The Power of Counting Logics on Restricted Classes of Finite Structures*

Anuj Dawar[1] and David Richerby[2]**

[1] University of Cambridge Computer Laboratory, William Gates Building, J.J. Thomson Avenue, Cambridge, CB3 0FD, UK.
E-mail: Anuj.Dawar@cl.cam.ac.uk .

[2] Department of Mathematics, University of Athens, Panepistimioupolis, GR157-84 Athens, Greece. E-mail: davidr@math.uoa.gr .

Abstract. Although Cai, Fürer and Immerman have shown that fixed-point logic with counting (IFP + C) does not express all polynomial-time properties of finite structures, there have been a number of results demonstrating that the logic does capture **P** on specific classes of structures. Grohe and Mariño showed that IFP + C captures **P** on classes of structures of bounded treewidth, and Grohe showed that IFP + C captures **P** on planar graphs. We show that the first of these results is optimal in two senses. We show that on the class of graphs defined by a non-constant bound on the tree-width of the graph, IFP + C fails to capture **P**. We also show that on the class of graphs whose local tree-width is bounded by a non-constant function, IFP + C fails to capture **P**. Both these results are obtained by an analysis of the Cai–Fürer–Immerman (CFI) construction in terms of the treewidth of graphs, and cops and robber games; we present some other implications of this analysis. We then demonstrate the limits of this method by showing that the CFI construction cannot be used to show that IFP + C fails to capture **P** on proper minor-closed classes.

1 Introduction

The central open problem in descriptive complexity theory is whether there exists a logic that can express exactly the polynomial-time decidable properties of unordered structures. For some time it was conjectured that the extension of fixed-point logic with counting (IFP + C) would be such a logic, but this was shown not to be the case by a construction due to Cai, Fürer and Immerman [4], which we refer to below as the CFI construction. Nonetheless, IFP + C provides a natural level of expressiveness within the complexity class **P** which has

* Research partly carried during the programme 'Logic and Algorithms' at the Isaac Newton Institute for Mathematical Sciences, University of Cambridge, 20 Clarkson Road, Cambridge, CB3 0EH, UK.

** Funded by the European Social Fund and Greek National Resources — (EPEAEK II) PYTHAGORAS II.

been explored in its own right [20]. It has also been shown that, on certain restricted classes of structures, IFP + C is indeed powerful enough to express all polynomial-time properties. In particular, Immerman and Lander have shown that IFP + C defines exactly the polynomial-time properties of trees [16] and Grohe and Mariño [14] extended this to show that on any class of structures of bounded tree-width, IFP + C captures **P**. Grohe also showed that IFP + C captures **P** on the class of planar graphs [12] and, more generally, on classes of embeddable graphs [13]. In particular, these results imply that the CFI construction cannot be carried out when restricting ourselves to such classes of structures.

There is a growing body of work studying the finite model theory of restricted classes of structures, where the restrictions are essentially borrowed from graph structure theory. Such graph-theoretic restrictions, such as bounding the tree-width of a graph or restricting to planar graphs (or, more generally, proper minor-closed classes of graphs), often yield classes with good algorithmic properties and there has been an effort to explore whether these also correspond to interesting model-theoretic properties which may be tied to the good algorithmic behaviour. In many cases, the logical or model-theoretic view provides a clean general "explanation" of the algorithmic properties of a class. Examples of such meta-theorems are Courcelle's theorem [5], which shows that any property definable in monadic second-order logic is decidable in linear time on classes of bounded tree-width, and the result of Dawar *et al.* [7] that first-order definable optimization problems admit polynomial-time approximation schemes on proper minor-closed classes. At the same time, ever more expansive (i.e., less restricted) classes of structures have been studied such as classes of bounded local tree-width [11], classes that locally exclude a minor [6] and classes of bounded expansion [18]. Our aim in this paper is to explore the boundary of the classes where IFP + C captures **P**. In particular, we wish to determine on which of these various classes the CFI construction can be carried out.

The CFI construction relies on the fact that every formula of IFP + C is equivalent to one of $\mathcal{C}^{\omega}_{\infty\omega}$, the infinitary logic with counting. A separator of a graph $G = (V, E)$ is a set $S \subseteq V$ of vertices whose deletion from the graph leaves no connected component with more than $|V|/2$ vertices. Cai *et al.* show that for each graph G, we can construct two graphs $X(G)$ and $\widetilde{X}(G)$ such that, if G has no separator of size k, then $X(G)$ and $\widetilde{X}(G)$ cannot be distinguished by any formula of $\mathcal{C}^{k}_{\infty\omega}$, the k-variable fragment of $\mathcal{C}^{\omega}_{\infty\omega}$. Since $X(G)$ and $\widetilde{X}(G)$ are distinguished by a polynomial-time algorithm, it follows that IFP + C does not capture **P** on any class of graphs that includes both $X(G)$ and $\widetilde{X}(G)$ for graphs G with arbitrarily large minimal size separators. As Cai *et al.* already noted, this includes the class of graphs with degree bounded by 3. We show (in Section 3) that the assumption that G has no separator of size k can be replaced by the weaker requirement that the tree-width of G is at least k. This is established by a game construction that combines the cops-and-robber game of Seymour and Thomas [24] with the bijection game of Hella [15] (see [1] for another application of the same idea).

An immediate consequence is that IFP + C does not capture **P** on any class of graphs that includes $X(G)$ and $\widetilde{X}(G)$ for graphs of unbounded tree-width. As a corollary, we show that the result of Grohe and Mariño is, in a sense, optimal. For a function $f : \mathbb{N} \to \mathbb{N}$, let TW_f be the class of all graphs G such that the tree-width of G is at most $f(|G|)$. Grohe and Mariño show that, if f is bounded above by a constant, then IFP + C captures **P** on TW_f. On the other hand, we show that, no matter how slowly f grows, if it is unbounded, then IFP + C does not capture **P** on TW_f. Note that this does not show that IFP + C fails to capture **P** on *any* class of graphs of unbounded tree-width. Indeed, planar graphs have unbounded tree-width but IFP + C capture **P** on this class. However, if the class contains all graphs of tree-width bounded by f, we show that the CFI construction applies.

Instead of restricting tree-width as a function of the order of the graph, we can consider graphs where tree-width grows as a function of the diameter. Recall that the r-neighbourhood of a vertex v in a graph is the subgraph induced by the vertices within distance r of v. For a non-decreasing function $f : \mathbb{N} \to \mathbb{N}$, let LTW_f be the class of graphs G such that, for all $r \geqslant 1$, the r-neighbourhood of every vertex in G has tree-width at most $f(r)$. (Eppstein introduces these classes as graphs with the 'diameter-treewidth property' [10] and the restriction is termed bounded local tree-width in [11]). For any such graph, we have $\mathrm{tw}(G) \leqslant f(|G|)$ so $\mathrm{LTW}_f \subseteq \mathrm{TW}_f$. We show, in Section 4 that, analogous to the case for global tree-width, IFP + C captures **P** on LTW_f if, and only if, $f = \mathcal{O}(1)$. Thus, the result of Grohe and Mariño is optimal in a stronger sense.

Grohe [13] has conjectured that IFP + C captures **P** on any proper minor-closed class of finite graphs. We show, in Section 5 that the CFI construction cannot be used to refute this conjecture. That is, we show that for any graph G and any graph H of sufficient tree-width G is a minor of both $X(H)$ and $\widetilde{X}(H)$. Thus, if a class of graphs forbids G as a minor, it excludes $X(H)$ and $\widetilde{X}(H)$ for all graphs H except those of some fixed tree-width.

There are several generalizations of the concept of tree-width to directed graphs including that of directed tree-width [17], DAG-width [2, 19] and entanglement [3]. In each of these measures, the class of directed acyclic graphs (DAGs) has width 1. Since the CFI construction works in the class of DAGs (see Section 3) it follows that our results do not extend to these measures.

Acknowledgment. We would like to thank Stephan Kreutzer for fruitful discussions, especially in connection with the construction in Section 5.

2 Background

The notion of a relational structure $\mathfrak{A} = (A, R_1^{\mathfrak{A}}, \ldots, R_s^{\mathfrak{A}})$ over vocabulary $\sigma = (R_1^{r_1}, \ldots, R_s^{r_s})$ is standard. All structures and graphs in this paper are finite; we treat graphs as structures with a single binary relation symbol E, interpreted by an irreflexive relation that, in the case of undirected graphs, is also symmetric. All graphs mentioned in this paper are undirected unless specifically stated to be directed.

2.1 Counting logics

IFP + C is the extension of first-order logic with inflationary fixed-points and a mechanism for counting. For formal definitions, which we will not need in this paper, we refer the reader to [9]. It is known that every class of structures definable in IFP + C is decidable in polynomial time.

The formulas of the logic $C_{\infty\omega}$ are obtained from the atomic formulas using negation, infinitary conjunction and disjunction, and counting quantifiers ($\exists^i x\, \varphi$ for every integer $i \geqslant 0$). The fragment $\mathcal{C}^k_{\infty\omega}$ consists of those formulas of $C_{\infty\omega}$ in which only k distinct variables appear and $\mathcal{C}^\omega_{\infty\omega} = \bigcup_{k\in\omega} \mathcal{C}^k_{\infty\omega}$. The significance of $\mathcal{C}^\omega_{\infty\omega}$ for our purposes is every formula of IFP + C is equivalent to one of $\mathcal{C}^\omega_{\infty\omega}$.

Hella shows that definability in $\mathcal{C}^k_{\infty\omega}$ is characterized by the *k-pebble bijection game* [15]. The game is played on structures $\mathfrak{A}$ and $\mathfrak{B}$ by two players, the spoiler and the duplicator, using pebbles $a_1, \dots, a_k$ on $\mathfrak{A}$ and $b_1, \dots, b_k$ on $\mathfrak{B}$. If $|A| \neq |B|$, the spoiler wins immediately; otherwise, each move is made as follows:

- the spoiler chooses a pair of pebbles a_i and b_i;
- the duplicator chooses a bijection $h : A \to B$ such that for pebbles a_j and b_j $(j \neq i)$, $h(a_j) = b_j$; and
- the spoiler chooses $a \in A$ and places a_i on a and b_i on $h(a)$.

If, after this move, the map $a_1 \dots a_k \mapsto b_1 \dots b_k$ is not a partial isomorphism $\mathfrak{A} \to \mathfrak{B}$, the game is over and the spoiler wins; the duplicator wins all infinite plays. Hella shows that the duplicator has a winning strategy in the k-pebble bijection game on $\mathfrak{A}$ and $\mathfrak{B}$ if, and only if, the two structures agree on every formula of $\mathcal{C}^k_{\infty\omega}$, in which case, we write $\mathfrak{A} \equiv^{\mathcal{C}^k_{\infty\omega}} \mathfrak{B}$. In order to show that a class Q of structures is not definable in $\mathcal{C}^\omega_{\infty\omega}$ (and, hence, not definable in IFP + C), it suffices to demonstrate, for each $k \geqslant 1$, structures $\mathfrak{A}_k \in Q$ and $\mathfrak{B}_k \notin Q$ on which the duplicator has a winning strategy in the k-pebble bijection game.

By a result of Otto [20, Theorem 4.22] we have the following:

Theorem 1 (Otto). *If* IFP + C *captures* **P** *on a class $\mathcal{C}$ of structures that is closed under disjoint unions, then there is a k such that $\equiv^{\mathcal{C}^{\mathrm{k}}_{\infty\omega}}$ coincides with isomorphism on $\mathcal{C}$.*

Thus, to show that IFP + C does not capture **P** on a class of structures $\mathcal{C}$, it suffices to show that for every k $\mathcal{C}$ contains a pair of non-isomorpic structures H and H', such that $H \equiv^{\mathcal{C}^{\mathrm{k}}_{\infty\omega}} H'$.

2.2 Tree-width

Tree-width was introduced by Robertson and Seymour [21] as a key component of the proof of the Graph Minor Theorem. A *tree decomposition* of a graph $G = (V, E)$ is a pair $(T, \{ B_t : t \in T \})$, where T is a tree, $B_t \subseteq V$ and

- $\bigcup_{t\in T} B_t = V$;
- if there is an edge $uv \in E$ then $\{ u, v \} \subseteq B_t$ for some t; and
- for each $v \in V$, the set $\{ t : v \in B_t \}$ is connected in T.

The *width* of a tree decomposition is $\max\{\,|B_t| : t \in T\,\} - 1$ and the *tree-width* of G is $\mathrm{tw}(G)$, the least k for which G has a tree decomposition of width k.

We will use the following game-theoretic characterization of tree width, due to Seymour and Thomas [24]. The k cops and robber game is played by two players, the cops and the robber, on a graph $G = (V, E)$. At each move, the cops player either removes a cop from the graph or takes a cop not currently on the graph and places him on some vertex v. The robber may then move along any cop-free path in the graph. If the cops' move was to place a cop on v, that vertex counts as cop-free for this turn. If a cop moves to the vertex occupied by the robber and the robber has no non-trivial legal move, the cops win; the robber wins if he can stay on the run indefinitely. Seymour and Thomas show that the cops have a winning strategy in the k cops and robber game on G if, and only if, $\mathrm{tw}(G) \leqslant k - 1$.

For a positive integer k, we write TW_k for the class of graphs of tree width at most k. For a function $f : \mathbb{N} \to \mathbb{N}$, TW_f denotes the class of all graphs G such that the tree-width of G is at most $f(|G|)$.

Given a graph G and $r \geqslant 0$, for each vertex $v \in G$, let $N^r_G(x)$ be the subgraph of G induced by the vertices at distance at most r from v. The *local tree-width* of a graph [11] is the function

$$\mathrm{ltw}_G(r) = \max\{\,\mathrm{tw}(N^r_G(v)) : v \in G\,\}\,.$$

A class $\mathcal{G}$ of graphs is said to have *bounded local tree-width* if there is a (non-decreasing) function $f : \mathbb{N} \to \mathbb{N}$ such that, for all $G \in \mathcal{G}$ and all $r \geqslant 0$, $\mathrm{ltw}_G(r) \leqslant f(r)$. Classes of graphs of bounded local tree-width are introduced by Eppstein [10], who refers to such classes as having the 'diameter-treewidth property'.

2.3 Graph Minors

We say that a graph G is a *minor* of H, and write $G \preccurlyeq H$, if there is a map that associates with each vertex v of G a non-empty, connected subgraph H_v of H such that H_u and H_v are disjoint for $u \neq v$ and if there is an edge between u and v in G then there is an edge in H between some vertex in H_u and some vertex in H_v. We refer to the sets H_v as the branch sets witnessing that G is a minor of H. An equivalent characterization (see [8]) states that G is a minor of H if G can be obtained from a subgraph of H by contracting edges. The contraction of an edge consists of identifying its two endpoints into a single vertex and removing the resulting loop.

We collect here a few facts about graph minors that we will need. All of these can be found in [8]. Note that if $G \preccurlyeq H$ then $\mathrm{tw}(G) \leqslant \mathrm{tw}(H)$. By the well-known Kuratowski–Wagner theorem, a graph G is planar if, and only if, neither K_5 nor $K_{3,3}$ is a minor of G. Robertson and Seymour [23] showed that any class of graphs that is closed under taking minors and is not the class of all graphs is characterized by a finite set of forbidden minors. We call such a class of graphs a *proper minor-closed class.*

For any $n > 1$, let G_n be the $n \times n$ grid graph, i.e., the graph with vertex set $\{\,1, \ldots, n\,\}^2$ and all edges of the form $\{\,(i, j), (i + 1, j)\,\}$ and $\{\,(i, j), (i, j + 1)\,\}$.

$\Delta(G_n) = 4$ (where $\Delta(G)$ denotes the maximum degree of any vertex in G) and it is easy to see from the cops and robber game that $\mathrm{tw}(G_n) = n$. Also, it can be shown that for any planar graph G, there is an n such that $G \preccurlyeq G_n$.

2.4 CFI graphs

The graphs we describe in this section are a minor variation on the graphs used by Cai *et al.* to separate IFP + C from **P** [4] and proofs of all the results we quote here can be found there. The difference is that Cai *et al.* do not have the 'c' and 'd' vertices but, instead, colour the vertices on the understanding that the colours can be replaced by appropriate gadgets. The gadgets are simple and we will use some of their properties later on so we prefer to make them explicit.

Let $G = (V, E)$ be a graph in which every vertex has degree at least two. In the following discussion, we will assume that G is connected but there are easy component-wise extensions in the case where G is not connected. For each $v \in V$ let $\Gamma(v) = \{\, u : uv \in E \,\}$ and let $\widehat{v}$ be the set of new vertices,

$$\begin{aligned}\widehat{v} = \{\, & a_{vw}, b_{vw}, c_{vw}, d_{vw} : w \in \Gamma(v) \,\} \\ & \cup \{\, v^X : X \subseteq \Gamma(v) \text{ and } |X| \equiv 0 \pmod 2 \,\}\,.\end{aligned}$$

Call the v^X *inner vertices* and the other members of $\widehat{v}$ *outer vertices*. Let $X_\emptyset(G)$ be the graph with vertices $\bigcup_{v \in V} \widehat{v}$ and edges as follows:

- edges $a_{vw}c_{vw}$, $b_{vw}c_{vw}$ and $c_{vw}d_{vw}$ for each edge $vw \in E$;
- an edge $a_{vw}v^X$ whenever $w \in X$;
- an edge $b_{vw}v^X$ whenever $w \in \Gamma(v) \setminus X$; and
- edges $a_{vw}a_{wv}$ and $b_{vw}b_{wv}$ for each edge $vw \in E$.

The subgraph of $X_\emptyset(G)$ induced by $\widehat{v}$ for a vertex v of G with three neighbours w_1, w_2, w_3 is illustrated in Fig. 1, where the dashed lines indicate edges connecting this subgraph to the rest of $X_\emptyset(G)$.

For any $S \subseteq E$, let $X_S(G)$ be $X_\emptyset(G)$ with the edges $a_{vw}a_{wv}$ and $b_{vw}b_{wv}$ deleted and edges $a_{vw}b_{wv}$ and $b_{vw}a_{wv}$ added, for every edge $vw \in S$. We say that the edges in S have been *twisted*. Cai *et al.* show that $X_S(G) \cong X_T(G)$ if, and only if, $|S| \equiv |T| \pmod 2$. This being the case, we write $X(G)$ for the graph $X_\emptyset(G)$ and write $\widetilde{X}(G)$ for $X_{\{\,e\,\}}(G)$ for any edge e and call these, respectively, the *untwisted* and *twisted CFI graphs of* G.

For distinct edges e and f of G, we can obtain an isomorphism between $X_{\{\,e\,\}}(G)$ and $X_{\{\,f\,\}}(G)$ as follows. Note that, for each $v \in V$ and $N \subseteq \Gamma(v)$ with $|N| \equiv 0 \pmod 2$, there is an automorphism $\eta_{v,N}$ of the subgraph of $X_S(G)$ induced by $\widehat{v}$ that exchanges a_{vw} and b_{vw}, for each $w \in N$ (and there is no such automorphism if $|N| \equiv 1 \pmod 2$). Let e be the edge uv and f be the edge xy. If $v = x$, then the required isomorphism is just the map $\eta_{v,\{\,u,y\,\}}$. Otherwise, if the four vertices are distinct then, by the assumption that G is connected, there is a simple path from one endpoint of e to an endpoint of f that does not pass through the other endpoints. Without loss of generality let $v_1 \ldots v_\ell$ be a simple

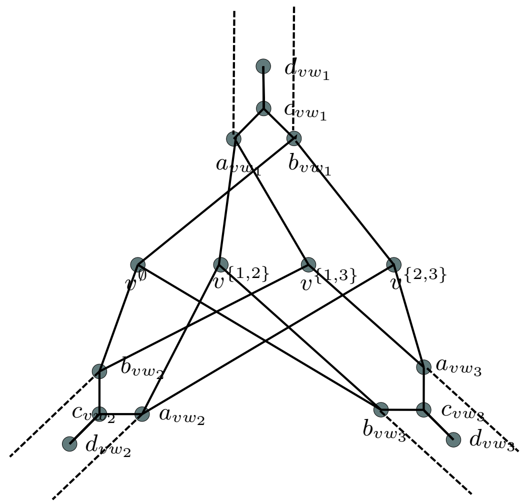

Fig. 1. The graph on the vertices $\widehat{v}$ in $X_{\emptyset}(G)$.

path with $v = v_1$ and $x = v_\ell$ such that neither u nor y occurs on the path. Then, the required isomorphism from $X_{\{e\}}(G)$ to $X_{\{f\}}(G)$ is

$$\eta = \eta_{v_1,\{u,v_2\}} \circ \eta_{v_2,\{v_1,v_3\}} \circ \cdots \circ \eta_{v_{\ell-1},\{v_{\ell-2},v_k\}} \circ \eta_{v_\ell,\{v_{\ell-1},y\}} .$$

3 Tree-width

A *separator* of a graph $G = (V, E)$ is a set $S \subseteq V$ such that every connected component of $G - S$ has at most $|V|/2$ vertices. Cai, Fürer and Immerman prove that, if G is connected, has minimal degree $\delta(G) \geqslant 2$ and has no separator of size k, then $X(G) \equiv^{\mathcal{C}^k_{\infty\omega}} \widetilde{X}(G)$. It follows that IFP + C does not define all polynomial-time queries on graphs, for instance by Theorem 1.

The following lemma on the relationship between separators and tree-width is a special case of [21, Theorem 2.5].

Lemma 2. *Every graph G of tree-width k has a separator of size at most $k+1$.*

So, any graph that has no separator of size k must have tree-width at least k. On the other hand, for all k there are connected graphs of tree-width k that have separators of size one (and, of course, disconnected graphs of tree-width k with $\emptyset$ as a separator): take any order-n graph G of tree-width k, choose a vertex $v \in G$ and add n new vertices and an edge from each new vertex to v. The resulting graph still has tree-width k but $\{v\}$ is a separator. Therefore, requiring that G have tree-width at least $k - 1$ is weaker than requiring it to have no separator of size k.

Theorem 3. *Let G be any connected graph with $\delta(G) \geqslant 2$ and $\mathrm{tw}(G) \geqslant k$. $X(G) \equiv^{\mathcal{C}^k_{\infty\omega}} \widetilde{X}(G)$.*

Proof. We exhibit a winning strategy for the duplicator in the k-pebble bijection game on $X(G)$ and $\widetilde{X}(G)$.

Given $u, v \in V(G)$, let $\sigma[u, v]$ be the permutation of $V(G)$ that exchanges u and v and fixes every other vertex. For a vertex u of G, we say that a bijection $h : X(G) \to \widetilde{X}(G)$ is *good except at* u if it satisfies the following conditions:

- for every vertex v of G, $h\widehat{v} = \widehat{v}$;
- h maps inner vertices to inner vertices and outer vertices to outer vertices;
- h is an isomorphism between the graphs $X(G) \setminus u_{\mathrm{I}}$ and $\widetilde{X}(G) \setminus u_{\mathrm{I}}$, where u_{I} is the set of inner vertices in $\widehat{u}$; and
- for every pair a_{uv}, b_{uv} in $\widehat{u}$, $h \circ \sigma[a_{uv}, b_{uv}]$ is an isomorphism from $X(G)[\widehat{u}]$ to $\widetilde{X}(G)[\widehat{u}]$, where $X(G)[\widehat{u}]$ is the subgraph of $X(G)$ induced by $\widehat{u}$.

For concreteness, say $\widetilde{X}(G)$ is the graph $X_{\{uv\}}(G)$. Then $\sigma[a_{uv}, b_{uv}]$ is a bijection that is good except at u; similarly, $\sigma[a_{vu}, b_{vu}]$ is good except at v. Note that if η is an automorphism of $\widetilde{X}(G)$ that fixes $\widehat{v}$ (set-wise) for every $v \in G$ and h is a bijection that is good except at u, then $h \circ \eta$ is also good except at u. We claim that, if h is a bijection that is good except at u and there is a simple path P from u to v, then there is a bijection h' that is good except at v such that for all vertices w not in P and all $x \in \widehat{w}$, $h'(x) = h(x)$.

To prove the claim, let the path P be $v_1 \dots v_\ell$ with $u = v_1$ and $v = v_\ell$. Let η_P be the permutation

$$\eta_P = \sigma[a_{uv_1}, b_{uv_1}] \circ \eta_{v_2,\{v_1,v_3\}} \circ \cdots \circ \eta_{v_{l-1},\{v_{l-2},v_l\}} \circ \sigma[a_{v,v_{l-1}}, b_{v,v_{l-1}}] \,.$$

The properties of the graphs $X(G)$ and $\widetilde{X}(G)$ then ensure that taking $h' = h \circ \eta_P$ satisfies the claim.

The duplicator's strategy can now be described as always playing a bijection that is good except at u for some u. The vertex u is given by the position of the robber in the cops and robber game played on G where the positions of the cops are $v_1 \dots v_k$ when the pebbles in the bijection game are in the sets $\widehat{v}_1 \dots \widehat{v}_k$.

Initially, the duplicator plays a bijection that is good except at u for one of the endpoints u of the twisted edge in $\widetilde{X}(G)$. At the same time, she initiates a cops and robber game on G with the robber initially at u. At each subsequent move, when the spoiler moves one of the pebbles, the duplicator moves the corresponding cop in the cops and robber game. This yields a path P for the robber to move along to a vertex v. By the claim, this yields a bijection h' that the duplicator can play. Since P, by definition, does not go through any of the cop positions, this means that h' agrees with h on all currently pebbled positions in the bijection game as required. Also, since h is, at all times, an isomorphism everywhere except at the inner vertices of $\widehat{v}$, for the current robber position v, it follows that it must be a partial isomorphism on the pebbled vertices, as required. □

The following corollary is immediate from Theorem 3 and Theorem 1.

Corollary 4. IFP + C *does not capture* **P** *on any class of graphs containing* $X(G)$ *and* $\widetilde{X}(G)$ *for graphs* G *of unbounded tree width.*

Let $\Delta(G)$ be the maximal degree of any vertex in a given graph G. The following lemma shows that $\mathrm{tw}(X(G))$ can be bounded in terms of tree-width and maximal degree of G. The same bounds apply to $\mathrm{tw}(\widetilde{X}(G))$.

Lemma 5. $\mathrm{tw}(G) \leqslant \mathrm{tw}(X(G)) \leqslant (4\Delta(G) + 2^{\Delta(G)-1})\mathrm{tw}(G)$.

Proof. The first inequality holds because $G \preccurlyeq X(G)$: contracting along every edge in each $\widehat{v}$ in $X(G)$ gives G. The second holds because if $(T, \{\, B_t : t \in T \,\})$ is a tree decomposition of G then $(T, \{\, (\bigcup_{v \in B_t} \widehat{v}) : t \in T \,\})$ is a tree decomposition of $X(G)$. □

For any function $f : \mathbb{N} \to \mathbb{N}$, let $\mathrm{TW}_f = \{\, G : \mathrm{tw}(G) \leqslant f(|G|) \,\}$. Grohe and Mariño show that, if f is any constant function, IFP + C captures **P** on TW_f [14]. We show that this result is, in a sense, optimal: if f is unbounded then IFP + C does not capture **P** on TW_f.

Recall that G_n is the $n \times n$ grid graph and that $\mathrm{tw}(G_n) = n$. For $n \geqslant 3$, $\Delta(G_n) = 4$.

Theorem 6. IFP + C *captures* **P** *on* TW_f *if, and only if,* $f = \mathcal{O}(1)$.

Proof. The 'if' direction is the Grohe–Mariño theorem. Conversely, if $f \neq \mathcal{O}(1)$ then, for any n, there is some $k > |X(G_n)|$ such that $f(k) \geqslant 24n$. Since $\mathrm{tw}(X(G_n)) \leqslant 24n$ (Lemma 5), it follows that TW_f contains $X(G)$ and $\widetilde{X}(G)$ (possibly padded with some number of isolated vertices) for graphs of arbitrary tree width and so, by Corollary 4, IFP + C does not capture **P** on TW_f. □

Notice that we do not claim that IFP + C fails to capture **P** on any class of graphs containing graphs of unbounded tree-width. For example, the complete graph on n vertices has tree-width $n - 1$ so the class of all complete graphs contains graphs of arbitrarily high tree-width but IFP + C does capture **P** on this class. Similarly, the class of planar graphs contains graphs of unbounded tree-width (it contains G_n for all n), but Grohe has shown that IFP + C captures **P** on this class [12]. However, if $\mathcal{C}$ contains all graphs whose tree-width is bounded by the function f, then the CFI construction applies.

Lemma 7. *If* G *is bipartite,* $X(G)$ *and* $\widetilde{X}(G)$ *are bipartite.*

Proof. Let $G = (V, E)$ with bipartition V_0, V_1. Then $X(G)$ and $\widetilde{X}(G)$ have bipartition W_0, W_1, where W_i consists of all inner and 'c' vertices corresponding to elements of V_i and all 'a', 'b' and 'd' vertices from V_{1-i}. □

Corollary 8. IFP + C *does not capture* **P** *on the class of bipartite graphs.*

Proof. For any n, G_n is bipartite so, by Lemmata 5 and 7, the class of bipartite graphs contains $X(G)$ and $\widetilde{X}(G)$ for graphs of unbounded tree width. □

Moreover, IFP + C does not capture **P** on the class of graphs of chromatic number k, for any fixed $k \geqslant 2$, as the disjoint union of $X(G_n)$ and K_k has chromatic number k but no formula of IFP + C can distinguish it from $\widetilde{X}(G_n) \cup K_k$ if n is large enough.

One might hope that Theorem 6 could be extended to measures of graph connectivity on directed graphs such as directed tree-width [17], DAG-width [2] or entanglement [3] but this is not the case. All directed acyclic graphs (DAGs) have low width in all of these measures (directed tree-width zero, DAG-width one and entanglement zero) but there are polynomial-time queries on DAGs not definable in IFP + C.

Theorem 9. IFP + C *does not capture* **P** *on the class of DAGs.*

Proof. Let D be a directed graph. Define $X'(D)$ and $\widetilde{X}'(D)$ in the same way as for undirected graphs but with the following directions on the edges:

- edges from inner vertices to outer vertices are directed that way;
- edges between outer vertices in the same $\widehat{v}$ are directed ac, bc and dc;
- any edges between a_{uv} or b_{uv}, and a_{vu} or b_{vu} have the same direction as the corresponding edge in D between u and v.

Note that, if D contains edges uv and vu then $\widehat{u}$ will contain two sets of outer vertices associated with v: one for each edge. Observe that $X'(D)$ and $\widetilde{X}'(D)$ are DAGs. Clearly, there is a polynomial-time algorithm that distinguishes $X'(D)$ from $\widetilde{X}'(D)$ — just forget the orientation of the edges and use the algorithm that distinguishes $X(G)$ from $\widetilde{X}(G)$. Suppose the query $\{ X'(D) : D \text{ is a DAG} \}$ is defined by some sentence $\varphi \in$ IFP + C for DAGs.

Fix any vertex $v \in G_n$. There are no edges in G_n between vertices at the same distance from v so the orientation $D(G_n, v)$ of G_n that orients every edge from its end further from v to the end nearer v is a DAG. There is an IFP formula $\psi(xy, v)$ that, given some vertex $v \in X(G_n)$ (respectively, $\widetilde{X}(G_n)$) as a parameter, defines the edge relation of $X'(D(G_n, v))$ (respectively, $\widetilde{X}'(D(G_n, v))$). Let $\chi \equiv \exists v\, \varphi[\psi(xy, v)/E(xy)]$, where $\varphi[\cdots]$ is the result of replacing every subformula $E(xy)$ with $\psi(xy, v)$. Then χ distinguishes $X(G_n)$ from $\widetilde{X}(G_n)$ for all n, contradicting Theorem 3. □

Corollary 8 and Theorem 9 are to be expected. The relation $\equiv^{\mathcal{C}^k_{\infty\omega}}$ can be tested in polynomial time by means of a colour-refinement algorithm (see, e.g., [20]). Therefore, by Theorem 1, it follows that if, IFP + C captures **P** on a class of structures $\mathcal{C}$ (closed under disjoint unions), then $\mathcal{C}$ admits a polynomial-time isomorphism test. It is not difficult to see that bipartite graphs and DAGs admit such a test if, and only if, all graphs do. Indeed, given an undirected graph $G = (V, E)$, let G' be the directed graph whose vertex set is $V \cup E$ and with a directed edge from $v \in V$ to $e \in E$ exactly when v is one of the ends of e in G. Clearly, G is acyclic and $G \cong H$ if, and only if, $G' \cong H'$. The undirected version of G' (known as the *incidence graph* of G) is bipartite.

4 Bounded local tree-width

Given a non-decreasing function $f : \mathbb{N} \to \mathbb{N}$, let LTW_f be the class of graphs whose local tree-width is bounded by f. In this section, we extend the results of Section 3 to show that IFP + C captures **P** on LTW_f if, and only if, $f = \mathcal{O}(1)$.

For a graph $G = (V, E)$ and a positive integer r, we define the graph $r(G)$ to have vertex set $V \cup \{(u, v, i) : u, v \in V, 1 \leqslant i \leqslant r\}$ and edges:

- $\{ u, (u, v, 1) \}$ for all $u, v \in V$;
- $\{ (u, v, i), (u, v, i + 1) \}$ for all $u, v \in V$ and $1 \leqslant i < r$; and
- $\{ (u, v, r), (v, u, r) \}$ for all edges $uv \in E$.

Recall that a *subdivision* of a graph G is any graph H formed by replacing each edge $uv \in G$ with a u–v path, such that the paths in H corresponding to distinct edges in G are internally disjoint. We can think of $r(G)$ as a graph that is obtained from a subdivision of G (where each edge is replaced by a path of length $2r$) by further adding, for each pair u, v that is not an edge in G, two simple paths of length r — one originating at u and one at v — that do not meet.

The properties of the graphs $r(G)$ that we need are established in the following lemmata.

Lemma 10. $\mathrm{tw}(r(G)) = \mathrm{tw}(G)$.

Proof. $\mathrm{tw}(G) \leqslant \mathrm{tw}(r(G))$ because $G \preccurlyeq r(G)$. For the converse, if $\mathrm{tw}(G) = 1$, then G is a forest and $r(G)$ is a forest as well, so $\mathrm{tw}(H) = 1$, as required. Now, suppose $\mathrm{tw}(G) = k > 1$. Let $(T, \{ B_t : t \in T \})$ be a width-k tree decomposition of G. We construct a width-k tree decomposition of H as follows. For each edge $uv \in G$ there must, by definition, be some $t \in T$ such that $\{ u, v \} \subseteq B_t$. Add to T a path $tt_1 \dots t_{2r}$ and set $B_{t_1} = \{ u, v, (u, v, 1) \}$; for $2 \leqslant i \leqslant r$ set $B_{t_i} = \{ v, (u, v, i - 1), (u, v, i) \}$; $B_{t_{r+1}} = \{ v, (u, v, r), (v, u, r) \}$; and for $r + 2 \leqslant i \leqslant 2r$ set $B_{t_i} = \{ v, (v, u, 2r - i + 2), (v, u, 2r - i + 1) \}$. Finally, if uv is not an edge in G, choose any $t \in T$ such that $u \in B_t$ and add a path $tt_1 \dots t_{r-1}$ to T with $B_{t_i} = \{ u, (u, v, i), (u, v, i + 1) \}$. □

For any vertex w in $r(G)$, we write $\pi w = w$, if $w \in V(G)$, and $\pi(u, v, i) = u$.

Lemma 11. *If* $G \equiv^{\mathcal{C}^k_{\infty\omega}} H$*, then* $r(G) \equiv^{\mathcal{C}^k_{\infty\omega}} r(H)$.

Proof. By the assumption $G \equiv^{\mathcal{C}^k_{\infty\omega}} H$, the duplicator has a winning strategy in the k-pebble bijection game played on these two graphs. This winning strategy is easily adapted to a winning strategy on the pair of graphs $r(G)$ and $r(H)$. For any bijection h between the vertices of G and the vertices of H, consider the extension h_r of h to the vertices of of $r(G)$ given by

$$h_r(w) = \begin{cases} h(w) & \text{if } w \in V^G \\ (h(u), h(v), i) & \text{if } w = (u, v, i). \end{cases}$$

The duplicator's strategy is to maintain the condition that, at any point in the game, if the pebbles are on the vertices $s_1, \ldots, s_k$ in $r(G)$ and $t_1, \ldots, t_k$ in $r(H)$, then $\pi s_1, \ldots, \pi s_k$ and $\pi t_1, \ldots, \pi t_k$ is a winning position in the game played on G and H. It is now easily verified that, if the duplicator's strategy called for playing the bijection h in the latter game, then the bijection h_r will maintain this condition in the game on $r(G)$ and $r(H)$. $\square$

We are now ready to prove the strengthening of Theorem 6.

Theorem 12. *Let $f : \mathbb{N} \to \mathbb{N}$ be any non-decreasing function.* IFP + C *captures* **P** *on* LTW_f *if, and only if,* $f = \mathcal{O}(1)$.

Proof. $\mathrm{LTW}_f \subseteq \mathrm{TW}_f$ so, if $f = \mathcal{O}(1)$, then IFP + C defines all polynomial-time properties over LTW_f by the Grohe–Mariño theorem.

Suppose $f \neq \mathcal{O}(1)$. For any k, let G by a graph with $tw(G) \geqslant k$ and $\delta(G) \geqslant 2$. Now, there is some r such that $f(2r) \geqslant \mathrm{tw}(X(G)) = \mathrm{tw}(\widetilde{X}(G))$. Let $H = r(X(G))$ and $H' = r(\widetilde{X}(G))$. By Lemma 10, $\mathrm{tw}(H) = \mathrm{tw}(H') = \mathrm{tw}(X(G))$. Notice that ltw_H and $\mathrm{ltw}_{H'}$ are bounded by the function

$$h(x) = \begin{cases} 1 & \text{if } x < 2r \\ \mathrm{tw}(X(G)) & \text{otherwise,} \end{cases}$$

and that this function is, in turn, bounded by f. Therefore, $H, H' \in \mathrm{LTW}_f$. Since, $\mathrm{tw}(G) \geqslant k$ we have (by Theorem 3) that $H \equiv^{\mathcal{C}^{\mathrm{k}}_{\infty\omega}} H'$. The result now follows from Theorem 1, $\square$

5 Graph minors

Grohe has conjectured that IFP + C captures **P** on any proper minor-closed class of graphs [13], i.e., any minor-closed class except the class of all graphs. In this section, we show that this conjecture cannot be refuted by the CFI construction. Specifically, we show that any class $\mathcal{C}$ of graphs containing at least one of $X(G)$ and $\widetilde{X}(G)$ for graphs G of unbounded tree-width has no forbidden minors. Since any proper minor-closed class must have at least one forbidden minor, it follows that $\mathcal{C}$ is either the class of all graphs or is not minor-closed. Note that the requirement that $\mathcal{C}$ contain CFI graphs derived from graphs of unbounded tree-width is crucial here: it does not suffice to require merely that $\mathcal{C}$ contain graphs of unbounded tree-width. For example, the class of planar graphs does not have bounded tree-width.

We wish to show that, for any graph G, if H is a graph with $\mathrm{tw}(H)$ large enough, relative to G, then $X(H)$ and $\widetilde{X}(H)$ contain G as a minor. To do this, we will first produce a planar graph G' such that $G' \preccurlyeq H$. The graph G' is obtained from a plane drawing of G by inserting new vertices at crossing points of edges. The assumption on the tree-width of H will ensure that any such planar graph is a minor of H. The paths in G' corresponding to distinct edges in G will be edge-disjoint but not necessarily independent: two of the paths may cross at some vertex. To show that $X(H)$ and $\widetilde{X}(H)$ contain G as a minor, we need to show that, even if edge-disjoint paths P_1 and P_2 meet at a vertex u, $X(H)$ and $\widetilde{X}(H)$ contain corresponding independent paths.

Lemma 13. *Let $u \in G$ be a vertex of degree 4, with neighbours w, x, y, z. For each $v \in \{ w, x, y, z \}$, choose $v' \in \{ a_{uv}, b_{uv} \}$. $\widehat{u}$ contains vertices v_1 and v_2 such that $X(G)$ and $\widetilde{X}(G)$ contain disjoint paths $w'v_1x'$ and $y'v_2z'$.*

Proof. Note that $X(G)$ and $\widetilde{X}(G)$ have the same edges within each $\widehat{u}$. The values of v_1 and v_2 are given in the following table.

w'	x'	y'	z'	v_1	v_2
a_{uw}	a_{ux}	a_{uy}	a_{uz}	$v^{\{ w,x \}}$	$v^{\{ y,z \}}$
a_{uw}	a_{ux}	a_{uy}	b_{uz}	$v^{\{ w,x,y,z \}}$	$v^{\{ x,y \}}$
a_{uw}	b_{ux}	a_{uy}	b_{uz}	$v^{\{ w,z \}}$	$v^{\{ x,y \}}$

The other cases are symmetric, either by permutations of $\{ w, x, y, z \}$ or the automorphisms of $\widehat{u}$ that exchange the 'a' and 'b' vertices for even-cardinality subsets of $\{ w, x, y, z \}$. □

We first restrict attention to minors of CFI graphs of grids. Recall that G_r is the $r \times r$ grid graph.

Theorem 14. *Let G be any graph. For sufficiently large grids G_r, $G \preccurlyeq X(G_r)$ and $G \preccurlyeq \widetilde{X}(G_r)$.*

Proof. Let $V(G) = \{ v_1, \ldots, v_n \}$ and $E(G) = \{ e_1, \ldots, e_m \}$. We first produce a drawing G^* of G. Choose a set $V^* = \{ v_1^*, \ldots, v_n^* \}$ of distinct points in $\mathbb{R}^2$ to represent the vertices of G and a set $E^* = \{ e_1^*, \ldots, e_m^* \}$ of distinct, simple, piecewise-linear curves to represent the edges, such that:

- if $e_i = v_j v_k$ then the endpoints of e_i^* are v_j^* and v_k^*;
- no e_i^* contains any v_j^* except its endpoints;
- for $i \neq j$, $e_i^* \cap e_j^*$ is finite; and
- no point in $\mathbb{R}^2 \setminus V^*$ appears in more than two of the e_i^*.

We can now produce a planar graph G' whose vertices are the points of intersection of the e_i^* (i.e., $V(G') = \bigcup_{1 \leqslant i < j \leqslant m} (e_i^* \cap e_j^*)$, including V) and whose edges are precisely those pairs $\{ x, y \}$ such that G^* contains an x–y curve that passes through no other points in $V(G')$.

Since G' is planar, we have $G' \preccurlyeq G_r$ for large enough r. Unless G is, itself, planar, we cannot have $G \preccurlyeq G_r$; however, we claim that $G \preccurlyeq X(G_r)$ and $G \preccurlyeq \widetilde{X}(G_r)$. To this end, let H be $X(G_r)$ or $\widetilde{X}(G_r)$.

Let $\{ V_x : x \in G' \}$ be the branch sets witnessing that G' is a minor of G_r. For each $x \in G'$, let $\widehat{V}_x = \bigcup_{y \in V_x} \widehat{y} \subseteq V(H)$. To show that $G \preccurlyeq H$, we proceed as follows. First, for each $x \in G$, contract all the edges in the subgraph induced by $\widehat{V}_x$, calling the resulting vertex v_x. We now show that there is a system of independent paths P_{xy}, from v_x to v_y, for each edge $xy \in G$.

For each $xy \in G$, let Q_{xy} be the x–y path in G' corresponding to the edge $xy \in G^*$. These paths are not necessarily independent: in particular, they share the vertices of $V(G') \setminus V(G)$. Lemma 13 shows that, even if Q_{wx} and Q_{yz} meet at vertex u, H contains paths $\widehat{w}$–$\widehat{x}$ and $\widehat{y}$–$\widehat{z}$ that are independent. This completes the proof. □

The significance of grids is given by the following theorem of Robertson and Seymour. (Note that the tree-width required grows rapidly with r: Diestel shows that $r^{4r^4(r+2)}$ suffices [8].)

Theorem 15 ([22]). *For every $r > 1$, every graph of sufficiently high tree-width contains G_r as a minor.*

Theorem 16. *The only minor-closed class of graphs that contains $X(G)$ or $\widetilde{X}(G)$ for graphs G of unbounded tree-width is the class of all graphs.*

Proof. Let $\mathcal{C}$ be a minor-closed class of graphs containing $X(G)$ or $\widetilde{X}(G)$ for graphs G of unbounded tree-width and let H be any graph. By Theorem 14, $X(G_r)$ and $\widetilde{X}(G_r)$ contain H as a minor, for large enough grids G_r. By Theorem 15, any graph of large enough tree-width contains G_r as a minor. Therefore, $\mathcal{C}$ contains a graph $X \in \{ X(G), \widetilde{X}(G) \}$ for some graph G containing G_r as a minor, so $H \preccurlyeq X$. But $\mathcal{C}$ is minor-closed so $H \in \mathcal{C}$. □

A consequence of this theorem is that any attempt to refute Grohe's conjecture that IFP + C captures **P** on all non-trivial minor-closed classes of graphs cannot rely on Cai–Fürer–Immerman graphs. For, to use the CFI construction (in the form in Theorem 3), we need precisely to find for each k, a graph G of tree-width at least k such that $X(G)$ and $\widetilde{X}(G)$ are both in the class.

References

1. A. Atserias, A. Bulatov, and A. Dawar. Affine systems of equations and counting infinitary logic. In *Proc. 34th International Colloquium on Automata, Languages and Programming*, 2007. to appear.
2. D. Berwanger, A. Dawar, P. Hunter, and S. Kreutzer. DAG-width and parity games. In *Proc. 23rd Annual Symposium on Theoretical Aspects of Computer Science*, volume 3884 of *LNCS*, pages 524–536. Springer-Verlag, 2006.
3. D. Berwanger and E. Grädel. Entanglement: A measure for the complexity of directed graphs with applications to logic and games. In *Proc. 11th International Conference on Logic for Programming, Artificial Intelligence and Reasoning*, volume 3452 of *LNCS*, pages 209–223. Springer-Verlag, 2004.
4. J.-Y. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410, 1992.
5. B. Courcelle. Graph rewriting: An algebraic and logic approach. In J. van Leeuwan, editor, *Handbook of Theoretical Computer Science, Volume B: Formal Models and Sematics (B)*, pages 193–242. Elsevier, 1990.

6. A. Dawar, M. Grohe, and S. Kreutzer. Locally excluding a minor. In *Proc. 22nd IEEE Annual Symposium on Logic in Computer Science*. IEEE Computer Society, 2007. To appear.
7. A. Dawar, M. Grohe, S. Kreutzer, and N. Schweikardt. Approximation schemes for first-order definable optimisation problems. In *Proc. 21st IEEE Annual Symposium on Logic in Computer Science*, pages 411–420. IEEE Computer Society, 2006.
8. R. Diestel. *Graph Theory*. Springer-Verlag, 3rd edition, 2005.
9. H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer-Verlag, 2nd edition, 1999.
10. D. Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27:275–291, 2000.
11. M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM*, 48:1184–1206, 2001.
12. M. Grohe. Fixed-point logics on planar graphs. In *Proc. 13th IEEE Annual Symposium on Logic in Computer Science*, pages 6–15. IEEE Computer Society, 1998.
13. M. Grohe. Isomorphism testing for embeddable graphs through definability. In *Proc. 32nd ACM Symposium on Theory of Computing*, pages 63–72. ACM, 2000.
14. M. Grohe and J. Mariño. Definability and descriptive complexity on databases of bounded tree-width. In *Proc. 7th International Conference on Database Theory*, volume 1540 of *LNCS*, pages 70–82. Springer-Verlag, 1999.
15. L. Hella. Logical hierarchies in PTIME. *Information and Computation*, 129(1):1–19, 1996.
16. N. Immerman and E.S. Lander. Describing graphs: A first-order approach to graph canonization. In A. Selman, editor, *Complexity Theory Retrospective*, pages 59–81. Springer-Verlag, 1990.
17. T. Johnson, N. Robertson, P.D. Seymour, and R. Thomas. Directed tree-width. *Journal of Combinatorial Theory Series B*, 82:138–155, 2001.
18. J. Nešetřil and P. Ossona de Mendez. The grad of a graph and classes with bounded expansion. In A. Raspaud and O. Delmas, editors, *International Colloquium on Graph Theory*, volume 22 of *Electronic Notes in Discrete Mathematics*, pages 101–106. Elsevier, 2005.
19. J. Obdržálek. DAG-width: Connectivity measure for directed graphs. In *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 814–821, 2006.
20. M. Otto. *Bounded Variable Logics and Counting — A Study in Finite Models*, volume 9 of *Lecture Notes in Logic*. Springer-Verlag, 1997.
21. N. Robertson and P.D. Seymour. Graph minors II: Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):307–322, 1986.
22. N. Robertson and P.D. Seymour. Graph minors V: Excluding a planar graph. *Journal of Combinatorial Theory Series B*, 41(1):91–114, 1986.
23. N. Robertson and P.D. Seymour. Graph minors XX: Wagner's conjecture. *Journal of Combinatorial Theory Series B*, 92(2):325–357, 2004.
24. P.D. Seymour and R. Thomas. Graph searching and a min-max theorem for tree-width. *Journal of Combinatorial Theory Series B*, 58:22–33, 1993.