

Logic, Computation and Constraint Satisfaction

Barnaby D. Martin

University of Leicester

Submitted for the degree of Doctor of Philosophy

November 2005

UMI Number: U537212

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U537212

Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.



ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

We study a class of non-deterministic program schemes with while loops: firstly, augmented with a priority queue for memory; secondly, augmented with universal quantification; and, thirdly, augmented with universal quantification and a stack for memory. We try to relate these respective classes of program schemes to well-known complexity classes and logics.

We study classes of structure on which path system logic coincides with polynomial time P.

We examine the complexity of generalisations of non-uniform boolean constraint satisfaction problems, where the inputs may have a bounded number of quantifier alternations (as opposed to the purely existential quantification of the CSP). We prove, for all bounded-alternation prefixes that have some universal quantifiers to the outside of some existential quantifiers (*i.e.* Π_2 and above), that this generalisation of boolean CSP respects the same dichotomy as that for the non-uniform boolean quantified constraint satisfaction problem.

We study the non-uniform QCSP, especially on digraphs, through a combinatorial analog – the alternating-homomorphism problem – that sits in relation to the QCSP exactly as the homomorphism problem sits with the CSP. We establish a trichotomy theorem for the non-uniform QCSP when the template is restricted to antireflexive, undirected graphs with at most one cycle. Specifically, such templates give rise to QCSPs that are either tractable, NP-complete or Pspace-complete.

We study closure properties on templates that respect QCSP hardness or QCSP equality. Our investigation leads us to examine the properties of first-order logic when deprived of the equality relation.

We study the non-uniform QCSP on tournament templates, deriving sufficient conditions for tractability, NP-completeness and Pspace-completeness. In particular, we prove that those tournament templates that give rise to tractable CSP also give rise to tractable QCSP.

Acknowledgements

I would like to thank my global and local supervisors, Iain Stewart and Rick Thomas, respectively, for their help, guidance and advice throughout the last four years. I would like to thank my mentor Florent Madelaine for his assistance in the field of constraint satisfaction problems. I would also like to thank Florent and his wife Sukhi for their hospitality in Newcastle on many occasions.

I would like to thank my examiners, Dave Cohen and Vincent Schmitt, for undertaking the arduous task of reading this thesis, and for their many helpful suggestions.

I am very grateful to the Engineering and Physical Sciences Research Council (EPSRC) for funding this research.

Contents

1	Introduction	5
2	Program Schemes	10
2.1	Structures and Logic	11
2.1.1	Graphs and Transitive Closure	11
2.1.2	Alternating graphs and Alternating Reachability	13
2.1.3	Hypergraphs and Path Systems	15
2.1.4	Least Fixed Point logic	16
2.1.5	Stratified Fixed Point logic	17
2.2	Program Schemes	17
2.2.1	Introducing NPS	18
2.2.2	Shorthands	21
2.2.3	Shorthands on successor structures	23
2.2.4	Adding a stack: introducing NPSS	24
2.3	Turing Machines	25
3	Adding a Priority Queue	27
3.1	A single weight: $k = 0$	32
3.2	The Hamilton Path problem is in NPSPQ	34
3.3	$\text{NPSPQ}_s^b \subseteq \text{NPspace}$	36
3.4	Expanding alphabets	37
3.5	Pushing and Popping Numbers in $\text{NPSPQ}_s^b / \text{NPSPQ}_s^u$	39
3.5.1	$\text{NPSPQ}_s^{u+}(k) = \text{NPSPQ}_s^u(k)$	39
3.5.2	$\text{NPSPQ}_s^{b+} = \text{NPSPQ}_s^b$	40
3.6	$\text{NPSPQ}^b(k) \subseteq \text{NPSPQ}^u(k)$	41

3.7	$\text{NPSPQ}_s^u(k) \subseteq \text{NPSPQ}^u(k+2)$.	43
3.7.1	Start-routine: τ_{start}	43
3.7.2	Simulation of ρ : τ_ρ .	44
3.7.3	End-routine: τ_{end}	45
3.8	$\text{NPspace} \subseteq \text{NPSPQ}_s^u$	47
3.8.1	Preparation: $\tau_{bin(A)}$	49
3.8.2	Simulation: τ_{sim}	51
3.8.3	Verification: τ_{end}	52
3.9	A polynomial-time restriction of NPSPQ^u .	54
4	Adding Universal Quantification	56
4.1	Introducing $\text{APS}(1)$	57
4.2	Introducing $\text{APSS}(1)$	61
4.3	The ACCEPT instruction.	63
4.4	$\text{APSS}_s(1) = \text{APSS}(1)$	64
4.4.1	Simulation of ρ ; τ_ρ	65
4.5	$\text{NPspace} \subseteq \text{APSS}_s(1)$	66
4.6	Summary	72
5	Classes of Structure on which $P = \pm\text{PS}^k[\text{FO}]$	73
5.1	Finitely generated sets	74
5.1.1	F contains a single k -ary partial function f_0 .	76
5.1.2	F contains multiple partial functions.	78
5.1.3	An application: finitely generated groups	78
5.2	Hamiltonian Outerplanar graphs	79
6	Dichotomies in Boolean Constraint Satisfaction	83
6.1	Introduction	83
6.2	Technical results	88
6.3	A dichotomy theorem for $\Pi_2\text{-SAT}_{\text{NC}}$	91
6.3.1	Case 1 : R^A is 0-valid and not 1-valid.	92
6.3.2	Case 2 : R^A is 1-valid and not 0-valid.	92
6.3.3	Case 3 : R^A is 0-valid and 1-valid, but not complimentary.	93
6.3.4	Case 4 : R^A is 0-valid and 1-valid, and complimentary.	94

7	Quantified Constraints on Graphs	96
7.1	Introduction	96
7.2	Preliminaries	99
7.2.1	Structures and Logic.	99
7.2.2	Alternating-homomorphism problems.	100
7.2.3	QCSP versus ALT-HOM.	103
7.2.4	Alternating-homomorphisms as winning strategies.	104
7.2.5	Graphs	106
7.3	Basic graph results	107
7.3.1	Restricting partitions	107
7.3.2	Basic results.	110
7.3.3	Non-connected templates.	112
7.4	Quantified \mathcal{H} -colouring	114
7.4.1	Bipartite templates.	114
7.4.2	Odd Catherine Wheels	116
7.4.3	A trichotomy theorem	122
7.5	Closure properties	123
7.5.1	Indicator construction	123
7.5.2	Adding a vertex to a core.	126
7.5.3	A sufficient condition for ALT-HOM problem equivalence.	129
7.5.4	Why that condition is not necessary: equivalence in fragments of FO	131
7.6	Results concerning tournament templates	137
7.6.1	Template is a directed cycle.	137
7.6.2	Template is a digraph with source and sink.	138
7.6.3	Tractable tournament ALT-HOM problems.	139
7.6.4	NP-complete tournament ALT-HOM problems.	143
7.6.5	Pspace-complete tournament ALT-HOM problems.	143
8	Conclusions and Further Work	146
8.1	Program Schemes	146
8.2	Classes of Structure on which $P = \pm PS^k[\mathbf{FO}]$	146
8.3	Quantified Constraints on Graphs	147

Chapter 1

Introduction

Structural Complexity is that part of the study of Computational Complexity that concerns itself with the intrinsic computational difficulty of decidable problems. Perhaps its main thrust is an attempt to classify problems into *complexity classes* by various upper and lower bounds on their computational complexity. Since its inception, logic has impinged on Computational Complexity in a variety of ways: in the first instance, many of the problems that are amongst the hardest of many natural complexity classes have been problems in logic. These problems, known as *complete* for the given complexity class, include the following problems in the Propositional Calculus: Circuit Value, complete for P; Satisfiability, complete for NP; and Quantified Satisfiability, complete for Pspace.

Another intersection between logic and complexity is in the field of Descriptive Complexity – in which finding an algorithm for a computational problem is seen as a question of expression – where complexity classes translate to classes of expressions, *i.e.* logics. Indeed, the complexity classes, defined through Turing Machines, may be seen as logics already: for example, if P is defined as those Turing Machines T , such that there exists k , such that T accepts an input x of size n iff T accepts x within time n^k , then this class of Turing Machines is a logic of sorts. However, it is a logic not of a form that lends itself to study by what are usually considered the tools of logic: it has a cumbersome syntax and is interpreted over strings and not first-order structures. We may remedy the second of these problems by considering certain standard binary string encodings of a structure. Since, for us, a decision problem is always a subset of finite structures, we consider a problem to be in P iff the language of all binary encodings of these structures is in P, as defined previously. In this way, we can talk of conventional logics capturing Turing complexity classes: a logic captures a complexity class iff the set of problems expressible in each coincides. It is in the translation between computational problems over strings and expression problems over finite structures that Descriptive Complexity is concerned. As such, it is very much a part of Finite Model Theory.

Perhaps the greatest hope for Descriptive Complexity was that known methods for separating logics might be brought to bear on complexity classes; that hard questions on the (non)-equivalence of complexity classes might become easier questions on the separation of logics. Among the major results of Descriptive

Complexity are the proven equivalence of: existential second-order logic $\exists\text{SO}$ and NP (Fagin, 1974 [17]); least fixed-point logic LFP (with successor) and P (Immerman/Vardi 1982 [28, 44]); partial fixed point logic PFP (with successor) and Pspace (Vardi 1982 [44]); and transitive closure logic TC (with successor) and NL (Immerman 1983/1988 [28, 30]). Despite these results, few advances have been made in the use of techniques such as Ehrenfeucht-Fraïssé games to separate these logics, and, consequently, their complexity classes¹. Partly, this can be explained by the somewhat artificial inclusion of the successor function into many of these logics (not $\exists\text{SO}$), since Ehrenfeucht-Fraïssé games are notoriously hard to win on structures with successor. However, the equivalence of LFP and PFP, both with successor, is known to be consequent on their equivalence without successor [1], yet still a proof resists that $P \neq \text{Pspace}$. One ray of sunshine in this field was Immerman's proof, through Transitive Closure logic, that $\text{NL} = \text{co-NL}$ [30].

The syntax of a logic is exactly its set of well-formed formulae; a logic is said to have recursive syntax iff its syntax is decidable. The inclusion of successor generally precludes the possibility that the resultant logic has recursive (or even recursively enumerable) syntax [24], a property which is certainly desirable, and is thought by some authors (e.g. Gurevich [24], Otto [36]) to be necessary, if the name 'logic' is to be bestowed.

In the first part of this thesis we study various classes of non-deterministic Program Schemes with while loops (based on those in [2, 41]), which are logics in Gurevich's sense, but which appear well-suited for computation. We attempt to relate these logics to standard complexity classes, preferably in the absence of that built-in successor.

Chapter 2 introduces these program schemes, and discusses some known results involving them [2]. The situation where a stack is available for memory [2] is considered.

In Chapter 3, we introduce some new work investigating the addition of a priority queue as a memory device available to these program schemes. We prove that the priority queue is sufficiently powerful to simulate a successor function:

¹Indeed, it is not known that NL, P, NP and Pspace are not equivalent.

thus we define two ‘logics’ with recursive syntax that subsume NPspace^2 and NP , respectively. These logics actually have identical syntax, and differ only in their semantics.

In Chapter 4, we introduce some new work examining the explicit introduction of universal quantification to our program schemes (since they are non-deterministic, existential quantification is already insinuated), both in the absence of any memory and with the benefit of a stack. With no added memory, we tie the ensuing logic to least fixed-point logic LFP . With the addition of a stack we, once again, are able to simulate a successor: this enhanced logic subsumes NPspace .

Chapter 5 uses some known results utilising program schemes [42] to study various classes of structure on which the infinite hierarchy which constitutes Path System logic collapses and captures P . We give a brief overview of known results and introduce some new ones.

A further confluence of logic and complexity, and, indeed, combinatorics, is in the study of the Constraint Satisfaction Problem, CSP , and its generalisations. In terms of complexity classification, much has been made of the conjectured dichotomy of the non-uniform CSP on finite templates [1]: it seems as though, for any \mathcal{T} , $\text{CSP}(\mathcal{T})$ is either tractable or NP -complete. This is remarkable given the breadth of CSP problems, together with Ladner’s result [35] that such a dichotomy will not hold over all NP (unless we actually have $\text{P} = \text{NP}$). The non-uniform CSP s, and their generalisations, lend themselves to dual interpretations: one in which they are model-checking problems over restricted logics; and one in which they are combinatorial problems between two structures. In particular, the non-uniform CSP may be seen as both a model-checking problem in existential positive conjunctive first-order logic and as the homomorphism problem. This duality is perhaps at its most obvious on graphs, and it is on these that we dwell most.

Chapter 6 concerns itself with the dichotomy of alternation-bounded generalisations of the non-uniform CSP on boolean templates. (These results have been obtained independently in [27] and, to a lesser extent, in [18].) We prove, for all bounded-alternation prefixes that have some universal quantifiers to the outside of

²We remind the reader that NPspace and Pspace coincide (e.g. [37]). Even so, we use both classes in this thesis, depending on which appears the more natural in a given situation.

some existential quantifiers (*i.e.* Π_2 and above), that our generalisation of boolean CSP respects the same dichotomy as that for boolean quantified constraint satisfaction problems.

Chapter 7 examines the non-uniform Quantified Constraint Satisfaction Problem, QCSP, on graph templates. We study the QCSP through a combinatorial analog, the so-called Alternating-Homomorphism problem ALT-HOM. We study a variety of graph templates that give rise to tractable, NP-complete or Pspace-complete QCSPs, culminating in a complete classification to those classes – a trichotomy theorem – when the template ranges over undirected antireflexive graphs with at most one cycle.

We consider two problems to be equal exactly when the respective subsets of structures that they define coincide. It is well-known that the two problems $\text{CSP}(\mathcal{T})$ and $\text{CSP}(\mathcal{T}')$ are equal iff the templates \mathcal{T} and \mathcal{T}' are homomorphically equivalent (which is exactly the condition that they have isomorphic cores). We study a similar condition on templates $\mathcal{T}, \mathcal{T}'$ that is sufficient to guarantee the equality of $\text{QCSP}(\mathcal{T})$ and $\text{QCSP}(\mathcal{T}')$. However, we find that this condition is not necessary, and that it has a closer relationship with first-order logic without equality than the logic we associate with QCSP (positive-conjunctive first-order logic).

Finally, we study the non-uniform QCSP on tournament templates, deriving sufficient conditions for tractability, NP-completeness and Pspace-completeness. In particular, we prove that those tournament templates that give rise to tractable CSP also give rise to tractable QCSP.

Chapter 2

Program Schemes

2.1 Structures and Logic

We will only consider finite relational structures, of at least two elements, over a given signature σ . We denote this set $\text{STRUC}(\sigma)$. If \mathcal{A} is a structure, then $|\mathcal{A}|$ is the universe, or domain, of the structure, and $||\mathcal{A}||$ is the cardinality of that domain. If R is a relation symbol of σ then $R^{\mathcal{A}}$ is the interpretation of R over \mathcal{A} . When the structure \mathcal{A} is clear, we may abuse notation by dropping it as the superscript, thus identifying R with both the relation *symbol* and the relation *actual*.

We will also consider the situation where a successor is available to us, built-in to the signature σ . We consider a successor to be a binary relation *succ*, whose realisation as a graph is a directed path, together with two constants *min* and *max*, whose interpretations are the first and last vertices of that path. This is equivalent to considering the restricted class of structures over $\sigma \uplus \{\text{succ}, \text{min}, \text{max}\}$ in which the interpretations of *succ*, *min* and *max* satisfy the properties given. Throughout, when we consider the restriction to structures that have a successor relation, we add the subscript *s* to our logic or class, for example FO_s . When we consider logics in which we have a successor, we will insist on a further *semantic* restriction on their formulae, namely, that a formula may only be in that logic if its truth is independent of the actual successor function used. For example, consider the formula $E(\text{min}, \text{max})$, ostensibly of FO_s , interpreted on the directed 3-path – the graph with vertices $\{0, 1, 2\}$ and edge set $\{(0, 1), (1, 2)\}$. The truth of this formula is not independent of the ordering we choose on the graph – if the successor is $\{(0, 2), (2, 1)\}$, it is true; if the successor is $\{(0, 1), (1, 2)\}$, it is false. We conclude that $E(\text{min}, \text{max})$ is not a formula of FO_s . Given such a formula, ostensibly of FO_s , establishing whether it has this property of order-independence is, in general, undecidable [24].

2.1.1 Graphs and Transitive Closure

A graph, or digraph, \mathcal{G} is a structure over signature $\sigma_2 = \langle E \rangle$, where E is a binary relation symbol. There is a (directed) path in \mathcal{G} from vertex x to vertex y iff either: $x = y$, $E(x, y)$, or there is a sequence of vertices z_1, \dots, z_r such that $E(x, z_1)$, $E(z_i, z_{i+1})$, for $1 \leq i < r$, and $E(z_r, y)$. This is equivalent to the inductive definition

that there is a path from x to y iff:

- $x = y$, or
- there is a z such that $E(x, z)$, and there is a path from z to y .

Definition. Define TC to be the global binary relation on graphs expressing reachability. Specifically:

$$\text{TC} = \{ \{(x, y) : \text{there is a path in } \mathcal{G} \text{ from } x \text{ to } y\} : \mathcal{G} \in \text{STRUC}(\sigma_2) \}$$

Let ψ be some formula whose only free variables are among those of the j -tuples \bar{x} and \bar{y} . A formula $\text{TC}[\lambda\bar{x}, \bar{y}\psi](\bar{u}, \bar{v})$ is interpreted as true on a structure \mathcal{A} in the case that, in the graph of $||\mathcal{A}||^j$ vertices with edge set specified by $\psi(\bar{x}, \bar{y})$, there is a path from vertex \bar{u} to vertex \bar{v} . (It is usual to allow additional free variables, *i.e.* other than \bar{x} and \bar{y} , in ψ . However, this does not increase our expressive power, since i such variables can be moved so they only appear free in the end-point $(j+i)$ -tuples \bar{u}' and \bar{v}' of some new Transitive Closure formula over a graph of size $||\mathcal{A}||^{j+i}$ (see [16]). We forbid such additional free variables for the sake of a simpler exposition.)

In the fashion described, the global relation TC has given us a uniform sequence of vectorised quantifiers of the same name. This sequence is derived from the arity of \bar{x} and \bar{y} . The first quantifier in the sequence corresponds to the arity of $\bar{x} = (x_1)$ and $\bar{y} = (y_1)$ being 1, and binds the 2 variables, x_1 and y_1 . The i th quantifier in this sequence corresponds to the arity of $\bar{x} = (x_1, \dots, x_i)$ and $\bar{y} = (y_1, \dots, y_i)$ being i , and binds the $2i$ variables x_1, \dots, x_i and y_1, \dots, y_i . This sequence of quantifiers is uniform in semantics and syntax, and is an example of a sequence of Lindström quantifiers (see, *e.g.*, [16]).

Definition. Let \bar{x}, \bar{y} be j -tuples of variables. Let $j', j'' \leq j$ and $u_1, \dots, u_{j'}, v_1, \dots, v_{j''}$ be variables, and $u_{j'+1}, \dots, u_j, v_{j''+1}, \dots, v_j$ be variables or constant symbols. Define:

- $\pm\text{TC}^1[\mathbf{FO}]$ to be the set of formulae of the form

$$\exists u_1 \dots u_{j'} \exists v_1 \dots v_{j''} \text{TC}[\lambda\bar{x}, \bar{y}\psi](\bar{u}, \bar{v})$$

where ψ is quantifier-free, and

- $\pm\text{TC}^{m+1}[\mathbf{FO}]$ to be the set of formulae of the form

$$\exists u_1 \dots u_j \exists v_1 \dots v_j \text{TC}[\lambda \bar{x}, \bar{y} \psi](\bar{u}, \bar{v})$$

where ψ is in the closure under boolean operators of formulae in $\pm\text{TC}^m[\mathbf{FO}]$.

In the presence of two distinct constants, any formula in $\text{TC}^{m+1}[\mathbf{FO}]$ is equivalent to some formula of the form $\text{TC}[\lambda \bar{x}, \bar{y} \psi](\bar{u}, \bar{v})$, with $\psi \in \text{TC}^m[\mathbf{FO}]$, i.e. without the need for existential quantification outside the TC operator [21]. However, we do not wish to restrict ourselves only to structures with such constants. We define $\pm\text{TC}^*[\mathbf{FO}]$ to be $\bigcup_{i \in \omega} \pm\text{TC}^i[\mathbf{FO}]$.

Recall that the subscript s denotes a built-in successor. The following gives us an idea as to the power of Transitive Closure logic.

Proposition 1 (Immerman 1983/1988, [29, 30]). $\text{TC}_s^1[\mathbf{FO}] = \text{TC}_s^*[\mathbf{FO}] = \text{NL}$.

Remark. It may be noticed that we are rather liberal with notation such as TC, allowing it to denote a global relation, an operator and a logic. Hopefully, the meaning should be clear by context.

2.1.2 Alternating graphs and Alternating Reachability

An *alternating graph* \mathcal{A} is a structure over the signature $\sigma_{21} = \langle E, U \rangle$, where the relation symbols E and U are binary and unary, respectively. In an alternating graph the unary relation U partitions the vertices into those that are existential ($\neg U$), and those that are universal (U). There is an alternating path in an alternating graph \mathcal{A} , from vertex x to vertex y , iff:

- $x = y$, or
- x is existential, and there is a z such that $E(x, z)$, and there is an alternating path from z to y , or
- x is universal, and, for all z such that $E(x, z)$, there is an alternating path from z to y .

Definition. Define AR to be the global binary relation on graphs expressing alternating reachability. Specifically:

$$\text{AR} = \{(x, y) : \text{there is an alternating path in } \mathcal{A} \text{ from } x \text{ to } y\} : \mathcal{A} \in \text{STRUC}(\sigma_{21})\}$$

Let ψ be some formula whose only free variables are among those of \bar{x} and \bar{y} . A formula $\text{AR}[\lambda\bar{x}, \bar{y}\psi](\bar{u}, \bar{v})$ is interpreted as true in the case that, in the graph specified by $\psi(\bar{x}, \bar{y})$, there is an alternating path from \bar{u} to \bar{v} . (Again, it is customary to permit additional free variables in ψ . Again, it is unnecessary for the same reason as given for TC.)

Definition. Let \bar{x}, \bar{y} , be j -tuples of variables. Let $j', j'' \leq j$ and $u_1, \dots, u_{j'}, v_1, \dots, v_{j''}$ be variables, and $u_{j'+1}, \dots, u_j, v_{j''+1}, \dots, v_j$ be variables or constant symbols. Define:

- $\pm\text{AR}^1[\mathbf{FO}]$ to be the set of formulae of the form:

$$\exists u_1 \dots u_{j'} \exists v_1 \dots v_{j''} \text{AR}[\lambda\bar{x}, \bar{y}\psi](\bar{u}, \bar{v})$$

where ψ is quantifier-free, and

- $\pm\text{AR}^{m+1}[\mathbf{FO}]$ to be the set of formulae of the form

$$\exists u_1 \dots u_{j'} \exists v_1 \dots v_{j''} \text{AR}[\lambda\bar{x}, \bar{y}\psi](\bar{u}, \bar{v})$$

where ψ is in the closure under boolean operators of formulae in $\pm\text{AR}^m[\mathbf{FO}]$.

We define $\pm\text{AR}^*[\mathbf{FO}]$ to be $\bigcup_{i \in \omega} \pm\text{AR}^i[\mathbf{FO}]$. In the presence of two distinct constants, $\pm\text{AR}^{m+1}[\mathbf{FO}]$ collapses to the class of formulae of the form $\text{AR}[\lambda\bar{x}, \bar{y}\psi](\bar{u}, \bar{v})$ for $\psi \in \pm\text{AR}^m[\mathbf{FO}]$ (proof similar to that for TC).

Recall that the subscript s denotes a built-in successor. The following gives us an idea as to the power of Alternating Reachability logic.

Proposition 2 (Immerman 1983, [29]). $\text{AR}_s^1[\mathbf{FO}] = \text{AR}_s^*[\mathbf{FO}] = \text{P}$.

2.1.3 Hypergraphs and Path Systems

We consider a *hypergraph*¹ \mathcal{H} to be a structure over the signature $\sigma_3 = \langle R \rangle$, where R is a ternary relation symbol. A vertex y is said to be *R-accessible* (or just *accessible*) from a vertex x iff:

- $x = y$, or
- there exist z_1, z_2 , both accessible from x , such that $R(z_1, z_2, y)$.

A hypergraph \mathcal{H} is said to be *commutative* exactly when, for all x, y, z , we have $R(x, y, z)$ iff $R(y, x, z)$. It is said to be *deterministic* iff, for all x, y , there exists at most one z such that $R(x, y, z)$.

Definition. Define PS to be the global binary relation on commutative hypergraphs expressing accessibility. Specifically:

$$\text{PS} = \{ \{ (x, y) : y \text{ is accessible in } \mathcal{H}, \text{ from } x \} : \mathcal{H} \text{ is a commutative hypergraph} \}$$

Let ψ be some formula whose only free variables are among those of \bar{x} , \bar{y} and \bar{z} . A formula $\text{PS}[\lambda\bar{x}, \bar{y}, \bar{z}\psi](\bar{u}, \bar{v})$ is interpreted as true in the case that, in the commutative hypergraph specified by $\psi(\bar{x}, \bar{y}, \bar{z})$, \bar{v} is accessible from \bar{u} . (Again, it is customary to permit additional free variables in ψ . Again, it is unnecessary for the same reason as given for TC.)

Definition. Let \bar{x}, \bar{y} , be j -tuples of variables. Let $j', j'' \leq j$ and $u_1, \dots, u_{j'}, v_1, \dots, v_{j''}$ be variables, and $u_{j'+1}, \dots, u_j, v_{j''+1}, \dots, v_j$ be variables or constant symbols. Define:

- $\pm\text{PS}^1[\mathbf{FO}]$ to be the set of formulae of the form

$$\exists u_1 \dots u_{j'} \exists v_1 \dots v_{j''} \text{PS}[\lambda\bar{x}, \bar{y}, \bar{z}\psi](\bar{u}, \bar{v})$$

where ψ is quantifier-free, and

¹What we refer to as a hypergraph would perhaps be better described as a *directed 3-uniform hypergraph*, taking into consideration more standard definitions.

- $\pm\text{PS}^{m+1}[\mathbf{FO}]$ to be the set of formulae of the form

$$\exists u_1 \dots u_{j'} \exists v_1 \dots v_{j''} \text{PS}[\lambda \bar{x}, \bar{y}, \bar{z} \psi](\bar{u}, \bar{v})$$

where ψ is in the closure under boolean operators of formulae in $\pm\text{PS}^m[\mathbf{FO}]$.

We define $\pm\text{PS}^*[\mathbf{FO}]$ to be $\bigcup_{i \in \omega} \pm\text{PS}^i[\mathbf{FO}]$. In the presence of two distinct constants, $\pm\text{PS}^{m+1}[\mathbf{FO}]$ collapses to the class of formulae of the form $\text{PS}[\lambda \bar{x}, \bar{y} \psi](\bar{u}, \bar{v})$ for $\psi \in \pm\text{PS}^m[\mathbf{FO}]$ (proof similar to that for TC).

Recall that the subscript s denotes a built-in successor. The following gives us an idea as to the power of Path System logic.

Proposition 3 ([40]). $\text{PS}_s^1[\mathbf{FO}] = \text{PS}_s^*[\mathbf{FO}] = \text{P}$.

2.1.4 Least Fixed Point logic

Let $\psi(P, \bar{x})$ be a first-order formula with free j -ary relation symbol P whose only free variables are those of the j -tuple \bar{x} . Then, over a structure \mathcal{A} , ψ may be seen as a function $f_{\mathcal{A}} : \mathcal{P}(|\mathcal{A}|^j) \rightarrow \mathcal{P}(|\mathcal{A}|^j)$ defined by:

$$f_{\mathcal{A}}(R) = \{\bar{x} : \mathcal{A} \models \psi(R, \bar{x})\}$$

If ψ does not contain negated instances of the free relation symbol P (i.e. is P -positive), then the function $f_{\mathcal{A}}$ is monotone, satisfying $R \subseteq f_{\mathcal{A}}(R)$. Given a P -positive ψ , we define inductively: $\psi_{\mathcal{A}}^0 = \emptyset$, and thereafter $\psi_{\mathcal{A}}^{k+1} = f_{\mathcal{A}}(\psi_{\mathcal{A}}^k)$. Since f is monotone and \mathcal{A} is finite, we are guaranteed that this sequence of relations must reach a fixed-point K where $\psi_{\mathcal{A}}^K = \psi_{\mathcal{A}}^i$ (for all $i \geq K$). This relation is denoted $\psi_{\mathcal{A}}^{\infty}$.

Definition. Given a formula $\psi(P, \bar{x})$ with free j -ary relation symbol P and including free variables of the j -tuple \bar{x} , and another j -tuple of variables or constants \bar{u} , we may apply the Least Fixed Point operator LFP to generate the formula $\text{LFP}[\lambda P \bar{x} \psi](\bar{u})$. This formula's free variables are those free in ψ that are not in \bar{x} , and those of \bar{u} . The formula is interpreted as true on a structure \mathcal{A} (under some valuation of its free variables) exactly when $\bar{u} \in \psi_{\mathcal{A}}^{\infty}$.

Least Fixed Point logic $\text{LFP}[\mathbf{FO}]$ is the closure of \mathbf{FO} under the Least Fixed Point operator.

Remark. It may be noted that we have allowed free variables in ψ that are not among the variables of \bar{x} , in contrast with the situation with the Lindström logics of the previous sections. It seems particularly unnatural to specify LFP with such free variables forbidden, moreover, we will make use of them in later chapters. It suffices to say that these additional variables could be forbidden by being forced into the outer tuple \bar{u} , as described for Transitive Closure logic.

2.1.5 Stratified Fixed Point logic

Definition. Let R be a free j -ary relation, and \bar{x} be a j -tuple of variables. Let $j' \leq j$ and $u_1, \dots, u_{j'}$ be variables, and $u_{j'+1}, \dots, u_j$ be variables or constant symbols. Define:

- $\exists\text{LFP}^1[\mathbf{FO}]$ to be the set of formulae of the form $\exists u_1 \dots u_{j'} \text{LFP}[\lambda R \bar{x} \psi](\bar{u})$ where ψ is first-order with no universal quantifiers and with negation only of atomic formulae, and
- $\exists\text{LFP}^{m+1}[\mathbf{FO}]$ to be the set of formulae of the form $\exists u_1 \dots u_{j'} \text{LFP}[\lambda R \bar{x} \psi](\bar{u})$ where ψ is first-order with no universal quantifiers but may contain positive or negative occurrences of formulae of $\exists\text{LFP}^m[\mathbf{FO}]$ that do not contain R .

We naturally define $\exists\text{LFP}^*[\mathbf{FO}]$ to be $\bigcup_{i \in \omega} \exists\text{LFP}^i[\mathbf{FO}]$. In the presence of two distinct constants, $\exists\text{LFP}^{m+1}[\mathbf{FO}]$ collapses to the class of formulae of the form $\exists\text{LFP}[\lambda R \bar{x} \psi](\bar{u}, \bar{v})$ for $\psi \in \exists\text{LFP}^m[\mathbf{FO}]$ [22].

We note that $\exists\text{LFP}^*[\mathbf{FO}]$ is often known as *Stratified Fixed Point logic SFP*. The following gives us an idea as to the power of $\exists\text{LFP}$.

Lemma 4 (Grohe 1997, [22]). $\pm\text{PS}^m[\mathbf{FO}] = \exists\text{LFP}^m[\mathbf{FO}]$.

2.2 Program Schemes

We will examine several classes of non-deterministic program schemes with while loops, originally seen in [2]. These program schemes were born of an attempt to

imbue logic with the tools of computation, whilst keeping that logic well-behaved, e.g., with recursive syntax. Unlike Turing Machines, which compute on strings encoding some structure, these schemes compute on a structure, in a similar manner to a formula of logic being interpreted on that structure. However, the syntax of computation is often more easily followed, and this may have advantages in simplifying proofs. For example, the recursion of while loops may be considered more natural than that of fixed point logics. Such advantages are largely cosmetic, but, in studying objects of computation, forms of memory can be added that would be most bizarre added directly to conventional logic. In doing this, new logics can be defined without obvious parallel in conventional logic. However, that which is not obvious is not necessarily untrue, and several results are known tying these new logics with their better known, conventional counterparts.

2.2.1 Introducing NPS

Definition (Syntax of NPS [2]). Each program scheme $\rho \in \text{NPS}(1)$, over signature σ , involves a set of input-output variables V_{io} , a set of free variables V_f , and a finite sequence of $|\rho|$ instructions, where each instruction, other than the first and last, is of one of the following forms:

- an assignment instruction of the form ‘ $v := q$ ’, where $v \in V_{io}$ and $q \in V_{io} \cup V_f \cup \{c : c \text{ is a constant symbol of } \sigma\}$, or
- a guess instruction of the form ‘GUESS v ’, where $v \in V_{io}$, or
- while loops of the form ‘WHILE t DO τ OD’, where t is quantifier-free $\text{FO}(\sigma)$ with free variables among $V_{io} \cup V_f$, and where τ is a sequence of instructions of one of the forms listed.

The first instruction is $\text{INPUT}(V_{io})$, and the last $\text{OUTPUT}(V_{io})$. All instructions begin a new line, and all, except while loops, take up only one line. While loops take up $1 + |\tau|$ lines, where $|\tau|$ is the number of lines in τ , in the obvious way. We consider sub-routines τ to be sequences of instructions of the types in the list, i.e. program schemes without an input instruction at the beginning and an output instruction at the end.

The program schemes $\rho \in \text{NPS}(m+1)$ are defined exactly as the schemes of $\text{NPS}(1)$, except that schemes $\rho' \in \text{NPS}(m')$ (for $m' \leq m$) may take the place of extensional relations in the tests in while loops. We define NPS to be $\bigcup_{i \in \omega} \text{NPS}(i)$.

In terms of semantics, the assignment instructions and while loops behave in the obvious way, and the guess instruction non-deterministically assigns an element of the universe to the variable in question. At the start of computation, the input-output variables are GUESSED, as just specified. A computation is deemed accepting if, and only if, it reaches the (final) OUTPUT line. It follows that non-accepting computations are forever trapped in while loops. Suppose a program scheme $\rho \in \text{NPS}(m)$ involves precisely i free variables z_1, \dots, z_i . Then, computing on a structure \mathcal{A} , we write $(\mathcal{A}, a_1, \dots, a_i) \models \rho$, or $\mathcal{A} \models \rho(a_1, \dots, a_i)$, iff ρ makes it to the output instruction when computing on \mathcal{A} under the free-variable assignment $(z_1, \dots, z_i) := (a_1, \dots, a_i) \in |\mathcal{A}|^i$.

Note that free variables may not be ‘used’ during computation, in that they can not have values assigned to them. However, input-output variables of schemes in $\text{NPS}(m)$ may appear as free variables in schemes of strictly lower strata that appear in tests in their while loops. In this manner, program schemes of $\text{NPS}(m)$ are evaluated ‘top-down’, entering sub-routines to evaluate any required tests involving such schemes of $\text{NPS}(m')$ (where $m' < m$).

Definition ([2]). Let the lines on \bar{z} denote an i -tuple, and the line on \bar{v} denote a j -tuple. Suppose the program scheme $\rho \in \text{NPS}(1)$ involves i free variables \bar{z} and j input-output variables \bar{v} . Then a configuration of ρ , computing on a structure \mathcal{A} , is an $(i+1+j)$ -tuple (\bar{z}, l, \bar{v}) giving the values of the free variables, the number of the line just executed, and the values of the input-output variables.

Each such program scheme $\rho \in \text{NPS}(1)$, computing over a structure \mathcal{A} of size n , gives rise to a graph:

- whose vertices are the $|\rho|.n^{(i+j)}$ possible configurations,
- and in which there is an edge (c, c') iff ρ , executing a single instruction, can move from configuration c to configuration c' .

It may be asking too much to specify this graph in quantifier-free **FO**, especially since \mathcal{A} may not have $|\rho|$ distinct constants to play the part of the line numbers. We

will actually specify a variant of it, namely the graph $\mathcal{G}_{\mathcal{A}}^{\rho}$, with $n^{i+|\rho|+j}$ vertices, where the $|\rho|$ -sub-tuple $\widehat{w} = (w_1, \dots, w_{|\rho|})$ represents a certain line according to the following scheme:

- if $w_1 = w_2$ then \widehat{w} represents line 1,
- if $w_1 \neq w_2$ but $w_2 = w_3$ then \widehat{w} represents line 2,
- if $w_1 \neq w_2, w_2 \neq w_3$ but $w_3 = w_4$ then \widehat{w} represents line 3,
- \vdots
- if $w_1 \neq w_2, \dots, w_{|\rho|-2} \neq w_{|\rho|-1}$ but $w_{|\rho|-1} = w_{|\rho|}$ then \widehat{w} represents line $|\rho| - 1$, and
- if $w_1 \neq w_2, \dots, w_{|\rho|-1} \neq w_{|\rho|}$ then \widehat{w} represents line $|\rho|$.

Let the the lines on \bar{z} denote an i -tuple, the line on \bar{u}, \bar{v} denote j -tuples, the hat on \widehat{w} denote a $|\rho|$ -tuple, and the line on \bar{x}, \bar{y} denote $(i + |\rho| + j)$ -tuples.

Proposition 5 ([2]). *Suppose $\rho \in \text{NPS}(1)$ is as in the definition, and that $\psi(\bar{x}, \bar{y})$ is a quantifier-free first order formula expressing the edge relation of $\mathcal{G}_{\mathcal{A}}^{\rho}$. The following are equivalent:*

- $\mathcal{A} \models \rho(\bar{z})$
-

$$\begin{aligned} \mathcal{A} \models & \exists w_1, \dots, w_{|\rho|} w_1 = w_2 \wedge \\ & \exists w'_1, \dots, w'_{|\rho|} w'_1 \neq w'_2 \wedge \dots \wedge w'_{|\rho|-1} \neq w'_{|\rho|} \wedge \\ & \exists \bar{u}, \bar{v} \text{ TC}[\lambda \bar{x}, \bar{y} \psi]((\bar{z}, w_1, \dots, w_{|\rho|}, \bar{u}), (\bar{z}, w'_1, \dots, w'_{|\rho|}, \bar{v})) \end{aligned}$$

Proof. Follows immediately from the definition of $\mathcal{G}_{\mathcal{A}}^{\rho}$, together with the existential semantics of $\text{NPS}(1)$. Note that the bizarre constraints on the w s are simply our means of encoding the first and last lines. As can be seen, we are not too interested in what the input-output variables are at the start and end of the computation, i.e. \bar{u} and \bar{v} , respectively. \square

Recall that the class $\text{NPS}_s(m)$ is as $\text{NPS}(m)$, but with a built-in successor available. The following gives us an idea as to the power of NPS.

Theorem 6 ([2]). *For $m \geq 1$, $\text{TC}^m[\mathbf{FO}] = \text{NPS}(m)$, and, consequently, $\text{TC}^*[\mathbf{FO}] = \text{NPS}$. Furthermore, for $m \geq 1$, $\text{TC}_s^m[\mathbf{FO}] = \text{NPS}_s(m) = \text{TC}_s^*[\mathbf{FO}] = \text{NPS}_s = \text{NL}$.*

Proof. The first part follows from the fact that there is a program scheme $\rho_{\text{TC}} \in \text{NPS}(1)$ that expresses the relation TC, combined with the previous proposition, by induction. The second part follows from the NL-completeness of the Transitive Closure problem. \square

Remark. The class NPS_s appears to be devoid of any memory, and it may seem surprising, in that light, that $\text{NPS}_s = \text{NL}$. However, NPS has memory, in the form of the constant number of input-output variables. Moreover, this constant number of variables may collectively attain $n^{|V_{io}|}$ values, computing on a structure \mathcal{A} of size n . This is of similar order to the number of different tape configurations on an NL-Turing Machine, which is $\log n \cdot |Q| \cdot |\Sigma|^{\log n}$, where Q is the set of states and Σ the alphabet. The NL-Turing Machine has constant alphabet and logarithmically-bounded number of tape squares, while the class NPS_s has linearly-sized alphabet and constant number of memory-variables.

2.2.2 Shorthands

We can build other useful instructions from those that we have, possibly requiring the introduction of additional new variables. Specifically:

- If \bar{p}, \bar{q} are j -tuples of variables or constants, then consider $\bar{p} = \bar{q}$ to be shorthand for $p_1 = q_1 \wedge \dots \wedge p_j = q_j$.
- If \bar{v} is a j -tuple of variables and \bar{q} is a j -tuple of variables or constants, then consider $\bar{v} := \bar{q}$ to be shorthand for $v_1 := q_1 ; \dots ; v_j := q_j$.
- Consider LOOP FOREVER to be shorthand for:

WHILE $v_1 = v_1$ DO OD.

- Consider IF t THEN DO τ FI to be shorthand for:

```

GUESS  $v_1, v_2$ 
WHILE  $v_1 = v_2$  DO LOOP FOREVER OD
WHILE  $v_1 \neq v_2 \wedge t$  DO  $\tau$ ;  $v_1 := v_2$  OD

```

Of course, it may come to pass that a computation entering an IF statement gets trapped in an endless loop. This may seem undesirable, but it does not affect us: owing to our existential semantics, we only require that *some* path leads through the conditional.

- Consider $v' \neq v$ (where v, v' are distinct variables) to be shorthand for:

```

GUESS  $v'$ 
IF  $v = v'$  THEN DO LOOP FOREVER FI

```

Sometimes we will want the computation to evaluate the disjunction of a fixed collection of possibilities. It may not be possible to write these directly in quantifier-free tests in WHILE loops. In the following, the labels $\text{word}_1, \dots, \text{word}_j$ act as local dummy ‘variables’.

- Let $\text{word}_1, \dots, \text{word}_j$ be words representing certain possibilities. Consider:

```

EITHER( $\text{word}_1, \dots, \text{word}_j$ )
IF  $\text{word}_1$  THEN DO  $\tau_1$  FI
  ⋮
IF  $\text{word}_j$  THEN DO  $\tau_j$  FI

```

to be shorthand for:

```

GUESS  $v_1, \dots, v_j$ 
IF  $v_1 = v_2$  THEN DO  $\tau_1$  FI
IF  $v_1 \neq v_2 \wedge v_2 = v_3$  THEN DO  $\tau_2$  FI
  ⋮
IF  $v_1 \neq v_2 \wedge \dots \wedge v_{j-2} \neq v_{j-1} \wedge v_{j-1} = v_j$  THEN DO  $\tau_{j-1}$  FI
IF  $v_1 \neq v_2 \wedge \dots \wedge v_{j-1} \neq v_j$  THEN DO  $\tau_j$  FI

```

The EITHER construction allows us to choose between any finite number of possibilities. Note that, in the EITHER shorthand, we have no need for an ‘Else’ construction, since all possibilities for the antecedent are covered. In all use of shorthands we will require that the variables we introduce in the longhand do not appear elsewhere in the program schemes involving those shorthands, lest we lose their information. This may ultimately require the introduction of new variables to our program schemes. We only need a fixed number of new variables for this, and we will usually be sloppy, omitting these variables when writing out program schemes involving shorthand.

2.2.3 Shorthands on successor structures

In the presence of a successor relation *succ*, we will use the following shorthands:

- $v' := cyc.succ(v)$ to be shorthand for:

```

IF  $v = max$  THEN DO  $v' = min$  OD
IF  $v \neq max$  THEN DO
  GUESS  $v'$ 
  IF  $v' \neq succ(v)$  THEN DO LOOP FOREVER FI FI

```

In contrast to *succ*, which is a partial function, *cyc.succ* is a total function. Moreover, it is a bijection.

- $v' := inv.cyc.succ(v)$ to be shorthand for:

```

GUESS  $v'$ 
IF  $v \neq cyc.succ(v')$  THEN DO LOOP FOREVER FI

```

- For variable j -tuples \bar{v}, \bar{v}' , consider $\bar{v}' := cyc.succ(\bar{v})$ to be shorthand for:

```

IF  $v_j \neq max$  THEN DO
   $(v'_1, \dots, v'_{j-1}) := (v_1, \dots, v_{j-1})$ 
   $v'_j := cyc.succ(v_j)$  FI

```

```

IF  $(v_j = \mathit{max}) \wedge (v_{j-1} \neq \mathit{max})$  THEN DO
   $(v'_1, \dots, v'_{j-2}) := (v_1, \dots, v_{j-2})$ 
   $v'_{j-1} := \mathit{cyc.succ}(v_{j-1})$ 
   $v'_j := \mathit{min}$  FI
  :
  :
IF  $(v_j = \mathit{max}) \wedge \dots \wedge (v_1 = \mathit{max})$  THEN DO
   $(v'_1, \dots, v'_j) := (\mathit{min}, \dots, \mathit{min})$  FI

```

- For variable j -tuples \bar{v}, \bar{v}' , consider $\bar{v}' := \mathit{inv.cyc.succ}(\bar{v})$ to be shorthand for:

```

IF  $v_j \neq \mathit{min}$  THEN DO
   $(v'_1, \dots, v'_{j-1}) := (v_1, \dots, v_{j-1})$ 
   $v'_j := \mathit{inv.cyc.succ}(v_j)$  FI
IF  $(v_j = \mathit{min}) \wedge (v_{j-1} \neq \mathit{min})$  THEN DO
   $(v'_1, \dots, v'_{j-2}) := (v_1, \dots, v_{j-2})$ 
   $v'_{j-1} := \mathit{inv.cyc.succ}(v_{j-1})$ 
   $v'_j := \mathit{max}$  FI
  :
  :
IF  $(v_j = \mathit{min}) \wedge \dots \wedge (v_1 = \mathit{min})$  THEN DO
   $(v'_1, \dots, v'_j) := (\mathit{max}, \dots, \mathit{max})$  FI

```

2.2.4 Adding a stack: introducing NPSS

We can increase the power of our program schemes by introducing certain types of memory. In [2], the authors considered adding a stack.

Definition (Syntax of NPSS[2]). The syntax of NPSS(1) is as that of NPS(1), with the addition of two new instructions:

- a push instruction 'PUSH v ', where

$$v \in V_{io} \cup V_f \cup \{c : c \text{ is a constant symbol of } \sigma\}, \text{ and}$$

- a pop instruction ‘ $v := \text{POP}$ ’, where $v \in V_{io}$.

Again, the program schemes of $\text{NPSS}(m+1)$ are those whose tests in while loops may include schemes from strictly lower strata as extensional relations.

For semantics, the push instruction should be viewed as pushing the value of the given variable (or constant) to the stack, and the pop instruction should be viewed as an assignment instruction removing the current top element of the stack. If the stack is empty, the pop instruction leaves its variable unchanged.

The following gives us an idea as to the power of NPSS.

Theorem 7 ([2]). *For $m \geq 1$, $\text{PS}^m[\mathbf{FO}] = \text{NPSS}(m)$, and $\text{PS}^*[\mathbf{FO}] = \text{NPSS}$. Furthermore, $\text{PS}_s^m[\mathbf{FO}] = \text{NPSS}_s(m) = \text{PS}_s^*[\mathbf{FO}] = \text{NPSS}_s = \text{P}$.*

2.3 Turing Machines

Turing Machines compute on strings and not structures. In order that we can consider the Turing complexity of problems $\Omega \subseteq \text{STRUC}(\sigma)$, we will need to have a standard encoding of structures over a signature σ . Let $\sigma = \langle R_1, \dots, R_j, c_1, \dots, c_{j'} \rangle$, where the arities of R_1, \dots, R_j are a_1, \dots, a_j respectively.

Over an ordered structure $\mathcal{A} \in \text{STRUC}(\sigma)$ of size n , we will code each R_i by a string $\text{bin}(R_i)$ over $\{0, 1\}$ of length n^{a_i} . For a number $0 \leq r \leq n^{a_i} - 1$, let \bar{r} be the a_i -tuple that represents r in n -ary. Since \mathcal{A} is ordered, this \bar{r} represents an a_i -tuple $\bar{r}_{\mathcal{A}}$ over \mathcal{A} . Let the r th² entry of $\text{bin}(R_i)$ be a 1 if $\bar{r}_{\mathcal{A}} \in R_i$, and a 0 otherwise. We code each c_i by a string $\text{bin}(c_i)$ of length n , as if c_i were a unary relation with one member. Finally, we consider $\text{bin}(\mathcal{A})$ to be the concatenation $\text{bin}(R_1) \dots \text{bin}(R_j) \text{bin}(c_1) \dots \text{bin}(c_{j'})$.

We consider Turing Machines to have a one-way infinite tape, finite state set Q and uniform alphabet $\Sigma = \{\text{zero}, \text{one}, \text{blank}\}$. The read/write head is initially over square 1. Given a (non-deterministic) Turing Machine T and a string $w \in \{0, 1\}^*$,

²This should really be $r+1$, since otherwise we would be considering the first entry of $\text{bin}(R_i)$ to be indexed by the number zero. This is an occupational hazard of variously considering the set \mathbb{Z}_n to be $\{1, \dots, n\}$ or $\{0, \dots, n-1\}$. We largely use the former for the chapters on program schemes, and the latter for the chapters on constraint satisfaction.

we write $T \downarrow w$, iff T enters the accept state, at some point in its computation, when it is given input w over squares 1 to $|w|$ with all other squares blank.

We say that a (non-deterministic) Turing Machine T accepts a problem $\Omega \subseteq \text{STRUC}(\sigma)$ iff, for all structures $\mathcal{A} \in \text{STRUC}(\sigma)$, and for all orderings of \mathcal{A} , we have:

$$T \downarrow \text{bin}(\mathcal{A}) \Leftrightarrow \mathcal{A} \in \Omega$$

Chapter 3

Adding a Priority Queue

We now consider the situation where we have a priority queue for memory. A priority queue allows us to send elements to memory tagged with a numerical weight. We are free to choose from a range of weights polynomially-bounded in the size of the structure on which we are computing, but we may only retrieve from the maximal (non-empty) weight. We will consider a variety of semantic variations, and will, therefore, be no more specific at this point as to the properties of the priority queue. However, we are in a very different situation from before, because now we deal with both elements of structures and numbers. Such is the power of the inclusion of numbers, that we will find ourselves dealing with Turing Machines and complexity classes directly, as opposed to Lindström logics that capture complexity classes only on ordered structures. We needed free variables in NPS and NPSS in order to build the stratification within those hierarchies. We do not need that variety of stratification here. Therefore, since we are only concerned with decision problems, we will have no need for free variables here, and we dispense with them for the sake of a simpler exposition.

Since we will deal in a range of queue weights that is polynomially-bounded in the size of the structure \mathcal{A} on which we are computing, we will have interest in the numbers $1, \dots, n$, where $n = \|\mathcal{A}\|$. We allow ourselves the first and last of these, 1 and n , as constants that we may refer to by name.

Definition. For each $k \geq 0$, the program schemes of NPSPQ(k), over a signature σ , involve two finite sets of variables, a set V of element variables and a set N of numeric variables. A program scheme $\rho \in \text{NPSPQ}(k)$ consists of a finite sequence of instructions, where each instruction, other than the first and last, is of one of the following forms:

- an assignment instruction of the form ‘ $p := q$ ’, where $p \in V$ and $q \in V \cup \{c : c \text{ is a constant symbol of } \sigma\}$ or $p \in N$ and $q \in N \cup \{1, n\}$
- a guess instruction of the form ‘GUESS v ’, where $v \in V$
- an increase (numeric successor) instruction ‘INCR m ’, where $m \in N$
- a push instruction ‘PUSH v, m_{i1}, \dots, m_{ik} ’¹ where, $v \in V$ and $m_{i1}, \dots, m_{ik} \in N$

¹When $k = 0$ there are no m 's.

- a pop instruction ‘ $v := \text{POP}$ ’ where $v \in V$.
- while loops of the form ‘ $\text{WHILE } t \text{ DO } \tau \text{ OD}$ ’, where t is quantifier-free $\text{FO}(\sigma)$ with free variables among V or quantifier-free $\text{FO}(\langle 1, n \rangle)$ whose free variables are among N , and where τ is a sequence of instructions of one of the forms listed.

The first instruction is $\text{INPUT}(V, N)$, and the last $\text{OUTPUT}(V, N)$. We further define NPSPQ to be $\cup_{k \in \omega} \text{NPSPQ}(k)$.

As hinted at before, the stratification here – dimension k of the queue – is quite different from the stratification we have seen thus far, which was based on nestings of negation. With a priority queue, we will have sufficient computational power [at what would have been the first level of that nesting] to not require stratification.

The assignment and guess instructions, and the while loops, behave as before. Observe that in each case there are two modes of use: one relates to elements, the other to numbers. We do not allow the guessing of numeric variables simply because it is unlikely to be useful. The instruction $\text{INCR } m$ increases the number m by one, under the convention that $\text{INCR } n$ is 1. This ensures that INCR is a function, like cyc.succ , and in contrast to succ . The push instruction sends the element in question to the priority queue tagged with weight k -tuple (m_{i1}, \dots, m_{ik}) , *i.e.* the current value of those numeric variables. It is for this reason that k is considered the dimension of the queue. We will consider a number of alternative semantics for the pop instruction:

- u The pop removes, deterministically, the last element to be sent to the queue at whatever is the maximal non-empty weight. This semantics leads to a potentially unbounded queue size, and hence will be referred to as semantics ‘ u ’.
- b The pop removes, deterministically, the last element to be sent to the queue at whatever is the maximal non-empty weight, and then scrubs all other entries at that weight. This is equivalent to the condition that the queue has only one space at each weight, *i.e.* new pushes would overwrite. This semantics leads to a (polynomially-)bounded queue size, and hence will be referred to as semantics ‘ b ’.

$u+$ As with ‘ u ’, but the maximal weight is also returned. This requires pop syntax ‘ $p, m_{i1}, \dots, m_{ik} =: \text{POP}$ ’. We refer to this as semantics ‘ $u+$ ’.

$b+$ As with ‘ b ’, but the maximal weight is also returned. This also requires pop syntax ‘ $p, m_{i1}, \dots, m_{ik} =: \text{POP}$ ’. We refer to this as semantics ‘ $b+$ ’.

As before, the pop instruction leaves its variable unchanged if the queue is empty. Also as before, the INPUT instruction non-deterministically assigns elements of the structure to V . The numeric variables N are set initially to 1. Again, we consider an accepting computation of a program scheme ρ on a structure \mathcal{A} to be any one that reaches OUTPUT, and we denote this $\mathcal{A} \models \rho$. We refer to each of the four alternative semantics above specifically by superscript, e.g. $\text{NPSPQ}^{b+}(k)$. Again, we will refer to the classes endowed with successor with the subscript s , e.g. $\text{NPSPQ}_s^{b+}(k)$.

We will use the following shorthands, specific to schemes of NPSPQ:

- Consider FOR $m = m'$ TO m'' DO τ NEXT to be shorthand for:

```

 $m := m'$ 
WHILE  $m \neq m''$  DO
   $\tau$ 
  INCR  $m$  OD
 $\tau$ .

```

- Consider DECR m to be shorthand for:

```

 $m' := 1; m'' := m'$ 
INCR  $m'$ 
WHILE  $m' \neq m$  DO
   $m'' := m'$ 
  INCR  $m'$  OD
 $m := m''$ .

```

- Consider $\text{FOR } m = m' \text{ DOWNTO } m'' \text{ DO } \tau \text{ NEXT}$ to be shorthand for:

```

 $m := m'$ 
WHILE  $m \neq m''$  DO
   $\tau$ 
  DECR  $m$  OD
 $\tau$ .

```

Note that FOR loops are inclusive with respect to their limits.

- Consider $m := m' + m''$ to be shorthand for:

```

 $m := m'$ 
FOR  $m''' = 1$  TO  $m''$  DO INCR  $m$  OD.

```

- Consider $m := m' - m''$ to be shorthand for:

```

 $m := 1$ 
DECR  $m$ ; DECR  $m$ 
FOR  $m''' = m'$  DOWNTO  $m''$  DO INCR  $m$  OD.

```

Henceforth, we will feel free to put arithmetic terms such as $m' - m''$ as limits in FOR loops.

Let \bar{m} be a j -tuple (m_1, \dots, m_j) of numeric variables. Consider $\text{INCR } \bar{m}$ to be shorthand for:

```

IF  $(m_j = n) \wedge \dots \wedge (m_2 = n)$  THEN DO INCR  $m_1; \dots; \text{INCR } m_j$  OD
IF  $(m_j = n) \wedge \dots \wedge (m_3 = n) \wedge (m_2 \neq n)$  THEN DO INCR  $m_2; \dots; \text{INCR } m_j$  OD
 $\vdots$ 
IF  $m_j \neq n$  THEN DO INCR  $m_j$  OD

```

Let 1^j and n^j be the j -tuples of 1s and n s, respectively. We have that $\text{INCR } \bar{m}$ returns the lexicographic next number, subject to the convention that $\text{INCR } (n^j) = (1^j)$. Define $\text{DECR } \bar{m}$ analogously. Sums and differences of j -tuples \bar{m} and \bar{m}' are defined in the natural way. However, we will insist that we never attempt the sum or difference of a j -tuple and j' -tuple when $j \neq j'$.

Computing over a structure \mathcal{A} , with $||\mathcal{A}|| = n$, we find we have been granted basic modulo n arithmetic. This is ostensibly weaker than an ordering of the elements of \mathcal{A} , but it will ultimately allow us to build such an order.

Remark. For each j , 1^j represents the number 1, in modulo n^j arithmetic, and n^j represents the additive identity (zero). So, for example, $\text{INCR } 1^j = 1^j + 1^j = (1, \dots, 1, 2)$, where 2 is $\text{INCR } 1$.

3.1 A single weight: $k = 0$

The bottom level in our apparent hierarchy merits brief attention. In the presence of a single weight, it is apparent that $b+$ (respectively, $u+$) is no stronger than b (respectively, u).

Lemma 8. $\text{NPSPQ}_s^b(0) = \text{NPS}_s(1) = \text{NL}$

Proof. We already have the second equality; we prove the first.

$(\text{NPS}_s(1) \subseteq \text{NPSPQ}_s^b(0))$. Trivially, we will have for any $\rho \in \text{NPS}_s(1)$, that also $\rho \in \text{NPSPQ}_s^b(0)$.

$(\text{NPSPQ}_s^b(0) \subseteq \text{NPS}_s(1))$. The priority queue may hold only one element at any time, and, as such, behaves like an extra element variable. Furthermore, the ability to count in $\text{NPSPQ}_s^b(0)$ may be simulated by the successor relation of $\text{NPS}_s(1)$. Specifically, if $\rho \in \text{NPSPQ}_s^b(0)$ involves $|V|$ element variables and $|N|$ numeric variables, then we construct $\rho' \in \text{NPS}_s(1)$ with $|V| + |N| + 3$ variables. Our simulation is made somewhat more complicated by our convention that popping from an empty queue leaves the variable unchanged: this is why we need the extra variables v', v'' (we use $v' = v''$ to signify that the queue is non-empty). We construct ρ' thus:

INPUT $(v_1, \dots, v_{|V|}, v_{|V|+1}, \dots, v_{|V|+|N|}, v_{\text{queue}}, v', v'')$
 $v' \neq v''$
 τ_{sim}
 OUTPUT $(v_1, \dots, v_{|V|}, v_{|V|+1}, \dots, v_{|V|+|N|}, v_{\text{queue}}, v', v'')$

Where τ_{sim} is the body of ρ (i.e. with the input and output lines removed), with the following substitutions:

- Convert all instances of variables m_i to variables $v_{|V|+i}$.
- Convert all instances of the numeric constant 1 (respectively, n) to the element constant min (respectively, max).
- Convert all instances of ‘INCR m_i ’ to: ‘ $v_{|V|+i} := cyc.succ(v_{|V|+i})$ ’.
- Convert all instances of ‘PUSH v_i ’ to: ‘ $v_{queue} := v_i ; v' := v''$ ’.
- Convert all instances of ‘ $v_i := POP$ ’ to:

IF $v' = v''$ THEN DO
 $v_i := v_{queue} ; v' := v''$ FI.

It should be clear that we have, for all structures \mathcal{A} , $\mathcal{A} \models \rho$ iff $\mathcal{A} \models \rho'$. □

Lemma 9. $NPSPQ_s^u(0) = NPSS_s(1) = P$

Proof. We already have the second equality; we prove the first.

($NPSS_s(1) \subseteq NPSPQ_s^u(0)$). Trivially, any $\rho \in NPSS_s(1)$ is such that $\rho \in NPSPQ_s^u(0)$.

($NPSPQ_s^u(0) \subseteq NPSS_s(1)$). The priority queue’s single weight here acts as a stack. We may use a similar, though simpler, reduction to that of the previous lemma: we no longer need the variables v', v'' in any capacity, and we leave instances of ‘PUSH v_i ’ and ‘ $v_i := POP$ ’ in ρ unchanged in ρ' . □

Remark. The previous lemmas are somewhat misleading. We had provision for free variables in NPS and NPSS, but we have none in NPSPQ. The previous results, therefore, can only authoritatively refer to *sentences* of NPS_s and $NPSS_s$, i.e. those schemes without free variables. The only reason we omit free variables from NPSPQ is to simplify our exposition. The previous lemmas would hold in generality, if we were to allow free variables in NPSPQ.

When we are deprived of the successor relation, we find $NPSPQ^b(0) \not\subseteq NPS(1)$ (respectively, $NPSPQ^u(0) \not\subseteq NPSS(1)$), since the parity problem may be expressed in the former, through counting, but not in the latter. Of course, we will have the inclusions

- $NPSPQ^b(0) \subseteq NPSPQ_s^b(0) = NPS_s(1)$ and

- $\text{NPSPQ}^u(0) \subseteq \text{NPSPQ}_s^u(0) = \text{NPSS}_s(1)$.

We conjecture that these inclusions are proper, and in particular that $\text{NPSPQ}^u(0)$ is contained within $\text{LFP} + \text{COUNT}[\mathbf{FO}]$, which is known to be strictly contained in P [31].

3.2 The Hamilton Path problem is in NPSPQ

We proceed by examining the power of the program schemes of NPSPQ and, in particular, one of their number ρ_{HP} with the ability to accept the NP-complete Hamilton Path problem. The Hamilton Path problem² HP is exactly the class of digraphs that have a directed path containing each vertex exactly once. The following is part of the program scheme $\rho_{HP} \in \text{NPSPQ}^u(2)$ that non-deterministically builds an order on such a structure.

1. INPUT(v_1, v_2, m_1, m_2, m_3)
2. FOR $m_1 = 1$ TO n DO
3. GUESS v_1
4. FOR $m_2 = 1$ TO n DO
5. PUSH v_1, m_2, m_1 NEXT NEXT

Our method is simple enough: we produce n copies of n guessed vertices, each copy occupying weights $(i, 1)$ to (i, n) for $1 \leq i \leq n$. These could be genuine orders, but only if we haven't picked some element twice. After line 5, the queue

²In contrast to TC, AR, PS etc., which we initially defined as global relations, we define HP as a decision problem.

looks like this (entry followed by weight):

$$\begin{array}{c}
 \underbrace{x_n}_{(n,n)} \\
 \vdots \\
 \underbrace{x_1}_{(n,1)} \\
 \vdots \\
 \vdots \\
 \underbrace{x_n}_{(1,n)} \\
 \vdots \\
 \underbrace{x_1}_{(1,1)}
 \end{array}$$

At line 6, we proceed by consuming $n - 1$ copies of our n guessed elements to see if some element is repeated. First we look at the last copy and last element, x_n , stored at weight (n, n) , then we look through $(n, n - 1)$ to $(n, 1)$, elements x_{n-1}, \dots, x_1 , to see if it is repeated. Next, lines 11-13, we remove the unneeded (already checked) element x_n at weight $(n - 1, n)$ and repeat the process for $(n - 1, n - 1)$ to $(n - 1, 1)$. If we do this $n - 1$ times, finding no element guessed twice, then we know we do indeed have a genuine order left in weights $(1, 1)$ to $(1, n)$. (If the first element had been repeated we would have already discovered that; we only need $n - 1$ iterations here.)

6. FOR $m_1 = 1$ TO $n - 1$ DO
7. $v_2 := \text{POP}$
8. FOR $m_2 = m_1$ TO $n - 1$ DO
9. $v_1 := \text{POP}$
10. IF $v_1 = v_2$ THEN DO LOOP FOREVER FI NEXT
11. IF $m_1 \neq n - 1$ THEN DO
12. FOR $m_3 = 1$ TO m_1 DO
13. $v_2 := \text{POP}$ FI NEXT NEXT NEXT

For any computation that gets past line 13 the queue will look like

$$\begin{array}{c} \underbrace{x_n}_{(1, n)} \\ \vdots \\ \underbrace{x_1}_{(1, 1)} \end{array}$$

where we know that x_1, \dots, x_n is an ordering of the vertices. We will now search along it for a Hamilton path:

14. $v_1 := \text{POP}$
15. FOR $m_1 = 1$ TO $n - 1$ DO
16. $v_2 := v_1$
17. $v_1 := \text{POP}$
18. IF $\neg E(v_1, v_2)$ THEN DO LOOP FOREVER FI NEXT
19. OUTPUT(v_1, v_2, m_1, m_2, m_3)

It is because we can non-deterministically guess all orderings that there will be an accepting computation if, and only if, the structure has a Hamilton path. For all digraphs \mathcal{G} , we will have $\mathcal{G} \models \rho_{HP}$ iff $\mathcal{G} \in HP$.

The scheme ρ_{HP} computes in such a way that, on all inputs \mathcal{G} , it only uses any weight at most once. Consequently, ρ_{HP} also accepts the Hamilton Path problem under semantics b . Clearly, ρ_{HP} can undergo minor syntactic changes to produce a program scheme that accepts HP for semantics $u+$ and $b+$, too.

3.3 $\text{NPSPQ}_s^b \subseteq \text{NPspace}$

With the polynomially-bounded memory of NPSPQ_s^b , the following is almost immediate.

Proposition 10. $\text{NPSPQ}_s^b \subseteq \text{NPspace}$.

Sketch Proof. The proof is by simulation. For $\rho \in \text{NPSPQ}_s^b$ we will construct a non-deterministic Turing Machine T , together with an exhibited bound l , such

that, for all structures \mathcal{A} (of size n), and all orderings of \mathcal{A} , the following are equivalent:

- $\mathcal{A} \models \rho$.
- $T \downarrow \text{bin}(\mathcal{A})$.
- $T \downarrow \text{bin}(\mathcal{A})$ with the read/write head never leaving the first n^l squares.

Note that equivalence of the last two guarantees that T is an NPspace machine since there exists some l' (dependent on the maximum relation arity of the signature σ) s.t. $|\text{bin}(\mathcal{A})| = O(n^{l'})$.

If $\rho \in \text{NPSPQ}_s^b(k)$ and involves j program scheme variables (element or numeric) then we need to record at most $n^k + j$ items, corresponding to the entries on the priority queue and the assignments of the variables of ρ , at any point of the simulation. Each of these $n^k + j$ items may take at most n possible values, so each of these items may be written on T 's tape in $\log(n)$ squares. We do not give full details of T 's simulation, but note that the amount of tape space required to hold all these $n^k + j$ items is $O((n^k + j) \log(n))$. It follows that we may take $l := k + j + 1$. □

Corollary. $\text{NPSPQ}^b \subseteq \text{NPspace}$.

Proof. The inclusion $\text{NPSPQ}^b \subseteq \text{NPSPQ}_s^b$ is trivial. □

3.4 Expanding alphabets

At present there are precisely n distinct symbols that we can send to the queue, namely the elements of the structure on which we are computing. However, we can expand this alphabet by always pushing and popping j -tuples, instead of single variables. In this way we potentially increase our working alphabet to n^j symbols.

Let \bar{v} be a j -tuple of variables. We work in semantics u , but our results apply to semantics $u+$, and also to NPSS. A similar method may be used for semantics b and $b+$, although at the cost of more weights. The method by which these results for semantics u transfer to semantics b will be explored later.

- Consider ‘ $\text{PUSH } \bar{v}, \bar{m}$ ’ to be shorthand for ‘ $\text{PUSH } v_j, \bar{m} ; \dots ; \text{PUSH } v_1, \bar{m}$ ’.
- Consider ‘ $\bar{v} := \text{POP}$ ’ to be shorthand for ‘ $v_1 := \text{POP} ; \dots ; v_j := \text{POP}$ ’ (note the reverse order).

By these methods, we can push and pop tuples as if they were single elements. We can now set up special symbols by the use of a certain convention. Suppose we want i special symbols M_1, \dots, M_i , then we can achieve this, in a rather sloppy manner, by always pushing and popping $(i + 1)$ -tuples (v_1, \dots, v_{i+1}) , using the convention:

- (v_1, \dots, v_{i+1}) where $v_1 = v_2$ is the element v_1 .
- (v_1, \dots, v_{i+1}) where $v_1 \neq v_2 \wedge v_2 = v_3$ is the symbol M_1 .
- \vdots
- (v_1, \dots, v_{i+1}) where $v_1 \neq v_2 \wedge \dots \wedge v_{i-1} \neq v_i \wedge v_i = v_{i+1}$ is the symbol M_{i-1} .
- (v_1, \dots, v_{i+1}) where $v_1 \neq v_2 \wedge \dots \wedge v_i \neq v_{i+1}$ is the symbol M_i .

Note that all $(i + 1)$ -tuples are defined. Henceforth we will assume a finite set of special symbols at our disposal.

Given a program scheme in which we are always pushing and popping j -tuples, we may drop the line over the variables, and use that line only when referring to some j' -tuple of ‘variables’ each of which is actually a j -tuple of real variables. This should not cause too much confusion. This will result in our having variables v that can hold values that do not represent actual elements of the universe on which we are computing. Such special characters will constitute the symbol set Λ .

3.5 Pushing and Popping Numbers in $\text{NPSPQ}_s^b / \text{NPSPQ}_s^u$.

For program schemes of NPSPQ_s^u or NPSPQ_s^b , consider ' $v := \text{element}(m)$ ' to be shorthand for

```

v := min
FOR m' = 1 TO m - 1 DO
  GUESS v'
  IF v' ≠ succ(v) THEN DO LOOP FOREVER FI
  v := v' NEXT

```

and ' $m := \text{position}(v)$ ' to be shorthand for

```

v' := min
m := 1
WHILE v' ≠ v DO
  GUESS v''
  IF v'' ≠ succ(v') THEN DO LOOP FOREVER FI
  INCR m
  v' := v'' OD

```

The instruction $v := \text{element}(m)$ assigns to v the m th element of the universe, conversely the instruction $m := \text{position}(v)$ assigns to m the position of the element v in that order.

For j -tuples $\bar{m} = (m_1, \dots, m_j)$ and $\bar{v} = (v_1, \dots, v_j)$:

- Consider $\bar{v} := \text{element}(\bar{m})$ to be shorthand for $v_1 := \text{element}(m_1) ; \dots ; v_j := \text{element}(m_j)$.
- Consider $\bar{m} := \text{position}(\bar{v})$ to be shorthand for $m_1 := \text{position}(v_1) ; \dots ; m_j := \text{position}(v_j)$.

3.5.1 $\text{NPSPQ}_s^{u+}(k) = \text{NPSPQ}_s^u(k)$.

Lemma 11. $\text{NPSPQ}_s^{u+}(k) \subseteq \text{NPSPQ}_s^u(k)$.

Proof. The proof is by simulation. For all $\rho \in \text{NPSPQ}_s^{u+}(k)$ we construct a $\rho' \in \text{NPSPQ}_s^u(k)$ such that, for all structures \mathcal{A} , we have $\mathcal{A} \models \rho$ iff $\mathcal{A} \models \rho'$.

The program scheme ρ' will involve all the variables of ρ together with a new k -tuple of element variables $\overline{v_m}$. Where ρ pushes and pops single variables, ρ' will always push and pop $(k+1)$ -tuples of variables (of which the trailing k -tuple contains the weight).

- Convert all instances of ‘PUSH v, \overline{m} ’ , in ρ , to the following in ρ' :

$$\begin{aligned} \overline{v_m} &:= \text{element}(\overline{m}) \\ \text{PUSH } (v, \overline{v_m}), \overline{m} \end{aligned}$$

- Convert all instances of ‘ $v, \overline{m} := \text{POP}$ ’ , in ρ , to the following in ρ' :

$$\begin{aligned} (v, \overline{v_m}) &:= \text{POP} \\ \overline{m} &:= \text{position}(\overline{v_m}) \end{aligned}$$

□

Corollary. $\text{NPSPQ}_s^{u+}(k) = \text{NPSPQ}_s^u(k)$

Proof. The converse inclusion $\text{NPSPQ}_s^u(k) \subseteq \text{NPSPQ}_s^{u+}(k)$ is trivial. □

3.5.2 $\text{NPSPQ}_s^{b+} = \text{NPSPQ}_s^b$.

Lemma 12. $\text{NPSPQ}_s^{b+}(k) \subseteq \text{NPSPQ}_s^b(2k)$.

Proof. The proof is broadly similar to that of the previous lemma, but we will require more than single weights to store the $(k+1)$ -tuples of that proof. For all $\rho \in \text{NPSPQ}_s^{b+}(k)$ we construct a $\rho' \in \text{NPSPQ}_s^b(k)$ such that, for all structures \mathcal{A} , we have $\mathcal{A} \models \rho$ iff $\mathcal{A} \models \rho'$.

The program scheme ρ' will involve all the variables of ρ together with a new k -tuple of element variables $\overline{v_m} = (v_m^1, \dots, v_m^k)$.

- Convert all instances of ‘PUSH v, \bar{m} ’ , in ρ , to the following in ρ' :

PUSH $v, (\bar{m}, 1^k)$
PUSH $v_m^k, (\bar{m}, 1^{k-1}, n)$
 \vdots
PUSH $v_m^1, (\bar{m}, n, 1^{k-1})$

- Convert all instances of ‘ $v, \bar{m} := \text{POP}$ ’ , in ρ , to to the following in ρ' :

$(v, \bar{v}_m) := \text{POP}$
 $\bar{m} := \text{position}(\bar{v}_m)$

□

Corollary. $\text{NPSPQ}_s^{b+} = \text{NPSPQ}_s^b$

Proof. The inclusion $\text{NPSPQ}_s^b(k) \subseteq \text{NPSPQ}_s^{b+}(k)$ is trivial. □

3.6 $\text{NPSPQ}^b(k) \subseteq \text{NPSPQ}^u(k)$

Intuitively, semantics u appears at least as strong as semantics b . It is relatively straightforward to prove this.

Lemma 13. $\text{NPSPQ}^b(k) \subseteq \text{NPSPQ}^u(k)$.

Proof. The proof is by simulation. For all $\rho \in \text{NPSPQ}^b(k)$ we construct a $\rho' \in \text{NPSPQ}^u(k)$ such that, for all structures \mathcal{A} , we have $\mathcal{A} \models \rho$ iff $\mathcal{A} \models \rho'$.

The program scheme ρ' will involve all the variables of ρ together with two new k -tuples of numeric variables \bar{m}', \bar{m}'' and a new element variable v' . Assume that M is a special marker symbol not used by ρ . Build ρ' from ρ by adding the following lines to the beginning (after INPUT):

FOR $\bar{m}' = 1^k$ TO n^k DO
PUSH M, \bar{m}' NEXT

This sends a copy of the marker M to every weight on the queue. Finally, convert all instances of ‘ $v, \bar{m} := \text{POP}$ ’ , in ρ , to the following in ρ' .

```

 $\bar{m}' := n^k$ 
 $v' := \text{POP}$ 
WHILE  $\bar{m}' \neq 1^k \wedge v' = M$  DO
  DECR  $\bar{m}'$ 
   $v' := \text{POP}$  OD
IF  $v' \neq M$  THEN DO
   $v := v'$ 
  WHILE  $v' \neq M$  DO
     $v' := \text{POP}$  OD FI
FOR  $\bar{m}'' = \bar{m}'$  TO  $n^k$  DO
  PUSH  $M, \bar{m}''$  NEXT

```

The given subroutine counts, top-down, the number of empty weights in the queue of ρ – these contain just M in the queue of ρ' . When it finds something other than an M , it stores this in v then removes everything else at that weight, *i.e.* until it reaches another M . The situation where the queue of ρ is empty is dealt with by the conditional $v' \neq M$ in the sixth line. Finally, an M is returned to each of the weights of the queue of ρ' above and including the weight of the retrieved element. \square

Corollary.

- $\text{NPSPQ}_s^b(k) \subseteq \text{NPSPQ}_s^u(k)$
- $\text{NPSPQ}^{b^+}(k) \subseteq \text{NPSPQ}^{u^+}(k)$
- $\text{NPSPQ}_s^{b^+}(k) \subseteq \text{NPSPQ}_s^{u^+}(k)$

Proof. Our proof is equally valid for these statements. \square

Remark. It may seem that our method is unnecessarily complicated. In simulating semantics b with semantics u , when popping from the queue at a certain weight, why do we not simply then pop everything else off at that same weight (foregoing any need for the marker M)? This method would generate very simple proofs for the last two statements of the corollary, and a relatively easy proof of the first statement of the corollary. However, it would not easily be applied in the case of the statement of the theorem.

3.7 NPSPQ_s^u(k) ⊆ NPSPQ^u(k + 2).

Lemma 14. $NPSPQ_s^u(k) \subseteq NPSPQ^u(k+2)$.

Proof. We prove the inclusion by simulating the successor relation. We take any scheme $\rho \in NPSPQ_s^u(k)$, and construct a scheme $\rho' \in NPSPQ^u(k+2)$ such that, for all structures \mathcal{A} , $\mathcal{A} \models \rho$ iff $\mathcal{A} \models \rho'$.

Assume, without loss of generality, that ρ involves element variables V and numeric variables N , and that $v_1, v_2 \notin V$ and $m_1, m_2, m_3 \notin N$. Given ρ we will construct ρ' by adding a special start-routine, a special end-routine, and amending push and pop instructions, as well as successor tests in while loops.

Let $V' = V \cup \{v_1, v_2\}$ and $N' = N \cup \{m_1, m_2, m_3\}$. Then ρ' will be:

```

INPUT( $V', N'$ )
 $\tau_{start}; \tau_\rho; \tau_{end}$ 
OUTPUT( $V', N'$ )

```

We will now meet the sub-routines $\tau_{start}, \tau_\rho, \tau_{end}$, and explain why each one performs the function that will be claimed of it.

3.7.1 Start-routine: τ_{start}

We will add a sub-routine τ_{start} to the start, that builds an order over \mathcal{A} 's n elements. We will simply guess an order, as we did in ρ_{HP} , and we will put this putative order in the weights $(1^k, 1, 1)$ to $(1^k, 1, n)$. τ_{start} will first send a special marker symbol M to each of these weights via:

```

FOR  $m_1 = 1$  TO  $n$  DO
  PUSH  $M, (1^k, 1, m_1)$  NEXT

```

We then add lines 2-13 of the scheme ρ_{HP} that solved the Hamilton Path problem, with the proviso that weight tuples (m_1, m_2) in ρ_{HP} become $(1^k, m_1, m_2)$ in τ_{start} .

Any computation that gets through τ_{start} will leave the queue looking like:

$$\begin{array}{l} \overbrace{x_n} \quad (1^k, 1, n) \\ M \quad (1^k, 1, n) \\ \vdots \\ x_1 \quad (1^k, 1, 1) \\ \underbrace{M} \quad (1^k, 1, 1) \end{array}$$

where x_1, \dots, x_n is an ordering of the elements of \mathcal{A} .

3.7.2 Simulation of ρ : τ_ρ .

The main body of ρ' , the sub-routine τ_ρ , is that bit that actually simulates ρ . It will use higher weights of the form (n, n, \bar{m}) , where the line on \bar{m} specifies a k -tuple. Before we get to the main simulation, we will push a special marker symbol M' to weight $(n, n, 1^k)$ to ensure that we never stray into the lower weights, in which the order is contained, during the simulation. Thus:

PUSH $M', (n, n, 1^k)$

For the actual simulation:

- Convert all instances of ‘PUSH v, \bar{m} ’ to: ‘PUSH $v, (n, n, \bar{m})$ ’.
- Convert all instances of ‘ $v := \text{POP}$ ’ to the following in ρ' :

$v_1 := v$
 $v := \text{POP}$
 IF $v = M'$ THEN DO PUSH $v, (n, n, 1^k)$ FI
 $v := v_1$

The simulation of pop is rather complicated because the original program scheme ρ must leave a pop unchanged when the queue is empty. But when the queue associated with ρ is empty, the queue associated with ρ' still contains entries in the lower weights beneath M' .

We may also have to evaluate quantifier-free successor queries of the form $v' = succ(v)$, that might appear in a test for a while loop, immediately before the test of that while loop. Let Φ be a propositional formula that involves the atom $v' = succ(v)$:

- Convert all instances of **WHILE** $\Phi(v' = succ(v))$ **DO** τ **OD** to:

τ_{succ} ; **WHILE** $\Phi(m_2 = n)$ **DO** τ ; τ_{succ} **OD**.

Where τ_{succ} is the sub-routine:

```

 $m_1 := 1; m_2 := 1$ 
WHILE  $m_1 \neq n$  DO
  GUESS  $v_1$ 
  PUSH  $v_1, (1^k, 1, m_1)$ 
  IF  $v_1 = v$  THEN DO
    INCR  $m_1$ 
    GUESS  $v_1$ 
    PUSH  $v_1, (1^k, 1, m_1)$ 
    IF  $v_1 = v'$  THEN DO  $m_2 := n$  FI FI
  INCR  $m_1$  OD

```

Observe that v' and v are free in the sub-routine. What is happening in the while loop in the added sub-routine is that we are guessing what we hope to be an order. If it is the order that we guessed at the start, then $m_2 = n$ iff $v' = succ(v)$. We will check later that all these guessed ‘orders’ are not only genuine orders, but also the same as the first. In this manner, each instance of $v' = succ(v)$ in Φ becomes a test of $m_2 = n$. There may be any constant number of tests of the form $v' = succ(v)$, involving different variable pairs: each one of these will cause its own copy of the τ_{succ} sub-routine to appear before, and in, the while loop.

3.7.3 End-routine: τ_{end}

Once the simulation of ρ is accomplished we will want access to the lower weights to verify that these ‘orders’ we have been guessing are uniformly the same. We

will want to pop everything on the queue down to, and including, the marker M' . τ_{end} will therefore begin:

WHILE $v_1 \neq M'$ DO $v_1 := \text{POP OD}$.

At this point the queue will look like:

$$\begin{array}{c}
 \overbrace{y_{n,s}} \\
 \vdots \\
 y_{n,1} \quad (1^k, 1, n) \\
 x_n \\
 \underbrace{M} \\
 \vdots \\
 \vdots \\
 \overbrace{y_{1,s}} \\
 \vdots \\
 y_{1,1} \quad (1^k, 1, 1) \\
 x_1 \\
 \underbrace{M}
 \end{array}$$

where s is the number of times that we needed to check successor queries in while loops. We already know that x_1, \dots, x_n is an order of the elements – what we must now check is that:

- $x_1 = y_{1,1} = \dots = y_{1,s}$
- \vdots
- $x_n = y_{n,1} = \dots = y_{n,s}$

so τ_{end} concludes:

```

FOR  $m_1 = 1$  TO  $n$ 
   $v_1 := \text{POP}$ 
  WHILE  $v_1 \neq M$  DO
     $v_2 := v_1$ 
     $v_1 := \text{POP}$ 
  IF  $v_1 \neq v_2$  THEN DO LOOP FOREVER FI OD NEXT

```

□

Corollary. $\text{NPSPQ}^{u+} = \text{NPSPQ}_s^{u+} = \text{NPSPQ}_s^u = \text{NPSPQ}^u$.

Proof. $\text{NPSPQ}^{u+} \subseteq \text{NPSPQ}_s^{u+}$ is trivial; $\text{NPSPQ}_s^{u+} \subseteq \text{NPSPQ}_s^u$ was proved in Lemma 11; $\text{NPSPQ}_s^u \subseteq \text{NPSPQ}^u$ was proved in the previous lemma; and, $\text{NPSPQ}^u \subseteq \text{NPSPQ}^{u+}$ is trivial. □

Remark. Whilst we have $\text{NPSPQ}_s^u \subseteq \text{NPSPQ}^u$, there is no reason to think that $\text{NPSPQ}_s^b \subseteq \text{NPSPQ}^b$. NPSPQ^b can simulate NPSPQ_s^u up to a point, as we will see, but if there is a super-polynomial number of successor calls in a scheme of NPSPQ_s^b , then we can not use our method to simulate in NPSPQ^b .

3.8 NPspace \subseteq NPSPQ_s^u

Let $\Omega \subseteq \text{STRUC}(\sigma)$ be some problem in NPspace. Then there exists a positive integer k and a non-deterministic Turing machine T such that, for all structures \mathcal{A} (of size n), and all orderings of \mathcal{A} , the following are equivalent:

- $T \downarrow \text{bin}(\mathcal{A})$.
- $T \downarrow \text{bin}(\mathcal{A})$ with the read/write head never leaving the first n^k squares.
- $\mathcal{A} \in \Omega$.

Let Q be the set of states of T , including start state q_s and accept state q_a .

In addition to variables ranging over the elements of \mathcal{A} , we will want to enlarge our alphabet such that we also have:

- The set of pairs $\Pi = \{(zero, q), (one, q), (blank, q) : q \in Q\}$.
- The special symbols $L, R,$ and U . These will track the movement of T 's read/write head.
- The marker symbol M .

Since Q is fixed this will not be a problem.

Let $\Gamma \subseteq \Pi^2$ be such that $((y_1, q_1), (y_2, q_2)) \in \Gamma$ iff $y_1 = y_2$. Γ appears to be a rather unusual set, but we will need to verify such pairs on the queue in our given simulation.

Let $\Delta \subseteq \Pi^2 \cup (\Pi \times \{L, R\})$ be such that:

- $((y_1, q_1), (y_2, q_2)) \in \Delta$ if there is a transition rule of T from (y_1, q_1) to (y_2, q_2) .
- $((y, q), R) \in \Delta$ if there is a transition rule that moves the read/write head Right from (y, q) .
- $((y, q), L) \in \Delta$ if there is a transition rule that moves the read/write head Left from (y, q) .

Δ is, therefore, our visualisation of T 's transition rules.

Theorem 15. $\text{NPspace} \subseteq \text{NPSPQ}_s^u$.

Proof. We aim to prove this by simulation. We will construct a program scheme $\rho_\Omega \in \text{NPSPQ}_s^u(k+1)$ such that $\mathcal{A} \models \rho_\Omega$ if, and only if, $\mathcal{A} \in \Omega$. The n^k weights of the form $(1, \bar{m})$ will mimic the n^k squares of the Turing machine T . The line on \bar{m} will always refer to a k -tuple. ρ_Ω will be:

INPUT $(\bar{v}, v_q, v, v', v'', \bar{m}, \overline{m_{r/w}})$
 $\tau_{bin(\mathcal{A})}; \tau_{sim}; \tau_{end}$
 OUTPUT $(\bar{v}, v_q, v, v', v'', \bar{m}, \overline{m_{r/w}})$

We will now meet the sub-routines $\tau_{bin(\mathcal{A})}$, τ_{sim} , and τ_{end} , and explain why each one performs the function claimed of it.

3.8.1 Preparation: $\tau_{bin(\mathcal{A})}$

First we will write the marker symbol M to the weights $(1, 1^k)$ to $(1, n^k)$. Before we can simulate the computation of T we must write $bin(\mathcal{A})$ to the queue. We will do this by randomly writing *zero*, *one* or *blank*, together with the start state q_s , simultaneously to the weight ranges $(1, 1^k)$ to $(1, n^k)$ and $(n, 1^k)$ to (n, n^k) . The n^k entries in the range $(1, 1^k)$ to $(1, n^k)$ will represent the n^k squares of the Turing tape at the start of computation. We let the variable \bar{m} range over these tape squares in the following.

```

FOR  $\bar{m} = (1^k)$  TO  $(n^k)$  DO
  PUSH  $M, (1, \bar{m})$ 
  EITHER(Zero, One, Blank)
  IF Zero THEN DO PUSH  $(zero, q_s), (1, \bar{m})$ ; PUSH  $(zero, q_s), (n, \bar{m})$  FI
  IF One THEN DO PUSH  $(one, q_s), (1, \bar{m})$ ; PUSH  $(one, q_s), (n, \bar{m})$  FI
  IF Blank THEN DO PUSH  $(blank, q_s), (1, \bar{m})$ ; PUSH  $(blank, q_s), (n, \bar{m})$  FI
NEXT

```

This will leave the queue looking like:

$$\begin{array}{c}
 \overbrace{(y_{n^k}^0, q_s)} \quad (n, n^k) \\
 \vdots \\
 \overbrace{(y_{1^k}^0, q_s)} \quad (n, 1^k) \\
 \overbrace{(y_{n^k}^0, q_s)} \quad (1, n^k) \\
 M \quad (1, n^k) \\
 \vdots \\
 \overbrace{(y_{1^k}^0, q_s)} \quad (1, 1^k) \\
 \overbrace{M} \quad (1, 1^k)
 \end{array}$$

where each $y \in \{zero, one, blank\}$. We will consume the top copy in weights $(n, 1^k)$ to (n, n^k) to check that $y_{1^k}^0, \dots, y_{n^k}^0$ is an encoding $bin(\mathcal{A})$.

If σ is a signature with relations R_1, R_2, \dots, R_j , of arities a_1, a_2, \dots, a_j then the coding of R_1 will take the weights $(n, 1^k)$ to $(n, 1^k + n^{a_1})$, the coding of R_2 will take the weights $(n, 1^k + n^{a_1} + 1)$ to $(n, 1^k + n^{a_1} + n^{a_2} + 1)$ etc. For $i \leq k$, note that the i 'th power of n is represented by the k -ary vector $\overline{n^i}$ that has a 1 in positions $1 \leq i' \leq k - i$ and an n in positions $k - i < i' \leq k$.

We will explicitly give the method when $\sigma = \sigma_2 = \langle E^2 \rangle$, i.e. on graphs. In the sequence $y_{1^k}^0, \dots, y_{n^k}^0$, we must ensure that all apart from the first n^2 entries are *blank*. We must then ensure that the first n^2 entries code the edge relation of the graph. Recall that $\overline{n^2} = (1, \dots, 1, n, n)$ ³.

```

FOR  $\overline{m} = n^k$  DOWNTO  $1^k + \overline{n^2} + 1^k$  DO
   $v := \text{POP}$ 
  IF  $v \neq \text{blank}$  THEN DO LOOP FOREVER FI NEXT
FOR  $\overline{m} = \overline{n^2}$  DOWNTO  $1^k$  DO
   $v := \text{POP}$ 
  IF  $v = \text{blank}$  THEN DO LOOP FOREVER FI
   $(v_1, \dots, v_k) := \text{element}(m_1, \dots, m_k)$ 
  IF  $E(v_{k-1}, v_k) \wedge v = \text{zero}$  THEN DO LOOP FOREVER FI
  IF  $\neg E(v_{k-1}, v_k) \wedge v = \text{one}$  THEN DO LOOP FOREVER FI NEXT

```

Any computation that gets through that will leave the queue looking like:

$$\begin{array}{l}
\overbrace{y_{n^k}^0} \quad (1, n^k) \\
M \quad (1, n^k) \\
\vdots \\
y_{1^k}^0 \quad (1, 1^k) \\
\overbrace{M} \quad (1, 1^k)
\end{array}$$

where $y_{1^k}^0, \dots, y_{n^k}^0$ is necessarily a copy of $\text{bin}(\mathcal{A})$ – which we consider to be T 's tape on input. The superscript 0 refers to time 0.

³We retain the line on the k -ary $\overline{n^2}$ to distinguish it from the binary $n^2 = (n, n)$.

3.8.2 Simulation: τ_{sim}

Throughout the simulation we will keep track of the position of T 's read/write head in a numeric variable k -tuple $\overline{m_{r/w}}$, and the state will be remembered in a single variable v_q .

In simulating the i th step of T we first guess what type of move T will perform at that stage. We will verify later that these were valid choices in the computation. There are two basic cases: either moving the read/write head; or changing the entry at the read/write head's current position. We can not move left from position 1^k and if we move right from position n^k we may assume we do not have an accepting computation.

In the first case we write the symbol R or L to the weight $(1, \overline{m_{r/w}})$, depending on whether the read/write head is to move right or left. We amend the position of the read/write head as stored in $\overline{m_{r/w}}$, either adding one, or subtracting one. Afterwards we guess what will be the entries of T 's tape at time $i + 1$ and write them, together with the current state stored in v_q , to all the weights $(1, 1^k)$ to $(1, n^k)$. We will want the tape-entries we have guessed to be exactly the same as at time i , written beneath them on the queue (except for the introduction somewhere of a symbol R or L). We will only verify that this is the case at the end of the simulation.

In the second case we write the symbol U to the weight $(1, \overline{m_{r/w}})$. We then choose a new state to go into, amending v_q accordingly. We then guess the entries of T 's tape at time $i + 1$ and write them, together with the new state stored in v_q , to all the weights $(1, 1^k)$ to $(1, n^k)$. In this case we will want the tape-entries to be the same as at time i except possibly for the weight $(1, \overline{m_{r/w}})$, *i.e.*, for entries split by a U symbol. We will verify this at the end of the computation.

This simulation will continue until we guess that we go into the accept state q_a .

```

WHILE  $v_q \neq q_a$  DO
  EITHER(Right, Left, Unmoved)
  IF Right THEN DO
    IF  $\overline{m_{r/w}} = n^k$  THEN DO LOOP FOREVER FI
    PUSH  $R, (1, \overline{m_{r/w}})$ 
    INCR  $\overline{m_{r/w}}$  FI
  IF Left THEN DO
    IF  $\overline{m_{r/w}} = 1^k$  THEN DO LOOP FOREVER FI
    PUSH  $L, (1, \overline{m_{r/w}})$ 
    DECR  $\overline{m_{r/w}}$  FI
  IF Unmoved THEN DO
    PUSH  $U, (1, \overline{m_{r/w}})$ 
    GUESS  $v_q$ ; IF  $v_q \notin Q$  THEN DO LOOP FOREVER FI
  FOR  $\overline{m} = (1^k)$  TO  $(n^k)$  DO
    EITHER(Zero, One, Blank)
    IF Zero THEN DO PUSH  $(zero, v_q), (1, \overline{m})$  FI
    IF One THEN DO PUSH  $(one, v_q), (1, \overline{m})$  FI
    IF Blank THEN DO PUSH  $(blank, v_q), (1, \overline{m})$  FI NEXT OD

```

3.8.3 Verification: τ_{end}

We now move into the verification, in which we check that we have effected a legitimate computation.

If t is the length of the simulated computation, then at this point each weight $(1, \overline{m})$ of the queue, representing the \overline{m} th square of the Turing tape, will have a

stack on it looking something like⁴:

$$\begin{array}{c}
 (y_{\bar{m}}^t, q_a) \\
 (y_{\bar{m}}^{t-1}, q^{t-1}) \\
 \vdots \\
 (y_{\bar{m}}^{g+1}, q^{g+1}) \\
 R \\
 (y_{\bar{m}}^g, q^g) \\
 \vdots \\
 (y_{\bar{m}}^{h+1}, q^{h+1}) \\
 U \\
 (y_{\bar{m}}^h, q^h) \\
 \vdots \\
 (y_{\bar{m}}^1, q^1) \\
 (y_{\bar{m}}^1, q_s) \\
 M
 \end{array}$$

Note that entries R , L , U , or, indeed, M may never be adjacent. We will read these entries off such that we can consider three adjacent at once. At any point the variables v, v', v'' will hold descending successive entries on the stack. τ_{end} will be:

```

FOR  $\bar{m} = n^k$  DOWNTO  $1^k$  DO
 $v := \text{POP}; v' := \text{POP}; v'' := \text{blank}$ 
WHILE  $v'' \neq M$  DO
 $v'' := \text{POP}$ 
 $\tau_{check}$ 
 $v := v'; v' := v''$  OD NEXT

```

In the case that $v, v' \notin \{L, R, U\}$, we will simply check that the tape-entry in v

⁴The q_s with superscript should be considered as representative of some state in Q , just as the y_s with superscript are representative of one of $\{zero, one, blank\}$. The q_s with subscript, e.g. q_s, q_a , are actual states.

is the same as in v' . This is not quite the condition $v = v'$, since each such entry on the queue is a pair of tape entry and state, but it is the condition $(v, v') \in \Gamma$, i.e. the tape entries contained in v and v' are the same – even if the states are different.

In the case that $v \in \{L, R, U\}$ we do nothing.

Where v' is the symbol R we check that (y^{i+1}, q^{i+1}) and (y^i, q^i) (in v and v'' , respectively) are the same. It is actually consequent on our simulation method that $q^{i+1} = q^i$. We must also check whether T has a transition rule in state q^i reading y^i to move right.

We do analogously when v' is L .

Where v' is the symbol U we check that T has a transition rule $((y^i, q^i), (y^{i+1}, q^{i+1}))$ (stored in (v, v'')) in Δ .

Thus τ_{check} will be:

```

IF  $v, v' \notin \{L, R, U\}$  THEN DO
  IF  $(v, v') \notin \Delta$  THEN DO LOOP FOREVER FI FI
IF  $v' = R$  THEN DO
  IF  $v \neq v''$  THEN DO LOOP FOREVER FI
  IF  $(v, R) \notin \Delta$  THEN DO LOOP FOREVER FI
IF  $v' = L$  THEN DO
  IF  $v \neq v''$  THEN DO LOOP FOREVER FI FI
  IF  $(v, L) \notin \Delta$  THEN DO LOOP FOREVER FI
IF  $v' = U$  THEN DO
  IF  $(v, v'') \notin \Delta$  THEN DO LOOP FOREVER FI FI

```

The result follows. □

Corollary. $\text{NPspace} \subseteq \text{NPSPQ}^u$

Proof. Recall $\text{NPSPQ}_s^u = \text{NPSPQ}^u$. □

3.9 A polynomial-time restriction of NPSPQ^u .

Definition. A program scheme $\rho \in \text{NPSPQ}^u(k)$ is said to be *polynomially step-bounded* if there exists a j such that, for all structures \mathcal{A} , ρ accepts \mathcal{A} if, and only if, ρ accepts \mathcal{A} within n^j steps. Let:

- $\text{NPSPQ}^u(k)_{poly} :=$

$\{ \rho : \rho \in \text{NPSPQ}^u(k) \text{ and } \rho \text{ is polynomially step-bounded} \}$

- $\text{NPSPQ}_{poly}^u := \cup_{k \geq 0} \text{NPSPQ}^u(k)_{poly}$.

Proposition 16. $\text{NPSPQ}_{poly}^u \subseteq \text{NPSPQ}^b$.

Proof. We prove this by simulation. The idea is that we can never attempt to use a weight more than once. Given some $\rho \in \text{NPSPQ}^u(k)_{poly}$, and the j that is the polynomial power of its step bound, we will construct a $\rho' \in \text{NPSPQ}^b(k+j)$, such that, for all structures \mathcal{A} , $\mathcal{A} \models \rho$ iff $\mathcal{A} \models \rho'$.

Let the line on \bar{m} indicate a j -tuple. Assume \bar{m} is a numeric variable tuple not involved in ρ .

Given ρ we construct ρ' by

- adding, after every line, except the last, the instructions

INCR \bar{m}

IF $\bar{m} = n^j$ THEN DO LOOP FOREVER FI.

(\bar{m} will act as a step-counter in ρ'), and

- converting all instances of PUSH $v, (m_1, \dots, m_k)$ to PUSH $v, (m_1, \dots, m_k, \bar{m})$.

□

Corollary. $\text{NP} \subseteq \text{NPSPQ}^b$

Proof. The simulation method we used in proving $\text{NPspace} \subseteq \text{NPSPQ}^u$ will also prove $\text{NP} \subseteq \text{NPSPQ}_{poly}^u$. The result follows from the previous lemma. □

Chapter 4

Adding Universal Quantification

4.1 Introducing APS(1)

The schemes of NPS have existential quantification built-in through their guess instruction. NPS(1) is devoid of any notion of universal quantification. The higher strata, NPS(m), have some notion of universal quantification, through negation of existential quantification, but have no facility to combine both types of quantification within while-loop recursion. We consider the effect of explicitly adding universal quantification. We are, once more, without the stack.

Definition (Syntax of APS(1)). The syntax of APS(1) is as that of NPS(1), except the extant GUESS instruction is renamed \exists GUESS, and a new instruction \forall GUESS is added, with identical syntax.

The schemes of NPS(1) accepted a structure, expanded with values for the free variables, iff there existed some accepting computation, *i.e.* at each point the program went through an \exists GUESS v , there existed an assignment to v such that thenceforth the scheme made it to output. The schemes of APS(1) accept an expanded structure iff:

- at each point the program goes through a \exists GUESS v , there exists an assignment to v such that thenceforth the computation makes it to output, and
- at each point the program goes through a \forall GUESS v , we have that for all assignments to v the computation thenceforth makes it to output.

These instructions have an appealing semantic characterisation in terms of the configurations of a scheme $\rho \in \text{APS}(1)$. When computing on a structure \mathcal{A} , we can construct an alternating graph $\mathcal{AG}_{\mathcal{A}}^{\rho}$ just as we constructed $\mathcal{G}_{\mathcal{A}}^{\rho}$, but with the additional information that a configuration $(\bar{z}, \hat{w}, \bar{u})$ is *universal* iff \hat{w} represents line l , and the instruction on line $(l+1)$ is a \forall GUESS. Observe that, for the edge relation of $\mathcal{AG}_{\mathcal{A}}^{\rho}$, there is no difference between \exists GUESS and \forall GUESS, since in each case the configuration can move to any configuration that is identical excepting the guess for the pertinent variable. Let the the lines on \bar{z} denote an i -tuple, the line on \bar{u}, \bar{v} denote j -tuples, the hat on \hat{w} denote a $|\rho|$ -tuple, and the line on \bar{x}, \bar{y} denote $(i + |\rho| + j)$ -tuples.

Proposition 17. *Suppose $\rho \in \text{APS}(1)$ has i free variables and j input-output variables, and that $\psi(\vec{x}, \vec{y})$ is a quantifier-free first order formula expressing the edge relation of $\mathcal{AG}_{\mathcal{A}}^{\rho}$, then the following are equivalent:*

- $\mathcal{A} \models \rho(\vec{z})$

-

$$\begin{aligned} \mathcal{A} \models \exists w_1, \dots, w_{|\rho|} w_1 = w_2 \wedge \\ \exists w'_1, \dots, w'_{|\rho|} w'_1 \neq w'_2 \wedge \dots \wedge w'_{|\rho|-1} \neq w'_{|\rho|} \wedge \\ \exists \vec{u}, \vec{v} \text{ AR}[\lambda \vec{x}, \vec{y} \psi](\vec{z}, w_1, \dots, w_{|\rho|}, \vec{u}), (\vec{z}, w'_1, \dots, w'_{|\rho|}, \vec{v})) \end{aligned}$$

Proof. Follows immediately from the semantics of $\text{APS}(1)$ and the definition of $\mathcal{AG}_{\mathcal{A}}^{\rho}$. Recall that the bizarre constraints on the w s are our encoding of the first and last lines. \square

Just as acceptance in $\text{NPS}(1)$ is a reachability (transitive closure) problem, so acceptance in $\text{APS}(1)$ is an alternating reachability problem.

Corollary. $\text{APS}(1) \subseteq \pm \text{AR}^1[\text{FO}]$.

Not only can the schemes of $\text{APS}(1)$ be recast as formulae of $\text{AR}^1[\text{FO}]$, but a scheme of $\text{APS}(1)$ can express the Alternating Reachability relation. In order to prove this, we will have use for another shorthand that is available to us in the presence of our new instruction $\forall \text{GUESS}$.

- Let $\text{word}_1, \dots, \text{word}_j$ be words representing certain possibilities. Consider:

```

ALL(word1, ..., wordj)
  IF word1 THEN DO  $\tau_1$  FI
  ⋮
  IF wordj THEN DO  $\tau_j$  FI

```

to be shorthand for:

```

 $\forall \text{GUESS } v_1, \dots, v_j$ 
  IF  $v_1 = v_2$  THEN DO  $\tau_1$  FI
  IF  $(v_1 \neq v_2) \wedge (v_2 = v_3)$  THEN DO  $\tau_2$  FI

```

$$\begin{array}{l} \vdots \\ \text{IF } (v_1 \neq v_2) \wedge \dots \wedge (v_{j-2} \neq v_{j-1}) \wedge (v_{j-1} = v_j) \text{ THEN DO } \tau_{j-1} \text{ FI} \\ \text{IF } (v_1 \neq v_2) \wedge \dots \wedge (v_{j-1} \neq v_j) \text{ THEN DO } \tau_j \text{ FI} \end{array}$$

ALL is the universal counterpoint to the existential EITHER. When a program scheme meets an EITHER instruction it will accept iff one of those choices leads to acceptance; when a program scheme meets the ALL instruction, it will accept iff all of the choices lead to acceptance.

Proposition 18. *There is a program scheme $\rho_{\text{AR}}(u, v) \in \text{APS}(1)$ with two free variables that expresses the relation AR. Formally, for all alternating graphs \mathcal{A} , and vertices $a, a' \in \mathcal{A}$:*

$\mathcal{A} \models \rho_{\text{AR}}(a, a')$ iff $\mathcal{A} \models \text{AR}(a, a')$ (there is an alt. path in \mathcal{A} from a to a')

Proof. We will construct ρ_{AR} . First we note that the relation $\text{AR}(u, v)$ may be written in LFP[FO] as $\text{LFP}[\lambda Pxy\psi](u, v)$, where $\psi(P, x, y) :=$

$$(x = y) \vee (\exists s P(x, s) \wedge \neg U(s) \wedge E(s, y)) \vee (\exists s P(x, s) \wedge U(s) \wedge [\forall r E(s, r) \rightarrow P(r, y)])$$

This can be re-written as $\psi(P, x, y) \equiv$

$$(x = y) \vee \exists s P(x, s) \wedge ([\neg U(s) \wedge E(s, y)] \vee [U(s) \wedge \forall r (\neg E(s, r) \vee P(r, y))])$$

Note that the $\exists s$ quantifies everything to its right. We will denote the two conjuncts after that quantification as Left and Right. Thus:

- Left is $P(x, s)$, and
- Right is $([\neg U(s) \wedge E(s, y)] \vee [U(s) \wedge \forall r (\neg E(s, r) \vee P(r, y))])$.

Let $\tau_{\text{AR}}(u, v, x, y, s, r)$ be the sub-routine involving free variables u, v :

$$\begin{array}{l} x := u; y := v \\ \text{WHILE } x \neq y \text{ DO} \\ \quad \exists \text{GUESS } s \\ \quad \text{ALL}(\text{Left}, \text{Right}) \end{array}$$

```

IF Left THEN DO  $y := s$  FI
IF Right THEN DO
  IF  $\neg U(s) \wedge E(s, y)$  THEN DO  $x := y$  FI
  IF  $\neg(\neg U(s) \wedge E(s, y)) \wedge \neg U(s)$  THEN DO LOOP FOREVER FI
   $\forall$ GUESS  $r$ 
  IF  $\neg(\neg U(s) \wedge E(s, y)) \wedge \neg E(s, r)$  THEN DO  $x := y$  FI
  IF  $\neg(\neg U(s) \wedge E(s, y)) \wedge E(s, r)$  THEN DO  $x := r$  FI OD

```

We now set $\rho_{\text{AR}}(u, v)$ to be:

- INPUT(x, y, r, s); τ_{AR} ; OUTPUT(x, y, r, s)

ρ_{AR} evaluates whether (u, v) is in AR from the outside-in, hence x and y are initially set to u and v , respectively. Each path of the computation succeeds only when the variables x and y become equal. ρ_{AR} mimics exactly $\text{LFP}[\lambda Pxy\psi](u, v)$: indeed if the rank of (u, v) in $\text{LFP}[\lambda Pxy\psi](u, v)$ is j , i.e. $(u, v) \in \psi^j$ but $(u, v) \notin \psi^{j-1}$, then τ_{AR} will go through the while loop [a maximum of] j times. \square

Proposition 19. $\text{AR}^1[\mathbf{FO}] \subseteq \text{APS}(1)$

Proof. Take any formula $\varphi \in \text{AR}^1[\mathbf{FO}]$. Then φ is of the form $\exists u_1 \dots u_{j'} \exists v_1 \dots v_{j''} \text{AR}[\lambda \bar{x}, \bar{y}\psi](\bar{u}, \bar{v})$, where ψ is quantifier-free. We construct $\rho_\varphi \in \text{APS}(1)$ such that, for all structures \mathcal{A} , $\mathcal{A} \models \rho_\varphi$ iff $\mathcal{A} \models \varphi$.

Let ρ_φ be:

```

INPUT( $\bar{x}, \bar{y}, \bar{s}, \bar{r}, u_1, \dots, u_{j'}, v_1, \dots, v_{j''}$ )
 $\exists$ GUESS  $u_1, \dots, u_{j'}, v_1, \dots, v_{j''}$ 
 $\tau_{\text{AR}}(\bar{u}, \bar{v}, \bar{x}, \bar{y}, \bar{s}, \bar{r})$ 
OUTPUT( $\bar{x}, \bar{y}, \bar{s}, \bar{r}, u_1, \dots, u_{j'}, v_1, \dots, v_{j''}$ )

```

\square

Theorem 20. $\text{APS}(1) = \text{AR}^1[\mathbf{FO}]$

Proof. Follows from the previous two propositions. \square

Corollary. $\text{APS}(1) = \text{AR}^1[\mathbf{FO}] = \text{LFP}[\mathbf{FO}] = \text{AR}^*[\mathbf{FO}]$.

Proof. $\text{AR}^1[\mathbf{FO}] = \text{LFP}[\mathbf{FO}] = \text{AR}^*[\mathbf{FO}]$ is proved in [16]. \square

4.2 Introducing APSS(1)

Here we consider the situation where we augment the schemes of APS(1) with a stack for memory. We will find that we can quantify over the stack in a way that was not possible with NPSS. Consequently, order will not be a problem, and we quickly establish that we subsume NPspace. We will have no need for free variables to generate stratification: as with NPSPQ, we dispense with them.

Definition (Syntax of APSS(1)). Notwithstanding the forbidding of free variables, the syntax is that of APS(1), with the PUSH and POP instructions of NPSS(1).

Definition. Suppose the program scheme $\rho \in \text{APSS}(1)$ involves j variables. Then a configuration of ρ , computing on a structure \mathcal{A} , is a sequence (\bar{v}, l, w) giving the values of the variables, the number of the line just executed, and the contents of the stack ($w \in |\mathcal{A}|^*$).

Each such program scheme ρ , computing on a structure \mathcal{A} , gives rise to an infinite alternating graph $\mathcal{A}\mathcal{G}_\rho^p$, defined as in the previous section. We say there is a *finite alternating path* between configurations c, c' in $\mathcal{A}\mathcal{G}_\rho^p$, if there is an $i \in \omega$ such that $(c, c') \in \psi^i(P, x, y)$, where ψ is as in the proof to Proposition 18.

For some structure \mathcal{A} , let $\Gamma_{\mathcal{A}}$ be some subset of $|\mathcal{A}|^*$. Then for some signature σ , let Γ be the global set $\{\Gamma_{\mathcal{A}} : \mathcal{A} \in \text{STRUC}(\sigma)\}$.

Definition (Recognising Stacks). We say that the global set Γ is *recognisable* iff there is a $\rho_\Gamma \in \text{APSS}(1)$ such that for all \mathcal{A} the following are equivalent:

- For all \bar{v} , there exists \bar{v}' and $w' \in |\mathcal{A}|^*$ such that there is an alternating path in $\mathcal{A}\mathcal{G}_{\rho_\Gamma}^p$ from configuration $(\bar{v}, 1, w)$ to configuration $(\bar{v}', |\rho_\Gamma|, w')$.
- $w \in \Gamma_{\mathcal{A}}$.

Suppose that τ_Γ is the subroutine constructed from ρ_Γ by removing the input and output instructions. We are stating that, for each \mathcal{A} , the uniform subroutine τ_Γ , when confronted with a stack w , finishes (i.e. does not loop forever) if, and only if, $w \in \Gamma_{\mathcal{A}}$. This is independent of the values of all input-output variables going into τ_Γ . Essentially, τ_Γ recognises w .

If a subroutine, computing on \mathcal{A} , recognises a stack with contents $w \in |\mathcal{A}|^*$, without ever popping off more than the top entries w' ($|w'| < |w|$), then it follows that that subroutine will recognise any word in $\{w'\}.|\mathcal{A}|^*$. This suggests that sometimes recognition only relates to the top portion of the contents of a stack. This motivates the cartesian product in the following:

Lemma 21. *Let M_1 be a special marker symbol. The following is recognisable as the stack:*

$$\{\{M_1, x_1, \dots, x_n, M_1 : x_1, \dots, x_n \text{ is an ordering of } |\mathcal{A}| \}.|\mathcal{A}|^* : \mathcal{A} \in \text{STRUC}(\sigma)\}$$

where the ' $.|\mathcal{A}|^*$ ' indicates cartesian product¹

Proof. Recall that Λ is our set of additional special symbols. Therefore $x \notin \Lambda$ iff x represents a *bona fide* element of the structure on which we are computing. We begin by defining τ_{order} :

```

v1 := POP ; IF v1 ≠ M1 THEN DO LOOP FOREVER FI
∀GUESS v2
IF v2 ∉ Λ THEN DO
  v1 := POP
  WHILE v1 ≠ v2 DO
    v1 := POP OD
  v1 := POP
  IF v1 = v2 THEN DO LOOP FOREVER FI
  WHILE v1 ≠ v2 ∧ v1 ≠ M1 DO
    v1 := POP
  IF v1 = v2 THEN DO LOOP FOREVER FI OD FI

```

The subroutine works by checking that every *bona fide* element appears once (lines 4-6), and only once, *i.e.* not again (lines 7-11), between two markers M_1 .

The following scheme ρ accepts the global set of the lemma:

¹There is potential for ambiguity here: By $|\mathcal{A}|^*$ we mean any possible finite string of *real* elements of \mathcal{A} , *i.e.* anything that could possibly be pushed to the stack, as opposed to just those symbols (encoded as tuples) that *represent* elements of \mathcal{A} .

INPUT(v_1, v_2)
 τ_{order}
 OUTPUT(v_1, v_2).

□

We will also define the following subroutine τ_{push} , which pushes a random (non-deterministic) number of random (non-deterministic) choices (except M_1) to the stack:

\exists GUESS v', v''
 WHILE $v' \neq v''$ DO
 \exists GUESS w'
 IF $w' = M_1$ THEN LOOP FOREVER FI
 PUSH w'
 \exists GUESS v', v'' OD

4.3 The ACCEPT instruction.

With the inclusion of a universal side to our semantics, we will have need of an ACCEPT instruction which, as its name suggests, tells the computation to immediately accept. Any program scheme ρ' that involves an ACCEPT instruction should be considered shorthand for a scheme $\rho \in \text{APSS}(1)$ in the following way. Assume, w.l.o.g., that ρ' involves variable set V and that $v_1, v_2 \notin V$. Let τ' be ρ' without the input and output instructions. Construct τ from τ' via the substitutions:

- All tests t in while loops in τ' become tests $v_1 \neq v_2 \wedge t$ in τ .
- All instances of ACCEPT in τ' become $v_1 := v_2$ in τ .

Then ρ should be considered as:

INPUT(V, v_1, v_2)
 $v_1 := v_2$
 WHILE $v_1 \neq v_2$ DO τ OD
 OUTPUT(V, v_1, v_2)

Note that, once $v_1 = v_2$, the program can never get trapped in an infinite loop, and, consequently, must make it to output.

4.4 $\mathbf{APSS}_s(1) = \mathbf{APSS}(1)$

Lemma 22. $\mathbf{APSS}_s(1) \subseteq \mathbf{APSS}(1)$.

Proof. We prove the inclusion by simulating the successor relation. We take any scheme $\rho \in \mathbf{APSS}_s(1)$, and construct a scheme $\rho' \in \mathbf{APSS}(1)$ such that, for all structures \mathcal{A} , $\mathcal{A} \models \rho$ iff $\mathcal{A} \models \rho'$.

Assume, without loss of generality, that ρ involves element variables V , with $v_1, v_2, v_3 \notin V$, and does not use the marker symbol M_1 . Given ρ , we will construct ρ' by adding a special start-routine, and amending pop instructions as well as successor tests in while loops.

Let $V' = V \cup \{v_1, v_2, v_3\}$. Then ρ' will be:

```

INPUT( $V'$ )
PUSH  $M_1$ 
 $\tau_{push}$ 
PUSH  $M_1$ 
ALL(CheckOrder, Continue)
IF CheckOrder THEN DO  $\tau_{order}$ ; ACCEPT FI
IF Continue THEN DO FI
 $\tau_\rho$ 
OUTPUT( $V'$ )

```

where τ_ρ is the, as yet undefined, subroutine that actually mimics ρ . Observe how we are using the ALL instruction to use the stack twice, once for verification of the order, and again for whatever we want to do in the rest of the computation. The stack is no longer readable only once, as it was with NPSS. Note that the ALL choice ‘Continue’ is a dummy, in that any computation that follows that path will continue through the rest of the program.

We will now meet τ_ρ .

4.4.1 Simulation of ρ ; τ_ρ

Any computation that gets to this sub-routine will have M_1 as the top element of the stack. Assuming M_1 is not a symbol of ρ , we can use it to ensure that we never stray into the bottom part of the stack, where the putative order is held, during our simulation of ρ . In constructing τ_ρ , we first remove the input and output lines (of ρ). Next we,

- convert all instances of ' $v := \text{POP}$ ' to:

```

 $v_1 := v$ 
 $v := \text{POP}$ 
IF  $v = M_1$  THEN DO PUSH  $v$  ;  $v := v_1$  FI

```

We may also have to evaluate quantifier-free successor queries of the form $v' = \text{succ}(v)$, that might appear in a test for a while loop, immediately before the test of that while loop. Let Φ be a propositional formula:

- Convert all instances of ' $\text{WHILE } \Phi(v' = \text{succ}(v)) \text{ DO } \tau \text{ OD}$ ' to:

```

 $\tau_{succ}$ 
WHILE  $\Phi(v_1 = v_2)$  DO  $\tau$ ;  $\tau_{succ}$  OD

```

Where τ_{succ} is:

```

ALL(CheckSuccv, Continue)
IF CheckSuccv THEN DO
  WHILE  $v_1 \neq M_1$  DO  $v_1 := \text{POP}$  OD
  WHILE  $v_1 \neq v$  DO  $v_1 := \text{POP}$  OD
   $v_2 := \text{POP}$ 
  IF  $v_2 = v'$  THEN DO ACCEPT FI
  LOOP FOREVER FI
IF Continue THEN DO FI

```

Observe that v' and v are free in the sub-routine.

The subroutine works by splitting the computation, both checking that $v' = \text{succ}(v)$ (consuming the stack in the process) and continuing the computation with the stack intact. \square

Corollary. $\text{APSS}_s(1) = \text{APSS}(1)$

Proof. The converse $\text{APSS}(1) \subseteq \text{APSS}_s(1)$ is trivial. \square

4.5 NPspace \subseteq $\text{APSS}_s(1)$

We will ultimately prove this by simulation of a non-deterministic Turing Machine that uses no more than n^k tape squares (for some k), on input $\text{bin}(\mathcal{A})$, where $|\mathcal{A}| = n$. First, we will need some technical lemmas, which are stated for the case when the signature is σ_2 , *i.e.* for graphs. Similar lemmas may be obtained for other signatures. Let m be such that $1 \leq m \leq n^k$, then we identify \bar{m} with the lexicographic m th variable k -tuple \bar{v} , with respect to the built-in successor.

Lemma 23 (Recognise $\text{bin}(\mathcal{G})$). *Let M_1 and M_2 be special marker symbols. The following is recognisable as the stack (the brackets are synthetic, and appear, as the commas, purely for clarity):*

$$\{ \{ (\bar{1}, \alpha_{\bar{1}}, M_2), (\bar{2}, \alpha_{\bar{2}}, M_2), \dots, (\bar{n^k}, \alpha_{\bar{n^k}}, M_2), M_1 \} \cdot |\mathcal{G}|^* \\ : \mathcal{G} \in \text{STRUC}_s(\sigma_2), \alpha_{\bar{1}} \dots \alpha_{\bar{n^2}} = \text{bin}(\mathcal{G}), \alpha_{\bar{n^2+1}}, \dots, \alpha_{\bar{n^k}} = \text{blank} \}$$

Proof. Letting $\bar{v} = (v_1, \dots, v_k)$, we define τ_{bin} :

```

( $\bar{v}, v', v''$ ) := POP
IF  $\bar{v} \neq \bar{1} \vee v'' \neq M_2$  THEN DO LOOP FOREVER FI
 $\bar{w} := \bar{v}$ 
WHILE  $\bar{v} \neq \bar{n^k}$  DO
  ( $\bar{v}, v', v''$ ) := POP
  IF  $\bar{v} \neq \text{succ}(\bar{w})$  THEN DO LOOP FOREVER FI
  IF  $v'' \neq M_2$  THEN DO LOOP FOREVER FI
  IF  $\bar{v} > \bar{n^2} \wedge v' \neq \text{blank}$  THEN DO LOOP FOREVER FI
  IF  $E(v_{k-1}, v_k) \wedge v' = \text{zero}$  THEN DO LOOP FOREVER FI
  IF  $\neg E(v_{k-1}, v_k) \wedge v' = \text{one}$  THEN DO LOOP FOREVER FI
   $\bar{w} := \bar{v}$  OD
 $v'' := \text{POP}$ ; IF  $v'' \neq M_1$  THEN DO LOOP FOREVER FI

```

Then ρ , as $\text{INPUT}(\bar{v}, v', v'', \bar{w}); \tau_{bin}; \text{OUTPUT}(\bar{v}, v', v'', \bar{w})$, recognises the global set of the lemma. \square

Lemma 24 (Read/Write Head Right). *Let M_1 and M_2 be special marker symbols. The following is recognisable as the stack (the brackets are synthetic, and appear as the commas, purely for clarity):*

$$\begin{aligned} & \{ \{ (\bar{m}, \alpha_{\bar{m}}, M_2), \dots, (\bar{n}^k, \alpha_{\bar{n}^k}, M_2), (\bar{1}, \alpha_{\bar{1}}, M_2), \dots, (\overline{m-1}, \alpha_{\overline{m-1}}, M_2), M_1, \\ & (\overline{m+1}, \alpha_{\overline{m+1}}, M_2), \dots, (\bar{n}^k, \alpha_{\bar{n}^k}, M_2), (\bar{1}, \alpha_{\bar{1}}, M_2), \dots, (\bar{m}, \alpha_{\bar{m}}, M_2), M_1 \} \cdot |\mathcal{G}|^* \\ & : \mathcal{G} \in \text{STRUC}_s(\sigma_2), \alpha_{\bar{1}}, \dots, \alpha_{\bar{n}^k} \in \Sigma \} \end{aligned}$$

Proof. We define $\tau_{r/w-right}$ in two parts. One part will check that the numbers \bar{m} to $\overline{m-1}$, and $\overline{m+1}$ to \bar{m} behave correctly; that the markers are placed properly; and that the α s are in Σ . The other part will check that the α s match in the two tape lists, i.e. each $\alpha_{\bar{j}}$ that appears before the first M_1 is equal to the $\alpha_{\bar{j}}$ that appears between the first and second M_1 s. In the following, the variables $\overline{u_{top}}$ will hold the number of the first entry of the first tape (\bar{m}), and $\overline{u_{bottom}}$ will hold the number of the last entry of that tape ($\overline{m-1}$).

Define $\tau_{checkform}$ to be:

```

( $\bar{v}, v', v''$ ) := POP
IF ( $v_1 \in \Lambda$ )  $\vee$  ...  $\vee$  ( $v_k \in \Lambda$ )  $\vee$  ( $v' \notin \Sigma$ )  $\vee$  ( $v'' \neq M_2$ ) THEN DO
  LOOP FOREVER FI
 $\overline{u_{top}}$  :=  $\bar{v}$ 
 $\overline{u_{bottom}}$  := inv.cyc.succ( $\overline{u_{top}}$ )
( $\bar{w}, w', w''$ ) := POP
WHILE  $\bar{v} \neq \overline{u_{bottom}}$  DO
  ( $\bar{v}, v', v''$ ) := POP
  IF  $\bar{w} \neq \text{cyc.succ}(\bar{v})$  THEN DO LOOP FOREVER FI.
  IF  $v' \notin \Sigma \vee v'' \neq M_2$  THEN DO LOOP FOREVER FI
   $\bar{w} := \bar{v}$  OD
 $v'' := \text{POP}$ ; IF  $v'' \neq M_1$  THEN DO LOOP FOREVER FI

```

```

( $\bar{v}, v', v''$ ) := POP
IF ( $v_1 \in \Lambda$ )  $\vee$  ...  $\vee$  ( $v_k \in \Lambda$ )  $\vee$  ( $v' \notin \Sigma$ )  $\vee$  ( $v'' \neq M_2$ ) THEN DO
  LOOP FOREVER FI
IF  $\bar{v} \neq \text{cyc.succ}(\overline{u_{top}})$  THEN DO LOOP FOREVER FI
WHILE  $\bar{v} \neq \overline{u_{top}}$  DO
  ( $\bar{v}, v', v''$ ) := POP
  IF  $\bar{w} \neq \text{cyc.succ}(\bar{v})$  THEN DO LOOP FOREVER FI.
  IF  $v' \notin \Sigma \vee v'' \neq M_2$  THEN DO LOOP FOREVER FI
   $\bar{w} := \bar{v}$  OD
 $v'' := \text{POP}$ ; IF  $v'' \neq M_1$  THEN DO LOOP FOREVER FI

```

The first half of the sub-routine (12 lines) checks the form of the stack up to, and including, the first M_1 . The second half (last 9 lines) does the same up to, and including, the second M_1 .

Define $\tau_{\text{checkcontent}}$ to be:

```

 $\forall$ GUESS  $\bar{w}$ 
IF  $w_1 \in \Lambda \vee \dots \vee w_k \in \Lambda$  THEN DO ACCEPT FI
( $\bar{v}, v', v''$ ) := POP
 $\overline{u_{top}} := \bar{v}$ 
 $\overline{u_{bottom}} := \text{inv.cyc.succ}(\overline{u_{top}})$ 
IF  $\bar{v} = \bar{w}$  THEN DO  $v_{\text{first}\alpha} := v'$  FI
WHILE  $\bar{v} \neq \bar{w}$  DO
  ( $\bar{v}, v', v''$ ) := POP
   $v_{\text{first}\alpha} := v'$  OD

WHILE  $\bar{v} \neq \overline{u_{bottom}}$  DO
  ( $\bar{v}, v', v''$ ) := POP OD
 $v'' := \text{POP}$ 

```

```

 $(\bar{v}, v', v'') := \text{POP}$ 
IF  $\bar{v} = \bar{w}$  THEN DO  $v_{second\alpha} := v'$  FI
WHILE  $\bar{v} \neq \bar{w}$  DO
   $(\bar{v}, v', v'') := \text{POP}$ 
   $v_{second\alpha} := v'$  OD

IF  $v_{first\alpha} \neq v_{second\alpha}$  THEN DO LOOP FOREVER FI

```

The three lines culminating in the middle ' $v'' := \text{POP}$ ' remove down to, and including, the first M_1 .

We now give $\tau_{r/w-right}$:

```

ALL(CheckForm, CheckContent)
IF CheckForm THEN DO  $\tau_{checkform}$  FI
IF CheckContent THEN DO  $\tau_{checkcontent}$  FI

```

Now, ρ , as:

```

INPUT  $(\bar{v}, v', v'', \bar{w}, w', w'', v_{first\alpha}, v_{second\alpha})$ 
 $\tau_{r/w-right}$ 
OUTPUT  $(\bar{v}, v', v'', \bar{w}, w', w'', v_{first\alpha}, v_{second\alpha})$ 

```

recognises the global set of the lemma. □

Lemma 25 (Read/Write Head Left). *Let M_1 and M_2 be special marker symbols. The following is recognisable as the stack (the brackets are synthetic, and appear as the commas, purely for clarity):*

$$\{ \{ (\bar{m}, \alpha_{\bar{m}}, M_2), \dots, (\bar{n}^k, \alpha_{\bar{n}^k}, M_2), (\bar{1}, \alpha_{\bar{1}}, M_2), \dots, (\overline{m-1}, \alpha_{\overline{m-1}}, M_2), M_1, (\overline{m-1}, \alpha_{\overline{m-1}}, M_2), \dots, (\bar{n}^k, \alpha_{\bar{n}^k}, M_2), (\bar{1}, \alpha_{\bar{1}}, M_2), \dots, (\overline{m-2}, \alpha_{\overline{m-2}}, M_2), M_1 \} \cdot |\mathcal{G}|^* : \mathcal{G} \in \text{STRUC}_s(\sigma_2), \alpha_{\bar{1}}, \dots, \alpha_{\bar{n}^k} \in \Sigma \}$$

Proof. We construct $\tau_{r/w-left}$ in a similar manner to $\tau_{r/w-right}$. □

Proposition 26. $\text{NPspace} \subseteq \text{APSS}_s(1)$

Proof. We aim to prove this by simulation. As before, and w.l.o.g., we assume that $\Omega \in \text{NPspace}$ is a graph problem. Similar lemmas to those previous may be obtained for other signatures. Suppose $\Omega \in \text{NPspace}$ is accepted by the Turing Machine T , with space bound of n^k on input $\text{bin}(\mathcal{A})$, where $|\mathcal{A}| = n$. We construct ρ_Ω such that, for all graphs \mathcal{G} , and for all orderings of \mathcal{G} , $\text{bin}(\mathcal{G}) \in \Omega$ iff $\mathcal{G} \models \rho_\Omega$.

We will consider our alphabet expanded to include the $|Q| + |\Sigma|$ symbols representing T 's states and alphabet; we also assume the additional marker symbols M_1 and M_2 . We will store T 's state in one variable v_q : let q_s and q_a be the distinguished start and accept states. Once again Δ is how we envisage T 's transition rules (cf. section 3.8). ρ_Ω will be:

```

INPUT( $\bar{v}, v', v'', \bar{w}, w', w'', v_{first\alpha}, v_{second\alpha}, v_q, v'_q$ )
 $v_q := q_s$ 
PUSH  $M_1$ 
 $\tau_{push}$ 
PUSH  $M_1$ 
ALL(CheckBin,Continue)
IF CheckBin THEN DO  $\tau_{bin}$  ; ACCEPT FI
IF Continue THEN DO FI
WHILE  $v_q \neq q_a$  DO
  ( $\bar{v}, v', v''$ ) := POP; PUSH ( $\bar{v}, v', v''$ )
  EITHER(Right,Left,Unmoved)
  IF Right THEN DO
    IF  $(v_q, v', R) \notin \Delta$  THEN DO LOOP FOREVER FI
    PUSH  $M_1$ ;  $\tau_{push}$ ; PUSH  $M_1$ 
    ALL(Verify,Continue2)
    IF Verify THEN DO  $\tau_{r/w-right}$  ; ACCEPT FI
    IF Continue2 THEN DO FI FI

```

```

IF Left THEN DO
  IF  $(v_q, v', L) \notin \Delta$  THEN DO LOOP FOREVER FI
  PUSH  $M_1$ ;  $\tau_{push}$ ; PUSH  $M_1$ 
  ALL(Verify, Continue2)
  IF Verify THEN DO  $\tau_{r/w-left}$ ; ACCEPT FI
  IF Continue2 THEN DO FI FI
IF Unmoved THEN DO
   $(\bar{v}, v', v'') := \text{POP}$ 
   $\exists$ GUESS  $w', v'_q$ 
  IF  $(v_q, v', v'_q, w') \notin \Delta$  THEN DO LOOP FOREVER FI
   $v_q := v'_q$ ; PUSH  $(\bar{v}, w', v'')$  FI
OD
OUTPUT  $(\bar{v}, v', v'', \bar{w}, w', w'', v_{first\alpha}, v_{second\alpha}, v_q, v'_q)$ 

```

□

4.6 Summary

Below, we summarise the results of this chapter, and the previous.

$\text{NPSPQ}^b(0)$	\subseteq	$\text{NPSPQ}_s^b(0)$	$=$	$\text{NPS}_s(1)$	$=$	NL
$\text{NPSPQ}^u(0)$	\subseteq	$\text{NPSPQ}_s^u(0)$	$=$	$\text{NPSS}_s(1)$	$=$	P
		$\text{NPSPQ}^b(k)$	\subseteq	$\text{NPSPQ}^u(k)$		
		$\text{NPSPQ}_s^b(k)$	\subseteq	$\text{NPSPQ}_s^{b+}(k)$		
		$\text{NPSPQ}_s^{b+}(k)$	\subseteq	$\text{NPSPQ}_s^b(2k)$		
		NPSPQ_s^{b+}	$=$	NPSPQ_s^b		
		$\text{NPSPQ}_s^{u+}(k)$	$=$	$\text{NPSPQ}_s^u(k)$		
		NPSPQ_s^{u+}	$=$	NPSPQ_s^u		
NPSPQ^{u+}	$=$	NPSPQ_s^{u+}	$=$	NPSPQ_s^u	$=$	NPSPQ^u
NPspace	\subseteq	NPSPQ^u				
NP	\subseteq	NPSPQ_{poly}^u	\subseteq	NPSPQ^b	\subseteq	NPspace
LFP	$=$	$\text{APS}(1)$				
P	$=$	$\text{APS}_s(1)$				
NPspace	\subseteq	$\text{APSS}_s(1)$	$=$	$\text{APSS}(1)$		

Figure 4.1: Summary of Results

Chapter 5

Classes of Structure on which

$$P = \pm\text{PS}^k[\text{FO}]$$

In [42], various classes of structure \mathbf{C} were studied, on which, for some k , $\text{NPSS}(k) = \text{PS}^k[\text{FO}]$ captures exactly P . The method used in the proofs involved building a canonical order in $\text{NPSS}(k-1)$, whereupon, since $\text{NPSS}(k) = P$ on ordered structures, the result followed. The following is a consequence of that work:

Proposition 27 ([42]). *Let \mathbf{C} be any class of structures, and let $\bar{x}, \bar{y}, \bar{z}$ be variable j -tuples. Suppose there are formulae $R(w_1, \dots, w_m, \bar{x}, \bar{y}, \bar{z}) \in \pm\text{PS}^k[\text{FO}]$ and $\Psi(w_1, \dots, w_m) \in \text{PS}^l[\text{FO}]$ such that, for all $\mathcal{A} \in \mathbf{C}$:*

- R is commutative in \bar{x} and \bar{y} , and deterministic in \bar{z} , i.e.,

$$\mathcal{A} \models \forall w_1, \dots, w_m \bar{x} \bar{y} \bar{z} R(w_1, \dots, w_m, \bar{x}, \bar{y}, \bar{z}) \leftrightarrow R(w_1, \dots, w_m, \bar{y}, \bar{x}, \bar{z})$$

$$\mathcal{A} \models \forall w_1, \dots, w_m \bar{x} \bar{y} \exists \bar{z} R(w_1, \dots, w_m, \bar{x}, \bar{y}, \bar{z}) \rightarrow \exists ! \bar{z} R(w_1, \dots, w_m, \bar{x}, \bar{y}, \bar{z})$$

- $\mathcal{A} \models \Psi(w_1, \dots, w_m)$ if, and only if

in the deterministic, commutative Hypergraph specified by

$R(w_1, \dots, w_m, \bar{x}, \bar{y}, \bar{z})$ on $|(\mathcal{A}, w_1, \dots, w_m)|^j$, we have, for all $u \in |\mathcal{A}|$, $u^j = (u, \dots, u)$ is accessible from w_1^j .

- $\mathcal{A} \models \exists w_1, \dots, w_m \Psi(w_1, \dots, w_m)$.

Then $P = \pm\text{PS}^{\max\{k,l\}+1}[\mathbf{FO}]$ on the class \mathbf{C} .

Any tuple (w_1, \dots, w_m) s.t. $\mathcal{A} \models \psi(w_1, \dots, w_m)$ may be considered a *generating* tuple for \mathcal{A} . Generating tuples will be denoted (g_1, \dots, g_m) , and their underlying generating set $\{g_1, \dots, g_m\}$ as G .

The principle results of [42] were:

- On the class of locally-ordered strongly connected digraphs, $P = \pm\text{PS}^1[\mathbf{FO}]$.
- On the class of planar triangulations, $P = \pm\text{PS}^2[\mathbf{FO}]$.

In these cases, the construction of ψ , as in the proposition, is fairly straightforward. However, if we are prepared to sacrifice a few levels in $\pm\text{PS}^*[\mathbf{FO}]$, we can disregard ψ altogether.

Given some relation $R \in \pm\text{PS}^k[\mathbf{FO}]$, as in the proposition, there must necessarily be some $\psi \in \text{PS}^{k+3}[\mathbf{FO}]$ that will satisfy the required conditions. We may take $\psi(w_1, \dots, w_m) :=$

$$\forall u \text{ PS}[\lambda \bar{x}, \bar{y}, \bar{z}, R(w_1, \dots, w_m, \bar{x}, \bar{y}, \bar{z})](w_1^j, u^j)$$

Since we may write the $\forall u$ as $\neg \exists u \neg$, $R \in \pm\text{PS}^k[\mathbf{FO}]$ indeed implies that $\psi \in \pm\text{PS}^{k+3}[\mathbf{FO}]$.

The following is now immediate:

Corollary. *Let \mathbf{C} be any class of structures, and let $\bar{x}, \bar{y}, \bar{z}$ be variable j -tuples. Suppose there is a formula $R(w_1, \dots, w_m, \bar{x}, \bar{y}, \bar{z}) \in \pm\text{PS}^k[\mathbf{FO}]$ such that, for all $A \in \mathbf{C}$, there exists a generating tuple $(g_1, \dots, g_m) \in |A|^m$ such that:*

- $R(g_1, \dots, g_m, \bar{x}, \bar{y}, \bar{z})$ is commutative in \bar{x} and \bar{y} , and deterministic in \bar{z}
- For all $u \in |A|$, u^j is R -accessible from $(g_1)^j$.

Then $P = \pm\text{PS}^{k+4}[\mathbf{FO}]$ on the class \mathbf{C} .

5.1 Finitely generated sets

Intuitively, a set A on which some partial functions are defined is described as m -generated if it has a (generating) subset of m elements such that all elements of

A may be obtained by (possibly nested) applications of these partial functions on the elements of this subset.

When F is a finite set of partial functions, each of some finite arity, we will want to define the set F^* of all functions that can be created from those in F by repeated relabelling and substitution. So long as F contains a non-unary function, F^* must be infinite (even under equivalent relabellings), since it will have functions of all arities. Throughout this chapter we will use a bracketed superscript to indicate the arity of variable tuples or partial functions. Thus, whilst for an element x , x^k denotes *the* k -tuple of x s, the notation $\bar{v}^{(k)}$ specifies a variable k -tuple, whose different positions may hold different values.

Definition. For some finite structure \mathcal{A} , let $F = \{f_1, \dots, f_j\}$ be partial functions of respective arities a_1, \dots, a_j (i.e. $f_1^{(a_1)} : |\mathcal{A}|^{a_1} \rightarrow |\mathcal{A}|$, \dots , $f_j^{(a_j)} : |\mathcal{A}|^{a_j} \rightarrow |\mathcal{A}|$). Then F^* is defined inductively via:

- $f_1(v_1, \dots, v_{a_1}), \dots, f_j(v_1, \dots, v_{a_j}) \in F^*$.
- (Projection/Reordering.) If $f \in F^*$ of arity a , and $a' \leq a$, then let $\{n_1, \dots, n_a\}$ and $\{n'_1, \dots, n'_{a'}\}$ be subsets of \mathbb{Z} of order a and a' , respectively. If we have a function $p : \{n_1, \dots, n_a\} \rightarrow \{n'_1, \dots, n'_{a'}\}$, then $f' \in F^*$, of arity a' , where:

$$f'(v_{n_1}, \dots, v_{n_{a'}}) := f(v_{p(n_1)}, \dots, v_{p(n_a)})$$

- (Composition.) If $f, f' \in F^*$, with respective arities a, a' , then $f'' \in F^*$, of arity $(a - 1 + a')$, where:

$$f''(v_1, \dots, v_{a-1}, v_a, \dots, v_{a+a'-1}) := f(v_1, \dots, v_{a-1}, f'(v_a, \dots, v_{a+a'-1}))$$

Projection/Reordering is nothing more than relabelling of the variables. Because of the Reordering rule, we have no need to explicitly mention compositions that occur other than at the right hand end of the outer partial function. The minimum depth of nestings of Composition in a partial function $f \in F^*$ will be known as the *rank* of f in F^* .

Definition. Let $F = \{f_1, \dots, f_j\}$ be a set of functions of respective arities a_1, \dots, a_j . Let $\sigma_F = \langle f_1, \dots, f_j \rangle$ be the associated signature. Then:

- $\text{STRUC}(\sigma_F, m)$ is the class of finite structures over σ_F , such that, for all $\mathcal{A} \in \text{STRUC}(\sigma_F, m)$, there exists a generating subset $G = \{g_1, \dots, g_m\} \subseteq |\mathcal{A}|$, such that, for every $u \in |\mathcal{A}|$, there exists an arity r , a $f^{(r)} \in F^*$, and a $\bar{w}^{(r)} = (w_1, \dots, w_r) \in G^r$, such that $u = f^{(r)}(\bar{w}^{(r)})$.

$\text{STRUC}(\sigma_F, m)$ is said to be the class of structures that can be *m-generated* by the set of partial functions F .

5.1.1 F contains a single k -ary partial function f_0 .

Theorem 28. *For each m , we have that $\text{P} = \pm\text{PS}^4[\mathbf{FO}]$ on the class $\text{STRUC}(\sigma_{\{f_0\}}, m)$.*

Proof. Let k be the arity of the partial function f_0 , and let $\bar{x}, \bar{y}, \bar{z}$ be variable $(k+1)$ -tuples. We will define a deterministic, commutative Hypergraph relation $R(w_1, \dots, w_m, \bar{x}, \bar{y}, \bar{z})$, in quantifier-free \mathbf{FO} , such that, for all structures $\mathcal{A} \in \text{STRUC}(\sigma_{\{f_0\}}, m)$, there exists $g_1, \dots, g_m \in |\mathcal{A}|$ such that, for all $u \in |\mathcal{A}|$, u^{k+1} is accessible from g_1^{k+1} . We may then appeal to the Corollary of Proposition 27.

We will specify $R(w_1, \dots, w_m, \bar{x}, \bar{y}, \bar{z})$ as $R(\bar{x}, \bar{y}, \bar{z})$, where the entries of $\bar{x}, \bar{y}, \bar{z}$ may be among the variables w_1, \dots, w_m . We will define R over $(k+1)$ -tuples from $(\{P, Q, S, \sqcup\} \uplus |\mathcal{A}|)^{k+1}$.

The symbols P , Q , and S are used for switching rules, and \sqcup represents blank. (Note that we can enlarge our alphabet to include these special symbols in precisely the manner we did in Section 3.4. Thus, each ‘variable’ we discuss here, will, in point of fact, be a quintuple of actual variables.)

We begin with the ‘start’ rules:

- $R[(w_1^{k+1}), (w_1^{k+1}), (P, \sqcup^{k-1}, w_1)]$.
- $R[(P, \sqcup^{k-1}, w_i), (P, \sqcup^{k-1}, w_i), (P, \sqcup^{k-1}, w_{i+1})]$ for $1 \leq i < m$.

We now progress to the ‘active’ rules:

Switching:

- $R[(P, \sqcup^{k-i}, \bar{x}^{(i)}), (P, \sqcup^{k-i}, \bar{x}^{(i)}), (Q, \sqcup^{k-i}, \bar{x}^{(i)})]$ for $\bar{x}^{(i)} \in |\mathcal{A}|^i$, and $i \leq k$.
- $R[(Q, \sqcup^{k-i}, \bar{x}^{(i)}), (Q, \sqcup^{k-i}, \bar{x}^{(i)}), (S, \sqcup^{k-i}, \bar{x}^{(i)})]$ for $\bar{x}^{(i)} \in |\mathcal{A}|^i$, and $i \leq k$.

Concatenation:

- $R[(P, \sqcup^{k-i}, \bar{x}^{(i)}), (Q, \sqcup^{k-j}, \bar{y}^{(j)}), (P, \sqcup^{k-j-i}, \bar{x}^{(i)}, \bar{y}^{(j)})]$ for $\bar{x}^{(i)} \in |\mathcal{A}|^i$, $\bar{y}^{(j)} \in |\mathcal{A}|^j$, and $i + j \leq k$.

Production:

- $R[(Q, \sqcup^i, \bar{x}^{(k-i)}), (S, \sqcup^j, \bar{y}^{(k-j)}), (P, \sqcup^{k-1}, z)]$ for $i + j = k$, and $f_0(\bar{x}, \bar{y}) = z$.

and the ‘finish’ rule:

- $R[(P, \sqcup^{k-1}, x), (S, \sqcup^{k-1}, x), (x^{k+1})]$ for $x \in |\mathcal{A}|$.

Finally, we consider R to be the symmetric closure of the above rules, *i.e.* for all $\bar{x}^{(k+1)}, \bar{y}^{(k+1)}, \bar{z}^{(k+1)}$,

$$R(\bar{x}^{(k+1)}, \bar{y}^{(k+1)}, \bar{z}^{(k+1)}) \Rightarrow R(\bar{y}^{(k+1)}, \bar{x}^{(k+1)}, \bar{z}^{(k+1)}).$$

This ensures the commutativity of R .

R is clearly deterministic in $\bar{z}^{(k+1)}$, and can be written in quantifier-free **FO**. We will now prove that, for all $u \in |\mathcal{A}|$, u^{k+1} is accessible from g_1^{k+1} . It follows, from the start and finish rules, that this is equivalent to the question of whether (P, \sqcup^{k-1}, u) is accessible from the collection $(P, \sqcup^{k-1}, g_1), \dots, (P, \sqcup^{k-1}, g_m)$.

We know that, for each such u , there exists a partial function $f^{(r)} \in \{f_0\}^*$, and tuple $(w_1, \dots, w_r) \in \{g_1, \dots, g_m\}^r$, such that $u = f^{(r)}(w_1, \dots, w_r)$. We prove (P, \sqcup^{k-1}, u) is accessible by induction on the rank of $f^{(r)}$.

(Base Case.) When the rank of $f^{(r)}$ is 0, then $r \leq k$, and it follows that $u = f_0(w_1, \dots, w_r)$ for some $w_1, \dots, w_r \in \{g_1, \dots, g_m\}$. We may access (P, \sqcup^{k-1}, u) from $(P, \sqcup^{k-1}, g_1), \dots, (P, \sqcup^{k-1}, g_m)$ by repeated use of Switching, then repeated Concatenation, and finally a single application of Production.

(Inductive Step). Assuming it works for rank δ , we prove it works for rank $\delta + 1$. If $f^{(r)}$ is of rank $\delta + 1$, then it follows from the definition of rank, and the inductive hypothesis, that $u = f_0(w_1, \dots, w_r)$, where $(P, \sqcup^{k-1}, w_1), \dots, (P, \sqcup^{k-1}, w_r)$ have been accessed (since w_1, \dots, w_r are generated by partial functions of strictly lower rank). Again, we access (P, \sqcup^{k-1}, u) by repeated use of Switching, then repeated Concatenation, then a single application of Production. \square

5.1.2 F contains multiple partial functions.

Theorem 29. *Let F be a finite set of partial function symbols. For each m , we have that $P = \pm\text{PS}^4[\mathbf{FO}]$ on the class $\text{STRUC}(\sigma_F, m)$.*

Proof. We reduce this case to the previous. Suppose F contains j partial functions of respective arities a_1, \dots, a_j . Let $a = \max\{a_1, \dots, a_j\}$. We aim to construct a single partial function f_F , of arity $(a + j)$, that simulates all the functions in F . $\bar{w}^{(j)} = (w_1, \dots, w_j)$ will represent functions f_1 to f_j according to our ubiquitous scheme:

- if $w_1 = w_2$ then $\bar{w}^{(j)}$ represents f_1 ,
- if $w_1 \neq w_2$ but $w_2 = w_3$ then $\bar{w}^{(j)}$ represents f_2 ,
- \vdots
- if $w_1 \neq w_2, \dots, w_{j-2} \neq w_{j-1}$ but $w_{j-1} = w_j$ then $\bar{w}^{(j)}$ represents f_{j-1} , and
- if $w_1 \neq w_2, \dots, w_{j-1} \neq w_j$ then $\bar{w}^{(j)}$ represents f_j .

Suppose \hat{w} represents f_i whose arity is a_i . Consider $f_F(\bar{x}^{(a_i)}, \bar{y}^{(a-a_i)}, \hat{w})$ to be $f_i(\bar{x}^{(a_i)})$, if f_i is defined at $\bar{x}^{(a_i)}$, and undefined otherwise.

Since j is fixed, this construction of f_F can be specified in quantifier-free \mathbf{FO} . \square

5.1.3 An application: finitely generated groups

We say that a finite group H is m -generated if there exists a set of m generating elements $G = \{g_1, \dots, g_m\}$, such that for every $x \in H$ we have some $y \in G^*$ such

that $x \equiv_H y$ (where $*$ is the usual Kleene star). Clearly, all groups of order $\leq m$ are m -generated.

Corollary. *For each m , on the class of m -generated finite groups, $\pm\text{PS}^4[\mathbf{FO}] = \mathbf{P}$.*

Corollary. *On the class of finite simple groups, $\pm\text{PS}^4[\mathbf{FO}] = \mathbf{P}$.*

Proof. Recall that finite simple groups are 2-generated [3, 39]. □

Despite finitely generated groups being a paradigm for our systems, it should be noted that the results of the previous section go well beyond groups; beyond single functions, and beyond associativity.

5.2 Hamiltonian Outerplanar graphs

A graph is said to be *outerplanar* if it can be drawn in the plane with all its vertices on the outer face. Such a drawing will be called an OP-drawing.

Definition. A Hamilton cycle in a graph \mathcal{G} , where $|\mathcal{G}| = n$, is a sequence s_{cyc} of distinct vertices v_i (for $1 \leq i \leq n$) such that, for $1 \leq i < n$, $E^{\mathcal{G}}(v_i, v_{i+1})$, and also $E^{\mathcal{G}}(v_n, v_1)$.

We consider a *hamiltonian outerplanar graph* (HOP) to be an antireflexive, undirected, outerplanar graph with a Hamilton cycle. We start by noting some basic properties of outerplanar graphs that have a Hamilton cycle.

Lemma 30.

- (i) *Consider an HOP graph \mathcal{G} , with Hamilton cycle s_{cyc} . Then, in any OP-drawing of \mathcal{G} , s_{cyc} must be on the outer face.*
- (ii) *For any HOP graph \mathcal{G} , the subgraph given by any Hamilton cycle is unique.*
- (iii) *\mathcal{G} has a unique OP-drawing in the plane, up to combinatorial isomorphism.*

Proof. (i) Note that any OP-drawing of s_{cyc} is combinatorially equivalent to the n -gon ($n = |\mathcal{G}|$). Thus, s_{cyc} must appear on the outer face of any OP-drawing.

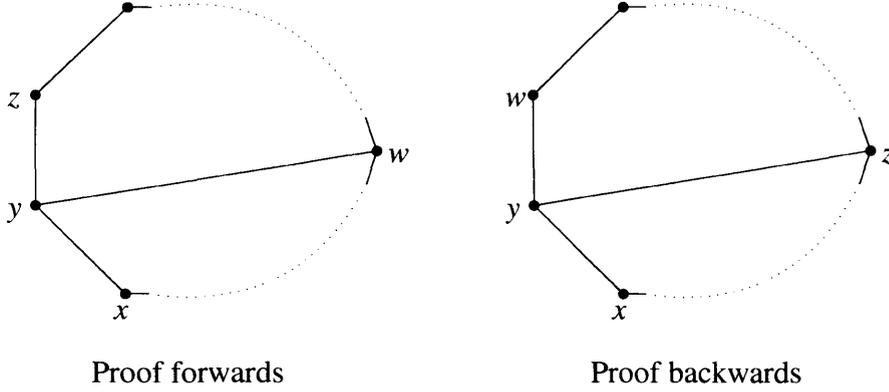


Figure 5.1: Diagrams for Lemma 31.

(ii) Consider a graph \mathcal{G} with two Hamilton cycles, s_{cyc} and s'_{cyc} , that give rise to different subgraphs. In any OP-drawing of \mathcal{G} , s_{cyc} and s'_{cyc} must be drawn as distinct n -gons over the same vertices. Yet not both can be on the outer face, violating part (i).

(iii) The unique Hamilton cycle subgraph dictates the unique OP-drawing. \square

We make use of this unique OP-drawing by now referring, unambiguously, to *the* outer face.

Lemma 31. *There is a formula $\varphi(x, y, z) \in \pm\text{PS}^3[\mathbf{FO}]$ that holds on an HOP \mathcal{G} if, and only if, x and z are the distinct neighbours of y , on the outer face.*

Proof. We first define $P(x, z, y, w) \in \pm\text{PS}^1[\mathbf{FO}]$, intended to mean that there is a path from x to z avoiding both y and w . We define P as the Transitive Closure (though in Path System logic) of the following formula θ :

$$\theta(p, q, y, w) := p \neq y \wedge p \neq w \wedge q \neq y \wedge q \neq w \wedge E(p, q)$$

Thus,

$$P(x, z, y, w) := \text{PS}[\lambda p, p, q\theta](x, z)$$

Now, $\varphi(x, y, z) :=$

$$E(x, y) \wedge E(y, z) \wedge x \neq z \wedge \\ \forall w[(w \neq x \wedge w \neq z \wedge E(y, w) \rightarrow \neg P(x, z, y, w))]$$

We now prove that x and z are the distinct neighbours of y , on the outer face of \mathcal{G} , if, and only if, $\mathcal{G} \models \varphi(x, y, z)$.

(Forwards.) If x and z are the distinct neighbours of y on the outer face, then the first three conjuncts of φ are clearly satisfied. Furthermore, the edge between any distinct w and y must cut across the OP-drawing of \mathcal{G} (see Figure 5.1). It follows that all paths from x to z must go through either y or w . Hence $\mathcal{G} \models \varphi(x, y, z)$.

(Backwards.) Suppose x and z are not the distinct neighbours of y on the outer face. If x and z are not distinct, or y is not adjacent to both, then we fail on one of the first three conjuncts of φ . So, assume x and z are distinct, and y is adjacent to both, but x and z are not the two neighbours of y on the outer face. If we choose some distinct w that is such a neighbour, then there is clearly a path from x to z avoiding both y and w (see Figure 5.1). In any case, $\mathcal{G} \not\models \varphi(x, y, z)$. \square

Theorem 32. *On the class HOP, $P = \pm\text{PS}^7[\mathbf{FO}]$.*

Proof. We will define a deterministic, commutative Hypergraph relation $R(w_1, w_2, w_3, x_1, x_2, y_1, y_2, z_1, z_2)$, in $\pm\text{PS}^3[\mathbf{FO}]$ s.t. for all structures $\mathcal{G} \in \text{HOP}$, there exists $g_1, g_2, g_3 \in |\mathcal{G}|$ such that, for all $u \in |\mathcal{G}|$, u^2 is accessible from g_1^2 . We may then appeal to the Corollary of Proposition 27, with $j = 2$.

The rules will be the symmetric closure of:

- $R[(w_1, w_1), (w_1, w_1), (w_1, w_2)]$
- $R[(w_2, w_2), (w_2, w_2), (w_2, w_3)]$
- $R[(w, x), (x, y), (y, z)]$ if $\varphi(w, x, y) \wedge \varphi(x, y, z)$ (where φ is as in Lemma 31).
- $R[(w, x), (w, x), (x, x)]$ if $w \neq x$.

That the rules are deterministic, commutative, and can be written in $\pm\text{PS}^3[\mathbf{FO}]$ is straightforward. It is also clear that, starting with any g_1, g_2, g_3 such that $\varphi(g_1, g_2, g_3)$, all vertices are accessible: from g_1 we access g_2 (*some* next vertex

on the outer face – which determines whether we are moving clockwise or anti-clockwise around a certain OP-drawing of \mathcal{G}), then g_3 (*the* next vertex on the outer face – now direction is set), then all the way round the outer face until we reach the final vertex on the Hamilton cycle (before we reach g_1 again). \square

Remark. We can easily extend our result to hamiltonian outerplanar graphs that are not undirected or not antireflexive, by considering their undirected, antireflexive versions. Specifically change all instances of $E(x,y)$, in the prior discourse, to $(E(x,y) \vee E(y,x)) \wedge x \neq y$.

Chapter 6

Dichotomies in Boolean Constraint Satisfaction

6.1 Introduction

Let σ range over all relational signatures. We define the relational class of boolean structures BOOL .

$$\text{BOOL} := \{\mathcal{A} : \mathcal{A} \text{ is a } \sigma\text{-structure and } \|\mathcal{A}\| = 2\}$$

We denote tuples of variables (resp. boolean constants) in bold, e.g. \mathbf{x} (resp. \mathbf{t}). These tuples are not of a uniform arity.

Definition. For $\mathcal{A} \in \text{BOOL}$, and each with the question as to whether $\mathcal{A} \models \varphi$,

- the problem $\text{SAT}_{\text{NC}}(\mathcal{A})$ has input $\varphi := \exists \mathbf{x} Q(\mathbf{x})$,
- the problem $\text{QSAT}_{\text{NC}}(\mathcal{A})$ has input, for some $n \geq 1$, of the form,

$$\varphi := \forall \mathbf{x}_1 \exists \mathbf{x}_2 \forall \mathbf{x}_3 \exists \mathbf{x}_4 \dots \forall \mathbf{x}_{2n+1} \exists \mathbf{x}_{2n+2} Q(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{2n+2})$$

- the problem $\Pi_{2n+1}\text{-SAT}_{\text{NC}}(\mathcal{A})$ has input,

$$\varphi := \forall \mathbf{x}_1 \exists \mathbf{x}_2 \forall \mathbf{x}_3 \exists \mathbf{x}_4 \dots \forall \mathbf{x}_{2n+1} Q(\mathbf{x}_1, \mathbf{x}_2 \dots \mathbf{x}_{2n+1})$$

- the problem $\Pi_{2n+2}\text{-SAT}_{\text{NC}}(\mathcal{A})$ has input,

$$\varphi := \forall \mathbf{x}_1 \exists \mathbf{x}_2 \forall \mathbf{x}_3 \exists \mathbf{x}_4 \dots \forall \mathbf{x}_{2n+1} \exists \mathbf{x}_{2n+2} Q(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{2n+2})$$

- the problem $\Sigma_{2n+1}\text{-SAT}_{\text{NC}}(\mathcal{A})$ has input,

$$\varphi := \exists \mathbf{x}_1 \forall \mathbf{x}_2 \exists \mathbf{x}_3 \forall \mathbf{x}_4 \dots \exists \mathbf{x}_{2n+1} Q(\mathbf{x}_1, \mathbf{x}_2 \dots \mathbf{x}_{2n+1})$$

- the problem $\Sigma_{2n+2}\text{-SAT}_{\text{NC}}(\mathcal{A})$ has input,

$$\varphi := \exists \mathbf{x}_1 \forall \mathbf{x}_2 \exists \mathbf{x}_3 \forall \mathbf{x}_4 \dots \forall \mathbf{x}_{2n+2} Q(\mathbf{x}_1, \mathbf{x}_2 \dots \mathbf{x}_{2n+1})$$

where, in each case, Q is a conjunction of positive atoms.

The problems $\text{SAT}_{\text{C}}(\mathcal{A})$ etc are defined analogously, but with the two boolean constants 0 and 1 built-in to the signature.

It is clear that $\Pi_{2n+2}\text{-SAT}_{\text{NC}}(\mathcal{A})$ (respectively, $\Sigma_{2n+1}\text{-SAT}_{\text{NC}}(\mathcal{A})$) is in the complexity class Π_{2n+2}^{P} (respectively, Σ_{2n+1}^{P}). In fact, it follows from [43] that they are complete for those classes, for certain \mathcal{A} . It is also clear that $\Pi_1\text{-SAT}_{\text{NC}}(\mathcal{A})$ is tractable, for all \mathcal{A} , since we may check each extensional relation independently, one-by-one, for an invalidating assignment. Indeed, if the maximum arity of a relation in \mathcal{A} is a , then the complexity of $\Pi_1\text{-SAT}_{\text{NC}}(\mathcal{A})$ is $O(n^a)$, where n is the size of the input. It follows, by similar argument [43], that $\Pi_{2n+1}\text{-SAT}_{\text{NC}}(\mathcal{A})$ (respectively, $\Sigma_{2n+2}\text{-SAT}_{\text{NC}}(\mathcal{A})$) is in the complexity class Π_{2n}^{P} (respectively, Σ_{2n+1}^{P}).

The comments of the previous paragraph apply equally to the problems Π_i - and $\Sigma_i\text{-SAT}_{\text{C}}$, *i.e.* in the situation where the Boolean constants are available.

Definition. For a relation R , of arity a , define:

- $\exists\text{-FORM}(R)$ to be the set of formulae formed from the closure of the atoms $R(\mathbf{x})$ (where \mathbf{x} is an a -tuple of not necessarily distinct variables), under conjunction and existential quantification.
- $\Pi_2\text{-FORM}(R)$ to be the set of formulae of the form $\forall \mathbf{x} \varphi(\mathbf{x}, \mathbf{y})$, where $\varphi \in \exists\text{-FORM}(R)$.

- \forall/\exists -FORM(R) to be the set of formulae formed from the closure of the atoms $R(\mathbf{x})$, under conjunction, existential quantification, and universal quantification.

We define \exists -REL(R) to be the set of relations expressible by formulae in \exists -FORM(R), when reading the variables lexicographically. We do likewise for Π_2 -REL(R) and \forall/\exists -REL(R). These sets are sometimes known as relational clones [14].

We may refer to boolean relations by some propositional formula that expresses them, reading the propositional variables lexicographically, e.g. $[A \vee B]$ expresses $\{(0, 1), (0, 1), (1, 1)\}$; $[A \neq B]$ expresses $\{(0, 1), (1, 0)\}$.

Definition. A relation R , of arity a , is:

- (i) *0-valid* iff it contains the tuple (0^a) .
- (ii) *1-valid* iff it contains the tuple (1^a) .
- (iii) *horn* iff it may be expressed by a propositional formula in CNF where each clause has at most one positive literal.
- (iv) *dual horn* iff it may be expressed by a propositional formula in CNF where each clause has at most one negative literal.
- (v) *bijunctive* iff it may be expressed by a propositional formula in 2-CNF.
- (vi) *affine* iff it may be expressed by a propositional formula that is the conjunction of linear equations over \mathbb{Z}_2 .

Given a template \mathcal{A} , over signature involving relations R_1, \dots, R_j , of respective arities a_1, \dots, a_j , we construct the relation $R^{\mathcal{A}}$ thus:

- If \mathcal{A} has $j' \leq j$ non-empty relations, then let R'_i be the i th non-empty relation of \mathcal{A} .
- Let $R^{\mathcal{A}} = R'_1 \times \dots \times R'_{j'}$.

This construction will enable us to consider signatures with multiple relations, as though they only had one. This is because each of the six attributes from the previous definition hold over all the relations of \mathcal{A} if, and only if, they hold for $R^{\mathcal{A}}$. Note that all the relations of \mathcal{A} , except possibly the empty relation ϕ , are present in $\exists\text{-REL}(R_{\mathcal{A}})$.

Theorem 33 (I-III: Schaefer [38], and IV: Dalmau/Creignou et al [15, 14]).

- I. $\text{SAT}_{\text{C}}(\mathcal{A})$ is tractable if $R^{\mathcal{A}}$ satisfies any of conditions (iii) – (vi), and is NP-complete otherwise.
- II. $\text{SAT}_{\text{NC}}(\mathcal{A})$ is tractable if $R^{\mathcal{A}}$ satisfies any of conditions (i) – (vi), and is NP-complete otherwise.
- III. $\text{QSAT}_{\text{C}}(\mathcal{A})$ is tractable if $R^{\mathcal{A}}$ satisfies any of conditions (iii) – (vi), and is Pspace-complete otherwise.
- IV. $\text{QSAT}_{\text{NC}}(\mathcal{A})$ is tractable if $R^{\mathcal{A}}$ satisfies any of conditions (iii) – (vi), and is Pspace-complete otherwise.

We will briefly consider the methods involved in proving these dichotomies. When the boolean constants are present, Schaefer was able to take any $R^{\mathcal{A}}$ not in classes (i) – (iv), and construct the ternary boolean not-all-equal relation, which is known to give rise to an NP-complete SAT, and Pspace-complete QSAT. When constants are not present, there are the degenerate cases of 0- and 1-validity, which become trivial. For the other tractable sub-classes of templates \mathcal{A} , we clearly have that $\text{SAT}_{\text{NC}}(\mathcal{A})$ is polynomially reducible to $\text{SAT}_{\text{C}}(\mathcal{A})$, guaranteeing its tractability in the no-constants scenario. It remained for him to prove that $\text{SAT}_{\text{C}}(\mathcal{A})$ is polynomially reducible to $\text{SAT}_{\text{NC}}(\mathcal{A})$ for those \mathcal{A} not in classes (i) – (vi). He did this by simulating the boolean constants. Call a relation $R_{\mathcal{A}}$ *complementative* if, for all tuples \mathbf{x} in $R_{\mathcal{A}}$, the tuple \mathbf{x}' , obtained from \mathbf{x} by swapping the 0s and 1s, is also in $R_{\mathcal{A}}$. Schaefer proved the following:

Lemma 34 ([38]). *For some $R_{\mathcal{A}}$, not in classes (i) – (vi), either $[A], [\neg A] \in \exists\text{-REL}(R_{\mathcal{A}})$, or $[A \neq B] \in \exists\text{-REL}(R_{\mathcal{A}})$ and $R_{\mathcal{A}}$ is complementative.*

Before going further, we will need the following lemma.

Lemma 35 ([15]). *If $R_{\mathcal{A}}$ is complementative, then all relations in $\forall/\exists\text{-REL}(R_{\mathcal{A}})$ are complementative.*

Proof. We prove this by induction on the term-complexity of $\varphi \in \forall/\exists\text{-REL}(R_{\mathcal{A}})$. The base case is trivial. For the inductive step, note that:

- $R(\mathbf{x})$ and $R'(\mathbf{x}')$ complementative, implies $R \wedge R'(\mathbf{x}, \mathbf{x}')$ complementative.
- $R(\mathbf{x})$ complementative, implies $\exists x_1 R(\mathbf{x})$ is complementative.
- $R(\mathbf{x})$ complementative, implies $\forall x_1 R(\mathbf{x})$ is complementative.

□

We can now sketch Schaefer's result, and method.

Proposition 36. *If $R^{\mathcal{A}}$ is in none of the classes (i) – (vi) above, then $\text{SAT}_{\mathcal{C}}(\mathcal{A}) \leq_{\text{P}} \text{SAT}_{\text{NC}}(\mathcal{A})$.*

Proof. By Lemma 34, we need to consider two cases.

(Case 1.) We have $\exists \mathbf{y}_1 Q_1(\mathbf{y}_1, b), \exists \mathbf{y}_0 Q_0(\mathbf{y}_0, a) \in \exists\text{-FORM}(R^{\mathcal{A}})$ expressing $[X], [\neg X]$, where Q_1 and Q_0 are positive conjunctive. We are now in a position to simulate the constants 0 and 1, for, given an input $\exists \mathbf{x} \varphi(\mathbf{x}, 0, 1)$ for $\text{SAT}_{\mathcal{C}}(\mathcal{A})$, we know,

$$\exists \mathbf{x} \varphi(\mathbf{x}, 0, 1) \Leftrightarrow \exists \mathbf{x} \exists a \exists b \varphi(\mathbf{x}, a, b) \wedge \exists \mathbf{y}_1 Q_1(\mathbf{y}_1, b) \wedge \exists \mathbf{y}_0 Q_0(\mathbf{y}_0, a)$$

The latter formula is an input for $\text{SAT}_{\text{NC}}(\mathcal{A})$, when the inner existential quantifiers are drawn out, putting it in prenex form.

(Case 2.) We have $\exists \mathbf{y} Q'(\mathbf{y}, a, b) \in \exists\text{-FORM}(R^{\mathcal{A}})$ that expresses $[A \neq B]$, and $R^{\mathcal{A}}$ is complementative. It follows that,

$$\exists \mathbf{x} \varphi(\mathbf{x}, 0, 1) \in \text{SAT}_{\mathcal{C}}(\mathcal{A}) \Leftrightarrow \exists \mathbf{x} \exists a \exists b \varphi(\mathbf{x}, a, b) \wedge \exists \mathbf{y} Q'(\mathbf{y}, a, b) \in \text{SAT}_{\text{NC}}(\mathcal{A})$$

since φ must be complementative (by Lemma 35). □

The problem of removing the constants in $\text{QSAT}_{\mathcal{C}}$ was not attended to by Schaefer. It was finally settled many years later by Dalmau [15], and, independently, by Creignou et al. [14].

6.2 Technical results

Before progressing, we will need a number of technical lemmas.

Lemma 37 (Quantifier Re-ordering). *Let the variables a, b not appear in \mathbf{x} . The following are equivalent on all boolean structures for all conjunctive positive Q :*

$$\exists \mathbf{x} \exists a \forall b B(a, b) \wedge Q(a, \mathbf{x})$$

$$\forall b \exists \mathbf{x} \exists a B(a, b) \wedge Q(a, \mathbf{x})$$

If, and only if:

- B is \emptyset (empty), singleton, $\{(0,0), (1,0)\}$, $\{(1,1), (0,1)\}$, or
- B contains $(0,0)$ and $(0,1)$.
- B contains $(1,0)$ and $(1,1)$.

Proof. If B is \emptyset , singleton, $\{(0,0), (1,0)\}$, or $\{(1,1), (0,1)\}$ then both sentences will be false irrespective of Q .

If B contains both $(0,0)$ and $(0,1)$, it may easily be verified that both sentences are equivalent. The case where B contains both $(1,0)$ and $(1,1)$ is symmetric.

The remaining possibilities for B are $\{(0,0), (1,1)\}$ and $\{(1,0), (0,1)\}$, which will be each false in the former sentence, but may be true in the latter (e.g. if Q is logically valid). \square

Given the boolean k -tuples $\mathbf{t}_1 = (t_1^1, \dots, t_1^k)$ and $\mathbf{t}_2 = (t_2^1, \dots, t_2^k)$ we define $\mathbf{t}_1 \oplus \mathbf{t}_2$ to be $(t_1^1 + t_2^1, \dots, t_1^k + t_2^k)$ where the addition is modulo 2.

Lemma 38 (0-affine case. [13]). *Let R be a boolean relation of arity k . The following are equivalent:*

- (a) R is 0-valid and affine.
- (b) $0^k \in R$, and, for all assignments $\mathbf{t}_1, \mathbf{t}_2 \in R$, we have $\mathbf{t}_1 \oplus \mathbf{t}_2 \in R$.

Definitions. For a relation R of arity k , and any set $T = \{i_1, \dots, i_j\}$ of j positions $0 \leq i_1 < \dots < i_j \leq k$, we define $R|T$ to be the j -ary relation $\exists x_{l_1} \dots \exists x_{l_{k-j}} R(x_1, \dots, x_k)$, where $\{l_1, \dots, l_{k-j}\} = \{1, \dots, k\} - T$. Observe that $R|T$ is in $\exists\text{-REL}(R)$.

For any $\mathbf{t} \in \{0, 1\}^k$ and $T \subset \{1, \dots, k\}$, we define $\mathbf{t}|T$ as the assignment $\mathbf{t}' \in \{0, 1\}^{|T|}$ that agrees with \mathbf{t} in the positions indexed by T .

Let R be a relation of arity k , and $\mathbf{t} \in \{0, 1\}^k$ an assignment. We say that \mathbf{t} is *j -compatible* (w.r.t. R) if, for every subset $T \subset \{1, \dots, k\}$ (of size $|T| < j$), we can find some assignment $\mathbf{t}' \in R$ such that \mathbf{t} and \mathbf{t}' agree in the positions indexed by T . This is equivalent to the condition that any j -ary sub-tuple of \mathbf{t} can be extended to some k -ary \mathbf{t}' in R . Clearly this is trivially true when \mathbf{t} is itself in R ; the interesting cases are when it is not.

The notion of j -compatibility is key in characterising horn logical relations.

Lemma 39 (horn case. [15]). *Let R be a boolean relation of arity k . The following are equivalent:*

- (a) R is horn.
- (b) for all $\mathbf{t}_1, \mathbf{t}_2 \in R$, we have $\mathbf{t}_1 \wedge \mathbf{t}_2 \in R$.
- (c) for every $T \subset \{1, \dots, k\}$, and for every $|T|$ -compatible (w.r.t. $R|T$) assignment $\mathbf{t} \in \{0, 1\}^{|T|}$ not in $R|T$, we have that \mathbf{t} contains at most a single 0.

Lemma 40 (Adapted from [15]). *If R^A is 0-valid and non-Horn, then there is a relation S_λ^4 definable in $\exists\text{-REL}(R^A)$ such that:*

$$\begin{array}{cc} \in S_\lambda^4 & \notin S_\lambda^4 \\ (0, 0, 0, 0) & (0, 0, 0, 1) \\ (0, 1, 0, 1) & \\ (0, 0, 1, 1) & \end{array}$$

Proof. Since R^A is non-horn, we may guarantee to break part (c) of the previous lemma. This implies that there is a subset of indices $T \subseteq \{1, \dots, k\}$ and a $|T|$ -compatible assignment \mathbf{t} not in $R|T$ that contains at least two zeros. We may benefit from dwelling on what exactly this means. It guarantees us some $|T|$ -ary relation $R|T$, in $\exists\text{-REL}(R^A)$, and a $|T|$ -tuple \mathbf{t} s.t.

- $\mathbf{t} \notin R|T$,
- for any \mathbf{t}' that agrees with \mathbf{t} in all but one position, $\mathbf{t}' \in R|T$, and
- \mathbf{t} contains at least two zeros.

We therefore consider $R|T(v_1, \dots, v_{|T|})$ and two indices $\alpha, \beta \in \mathbb{Z}_{|T|}$ at which \mathbf{t} has zeros. Note that \mathbf{t} can not be all zeros, since R is 0-valid yet $\mathbf{t} \notin R|T$. Let I' be the set of indices at which \mathbf{t} is one. Finally, let I'' be the set of indices, other than i, j , at which \mathbf{t} is zero¹. We obtain $S_\lambda^4(x, v_\alpha, v_\beta, y)$ from $R|T(v_1, \dots, v_{|T|})$ by substituting all variables v_i s.t. $i \in I'$ by the variable y and substituting all variables v_i s.t. $i \in I''$ by the variable x .

We already know that $(0, 0, 0, 1) \notin S_\lambda^4$, and we have assumed that R is 0-valid, hence $(0, 0, 0, 0) \in S_\lambda^4$. Since $(0, 0, 0, 1) \notin S_\lambda^4$, we will have $(0, 1, 0, 1), (0, 0, 1, 1) \in S_\lambda^4$ by the $|T|$ -compatibility, since these two assignments each only change the value of a single variable in $R|T$. \square

Lemma 41 (Adapted from [15]). *If R^A is 0-valid and non-affine, then there is a relation S_μ^4 definable in $\exists\text{-REL}(R^A)$ such that:*

$$\begin{array}{cc} \in S_\mu^4 & \notin S_\mu^4 \\ (0, 0, 0, 0) & (0, 1, 1, 0) \\ (0, 1, 0, 1) & \\ (0, 0, 1, 1) & \end{array}$$

Proof. Since R is 0-valid and affine, it follows from Lemma 38 that there are assignments $\mathbf{t}_1, \mathbf{t}_2 \in R$ such that $\mathbf{t}_1 \oplus \mathbf{t}_2 \notin R$. For $R(x_1, \dots, x_k)$, define:

- $V_{00} = \{v : v \in \{x_1, \dots, x_k\} \text{ } v \text{ is 0 in } t_1 \text{ and 0 in } t_2 \}$
- $V_{01} = \{v : v \in \{x_1, \dots, x_k\} \text{ } v \text{ is 0 in } t_1 \text{ and 1 in } t_2 \}$
- $V_{10} = \{v : v \in \{x_1, \dots, x_k\} \text{ } v \text{ is 1 in } t_1 \text{ and 0 in } t_2 \}$
- $V_{11} = \{v : v \in \{x_1, \dots, x_k\} \text{ } v \text{ is 1 in } t_1 \text{ and 1 in } t_2 \}$

¹It is possible that I'' is empty, in which case S_λ^4 will actually be a ternary relation. This will cause no problems, and will come out in the wash in Lemma 42.

Let $S_\mu^4(y_{00}, y_{01}, y_{10}, y_{11})$ be $R(x_1, \dots, x_k)$ with the substitutions y_{00} for V_{00} , y_{01} for V_{01} , y_{10} for V_{10} , and y_{11} for V_{11} . The claimed properties follow immediately. \square

Lemma 42. *If R^A is 0-valid and non-Horn and non-affine, then there is a relation $S_{\lambda\mu}^4$ definable in $\exists\text{-REL}(R^A)$ such that:*

$$\begin{array}{cc} \in S_{\lambda\mu}^4 & \notin S_{\lambda\mu}^4 \\ (0, 0, 0, 0) & (0, 0, 0, 1) \\ (0, 1, 0, 1) & (0, 1, 1, 0) \\ (0, 0, 1, 1) & \end{array}$$

Proof. $S_{\lambda\mu}^4 := S_\lambda^4 \wedge S_\mu^4$. \square

6.3 A dichotomy theorem for $\Pi_2\text{-SAT}_{\text{NC}}$

It follows from Schaefer's work and [43] that $\Pi_2\text{-SAT}_{\text{C}}(\mathcal{A})$ is tractable if R^A is in any of the classes (iii) – (vi), and Π_2^{P} -complete otherwise. Borrowing much from Dalmau, we will show that $\Pi_2\text{-SAT}_{\text{NC}}$ exhibits the same dichotomy.

Our proof rests on the following:

Proposition 43. *Let $\mathcal{A} \in \mathcal{B}$. If R^A is neither horn, dual horn, affine, nor bijunctive, then $\Pi_2\text{-SAT}_{\text{C}}(\mathcal{A})$ is polynomially reducible to $\Pi_2\text{-SAT}_{\text{NC}}(\mathcal{A})$.*

If R^A is neither 0-valid nor 1-valid, we may appeal to Schaefer's method for simulating the constants. This may only result in more existential quantifiers on the inside of the input instance, which will not jeopardise our being in Π_2^{P} . However, if we need formulae with universal quantifiers to simulate the constants, then we find ourselves potentially outside Π_2^{P} , with more than a single alternation of quantifiers in the input instance.

Recall that we are only concerned with R^A that are non-horn, non-dual-horn, non-bijunctive, non-affine, and either 1-valid or 0-valid. We will consider four cases.

6.3.1 Case 1 : R^A is 0-valid and not 1-valid.

In this case we have the constant 0 for free, since $R^A(a, \dots, a)$ expresses $[\neg A]$.

Let S_3^2 be the boolean relation $\{(0,0), (1,0), (1,1)\}$.

Lemma 44. *If R^A is 0-valid, non-horn, non-affine, and not 1-valid, then S_3^2 is definable in $\exists\text{-REL}(R^A)$.*

Proof. We consider two further possibilities for the relation $S_{\lambda\mu}^4$ above.

- If $S_{\lambda\mu}^4$ also contains $(0,0,1,0)$, then $S_3^2 = \exists a S_{\lambda\mu}^4(a, a, b, c) \wedge R^A(a, \dots, a)$.
- If $S_{\lambda\mu}^4$ does not contain $(0,0,1,0)$, then $S_3^2 = \exists a \exists a' S_{\lambda\mu}^4(a, a', c, b) \wedge R^A(a, \dots, a)$.

□

In both cases S_3^2 is of the form $\exists \mathbf{w} Q_S(\mathbf{w}, b, c)$, where Q_S is positive conjunctive. Note that $[A]$ is expressed by $\forall c S_3^2(b, c)$.

Lemma 45. *If R^A is 0-valid, non-horn, non-affine, and not 1-valid, then $\Pi_2\text{-SAT}_C(\mathcal{A})$ polynomially reduces to $\Pi_2\text{-SAT}_{NC}(\mathcal{A})$.*

Proof. Given an input $\forall \mathbf{x}_1 \exists \mathbf{x}_2 Q(\mathbf{x}_1, \mathbf{x}_2, 0, 1)$ for $\Pi_2\text{-SAT}_C(\mathcal{A})$, observe,

$$\begin{aligned}
& \forall \mathbf{x}_1 \exists \mathbf{x}_2 && Q(\mathbf{x}_1, \mathbf{x}_2, 0, 1) \\
\Leftrightarrow & \forall \mathbf{x}_1 \exists \mathbf{x}_2 \exists a \exists b && Q(\mathbf{x}_1, \mathbf{x}_2, a, b) \wedge \forall c S_3^2(b, c) \\
& && \wedge R^A(a, \dots, a) \\
\Leftrightarrow & \forall \mathbf{x}_1 \exists \mathbf{x}_2 \exists a \exists b && Q(\mathbf{x}_1, \mathbf{x}_2, a, b) \wedge \forall c \exists \mathbf{w} Q_S(\mathbf{w}, b, c) \\
& && \wedge R^A(a, \dots, a) \\
\Leftrightarrow & \forall \mathbf{x}_1 \forall c \exists \mathbf{x}_2 \exists a \exists b && Q(\mathbf{x}_1, \mathbf{x}_2, a, b) \wedge \exists \mathbf{w} Q_S(\mathbf{w}, b, c) \\
& && \wedge R^A(a, \dots, a)
\end{aligned}$$

Note that the final line is a valid input for $\Pi_2\text{-SAT}_{NC}(\mathcal{A})$. The final equivalence holds by the Quantifier Re-ordering Lemma, with $B := \exists \mathbf{w} Q_S(\mathbf{w}, b, c)$. □

6.3.2 Case 2 : R^A is 1-valid and not 0-valid.

This is the symmetric case of the previous (where zero is replaced with one, and vice-versa).

6.3.3 Case 3 : R^A is 0-valid and 1-valid, but not complimentary.

Since R^A is 0-valid, 1-valid, and yet not complimentary, there exists a tuple \mathbf{t} in R^A s.t \mathbf{t} 's complement is not in R^A . Let I index the set of positions at which \mathbf{t} is zero and let J index those positions at which \mathbf{t} is one. If $R^A(v_1, \dots, v_k)$ is a k -ary relation, consider $Q_{atom}(a, b)$ to be R^A under the substitution a for all variables indexed by I and b for all variables indexed by J . Q_{atom} is atomic, and it expresses S_3^2 . We now have that $\forall b S_3^2(a, b)$ expresses $[A]$, and $\forall a S_3^2(a, b)$ expresses $[\neg A]$.

Lemma 46. *If R^A is 0-valid, non-horn, non-affine, 1-valid, but not complimentary, then $\Pi_2\text{-SAT}_C(\mathcal{A})$ polynomially reduces to $\Pi_2\text{-SAT}_{NC}(\mathcal{A})$.*

Proof. Given an input $\forall \mathbf{x}_1 \exists \mathbf{x}_2 Q(\mathbf{x}_1, \mathbf{x}_2, 0, 1)$ for $\Pi_2\text{-SAT}_C(\mathcal{A})$, observe,

$$\begin{aligned}
& \forall \mathbf{x}_1 \exists \mathbf{x}_2 && Q(\mathbf{x}_1, \mathbf{x}_2, 0, 1) \\
\Leftrightarrow & \forall \mathbf{x}_1 \exists \mathbf{x}_2 \exists a \exists b && Q(\mathbf{x}_1, \mathbf{x}_2, a, b) \wedge \forall b' S_3^2(a, b') \\
& && \wedge \forall a' S_3^2(a', b) \\
\Leftrightarrow & \forall \mathbf{x}_1 \exists \mathbf{x}_2 \exists a \exists b && Q(\mathbf{x}_1, \mathbf{x}_2, a, b) \wedge \forall b' Q_{atom}(a, b') \\
& && \wedge \forall a' Q_{atom}(a', b) \\
\Leftrightarrow & \forall \mathbf{x}_1 \forall b' \forall a' \exists \mathbf{x}_2 \exists a \exists b && Q(\mathbf{x}_1, \mathbf{x}_2, a, b) \wedge Q_{atom}(a, b') \\
& && \wedge Q_{atom}(a', b)
\end{aligned}$$

The final line is a valid input for $\Pi_2\text{-SAT}_{NC}(\mathcal{A})$. The final equivalence holds via two applications of the Quantifier Re-ordering Lemma. \square

6.3.4 Case 4 : R^A is 0-valid and 1-valid, and complementative.

In this case $S_{\lambda\mu}^4$ will look like:

$$\begin{array}{ll}
 \in S_{\lambda\mu}^4 & \notin S_{\lambda\mu}^4 \\
 (0, 0, 0, 0) & (0, 0, 0, 1) \\
 (0, 1, 0, 1) & (0, 1, 1, 0) \\
 (0, 0, 1, 1) & \\
 (1, 1, 1, 1) & (1, 1, 1, 0) \\
 (1, 0, 1, 0) & (1, 0, 0, 1) \\
 (1, 1, 0, 0) &
 \end{array}$$

Note that $\forall d \exists a S_{\lambda\mu}^4(a, b, c, d)$ expresses $[B \neq C]$. Since $S_{\lambda\mu}^4 \in \exists\text{-REL}(R^A)$, it follows that $[A \neq B]$ is expressed by $\forall d \exists a \exists \mathbf{w} Q_{\lambda\mu}(\mathbf{w}, a, b, c, d)$, where $Q_{\lambda\mu}$ is positive conjunctive.

Lemma 47. *If R^A is 0-valid, non-horn, non-affine, 1-valid, and complementative, then $\Pi_2\text{-SAT}_C(\mathcal{A})$ polynomially reduces to $\Pi_2\text{-SAT}_{\text{NC}}(\mathcal{A})$.*

Proof. Given an input $\forall \mathbf{x}_1 \exists \mathbf{x}_2 Q(\mathbf{x}_1, \mathbf{x}_2, 0, 1)$ for $\Pi_2\text{-SAT}_C(\mathcal{A})$, observe,

$$\begin{aligned}
 & \forall \mathbf{x}_1 \exists \mathbf{x}_2 \quad Q(\mathbf{x}_1, \mathbf{x}_2, 0, 1) \\
 \Leftrightarrow & \forall \mathbf{x}_1 \exists \mathbf{x}_2 \exists b \exists c \quad Q(\mathbf{x}_1, \mathbf{x}_2, b, c) \wedge \forall d \exists a S_{\lambda\mu}^4(a, b, c, d) \\
 \Leftrightarrow & \forall \mathbf{x}_1 \exists \mathbf{x}_2 \exists b \exists c \quad Q(\mathbf{x}_1, \mathbf{x}_2, b, c) \wedge \forall d \exists a \exists \mathbf{w} Q_{\lambda\mu}(\mathbf{w}, a, b, c, d)
 \end{aligned}$$

Now this is not, in general, equivalent to:

$$\forall \mathbf{x}_1 \forall d \exists \mathbf{x}_2 \exists b \exists c \quad Q(\mathbf{x}_1, \mathbf{x}_2, b, c) \wedge \exists a \exists \mathbf{w} Q_{\lambda\mu}(\mathbf{w}, a, b, c, d)$$

because, when $b = c$, that formula may be true, but the previous ones are always false. However, we claim that:

$$\forall \mathbf{x}_1 \exists \mathbf{x}_2 \exists b \exists c \quad Q(\mathbf{x}_1, \mathbf{x}_2, b, c) \wedge \forall d \exists a \exists \mathbf{w} Q_{\lambda\mu}(\mathbf{w}, a, b, c, d)$$

is equivalent to

$$\forall \mathbf{x}_1 \forall d \forall d' \exists \mathbf{x}_2 \exists b \exists c Q(\mathbf{x}_1, \mathbf{x}_2, b, c) \wedge \exists a \exists \mathbf{w} Q_{\lambda\mu}(\mathbf{w}, a, b, c, d) \wedge \exists a' \exists \mathbf{w} Q_{\lambda\mu}(\mathbf{w}, a', b, c, d')$$

which is an input for $\Pi_2\text{-SAT}_{\text{NC}}(\mathcal{A})$. It remains for us to prove this equivalence.

(forwards.) This direction is trivial. For each given \mathbf{x}_1 in both formulae: any b, c, \mathbf{x}_2 that witness the first formula will also witness the second .

(backwards.) For each given \mathbf{x}_1 in both formulae: if $d \neq d'$, it follows that any true valuation of the second formula has $b \neq c$. This ensures that, if the second formula is true, that the first formula will also be, witnessed by some $b \neq c$. \square

Theorem 48. $\Pi_2\text{-SAT}_{\text{NC}}(\mathcal{A})$ is tractable if $R^{\mathcal{A}}$ is horn, dual horn, bijunctive, or affine, and is Π_2^{P} -complete otherwise.

Proof. We know $\Pi_2\text{-SAT}_{\text{C}}$ has the proposed dichotomy. Trivially, the tractability of $\Pi_2\text{-SAT}_{\text{C}}(\mathcal{A})$ implies the tractability of $\Pi_2\text{-SAT}_{\text{NC}}(\mathcal{A})$. Furthermore, we have proved, for $R^{\mathcal{A}}$ outside the listed classes, that $\Pi_2\text{-SAT}_{\text{C}}(\mathcal{A})$ polynomially reduces to the Pspace-complete $\Pi_2\text{-SAT}_{\text{NC}}(\mathcal{A})$. The result follows. \square

Corollary. For $i > 2$, $\Pi_i\text{-SAT}_{\text{NC}}$ and $\Sigma_i\text{-SAT}_{\text{NC}}$ exhibit the same dichotomy as $\Pi_2\text{-SAT}_{\text{NC}}$ and QSAT_{NC} .

Proof. Let $j > 1$. Our manipulation of the innermost universal quantifiers in the pertinent Π_3 or Π_4 formulae, such that we build equivalent ones in Π_2 , will clearly also work on Π_{2j+1} or Π_{2j+2} (resp. Σ_{2j+1} or Σ_{2j}) formulae to obtain equivalent ones in Π_{2j} (resp. Σ_{2j-1}). Consequently, our proof is equally valid for these problems. \square

Remark. We are left with the class of problems $\Sigma_2\text{-SAT}_{\text{NC}}$. As noted before, these are in NP, and they exhibit the same dichotomy as SAT_{NC} .

Remark. A similar proof to this dichotomy theorem appears in [27]. The result is also inferred in [18].

Some recent work has been undertaken in alternation-bounded QCSP, by Chen [12]. He studies certain templates for which the complexity of the problem collapses to co-NP-completeness for all levels of the polynomial hierarchy above or equal to Π_2 .

Chapter 7

Quantified Constraints on Graphs

7.1 Introduction

The *uniform constraint satisfaction problem*, as used in Artificial Intelligence, is usually defined as follows (see e.g. [32]).

- Input: a finite set of variables A , a finite domain of values T , and a set of constraints $\{C(S_1), \dots, C(S_c)\}$ where each S_i is an a_i -tuple of (not necessarily distinct) variables from A and each $C(S_i)$ is an a_i -ary relation over T .
- Question: is there an assignment to the variables over the domain that mutually satisfies all of the constraints?

It is clear that T , together with the relations S_1, \dots, S_c , is a first-order structure \mathcal{T} (over some signature σ of the form $\langle S_1^{a_1}, \dots, S_c^{a_c} \rangle$). It is also clear that the question we are posing of this structure concerns the existence of a simultaneous solution to a conjunction of atomic relational constraints. Therefore, we will prefer to use the following formulation of the uniform CSP (see e.g. [5]).

- Input: a structure \mathcal{T} and a sentence $\varphi = \exists \mathbf{x} Q(\mathbf{x})$, where Q is a conjunction of positive atoms.
- Question: does $\mathcal{T} \models \varphi$?

In this thesis we will be concerned only with the *non-uniform* variant of CSP, which is a family of problems parameterised by the template \mathcal{T} . Thus, for each template \mathcal{T} , $\text{CSP}(\mathcal{T})$ is the decision problem with:

- Input: a sentence $\varphi = \exists \mathbf{x} Q(\mathbf{x})$, where Q is a conjunction of positive atoms.
- Question: does $\mathcal{T} \models \varphi$?

Note that the well-known NP-complete problems 3-SAT (satisfiability of a formula in conjunctive normal form with exactly three literals per clause) and 3-COL (graph 3-colourability) correspond to constraint satisfaction problems. The problem 3-QSAT is a popular generalisation of 3-SAT to quantified formulae, which is Pspace-complete. In this context, it makes sense to generalise constraint satisfaction problems to *quantified constraint satisfaction problems*.

Definition ([5]). The *non-uniform quantified constraint satisfaction problem* with template \mathcal{T} , denoted by $\text{QCSP}(\mathcal{T})$, is the decision problem with:

- Input: a sentence ψ of the form

$$\forall \mathbf{x}_1 \exists \mathbf{x}_2 \forall \mathbf{x}_3 \exists \mathbf{x}_4 \dots \forall \mathbf{x}_{2n+1} \exists \mathbf{x}_{2n+2} Q(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{2n+2})$$

(for some $n \geq 1$), where Q is a conjunction of positive atoms.

- Question: does $\mathcal{T} \models \psi$?

3-QSAT is easily recast as a QCSP, but we will be more interested in a variant problem. Let $\mathcal{B}_{\text{NAE}}^3$ be the boolean structure with a single ternary not-all-equal relation

$$\text{NAE}^3 = \{(0,0,1), (0,1,0), (1,0,0), (1,1,0), (1,0,1), (0,1,1)\}$$

The problem $\text{QCSP}(\mathcal{B}_{\text{NAE}}^3)$ is known to be Pspace-complete [38].

Much effort has gone into identifying the \mathcal{T} for which $\text{CSP}(\mathcal{T})$ is tractable (e.g. [32, 34]) and NP-complete (e.g. [33]). It has been conjectured in [19] that $\text{CSP}(\mathcal{T})$ is always either tractable, or NP-complete. (Indeed, it has even been conjectured in [8] where this separation lies.) However, the grand classification

into dichotomy remains incomplete. Some partial results are known: many years ago Schaefer proved the dichotomy for \mathcal{T} ranging over boolean domains [38]. That was recently extended to domains of size 3, through methods of universal algebra, by Bulatov in [7]. Of greater interest to us is the dichotomy theorem for undirected, antireflexive graphs of Hell and Nešetřil. They prove in [25] that an undirected template \mathcal{T} gives rise to a $\text{CSP}(\mathcal{T})$ that is tractable, if \mathcal{T} is bipartite, and a $\text{CSP}(\mathcal{T})$ that is NP-complete otherwise. This dichotomy extends trivially to all undirected graphs, since templates with self-loops will give rise to a trivial CSP. Bang-Jenson, Hell and MacGillavray prove a similar dichotomy theorem for tournament templates in [4]. Specifically, they prove that $\text{CSP}(\mathcal{T})$ is tractable, if \mathcal{T} is a tournament with at most one (directed) cycle, and that $\text{CSP}(\mathcal{T})$ is NP-complete, if \mathcal{T} is any other tournament. Both of these graph dichotomy results are proved by non-constructive means.

Following on from Schaefer's work [38], Dalmau [15] and Creignou et al. [14] eventually proved a dichotomy (tractable or Pspace-complete) for QCSP on boolean domains. A trichotomy (tractable, NP-complete or Pspace-complete) has been proved for QCSP on templates where all graphs of permutations appear as relations [5]. A significant body of tractability results has been established for QCSP, largely along the same lines as for CSP, in [5, 11]. However, so far, no overarching polychotomy for QCSP has been conjectured.

It is well known from work by Chandra and Merlin [9], on the problem of *Conjunctive Query Containment* from database theory, that (existential positive) conjunctive queries are directly related to the existence of homomorphism between structures. Defining constraint satisfaction problems in terms of structure homomorphism became popular after the seminal paper by Feder and Vardi [19]. The *non-uniform homomorphism problem* with template \mathcal{T} , denoted $\text{HOM}(\mathcal{T})$, is the decision problem with:

- Input: a structure \mathcal{A} .
- Question: does $\mathcal{A} \xrightarrow{h} \mathcal{T}$?

If \mathcal{K}_3 is the 3-clique, then it is clear that $\text{HOM}(\mathcal{K}_3)$ is the problem of graph 3-colourability.

We define the *canonical query* $\varphi_{\mathcal{A}}$ associated with \mathcal{A} to be the existential quantification of the conjunction of the facts of \mathcal{A} . For example, \mathcal{K}_3 has the canonical query

$$\varphi_{\mathcal{K}_3} := \exists x \exists y \exists z E(x, y) \wedge E(y, x) \wedge E(y, z) \wedge E(z, y) \wedge E(z, x) \wedge E(x, z).$$

$\text{HOM}(\mathcal{T})$ and $\text{CSP}(\mathcal{T})$ are essentially two views of the same problem. Specifically, they are equivalent under the bijective (up to structural isomorphism and labelling of variables) reduction $r(\mathcal{A}) = \varphi_{\mathcal{A}}$, i.e. $\mathcal{A} \in \text{HOM}(\mathcal{T})$ iff $\varphi_{\mathcal{A}} \in \text{CSP}(\mathcal{T})$.

In this chapter, we introduce a new problem $\text{ALT-HOM}(\mathcal{T})$, which is to $\text{QCSP}(\mathcal{T})$ what $\text{HOM}(\mathcal{T})$ is to $\text{CSP}(\mathcal{T})$. It is defined in terms of *alternating-homomorphism* from a partitioned structure to a non-partitioned template. We also give a characterisation of this problem through the existence of winning strategies in a certain game. Such a method has been used independently by Chen (e.g. in [11, 10, 12]).

7.2 Preliminaries

7.2.1 Structures and Logic.

We consider only finite, non-empty structures. Let \mathcal{A} and \mathcal{T} be such structures over σ . We denote the universe, or domain, of \mathcal{A} by $|\mathcal{A}|$, and the cardinality of $|\mathcal{A}|$ by $||\mathcal{A}||$. For each relation R_i of σ , with arity a_i , $R_i^{\mathcal{A}} \subseteq |\mathcal{A}|^{a_i}$ is the interpretation of R_i over \mathcal{A} . When it does not lead to confusion we may be sloppy in identifying R_i and $R_i^{\mathcal{A}}$. A structure \mathcal{A} is *connected* if, and only if, it is *not* the disjoint union of some structures \mathcal{A}' and \mathcal{A}'' . An *isolated element* of a structure \mathcal{A} is one that does not appear in any tuple of any relation of \mathcal{A} .

A homomorphism from \mathcal{A} to \mathcal{T} is a function $h : |\mathcal{A}| \rightarrow |\mathcal{T}|$ such that, for all relations R_i of σ , with arity a_i , and for all $(x_1, \dots, x_{a_i}) \in |\mathcal{A}|^{a_i}$, we have that $R_i^{\mathcal{A}}(x_1, \dots, x_{a_i})$ implies $R_i^{\mathcal{T}}(h(x_1), \dots, h(x_{a_i}))$. If there exists a homomorphism from \mathcal{A} to \mathcal{T} , then we write $\mathcal{A} \xrightarrow{h} \mathcal{T}$. If we have both $\mathcal{A} \xrightarrow{h} \mathcal{T}$ and $\mathcal{T} \xrightarrow{h} \mathcal{A}$ then we describe \mathcal{A} and \mathcal{T} as homomorphically equivalent.

A quantifier-free first-order formula Q is *positive conjunctive* if it is a conjunc-

tion of positive atoms, *i.e.* of the form,

$$Q(\mathbf{x}) = R_{i_1}(\mathbf{x}_1) \wedge R_{i_2}(\mathbf{x}_2) \wedge \dots \wedge R_{i_n}(\mathbf{x}_n),$$

where, for every $1 \leq j \leq n$, R_{i_j} is a relational symbol from σ , and \mathbf{x}_j is a tuple of variables of suitable length (*i.e.*, of the same length as the arity of R_{i_j}). Note that a variable may occur more than once in a given tuple.

7.2.2 Alternating-homomorphism problems.

For $n \in \mathbb{N}$, let $\kappa_n = \{U_1, E_2, U_3, E_4, \dots, U_{2n+1}, E_{2n+2}\}$ be a set of unary symbols that do not occur in σ . Define an *n-partitioned structure* \mathfrak{A} over σ to be a finite structure over the signature $\sigma \cup \kappa_n$, such that the interpretation of the symbols from κ_n is a partition of the structure: *i.e.*,

- $|\mathfrak{A}| = \bigcup_{i=0}^n (|U_{2i+1}| \cup |E_{2i+2}|)$; and,
- for any $0 \leq i < j \leq n$, the sets U_{2i+1} , E_{2i+2} , U_{2j+1} and E_{2j+2} are pairwise disjoint.

We write $\mathcal{S}_{\mathfrak{A}}$ to denote the σ -structure underlying \mathfrak{A} . We write $\mathfrak{A}|_{U_i}$ (respectively, $\mathfrak{A}|_{E_i}$) to denote the substructure of $\mathcal{S}_{\mathfrak{A}}$ induced by U_i (respectively, E_i). When this does not cause confusion, we write U_i (respectively, E_i) for the sake of brevity.

We say there is an *alternating-homomorphism* from the *n-partitioned structure* \mathfrak{A} over σ to the (non-partitioned) σ -structure \mathcal{T} , and we write $\mathfrak{A} \xrightarrow{alt} \mathcal{T}$ if, and only if,

- for all functions $f_{U_1} : U_1 \rightarrow |\mathcal{T}|$,
- there exists a function $f_{E_2} : E_2 \rightarrow |\mathcal{T}|$, such that,
- \vdots
- for all functions $f_{U_{2n+1}} : U_{2n+1} \rightarrow |\mathcal{T}|$,
- there exists a function $f_{E_{2n+2}} : E_{2n+2} \rightarrow |\mathcal{T}|$, such that,
- $f_{U_1} \cup f_{E_2} \cup \dots \cup f_{U_{2n+1}} \cup f_{E_{2n+2}}$ is a homomorphism from $\mathcal{S}_{\mathfrak{A}}$ to \mathcal{T} .

A *partitioned structure* is one that is n -partitioned, for some n .

Definition (Alternating-homomorphism problem). The *non-uniform alternating-homomorphism problem* with template \mathcal{T} , denoted by $\text{ALT-HOM}(\mathcal{T})$, is the decision problem with:

- Input: a partitioned structure \mathfrak{P} .
- Question: does $\mathfrak{P} \xrightarrow{\text{alt}} \mathcal{T}$?

Examples. Consider the graph \mathcal{G} with vertices $\{a, b, c, d\}$ and edge set $\{(a, b), (b, a), (c, d), (d, c)\}$. We define three partitioned structures which have \mathcal{G} as their underlying graph:

- \mathfrak{P}_1 such that $U_1 = \{a\}, E_2 = \{b\}, U_3 = \{c\}, E_4 = \{d\}$.
- \mathfrak{P}_2 such that $U_1 = \{a\}, E_2 = \{b\}, U_1 = \{c\}, E_2 = \{d\}$.
- \mathfrak{P}_3 such that $U_1 = \{a\}, E_8 = \{b\}, U_1 = \{c\}, E_{10} = \{d\}$.

These partitioned structures are depicted in Figure 7.1.

The above partitioned structures are equivalent in the sense that, for any structure \mathcal{T} , if for any one of them there exists an alternating-homomorphism to \mathcal{T} , then there exists also an alternating-homomorphism to \mathcal{T} from the others. This leads us to define the following *rewrite* scheme to transform an n -partitioned structure \mathfrak{P} to a *rewrite-reduced* partitioned structure, denoted $\overline{\mathfrak{P}}$.

1. If all relations of $\mathcal{S}_{\mathfrak{P}}$ are empty, i.e. all elements of $\mathcal{S}_{\mathfrak{P}}$ are isolated, then set $\overline{\mathfrak{P}}$ to be the singleton with $|\mathcal{S}_{\overline{\mathfrak{P}}}| = \{0\}$, all relations of $\mathcal{S}_{\overline{\mathfrak{P}}}$ empty, and $E_2 = \{0\}$. *Otherwise:*
2. Remove all isolated elements of \mathfrak{P} .
3. Suppose \mathfrak{P} is the disjoint union of m connected substructures $\mathfrak{P}_1, \dots, \mathfrak{P}_m$. For each $1 \leq l \leq m$, construct $\overline{\mathfrak{P}}_l$ from \mathfrak{P}_l thus:
 - (a) while there is a minimal $i \leq n$ such that U_{2i+1} is empty, move every element of E_{2j+2} into E_{2j} , for all $i \leq j \leq n$ and every element of U_{2j+3} into U_{2j+1} , for all $i \leq j \leq n - 1$.

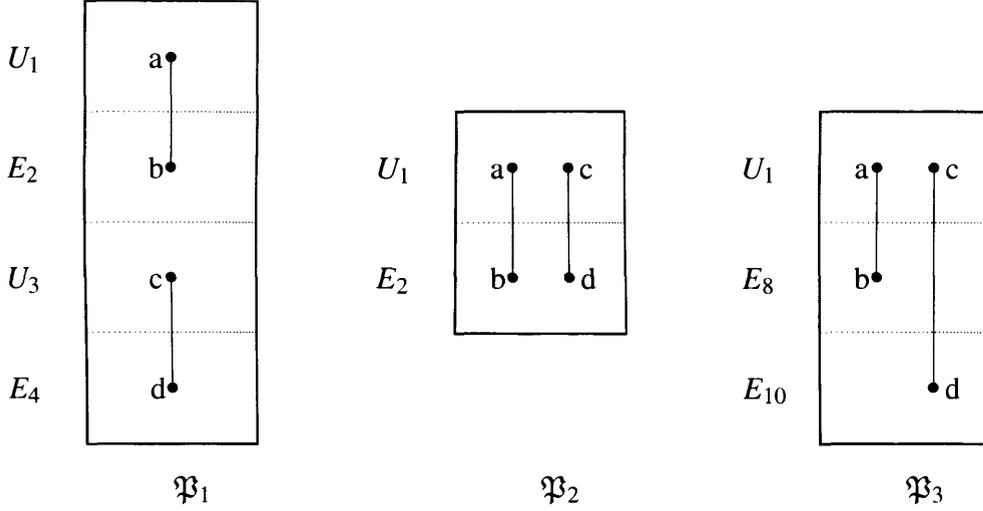


Figure 7.1: Three Partitioned Graphs.

(b) while there is a minimal $i \leq n$ such that E_{2i} is empty, move every element of E_{2j+2} into E_{2j} , for all $i \leq j \leq n$ and every element of U_{2j+1} into U_{2j-1} , for all $i \leq j \leq n$.

4. Set $\bar{\mathfrak{P}}$ to be $\bar{\mathfrak{P}}_1 \uplus \dots \uplus \bar{\mathfrak{P}}_m$.
5. Remove as many empty partitions as possible, so obtaining an n' -partitioned structure (for some $n' \leq n$).

The rewrite scheme is deterministic, up to the order in which the connected sub-structures are considered, and so $\bar{\mathfrak{P}}$ is well-defined. We say that two partitioned structures \mathfrak{P} and \mathfrak{P}' are *rewrite-equivalent* if $\bar{\mathfrak{P}} = \bar{\mathfrak{P}}'$. In figure 7.1, the structures \mathfrak{P}_1 , \mathfrak{P}_2 , and \mathfrak{P}_3 are all rewrite equivalent, and \mathfrak{P}_2 is rewrite-reduced. Note that, for any \mathfrak{P} , we can compute its rewrite-reduced $\bar{\mathfrak{P}}$ in polynomial time.

Two partitioned structures \mathfrak{P} and \mathfrak{P}' are said to be *problem-equivalent* if, for all templates \mathcal{T} , we have $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{T})$ iff $\mathfrak{P}' \in \text{ALT-HOM}(\mathcal{T})$.

Proposition 49. *Let \mathfrak{P} and \mathfrak{P}' be two partitioned structures. If \mathfrak{P} and \mathfrak{P}' are rewrite-equivalent then they are problem-equivalent.*

Proof. It is easy to see that the rewrite rules preserve the existence of alternating-homomorphism. \square

Note the converse does not hold as the following example shows.

Example. Consider the 1-partitioned digraphs:

- \mathfrak{P}_4 : with domain $\{x, y, z\}$, edge set $\{(x, y), (z, y)\}$, and partitions $U_1 = \{x\}$, $E_2 = \{y, z\}$.
- \mathfrak{P}_5 : with domain $\{x, y\}$, edge set $\{(x, y)\}$, and partitions $U_1 = \{x\}$, $E_2 = \{y\}$.

Whilst they are not rewrite equivalent, they are problem equivalent. \mathfrak{P}_4 is equivalent to the sentence $\forall x \exists y \exists z E(x, y) \wedge E(y, z)$ and \mathfrak{P}_5 is equivalent to the sentence $\forall x \exists y E(x, y)$. Both sentences have the same class of finite models.

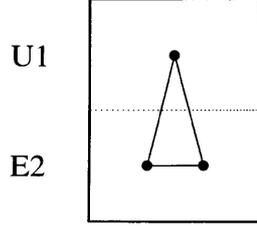
We note, for any \mathfrak{P} and its rewrite-reduced $\overline{\mathfrak{P}}$, that their underlying structures $\mathcal{S}_{\mathfrak{P}}$ and $\mathcal{S}_{\overline{\mathfrak{P}}}$ differ by possibly only some isolated elements. $\mathcal{S}_{\mathfrak{P}}$ and $\mathcal{S}_{\overline{\mathfrak{P}}}$ are homomorphically equivalent, and, for all templates \mathcal{T} , $\mathcal{S}_{\mathfrak{P}} \xrightarrow{h} \mathcal{T}$ iff $\mathcal{S}_{\overline{\mathfrak{P}}} \xrightarrow{h} \mathcal{T}$.

7.2.3 QCSP versus ALT-HOM.

In this section, we show that QCSP and ALT-HOM are essentially the same problem.

Theorem 50. *Let \mathcal{T} be a finite σ -structure. The problems QCSP(\mathcal{T}) and ALT-HOM(\mathcal{T}) are equivalent under logspace reduction.*

Proof. We will modify the bijective reduction $r(\mathcal{A}) = \varphi_{\mathcal{A}}$, mapping a structure to its canonical query, that proved the equivalence of HOM and CSP in the introduction to this chapter. From r we build the function s from partitioned structures to prenex quantified formulae whose quantifier-free part is positive conjunctive. Given a partitioned structure \mathfrak{P} , consider the canonical query $\varphi_{\mathcal{S}_{\mathfrak{P}}}$ of its underlying structure $\mathcal{S}_{\mathfrak{P}}$. Given this existential query $\varphi_{\mathcal{S}_{\mathfrak{P}}}$, we produce the query $\varphi_{\mathfrak{P}}$ by replacing all instances of $\exists x$, for variables x that correspond to elements of \mathfrak{P} in a universal partition, by $\forall x$. The map $s(\mathfrak{P}) = \varphi_{\mathfrak{P}}$ is bijective (up to isomorphism of the rewrite-reduced $\overline{\mathfrak{P}}$ and labelling of variables) and, along with its inverse, may



$$\forall x \exists y \exists z E(x, y) \wedge E(y, x) \wedge E(y, z) \wedge E(z, y) \wedge E(z, x) \wedge E(x, z)$$

Figure 7.2: Canonical sentence of a Partitioned Structure.

be computed in logarithmic space. It follows directly from the definitions that s and s^{-1} are reductions between $\text{QCSP}(\mathcal{T})$ and $\text{ALT-HOM}(\mathcal{T})$. \square

Just as we refer to the canonical query of a non-partitioned structure, so we will refer to the canonical query of a partitioned structure as being the sentence it reduces to, as in the previous theorem.

Example. Let \mathcal{T} be any graph. The partitioned structure and sentence of Figure 7.2 give rise to equivalent instances of, respectively, $\text{ALT-HOM}(\mathcal{T})$ and $\text{QCSP}(\mathcal{T})$, that reduce to one another in logarithmic space. The sentence is the canonical query of the partitioned structure, as just defined.

7.2.4 Alternating-homomorphisms as winning strategies.

We give a game characterisation of QCSP . The game we are about to define corresponds exactly to a standard model-checking game, also known as a Hintikka game [20]. We define this game in order to use the game parlance in subsequent proofs.

Definition (Game for QCSP). Let \mathfrak{P} be an n -partitioned structure and \mathcal{T} a (non-partitioned) template. The $(\mathfrak{P}, \mathcal{T})$ -game goes as follows. Opponent plays on the universal partitions and Proponent plays on the existential partitions. They play

alternate partitions, in ascending order, until all the partitions have been played.

For $0 \leq i \leq n$:

- Opponent (U -move): for every element in partition U_{2i+1} , Opponent chooses an element in \mathcal{T} , that is, Opponent gives a function $opp_{2i+1} : U_{2i+1} \rightarrow |\mathcal{T}|$.
- Proponent (E -move): for every element in partition E_{2i+2} , Proponent chooses an element in \mathcal{T} , that is, Proponent gives a function $pro_{2i+2} : E_{2i+2} \rightarrow |\mathcal{T}|$.

If, at any stage of the game, the function defined by the union of the moves of both players, $opp_1 \cup pro_2 \cup opp_3 \cup \dots$, is not a partial homomorphism from $\mathfrak{S}_{\mathfrak{P}}$ to \mathcal{T} , then Opponent wins. Otherwise this finite game will finish with some homomorphism from $\mathfrak{S}_{\mathfrak{P}}$ to \mathcal{T} having been constructed, and Proponent wins. It is Proponent's aim to construct such a homomorphism, and it is Opponent's aim to stop her. (In deference to the conventions of Ehrenfeucht-Fraïssé games, Proponent is considered female, and Opponent male.) Note that, if at some point the partial function defined by the play can not be extended to a homomorphism no matter how either side plays, then Opponent must necessarily win the game.

A strategy specifies how Opponent or Proponent are to play, given what has been played before:

- (Proponent) A strategy for partition E_{2i+2} in the $(\mathfrak{P}, \mathcal{T})$ -game is a function $\sigma_{E_{2i+2}} : E_{2i+2} \times \prod_{\lambda < i} (E_{2\lambda+2} \times \mathcal{T}) \times (U_{2\lambda+1} \times \mathcal{T}) \rightarrow \mathcal{T}$.
- A strategy σ for Proponent in the $(\mathfrak{P}, \mathcal{T})$ -game is the union of her strategies for all the existential partitions, viz $\bigcup_{\lambda \leq n} \sigma_{E_{2\lambda+2}}$.
- (Opponent) A strategy for partition U_{2i+1} in the $(\mathfrak{P}, \mathcal{T})$ -game is a function $\tau_{U_{2i+1}} : U_{2i+1} \times \prod_{\lambda < i} (E_{2\lambda} \times \mathcal{T}) \times (U_{2\lambda+1} \times \mathcal{T}) \rightarrow \mathcal{T}$.
- A strategy τ for Opponent in the $(\mathfrak{P}, \mathcal{T})$ -game is the union of his strategies for all the universal partitions, viz $\bigcup_{\lambda \leq n} \tau_{U_{2\lambda+1}}$.

A *winning strategy* for Proponent is a strategy σ that beats all Opponent strategies τ .

Theorem 51. *Let \mathfrak{P} be a partitioned structure, and $\phi_{\mathfrak{P}}$ the corresponding canonical query. The following are equivalent.*

(i) $\mathfrak{A} \xrightarrow{alt} \mathcal{T}$.

(ii) *Proponent has a winning strategy in the $(\mathfrak{A}, \mathcal{T})$ -game.*

(iii) $\mathcal{T} \models \varphi_{\mathfrak{A}}$.

Proof. The equivalence of (i) and (iii) follows from Theorem 50. It is well known that Proponent has a winning strategy in the **FO**-model-checking game on $(\varphi_{\mathfrak{A}}, \mathcal{T})$ if, and only if, $\mathcal{T} \models \varphi_{\mathfrak{A}}$ [20]. The game we define is the model-checking game restricted to sentences in prenex form whose quantifier-free part is positive conjunctive. The equivalence of (ii) and (iii) follows. \square

7.2.5 Graphs

A digraph is a structure over the signature containing a single binary relation E . An undirected graph is one whose edge relation is symmetric.

Definitions.

(Cliques.) Let $n \geq 1$. Let \mathcal{K}_n be the (*antireflexive*) *n-clique*, that is the graph with vertices $\{0, 1, \dots, n-1\}$ such that *all distinct* vertices are adjacent. Let \mathcal{K}_n^{ref} be the *reflexive n-clique*, that is the graph with vertices $\{0, 1, \dots, n-1\}$ such that *all* vertices are adjacent (when $i = j$, we call the corresponding edge a *self-loop*).

(Paths.) Let \mathcal{P}_n be the undirected antireflexive *n-path*, *i.e.* with vertices $\{0, 1, \dots, n-1\}$ such that $E(i, j)$ iff $j = i + 1$ or $j = i - 1$. It follows that the 1-path \mathcal{P}_1 is \mathcal{K}_1 and the 2-path \mathcal{P}_2 is \mathcal{K}_2 .

(Cycles.) Let \mathcal{DC}_n be the directed antireflexive *n-cycle*, *i.e.* with vertices $\{0, 1, \dots, n-1\}$ such that $E(i, j)$ iff $j = i + 1 \pmod n$. Let \mathcal{C}_n be the undirected antireflexive *n-cycle*, *i.e.* the symmetric closure of \mathcal{DC}_n . It follows that \mathcal{C}_3 is the 3-clique \mathcal{K}_3 .

It is proved in [5] that, for $n \geq 3$, $\text{QCSP}(\mathcal{K}_n)$ is Pspace-complete. It follows immediately that $\text{ALT-HOM}(\mathcal{K}_n)$ is also Pspace-complete, for $n \geq 3$.

An induced sub-digraph $\mathcal{G}' \subseteq \mathcal{G}$ is a *retract* of \mathcal{G} iff there is a homomorphism $h : \mathcal{G} \rightarrow \mathcal{G}'$ s.t. \mathcal{G}' is the image of h . A graph \mathcal{G}' is a *core* if it contains no proper retracts. For an arbitrary digraph \mathcal{G} , we define a *core of \mathcal{G}* to be any minimal (w.r.t.

size) retract that is itself a core. It is well-documented (e.g. [26]) that the core of a digraph is unique, up to isomorphism.

A bipartite graph is an undirected graph that is 2-colourable. A graph is bipartite iff it has either \mathcal{K}_1 or \mathcal{K}_2 as its core. Note that bipartite graphs are antireflexive.

7.3 Basic graph results

Most of these results are given for digraphs; we will specifically consider undirected graphs in the next section. When we discuss edges *between* vertices x and y , these may be oriented either way, or not at all (double edges).

7.3.1 Restricting partitions

We define restrictions on the input partitioned structure which will ultimately lead to tractability.

Definitions (restricting partitions). Let \mathfrak{P} be a partitioned digraph. We say that \mathfrak{P} is in Σ_1 -form (respectively, Π_2 -form), if the only non-empty partition is E_2 (respectively, if the only non-empty partitions are among $\{U_1, E_2\}$). If \mathfrak{P} is in Π_2 -form and there is at most one vertex in U_1 , then we say that \mathfrak{P} is in Π_2 -fan form. If, moreover, the vertex $x \in U_1$ exists, and is adjacent to some vertex $y \in E_2$, then we say that \mathfrak{P} is in *strict* Π_2 -fan form. Finally, we say that \mathfrak{P} is in Π_2 -multifan form, if \mathfrak{P} is the finite disjoint union of structures in Π_2 -fan form.

Note that, if \mathfrak{P} is in Σ_1 -form, then \mathfrak{P} is *a fortiori* in Π_2 -fan form. Any \mathfrak{P} in Π_2 -fan form, but not in strict Π_2 -fan form, has a rewrite-reduced $\overline{\mathfrak{P}}$ in Σ_1 -form.

Proposition 52 (Π_2 -multifan form). *Let \mathcal{T} be a digraph. The restriction of $\text{ALT-HOM}(\mathcal{T})$ to inputs in Π_2 -multifan form is NP-complete, whenever $\text{HOM}(\mathcal{T})$ is NP-complete.*

Proof. Let \mathfrak{P} be the disjoint union of $\mathfrak{P}_1, \dots, \mathfrak{P}_m$ all in Π_2 -fan form. Note that $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{T})$ iff $\mathfrak{P}_i \in \text{ALT-HOM}(\mathcal{T})$, for $1 \leq i \leq m$.

(Membership of NP) For each $1 \leq i \leq m$, if \mathfrak{P}_i is not in strict Π_2 -fan form, then it is equivalent to $\overline{\mathfrak{P}_i}$ in Σ_1 -form, and we may simply guess a homomorphism

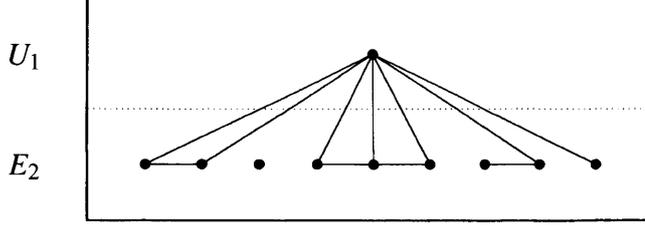


Figure 7.3: A Partitioned Graph in strict Π_2 -fan form.

and verify in polynomial time. If \mathfrak{P}_i is in strict Π_2 -fan form, we test all possible maps for the single element in U_1 , guessing the rest of the homomorphism and verifying in polynomial time.

(NP-hardness) If \mathfrak{P} is in Σ_1 -form, then $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{T})$ if, and only if, $\mathfrak{S}_{\mathfrak{P}} \in \text{HOM}(\mathcal{T})$. Hence, $\text{ALT-HOM}(\mathcal{T})$ is NP-hard provided that $\text{HOM}(\mathcal{T})$ is NP-complete. □

We will find, for a wide range of templates \mathcal{T} , that every input which is not in Π_2 -multifan form can be discarded. This will be because inputs \mathfrak{P} not in Π_2 -multifan form are either easily seen to be no-instances of $\text{ALT-HOM}(\mathcal{T})$, or to be equivalent to the rewrite-reduced $\overline{\mathfrak{P}}$ which is in Π_2 -multifan form. Further, we will find that we can split up inputs $\overline{\mathfrak{P}}$ in Π_2 -multifan form into their constituent Π_2 -fan components (as in the previous proof). Thus, structures in Π_2 -fan form are central to our discourse. Such a structure appears in Figure 7.3.

Remark. The ‘converse’ of Theorem 52, that the restriction of $\text{ALT-HOM}(\mathcal{T})$ to Π_2 -multifan form being NP-complete implies $\text{HOM}(\mathcal{T})$ is NP-complete, does not in general hold. For example, take \mathcal{T} to be (the disjoint union) $\mathcal{K}_3 \uplus \mathcal{K}_1^{ref}$. The self-loop \mathcal{K}_1^{ref} makes $\text{HOM}(\mathcal{K}_3 \uplus \mathcal{K}_1^{ref})$ trivial (every instance is a yes-instance). However:

Proposition 53. *The restriction of $\text{ALT-HOM}(\mathcal{K}_3 \uplus \mathcal{K}_1^{ref})$ to Π_2 -multifan form is NP-complete.*

Proof. Membership of NP follows as in the first part of the proof to the previous

proposition. For completeness, we give a reduction from the NP-complete 3-colourability problem $\text{HOM}(\mathcal{K}_3)$.

Let \mathcal{G} be an input for the problem $\text{HOM}(\mathcal{K}_3)$. Let $\mathcal{G}_1, \dots, \mathcal{G}_m$ be the connected components of \mathcal{G} and let x_1, \dots, x_m be some sequence of vertices in these respective components.

We construct an input for $\text{ALT-HOM}(\mathcal{K}_3 \uplus \mathcal{K}_1^{ref})$ thus:

- For each component \mathcal{G}_i , construct a partitioned graph \mathfrak{P}_i (in Π_2 -fan form) whose underlying graph has vertices $|\mathcal{G}_i| \uplus \{y_i\}$ (where y_i is a new vertex) and edge set $E^{\mathcal{G}_i} \uplus \{(x_i, y_i), (y_i, x_i)\}$, and whose partitions are $U_1 := \{y_i\}$ and $E_2 := |\mathcal{G}_i|$.
- Set \mathfrak{P} to be the disjoint union $\mathfrak{P}_1 \cup \dots \cup \mathfrak{P}_m$.

Clearly \mathfrak{P} is in Π_2 -multifan form: we claim that $\mathcal{G} \in \text{HOM}(\mathcal{K}_3)$ if, and only if, $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{K}_3 \uplus \mathcal{K}_1^{ref})$.

(forwards.) Suppose $\mathcal{G} \in \text{HOM}(\mathcal{K}_3)$. For each of the connected components \mathcal{G}_i there must be a homomorphism h_i to \mathcal{K}_3 . It suffices to show that there is a winning strategy for Proponent for each of the games $(\mathfrak{P}_i, \mathcal{K}_3 \uplus \mathcal{K}_1^{ref})$. If Opponent plays y_i to the self-loop of \mathcal{K}_1^{ref} , then Proponent may play the remainder of \mathfrak{P}_i to the same self-loop to win. If Opponent plays y_i to one of the vertices of the triangle \mathcal{K}_3 , then Proponent plays x_i to an adjacent vertex on the triangle, and may play the remainder of \mathfrak{P}_i according to [a cyclic permutation of] the homomorphism h_i to win.

(backwards.) Suppose $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{K}_3 \uplus \mathcal{K}_1^{ref})$: which implies that for each \mathfrak{P}_i there is a winning strategy for Proponent in the game on $(\mathfrak{P}_i, \mathcal{K}_3 \uplus \mathcal{K}_1^{ref})$. It suffices to show that this must imply the existence of a homomorphism from each \mathcal{G}_i to \mathcal{K}_3 . This is immediate, for suppose Opponent plays the y_i to some vertex in the triangle \mathcal{K}_3 , then the remainder of \mathfrak{P}_i must be played to the triangle, since \mathfrak{P}_i is connected, and so the winning strategy provides the necessary homomorphism. \square

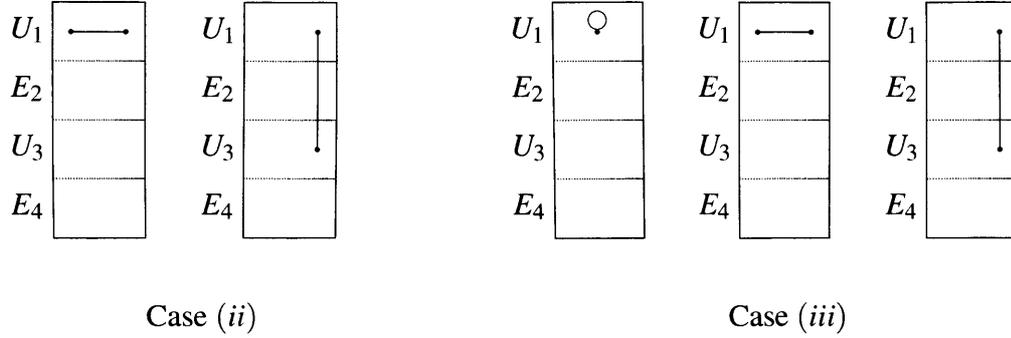


Figure 7.4: Types of forbidden edges in the last two cases of Proposition 55.

7.3.2 Basic results.

Proposition 54 (reflexive clique). *If \mathcal{T} is a reflexive clique, then $\text{ALT-HOM}(\mathcal{T})$ is trivial. Specifically: $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{T})$, for every \mathfrak{P} .*

Proof. If \mathcal{T} is a reflexive clique, then for any \mathfrak{P} , all functions from $\mathcal{S}_{\mathfrak{P}}$ to \mathcal{T} will be homomorphisms. □

Definition (e.g. [6]). A *dominating* vertex z in a digraph \mathcal{T} is one s.t. for all $w \in |\mathcal{T}|$, both $E^{\mathcal{T}}(w, z)$ and $E^{\mathcal{T}}(z, w)$ hold. (It follows that $(z, z) \in E^{\mathcal{T}}$.)

Proposition 55 (dominating vertex). *Let \mathcal{T} be a digraph in which there exists a dominating vertex z . Then $\text{ALT-HOM}(\mathcal{T})$ may be decided in logarithmic space.*

Proof. We consider three cases.

(i) If \mathcal{T} is a reflexive clique, then $\text{ALT-HOM}(\mathcal{T})$ is trivial (Proposition 54).

(ii) \mathcal{T} is not a reflexive clique, but is reflexive. In this case $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{T})$ iff \mathfrak{P} has no edges between distinct vertices $x \in U_i$ and $y \in U_j$ (for any i, j , see Figure 7.4). This property can clearly be checked in logarithmic space; we prove its correctness. Let a and b be distinct vertices of \mathcal{T} s.t. $\neg E^{\mathcal{T}}(a, b)$.

(\Rightarrow) By contraposition: if \mathfrak{P} has an edge between distinct vertices $x \in U_i$ and $y \in U_j$, then Opponent may play x on a and y on b to win, proving $\mathfrak{P} \notin \text{ALT-HOM}(\mathcal{T})$.

(\Leftarrow) If \mathfrak{P} has no edge between vertices $x \in U_i$ and $y \in U_j$, then Proponent may follow the strategy of playing all existential vertices to the dominating vertex z . This will overcome all Opponent strategies.

(iii) \mathcal{T} is not a reflexive clique, and is not reflexive. In this case $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{T})$ iff \mathfrak{P} has no edges between (not necessarily distinct) vertices $x \in U_i$ and $y \in U_j$ (for any i, j , see Figure 7.4). The proof proceeds as in part (ii), with all instances of the word ‘distinct’ dropped. \square

Lemma 56 (not a reflexive clique). *Let \mathcal{T} be a digraph that is not a reflexive clique, and let \mathfrak{P} be a partitioned digraph. If there is an edge in \mathcal{T} between distinct vertices $x \in U_i$ and $y \in U_j$ (for any i, j) then \mathfrak{P} is a no-instance of $\text{ALT-HOM}(\mathcal{T})$.*

Proof. By assumption, there are vertices a and b in \mathcal{T} that are not adjacent (a may be equal to b). Opponent plays a for x and b for y and wins. \square

Lemma 57 (antireflexivity). *Let \mathcal{T} be an antireflexive digraph. If there is an edge in \mathfrak{P} between nodes $x \in E_i$ and $y \in U_j$ (for $i < j$) then \mathfrak{P} is a no-instance of $\text{ALT-HOM}(\mathcal{T})$. If there is an edge in \mathfrak{P} between (not necessarily distinct) nodes $x \in U_i$ and $y \in U_j$ (for any i, j) then \mathfrak{P} is a no-instance of $\text{ALT-HOM}(\mathcal{T})$.*

Proof. In the first case, Proponent has chosen some vertex s in \mathcal{T} for x . Opponent also chooses s for y and wins. Similarly, in the second case, Opponent chooses the same vertex s for both x and y . \square

Lemma 58 (isolated vertex). *Let \mathcal{T} be a digraph with an isolated vertex s , and let \mathfrak{P} be a partitioned digraph. If there is an edge in \mathfrak{P} between $x \in U_i$ and $y \in E_j$ (for any i, j), or between $x \in U_i$ and $y \in U_j$ (any i, j), then \mathfrak{P} is a no-instance of $\text{ALT-HOM}(\mathcal{T})$.*

Proof. We prove the first case; the second may be done similarly. Regardless of what is played before, when Opponent plays x on s , there is no way that partial function can extend to homomorphism, regardless of where Proponent will map, or has mapped, y . \square

Proposition 59 (isolated vertex). *If \mathcal{T} is antireflexive and has an isolated vertex, then $\text{ALT-HOM}(\mathcal{T})$ and $\text{HOM}(\mathcal{T})$ are logspace equivalent.*

Proof. The reduction of $\text{HOM}(\mathcal{T})$ to $\text{ALT-HOM}(\mathcal{T})$ is trivial.

We reduce $\text{ALT-HOM}(\mathcal{T})$ to $\text{HOM}(\mathcal{T})$ as follows. Let \mathcal{N} be a fixed no-instance of $\text{HOM}(\mathcal{T})$ (say, \mathcal{T} augmented with one vertex adjacent to every vertex of \mathcal{T}). If \mathfrak{P} has an edge as in the previous lemma then we know that it is a no-instance and we reduce \mathfrak{P} to \mathcal{N} . If \mathfrak{P} has no such edge then every element in a universal partition is isolated. Thus, Opponent's moves have no bearing on Proponent's moves, and we may disregard every element occurring in a universal partition. Indeed, the rewritten-reduced graph $\overline{\mathfrak{P}}$ will be in Σ_1 -form. We reduce \mathfrak{P} to its underlying graph $\mathcal{S}_{\mathfrak{P}}$. \square

It is important that \mathcal{T} be antireflexive, to guarantee the existence of an \mathcal{N} in the previous proof. The following proposition is a cousin of the previous.

Proposition 60. *If \mathcal{T} has an isolated vertex, then $\text{ALT-HOM}(\mathcal{T})$ and $\text{HOM}(\mathcal{T})$ are equivalent under logspace Turing reductions.*

Proof. Again, the reduction from $\text{HOM}(\mathcal{T})$ to $\text{ALT-HOM}(\mathcal{T})$ is trivial.

We give the reduction $\text{ALT-HOM}(\mathcal{T})$ to $\text{HOM}(\mathcal{T})$. If \mathfrak{P} has an edge as in the previous lemma then we reject the input. If \mathfrak{P} has no such edge then we reduce it to the underlying graph $\mathcal{S}_{\mathfrak{P}}$. \square

Example. The problem $\text{ALT-HOM}(\mathcal{K}_3 \uplus \mathcal{K}_1)$ is NP-complete. It is equivalent to $\text{HOM}(\mathcal{K}_3)$, which in turn is the well-known NP-complete problem 3-COL.

7.3.3 Non-connected templates.

Lemma 61 (forbidden paths and non-connected templates). *Let \mathcal{T} be a non-connected graph. If there is a path in \mathfrak{P} between any $x \in E_i$ and $y \in U_j$ (for $i < j$), then \mathfrak{P} is a no-instance of $\text{ALT-HOM}(\mathcal{T})$. If there is a path in \mathfrak{P} between any $x \in U_i$ and $y \in U_j$ (for any i, j), then \mathfrak{P} is a no-instance of $\text{ALT-HOM}(\mathcal{T})$.*

Proof. We prove the first case; the second may be done similarly. Wherever Proponent plays x , Opponent need only play y in another connected component to win. \square

Lemma 62 (forbidden paths yield Π_2 -multifan form). *If there is no path in a partitioned digraph \mathfrak{P} between any $x \in E_i$ and $y \in U_j$ (for $i < j$), or between any $x \in U_i$ and $y \in U_j$ (for any i, j), then the rewrite-reduced $\overline{\mathfrak{P}}$ will be in Π_2 -multifan form.*

Proof. Consider \mathfrak{P} with all isolated vertices removed, and split into disjoint connected components. It suffices to prove that each of these is in Π_2 -fan form.

For any such component \mathfrak{P}' , let $0 < i \leq n$ be the largest integer such that U_{2i+1} is non-empty. It follows that there is an $x \in U_{2i+1}$, and that all other elements of \mathfrak{P}' are in existential partitions of index at least $2i + 2$, for otherwise there would be a path that violates our assumptions. The rewrite rules may be applied to move x to U_1 and all other vertices in the component to E_2 . The result follows. \square

Note that we can determine in polynomial time whether or not a digraph \mathcal{T} has any of the paths of the previous lemma.

Lemma 63 (non-connected). *Let \mathcal{T} be a digraph that is not connected, then $\text{ALT-HOM}(\mathcal{T})$ is in NP.*

Proof. Let \mathfrak{P} be a partitioned input digraph. If there are any of the paths in \mathfrak{P} as in Lemma 61 then we may reject the instance. Otherwise, it follows from the previous lemma that the rewrite-reduced $\overline{\mathfrak{P}}$ is in Π_2 -multifan form, and we can use the algorithm of Theorem 52. \square

Proposition 64 (non-connected). *Let \mathcal{T} be a digraph that is not connected. If $\text{HOM}(\mathcal{T})$ is NP-complete then $\text{ALT-HOM}(\mathcal{T})$ is NP-complete.*

Proof. By Lemma 63, $\text{ALT-HOM}(\mathcal{T})$ is in NP. $\text{HOM}(\mathcal{T})$ reduces trivially to $\text{ALT-HOM}(\mathcal{T})$, and completeness follows. \square

Remark. As in the remark after Theorem 52, the ‘converse’ of the previous proposition is not in general true: *i.e.* there are \mathcal{T} such that $\text{HOM}(\mathcal{T})$ is tractable but $\text{ALT-HOM}(\mathcal{T})$ is not. For example, when $\mathcal{T} = \mathcal{K}_3 \uplus \mathcal{K}_1^{ref}$, $\text{ALT-HOM}(\mathcal{T})$ is readily seen to be NP-complete. (Membership follows from Lemma 63 and completeness follows from Proposition 53.)

7.4 Quantified \mathcal{H} -colouring

7.4.1 Bipartite templates.

Lemma 65 (forbidden paths and bipartite template). *Let \mathcal{H} be a bipartite graph. If there is a path in \mathfrak{P} between any $x \in E_i$ and $y \in U_j$ (for $i < j$), then \mathfrak{P} is a no-instance of $\text{ALT-HOM}(\mathcal{H})$. If there is a path in \mathfrak{P} between any $x \in U_i$ and $y \in U_j$ (for any i, j), then \mathfrak{P} is a no-instance of $\text{ALT-HOM}(\mathcal{H})$.*

Proof. We prove the first case; the second may be done similarly. Let a be any vertex in \mathcal{H} on which Proponent plays x . If a is an isolated vertex, then Opponent wins (cf. proof of Lemma 58). Assume that a is not isolated. If the path in \mathfrak{P} between x and y is of even length, then Opponent plays y on b , where b is adjacent to a . A winning strategy for Proponent would imply the existence of an odd cycle in \mathcal{H} . This contradicts the fact that \mathcal{H} is bipartite, thus it follows that Opponent wins. If the path in \mathfrak{P} is of odd length, then Opponent plays y on a and wins by the same argument. \square

Lemma 66 (Π_2 -multifan form and bipartite). *Let \mathcal{H} be a bipartite graph. If \mathcal{H} has no isolated vertices then, for any \mathfrak{P} in Π_2 -multifan form, the following are equivalent:*

$$(i) \mathfrak{P} \xrightarrow{\text{alt}} \mathcal{H}$$

$$(ii) \mathfrak{P} \xrightarrow{\text{alt}} \mathcal{K}_2$$

$$(iii) \mathcal{S}_{\mathfrak{P}} \xrightarrow{h} \mathcal{K}_2$$

Proof. If \mathfrak{P} is the disjoint union of $\mathfrak{P}_1, \dots, \mathfrak{P}_m$ all in Π_2 -fan form, recall that $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{T})$ iff $\mathfrak{P}_i \in \text{ALT-HOM}(\mathcal{T})$, for $1 \leq i \leq m$.

For each \mathfrak{P}_i in Π_2 -fan form. When \mathfrak{P}_i has no vertex in U_1 , the result holds trivially. Otherwise, let x be the unique vertex in U_1 . Again, if x is isolated then \mathfrak{P} can be rewrite-reduced to $\overline{\mathfrak{P}}_i$ in Σ_1 -form. So, assume that \mathfrak{P}_i is in strict Π_2 -fan form, and that x is adjacent to some y in E_2 .

- (i) \Rightarrow (ii): Given a winning strategy σ for Proponent in the $(\mathfrak{P}_i, \mathcal{H})$ -game, we construct a winning strategy for Proponent in the $(\mathfrak{P}_i, \mathcal{K}_2)$ -game. Sup-

pose, w.l.o.g., that Opponent plays the x on the 1 in \mathcal{K}_2 . All remaining moves are Proponent's. So, Proponent chooses any homomorphism $h : \mathcal{H} \rightarrow \mathcal{K}_2$, and a vertex a in \mathcal{H} such that $h(a) = 1$. She then plays the rest of the vertices (all in E_2) according to the strategy $h \circ \sigma$ (where she assumes Opponent played the x to a in the oracle-game on $(\mathfrak{P}_i, \mathcal{H})$). Since h is a homomorphism, any outcome of the game on $(\mathfrak{P}_i, \mathcal{H})$ under strategy σ that is a homomorphism will lead to an outcome of the game on $(\mathfrak{P}_i, \mathcal{K}_2)$ under strategy $h \circ \sigma$ that is a homomorphism. We know that, under strategy σ , all outcomes of $(\mathfrak{P}_i, \mathcal{H})$ are homomorphisms, so the result follows.

- (ii) \Rightarrow (i): Given a winning strategy σ for Proponent in the $(\mathfrak{P}_i, \mathcal{K}_2)$ -game, we construct a winning strategy for Proponent in the $(\mathfrak{P}_i, \mathcal{H})$ -game. Suppose Opponent plays the x on a vertex a in \mathcal{H} . We know that a is not isolated, and has a distinct neighbour b . Let $h' : \mathcal{K}_2 \rightarrow \mathcal{H}$ be the homomorphism $\{(1, a), (2, b)\}$. All remaining moves are Proponent's. Proponent now plays the rest of the vertices (all in E_2) according to the strategy $h' \circ \sigma$ (where she assumes Opponent played the x to 1 in the oracle-game on $(\mathfrak{P}_i, \mathcal{K}_2)$). The argument concludes as before.
- (ii) \Leftrightarrow (iii): Since each \mathfrak{P}_i is in strict Π_2 -fan form, the result follows immediately from the symmetry of \mathcal{K}_2 .

□

Theorem 67 (bipartite). *Let \mathcal{H} be a bipartite graph. The problem $\text{ALT-HOM}(\mathcal{H})$ is tractable.*

Proof. We propose the following algorithm to solve $\text{ALT-HOM}(\mathcal{H})$.

The input \mathfrak{P} is first scanned to check whether it has any of the forbidden paths of Lemma 65. If there are any, then the input is rejected.

If there are none of the forbidden paths and \mathcal{H} has an isolated vertex, then we evaluate $\text{HOM}(\mathcal{H})$ on input $\mathfrak{S}_{\mathfrak{P}}$. That this is correct follows from Proposition 59; that it is tractable follows from Hell and Nešetřil's dichotomy theorem [25].

Otherwise, if there are none of the forbidden paths and \mathcal{H} has no isolated vertex, then we check whether $\mathfrak{S}_{\mathfrak{P}}$ is 2-colorable, and answer accordingly. This

is clearly polynomial: we prove its correctness. We know that if \mathfrak{P} has none of the forbidden paths, then it is rewrite-equivalent to the reduced $\overline{\mathfrak{P}}$ in Π_2 -multifan form, by Lemma 62. In particular, $\overline{\mathfrak{P}} \xrightarrow{alt} \mathcal{H}$ if, and only if, $\mathfrak{P} \xrightarrow{alt} \mathcal{H}$. By Lemma 66, we know that $\overline{\mathfrak{P}} \xrightarrow{alt} \mathcal{H}$ if, and only if, $\mathcal{S}_{\overline{\mathfrak{P}}} \xrightarrow{h} \mathcal{K}_2$. Moreover, by definition of $\overline{\mathfrak{P}}$, $\mathcal{S}_{\overline{\mathfrak{P}}}$ is the same as $\mathcal{S}_{\mathfrak{P}}$ up to possibly some isolated vertices. Hence, $\mathcal{S}_{\overline{\mathfrak{P}}} \xrightarrow{h} \mathcal{K}_2$ if, and only if, $\mathcal{S}_{\mathfrak{P}} \xrightarrow{h} \mathcal{K}_2$. It follows that $\mathfrak{P} \xrightarrow{alt} \mathcal{H}$ if, and only if, $\mathcal{S}_{\mathfrak{P}} \xrightarrow{h} \mathcal{K}_2$. \square

7.4.2 Odd Catherine Wheels

We have already met the problem $\text{QCSP}(\mathcal{B}_{\text{NAE}}^3)$, known to be Pspace-complete. Let $\mathcal{B}_{\text{NAE}}^n$ be the boolean structure with single n -ary not-all-equal relation

$$\text{NAE}^n := \{0, 1\}^n - \{(0^n), (1^n)\}$$

For all $n \geq 3$, $\text{QCSP}(\mathcal{B}_{\text{NAE}}^n)$ is Pspace-complete, by a trivial reduction from $\text{QCSP}(\mathcal{B}_{\text{NAE}}^3)$.

Definition. We consider an undirected graph \mathcal{W} to be an *odd catherine wheel* (OCW) if it is isomorphic to some graph \mathcal{G} constructed as follows. For some k , take the $(2k+1)$ -cycle \mathcal{C}_{2k+1} , together with $(2k+1)$ undirected paths $\mathcal{P}^0, \dots, \mathcal{P}^{2k}$ (each of any finite length, where \mathcal{K}_1 is considered the 0-path). Construct \mathcal{G} by identifying an end of each path \mathcal{P}^i with vertex i of \mathcal{C}_{2k+1} .

As in that construction, an OCW may be given an ordering over its $(2k+1)$ -cycle, which we will call a *listing*. An OCW may have up to $2 \cdot (2k+1)$ distinct listings, corresponding to orientation of the cycle, together with position of the zero (first) vertex. We will usually refer to a listing by a corresponding sequence of paths.

Definition. For an undirected graph \mathcal{G} , and a subset $A \subseteq |\mathcal{G}|$ define:

- $d(x, y)$ to be the length of the shortest path in \mathcal{G} from x to y ,
- $d(A, y) = \min\{d(x, y) : x \in A\}$, and,
- $D(\{p, q\}) = \max\{d(\{p, q\}, y) : y \in |\mathcal{G}|\}$.

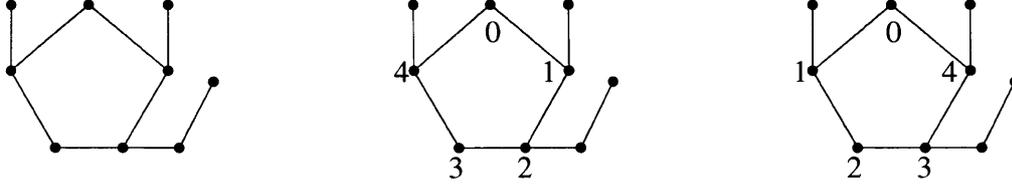


Figure 7.5: An OCW and its two D -minimal listings.

$D(\{p, q\})$ is minimal such that there is an m' -walk (for some $m' \leq D(\{p, q\})$) from $\{p, q\}$ to every vertex of \mathcal{G} . We will only be concerned with $D(\{p, q\})$ when p and q are adjacent vertices on the cycle of an OCW.

Definition. For any OCW \mathcal{W} , define:

$$m_{\mathcal{W}} = \min\{D(\{p, q\}) : p, q \text{ adjacent on the cycle of } \mathcal{W}\}$$

A D -minimal listing $\mathcal{P}^0, \dots, \mathcal{P}^{2k}$ of \mathcal{W} is one in which $D(\{k, k+1\}) = m_{\mathcal{W}}$.

A D -minimal listing is one in which the maximal distance from the vertices $\{k, k+1\}$ is minimised. These middle vertices k and $k+1$ will eventually play the role of TRUE and FALSE in a reduction from $\text{QCSP}(\mathcal{B}_{\text{NAE}}^{2k+1})$. It is the following property of D -minimal listings that is important.

Lemma 68. *Given a D -minimal listing $\mathcal{P}^0, \dots, \mathcal{P}^{2k}$ of an OCW \mathcal{W} , i.e. one in which $D(\{k, k+1\}) = m_{\mathcal{W}} = m$, there exists:*

- a $t \in |\mathcal{W}|$ s.t. there is an m -walk from t to k , but no m -walk from t to $k+1$,
and
- an $s \in |\mathcal{W}|$ s.t. there is an m -walk from s to $k+1$, but no m -walk from s to k .

Proof. Let the respective lengths of the paths \mathcal{P}^i be λ_i . It follows that:

$$\begin{aligned}
m &= D(\{k, k+1\}) \\
&= \max\{ && \lambda_0 + k, \\
&& \lambda_1 + [k-1], && \lambda_{2k} + [k-1], \\
&& \vdots && \vdots \\
&& \lambda_{k-1} + 1, && \lambda_{k+2} + 1, \\
&& \lambda_k, && \lambda_{k+1} && \}
\end{aligned}$$

We consider three cases.

(exists i , $1 \leq i \leq k$, s.t. $m = \lambda_i + [k-1]$.) Take such a branch i that has a vertex at maximal distance from $\{k, k+1\}$. This vertex will be at the end of the path \mathcal{P}^i . Label this vertex t , and its neighbour, the penultimate vertex along \mathcal{P}_i , s . (If the path \mathcal{P}^i is \mathcal{K}_1 , i.e. there is no path leaving the cycle, then consider $i+1$ on the cycle to be the ‘penultimate’ vertex). It follows that there will be an m -walk from t to k but not to $k+1$ (by maximality of m , together with the fact that t must be closer to k than $k+1$). It also follows that there will be an m -walk from s to $k+1$ but not k (we can not go the long way round the cycle, and any backstepping increases the walk by an even amount).

(exists i , $k+1 \leq i \leq 2k$, s.t. $m = \lambda_i + [i-k-1]$.) This case is symmetric to the previous.

(previous cases fail, and $m = \lambda_0 + k$.) This is the case in which the ultimate vertex of \mathcal{P}^0 is the *unique* vertex at maximal distance m from $\{k, k+1\}$. In this case, we make two claims:

- (i) there exists $1 \leq i \leq k$ s.t. $m-1 = \lambda_i + [k-1]$, and
- (ii) there exists $k+1 \leq i \leq 2k$ s.t. $m-1 = \lambda_i + [i-k-1]$.

For the first claim, if no such i exists, then we do not have a D -minimal listing, since $D(\{k+1, k+2\}) < D(\{k, k+1\})$. (If the vertex 0 is always considered at the top of a drawing of \mathcal{W} , then this represents rotating the wheel one place anticlockwise.) The proof of the second claim is symmetric. We may now take t and s to be the ultimate vertices on some paths that fulfill the the second and first criteria, respectively. The proof concludes as with the previous cases. \square

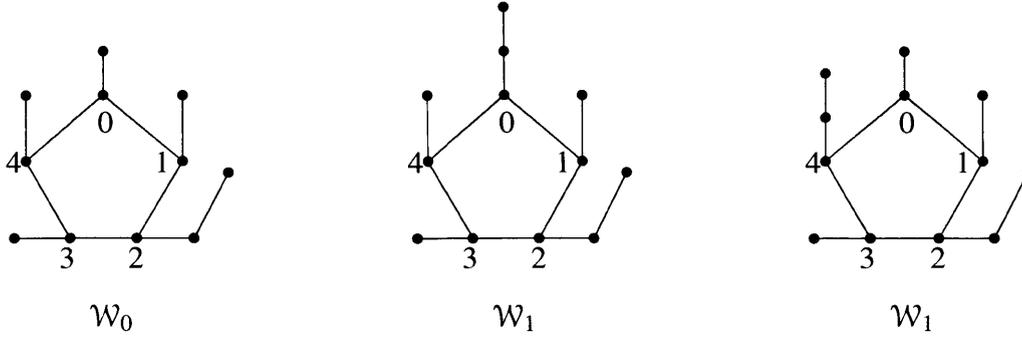


Figure 7.6: The listing on \mathcal{W}_0 is D -minimal, but only the rightmost listing of \mathcal{W}_1 is D -minimal.

Before we go on, we will need the following.

Lemma 69 (e.g. [25]). *For all vertices x of the $(2k + 1)$ -cycle \mathcal{C}_{2k+1} , there is no $(2k - 1)$ -walk from x to x , but there is a $(2k - 1)$ -walk from x to all distinct vertices y .*

Proof. That there is no walk from x to itself follows by a parity argument, together with the fact that the walk can not go round the entire cycle. We now construct a $(2k - 1)$ walk from vertex 0 to any vertex $1 \leq i \leq 2k$, whereafter we may appeal to symmetry. If i even, then walk backwards (anticlockwise) until i is attained, in an odd number of moves, and waste the even number of moves remaining walking between i and some neighbour. If i odd, then walk forwards (clockwise) until i is attained, in an odd number of moves, and waste the remaining moves as before. \square

Pspace-completeness.

Theorem 70. *If \mathcal{W} is an OCW, then $\text{ALT-HOM}(\mathcal{W})$ is Pspace-complete.*

Proof. Suppose \mathcal{W} has a $(2k + 1)$ -cycle, and let $m_{\mathcal{W}} = m$ be given. The proof is by direct reduction from $\text{QCSP}(\mathcal{B}_{\text{NAE}}^{2k+1})$. It is based on that given in [5]. Let φ be an instance of $\text{QCSP}(\mathcal{B}_{\text{NAE}}^{2k+1})$. Without loss of generality, we assume that φ has at least one universal quantifier: if there is none we can introduce a dummy¹.

¹This restriction is actually unnecessary in the reduction we use, but it saves us considering as a special case the situation where there are no universally quantified variables.

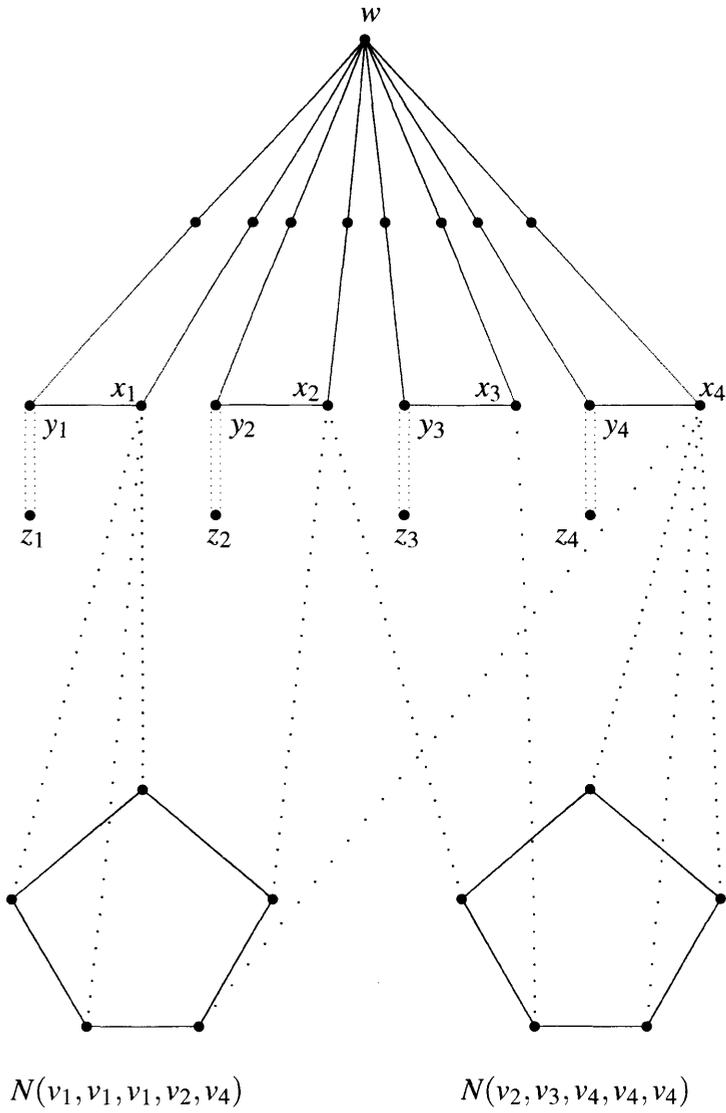


Figure 7.7: Underlying Graph in reduction from $\text{QCSP}(\mathcal{B}_{NAE}^{2k+1})$. The dotted lines are $(2k - 1)$ -paths; the double dotted lines are m -paths.

Suppose φ has ν variables and κ clauses: we will construct a partitioned graph \mathfrak{P} such that $\varphi \in \text{QCSP}(\mathcal{B}_{\text{NAE}}^{2k+1})$ iff $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{W})$.

To build the underlying graph $\mathcal{S}_{\mathfrak{P}}$, we first take ν copies of the $(2k+1)$ -gon \mathcal{C}_{2k+1} , one for each variable. Consider each of these $(2k+1)$ -gons \mathcal{C}_{2k+1}^i to have identified vertex w_i , and labelled twin vertices x_i, y_i farthest away in \mathcal{C}_{2k+1}^i from w_i (it is irrelevant which way round x_i and y_i are). Attach to each \mathcal{C}_{2k+1}^i an m -path from y_i , and label the end-vertex on this path z_i . Now identify the w_i s as a single vertex w . (A case involving pentagons, with $\nu = 4$, is shown in the top half of Figure 7.7.) We now take κ further copies of \mathcal{C}_{2k+1} , one for each clause. Each vertex of each of these $(2k+1)$ -gons represents a variable in the clause. For each variable v_i in such a clause, add a $(2k-1)$ -path from the vertex representing v_i to the x_i previously introduced. The case in Figure 7.7 has $\kappa = 2$, with clauses $(v_1, v_1, v_1, v_2, v_3)$ and $(v_2, v_3, v_4, v_4, v_4)$. It remains for us to partition the vertices of $\mathcal{S}_{\mathfrak{P}}$. There is nothing in partition U_1 , and w is on its own in partition E_2 . Now we read the quantifiers in φ , from the outside. For each existentially quantified variable v_i , we add z_i , its path, and all the rest of \mathcal{C}_{2k+1}^i , to the next strictly higher existential partition. For each universally quantified variable v_i , we add just z_i to the next available universal partition. We then add the rest of z_i 's path, and all the rest of \mathcal{C}_{2k+1}^i , to the next existential partition. When we have gone through all the quantifiers of φ , we add all of the remaining vertices, *i.e.* those in the clause $(2k+1)$ -gons, and in the paths that reach them, to the next available existential partition.

This construction is clearly polynomial. It remains for us to prove its correctness. Note that Proponent can not successfully play all the x_i associated with some clause $(2k+1)$ -gon to a single vertex of \mathcal{W} (if she plays off the cycle she clearly loses; if she plays on the cycle she loses by Lemma 69).

$(\varphi \in \text{QCSP}(\mathcal{B}_{\text{NAE}}^{2k+1}) \rightarrow \mathfrak{P} \in \text{ALT-HOM}(\mathcal{W}))$. We give Proponent's strategy in the game $(\mathfrak{P}, \mathcal{W})$. She should play w on some vertex 0 on the $(2k+1)$ -cycle of \mathcal{W} such that this gives rise to a D -minimal listing of \mathcal{W} . Whenever Opponent plays a z_i in \mathcal{W} , there will be an m -walk such that y_i may be played on one of vertices k or $k+1$ of the cycle of \mathcal{W} . These vertices represent TRUE and FALSE respectively. Since, no matter whether the universal variables are true or false, there is a valuation of the existential variables that gives the clauses a not-all-equal valuation,

Proponent may ensure that not all x_i associated with each clause are mapped to TRUE (respectively, FALSE). She should play this valuation, finally playing each clause $(2k + 1)$ -gon and the path to it according to Lemma 69, ensuring homomorphism.

$(\mathfrak{P} \in \text{ALT-HOM}(\mathcal{W}) \rightarrow \varphi \in \text{QCSP}(\mathcal{B}_{\text{NAE}}^{2k+1}))$. If w is not played to a vertex 0 in the $(2k + 1)$ -cycle of \mathcal{W} such that this gives a D -minimal listing, then Opponent may play any universal z_i (by assumption there is at least one) to some vertex in \mathcal{W} that does not have an m -walk to either k or $k + 1$, and Proponent loses. (Such a vertex exists by minimality of m .) Thus, in a winning strategy, w must be played to some 0 on the $(2k + 1)$ -cycle that gives rise to a D -minimal listing. Note that now Proponent must play each $\{x_i, y_i\}$ to $\{k, k + 1\}$, with which being played to which specifying truth or falsity of variable v_i . Thereafter, for any play of z_i , there is an m -walk to either k or $k + 1$ (TRUE or FALSE respectively), by the definition of m over a D -minimal listing. For certain z there is an m -walk to k but not to $k + 1$, and vice-versa, as guaranteed by Lemma 68. This ensures that Proponent must answer to all valuations of the universal variables. Finally, when the clause $(2k + 1)$ -gons are reached, if Proponent can extend to homomorphism, then not all the x_i s of each clause were played to k (respectively $k + 1$), and we have a not-all-equal assignment \square

Corollary. *Let \mathcal{G} be a (undirected antireflexive) connected graph that has a unique cycle, which is of odd length. Then $\text{ALT-HOM}(\mathcal{G})$ is Pspace-complete.*

Proof. Such a graph \mathcal{G} is of form similar to an OCW, but with trees affixed to the vertices of the odd cycle, instead of paths. The completeness result holds for such graphs under exactly the same reduction. (The length of such a tree should be considered the maximal depth from its root on the odd cycle.) \square

7.4.3 A trichotomy theorem

Theorem 71. *The class of antireflexive undirected graphs with at most one cycle exhibits ALT-HOM-trichotomy. Specifically:*

- *If \mathcal{H} is bipartite, then $\text{ALT-HOM}(\mathcal{H})$ is tractable.*

- If \mathcal{H} has an odd cycle and is not connected, then $\text{ALT-HOM}(\mathcal{H})$ is NP-complete.
- If \mathcal{H} has an odd cycle and is connected, then $\text{ALT-HOM}(\mathcal{H})$ is Pspace-complete.

Proof. We have just proved the final part. The first part is proved in Theorem 67. The second part is a consequence of Proposition 64, and Hell and Nešetřil's Dichotomy Theorem [25]. \square

Remark. The same trichotomy holds on the class of antireflexive undirected graphs with exactly one cycle. This is because bipartite graphs may contain even cycles.

7.5 Closure properties

We examine some closure properties on templates that may be used for proving Pspace-hardness. Later we look at the question of problem equivalence in QCSP.

7.5.1 Indicator construction

Hell and Nešetřil defined three graph constructions to prove their dichotomy theorem for undirected graphs in [25]. One of them is known as the *Indicator construction*. An Indicator is a digraph \mathcal{J} with two identified vertices i and j .

Definition (Hell and Nešetřil [25]). The indicator construction \mathcal{T}^* of a digraph \mathcal{J} with respect to Indicator (\mathcal{J}, i, j) is the graph with vertex set $|\mathcal{T}|$; and, edge set:

$$\{(a, b) : \text{exists hom. } h : \mathcal{J} \rightarrow \mathcal{T} \text{ s.t. } h(i) = a \text{ and } h(j) = b\}.$$

Remark. In the case of undirected graphs it makes sense to consider only Indicators (\mathcal{J}, i, j) that have an automorphism swapping i and j . This ensure that \mathcal{T}^* remains undirected when \mathcal{J} is undirected. For an example of this construction, see Figure 7.8.

Lemma 72 (Hell and Nešetřil [25]). *Let \mathcal{T} be an undirected graph and (\mathcal{J}, i, j) an indicator that has an automorphism swapping i and j . If $\text{HOM}(\mathcal{T}^*)$ is NP-complete then $\text{HOM}(\mathcal{T})$ is NP-complete.*

Their result readily extends to digraphs; we extend it to QCSP.

Theorem 73 (Indicator Construction). *Let \mathcal{G} be a digraph and (\mathcal{J}, i, j) an indicator.*

- *If $\text{CSP}(\mathcal{G}^*)$ is NP-complete then $\text{CSP}(\mathcal{G})$ is NP-complete.*
- *If $\text{QCSP}(\mathcal{G}^*)$ is Pspace-complete then $\text{QCSP}(\mathcal{G})$ is Pspace-complete.*

Proof. We prove the first claim by reducing $\text{CSP}(\mathcal{G}^*)$ to $\text{CSP}(\mathcal{G})$. The proof broadly follows that of Hell and Nešetřil. Let $m := \|\mathcal{J}\|$. Take the canonical query θ of \mathcal{J} and remove the two existential quantifiers for i and j . It is now of the form:

$$\theta(z, z') := \exists y_1 \exists y_2 \dots \exists y_{m-2} Q(z, z', y_1, y_2, \dots, y_{m-2}),$$

where Q is positive conjunctive. It follows directly from the definitions that $E^{\mathcal{G}^*}(a, b)$ holds if, and only if,

$$\mathcal{G} \models \theta(z/a, z'/b) = \exists y_1 \exists y_2 \dots \exists y_{m-2} Q(a, b, y_1, y_2, \dots, y_{m-2}).$$

Hence, given an instance of $\text{CSP}(\mathcal{G}^*)$, we can replace each occurrence of $E(z, z')$ by $\theta(z, z')$, ensuring that variables introduced are new variables. More precisely, let $z_{k_1}, \dots, z_{k_{2r}}$ be (not necessarily distinct) variables among \mathbf{z} , and let $\varphi := \exists \mathbf{z} \bigwedge_{i=1}^r E(z_{k_{2i-1}}, z_{k_{2i}})$ be an instance of $\text{CSP}(\mathcal{G}^*)$. We have $\mathcal{G}^* \models \varphi$ iff $\mathcal{G} \models \psi$, where

$$\psi := \exists \mathbf{z} \bigwedge_{i=1}^r (\exists y_1^i \exists y_2^i \dots \exists y_{m-2}^i) Q(z_{k_{2i-1}}, z_{k_{2i}}, y_1^i, y_2^i, \dots, y_{m-2}^i).$$

Note that ψ can be built from φ in polynomial time (remember that the indicator is a fixed graph). Both $\text{CSP}(\mathcal{G})$ and $\text{CSP}(\mathcal{G}^*)$ are in NP, so it follows immediately that NP-completeness of the latter implies NP-completeness of the former.

We use the same method to prove the second claim, by reducing $\text{QCSP}(\mathcal{G}^*)$ to $\text{QCSP}(\mathcal{G})$. We have seen that the edge relation of \mathcal{G}^* can be defined in existential

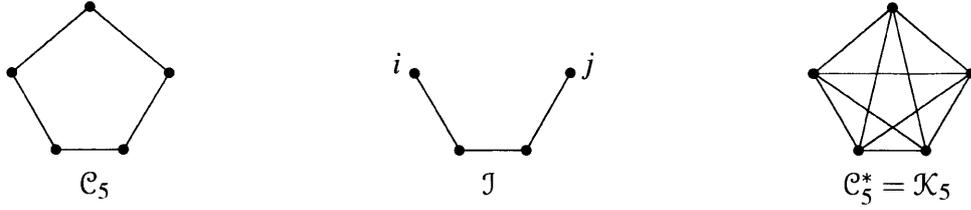


Figure 7.8: Example of the indicator construction.

positive conjunctive first-order logic on \mathcal{G} . Thus, by replacing in the same way each occurrence of $E(z, z')$ by $\theta(z, z')$, where every variable apart from z and z' is a new one, we get a quantified formula ψ . The proof concludes as before, but for Pspace instead of NP. \square

We now have an alternative proof, based on that in [25], of the following.

Corollary. *For every undirected cycle \mathcal{C}_{2k+1} ($k \geq 1$), the problem $\text{ALT-HOM}(\mathcal{C}_{2k+1})$ is Pspace-complete.*

Proof. Recall that $\text{ALT-HOM}(\mathcal{K}_{2k+1})$ is Pspace-complete. For $\mathcal{G} := \mathcal{C}_{2k+1}$, choose the indicator (\mathcal{J}, i, j) to be the undirected $(2k+1)$ -path from the vertex i to the vertex j . It follows from Lemma 69 that $\mathcal{G}^* = \mathcal{K}_{2k+1}$. The result follows from Theorem 73. The case $k = 2$ is depicted in Figure 7.8. \square

Remark. In [25], Hell and Nešetřil introduced two other graph constructions. When \mathcal{H} is a core, they defined the graphs \mathcal{H}^\sim [respectively, \mathcal{H}^\wedge] with respect to sub-indicator \mathcal{J} [respectively, edge-sub-indicator \mathcal{J}']. We do not go into the details of these constructions here. They proved that either of $\text{HOM}(\mathcal{H}^\sim)$ or $\text{HOM}(\mathcal{H}^\wedge)$ being NP-complete implies $\text{HOM}(\mathcal{H})$ is NP-complete. We note that this result is unlikely to extend to ALT-HOM (QCSP) in the case of the sub-indicator construction. That is, $\text{ALT-HOM}(\mathcal{H}^\sim)$ being Pspace-complete does not imply $\text{ALT-HOM}(\mathcal{H})$ is Pspace-complete, under the assumption that $\text{NP} \neq \text{Pspace}$. We do not prove this here.

7.5.2 Adding a vertex to a core.

Definition. Let \mathcal{H} be an antireflexive core (i.e. any core other than the self-loop \mathcal{K}_1^{ref}). Let \mathcal{H}^+ be \mathcal{H} with a new vertex γ added, adjacent to all vertices of \mathcal{H} , but not adjacent to itself (so it does not introduce a self-loop). Formally:

- $|\mathcal{H}^+| = |\mathcal{H}| \uplus \{\gamma\}$.
- $E^{\mathcal{H}^+} = E^{\mathcal{H}} \cup \{(x, \gamma), (\gamma, x) : x \in \mathcal{H}\}$.

We aim to establish that \mathcal{H}^+ is a core.

Lemma 74. *Any homomorphism $h : \mathcal{H}^+ \rightarrow \mathcal{H}^+$ is such that there is an automorphism i of \mathcal{H}^+ that swaps γ and $h(\gamma)$.*

Proof. We may assume $h(\gamma) \neq \gamma$. We prove that $h(\gamma)$ is [forward- and backward-]adjacent to all vertices of \mathcal{H}^+ except itself. If that is the case then the function that leaves all vertices unchanged, but swaps γ and $h(\gamma)$, will be an automorphism.

Suppose $h(\gamma)$ were not adjacent to everything in \mathcal{H}^+ (except itself), and that its neighbours constitute the proper subgraph $\mathcal{H}' \subset \mathcal{H}^+ - \{h(\gamma)\}$. Since h is a homomorphism, it follows that $h(\mathcal{H}) \subseteq \mathcal{H}'$, so we have:

$$h(\mathcal{H}) \subset \mathcal{H}^+ - \{h(\gamma)\}.$$

Now, γ itself may or may not be in the image $h(\mathcal{H})$. We consider both cases separately.

[$\gamma \notin h(\mathcal{H})$.] We have $h(\mathcal{H})$ is a proper subgraph of \mathcal{H} , and we are done since this contradicts \mathcal{H} being a core.

[$\gamma \in h(\mathcal{H})$.] See Figure 7.9. It follows that there is a homomorphism

$$h' : \mathcal{H} \longrightarrow h(\mathcal{H}) - \{\gamma\} \cup \{h(\gamma)\},$$

defined by $h'(x) := h(x)$ except when $h(x) = \gamma$, in which case $h'(x) := h(\gamma)$. Whilst $h(\mathcal{H})$ is not actually a subgraph of \mathcal{H} , $h'(\mathcal{H})$ is, and it is proper since it has the same cardinality as $h(\mathcal{H})$. This contradicts \mathcal{H} being a core, and we are done. \square

Lemma 75. \mathcal{H}^+ is a core.

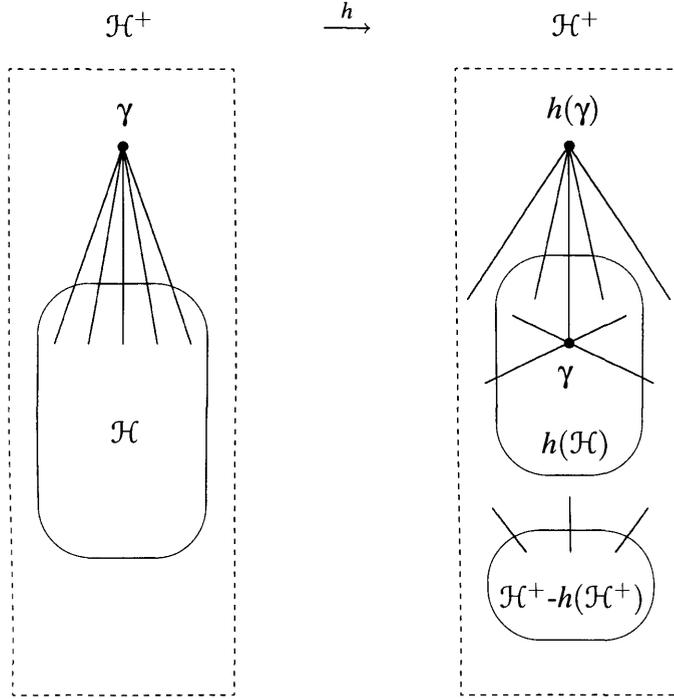


Figure 7.9: The case $\gamma \in h(\mathcal{H})$ in Lemma 74. Note that $\mathcal{H}^+ - h(\mathcal{H}^+)$ must be non-empty.

Proof. Suppose there were a homomorphism h from \mathcal{H}^+ to a proper subgraph $\mathcal{H}' \subset \mathcal{H}^+$. Since we know there is an automorphism of \mathcal{G}^+ that swaps γ and $h(\gamma)$, we may assume w.l.o.g. that $h(\gamma) = \gamma$. But that implies that h maps \mathcal{H} to a proper subgraph of itself, which contradicts \mathcal{H} being a core. \square

Theorem 76. *Let \mathcal{H} be a core. Then $\text{ALT-HOM}(\mathcal{H})$ is logspace reducible to $\text{ALT-HOM}(\mathcal{H}^+)$.*

Proof. Given an input \mathfrak{P} for $\text{ALT-HOM}(\mathcal{H})$, we construct \mathfrak{P}' as an input for $\text{ALT-HOM}(\mathcal{H}^+)$ such that $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{H})$ iff $\mathfrak{P}' \in \text{ALT-HOM}(\mathcal{H}^+)$. We construct \mathfrak{P}' from \mathfrak{P} by introducing a new existential partition E_0 before U_1 (we may renumber later). Into E_0 we place a copy of \mathcal{H}^+ , adding an edge from the γ of that \mathcal{H}^+ to all the existential vertices of \mathfrak{P} . Our proof rests on the equivalence of the following:

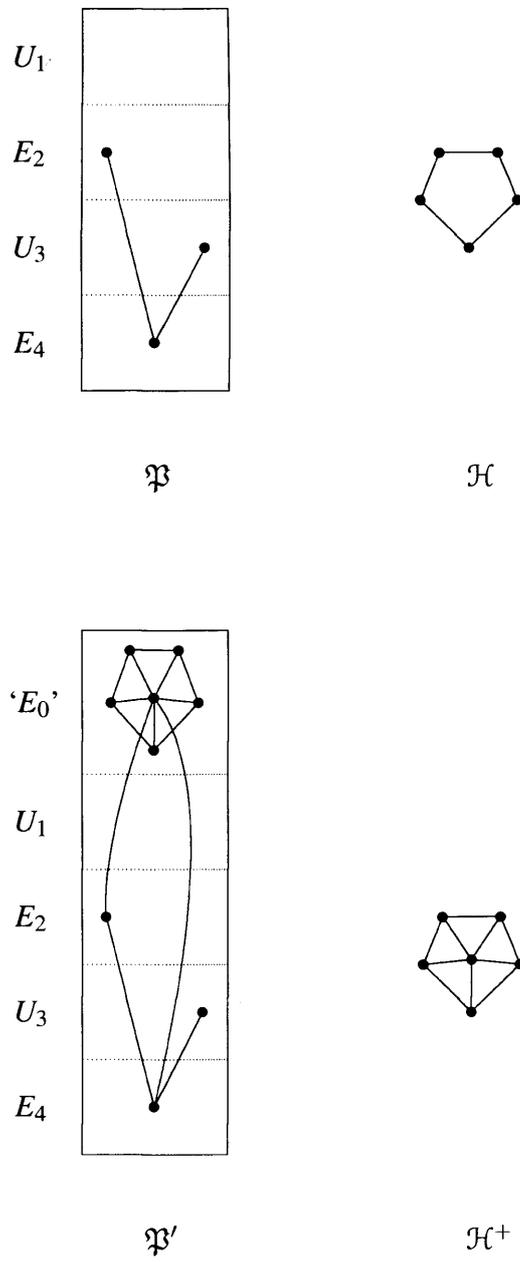


Figure 7.10: An example of the reduction used in Theorem 76, when $\mathcal{H} = \mathcal{C}_5$.

- (i) Proponent has a winning strategy in the game on $(\mathfrak{P}', \mathcal{H}^+)$.
- (ii) Proponent has a winning strategy in the game on $(\mathfrak{P}', \mathcal{H}^+)$ where Opponent is forbidden to play γ^2 .
- (iii) Proponent has a winning strategy in the game on $(\mathfrak{P}, \mathcal{H})$.

In the game $(\mathfrak{P}', \mathcal{H}^+)$, Proponent must play the copy of \mathcal{H}^+ in E_0 to itself in the template. Thereafter, Proponent may never play this γ on the template, since \mathcal{H}^+ is a core. The equivalence of (ii) and (iii) follows.

The equivalence of (i) and (ii) follows from the fact that γ is adjacent to everything in \mathcal{H}^+ , so Opponent gains no advantage in playing it. \square

Corollary. *The graphs associated with the $(2n + 1)$ -gonal pyramids give rise to Pspace-complete ALT-HOM problems.*

Proof. The graph associated with the $(2n + 1)$ -gonal pyramid is \mathcal{C}_{2n+1}^+ . The case of the pentagonal pyramid appears in Figure 7.10. \square

7.5.3 A sufficient condition for ALT-HOM problem equivalence.

It is well known that, for any digraph \mathcal{G} whose core is \mathcal{H} , we have $\text{HOM}(\mathcal{G}) = \text{HOM}(\mathcal{H})$, i.e., for all digraphs \mathcal{D} we have $\mathcal{D} \in \text{HOM}(\mathcal{G})$ iff $\mathcal{D} \in \text{HOM}(\mathcal{H})$. This result does not extend to ALT-HOM.

Example. Let \mathcal{G} be $\mathcal{K}_3 \uplus \mathcal{K}_3$, therefore its core \mathcal{H} will be \mathcal{K}_3 . Consider a partitioned graph \mathfrak{P} whose underlying graph is the directed 3-path \mathcal{DP}_3 . Placing the two end-nodes in partition U_1 and the middle node in partition E_2 , we will have $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{H})$ whilst $\mathfrak{P} \notin \text{ALT-HOM}(\mathcal{G})$. This is because, for $\varphi :=$

$$\forall x \forall z \exists y E(x, y) \wedge E(y, z),$$

we have $\mathcal{K}_3 \models \varphi$, but $\mathcal{K}_3 \uplus \mathcal{K}_3 \not\models \varphi$.

However, we propose a digraph construction that preserves the alternating-homomorphism problem.

²Note that γ is not necessarily well-defined in \mathcal{H}^+ until Proponent plays the copy of \mathcal{H}^+ in \mathfrak{P}' on to the template \mathcal{H}^+ .

Definition. Given a digraph \mathcal{G} and specified vertex g , we construct \mathcal{G}^{+g} by *duplicating* the vertex g . Specifically:

- $|\mathcal{G}^{+g}| = |\mathcal{G}| \cup \{g'\}$, and

-

$$E^{\mathcal{G}^{+g}} = E^{\mathcal{G}} \cup \{(g', x) : (g, x) \in E^{\mathcal{G}}\} \cup \{(x, g') : (x, g) \in E^{\mathcal{G}}\} \\ \cup \{(g, g'), (g', g), (g', g') : \text{iff } (g, g) \in E^{\mathcal{G}}\}.$$

Theorem 77. For all digraphs \mathcal{G} , and any $g \in \mathcal{G}$, the problems $\text{ALT-HOM}(\mathcal{G})$ and $\text{ALT-HOM}(\mathcal{G}^{+g})$ are equal, i.e. for all partitioned digraphs \mathfrak{P} we have $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{G})$ iff $\mathfrak{P} \in \text{ALT-HOM}(\mathcal{G}^{+g})$.

Proof.

- (Forwards) We prove that a winning strategy σ for Proponent in the $(\mathfrak{P}, \mathcal{G})$ -game can be translated to a winning strategy σ' for her in the $(\mathfrak{P}, \mathcal{G}^{+g})$ -game.

The strategy σ' will tell Proponent that, if Opponent ever plays g' in the game on $(\mathfrak{P}, \mathcal{G}^{+g})$, she should behave in the game on $(\mathfrak{P}, \mathcal{G})$ exactly as if he had played on g . Since g and g' are adjacent to *exactly* the same vertices in \mathcal{G}^{+g} , any play of the $(\mathfrak{P}, \mathcal{G})$ -game that results in homomorphism must yield a play of the $(\mathfrak{P}, \mathcal{G}^{+g})$ -game that results in homomorphism.

- (Backwards) We prove that a winning strategy σ for Proponent in the $(\mathfrak{P}, \mathcal{G}^{+g})$ -game can be translated to a winning strategy σ' for her in the $(\mathfrak{P}, \mathcal{G})$ -game. Indeed, if we take σ and substitute all instances of g' for g , then we will have such a strategy, for the same reason as before.

□

Definition. For the n -clique \mathcal{K}_n , define \mathcal{K}_n^- to be \mathcal{K}_n with any single edge removed.

Corollary. For all $n \geq 4$, the problem $\text{ALT-HOM}(\mathcal{K}_n^-)$ is Pspace-complete.

Proof. Observe that $\mathcal{K}_n^- = \mathcal{K}_{n-1}^{+g}$, for any $g \in \mathcal{K}_{n-1}$.

□

Corollary. *The graphs associated with the $(2n + 1)$ -gonal bipyramids give rise to Pspace-complete ALT-HOM problems.*

Proof. The graph associated with the $(2n + 1)$ -gonal bipyramid is $(\mathcal{C}_{2n+1}^+)^{+\gamma}$, where γ is the vertex of \mathcal{C}_{2n+1}^+ adjacent to everything but itself. \square

7.5.4 Why that condition is not necessary: equivalence in fragments of FO.

Definition.

- Let $\mathbf{FO} \setminus \{=\}$ be first-order logic **FO** deprived of the binary equality relation.
- Let *pos-conj-FO* be the fragment of **FO** involving formulae in prenex form, whose quantifier-free portion is positive conjunctive.
- Let \exists -*pos-conj-FO* be the existential fragment of *pos-conj-FO*.

Note the trivial containments \exists -*pos-conj-FO* \subseteq *pos-conj-FO* \subseteq $\mathbf{FO} \setminus \{=\}$ \subseteq **FO**. These containments are readily seen to be proper; we will return to this later.

\exists -*pos-conj-FO* is closely related to the problems CSP and HOM. Indeed, it follows from the definition of CSP that, for any two templates \mathcal{T} and \mathcal{T}' , $\text{CSP}(\mathcal{T}) = \text{CSP}(\mathcal{T}')$ iff \mathcal{T} and \mathcal{T}' agree on all sentences of \exists -*pos-conj-FO*. Similarly, it follows from the definition of QCSP that $\text{QCSP}(\mathcal{T}) = \text{QCSP}(\mathcal{T}')$ iff \mathcal{T} and \mathcal{T}' agree on all sentences of *pos-conj-FO*. We also have, from the definition of the HOM problems, that $\text{HOM}(\mathcal{T}) = \text{HOM}(\mathcal{T}')$ iff \mathcal{T} is homomorphically equivalent to \mathcal{T}' , i.e. we have both $\mathcal{T} \xrightarrow{h} \mathcal{T}'$ and $\mathcal{T}' \xrightarrow{h} \mathcal{T}$. This is equivalent to the condition that \mathcal{T} and \mathcal{T}' have isomorphic cores. The concept of core gives us a combinatorial characterisation of what it is to be \exists -*pos-conj-FO*-equivalent. Such a characterisation seems harder for the logic *pos-conj-FO*: in the world of QCSP and ALT-HOM. However, such a characterisation is at hand with the logic $\mathbf{FO} \setminus \{=\}$, but first we must define the pertinent Ehrenfeucht-Fraïssé game. Our exposition is based on that for the standard **FO**-game given in [31].

Definitions. Let \mathcal{A} and \mathcal{B} be digraphs. The $\mathbf{FO} \setminus \{=\}_k^m$ -game on a pair $(\mathcal{A}, \alpha_0, \mathcal{B}, \beta_0)$ is played by two players, Spoiler and Duplicator, with k pairs of pebbles over m

rounds. A position in such a game, a k -configuration, is a pair of partial functions (α, β) , where

- $\alpha : \{v_1, \dots, v_k\} \rightarrow |\mathcal{A}|$, and
- $\beta : \{v_1, \dots, v_k\} \rightarrow |\mathcal{B}|$,

and we further have that $\text{dom}(\alpha) = \text{dom}(\beta)$. The domains of α and β are the pebbles that have already been played. From some position (α_j, β_j) , for the next move, Spoiler picks some pebble i (where $1 \leq i \leq k$) and chooses to play it in either $|\mathcal{A}|$ or $|\mathcal{B}|$. If the former [resp. latter], then he adds the pair (v_i, a) to α_j [resp. (v_i, b) to β_j] and Duplicator adds the pair (v_i, b) to β_j [resp. (v_i, a) to α_j], so obtaining the new position $(\alpha_{j+1}, \beta_{j+1})$. Note that v_i may already have been in $\text{dom}(\alpha_j) = \text{dom}(\beta_j)$, i.e. the pebble i may already have been played, in which case some former pairs $(v_i, a') \in \alpha_j$ and $(v_i, b') \in \beta_j$ will have been removed. The initial position of the game is (α_0, β_0) . Spoiler wins if at any point the relation $\alpha^{-1} \circ \beta \subseteq |\mathcal{A}| \times |\mathcal{B}|$ does not satisfy:

$$(*) \quad E^{\mathcal{A}}(a, a') \Leftrightarrow E^{\mathcal{B}}(b, b') \quad \text{for all } (a, b), (a', b') \in \alpha^{-1} \circ \beta$$

If Spoiler does not win, then Duplicator wins.

The $\mathbf{FO} \setminus \{=\}$ -game on $(\mathcal{A}, \alpha_0, \mathcal{B}, \beta_0)$ is played similarly, but with an unbounded number of pebble pairs and an unbounded number of moves. This means that (α_0, β_0) is an ω -configuration with potentially infinite domain: though only if the initial position were infinite. Since we are concerned with finite digraphs, this will never happen.

The quantifier-rank qr of a formula $\mathbf{FO} \setminus \{=\}$ ³ is defined inductively thus: if φ is quantifier-free then $qr(\varphi) = 0$; if $\varphi = \neg\varphi'$ then $qr(\varphi) = qr(\varphi')$; if $\varphi = \varphi' \wedge \varphi''$ then $qr(\varphi) = \max\{qr(\varphi'), qr(\varphi'')\}$; and if $\varphi = \exists v\varphi'$ then $qr(\varphi) = qr(\varphi') + 1$.

Let $\mathbf{FO} \setminus \{=\}_k^m$ be that fragment of $\mathbf{FO} \setminus \{=\}$ whose formulae involve only the variables v_1, \dots, v_k and whose quantifier-rank is at most m . For (α, β) a k -configuration and $\varphi \in \mathbf{FO} \setminus \{=\}_k^m$ a formula whose free variables are among $\text{dom}(\alpha) = \text{dom}(\beta)$ we consider φ to be true on (\mathcal{A}, α) [resp. (\mathcal{B}, β)] if φ is true

³Quantifier-rank is defined identically in \mathbf{FO} .

on \mathcal{A} [resp. \mathcal{B}] under free variable assignment α [resp. β]. Letting (α, β) be a k -configuration, we write:

- $(\mathcal{A}, \alpha) \cong_k^m (\mathcal{B}, \beta)$ iff \mathcal{A} and \mathcal{B} agree on all formulae of $\mathbf{FO} \setminus \{=\}_k^m$ (whose free variables are among $\text{dom}(\alpha) = \text{dom}(\beta)$), and
- $(\mathcal{A}, \alpha) \sim_k^m (\mathcal{B}, \beta)$ iff Duplicator has a winning strategy for the $\mathbf{FO} \setminus \{=\}_k^m$ -game on $(\mathcal{A}, \alpha, \mathcal{B}, \beta)$.

Letting (α, β) be an ω -configuration, we write:

- $(\mathcal{A}, \alpha) \cong (\mathcal{B}, \beta)$ iff \mathcal{A} and \mathcal{B} agree on all formulae of $\mathbf{FO} \setminus \{=\}$ (whose free variables must be among $\text{dom}(\alpha) = \text{dom}(\beta)$), and
- $(\mathcal{A}, \alpha) \sim (\mathcal{B}, \beta)$ iff Duplicator has a winning strategy for the $\mathbf{FO} \setminus \{=\}$ -game on $(\mathcal{A}, \alpha, \mathcal{B}, \beta)$.

Lemma 78. *There are only finitely many inequivalent formulae of $\mathbf{FO} \setminus \{=\}_k^m$.*

Proof. We prove the result for digraphs: it is easily extended to arbitrary (finite, relational) signatures. We proceed by induction on m .

(Base Case.) For $m = 0$, the formulae we can write are boolean combinations of $E(v_i, v_j)$ (for $i, j \in \{1, \dots, k\}$). We may consider any $\varphi \in \mathbf{FO} \setminus \{=\}_k^0$ to be a propositional formula in these k^2 propositional variables. We may rewrite this in CNF to obtain a formula $\varphi \in \mathbf{FO} \setminus \{=\}_k^0$ s.t. for all digraphs \mathcal{G} , $\mathcal{G} \models \varphi \leftrightarrow \varphi'$. The number distinct clauses for a formula of $\mathbf{FO} \setminus \{=\}_k^0$ in CNF is bounded by 2^{2k^2} , so it follows that the number of inequivalent formulae in CNF is bounded by $2^{2^{2k^2}}$. The result for base case follows.

(Inductive step.) Assume it is true for m . Any formula $\varphi \in \mathbf{FO} \setminus \{=\}_k^{m+1}$ is a boolean combination of formulae of the form $\exists v_i \varphi'$ with $\varphi' \in \mathbf{FO} \setminus \{=\}_k^m$. It follows from the inductive hypothesis that the number of inequivalent such formulae φ is finite, say c , and that therefore the number of inequivalent formulae in $\mathbf{FO} \setminus \{=\}_k^{m+1}$ is bounded above by 2^{2^c} . \square

Proposition 79 (Methodology). *Let \mathcal{A} and \mathcal{B} be digraphs and let (α_0, β_0) be an initial k -configuration. Then the following are equivalent:*

$$(i) (\mathcal{A}, \alpha_0) \sim_k^m (\mathcal{B}, \beta_0)$$

$$(ii) (\mathcal{A}, \alpha_0) \cong_k^m (\mathcal{B}, \beta_0)$$

Proof. (Based on that for **FO** in [31].) We proceed by induction on m . For the base case, $m = 0$, Duplicator wins the zero-round game on $(\mathcal{A}, \alpha_0, \mathcal{B}, \beta_0)$ iff $\alpha_0^{-1} \circ \beta_0$ satisfies (*) iff (\mathcal{A}, α_0) and (\mathcal{B}, β_0) agree on all quantifier-free formulae φ (whose free variables, indeed, only variables, are in $\text{dom}(\alpha_0) = \text{dom}(\beta_0)$).

Inductive step: (i) \Rightarrow (ii) (by contraposition). Suppose the proposition is true for m , but that (\mathcal{A}, α_0) and (\mathcal{B}, β_0) disagree on some formula $\varphi \in \mathbf{FO} \setminus \{=\}_k^{m+1}$. If φ were of the form $\neg\varphi'$ [resp. $\varphi' \wedge \varphi''$] then they would disagree on φ' [resp. one of φ', φ''], so we may assume w.l.o.g. that $\varphi = \exists v_i \varphi'$ ($1 \leq i \leq k$) and that $(\mathcal{A}, \alpha_0) \models \varphi$ and $(\mathcal{B}, \beta_0) \models \neg\varphi$. In playing the $\mathbf{FO} \setminus \{=\}_k^{m+1}$ -game on $(\mathcal{A}, \alpha_0, \mathcal{B}, \beta_0)$, Spoiler begins by playing a witness for φ in \mathcal{A} , but, no matter where Duplicator replies, we will end up with a position (α_1, β_1) s.t. $(\mathcal{A}, \alpha_1) \models \varphi'$ and $(\mathcal{B}, \beta_1) \models \neg\varphi'$. Since the quantifier-rank of φ' is m , it follows from the inductive hypothesis that Spoiler wins the $\mathbf{FO} \setminus \{=\}_k^m$ -game on $(\mathcal{A}, \alpha_1, \mathcal{B}, \beta_1)$, and we are done.

Inductive step: (ii) \Rightarrow (i). Suppose that $(\mathcal{A}, \alpha_0) \cong_k^m (\mathcal{B}, \beta_0)$, and let Spoiler take his first move in the $\mathbf{FO} \setminus \{=\}_k^{m+1}$ -game on $(\mathcal{A}, \alpha_0, \mathcal{B}, \beta_0)$. Let him place a pebble i on an element of \mathcal{A} , so defining α_1 . Remembering that there are only finitely inequivalent formulae of $\mathbf{FO} \setminus \{=\}_k^m$, let Φ be the conjunction of all of these formulae that (\mathcal{A}, α_1) satisfies. We know that $(\mathcal{A}, \alpha_0) \models \exists v_i \Phi$: so by assumption $(\mathcal{B}, \beta_0) \models \exists v_i \Phi$. Let Duplicator play her pebble i on a witness for $\exists v_i \Phi$ in \mathcal{B} . Thus (\mathcal{A}, α_1) and (\mathcal{B}, β_1) both satisfy Φ . Since Φ is a complete description of everything satisfied by (\mathcal{A}, α_1) in $\mathbf{FO} \setminus \{=\}_k^m$, it follows that $(\mathcal{A}, \alpha_1) \cong_k^m (\mathcal{B}, \beta_1)$, and the result follows from the inductive hypothesis. \square

Corollary. *Let \mathcal{A} and \mathcal{B} be digraphs and let (α_0, β_0) be an initial ω -configuration. Then:*

- $(\mathcal{A}, \alpha_0) \sim (\mathcal{B}, \beta_0)$ iff $(\mathcal{A}, \alpha_0) \cong (\mathcal{B}, \beta_0)$.
- $(\mathcal{A}, \alpha_0) \sim_{\|\mathcal{A}\| + \|\mathcal{B}\|} (\mathcal{B}, \beta_0)$ iff $(\mathcal{A}, \alpha_0) \sim (\mathcal{B}, \beta_0)$.

Proof. The first part follows immediately from the previous proposition and the definitions. The second part follows from the fact that Duplicator need only find

an answer for the $||\mathcal{A}|| + ||\mathcal{B}||$ positions Spoiler can play: thereafter she may copy previous replies. \square

We now introduce the converse of the vertex duplication that we have already seen.

Definition. If a digraph \mathcal{A} possesses two vertices a, a' such that $\{x : E(a, x) \in E^{\mathcal{A}}\} = \{x : E(a', x) \in E^{\mathcal{A}}\}$ and $\{x : E(x, a) \in E^{\mathcal{A}}\} = \{x : E(x, a') \in E^{\mathcal{A}}\}$, then \mathcal{A} may be *folded* to the graph \mathcal{A}^{-x} by collapsing the vertices x and x' to a single vertex x' (alternatively, removing vertex v).

A digraph that has no potential folds is said to be *stiff*. (Similar definitions for fold, and stiff graph, appear, e.g., in [26].)

Theorem 80. *The following are equivalent, for all digraphs \mathcal{A}, \mathcal{B} :*

- (i) \mathcal{A} and \mathcal{B} agree on all sentences of $\mathbf{FO} \setminus \{=\}$.
- (ii) Duplicator wins the $\mathbf{FO} \setminus \{=\}$ -game on $(\mathcal{A}, \mathcal{B})$.
- (iii) There exists a stiff \mathcal{C} such that both \mathcal{A} and \mathcal{B} may be put through a sequence of folds to derive [isomorphic copies of] \mathcal{C} .

Proof. We already have the equivalence of (i) and (ii).

((iii) \rightarrow (ii).) Suppose $f_{\mathcal{A}} : \mathcal{A} \rightarrow \mathcal{C}$ and $f_{\mathcal{B}} : \mathcal{B} \rightarrow \mathcal{C}$ are the surjective collapsing functions for the respective sequences of folds. For the $\mathbf{FO} \setminus \{=\}$ -game on $(\mathcal{A}, \mathcal{B})$, if Spoiler plays in $|\mathcal{A}|$ [resp. $|\mathcal{B}|$] then Duplicator should answer with any vertex $b \in |\mathcal{B}|$ s.t. $f_{\mathcal{B}}(b) = f_{\mathcal{A}}(a)$ [resp. any vertex $a \in |\mathcal{A}|$ s.t. $f_{\mathcal{A}}(a) = f_{\mathcal{B}}(b)$] (the existence of such vertices is guaranteed by surjectivity of $f_{\mathcal{A}}$ and $f_{\mathcal{B}}$). This is a winning strategy by the definition of folding and, its inverse, duplication.

((ii) \rightarrow (iii).) Suppose Spoiler plays all $||\mathcal{A}|| + ||\mathcal{B}||$ distinct vertices. Duplicator must be able to answer. Suppose $a_1, \dots, a_{||\mathcal{A}||}$ was answered with $a'_1, \dots, a'_{||\mathcal{A}||}$ and $b_1, \dots, b_{||\mathcal{B}||}$ was answered with $b'_1, \dots, b'_{||\mathcal{B}||}$. In \mathcal{A} repeatedly collapse vertices a_i, a_j to a single vertex iff $a'_i = a'_j$. Continuing until there are no more vertices to collapse, we ultimately build a stiff $\mathcal{C}_{\mathcal{A}}$. In \mathcal{B} repeatedly collapse vertices b_i, b_j to a single vertex iff $b'_i = b'_j$, so obtaining a stiff $\mathcal{C}_{\mathcal{B}}$. It follows by transitivity that Duplicator has a winning strategy in the $\mathbf{FO} \setminus \{=\}$ -game on $(\mathcal{C}_{\mathcal{A}}, \mathcal{C}_{\mathcal{B}})$. Let Spoiler

play all the positions in \mathcal{C}_A and let Duplicator make her reply. The so obtained $\alpha^{-1} \circ \beta$ that satisfies (*) must also satisfy:

$$\text{injectivity: } \forall a, a' \in \mathcal{C}_A \forall b \in \mathcal{C}_B \quad (a, b), (a', b) \in \alpha^{-1} \circ \beta \Rightarrow a = a'$$

(for otherwise \mathcal{A}_e is not stiff since a may be folded to a'), and

$$\text{surjectivity: } \forall b \in \mathcal{C}_B \exists a \in \mathcal{C}_A \quad (a, b) \in \alpha^{-1} \circ \beta$$

(for otherwise \mathcal{C}_B is not stiff). It follows that $\alpha^{-1} \circ \beta$ is an isomorphism, and the result follows. \square

Corollary ([26]). *Every digraph \mathcal{G} has a unique [up to isomorphism] stiff induced subgraph that it can be put through a sequence of folds to obtain.*

Proof. Take $\mathcal{A} = \mathcal{B} = \mathcal{G}$ in the previous proof. \square

We will now unambiguously refer to the *stiff-graph-within* \mathcal{G} as the one which \mathcal{G} reaches through a maximal sequence of folds. We have seen that stiff-graphs-within characterise $\mathbf{FO} \setminus \{=\}$ -equivalence in exactly the way that cores characterise \exists -*pos-conj-FO*-equivalence. This gives us a new proof of Theorem 77, in which we proved that, for all digraphs \mathcal{G} and $g \in \mathcal{G}$, $\text{ALT-HOM}(\mathcal{G}) = \text{ALT-HOM}(\mathcal{G}^{+g})$. Since \mathcal{G}^{+g} may be folded to \mathcal{G} , it follows that they share the same stiff-graph-within. It is now clear that \mathcal{G} and \mathcal{G}^{+g} agree on all sentences $\mathbf{FO} \setminus \{=\}$, which certainly includes all sentences of *pos-conj-FO*.

It remains for us to ask whether or not stiff-graphs-within capture equivalence in QCSP, *i.e.* whether *pos-conj-FO* actually coincides with $\mathbf{FO} \setminus \{=\}$. It turns out not to be so; we demonstrate the proper containments \exists -*pos-conj-FO* \subset *pos-conj-FO* \subset $\mathbf{FO} \setminus \{=\}$ \subset \mathbf{FO} .

Examples.

\mathcal{K}_3 and $\mathcal{K}_3 \uplus \mathcal{K}_3$ give rise to the same CSP problem, *i.e.* are equivalent in \exists -*pos-conj-FO*, but do not give rise to the same QCSP problem, since $\forall x \forall z \exists y E(x, y) \wedge E(y, z)$ is true in the former, yet false in the latter. It follows that \exists -*pos-conj-FO* can not express that property.

$\mathcal{K}_2 = \mathcal{P}_2$ and \mathcal{P}_4 give rise to the same QCSP problems. (In fact, it follows from Lemmas 58 and 66 that there are precisely three classes of ALT-HOM = QCSP problem for bipartite template \mathcal{T} : specifically, \mathcal{T} having no edges; or \mathcal{T} has an edge and an isolated vertex; or \mathcal{T} has an edge and no isolated vertices. \mathcal{P}_2 and \mathcal{P}_4 are both in the last class.) Consider the sentence $\exists w, x, y, z E(w, x) \wedge E(x, y) \wedge E(y, z) \wedge \neg E(z, x)$. This is false in the former, but true in the latter. It follows that the property can not be expressed in *pos-conj-FO*.

Finally, consider the query $\forall x \forall y E(x, y) \vee x = y$. This can not be expressed in $\mathbf{FO} \setminus \{=\}$ since \mathcal{K}_3 and \mathcal{K}_4^- disagree on it, yet agree on all sentences of $\mathbf{FO} \setminus \{=\}$.

7.6 Results concerning tournament templates

7.6.1 Template is a directed cycle.

We consider the case where the template \mathcal{T} is a directed n -cycle \mathcal{DC}_n . Such a graph is a tournament only when $n = 3$, but the method easily generalises.

Definition. An *oriented path* is a list of vertices v_1, \dots, v_m and, for $1 \leq i < m$, exactly one of the edges $E(v_i, v_{i+1})$ or $E(v_{i+1}, v_i)$. The *net length* of this oriented path is the number of instances of edges $E(v_i, v_{i+1})$ (forward-edges) minus the number of instances of edges $E(v_{i+1}, v_i)$ (backward-edges). An oriented path in a digraph \mathcal{G} is a (not necessarily induced) subgraph of \mathcal{G} that is an oriented path.

In a directed n -cycle any oriented path between a vertex and itself must have net length $0 \pmod n$. Furthermore, any path between a vertex and its forward-neighbour must have net length $1 \pmod n$, and every vertex *has* a forward-neighbour. These facts will allow us once again to consider only partitioned inputs in Π_2 -multifan form since:

Lemma 81. For $n \geq 3$,

- if there is a path in \mathfrak{B} between any $x \in E_i$ and $y \in U_j$ (for $i < j$), then $\mathfrak{B} \notin \text{ALT-HOM}(\mathcal{DC}_n)$, and
- if there is a path in \mathfrak{B} between any $x \in U_i$ and $y \in U_j$ (any i, j), then $\mathfrak{B} \notin \text{ALT-HOM}(\mathcal{DC}_n)$.

Proof. We prove the first claim, the proof of the second is similar. If the path has net length $0 \bmod n$, then, if Proponent plays a for x , then Opponent plays the forward-neighbour b of a for y , and wins. If the path has net length other than $0 \bmod n$, then, if Proponent plays a for x , then Opponent also plays a for x , and again wins. \square

Theorem 82. *If \mathcal{T} is a directed n -cycle then $\text{ALT-HOM}(\mathcal{T})$ is tractable.*

Proof. If the input \mathfrak{P} has any of the paths of the previous lemma, then we have a no-instance. We may therefore assume that \mathfrak{P} has no such path, and is equivalent to the rewrite-reduced $\overline{\mathfrak{P}}$ in Π_2 -multifan form. We may further split $\overline{\mathfrak{P}}$ into its Π_2 -fan form components $\overline{\mathfrak{P}}_i$ ($1 \leq i \leq m$), for some m , and solve separately for each.

Since \mathcal{T} is rotationally symmetric, each Π_2 -fan structure $\overline{\mathfrak{P}}_i$ admits an alternating-homomorphism to \mathcal{T} iff the structure $\mathcal{S}_{\overline{\mathfrak{P}}_i}$ admits a homomorphism to \mathcal{T} . It is known that the problem $\text{HOM}(\mathcal{T})$ is tractable [4], and the result follows. \square

7.6.2 Template is a digraph with source and sink.

In a digraph, a source (respectively, sink) is a vertex with in-degree (respectively, out-degree) 0.

We consider the case where the template \mathcal{T} is a digraph with both a source s and a sink t . In such cases we need only consider inputs in Σ_1 -form since:

Lemma 83. *For any i, j ,*

- *if there is a forward-edge in \mathfrak{P} between any $x \in E_i$ and $y \in U_j$, then $\mathfrak{P} \notin \text{ALT-HOM}(\mathcal{T})$,*
- *if there is a backward-edge in \mathfrak{P} between any $x \in E_i$ and $y \in U_j$, then $\mathfrak{P} \notin \text{ALT-HOM}(\mathcal{T})$,*
- *if there is any forward-edge in \mathfrak{P} between any $x \in U_i$ and $y \in U_j$, then $\mathfrak{P} \notin \text{ALT-HOM}(\mathcal{T})$, and*
- *if there is any backward-edge in \mathfrak{P} between any $x \in U_i$ and $y \in U_j$, then $\mathfrak{P} \notin \text{ALT-HOM}(\mathcal{T})$.*

Proof. For the first and third claims, Opponent plays s for y and wins. For the second and fourth claims, Opponent plays t for y and wins. \square

The following should be seen as a generalisation of Proposition 59.

Proposition 84. *If \mathcal{T} is an antireflexive digraph with both a source and a sink, then $\text{ALT-HOM}(\mathcal{T})$ is logspace equivalent to $\text{HOM}(\mathcal{T})$.*

Proof. Use the reduction of Proposition 59 (we have the same forbidden edges here as we did there). \square

Definition. An n -tournament is a digraph \mathcal{G} with vertex set $\{0, \dots, n-1\}$, such that, for all $i, j \in \mathbb{Z}_n$, exactly one of $E(i, j)$ or $E(j, i)$ is an edge of \mathcal{G} . The unique n -tournament which contains no directed cycle as a subgraph is known as the transitive n -tournament, and will be denoted \mathcal{T}_n^t .

Corollary. *If \mathcal{T}_n^t is the transitive n -tournament then $\text{ALT-HOM}(\mathcal{T}_n^t)$ is tractable.*

Proof. \mathcal{T}_n^t has both a source and a sink. So, the result follows from the previous theorem, and the fact that $\text{HOM}(\mathcal{T}_u^t)$ is tractable [4]. \square

7.6.3 Tractable tournament ALT-HOM problems.

The tournaments \mathcal{T}_{m+3}^u

We examine the tournaments \mathcal{T}_{m+3}^u which are constructed from the directed 3-cycle by repeatedly adding a source m times. (The superscript u suggests this unique cycle.)

Definition. We define \mathcal{T}_{m+3}^u inductively:

- Let $\mathcal{T}^{(0)} := \mathcal{T}_3^u = \mathcal{DC}_3$, the directed 3-cycle..
- From $\mathcal{T}^{(r)}$ build $\mathcal{T}^{(r+1)}$ by adding a new source, i.e.,
 $|\mathcal{T}^{(r+1)}| := |\mathcal{T}^{(r)}| \uplus \{r+3\}$ ⁴ and $E^{\mathcal{T}^{(r+1)}} := E^{\mathcal{T}^{(r)}} \uplus \{(r+3, i) : i \in |\mathcal{T}^{(r)}|\}$.
- Let $\mathcal{T}_{m+3}^u := \mathcal{T}^{(m)}$.

⁴ $\mathcal{T}^{(r)}$, being a tournament with $r+3$ vertices, will already have vertex numbers 0 to $r+2$.

Since we have dealt with the case of the directed 3-cycle, we consider $m > 0$, i.e. when \mathcal{T}_{m+3}^u has a source.

Lemma 85. *For $m > 0$:*

- *If there is a directed edge in \mathfrak{P} between $x \in E_i$ and $y \in U_j$ ($i < j$), then $\mathfrak{P} \notin \text{ALT-HOM}(\mathcal{T}_{m+3}^u)$.*
- *If there is a directed edge in \mathfrak{P} between $x \in U_i$ and $y \in U_j$ (any i, j), then $\mathfrak{P} \notin \text{ALT-HOM}(\mathcal{T}_{m+3}^u)$.*
- *If there is a directed edge in \mathfrak{P} from $x \in E_j$ to $y \in U_i$ ($i < j$), then $\mathfrak{P} \notin \text{ALT-HOM}(\mathcal{T}_{m+3}^u)$.*

Proof. The first two parts follow from the antireflexivity of \mathcal{T}_{m+3}^u , by Lemma 57. For the final part, if Opponent plays y to the source of \mathcal{T}_{m+3}^u , Proponent can have no reply for x . \square

Let \mathfrak{P} be a partitioned digraph, we define its cousin $\tilde{\mathfrak{P}}$ inductively:

- $\mathfrak{P}^{(0)} := \mathfrak{P}$.
- From $\mathfrak{P}^{(r)}$ we build $\mathfrak{P}^{(r+1)}$ by removing *all* sources that are in existential partitions.
- $\tilde{\mathfrak{P}} := \mathfrak{P}^{(m)}$.

Let $Ex(\mathfrak{P} - \tilde{\mathfrak{P}})$ be those existential vertices in \mathfrak{P} that are not in $\tilde{\mathfrak{P}}$, and let $Ex(\tilde{\mathfrak{P}})$ be those existential vertices in \mathfrak{P} that are also in $\tilde{\mathfrak{P}}$ (since $\tilde{\mathfrak{P}} \subseteq \mathfrak{P}$ these are the existential vertices of $\tilde{\mathfrak{P}}$). Let $PrEx(\mathfrak{P})$ be those vertices of \mathfrak{P} in existential partitions *to* which there is a directed path *from* some vertex in a universal partition. Let $Un(\mathfrak{P})$ be the set of universal vertices of \mathfrak{P} . We refer to the vertices of \mathcal{T}_{m+3}^u that are not in the 3-cycle as the *tail* of \mathcal{T}_{m+3}^u .

We will benefit from examining which vertices of the underlying graph $\mathcal{S}_{\mathfrak{P}}$ have been removed in the graph $\mathcal{S}_{\tilde{\mathfrak{P}}}$. It should be clear that vertices in $Un(\mathfrak{P})$ and $PrEx(\mathfrak{P})$ can never be removed, and are, therefore, *protected*. Let us consider the sub-partitioned-graph \mathfrak{P}_1 of \mathfrak{P} induced by the existential vertices that are not protected. \mathfrak{P}_1 may be put through our given inductive scheme, iteratively removing

sources m times, so obtaining $\widetilde{\mathfrak{P}}_1$. It should be clear that $\widetilde{\mathfrak{P}}$ is that subgraph of \mathfrak{P} induced by the set $Un(\mathfrak{P}) \cup PrEx(\mathfrak{P}) \cup |\mathcal{S}_{\widetilde{\mathfrak{P}}_1}|$. Apart from the universal vertices and those existential vertices that are protected, our construction is that given for proving the tractability of $HOM(\mathcal{T}_{m+3}^u)$ in [4]. All of the sets we have defined should now be considered as subsets of \mathfrak{P} (though some may be subsets of $\widetilde{\mathfrak{P}}$ too). Before going on we will benefit from the following lemmas.

Lemma 86. *In a winning strategy for Proponent on $(\mathfrak{P}, \mathcal{T}_{m+3}^u)$, if Opponent plays all his vertices to the 3-cycle, then Proponent must play all of the vertices of $Ex(\widetilde{\mathfrak{P}})$ [in \mathfrak{P}] to the 3-cycle.*

Proof. Again, let \mathfrak{P}_1 be the sub-partitioned-graph of \mathfrak{P} induced by those existential vertices that are not protected. Recall that $\widetilde{\mathfrak{P}}$ is the subgraph of \mathfrak{P} induced by $Un(\mathfrak{P}) \cup PrEx(\mathfrak{P}) \cup |\mathcal{S}_{\widetilde{\mathfrak{P}}_1}|$, as in the previous paragraph. So $Ex(\widetilde{\mathfrak{P}})$ is $PrEx(\mathfrak{P}) \cup |\mathcal{S}_{\widetilde{\mathfrak{P}}_1}|$. Since universal vertices are played to the 3-cycle, it follows that all vertices of $PrEx(\mathfrak{P})$ must be played to the 3-cycle. Furthermore, if any vertex of $|\mathcal{S}_{\widetilde{\mathfrak{P}}_1}|$ [in \mathfrak{P}] could be played to the tail of \mathcal{T}_{m+3}^u , then this could not be extended to a homomorphism from $\mathcal{S}_{\mathfrak{P}_1}$ to \mathcal{T}_{m+3}^u – by definition of $\widetilde{\mathfrak{P}}_1$ – so this could not be a winning strategy for Proponent on $(\mathfrak{P}, \mathcal{T}_{m+3}^u)$. The result follows. \square

Lemma 87. *Assume that \mathfrak{P} has none of the edges of Lemma 85. Then Opponent can win the game on $(\mathfrak{P}, \mathcal{T}_{m+3}^u)$ iff he can win it whilst never playing in the tail of \mathcal{T}_{m+3}^u .*

Proof. Since edges of \mathfrak{P} from universal partitions only point toward vertices x in higher existential partitions, if Opponent plays in the tail then he allows Proponent to answer x with anything on the 3-cycle, whereas, if he plays on the 3-cycle he limits Proponent to a single adjacent vertex of the 3-cycle. It is clear that Opponent gains nothing by playing in the tail. \square

Theorem 88. *The problems $ALT-HOM(\mathcal{T}_{m+3}^u)$ are tractable.*

Proof. We already have the result for $m = 0$. For $m > 0$ we will solve $ALT-HOM(\mathcal{T}_{m+3}^u)$ by taking any input \mathfrak{P} for that problem, and constructing a given \mathfrak{P}' . We will prove $\mathfrak{P} \in ALT-HOM(\mathcal{T}_{m+3}^u)$ iff $\mathfrak{P}' \in ALT-HOM(\mathcal{T}_3^u)$, whereupon we may appeal to the known tractability of $ALT-HOM(\mathcal{T}_3^u)$, and our result will follow.

If \mathfrak{P} has any of the edges of Lemma 85 then we define \mathfrak{P}' to be any set no-instance of $\text{ALT-HOM}(\mathcal{T}_3^u)$ (e.g. the transitive 4-tournament \mathcal{T}_4^u with all vertices in E_2). If \mathfrak{P} has none of those edges then we set \mathfrak{P}' to be $\tilde{\mathfrak{P}}$, via the construction already described. It remains for us to prove that this is correct. It is trivially correct if \mathfrak{P} has any of the edges of Lemma 85: we assume it does not.

($\mathfrak{P} \in \text{ALT-HOM}(\mathcal{T}_{m+3}^u) \Rightarrow \tilde{\mathfrak{P}} \in \text{ALT-HOM}(\mathcal{T}_3^u)$.) For a winning strategy σ for Proponent in the game on $(\mathfrak{P}, \mathcal{T}_{m+3}^u)$, we claim σ is also a winning strategy for her in the game on $(\tilde{\mathfrak{P}}, \mathcal{T}_3^u)$. This follows immediately from Lemma 86.

($\tilde{\mathfrak{P}} \in \text{ALT-HOM}(\mathcal{T}_3^u) \Rightarrow \mathfrak{P} \in \text{ALT-HOM}(\mathcal{T}_{m+3}^u)$.) From a winning strategy σ for Proponent in the game on $(\tilde{\mathfrak{P}}, \mathcal{T}_3^u)$, we construct a winning strategy σ' in the game on $(\mathfrak{P}, \mathcal{T}_{m+3}^u)$ where Opponent only plays in the 3-cycle. In that game on $(\mathfrak{P}, \mathcal{T}_{m+3}^u)$, when Opponent plays on the 3-cycle, then Proponent answers the vertices in $Ex(\tilde{\mathfrak{P}})$ according to σ , and then maps $Ex(\mathfrak{P} - \tilde{\mathfrak{P}})$ to the tail of \mathcal{T}_{m+3}^u . The result follows from Lemma 87. \square

The tournaments \mathcal{T}_{3+m}^u

These tournaments are analogous to the tournaments \mathcal{T}_{m+3}^u , but are constructed by the repeated addition of a sink, rather than a source. It follows by a similar argument that, for all m , $\text{ALT-HOM}(\mathcal{T}_{m+3}^u)$ is tractable.

The result

Theorem 89. *If \mathcal{T} is a tournament with at most one cycle then $\text{ALT-HOM}(\mathcal{T})$ is tractable.*

Proof. It follows from standard results about tournaments (see [4]) that \mathcal{T} is either transitive or is \mathcal{T}_3^u with a succession of sources and/or sinks added. If \mathcal{T} has both a source and a sink then we can reduce the problem to $\text{HOM}(\mathcal{T})$, which is known to be tractable [4]. If it has no sink, then it is one of the tournaments \mathcal{T}_{m+3}^u above. If it has no source, then it is one of the tournaments \mathcal{T}_{3+m}^u above. \square

7.6.4 NP-complete tournament ALT-HOM problems.

Bang-Jenson, Hell, and MacGillavray proved that, for any n -tournament \mathcal{T}_n that has at least two distinct cycles, $\text{HOM}(\mathcal{T}_n)$ is NP-complete [4].

Theorem 90. *Let \mathcal{T} be a tournament with more than one cycle and a source and sink, then $\text{ALT-HOM}(\mathcal{T})$ is NP-complete.*

Proof. Follows from Proposition 84 , together with [4]. □

7.6.5 Pspace-complete tournament ALT-HOM problems.

A *2-walk Tournament* (2WT) is a tournament \mathcal{T} in which, for all distinct $i, j \in \mathcal{T}$, there is a directed 2-walk from x to y .

Theorem 91. *For every 2WT \mathcal{T} , $\text{ALT-HOM}(\mathcal{T})$ is Pspace-complete.*

Proof. Note that $||\mathcal{T}|| > 2$. Also, since \mathcal{T} is a tournament, there can be no 2-walk from any vertex $x \in \mathcal{T}$ to itself. Using the directed 2-path from vertex i to vertex j as indicator, we find that \mathcal{T}^* is $\mathcal{K}_{||\mathcal{T}||}$. Pspace-completeness follows from Theorem 73. □

We conclude by proving that the class of 2WTs is infinite.

Definition. For $m \geq 5$, define the Tournament \mathcal{T}_{2m+1}^p thus:

- $|\mathcal{T}_{2m+1}^p| = \{0, \dots, 2m\}$
- $E^{\mathcal{T}_{2m+1}^p} =$
 - $\{(i, j) : i, j \in \mathbb{Z}_{2m+1} \text{ s.t. } i + 1 = j \text{ mod } 2m + 1\} \cup$
 - $\{(i, j) : i, j \in \mathbb{Z}_{2m+1} \text{ s.t. } i - 2 = j \text{ mod } 2m + 1\} \cup$
 - $\{(i, j) : i, j \in \mathbb{Z}_{2m+1} \text{ s.t. } i + 3 = j \text{ mod } 2m + 1\} \cup$
 - $\{(i, j) : i, j \in \mathbb{Z}_{2m+1} \text{ s.t. } i + 4 = j \text{ mod } 2m + 1\} \cup$
 - $\{(i, j) : i, j \in \mathbb{Z}_{2m+1} \text{ s.t. } i + 5 = j \text{ mod } 2m + 1\} \cup$
 - \vdots
 - $\{(i, j) : i, j \in \mathbb{Z}_{2m+1} \text{ s.t. } i + m = j \text{ mod } 2m + 1\}$

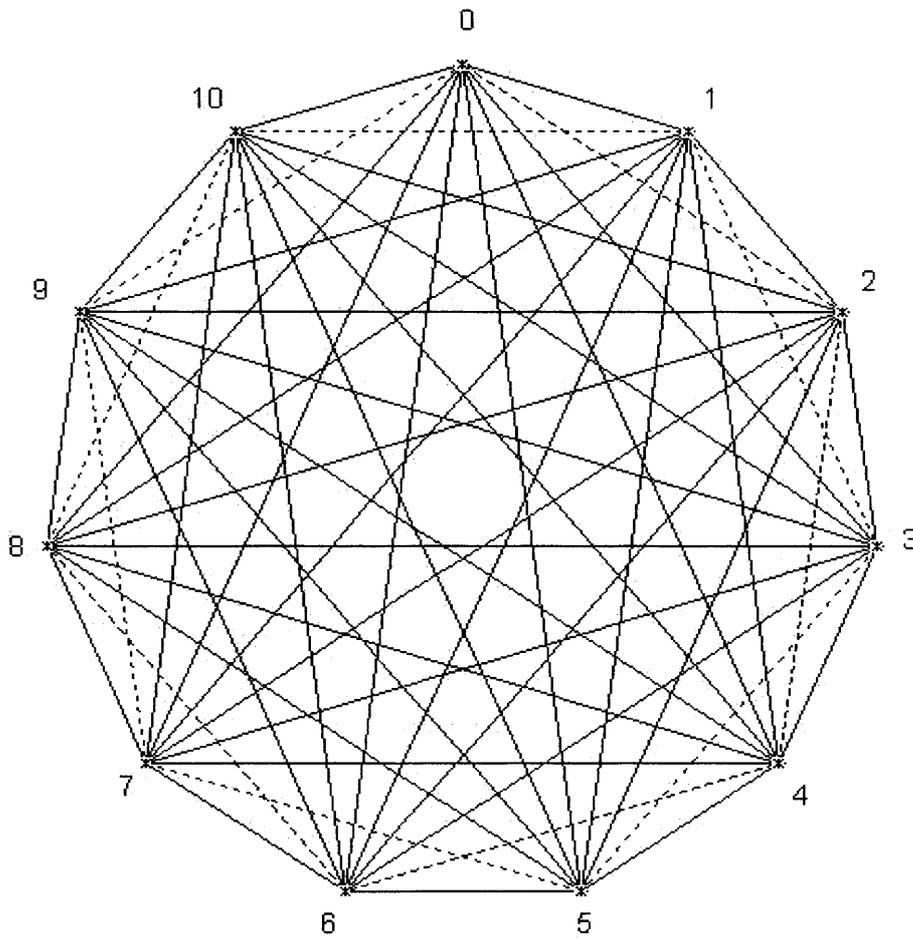


Figure 7.11: The tournament \mathcal{T}_{11}^p . The 2-jump dotted edges point anticlockwise; all other edges point clockwise.

It may easily be verified that \mathcal{T}_{2m+1}^p is a tournament. Note that this is partly a consequence of the odd number of vertices: if we had an even number of vertices under a similar construction we would either have vertices not joined by an edge or vertices joined by a double edge. Observe the aberration of the 2-jumps: if we draw edges on a regular $(2m+1)$ -gon with the vertices enumerated clockwise, then all edges point in a clockwise direction, except the 2-jumps which point anticlockwise (see Figure 7.11.)

Lemma 92. *For all distinct vertices $i, j \in \mathcal{T}_{2m+1}^p$, there is a directed 2-walk from i to j , i.e. \mathcal{T}_{2m+1}^p is a 2WT.*

Proof. We will prove that there is a directed 2-walk from vertex 0 to each of the vertices $1, \dots, 2m$. We may then appeal to symmetry.

It will suffice to show that every number $1, \dots, 2m$ is the sum (mod $2m+1$) of exactly two elements of the set $\{1, -2, 3, 4, 5, \dots, m\}$. So: $1 = 3 - 2$; $2 = 4 - 2$; $3 = 5 - 2$ (this is why $m \geq 5$), henceforth we may use the positive numbers $\{3, \dots, m\}$ for $6 = 3 + 3$ through to $2m = m + m$. \square

Chapter 8

Conclusions and Further Work

8.1 Program Schemes

Most of our results for program schemes augmented with priority queue are far from tight, and those that are tight are unsurprising. It is hard to see how a Turing machine simulation might prove that NPSPQ^u is contained in a space-bounded (or time-bounded) complexity class. This is because the size of the potential memory of the priority queue in NPSPQ^u seems to be unbounded. We give no better upper bound to NPSPQ^u than the class of recursively enumerable languages. For better lower bounds for NPSPQ^u , our simulation method can go no further than NPspace , since we rely on the fact that we can enumerate the tape squares through a constant number of variables (which may encode only a polynomial quantity of numbers). A similar problem arises in the case of a better lower bound for $\text{APSS}(1)$.

We suggest that an indirect method may be more likely to succeed, perhaps like that used to prove $\text{P} \subseteq \text{NPSS}_s(1)$ – via the path system problem – in [2].

8.2 Classes of Structure on which $\text{P} = \pm\text{PS}^k[\text{FO}]$

Grohe had proved in [23] that $\text{P} = \text{LFP}[\text{FO}]$ on the class of 3-connected planar graphs. Since triangulations are 3-connected planar graphs and it is known that $\text{P} = \pm\text{PS}[\text{FO}]$ on the class of triangulations [42], it is natural to question whether P

= LFP[FO] on the class of 3-connected planar graphs. Thus far, we have failed to adapt Grohe’s method to settle this question.

8.3 Quantified Constraints on Graphs

The method used to prove the Pspace-completeness of ALT-HOM(\mathcal{H}), where \mathcal{H} is an odd catherine wheel, may be applied to prove Pspace-completeness for similar templates. An obvious extension is for graphs that are constructed like odd catherine wheels, but where any bipartite graph (not just a tree) may be appended to each position on the odd cycle.

For quantified \mathcal{H} -colouring, we conjecture the following extension to Theorem 71.

Conjecture 93. The class of antireflexive undirected graphs exhibits ALT-HOM-trichotomy. Specifically:

- If \mathcal{H} is bipartite, then ALT-HOM(\mathcal{H}) is tractable.
- If \mathcal{H} is not bipartite, and is not connected, then ALT-HOM(\mathcal{H}) is NP-complete.
- If \mathcal{H} is not bipartite, and is connected, then ALT-HOM(\mathcal{H}) is Pspace-complete.

In order to prove this, it would remain for us to prove that ALT-HOM(\mathcal{H}) is Pspace-complete, when \mathcal{H} is antireflexive, undirected and connected, and has more than one odd cycle.

It is well-known that common cores characterise equivalent \exists -*pos-conj*-FO-theories (CSP/ HOM), i.e. two templates $\mathcal{T}, \mathcal{T}'$ have the same core iff they agree on all sentences of \exists -*pos-conj*-FO. Similarly, common stiff-graphs-within characterise equivalent FO\{=\}-theories. We know of no such characterisation for equivalence of *pos-conj*-FO-theories (QCSP/ ALT-HOM). It would be interesting to isolate some characteristic on two templates that exactly specifies whether they give rise to the same QCSP problem.

Bibliography

- [1] ABITEBOUL, S., AND VIANU, V. Generic computation and its complexity. In *32nd IEEE Symposium on FOCS* (1991).
- [2] ARRATIA-QUESEDA, A., CHAUHAN, S., AND STEWART, I. Hierarchies in classes of program schemes. *Journal of Logic and Computation* 9:6 (1999), 915–957.
- [3] ASCHBACHER, M., AND GURALNICK, R. Some applications of the first cohomology group. *Journal of Algebra* 90 (1984), 446–460.
- [4] BANG-JENSON, J., HELL, P., AND MAGILLIVRAY, G. The complexity of colourings by semi-complete digraphs. *SIAM Journal on Discrete Mathematics* 1:3 (1988).
- [5] BORNER, F., KROKHIN, A., BULATOV, A., AND JEAVONS, P. Quantified constraints and surjective polymorphisms. Tech. Rep. PRG-RR-02-11, Oxford University, 2002.
- [6] BREWSTER, R., FEDER, T., HELL, P., HUANG, J., AND MACGILLIVRAY, G. Near-unanimity functions and varieties of graphs. 2003.
- [7] BULATOV, A. A dichotomy theorem for constraints on a three-element set. In *FOCS'02* (2002).
- [8] BULATOV, A., KROKHIN, A., AND JEAVONS, P. Constraint satisfaction problems and finite algebras. In *Proceedings 27th International Colloquium on Automata, Languages and Programming, ICALP'00* (2000), vol. 1853 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 272–282.

- [9] 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.
- [10] CHEN, H. Collapsibility and consistency in quantified constraint satisfaction. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence* (2004).
- [11] CHEN, H. *The Computational Complexity of Quantified Constraint Satisfaction*. PhD thesis, Cornell University, August 2004.
- [12] CHEN, H. Existentially restricted quantified constraint satisfaction. Tech. Rep. cs.CC/0506059, ACM Computing Research Repository, 2005.
- [13] CREIGNOU, N., AND HERMANN, M. Complexity of generalized satisfiability counting problems. *Information and Computation* 125 (1996).
- [14] CREIGNOU, N., KHANNA, S., AND SUDAN, M. *Complexity classifications of Boolean Constraint Satisfaction Problems*. SIAM Monographs on Discrete Mathematics and Applications 7. 2001.
- [15] DALMAU, V. Some dichotomy theorems on constant-free quantified boolean formulas. *Machine Learning* 35:3 (1999).
- [16] EBBINGHAUS, H.-D., AND FLUM, J. *Finite Model Theory*. Perspective in Mathematical Logic. Springer-Verlag, 1995.
- [17] FAGIN, R. Generalized first-order spectra and polynomial-time recognizable sets. In *Complexity and Computation: SIAM-AMS 7* (1974).
- [18] FEDER, T., AND KOLAITIS, P. Closures and dichotomies for quantified constraints. To appear in *SIAM J. Discrete Math.*
- [19] 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).

- [20] GRADEL, E. Model checking games. *Electronic Notes in Theoretical Computer Science* 67 (2002).
- [21] GRADEL, E., AND MCCOLM, G. Hierarchies in transitive closure logic, stratified datalog and infinitary logic. *Annals of pure and applied logic* 77 (1996), 166–199.
- [22] GROHE, M. Existential least fixed-point logic and its relatives. *Journal of Logic and Computation* 7 (1997).
- [23] GROHE, M. Fixed-point logics on planar graphs. In *Logic and Computer Science: IEEE 13* (1998), pp. 6–15.
- [24] GUREVICH, Y. Logic and the challenge of computer science. In *Current trends in Theoretical Computer Science* (1988).
- [25] HELL, P., AND NEŠETŘIL, J. On the complexity of H-coloring. *J. Combin. Theory Ser. B* 48 (1990).
- [26] HELL, P., AND NEŠETŘIL, J. *Graphs and Homomorphisms*. OUP, 2004.
- [27] HESMAPANDRA, E. Dichotomy theorems for alternation-bounded quantified boolean formulas. Tech. Rep. cs.CC/0306134, ACM Computing Research Repository, 2003.
- [28] IMMERMANN, N. Relational queries computable in polynomial time. *Information and Control* 68 (1986).
- [29] IMMERMANN, N. Languages that capture complexity classes. *SIAM Journal of Computing* 16:4 (1987).
- [30] IMMERMANN, N. Nondeterministic space is closed under complementation. *SIAM Journal of Computing* 17:5 (1988).
- [31] IMMERMANN, N. *Descriptive Complexity*. Graduate Texts in Computer Science. Springer, 1998.

- [32] JEAVONS, P., COHEN, D., AND GYSSENS, M. A unifying framework for tractable constraints. In *Proceedings 1st International Conference on Constraint Programming, CP'95* (1995), vol. 976 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 276–291.
- [33] JEAVONS, P., COHEN, D., AND GYSSENS, M. A test for tractability. In *Proceedings 2nd International Conference on Constraint Programming—CP'96 (Boston, August 1996)* (1996), vol. 1118 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 267–281.
- [34] KOLAITIS, P. G., AND VARDI, M. Y. A game-theoretic approach to constraint satisfaction. In *AAAI-2000/IAAI-2000 Proceedings* (2000).
- [35] LADNER, R. E. On the structure of polynomial time reducibility. *J. ACM* 22 (1975), 155–171.
- [36] OTTO, M. *Bounded Variable Logics and Counting*. Lecture Notes in Logic 9. Springer-Verlag, 1997.
- [37] PAPADIMITRIOU, C. *Computational Complexity*. Addison-Wesley, 1994.
- [38] SCHAEFER, T. The complexity of satisfiability problems. In *STOC* (1978).
- [39] STEINBERG, R. Generators for simple groups. *Canad. J. Math.* 14 (1962), 277–283.
- [40] STEWART, I. Logical description of monotone NP problems. *Journal of Logic and Computation* 4 (1994), 337–357.
- [41] STEWART, I. Program schemes, arrays, lindstrom quantifiers and zero-one laws. *Theoretical Computer Science* 275 (2002), 283–310.
- [42] STEWART, I. Using program schemes to logically capture polynomial-time on certain classes of structures. *LMS Journal of Computation and Mathematics* 6 (2003), 40–67.
- [43] STOCKMEYER, L. The polynomial-time hierarchy. *Theoretical Computer Science* 3 (1977).