^2 and BDC_2^2 (as depicted in Figure 6.1):

here the top index denotes the type of edges involved in a forbidden pattern (on the figure, edges of type E_1 are drawn as solid lines and edges of type E_2 as dotted lines).

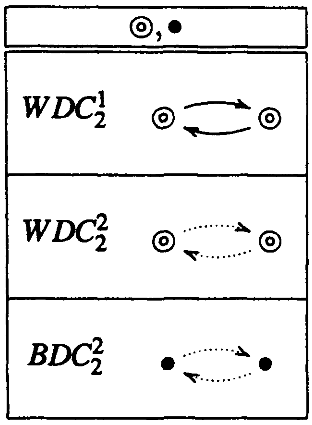

Figure 6.1: The representation $\mathfrak{S}$

Fact 6.1 *$FP(\mathfrak{S})$ is in* NL.

PROOF. Let W be a monadic predicate (standing for white) and x, y, z be some variables. WDC^1_2 corresponds to the following negated conjunct

$$\neg(E_1(x,y) \wedge E_1(y,x) \wedge W(x) \wedge W(y))$$

WDC^2_2 corresponds to the following negated conjunct

$$\neg(E_2(x,y) \wedge E_2(y,x) \wedge W(x) \wedge W(y))$$

and BDC^2_2 corresponds to

$$\neg(E_2(x,y) \wedge E_2(y,x) \wedge \neg W(x) \wedge \neg W(y)).$$

Hence the following sentence of MMSNP expresses exactly the problem $FP(\mathfrak{S})$.

$$\begin{aligned}\exists W \forall x \forall y \forall z \quad & \neg(E_1(x,y) \wedge E_1(y,x) \wedge W(x) \wedge W(y)) \\ & \wedge \neg(E_2(x,y) \wedge E_2(y,x) \wedge W(x) \wedge W(y)) \\ & \wedge \neg(E_2(x,y) \wedge E_2(y,x) \wedge \neg W(x) \wedge \neg W(y))\end{aligned}$$

Notice that this sentence has at most two occurences of the monadic predicate W in each negated conjunct, that is, it is in the fragment of second order logic known as ESO-Krom. By a result of Grädel, this logic is known to capture the complexity class NL. Hence the result follows. □

Fact 6.2 *$FP(\mathfrak{S})$ is hard for* NL.

PROOF. The restriction of SAT to formulas with at most two literals per conjunct, namely 2-SAT, is known to be complete for NL. We reduce 2-SAT to $FP(\mathfrak{S})$. For each variable y that occurs in some instance φ of 2-SAT, we put two elements v_y and $v_{\bar{y}}$, one for each literal. Moreover we set $E_2(v_y, v_{\bar{y}})$ and $E_2(v_{\bar{y}}, v_y)$ to hold. For each clause C of φ involving two literals ℓ_1 and ℓ_2, we set $E_1(v_{\ell_1}, v_{\ell_2})$ and $E_1(v_{\ell_2}, v_{\ell_1})$ to hold. Denote by G_φ this $\sigma_{2,2}$-structure. We claim that $\varphi \in$ 2-SAT if, and only if, $G_\varphi \in FP(\mathfrak{S})$. See white as false and black as true. A colouring of G_φ valid w.r.t. WDC_2^2 and BDC_2^2 corresponds exactly to an assignment of the variables of the formula φ, since these two forbidden patterns enforce that the vertices corresponding to opposite literals have opposite colours. If a colouring is also valid w.r.t. the forbidden pattern WDC_2^1 then the corresponding assignment for φ is valid; indeed, the forbidden pattern WDC_2^2 enforces that at least one of two vertices v_{ℓ_1} and v_{ℓ_2}, that corresponds to the literals of a clause C, is coloured black. Clearly, the converse also holds. It can be checked that this transformation can be achieved via a quantifier-free first-order reduction. □

Hence we obtain the following corollary using the theorem of Subsection 4.7.1.

Corollary 6.3 *$FP(\mathfrak{S})$ is* NL*-complete and is a forbidden patterns problem that is not a homomorphism problem.*

Notice that it is probably not true that all forbidden patterns problems that are in NL have a defining MMSNP sentence that is also in ESO-Krom. Indeed the important mechanism of being able to use the full power of second order logic is missing if we restrict ourselves to MMSNP where we use only monadic predicates. Here we used Grädel's result only to provide a quick proof of the complexity of our example.

6.1.2 A P-complete problem

The following example is an adaptation of an example of a P-complete problem from [21]. Consider the following signature $\sigma_c = (E_1, E_2, S^+, S^-, A)$ where the symbols are of respective arities $2, 2, 1, 1$ and 1. Define CVP to be the problem captured by the following sentence of MMSNP.

$$\begin{aligned}\exists T \exists F \forall x \forall y \forall z \quad & \neg(S^+(x) \wedge \neg T(x)) \wedge \neg(S^-(x) \wedge \neg F(x)) \\ & \wedge \neg(E_1(x,z) \wedge E_2(z,x) \wedge F(x) \wedge \neg T(z)) \\ & \wedge \neg(\mathtt{NAND}(x,y,z) \wedge T(x) \wedge T(y) \wedge \neg F(z)) \\ & \wedge \neg(T(x) \wedge F(x)) \wedge \neg(A(x) \wedge \neg T(x))\end{aligned}$$

where:

$$\mathtt{NAND}(x,y,z) = E_1(x,z) \wedge E_2(z,x) \wedge E_1(y,z) \wedge E_2(z,y) \wedge E_2(x,y) \wedge E_2(y,x).$$

Note that this sentence is in ESO-Horn. It follows that the problem CVP is in the class P. Moreover it is complete for this class, as it encodes the *circuit value problem*. The predicate S^+ corresponds to the positive inputs of the circuit; the predicate S^- to the negative inputs; and, the predicate A to the output of the circuit. Using the relations E_1 and E_2, we can encode Nand gates (Sheffer's stroke); put an edge of the first type between the input x of a gate and the output of a gate z and an edge of the second type from z to x; and, put edges of the second type between the input x and y of a gate. The monadic predicate T stands for true and the monadic predicate F for false. The first negated conjunct ensures that positive inputs are set to true. The second one that negative inputs are set to false. The third negated conjunct enforces that if one of the inputs of a NAND gate is false then its output is true. The fourth negated conjunct ensures that if both inputs of a gate are true then the output is false. The fifth negated conjunct enforces that we can not have a vertex set simultaneously to true and false. The last negated

conjunct states that the output is set to true. Note that we do not need the negated conjunct $\neg(\neg T(x) \wedge \neg F(x))$, as this can not occur in a σ_c-structure that encodes a circuit because of the first four negated conjuncts (this is the trick that allows us to have a sentence in ESO-Horn).
We complete the colouring and simplify the above sentence and get the following good sentence that is logically equivalent (*cf.* Proposition 3.3):

$$
\begin{aligned}
\exists T \exists F \forall x \forall y \forall z \quad & \neg(S^+(x) \wedge \neg T(x) \wedge F(x)) \wedge \neg(S^+(x) \wedge \neg T(x) \wedge \neg F(x)) \\
& \wedge \neg(S^-(x) \wedge T(x) \wedge \neg F(x)) \wedge \neg(S^-(x) \wedge \neg T(x) \wedge \neg F(x)) \\
& \wedge \neg(E_1(x,z) \wedge E_2(z,x) \wedge \neg T(x) \wedge F(x) \wedge \neg T(z) \wedge F(z)) \\
& \wedge \neg(E_1(x,z) \wedge E_2(z,x) \wedge \neg T(x) \wedge F(x) \wedge \neg T(z) \wedge \neg F(z)) \\
& \wedge \neg(\mathtt{NAND}(x,y,z) \wedge T(x) \wedge \neg F(x) \wedge T(y) \wedge \neg F(y) \wedge \neg T(z) \wedge \neg F(z)) \\
& \wedge \neg(\mathtt{NAND}(x,y,z) \wedge T(x) \wedge \neg F(x) \wedge T(y) \wedge \neg F(y) \wedge T(z) \wedge \neg F(z)) \\
& \wedge \neg(T(x) \wedge F(x)) \\
& \wedge \neg(A(x) \wedge \neg T(x) \wedge F(x)) \wedge \neg(A(x) \wedge \neg T(x) \wedge \neg F(x))
\end{aligned}
$$

We shall now build the representation that corresponds to this sentence; however, since the colour $(T(x) \wedge F(x))$ is not allowed, we directly remove it from the set of colours. We get a representation with three colours:

1. ◎ for $(\neg F(x) \wedge T(x))$;
2. • for $(F(x) \wedge \neg T(x))$; and
3. for $(\neg F(x) \wedge \neg T(x))$.

We write S^+(◎) to depict $\neg(S^+(x) \wedge \neg T(x) \wedge F(x))$ and proceed similarly for the other monadic predicates from σ_c. Let $\mathfrak{C}$ be the representation hence obtained. $\mathfrak{C}$ is depicted in Figure 6.2. Showing that the corresponding forbidden patterns problem is not in CSP requires to compute the normal form of the above. This is rather tedious as the fifth forbidden pattern has a homomorphic image that is not biconnected:

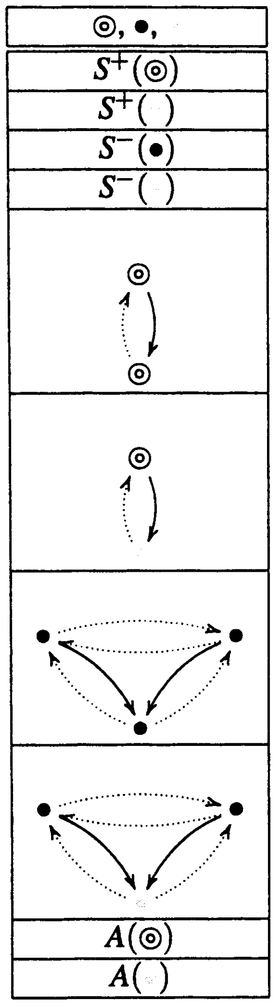

Figure 6.2: The representation $\mathfrak{C}$

The seventh has two such homomorphic images:  and 

The eighth has one such homomorphic image: 

However, after a Feder-Vardi transformation they do not yield any conform forbidden patterns. Hence, the normal form of the representation $\mathfrak{C}$ is not conform and we have the following.

Corollary 6.4 CVP *is* P*-complete and is a forbidden patterns problem that is not a homomorphism problem.*

6.1.3 An NP-complete problem

The problem NO-MONO-TRI was already considered in [16] as an example of an NP-complete problem in MMSNP but not in CSP, but they referred to [20] for completeness; as a matter of fact the problem considered in [20] involves colouring of the edges.

Proposition 6.5 *The problem* NO-MONO-TRI *is computationally equivalent to the problem* NAE-SAT*:*

- NO-MONO-TRI $\leq_{q.f.FO}$ NAE-SAT*; and*
- NO-MONO-TRI $\geq_{FO}$ NAE-SAT.

PROOF. First, we reduce an instance G of NO-MONO-TRI to NAE-SAT, that is a set U of variables and a collection C of clauses over U such that each clause $c \in C$ has length 3. (Recall that NAE-SAT asks the following question: is there a truth assignment for U such that each clause in C has at least one true literal and at least one false literal?). The traditional encoding for NAE-SAT involves a signature $\sigma_n = (C_0, C_1, C_2, C_3)$, where the C_i are ternary predicates. Hence a σ_n-structure U can be seen as an encoding of an instance of NAE-SAT; its universe is a set of variables, and if $C_i(x,y,z)$ holds, it means that there is a clause involving x, y and z, where the i first variable(s) appear as negative literal(s) and the other(s) positively. Let

$$\Pi = (\varphi_0, \varphi_1, \varphi_2, \varphi_3),$$

where:

$$\begin{aligned}\varphi_0(x,y,z) = (E(x,y) \vee E(y,x)) \wedge (E(y,z) \vee E(z,y)) \wedge (E(z,x) \vee E(x,z))\\ \varphi_1 = \mathbf{false}\\ \varphi_2 = \mathbf{false}\\ \varphi_3 = \mathbf{false}\end{aligned}$$

Π is an interpretation of σ_n in σ_2 of width one; and, clearly, $U \in$ NO-MONO-TRI if, and only if, $\Pi(U) \in$ NAE-SAT. Thus, NO-MONO-TRI $\leq_{q.f.FO}$ NAE-SAT.

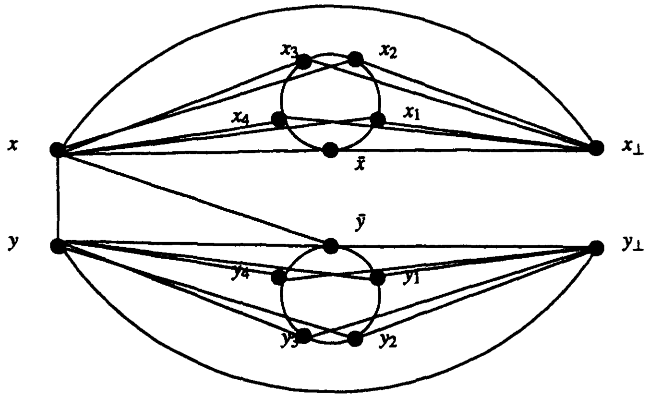

Figure 6.3: example of the reduction of one clause $\{\bar{y}, y, x\}$.

Now, we shall reduce NAE-SAT to NO-MONO-TRI via a FO-interpretation. We, first introduce the idea of the reduction in more traditional terms, and in a second time show that this reduction can be implemented via FO-interpretation. First, we need to define a graph, used as a gadget in the reduction. Let G_5 be the graph with vertices $\{x_\top, x_0, x_1, x_2, x_3, x_4, x_\perp\}$, and whose edges consist of the union of the following sets:

- $\{(x_0,x_1),(x_1,x_2),(x_2,x_3),(x_3,x_4),(x_4,x_0)\}$;
- $\{(x_\top,x_i)|i = 0,4\}$;
- $\{(x_\perp,x_i)|i = 0,4\} \cup \{(x_\top,x_\perp)\}$.

Note that there are only two possible 2-colourings of G_5 such that G_5 has no monochromatic triangle and, further, that these colourings set $x_\top$ and $x_\bot$ with the same colour, whereas the x_i's are set the other colour.
For every instance (U,C) of NAE-SAT, we construct the graph G as follows.

- G has a vertex x and a vertex $\bar{x}$ for each variable x in G; and,
- we add a copy of the gadget G_5 between any two such vertices x and $\bar{x}$, identifying x with $x_\top$ and $\bar{x}$ with x_0; and,
- for every clause $c \in C$ involving three literals ℓ_1, ℓ_2, ℓ_3, we add three special vertices $\ell_1^c, \ell_2^c, \ell_3^c$ and three copies of G_5 that enforce that the ℓ_i^c's and the ℓ_i's have opposite colours.
- Finally, the constraint given by the clause c between the literals ℓ_1, ℓ_2, ℓ_3 is enforced by adding a triangle between the three special vertices [1] $\ell_1^c, \ell_2^c, \ell_3^c$.

Suppose that the original instance is satisfiable: then colour in white one node corresponding to a literal assigned to false and in black a node corresponding to a literal assigned to true. Now, colour the gadget as follows, assign to $x_\bot$ the same colour as the one assigned to x, and assign the opposite colour to $x_1, \dots, x_4$. Clearly, this colouring does not introduce any monochromatic triangle and the graph belongs to NO-MONO-TRI. On the other hand, if the graph belongs to NO-MONO-TRI, the nodes added enforce that nodes x and $\bar{x}$ have an opposite colour and because every triangle corresponding to a clause is non-monochromatic, at least one literal per clause must have been assigned a value different from the other literals.

This reduction can be implemented via a FO-interpretation. We leave this as an exercise for the reader. □

We have proved in Section 2.4 that NO-MONO-TRI was not in CSP. We get the following.

Corollary 6.6 NO-MONO-TRI *is* NP*-complete and is a forbidden patterns problem that is not a homomorphism problem.*

[1] We can not add directly a triangle between ℓ_1, ℓ_2, ℓ_3, otherwise the interaction of such triangles may well lead to a triangle that does not correspond to a clause of the instance (U,C) of NAE-SAT.

6.2 Some restrictions ensuring tractability

There are well-known restrictions over instances of difficult graph problems which tend to give rise to tractable problems; restrict the girth of the instances, restrict the problem over trees, over planar graphs or over graphs of some suitable bounded degree. We briefly discuss these approaches in this section.

6.2.1 High girth

The first kind of obvious restriction for forbidden patterns problems whose normal form has no conform forbidden patterns like NO-MONO-TRI consists in restricting the instance to have sufficiently high girth such that none of the forbidden patterns can occur in any colouring. Hence clearly we have the following.

Fact 6.7 *Every σ_2 structure that encodes a graph with girth greater or equal to 4 belongs to* NO-MONO-TRI.

This can be generalised as follows.

Corollary 6.8 *Let $(T, \mathcal{M})$ be some normal connected representation. Let g be the largest cycle that embeds in a forbidden pattern from $\mathcal{M}$. If $CSP(T)$ is tractable then the problem $FP(T, \mathcal{M})$ restricted to instances of girth strictly greater than g is tractable.*

PROOF. Let A be some instance of girth greater than g. If $A \not\longrightarrow T$ then $A \notin FP(T, \mathcal{M})$; otherwise, any $A \xrightarrow{h} T$ is a valid colouring w.r.t. $(T, \mathcal{M})$. In other words, the problem reduces to $CSP(T)$. □

6.2.2 Bounded tree width

The approach restricting the instances of some difficult graph problems to trees (thus avoiding back-track) can be generalised to instances of bounded tree-width (thus avoiding back-track once it has been checked that an instance is locally satisfiable). For the constraints satisfaction problem, this has been investigated among others by Freuder [18, 19] and Dechter *et al.* [8, 10]. Recently, the latter has proposed a unifying framework based on the algorithmic aspect of this method:

bucket elimination [9]. A more formal generalisation is also known for problems in monadic second order logic. This general result was proved by Courcelle [5]. This leads to the following.

Corollary 6.9 *Let k be some fixed positive integer. When restricted to instances of tree-width at most k, a forbidden pattern problem is tractable* [2].

6.2.3 Bounded degree

A further way of restricting graph problems is well-known; it consists in considering only graphs of a certain bounded degree. We investigate here the case of NO-MONO-TRI.

Lemma 6.10 *Every σ_2-structure that encodes a graph of degree at most two is a yes-instance of* NO-MONO-TRI.

PROOF. There is an obvious algorithm to build valid colourings of such instances. Every connected component can be dealt with independently. So assume w.l.o.g. that the instance is connected. Pick up some vertex and colour it in white. Pick up the vertices it is adjacent to (there are at most two) and colour them black and so on. We have levels that correspond to each stage of the algorithm. There can not be any edges between two vertices that are at least two levels apart. Moreover there are at most two vertices per level. Hence the result clearly follows. □

6.2.4 Planar instances

Another way of restricting a forbidden patterns problem to obtain tractability would probably involve some concept near the concept of planarity for graphs. We shall use here the four colour theorem to prove that our main example NO-MONO-TRI becomes tractable (as a matter of fact it becomes trivial) when restricted to planar graphs.

[2] More precisely, in linear time: the problem is decidable in time linear in the structure size but also the solutions are computable in time linear in the structure size plus the size of the output by a recent generalisation of Courcelle's result [17].

Lemma 6.11 *Every σ_2-structure that encodes a planar graph is a yes-instance of* NO-MONO-TRI.

PROOF.This short and elegant argument has been proposed by Regis Barbanchon. Let A be a σ_2-structure that encodes a planar graph G. By the four-colour theorem, G is 4-colourable (in the restricted sense: adjacent edges have different colour). Consider some valid 4-colouring c^G of the vertices of G with $\{0,1,2,3\}$. Colour in 0 those vertices that have been coloured in 0 and 2 and in 1 otherwise. This colouring of G has no monochromatic triangle, otherwise c^G would not be a valid colouring. □

Hence we obtain the following.

Corollary 6.12 NO-MONO-TRI *is tractable (trivial) when restricted to instances encoding planar graphs.*

Recall that planar graphs can be defined in terms of forbidden minors. So, it would be interesting to investigate how sets of graphs defined in terms of forbidden minors compare with forbidden patterns problems.

Chapter 7

Algebra homomorphism problems

We show that the uniform constraint satisfaction problem where instances consist of pairs of unary functions can be solved in logspace. We also show that any analogous non-uniform problem is L-complete if the (fixed) template function does not contain a fixed point; otherwise it consists of all unary functions. There is a significant jump in complexity when we consider constraint satisfaction problems where the instances are pairs of pairs of unary functions: the uniform problem can trivially be solved in NP and we show that there exist non-uniform problems that are NP-complete. For information this chapter has been derived from a joined work (see [42]) with Iain Stewart that has been submitted for publication.

There are two outstanding and well-known results which illustrate the attempt of classifying the complexity of non-uniform constraint satisfaction problems. The first was established by Schaefer [52] who completely classified the complexity of a non-uniform constraint satisfaction problem when the template is a finite structure whose domain consists of two elements, *i.e.*, the template is a *Boolean* structure. He showed that if the template belongs to one of six specific classes of Boolean structures then the non-uniform constraint satisfaction problem is solvable in polynomial-time, otherwise it is NP-complete. Note the dichotomy here: a non-uniform constraint satisfaction problem with a Boolean template is either in Por is NP-complete (recall that, in general, if P$\neq$ NP then there is an infinite collection of distinct classes of polynomial-time equivalent problems between P and NP, by a result of Ladner). The second result is due to Hell and Nešetřil [23] who showed that if all structures involved are finite undirected graphs (without self-loops) then the non-uniform constraint satisfaction problem is solvable in polynomial-time if the template is bipartite, otherwise it is NP-complete (again, note the dichotomy). For more details on those results *cf.* Section2.2.

In this chapter, we look at the computational complexity of constraint satisfaction problems involving, first, finite structures consisting of one unary function, and, second, finite structures consisting of two unary functions. In the first case, the uniform constraint satisfaction problem can be solved in L and there exist non-uniform constraint satisfaction problems whose complexity is L-complete. Indeed, we obtain a rather severe dichotomy result for such non-uniform problems: we show that such a non-uniform constraint satisfaction problem is always L-complete unless the unary function template contains a fixed point when the problem consists of all unary functions (and so is trivial). In the second case, the uniform constraint satisfaction problem can trivially be solved in NP and we show that there exist non-uniform constraint satisfaction problems whose complexity is NP-complete. Our results add to the ongoing classification of constraint satisfaction problems and, as far as we know, provide the first classification of a natural class of non-uniform constraint satisfaction problems where the complexity measures are 'below' P (assuming L $\neq$ P).

7.1 Basic definitions

A *signature* consists of a finite collection of constant symbols, function symbols and relation symbols, and each function and relation symbol has an associated arity. A *finite structure* A over the signature σ, or σ-*structure*, consists of a finite set $|A|$, the *domain* or *universe*, together with a constant C^A (resp. function F^A, relation R^A) for every constant symbol C (resp. function symbol F, relation symbol R) of σ, with functions and relations being of the appropriate arity (we usually only include superscripts in the names of our constants, functions and relations when it may be unclear as to which structure we are dealing with). The *size* of a structure A is the size of the domain and is denoted $|A|$ also. A *homomorphism* $\varphi : A \to B$ of a σ-structure A to a σ-structure B is a map $\varphi : |A| \to |B|$ such that:

- any constant of A is mapped to the corresponding constant of B;
- if F is a function symbol of arity a then

$$F^A(u_1, u_2, \ldots, u_a) = v \Rightarrow F^B(\varphi(u_1), \varphi(u_2), \ldots, \varphi(u_a)) = \varphi(v),$$
$$\text{for all } u_1, u_2, \ldots, u_a, v \in |A|;$$

- if R is a relation symbol of arity b then

$$R^A(u_1, u_2, \ldots, u_b) \text{ holds } \Rightarrow R^B(\varphi(u_1), \varphi(u_2), \ldots, \varphi(u_b)) \text{ holds},$$
$$\text{for all } u_1, u_2, \ldots, u_b \in |A|.$$

If there exists a homomorphism of A to B then we write $A \to B$.

Let C be a class of finite structures. The *uniform constraint satisfaction problem* CSP_C has: as its instances pairs (A, B) of structures from C over the same signature; and as its yes-instances those instances (A, B) for which there exists a homomorphism of A to B. If all structures in C are over the same signature and $T \in C$ then the *non-uniform constraint satisfaction problem* $\mathrm{CSP}_C(T)$ has: as its instances structures $A \in C$; and as its yes-instances those instances A for which there exists a homomorphism of A to T. We should add that the individual tractability, for example, of an infinite collection of non-uniform constraint satisfaction problems $\{\mathrm{CSP}_C(T) : T \in C\}$ does not automatically yield the tractability of the uniform constraint satisfaction problem CSP_C; for it may be the case that the size of the template, whilst a constant in a non-uniform problem, might play an exponential role in some time bound (see [35] for an examination of this issue).

We shall be involved with problems solvable in L and complete for this complexity class. As regards completeness, the notion of reduction we work with comes from finite model theory and is the quantifier-free projection. Before giving a definition of a quantifier-free projection, we present an example of a quantifier-free projection from one problem to another. As it turns out, we will need this actual reduction later on. The reader is referred to, for example, [26, 27, 54] for more on quantifier-free projections and other logical reductions, and their relevance as low-resource reductions: we only sketch the issues here.

Let the signature σ_{2++} consist of the binary relation symbol E and the two constant symbols C and D. We can think of a σ_{2++}-structure as a digraph, possibly with self-loops, with two designated vertices (which may be identical). The problem $\mathrm{DTC}_{0,1}$ has: as its instances the class of σ_{2++}-structures which, when considered as digraphs with self-loops, have the property that every vertex has degree *at most* 1; and as its yes-instances those instances with the property that there is a path in the digraph from the vertex C to the vertex D. The problem DTC_1 has: as its instances the class of σ_{2++}-structures which, when considered as digraphs with self-loops, have the property that every vertex has degree *exactly* 1; and as its yes-instances those instances with the property that there is a path in the digraph from the vertex C to the vertex D.

We shall derive four quantifier-free formulae over the signature σ_{2++} and we shall use our formulae to describe, given an instance A of $\mathrm{DTC}_{0,1}$, an instance $\rho(A)$ of DTC_1: the first formula will define the vertex set of $\rho(A)$; the second formula will describe the edge relation of our instance; and the third and fourth formulae will describe the source and target vertices.

The domain of $\rho(A)$ is $|A|^2$. We assume that, regardless of the signature, we always have a binary relation *succ* at our disposal that is always interpreted as a *successor relation* on the domain of any structure, *i.e.*, as a relation of the form

$$\{(i_j, i_{j+1}) : j = 0, 1, \ldots, n-1\},$$

when the domain of a structure of size n is $\{i_0, i_1, \ldots, i_{n-1}\}$, and also two constant symbols, 0 and *max*, that are always interpreted as the least and greatest elements, respectively, of the successor relation *succ* (more of this successor relation later). Let us suppose for simplicity that the elements of $|A|$ are $\{0, 1, \ldots, n-1\}$ and abbreviate '$succ(u, v)$' by '$v = u + 1$'. The vertices of $\{(u, v) : v = 0, 1, \ldots, n-1\}$ will form a path $(u, 0), (u, 1), \ldots, (u, n-1)$ in $\rho(A)$, with a self-loop at $(u, n-1)$, except that:

- if (u, v) is an edge of E^A, where $u \neq v$, then there is no edge $((u, v), (u, v+1))$ in $\rho(A)$ nor self-loop $((u, n-1), (u, n-1))$, if $v = n-1$, but there is an edge $((u, v), (v, 0))$ in $\rho(A)$; and
- if (u, u) is an edge of E^A then there is no edge $((u, u), (u, u+1))$ in $\rho(A)$ but there is a self-loop $((u, u), (u, u))$.

The source vertex of $\rho(A)$ is the vertex $(C^A, 0)$ and the target vertex is $(D^A, 0)$. It is easy to see that an instance A of $\mathrm{DTC}_{0,1}$ is a yes-instance if, and only if, the instance $\rho(A)$ is a yes-instance of DTC_1 (as whenever $u \neq v$, there is an edge (u, v) in E^A if, and only if, there is a path from vertex $(u, 0)$ to vertex $(v, 0)$ in $\rho(A)$).

The formula ψ_0, ψ_E, ψ_C and ψ_D describing the above construction are as follows.

$$
\begin{aligned}
\psi_0(x_1, x_2) &\equiv x_1 = x_1 \\
\psi_E(x_1, x_2, y_1, y_2) &\equiv (x_1 = y_1 \wedge y_2 = x_2 + 1 \wedge \neg E(x_1, x_2)) \\
&\quad \vee (x_1 = y_1 \wedge x_2 = y_2 = max \wedge \neg E(x_1, max)) \\
&\quad \vee (x_1 \neq x_2 \wedge y_1 = x_2 \wedge y_2 = 0 \wedge E(x_1, x_2)) \\
&\quad \vee (x_1 = x_2 \wedge x_1 = y_1 \wedge x_2 = y_2 \wedge E(x_1, x_2)) \\
\psi_C(x_1, x_2) &\equiv x_1 = C \wedge x_2 = 0 \\
\psi_D(x_1, x_2) &\equiv x_1 = D \wedge x_2 = 0
\end{aligned}
$$

The formula $\psi_0(x_1, x_2)$ tells us that the vertex set of $\rho(A)$ is the whole of $|A|^2$ (it might have restricted the vertex set to be some appropriately defined subset of $|A|^2$ but in this case didn't); and ψ_E, ψ_C and ψ_D describe the edge relation, the source vertex and the target vertex of $\rho(A)$, respectively.

So, we can say that DTC_1 is a *quantifier-free first-order translation* of $\mathrm{DTC}_{0,1}$ (as the defining formulae are quantifier-free first-order); but we can actually say more. Note that the above formula ψ_E is of the following form.

$$\bigvee\{(\alpha_i \wedge \beta_i) : i = 1,2,\ldots,k\},$$

for some $k \geq 1$, where:

- each α_i is a conjunction of atoms and negated atoms not involving any relation or function symbols of the underlying signature (σ_{2++} in the illustration above);
- the α_i's are *mutually exclusive*, *i.e.*, for any valuation on the variables (and constants) of any α_i and α_j, where $i \neq j$, it is not the case that both α_i and α_j hold;
- each β_i is an atom or a negated atom (over the underlying signature).

Indeed, the formulae ψ_C and ψ_D are trivially of this form too; and, furthermore, ψ_0 is a quantifier-free first-order formula not invloving any relation or function symbols of the underlying signature. Hence, there is a *quantifier-free projection* from the problem $\mathrm{DTC}_{0,1}$ to the problem DTC_1. It was proven in [54] that $\mathrm{DTC}_{0,1}$ is complete for L via quantifier-free projections; and consequently DTC_1 is also complete for L via quantifier-free projections.

Quantifier-free projections are so called because the defining formulae are quantifier-free first-order and any 'bit' of a target instance, *e.g.*, edge of $\rho(A)$, above, depends only upon at most one 'bit' of the source structure, *e.g.*, edge of A, above. They are extremely restricted reductions between problems and can easily be translated into other restricted circuit-based or model-based reductions, *e.g.*, logtime-uniform NC^1-reductions, used in complexity theory (see [27]). The (built-in) successor relation and the two associated constants give us an ordering of our data which often enables us to model machine-based computations where all data (such as input strings and instantaneous descriptions) is ordered.

We have one final remark: in our example above, we used quantifier-free first-order formulae to describe an edge relation and two constants. We can equally well use such formulae to describe functions by treating an m-ary function F as an

$(m+1)$-ary relation R_F where for any elements $u_1, u_2, \ldots, u_m$, there exists exactly one v such that $R_F(u_1, u_2, \ldots, u_m, v)$ holds (constants, *i.e.*, 0-ary functions, are described above in this way).

7.2 One unary function

Let λ_1 be the signature consisting of one unary function symbol f. The decision problem *Hom-Alg*$_1$ has as its instances pairs (A,B) of λ_1-structures; and as its yes-instances instances (A,B) for which $A \to B$ (and so Hom-Alg$_1$ is the problem CSP$_C$, where C is the class of all λ_1-structures). The size of an instance is the maximum of the sizes of A and B. We assume that a unary function f is encoded for input to some Turing machine as a list of pairs of the form $(u, f(u))$.

Let A be a λ_1-structure. The *graph* of A is the σ_2-structure $\dot{A} = \langle |A|, E \rangle$, where $E(u,v)$ holds if, and only if, $f(u) = v$ (note that it may be the case that $E(u,u)$ holds in $\dot{A}$). The proof of the following lemma is trivial.

Lemma 7.1 Let A and B be λ_1-structures. Then $A \to B$ if, and only if, $\dot{A} \to \dot{B}$.

Proposition 7.2 The problem Hom-Alg$_1$ is in L.

PROOF.By Lemma 7.1, we can assume that we are given pairs of graphs of unary functions as instances rather than pairs of unary functions.

Let $\dot{A}$ be the graph of some unary function A. Then in general $\dot{A}$ consists of a collection of connected components where each component is an directed cycle, which may have any length greater than 0 (and so may be a self-loop), some of whose vertices are roots of in-trees. These components can be visualised as in Figure 7.2. We call these components cycles with pendant in-trees. We define the length of a cycle with pendant in-trees as the length of the directed cycle.

Let $(\dot{A},\dot{B})$ be a pair of graphs of unary functions where $\max\{|\dot{A}|, |\dot{B}|\}$ is n. Suppose that there is a homomorphism taking some connected component C of $\dot{A}$ to a connected component D of $\dot{B}$. If C is a cycle with pendant in-trees of length c then D must be a cycle with pendant in-trees of length d where d divides c. Furthermore, if C and D are cycles with pendant in-trees of lengths c and d,

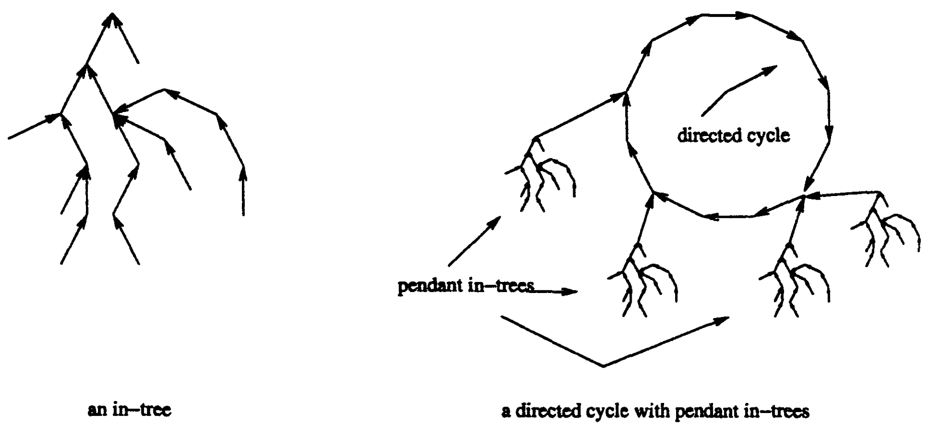

Figure 7.1: The components of the graph of a unary function.

respectively, and d divides c then there is a homomorphism of C to D. Hence, the following is a necessary and sufficient condition for a homomorphism of $\dot{A}$ to $\dot{B}$ to exist.

- For every cycle with pendant in-trees of length c in $\dot{A}$, there must exist a cycle with pendant in-trees of length d in $\dot{B}$ where d divides c.

This condition can easily be verified using $O(\log n)$ space (in n). For example, we can ascertain whether a vertex u lies on the cycle of a cycle with pendant in-trees in $\dot{A}$ by walking along the path emanating from u and stopping after n moves (when u doesn't lie on a cycle) or after we have returned to u (when u does lie on a cycle). By counting as we walk, we obtain the length of the cycle (if u lies on a cycle). We can then work through the vertices of $\dot{B}$ checking to see whether they lie on the cycle of a cycle with pendant in-trees in $\dot{B}$; and if a vertex does lie on the cycle of a cycle with pendant in-trees then we can check whether the length of this cycle divides c. Hence, the problem Hom-Alg$_1$ $\in$ L. □

Proposition 7.3 The problem Hom-Alg$_1$ is L-hard *(*via quantifier-free projections*)*.

PROOF. Let A be an instance of DTC$_1$. Define the unary function f_A as follows. The domain of f_A is $|A|^2 \times \{0,1\}$ and:

- if $C = D$ then:

  - $f((u,v,b)) = (C,C,0)$, for all $(u,v,b) \in |A|^2 \times \{0,1\}$;

- if $C \neq D$ then:
  - if $(u,v) \in E$ where $u \neq D$, $v \neq C$ and $u \neq v$ then $f_A((u,u,0)) = (u,v,0)$ and $f_A((u,v,0)) = (v,v,0)$
  - if $(u,u) \in E$ where $u \neq D$ then $f_A((u,u,0)) = (u,u,1)$ and $f_A((u,u,1)) = (u,u,0)$
  - $f_A((D,D,0)) = (C,C,0)$
  - for any element $(u,v,b) \in |A|^2 \times \{0,1\} \setminus \{(D,C,0)\}$ for which $f_A((u,v,b))$ is still undefined, define $f_A((u,v,b)) = (D,C,0)$, and define $f_A((D,C,0)) = (D,C,1)$.

Essentially, apart from the trivial case where $C = D$, the graph of f_A is obtained from the digraph whose edge relation is E as follows:

- take a copy of the digraph (with self-loops) whose edge relation is E, and replace any edge emanating from vertex D with the edge (D,C); and

- replace every edge (u,v), apart from the edge (D,C), by a pair of edges $(u,e_{u,v})$ and $(e_{u,v},v)$, where $e_{u,v}$ is a new vertex.

Other vertices are actually introduced in the formal constructive process (defined above), with two of these vertices being $(D,C,0)$ and $(D,C,1)$. The construction is completed by introducing edges from all vertices, apart from $(D,C,0)$, to $(D,C,0)$; and also an edge from $(D,C,0)$ to $(D,C,1)$. Now define g_A to have domain $\{0,1\}$ and to be such that $g_A(0) = 1$ and $g_A(1) = 0$. We claim that $A \in \mathrm{DTC}_1$ if, and only if, $(f_A, g_A) \notin$ Hom-Alg$_1$.

The trivial case is straightforward (note that if the graph of f_A has a self-loop then there is not a homomorphism of f_A to g_A): so suppose henceforth that $C \neq D$. Suppose that there is a path in the digraph whose edge relation is E from vertex C to vertex D. Then in the graph of f_A, there is a odd length cycle with pendant in-trees of length greater than 1. Hence, there is no homomorphism of f_A to g_A.

Suppose that there is not a path in the digraph whose edge relation is E from vertex C to vertex D. Then all components of the graph of f_A are even length cycles with pendant in-trees. Hence, there is a homomorphism of f_A to g_A.

The construction of the unary functions f_A and g_A from A can easily be described by quantifier-free projections (see, *e.g.*, [54] for concrete illustrations of logical formulae describing translations between problems) and so the result follows as DTC_1 is complete for L via quantifier-free projections (note that there are quantifier-free projections describing both the λ_1-structures f_A and g_A). $\square$

The following is now immediate from Propositions 7.2 and 7.3.

Theorem 7.4 The problem Hom-Alg_1 is L-complete *(*via quantifier-free projections*)*.

The problem Hom-Alg_1 is uniform in the sense that any unary function can appear as either the first or second component of an instance. We obtain non-uniform versions of Hom-Alg_1 by fixing the second component. The problem *$Hom\text{-}Alg_1(T)$*, for some λ_1-structure T, consists of all those λ_1-structures A for which $A \to T$ (and so $\text{Hom-Alg}_1(T)$ is the problem $\mathrm{CSP}_\mathrm{C}(T)$, where C is the class of all λ_1-structures).

The following is immediate from Propositions 7.2 and 7.3.

Theorem 7.5 Let T be the λ_1-structure corresponding to the unary function g whose domain is $\{0,1\}$ and $g(0) = 1$ and $g(1) = 0$. The problem $\text{Hom-Alg}_1(T)$ is L-complete *(*via quantifier-free projections*)*.

Hence, not only is the uniform problem Hom-Alg_1 L-complete, there are also non-uniform problems $\text{Hom-Alg}_1(T)$ that are L-complete (moreover, even when T has only two elements).

Actually, we can say more about non-uniform problems of the form $\text{Hom-Alg}_1(T)$. Whilst the proof of Proposition 7.3 is such that the template has a graph that is a cycle of length 2, we can actually replace this template with any λ_1-structure T so long as the graph of T has a cycle of pendant in-trees of length at least 2 as follows. Suppose that $\dot{T}$ has cycles of pendant in-trees of lengths $d_1, d_2, \ldots, d_k$, for some $k > 0$. Adopting the terminology of the proof of Proposition 7.3 and with reference to this proof, in our construction process when we replace an edge of the graph of f_A with a path of 2 edges, instead we replace the edge with a path of $d_1 d_2 \ldots d_k$ edges. So, if there is a path in the digraph whose edge relation is E from vertex C to vertex D then the graph of f_A has a cycle with pendant in-trees of length $c.d_1d_2\ldots d_k+1$, for some $c \geq 1$, and all other cycles with pendant in-trees have length divisible by $d_1d_2\ldots d_k$ (if there are any); and if there is no such path then the graph of f_A is such that every cycle with pendant in-trees has length divisible by $d_1d_2\ldots d_k$. Hence, we obtain the following corollary.

Corollary 7.6 Let T be any λ_1-structure without a fixed point. Then Hom-Alg$_1(T)$ is L-complete (via quantifier-free projections).

Trivially, if the λ_1-structure T has a fixed point then Hom-Alg$_1(T)$ consists of every λ_1-structure and is identical to the problem Hom-Alg$_1(F_0)$, where F_0 is the function whose domain has one element. Note that whereas the 'trivial' cases of Hom-Alg$_1(T)$ are identical to Hom-Alg$_1(F_0)$, so there is an analogous remark to be made about Hell and Nešetřil's dichotomy: the 'trivial' cases, here the cases where the problem is solvable in polynomial-time, are identical to the case where the template graph consists of a solitary edge.

7.3 Two unary functions

Let λ_2 be the signature consisting of the two unary function symbols f and g. The decision problem *Hom-Alg$_2$* has as its instances pairs (A,B) of λ_2-structures; and as its yes-instances instances (A,B) for which $A \to B$. As before, the size of an instance is the maximum of the sizes of A and B.

Let $\sigma_2 = \langle E \rangle$, where E is a binary relation symbol. We shall begin by explaining how we can transform any σ_2-structure G, which we regard as a simple undirected graph via 'there is an edge (u,v), for $u \neq v$, if, and only if, either $E(u,v)$ or $E(v,u)$ holds', into a λ_2-structure. The λ_2-structure $\lambda_2(G)$ is defined as follows.

- The domain of $\lambda_2(G)$ consists of

$$\{u : u \in |G|\} \cup \{u' : u \in |G|\} \cup \{e_{u,v}, e_{v,u} : E(u,v) \text{ or } E(v,u) \text{ holds and } u \neq v\}.$$

  Furthermore, we call the elements of $\{u : u \in |G|\}$ *straight elements*, the elements of $\{u' : u \in |G|\}$ *prime elements* and the elements of $\{e_{u,v}, e_{v,u} : E(u,v) \text{ or } E(v,u) \text{ holds and } u \neq v\}$ *edge elements*.

- For any straight element u, $f(u) = u$ and $g(u) = u'$; for any prime element u', $f(u') = u$ and $g(u') = u'$; and for any edge element $e_{u,v}$, $f(e_{u,v}) = v'$ and $g(e_{u,v}) = u$.

The above construction can be visualized in Figure 7.2.

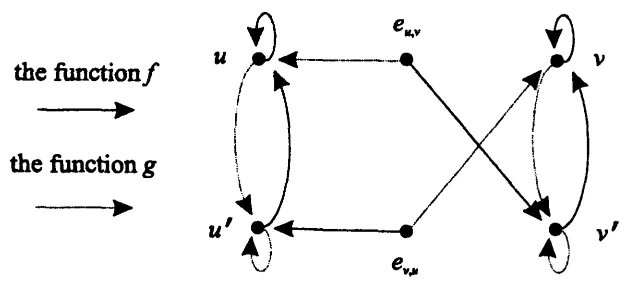

Figure 7.2: The construction of $\lambda_2(G)$ from G.

Lemma 7.7 Let G and H be undirected graphs. Then $G \to H$ if, and only if, $\lambda_2(G) \to \lambda_2(H)$.

PROOF. Suppose that $\psi : G \to H$ is a homomorphism. Define the map $\varphi : |\lambda_2(G)| \to |\lambda_2(H)|$ as follows:

- if u is a straight vertex of $\lambda_2(G)$ then $\varphi(u)$ is the straight vertex $\psi(u)$ of $\lambda_2(H)$;
- if u' is a prime vertex of $\lambda_2(G)$ then $\varphi(u')$ is the prime vertex $\psi(u)'$ of $\lambda_2(H)$;
- if $e_{u,v}$ is an edge vertex of $\lambda_2(G)$ then $\varphi(e_{u,v})$ is the edge vertex $e_{\psi(u),\psi(v)}$ of $\lambda_2(H)$.

That φ is a homomorphism is straightforward: for example, $f(e_{u,v}) = v'$ in $\lambda_2(G)$ and $f(\varphi(e_{u,v})) = f(e_{\psi(u),\psi(v)}) = \psi(v)' = \varphi(v)'$ in H.

Suppose that $\varphi : \lambda_2(G) \to \lambda_2(H)$ is a homomorphism. It is immediate that for any straight or prime vertex u, $\varphi(u)$ cannot be an edge vertex (as f maps every straight or prime vertex to itself but not so an edge vertex). Hence, define the map $\psi : |G| \to |H|$ as follows:

$\psi(u) = v$ if, and only if, φ maps the straight vertex u of $\lambda_2(G)$ to either the straight vertex v or the prime vertex $\lambda_2(v)'$ of $\lambda_2(H)$.

Suppose that (u, v) is an edge of G. Then $e_{u,v}$ and $e_{v,u}$ are vertices of $\lambda_2(G)$ and $\psi(e_{u,v}) = e_{a,b}$, for some vertex $e_{a,b}$ of $\lambda_2(H)$ where (a, b) is an edge of G. In $\lambda_2(G)$, $u = g(e_{u,v})$ and so:

$$\varphi(u) = \varphi(g(e_{u,v})) = g(\varphi(eu, v)) = g(e_{a,b}) = b.$$

Also, $v' = f(e_{u,v})$ in $\lambda_2(G)$, and so:

$$\varphi(v') = \varphi(f(e_{u,v})) = f(\varphi(e_{u,v})) = f(e_{a,b}) = b',$$

with $f(\varphi(v')) = f(b')$, *i.e.*, $f(\varphi(v)) = f(b)$, *i.e.*, $\psi(v) = b$. Hence, ψ is a homomorphism. □

Theorem 7.8 The problem Hom-Alg$_2$ is NP-complete.

PROOF. Let 3COL be the problem, over σ_2, whose instances are undirected graphs and whose yes-instances are instances that can be properly 3-coloured (this problem has long been know to be NP-complete [20]). The problem 3COL can be reformulated as those undirected graphs for which there is a homomorphism to the complete graph on 3 vertices. The result follows by Lemma 7.7. □

As before, we obtain non-uniform versions of Hom-Alg$_2$ by fixing the second component. The problem *Hom-Alg*$_2(T)$, for some λ_2-structure T, consists of all those λ_2-structures A for which $A \to T$.

Theorem 7.9 Let T be the λ_2-structure of the form $\lambda_2(G)$, where G is the complete undirected graph on 3 vertices. The problem Hom-Alg$_2(T)$ is NP-complete.

Hence, not only is the uniform problem Hom-Alg_2 NP-complete, there are also non-uniform problems Hom-$\text{Alg}_2(T)$ that are NP-complete. However, we have as yet been unable to obtain a classification of the non-uniform constraint satisfaction problems of the form Hom-$\text{Alg}_2(T)$. Our only comment is that we could have taken any NP-complete graph-problem that can be formulated as a non-uniform constraint satisfaction problem, and not just 3COL, to obtain an NP-complete problem of the form Hom-$\text{Alg}_2(T)$. Unfortunately, there are many λ_2-structures which are not the images of undirected graphs (under the map λ_2, above).

We have recently extended these results in a joint work with Tòmas Feder and Iain Stewart: the former had contemporary and independent related results for tractability of some related digraphs homomorphisms problems (*cf.* [14]).

Chapter 8

Conclusion

My main contribution is a theorem that characterises precisely the borderline between two classes of combinatorial problems: that of the constraint satisfaction problems (CSP) and that of the forbidden patterns problems (FP). The latter correspond exactly to Feder and Vardi's logic MMSNP, who proved in [16] that MMSNP and CSP are computationally equivalent. However, it is important to note that this equivalence relies on randomised polynomial-time reductions and that it remains open whether these reductions may be derandomised. These authors have also proved that CSP is strictly included within MMSNP; since their proof relies on counting arguments, we proved this for the same examples (and also for further examples) in a constructive manner in [43]. Initially, I wanted to show that it is not possible to derandomise Feder and Vardi's randomised reductions (or, more precisely, that there does not exist "fine" and "monotone enough" deterministic reductions: by fine, I mean a logical reduction like a *FO-reduction* and by "monotone enough", I mean a restriction of the former that would take into account the fact that CSP and FP are closed under inverse homomorphism). This led to my attempt of characterising those sentences of MMSNP that define problems that are not in CSP. The core of this thesis answers the latter whereas the former remains open. However, I hope that the reader has been convinced that this is slightly more than an intermediate result.

On the one hand, the proof of this main result is interesting in itself. Indeed, it presents an original generalisation of the notions of a structure and that of a homomorphism, namely, that of a representation and that of a recolouring, respectively. This proof adapts some techniques used by Feder and Vardi, which together with some novel techniques, yield a normal form for problems in FP that can be effectively computed. Furthermore, from this normal form, one can easily decide whether the corresponding problem is in CSP or not. Finally, if it is not the case then the proof gives an effective and generic method to construct a family of witnesses that prove that the problem currently investigated can not be a CSP, since each finite structure is shown to be unsuitable as a template for the problem by some witness.

On the other hand, this main result generalises a result by Tardif and Nešetřil (*cf.* [45]). These authors used an elegant correspondence between duality and density to derive a characterisation for duality pairs (which correspond in our settings to very restricted forbidden patterns problems: they have a single colour and also a single forbidden pattern). The generalisation of structures and homomorphisms mentioned above seems even more pertinent, since I successfully extend this correspondence in the case of representations and recolourings (this is achieved by highlighting the structure of a Heyting algebra for cores of representations).

In [42], we turned to questions only loosely related to the above. We noticed indeed that, whereas numerous results concerning the complexity of CSP were known in the case of structures, there seemed to be none in the case of functions. We concentrated on a very restricted case, that of unary algebras (that is structures for which the signature consists of unary functions only). We proved that in the case of a signature that consists of two unary symbols, the uniform problem is NP-complete (here, "uniform" means that an instance consists of a pair of algebra; and, the question is to decide whether there exists a homomorphism from the first algebra to the second). Moreover, in the case of a single unary symbol, we highlighted a rather severe dichotomy: the non-uniform problems are either trivial or L-complete (by opposition, "non-uniform" means that an instance consists of a single algebra; and, the question is to decide whether there exists a homomorphism from this algebra to some fixed algebra, called the template of the non-uniform problem). This gives the first known result of constraint satisfaction problems that are L-complete. We have proved more recently in [15] that it is at least as hard to prove a dichotomy result in the restricted case of two unary function symbols than it is in the classical case (*i.e.* for CSP).

The work presented in this thesis has inspired me some ideas and problems which I hope to solve in the near future.

First, I have not completely renounced to prove that, in some sense, the randomised reductions of Feder and Vardi's theorem can not be derandomised. For this, I plan to use the main result of this thesis as an intermediate result and to restrict myself to meaningful reductions in the context of CSP and FP.

Secondly, an interesting theoretic exercise consists in extrapolating the properties of CSP to build a hierarchy *à la* CSP above the complexity class NP. The mechanism to which I am thinking has been briefly sketched on the end of Chapter 5. The first level of the hierarchy has two layers: the first layer consists of the (non-uniform) constraint satisfaction problems (CSP) and the second layer consists of the forbidden patterns problems (FP). At the second level, the first layer consists of the non-uniform recolouring problem (that is, if we accept representation as reasonable encodings) and the second layer consists of the corresponding "forbidden patterns problems". The former correspond to the non-uniform version of the containment problem for forbidden patterns problems, under the assumption that Conjecture 5.17 holds, and the latter are to the former what FP is to CSP (and being less natural are not so easily motivated). Furthermore, I have the intuition that the notion of an exponential of representations, introduced in Chapter 5, may be used to prove Conjecture 5.17. Indeed, notice that a recolouring is in fact context-free: a colour of the first representation is mapped to another colour, no matter where it occurs in a coloured structure. We can imagine *contextual recolouring* that would recolour an element of a coloured structure differently according to some local information, given by the *context* of this element, which could be modelled by some bounded pattern that occurs around this element. Recall that the exponential $\mathfrak{R}_2^{\mathfrak{R}_1}$ is a representation with the set of maps from the colour set of $\mathfrak{R}_1$ to $\mathfrak{R}_2$ as its colour set and that its forbidden patterns have patterns that are patterns from forbidden patterns of $\mathfrak{R}_2$. Hence, when looking carefully at the definition of the exponential, we can see that the data carried by one forbidden pattern of the exponential $\mathfrak{R}_2^{\mathfrak{R}_1}$ is a context (given by the pattern) in which the generalised recolouring (given elementwise by a colour map) fails.

Thirdly, I think that some of the techniques, that I adapted from Feder and Vardi's proof to define my normal form, may be found useful to characterise tractable CSP. Indeed, we can consider a family of forbidden patterns problems that, on the one hand, belongs to CSP by the main result, and on the other hand

are easily seen to be in P according to their representation as forbidden patterns problems. For example, consider for the case of oriented graphs (signature σ_2), the family of monochrome forbidden pattern problems "no path of length n". For $n = 1$, the problem is clearly in CSP for which the oriented graph with a single element and no edge is clearly a template. For $n > 1$, we may use a Feder-Vardi transformation, "cutting" at the second element of the path. Thus, we get a representation with two colours and $\bullet$ and two compact forbidden patterns, one is a single edge, whereas the other is a path of length $n-1$, that is: $\{\ ,\bullet\} \longrightarrow$ and $\bullet \longrightarrow \{\ ,\bullet\} \longrightarrow \cdots \longrightarrow \{\ ,\bullet\}$. The former corresponds to the following two patterns: $\longrightarrow$ and $\bullet \longrightarrow$. The above pattern allows to simplify the expression of the latter compact forbidden pattern, and to simply ignore the colour . Finally, we get the representation with the following three forbidden patterns: $\longrightarrow$, $\bullet \longrightarrow$ and $\bullet \longrightarrow \bullet \longrightarrow \cdots \longrightarrow \bullet$.
This simple "chromatic calculus" proves that that T_n, the template of the problem of index n, consists of T_{n-1}, the template of the problem of index $n-1$ (which corresponds above to the $\bullet$ path), to which one new element x is added (which corresponds to the colour) that is linked to every element of T_{n-1} (since there is no loop around this element because of the first forbidden pattern and that there can not be any edge from T_{n-1} to x because of the second). Hence, we can inductively build the templates, since we know that T_1 is the graph:

$$\bullet$$

Hence, T_2 is the graph:

$$\bullet \longrightarrow \bullet$$

Then T_3 is as follows:

$$\begin{tikzcd} \bullet \arrow[r] \arrow[rr, bend right] & \bullet \arrow[r] & \bullet \end{tikzcd}$$

Finally, T_4 is the graph:

$$\begin{tikzcd} \bullet \arrow[r] \arrow[rr, bend right] \arrow[rrr, bend right] & \bullet \arrow[r] \arrow[rr, bend right] & \bullet \arrow[r] & \bullet \end{tikzcd}$$

We get the oriented graphs that corresponds to linear orders. This allows us to give a "fine" characterisation of the corresponding non-uniform constraint satisfaction problems w.r.t. their complexity. Indeed, if we suppose that the queries corresponding to these problems are given by sentences of MMSNP, it is clear that these sentence are first-order sentences. Hence the corresponding problems are in the complexity class L. Since various communities work on tractable classes of CSP, it is most probable that this class is not novel. However, we may be able to use a similar technique to give "good characterisations" (in the sense of [45]) of known tractable classes of CSP, which could lead to more efficient algorithm.

Appendix A

(Descriptive) complexity theory

Further definition and examples can be found in the following complexity theory textbooks [25], [39] or [46]. We refer further to [20] for NP-completeness and to [12] or [39] for descriptive complexity theory.

Complexity classes

The model of a computation used to define complexity classes relevant to this work is that of a (non-)deterministic Turing machine and throughout this work:

- L denotes the class of problem decidable in logarithmic space on a deterministic Turing machine;
- NL denotes the class of problem decidable in logarithmic space on a non-deterministic Turing machine;
- P denotes the class of problem decidable in polynomial time on a deterministic Turing machine; and
- NP the class of problem decidable in polynomial time on a non-deterministic Turing machine.

Problem

A *problem* is a class of structures that is closed under isomorphism.

Logics

- FO denotes first order logic.
- ESO denotes existential second order logic.

The definition of the above logics can be found in [12].

- ESO-Krom denotes a fragment of ESO.
- ESO-Horn denotes another fragment of ESO.

The above are defined in [21].

Reductions

Let **P** and **Q** be two problems and let r be a function from the set of instances of **P** to the set of instances of **Q**.
We say that r is a *polynomial-time reduction* from **P** to **Q** whenever:

- r can be computed in polynomial time; and
- for every instance A of **P**,

$$A \in \mathbf{P} \iff r(A) \in \mathbf{Q}.$$

We say that r is a *randomized polynomial-time reduction* from **P** to **Q** whenever:

- r can be computed in polynomial time; and
- for every instance A of **P**, the probability that

$$A \in \mathbf{P} \iff r(A) \in \mathbf{Q}$$

  is high (say strictly greater than $\frac{1}{2}$).

Interpretations and logical reductions

In the following $\mathcal{L}$ denotes some logic (typically some fragment of first-order logic). Let σ and τ be two relational signatures where τ consists of n relation symbols R_i of respective arity r_i ($1 \leq i \leq n$). Let k be a positive integer. Let $\varphi_1, \dots, \varphi_n$ be formulae from $\mathcal{L}(\sigma)$, where the free variables of φ_i are a subset of $\{x_1, \dots, x_{k.r_i}\}$.
$\Pi = (\varphi_1, \varphi_2, \dots, \varphi_n)$ induces a mapping from $STRUC(\sigma)$ to $STRUC(\tau)$ as follows. Let $A \in STRUC(\sigma)$. Then, the structure $\Pi(A) = B$ is the τ-structure with:

- universe $|B| := |A|^k$; and
- for every $1 \leq i \leq n$ and any $(t_1, t_2, \dots, t_{r_i}) \in |B|^{r_i}$, where:

  $$t_1 = (u_1, u_2, \dots, u_k), t_2 = (u_{k+1}, u_{k+2}, \dots, u_{2k}), \dots, t_{r_i} = (u_{kr_i-k+1}, \dots, u_{kr_i})$$

  $R_i^B(t_1, t_2, \dots, t_{r_i})$ holds if, and only if, $A \models \varphi_i(\bar{x}/\bar{u})$.

Π is called a *$\mathcal{L}$-interpretation of σ in τ of width k.*
Let $\mathbf{P} \subseteq STRUC(\sigma)$ and $\mathbf{Q} \subseteq STRUC(\tau)$ be two problems. We say that the problem $\mathbf{P}$ is *$\mathcal{L}$-reducible* to $\mathbf{Q}$ ($\mathbf{P} \leq_{\mathcal{L}} \mathbf{Q}$, for short) whenever:

- there exists a $\mathcal{L}$-interpretation Π of σ in τ; and
- for any σ-structures A,

$$A \in \mathbf{P} \iff \Pi(A) \in \mathbf{Q}.$$

If $\mathcal{L} = \text{FO}$ then we speak of a *FO-reduction.* When the FO-interpretation Π satisfies the following *projection condition,* we speak of *FO-projection* or *fop* for short. Every formula is of the form:

$$\alpha_1 \vee (\alpha_2 \wedge \ell_2) \vee \ldots \vee (\alpha_e \wedge \ell_e)$$

where:

- every α_i is free of any occurrence of relational symbols from the signature σ;

- the α_i's are mutually exclusive; and
- every ℓ_i consists of a single literal.

If, moreover, the formulas are quantifier-free, that is every α_i is quantifier-free, then we say that Π is a *quantifier-free projection* or *qfp* for short. Moreover, as usual with qfps, except if otherwise stated, we allow a built-in successor function Succ and two constants 0 and $\mathtt{max}$.

Appendix B

Category theory

For more detail and examples, we refer to [38].

Categories

A *diagram scheme* consists of a set O of *objects* and a set A of *arrows* together with two functions:

$$A \underset{\text{cod}}{\overset{\text{dom}}{\rightrightarrows}} O$$

For $a, b \in O$ and $f \in A$ such that dom $f = a$ and cod $f = b$, we write:

$$a \xrightarrow{f} b$$

In this graph, the set of composable pairs of arrows is the set:

$$A \times_O A = \{< g, f > | g, f \in A \text{ and dom } g = \text{cod } f\}$$

A *category* is a diagram scheme with two additional functions

$$\begin{array}{rcl} \text{id} : O & \longrightarrow & A \\ c & \mapsto & \text{id}_c \end{array} \qquad \begin{array}{rcl} \circ : A \times_O A & \longrightarrow & A \\ < g, f > & \mapsto & g \circ f \end{array}$$

called *identity* and *composition*, such that:

- for all objects $a \in O$ and all composable pairs of arrows $< g, f > \in A \times_O A$,

$$\text{dom }(\text{id }(a)) = \text{cod }(\text{id }(a)), \text{dom }(g \circ f) = \text{dom } f, \text{cod }(g \circ f) = \text{cod } g$$

  and;

- the composition is associative and the identity law holds; that is, for all objects a, b, c, d and arrows f, g, h, if $a \xrightarrow{f} b \xrightarrow{g} c \xrightarrow{h} d$ then:

$$f \circ (g \circ h) = (f \circ g) \circ h$$

$$\text{id}_b \circ f = f \text{ and } g \circ \text{id}_b = g$$

From now on, we write simply $a \in C$ for "a an object in C" and $f \in C$ for "f an arrow in C". We may also say *morphism* instead of "arrows".

EXAMPLE.

1. **Set** is the category whose objects are sets, and whose arrows are functions.
2. $STRUC(\sigma)$ is the category whose objects are σ-structures, and whose arrows are homomorphisms.
3. A partial order is a category (with the property that there is exactly one arrow between any two objects; and, that there is no cycle apart from self-loops when viewed as a directed graph).

▲

Let C be a category.

Duality

A very important feature of category theory is that of *duality*: given some notion, the dual notion is obtained by reversing all arrows. Indeed, a statement holds if, and only if, its dual holds.

Isomorphisms

An arrow $a \xrightarrow{e} b$ is *invertible* in C if there is an arrow $b \xrightarrow{e'} a$ in C with $e \circ e' = \mathrm{id}_a$ and $e' \circ e = \mathrm{id}_b$. If such an e' exists, it is unique, and is written $e' = e^{-1}$. Two objects a and b in the category C are *isomorphic* if there is an invertible arrow (an *isomorphism*) $a \xrightarrow{e} b$; we write $a \approx b$. The relation of isomorphism of objects is an equivalence relation.

Monomorphism

An arrow $a \xrightarrow{i} b$ is *monic* (or left cancelable) if for any two parallel arrows $d \xrightarrow{f_1} a$ and $d \xrightarrow{f_2} a$, the equality $i \circ f_1 = i \circ f_2$ implies $f_1 = f_2$. We also say *monomorphism* for "monic arrow" and write $a \xhookrightarrow{i} b$.

Epimorphism

An arrow $a \xrightarrow{s} b$ is *epi* (or right cancelable) if for any two parallel arrows $b \xrightarrow{g_1} c$ and $b \xrightarrow{g_2} c$, the equality $g_1 \circ s = g_2 \circ s$ implies $g_1 = g_2$. We also say *epimorphism* for "epi arrow" and write $a \overset{s}{\twoheadrightarrow} b$. Note that this is the dual notion of the above.

EXAMPLE.

1. In **Set**, the above three notions correspond respectively to the notion of a bijective, an injective and a surjective function.
2. In $STRUC(\sigma)$, these notions correspond respectively to a (structure) isomorphism, an embedding and a surjective homomorphism.
3. In a partial order, the only isomorphisms are the identity arrows (equality) and the fact that there exists a unique arrow between any two objects implies that every arrow is mono and every arrow is epi.

▲

Retraction

For an arrow $a \xrightarrow{h} b$, a *left inverse* is an arrow $a \xrightarrow{l} b$ with $l \circ h = \mathrm{id}_a$. A left inverse (which is usually not unique) is also called a *retraction* of h. Note that it follows that h is monic. Moreover any left inverse of h is epi.

EXAMPLE. Consider the category $STRUC(\sigma)$. Let $a \xrightarrow{l} b$ be a retraction of $a \xrightarrow{h} b$. Since h is an embedding, we can see a as a (not necessarily induced) substructure of b such that b can be mapped homomorphically onto a via l, leaving the vertices of a fixed. ▲

Terminal object

An object 1 is *terminal* in C if from each object $a \in C$ there is exactly one arrow $a \longrightarrow 1$. If 1 is terminal, the only arrow $1 \longrightarrow 1$ is the identity id_1 , and any two terminal objects of C are isomorphic in C.

Initial object

It is the dual of a terminal object. An object 0 is *initial* in C if to each object $a \in C$ there is exactly one arrow $0 \longrightarrow a$. If 0 is initial, the only arrow $0 \longrightarrow 0$ is the identity id_1 , and any two initial objects of C are isomorphic in C.

Equalizer

$d \xrightarrow{e} b$ forms an *equalizer* of $b \underset{g}{\overset{f}{\rightrightarrows}} c$ if $f \circ e = g \circ e$ and for any $c \xrightarrow{h} b$ such that $f \circ h = g \circ h$ there exists a unique h' such that $e \circ h' = h$.

EXAMPLE. In **Set**, take $d := \{x \in c \text{ such that } f(x) = g(x)\}$ and take for e the function that sends $x \in d$ to $x \in b$. ▲

Product

Let $a, b \in C$. An object $a \times b$ together with arrows $a \times b \xrightarrow{\pi_a} a$ and $a \times b \xrightarrow{\pi_b} b$ forms a *product* if for any object c, and any arrows $c \xrightarrow{f} a$ and $c \xrightarrow{g} b$, there exists a unique arrow h such that the following diagram commutes:

$$\begin{tikzcd}
 & c \arrow[dl, "f"'] \arrow[d, dashed, "h"] \arrow[dr, "g"] & \\
a & a \times b \arrow[l, "\pi_a"] \arrow[r, "\pi_b"'] & b
\end{tikzcd}$$

$a \times b$ is called the *product (object)* and the arrows π_a and π_b the *projections*. Note that the product of two objects is unique up to isomorphism.

EXAMPLE.

1. In **Set**, it corresponds to the Cartesian product.

2. In a partial order, it corresponds to the least upper bound.

▲

Coproduct

It is the dual of the above notion. An object $a+b$ together with arrows $a \xrightarrow{\iota_a} a+b$ and $b \xrightarrow{\iota_b} a+b$ forms a *coproduct* if for any object c, and any arrows $a \xrightarrow{f} c$ and $b \xrightarrow{g} c$, there exists a unique arrow h such that the following diagram commutes:

$$\begin{tikzcd}
& c & \\
a \arrow[ur, "f"] \arrow[r, "\iota_a"'] & a+b \arrow[u, dashed, "h"'] & b \arrow[ul, "g"'] \arrow[l, "\iota_b"]
\end{tikzcd}$$

$a+b$ is called the *coproduct (object)* and the arrows ι_a and ι_b the *injections* (though they are not required to be injective functions). Note again that the coproduct of two objects is unique up to isomorphism.

EXAMPLE.

1. In **Set**, it corresponds to the disjoint union of two sets.

2. In a partial order, it corresponds to the greatest lower bound.

▲

Adjoint functors

A *functor* is a morphism of categories: that is, a function that preserves objects, arrows, identity and composition. In detail, for categories C and B a functor $T : C \longrightarrow B$ with domain C and codomain B consists of two suitably related functions: the *object function* T, which assigns to each object $c \in C$ an object Tc of B and the *arrow function* (also written T) which assigns to each arrow $c \xrightarrow{f} c'$ of C an arrow $Tc \xrightarrow{Tf} Tc'$ of B in such a way that:

$$T(\mathrm{id}_c\,) = \mathrm{id}_{T(c)}\ ,\ T(g \circ f) = \mathrm{id}_{T(g) \circ T(f)}$$

the latter whenever the composite $g \circ f$ is defined in C. When the codomain and domain are the same, we speak of an *endofunctor*.

Given two objects a and b in C, we write $\hom(a,b)$ for the *set of arrows from a to b*.

Let C be a category. Let F and G be two endofunctors of C. Let φ be a function which assigns to each pair of objects a and c of C a bijection

$$\varphi_{a,c} : \hom(F(a),c) \longrightarrow \hom(a, G(c))$$
$$\varphi_{a,c} : \hom(F(a),c) \longrightarrow \hom(a, G(c))$$

which is natural in a and c: that is, for all $c \xrightarrow{k} c'$ and all $a \xrightarrow{h} a'$ both the diagrams:

$$\begin{tikzcd}
\hom(Fa,c) \arrow[r, "\varphi_{a,c}"] \arrow[d, "k_\star"'] & \hom(a,Gc) \arrow[d, "(Gk)_\star"] \\
\hom(Fa,c') \arrow[r, "\varphi_{a,c'}"] & \hom(a,Gc')
\end{tikzcd}
\qquad
\begin{tikzcd}
\hom(Fa,c) \arrow[r, "\varphi_{a,c}"] \arrow[d, "(Fh)^\star"'] & \hom(a,Gc) \arrow[d, "h^\star"] \\
\hom(Fa',c) \arrow[r, "\varphi_{a',c}"] & \hom(a',Gc)
\end{tikzcd}$$

will commute. Here $k_\star$ is short for $\hom(F(a),k)$ the operation of composition with k, and $h^\star = \hom(h, Gc)$. Then, we say that F and G are *adjoint functor.*

We call G *the right adjoint* of F (as a right adjoint of G is unique up to natural isomorphism).

EXAMPLE. If C and B are lattices then the pair of adjoint functors F and G are the operators of a Galois connection between those lattices. ▲

Cartesian closed categories

Let C be a category with products. Consider the following endofunctor of C:

$$\begin{array}{rcl} \_\times b : C & \longrightarrow & C \\ a & \longmapsto & a \times b \end{array}$$

If $\_ \times b$ has a right adjoint $\_^b$:

$$\begin{array}{rcl} \_^b : C & \longrightarrow & C \\ c & \longmapsto & c^b \end{array}$$

then we call the object c^b the *exponential* of c by b and we say that the category C is *cartesian closed.*

EXAMPLE. **Set** is a cartesian closed category; the exponential c^b is the set of functions from b to c. ▲

Pullback

Given in C a pair $b \xrightarrow{f} a, d \xrightarrow{g} a$ of arrows with a comon codomain a, a *pullback square* of $< f, g >$ is a commutative square,

$$\begin{tikzcd} p \arrow[r, "k"] \arrow[d, "h"'] & d \arrow[d, "g"] \\ b \arrow[r, "f"'] & a \end{tikzcd}$$

such that for every other commutative square built on f, g,

$$\begin{tikzcd} c \arrow[r, "q"] \arrow[d, "p"'] & d \arrow[d, "g"] \\ b \arrow[r, "f"'] & a \end{tikzcd}$$

there is a unique $c \xrightarrow{t} p$ such that:

$$\begin{tikzcd} c \arrow[dr, dashed, "t"] \arrow[drr, bend left, "q"] \arrow[ddr, bend right, "p"'] & & \\ & p \arrow[r, "k"] \arrow[d, "h"'] & d \arrow[d, "g"] \\ & b \arrow[r, "f"'] & a \end{tikzcd}$$

Subobject classifier

A *subobject classifier* for a category C with a terminal object 1 is defined to be a monomorphism $1 \xhookrightarrow{t} \Omega$ such that for every monomorphism $S \xhookrightarrow{m} X$ in C, there exists a unique $X \xrightarrow{\psi} \Omega$ such that the following is a pullback square:

$$\begin{tikzcd} S \arrow[r] \arrow[d, "m"'] & 1 \arrow[d, "t"] \\ X \arrow[r, "\psi"] & \Omega \end{tikzcd}$$

In this pullback square, the top horizontal arrow is the unique map to the terminal object 1, the lower horizontal arrow ψ acts as the "characteristic function" of the given subobject S, while the "universal" monomorphism $1 \xhookrightarrow{t} \Omega$ may be called "truth".

EXAMPLE. In **Set**, the terminal object is a singleton $1 = \{0\}$, $\Omega = \{0, 1\}$ and t is the injection such that $t(0) = 0$. ▲

Limit

We refer the reader to [38] for the definition of a limit as we shall never directly check for limits, but use the following corollary (*cf.* [38, corollary 1, page 113]).

Corollary B.1 *(Saunders Mac Lane)*
If a category C has a terminal object, equalizers of all pair of arrows, and products of all pair of objects, then C has all finite limits.

Topos

An (elementary) *topos* is defined to be a category E with the following properties:

(i) E has all finite limits;

(ii) E has a subobject classifier; and

(iii) E is cartesian closed.

Bibliography

[1] BAKER, K., AND PIXLEY, A. Polynomial interpolation and the chinese remainder theorem. *Mathematische Zeitschrift 143* (1975), 165–174.

[2] BJARELAND, M., AND JONSSON, P. Exploiting bipartiteness to identify yet another tractable subclass of CSP. In *Proceedings of the Fifth International Conference on Principles and Practice of Constraint Programming (CP'99)* (1999).

[3] BULATOV, A. A., KROKHIN, A. A., AND JEAVONS, P. Constraint satisfaction problems and finite algebras. *Lecture Notes in Computer Science 1853* (2000), 272–282. Automata, Languages and Programming. Proceedings of ICALP'2000.

[4] CHANDRA, A., AND MERLIN, P. Optimal implementation of conjunctive queries in relational databases. In *9th ACM Symposium on Theory of Computing* (1979), pp. 77–90.

[5] COURCELLE, B. *Handbook of Theoretical Computer Science*, vol. 2. Elsevier Science, 19990, ch. Graph rewriting: an algebraical and logic approach, pp. 194–242.

[6] COURCELLE, B., MAKOWSKY, J., AND ROTICS, U. Linear time solvable optimization problems on certain structured graph families. *Theoretical Computer Science 33* (2000), 125–150.

[7] CREIGNOU, N., KHANNA, S., AND SUDAN, M. *Complexity classifications of Boolean Constraint Satisfaction Problems*. SIAM Monographs on Discrete Mathematics and Applications 7. SIAM, 2001.

[8] DECHTER, R. From local to global consistency. *Artificial Intelligence 55* (1992), 87–107.

[9] DECHTER, R. Bucket elimination: a unifying framework for reasoning. *Artificial Intelligence 113* (1999), 41–85.

[10] DECHTER, R., AND PEARL, J. Tree clustering for contraint networks. *Artificial Intelligence 38* (1989), 353–356.

[11] DUNKIN, N., AND ALLEN, S. Frequency assignment problems: Representations and solutions. Tech. rep., Department of Computer Science, Royal Holloway, University of London, Egham, Surrey, UK, 1997.

[12] EBBINGHAUS, H.-D., AND FLUM, J. *Finite Model Theory*. Perspectives in Mathematical Logic. Springer-Verlag, 1995.

[13] FAGIN, R. Generalized first-order spectra and polynomial-time recognizable sets. In *Complexity of Computation* (1974), R. M. Karp, Ed., pp. 43–73.

[14] FEDER, T. Classification of homomorphisms to oriented cycles and of k-partite satisfiability. *SIAM Journal of Discrete Mathematics 14*, 4 (2001), 471–480.

[15] FEDER, T., MADELAINE, F., AND STEWART, I. A. Dichotomies for classes of homomorphism problems involving unary functions. (submitted), October 2002.

[16] FEDER, T., AND VARDI, M. Y. The computational structure of monotone monadic SNP and constraint satisfaction: a study through datalog and group theory. *SIAM J. Comput. 28* (1999), 57–104.

[17] FLUM, J., FRICK, M., AND GROHE, M. Query evaluation via tree-decompositions. *Journal of the ACM 49* (2002), 715–152.

[18] FREUDER, E. A sufficient condition for backtrack-free search. *Journal of the ACM 29* (1982), 24–32.

[19] FREUDER, E. A sufficient condition for backtrack-bounded search. *Journal of the ACM 32* (1985), 755–761.

[20] GAREY, M., AND JOHNSON, D. *Computers and intractability: a guide to NP-completeness*. Freeman, San Francisco, California, 1979.

[21] GRÄDEL, E. Capturing complexity classes by fragments of second-order logic. *Theoretical Computer Science 101* (1992), 35–57.

[22] HAHN, G., AND TARDIF, C. Graph homomorphisms: structure and symmetry. *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci. 497* (1997), 107–166. Kluwer Acad. Publ., Dordrecht.

[23] HELL, P., AND NEŠETŘIL, J. On the complexity of h-coloring. *J. Combin. Theory Ser. B 48* (1990), 92–110.

[24] HELL, P., AND NEŠETŘIL, J. The core of a graph. *Discrete Math 109* (1992), 117–126.

[25] HOPCROFT, J. E., AND ULLMAN, J. D. *Introduction to automata theory, languages and computation.* Addison-Wesley, 1979.

[26] IMMERMAN, N. *Descriptive Complexity.* Springer-Verlag, 1998.

[27] IMMERMAN, N., AND LANDAU, S. The complexity of iterated multiplication. *Information and Computation 116* (1995), 103–116.

[28] JEAVONS, P. On the algebraic structure of combinatorial problems. *Theoritical Computer Science 200* (1998), 185–204.

[29] JEAVONS, P., COHEN, D., AND COOPER, M. Characterising tractable constraints. *Artificial Intelligence 65*, 2 (1994).

[30] JEAVONS, P., COHEN, D., AND COOPER, M. Constraint, consistency, and closure. *Artificial Intelligence 101*, 1-2 (1998), 251–265.

[31] JEAVONS, P., COHEN, D., AND GYSSENS, M. Closure properties of constraints. *Journal of the ACM 44* (1997), 527–548.

[32] JEAVONS, P., COHEN, D., AND PEARSON, J. Constraints and universal algebra. *Annals of Mathematics and Artificial Intelligence 24* (1998), 51–67.

[33] JEAVONS, P., AND COOPER, M. Tractable constraints on ordered domains. *Artificial Intelligence 79*, 2 (1995), 327–339.

[34] JEAVONS, P., D., D. C., AND GYSSENS, M. A unifying framework for tractable constraints. In *Proceedings 1st International Conference on Constraint Programming* (1995), vol. 976, Lecture Notes in Computer Science, Springer Verlag.

[35] KOLAITIS, P., AND VARDI, M. Conjunctive-query containment and constraint satisfaction. In *Proc. 17th ACM Symp. on Principles of Database Systems* (1998), pp. 205–213.

[36] KOLAITIS, P. G., AND VARDI, M. Y. A game-theoretic approach to constraint satisfaction. In *AAAI-2000/IAAI-2000 Proceedings (Austin Texas)* (2000), MIT Press, Cambridge, MA,, pp. 175–181.

[37] LADNER, R. On the structure of polynomial time reducibility. *J.ACM 22* (1975), 155–171.

[38] LANE, S. M. *Categories for the working mathematician*, 2nd ed. Springer-Verlag, 1998.

[39] LASSAIGNE, R., AND DE ROUGEMONT, M. *Logique et complexité.* collection informatique. Hermès, 1996.

[40] LAU, D. Funktionenalgebren über endlichen Mengen. Universität Rostock, Fachbereich Mathematik (unpublished), 1998.

[41] MADELAINE, F. Clones and constraints. Master's thesis, Université de Caen and Technische Universität Dresden, September 1998. *mémoire de DEA* .

[42] MADELAINE, F., AND STEWART, I. A. On the complexity of homomorphism problems involving unary functions. Tech. rep., Department of Mathematics and Computer Science, University of Leicester, 2001.

[43] MADELAINE, F., AND STEWART, I. A. Some problems not definable using structures homomorphisms. *Ars Combinatoria LXVII* (2003).

[44] MCKENZIE, R., MCNULTY, G., AND TAYLOR, W. *Algebras, Lattices and Varieties*, vol. I. Wadsworth and Brooks, California, 1987.

[45] NEŠETŘIL, J., AND TARDIF, C. Duality theorems for finite structures (characterising gaps and good characterisations. *Journal of Combin. Theory Ser. B 80* (2000), 80–97.

[46] PAPADIMITRIOU, C. *Computational Complexity*. Addison-Wesley, 1994.

[47] POST, E. The two-valued iterative systems of mathematical logic. In *Annals of Math. studies*, vol. 5. Princeton, 1941.

[48] PÖSCHEL, R. A general galois theory for operations and relations and concrete characterization of related algebraic structures. Tech. rep., Akademie der Wissenschaften der DDR, Zentralinstitut für Mathematik und Mechanik, 1980.

[49] PÖSCHEL, R., AND KALUŽNIN, L. *Funktionen- und relationen-Algebren*. VEB Deutscher Verlag der Wissenschaften, 1979.

[50] SAUER, N. Hedetniemi's conjecture–a survey. *Discrete Mathematics 229* (2001), 261–292.

[51] SAUER, N., 2002. private discussion.

[52] SCHAEFER, T. The complexity of satisfiability problems. In *Proc* 10^{th} *ACM Symposium on theory of computing (STOC)* (1978), pp. 216–226.

[53] STEWART, I. A. Comparing the expressibility of languages formed using NP-complete operators. *J. Logic Computat. 3* (1991), 305–330.

[54] STEWART, I. A. Completeness of path-problems via logical reductions. *Information and Computation 121* (1995), 123–134.

[55] SZENDREI, A. *Clones in Universal Algebra*, vol. 99 of *Seminaires de Mathematiques Superieures*. University of Montreal, 1986.

Index