Lecture Notes in Computer Science 665
Edited by G. Goos and J. Hartmanis

Advisory Board: W. Brauer D. Gries J. Stoer

P. Enjalbert A. Finkel K.W. Wagner (Eds.)

STACS 93

10th Annual Symposium
on Theoretical Aspects of Computer Science
Würzburg, Germany, February 25-27, 1993
Proceedings

Springer-Verlag
Berlin Heidelberg New York
London Paris Tokyo
Hong Kong Barcelona
Budapest

Fixed-Parameter Intractability II
(Extended Abstract)

Karl A. Abrahamson[1], Rodney G. Downey[2] *, Michael R. Fellows[3] **

[1] Washington State University, Department of E.E. & C.S., Pullman, Washington 99164-1210, U.S.A. (karl@eecs.wsu.edu)
[2] Victoria University of Wellington, Department of Mathematics, P.O. Box 600, Wellington, New Zealand (downey@math.vuw.ac.nz)
[3] University of Victoria, Department of Computer Science, Victoria, British Columbia V8W 3P6, Canada (mfellows@csr.uvic.ca)

Abstract. We describe new results in parameterized complexity theory, including an analogue of Ladner's theorem, and natural problems concerning k-move games which are complete for parameterized problem classes that are analogues of P-space.

1 Introduction

The theory of NP-completeness provides an excellent vehicle for explaining the apparent asymptotic intractability of many algorithmic problems. Yet while many natural problems do behave intractably *in the limit*, the manner *by which* they arrive at this intractable behavior can vary considerably. The standard NP and other completeness models are often far too coarse to give insight into this variation. To be specific, many natural computational problems take input consisting of two or more parts.

Example 1. The Vertex Cover problem takes as input a pair (G, k) consisting of a graph G and integer k, and determines whether there is a set of k vertices in G such that every edge in G has at least one endpoint in this set.

Example 2. The Graph Genus problem takes as input a pair (G, k) as above, and determines whether the graph G embeds on the surface of genus k.

Example 3. The Planar Improvement problem takes as input a pair (G, k) and determines if G is a subgraph of a planar graph G' of diameter at most k.

Example 4. The Graph Linking Number problem takes as input a pair (G, k) as above, and determines whether G can be embedded in 3-space so that at most k disjoint cycles in G are topologically linked.

* Research supported by Victoria University IGC, by the United States / New Zealand Cooperative Science Foundation under grant INT 90-20558, by the University of Victoria, and by the Mathematical Sciences Institute at Cornell University.
** Research supported by the National Science and Engineering Research Council of Canada, and the U.S. National Science Foundation under grant MIP-8919312.

Example 5. The Dominating Set problem takes as input a pair (G, k) as above, and determines whether there is a set of k vertices in G having the property that every vertex of G either belongs to the set, or has a neighbor in the set.

With the exception of examples 3 and 4 for which this question is open, the above problems are known to be NP-complete. But what can be said when the parameter k is held fixed? For examples 1-4, there is a constant α such that for every fixed k the problem can be solved in time $O(n^\alpha)$. For examples 2 and 4 we may take $\alpha = 3$ by the deep results of Robertson and Seymour [17] [18]. Example 5 illustrates the contrasting situation where for fixed values of k we seem to be able to do no better than a brute force examination of all possible solutions. We are thus concerned with an issue that is very much akin to P *versus* NP. The previous papers of this series [7] [8] [9] [10], established a framework with which to address the apparent fixed-parameter intractability of problems such as example 5. We feel that our framework provides an important contribution to the analysis of complexity of combinatorial problems for the following reasons.

(i) Distinct from most other notions and classes inroduced since the original clarification of NP and $PSPACE$ completeness, our framework is applicable to a *wide* class of *practical* problems.

(ii) Our framework provides a refined measure to that helps to explain the apparent diversity of *actual* behavior of many hard problems, as well as having numerous other applications.

2 Preliminaries

A *parameterized problem* is a set $L \subseteq \Sigma^* \times \Sigma^*$. Typically, the second component represents a parameter $k \in N$ in unary. For $k \in \Sigma^*$ we write $L_k = \{y | (y, k) \in L\}$. Consideration of examples 1-4 of §1 leads to three flavours of tractability and reduction.

Definition 1. We say that a parameterized problem L is
(1) *nonuniformly fixed-parameter tractable* if there is a constant α and a sequence of algorithms Φ_x such that, for each $x \in N$, Φ_x computes L_x in time $O(n^\alpha)$;
(2) *uniformly fixed-parameter tractable* if there is a constant α and an algorithm Φ such that Φ decides if $(x, k) \in L$ in time $f(k)|x|^\alpha$ where $f : N \to N$ is an arbitrary function;
(3) *strongly uniformly fixed-parameter tractable* if L is uniformly fixed-parameter tractable with the function f recursive.

Example 1 is strongly uniformly f.p.tractable (as are most examples of f. p. tractability obtained without essential use of the Graph Minor Theorem). Example 2 can be shown to be strongly uniformly f.p. tractable by the methods of [14] (the GMT alone gives only give nonuniform tractability). Example 3 can be shown to be uniformly f.p. tractable by the method of [13] (since the technique of [14] is not presently known to apply, we do not know a strongly uniform algorithm). Example 4 is at present only known to be nonuniformly f.p. tractable. If $P = NP$ then example 5 is also f.p. tractable.

Definition 2. Let A, B be parameterized problems. We say that A is *uniformly P-reducible* to B if there is an oracle algorithm Φ, a constant α, and an arbitrary function $f : N \to N$ such that
(a) the running time of $\Phi(B; \langle x, k\rangle)$ is at most $f(k)|x|^\alpha$,
(b) on input $\langle x, k\rangle$, Φ only asks oracle questions of $B^{(f(k))}$ where

$$B^{(f(k))} = \bigcup_{j \le f(k)} B_j = \{\langle x, j\rangle : \langle x, j\rangle \in B\}$$

(c) $\Phi(B) = A$.

If A is uniformly P-reducible to B, write $A \le_T^u B$. We may say that $A \le_T^u B$ *via* f. If the reduction is many:1 (an *m-reduction*), write $A \le_m^u B$.

Definition 3. We say that A is *strongly uniformly P-reducible* to B if $A \le_T^u B$ via f where f is recursive. We write $A \le_T^s B$ in this case.

Definition 4. We say that A is *nonuniformly P-reducible* to B if $\exists$ a constant α, a function $f : N \to N$, and a collection of procedures $\{\Phi_k : k \in N\}$ such that $\Phi_k(B^{(f(k))}) = A_k$ for $k \in N$, and the running time of Φ_k is $f(k)|x|^\alpha$. Here we write $A \le_T^n B$.

Note that the above are good definitions since whenever $A < B$ with $<$ any of the reducibilities, if B is f.p. tractable so too is A. We will write $FPT(\le)$ for the f.p. tractable class corresponding to the reducibility $\le$.

Theorem 5. *[11] The ordering $\le_T^s$ partitions $FPT(\le_T^u)$ into infinitely many classes. Similarly $\le_T^u$ partitions $FPT(\le_T^n)$ into infinitely many classes.*

Fix attention on any of the above reducibilities, and call it *fp-reducibility*. We consider circuits (termed *mixed type*) in which some gates have bounded fan-in (*small gates*) and some have unrestricted fan-in (*large gates*).

Definition 6. The *depth* $d(C)$ of a circuit C is the maximum number of gates (small or large), on an input-output path in C. The *weft* $w(C)$ of a circuit C is the maximum number of large gates on an input-output path in C.

Definition 7. We say that a family of circuits F has *bounded depth* if $\exists$ a constant h such that $\forall C \in F$, $d(C) \le h$. We say that F has *bounded weft* if $\exists$ a constant t such that $\forall C \in F$, $w(C) \le t$. F is a *decision circuit family* if each circuit has a single output. A decision circuit C *accepts* an input vector x if the single output gate has value 1 on input x. The *weight* of a boolean vector x is the number of 1's in the vector.

Definition 8. Let F be a family of decision circuits. We allow that F may have many different circuits with a given number of inputs. To F we associate the parameterized circuit problem $L_F = \{(C, k) : C \in F$ and C accepts an input vector of weight $k\}$. A parameterized problem L belongs to $W[t]$ if L reduces to the parameterized decision circuit problem $L_{F(t,h)}$ for the family $F(t, h)$ of mixed type decision circuits of weft at most t, and depth at most h, for some constant h. A parameterized problem L belongs to $W[P]$ if L reduces to the circuit problem L_F where F is the set of all circuits (no restrictions).

The above leads to an interesting hierarchy of parameterized problem classes

$$FPT \subseteq W[1] \subseteq W[2] \subseteq \ldots \subseteq W[P]$$

If $P = NP$ then the hierarchy collapses. We conjecture that each of the containments is proper. Many natural problems are complete for various levels. For example, Independent Set is complete for $W[1]$ and Dominating Set is complete for $W[2]$. Determining whether there are k vertices in a graph covering all subgraphs of minimum degree 3 is complete for $W[P]$. For a catalogue of problem classifications see [8], [10].

3 Fixed Parameter Analogs of PSPACE

Several game problems are known to be PSPACE-complete. Typically, such problems ask whether the first player to move has a winning strategy. A natural parameterized version of the problem is whether the first player has a strategy that wins within at most k moves.

Some hard parameterized game problems are f. p. tractable. An example is the Alternating Hitting Set game [15] [19] restricted to sets of size 2, which is PSPACE-complete. Some game problems appear not to be f. p. tractable. Generalized Geography is a game played on a directed graph G with a distinguished start vertex [15], [19]. Players alternate choosing vertices, starting at the start vertex, in such a way that the chosen vertices form, in sequence, a simple directed path in G. The first player who is unable to choose a vertex loses. Determining whether player 1 has a winning strategy in a Generalized Geography game is PSPACE-complete. A good candidate for a game problem that is not f. p. tractable is *Short Geography*, in which it is asked whether player 1 has a strategy that wins a given game of Generalized Geography in at most k moves.

In order to address such questions we introduce the class AW and $AW[*]$, which plausibly contain problems that are not in FPT. We show that Short Geography is $AW[*]$-complete.

Like W, AW is the closure under fp-reductions of a kernel problem of such a general nature that it appears not to be fixed parameter tractable. This problem is a parameterized version of Quantified Boolean Formulas, defined as follows.

Definition 9. *PQBF* is the parameterized problem specified
Instance: A sequence $s_1, \ldots, s_r$ of pairwise disjoint sets of boolean variables, and a boolean formula p involving the variables in $s_1 \cup \cdots \cup s_r$.
Parameters: $r, k_1, \ldots, k_r$.
Question: Is it the case that there exists a size k_1 subset t_1 of s_1 such that for every size k_2 subset t_2 of s_2 there exists a size k_3 subset t_3 of s_3 such that

...(alternating quantifiers) such that, when the variables in $t_1 \cup \cdots \cup t_r$ are made true, and all other variables are made false, formula p is true?

Definition 10. AW is the set of all problems that fp-reduce to $PQBF$.

AW appears to be too large a class for our purposes. Instead, we concentrate on the class $AW[*]$.

Definition 11. $PQBF_t$ is the restriction of $PQBF$ in which the formula part must consist of t alternating layers of conjunctions and disjunctions, with negations applied only to variables, and the main operator a conjunction. For example, if $t = 2$ then the formula must be in conjunctive normal form. $AW[t]$ is the set of all parameterized problems that fp-reduce to $PQBF_t$, and $AW[*] = \bigcup_t AW[t]$.

Definition 12. A parameterized problem X is *$AW[t]$-complete* iff X is in $AW[t]$ and every problem in $AW[t]$ fp-reduces to X. X is *$AW[*]$-complete* iff X is in $AW[*]$ and every problem in $AW[*]$ fp-reduces to X.

Short Geography can be shown to be in $AW[3]$. This section sketches of a proof that $PQBF_t$ fp-reduces to Short Geography for every t. The reduction is actually from a restricted form of $PQBF_t$.

Definition 13. *Unitary* $PQBF_t$ is the restriction of $PQBF_t$ in which the parameters $k_1, \ldots, k_r$ are all 1.

The proof of the following lemma can be found in [1].

Lemma 14. *For $t > 0$, Unitary $PQBF_{2t}$ is $AW[2t]$-complete.*

Theorem 15. *Short Geography is $AW[*]$-complete. Hence, $AW[*] = AW[3]$.*

Proof sketch. We reduce $PQBF_{2t}$ to Short Geography for an arbitrary $t > 0$. Let $I = (r, k_1, \ldots, k_r, s_1, \ldots, s_r, p)$ be an instance of $PQBF_{2t}$, and assume that r is odd. The reduction is a modification of Schaefer's polynomial time reduction from QBF to Generalized Geography [19]. The graph on which the geography game is played has three parts: the choice component, the formula testing component and the literal testing component.

The choice component is similar to Shaefer's, and is designed so that player 1 chooses a member of s_1, then player 2 chooses a member of s_2, then player 1 chooses a member of s_3, etc. A total of $3r$ moves are made through this component, and it is player 2's move at the end, where the game enters the formula testing component.

The formula testing component is also similar to Shaefer's, but has $2t$ levels. It uses player 1's moves to simulate disjunctions, and player 2's moves to simulate conjunctions. A total of $2t$ moves are made through this component. Play ends at a *literal vertex* v, corresponding to a literal in formula p, with the move being player 2's.

A literal vertex v corresponding to a positive literal x has an edge to the vertex v_x in the choice component that corresponds to x. If v_x was chosen (variable x is true), then player 2 has no move, and player 1 wins. If v_x was not chosen (so x is false), then player 2 moves to v_x and wins.

If literal vertex v corresponds to a negative literal $\bar{x}$, then the literal testing component has edges (v, u_x) and (u_x, v_x), where u_x is a new vertex and v_x is the vertex corresponding to x in the choice component. Vertex u_x switches the initiative, and causes player 1 to win if v_x was not chosen, and player 2 to win if v_x was chosen.

In summary, player 1 wins if and only if I is a *yes* instance of $PQBF_t$. The total number of moves is at most $3r + 2t + 2$. Since t is fixed, the conditions of an fp-reduction are met. □

4 Density and other Structural Results

We begin with an analogue of Ladner's density theorem. At present we can prove this only for strong uniform reducibility.

Theorem 16. *If A and B are recursive sets with $A < B$ and $<$ either of $<^s_T$ or $<^s_m$ then there is a recursive set C such that $A < A \oplus C < B$.*

Proof sketch. We begin by recalling the proof of Ladner's [16] density theorem for the polynomial time degrees. Recall that this works as follows. We are given recursive sets $A \not< B$ (working with $\leq^p_T$, say). Let $\{z_n : n \in N\}$ be a standard P-time length/lexicographic ordering of Σ^*. We can assume that A and B are given as the range of P-time functions with domain N in unary notation. We write $A_s = \{f(1^0), \ldots, f(1^s)\}$ if $f(N) = A$ in this sense. We can also ask that if $|f(1^y)| > |f(1^{y-1})|$ then for all $z > y$, $|f(z)| \geq |f(1^y)|$. We call this a *P-standard enumeration*. So we will assume that we have such enumerations of A and B. Recall also for a reduction Δ on a set E , $u(\Delta(E; x))$ denotes the length of the longest element used in the computation. Let $\{\Phi'_e : e \in N\}$ denote a standard enumeration of all P-time T-procedures.

We must build C to satisfy $C \leq^p_m B$ and the requirements:

$$R'_{2e} \ : \ \Phi'_e(A \oplus C) \neq B$$

$$R'_{2e+1} \ : \ \Phi'_e(A) \neq C$$

For the sake of the R'_j we define a polynomial time relation $R(n)$ on $N = \{1\}^*$. Then we declare that $x \in C$ iff $R(|x|) = 0$ and $x \in B$. Clearly this makes $C \leq^p_m B$.

Now we meet the R'_j *in order* by 'delayed' diagonalization. So we begin with R'_0. At each stage s, set $R(s) = 1$ until a stage t is found where (i) - (iv) below hold. (Here we consider s, t etc as being in N.)

(i) $\Phi'_{0,t}(A_t \oplus \emptyset; z_n) \downarrow$ in less than t steps.
(ii) $A_t[q] = A[q]$ if $|q| < u(\Phi'_0(A_t \oplus \emptyset; z_n)$.

(iii) $B_t[z_n] = B[z_n]$.
(iv) $\Phi'_{0,t}(A_t \oplus \emptyset; z_n)(= \Phi'_0(A \oplus \emptyset; z_n)) \neq B(z_n) = B_t(z_n)$.

At stage t we say that we have diagonalized R'_0 at z_n, this being found by *looking back for an A- and a B- certified disagreement.* The idea is then to move to R'_1 and then to R'_2 etc. For R'_1 we set $R(t+1) = 0$, causing C to look like B locally. So we keep $R(u)$ for $u > t$ equal to zero until a stage v is found with some $m \leq v$ and $\Phi'_{0,v}(A_v; z_m) \neq C_v(z_m)$, via A- and B- certified computations. We then move to R'_2 setting $R(v+1)$ to be 1 again. Thus the set C so constructed looks like B with 'holes' in it.

The above forms a sort of inner strategy for the full construction. The next step in the journey is to consider the strategy used by the first two authors to prove the density theorem for the strong uniform reducibility $\leq^s_T$ (and $\leq^s_m$) in [11]. Now imagine we are given $A < B$ with $\leq$ either $\leq^s_T$ or $\leq^s_m$. Again we must construct C, now to meet the following requirements:

$R_{2\langle e,n\rangle}$: Either ϕ_e is not total,
or $(\exists k)(B_k \neq \Phi_e(A \oplus C^{(\phi_e(k))}))$
or $(\exists x, k)(\Phi_e(A \oplus C^{(\phi_e(k))}; \langle x, k\rangle)$
does not run in time $\phi_e(k)|x|^n$.

$R_{2\langle e,n+1\rangle}$: Either ϕ_e is not total,
or $(\exists k)(C_k \neq \Phi_e(A^{(\phi_e(k))}))$
or $(\exists x, k)(\Phi_e(A^{(\phi_e(k))}; \langle x, k\rangle)$
does not run in time $\phi_e(k)|x|^n$).

To aid the discussion we will use several conventions. First, if $\phi_{e,s}(k)\downarrow$, then the computation $\Phi_e(E^{(\phi_e(k))}; \langle x, k\rangle)$ cannot call any y of the form $\langle k', z\rangle$ for $k' > \phi_e(k)$. Also since we get a win for free if $\phi_{e,s}(k)\downarrow$ and the running time of $\Phi_e(E^{(\phi_e(k))}; \langle x, k\rangle)$ exceeds $\phi_e(k)|x|^n$, we shall assume that in the above the third option does not pertain to R_j and concentrate on the first two. This is because if the running time exceeds the bounds during the construction, we can *cancel* the relevant requirement. The argument to follow is a priority one with the Ladner strategy embedded.

Without loss of generality we can take ϕ_e to be strictly increasing. Again there will be long intervals with $C(\langle x, k\rangle)$ equal to $\emptyset$ and long intervals where it looks like B, for 'many' k. We have problems, since, for instance, we cannot decide if ϕ_e is total. We first focus on the satisfaction of a single $R_0 = R_{2\langle e,n\rangle}$. We then describe the basic module for an odd type requirement, and finally describe the coherence mechanism whereby we combine strategies.

The Basic R_0-Module.

To meet R_0 above, we perform the following cycle. We have a parameter $k(0, s)$ that is nondecreasing in s and such that $\lim_s k(0, s) = k(0)$ exists. This is meant to be the number of "rows" devoted to R_0. It remains constant until we change it.

1. (Initialization.) Pick $k(0,0) = 1$.
2. Wait until a stage s occurs with one of the following holding:

2(a). (Win.) "Looking back" we see a disagreement. That is, as with the Ladner argument, we see an $n < s$ with $z_n \in \{\langle x, j\rangle : j < k(0, s)\}$.

$$\Phi_{e,s}(A \oplus C^{(\phi_e(k(0,s)-1))}; z_n) \neq B(z_n)$$

via A- and B-certified computations, or

2(b). Not (2a) and $\phi_{e,s}(k(0, s))\downarrow$.

Comment If s does not occur then $\phi_e(k(0, s))\uparrow$ and hence ϕ_e is not total. In this case we call $k(0, s)$ a *witness to the nontotality* of ϕ_e.

If 2(a) pertains, we declare R_0 to be *satisfied* (forever) and end its effect (forever). If 2(b) pertains, then we perform the following action.

3. R_0 *asserts control of* $C^{(\phi_e(k(0,s)))}$. That is, R_0 asks that for all $t \geq s$, until 2(b) pertains, we promise to set $C^{(\phi_e(k(0,s)))}(y) = 0$ for all y with $|y| = t$ and $y \in (\Sigma^*)^{(\phi_e(k(0,s)))}$. This can be achieved *via* a restraint $r(n, k)$.

4. Reset $k(0, s+1) = k(0, s) + 1$ and go to 2.

The Outcomes of the Basic R_0 Module.

We claim that 2(b) cannot occur infinitely often and hence $\lim_s k(0, s) = k(0)$ exists. Note that we have only reset $k(0, s)$ if 2(b) pertains in step 3. So suppose $k(0, s) \to \infty$ and hence $\phi_e(k(0, s)) \to \infty$. Then for each q and almost all y , we have $C(\langle q, y\rangle) = 0$. We write $A =^* B$ to denote that the symmetric difference of A and B is finite. So $C_q =^* \emptyset$ for all q. Furthermore, for all q, we can compute a stage $h(q)$ where

$$[\forall t > h(q)](C_q(\langle y, q\rangle) = 0 \quad \text{for all } y \text{ with } |y| > h(q))$$

where $h(q)$ is the stage where R_0 asserts control of row q. Finally, we know that for all k,

$$\Phi_e((A \oplus C)^{(\phi_e(k))}) = B_k$$

This allows us to get a reduction $\Delta(A) = B$. For each input $\langle y, k\rangle$, Δ simply computes $B(\langle y, k\rangle)$ for all y with $|y| \leq h(k)$, and $C(\langle z, k'\rangle)$ for all k', z with $k' \leq \phi_e(k)$ and $|z| \leq h(k)$. Then Δ simulates $\Phi_e(A^{(\phi_e(k))}; \langle y, k\rangle)$ if $|y| > h(k)$ with the exception that, if Φ_e calls some $\langle r, k'\rangle$ with $|r| \leq h(k)$ (and necessarily $k' \leq \phi_e(k)$), then Δ uses the table of values for C to provide the answer.

Note that the computations of $\Delta(A; \langle x, k\rangle)$ and $\Phi_e(C; \langle x, k\rangle)$ must agree and hence $\Delta(A) = B$, a contradiction. Thus 2(b) can pertain only finitely often. It follows that there are two outcomes.

Outcome $(0, f)$: 2(a) occurs for some t. Then we win R_0 with finite effect. (*Comment:* Once R_0 is met in this way, say at stage t, then we are completely free to do what we like with all y for which $|y| > t$ without injuring R_0.)

Outcome $(0, \infty)$: 2(a) does not occur. Then ϕ_e is not total. Note that the effect of R_0 is in this case infinite and for some $k = \lim_s k(0, s) - 1$, we will have

$$C^{(\phi_e(k))} =^* \emptyset$$

and furthermore, there is a reduction Δ_0 with time bound $\phi_e(k)|x|^n$ for which

$$\Delta_0(A^{(\phi_e(k))}) = B^{(k)}$$

Note that for the basic module, Δ_0 is simply Φ_e.

The Basic Module for R_1.

This is essentially the same as for R_0 except that for R_1 we wish to set $C(\langle x,k\rangle) = B(\langle x,k\rangle)$. Herein is the basic conflict: an even-indexed requirement R_j asks that lots of rows look like $\emptyset$ and an odd-indexed R_j asks for them to look like B.

Combining Strategies.

We cannot perform a delayed diagonalization as in the proof of Ladner's theorem, since we cannot know if $\phi_e(k)$ is defined. The combination of strategies needs the priority method. Let us consider a module for R_1 that works in the outcomes of R_0. We cannot know if this outcome is $(0,f)$ or $(0,\infty)$. Instead we have a strategy based on a guess as to R_0's behavior. Basically R_0 always believes that $k(0,s)$ is $k(0)$, that is, that the current value is the final one. Let $e = e(0)$, $n = n(0)$, $f = e(1)$ and $m = n(1)$.

Whilst R_1 believes that $\phi_e(k(0,0)) \uparrow$, R_1 acts as if R_0 is not there. So if $k(0,0) = k(0)$ and $\phi_e(k(0,0)) \uparrow$ then we win R_1 for the same reasons as we did for R_0. On the other hand, if $\phi_e(0) \downarrow$ for some least stage s, then R_0 will assert control of $C^{(\phi_e(k(0,0)))}$. For the sake of R_1 we have probably been setting $C(0,x) = B(0,x)$ for all x with $|x| < s$. Since R_0 has higher priority than R_1, R_1 must release its control of C_0 (and indeed of C_j for $j \le \phi_e(k(0))$) until a stage, if any, occurs where 2(a) pertains to R_0 so that R_0 is satisfied and releases control forever (or it becomes inactive because of a time bound being exceeded). Note that if 2(a) pertains at t, then R_1 is free to reassert control of C_0 for all y of the form $\langle y,0\rangle$ with $|y| > t$. Also, in this case, as R_1 is the requirement of highest overall priority remaining, its control cannot be violated and hence it will be met.

On the other hand, while R_0 can hope that 2(a) will pertain to R_1, R_0 may have outcome $(0,\infty)$ and R_0 will never release control of C_0. The key idea at this point is that we begin anew with a version of R_1 believing that $k(0,s+1) = k(0)$. That is, R_0 will *never again act.*

This version of R_1 can only work with C_q for $q > \phi_e(k(0,s)) = \phi_e(k(0,0))$. Some care is needed since potentially we need all of B to meet R_1.

An elegant solution to this difficulty is to *shift* B into C above $\phi_e(k(0,s))$. Thus R_1 will ask that

$$C(\langle x,q\rangle) = B(\langle x, q-\phi_e(k(0,s))-1\rangle)$$

for $q > \phi_e(k(0,s))$. It does so until either $k(0,t)$ is reset again, or 2(a) pertains, or the time bounds are exceeded. In the latter cases, it reverts to the $(0,f)$-strategy. In the first case it begins anew on $q > \phi_e(k(0,t))$. Since this restart process only occurs finitely often, it follows that we eventually get a *final* version of R_1 whose actions will not be disturbed.

Thus there is a final version of R_1 that is met as follows. As $\lim_s k(0,s) = k(0)$ exists, there is a value r and a stage s_0 so that for $q \ge r$ and $s > s_0$, R_1 is not initialized at stage s and can assert control on C_q if it so desires. If R_0 has outcome $(0,f)$, then $r = 0$, otherwise $r = \phi_e(k(0)-1)+1$. So we know that if R_1 fails then for all j there is a stage $h(j)$ (computable from the parameters r and s_0) where for y with $|y| > h(j)$

$$C(\langle y, r+j\rangle) = B(\langle y,j\rangle) \text{ and}$$

$$\Phi_f(A; \langle y, r+j\rangle) = C(\langle y, r+j\rangle).$$

Thus if R_1 fails again we can prove there is a reduction $\Delta(A) = B$ with running time $O(|z|^m)$ and computable constants. This is a contradiction.

The outcomes for R_1 are thus either $(1,\infty)$ and $(1,f)$. In the former case we know that for a finite number of rows j and for almost all y, $C(\langle y,j\rangle) = B(\langle y,j\rangle)$. But we also know that for such rows there is a reduction Δ_f such that

$$\Delta_f(A; \langle y,j\rangle) = C(\langle y,j\rangle) \text{ in time } O(|y|^m) \\ \text{and computable constants.}$$

We continue in the obvious way with the inductive strategies. Consider e.g. R_2. It is confronted with at worst a finite number of rows permanently controlled by R_0 and a finite number by R_1. However, in each case we know that there is a reduction from a computable number of rows af A to these rows, and hence a reduction

$$\Psi_2(A; \langle y,j\rangle) = C(\langle y,j\rangle)$$

for all j cofinally under the control of either R_0 or R_1. Therefore to argue that R_2 is met, we get to use Ψ_2 to help construct a reduction from A to B. That is, for R_i, let $e = e(i)$ and $n = n(i)$. Then inductively we have a reduction and constants $p(2), m(2)$ and $r(2)$ with

$$\Psi_2(A^{m(2)}; \langle x,j\rangle) = C(\langle x,j\rangle)$$

for all $j \le p(q)$ running in time $m(2)|x|^{r(2)}$. Futhermore, we have a stage s_2 such that for all $k < 3$, R_k ceases further activity.

Thereafter R_2 is free to assert control over any row q of C for $q > p(2)$. If we suppose that R_2 fails, then for each such q, R_2 will eventually assert control of C_q at some stage $h(q)$ to make $C(\langle x,q\rangle) = 0$ for all x with $|x| > h_2(q)$ and we have $\Phi_{e(2)}(C) = B$.

Now to get a reduction Δ from A to B we go as for R_0 except that now if $\Phi_{e(2)}$ makes an oracle question of $\langle y,j\rangle$ for $j \le p(2)$,we use Ψ_2 to answer this question.Thus we get a reduction Δ_2 that runs in time $O(|x|^{r(2)+n(2)})$, with computable constants and correct use. Thus again $B \le A$, a contradiction. Further details can be found in [11]. □

Remark. The Ladner paper not only proved density and other theorems for polynomial time reducibilities but additionally introduced the technique of delayed diagonalization ("looking back") which has been central to most structural analyses of polynomial-time degree structure. In our case we similarly introduce new techniques that would seem to be central to any analyses of the local degree structure for these parameterized reducibilities.

While we can only prove the general density theorem for the strongly uniform reducibilities, we have been able to prove some weaker versions of the density theorem for the other less uniform reducibilities. For instance we have been able to show that if there is a row k such that $B_k \not\leq_T^u A$ but $A \leq_T^u B$ then there is a C with $A <_T^u A \oplus C <_T^u B$. This again uses an infinite injury priority argument.

Ladner's Theorem holds much more generally than simply for the polynomial time degrees of *recursive* sets. Shinoda (see [11] [20] [4] observed that for $q = T$ the theorem holds for *any* sets $A <_T^p B$. This is not true for our reducibilities.

Theorem 17. *There exist sets A, B such that $A >_T^s B$ and for all D it is not the case that $A >_T^s D >_T^s B$. (Similarly for the other reducibilities.)*

The proof of this result uses a tree of strategies priority argument, and can be found in [11]. With Peter Cholak, the second author has also analysed the complexity of the reducibility orderings. Using a coding similar to that used by Ambos-Spies and Nies [2], for the P-time m-degrees, and a coding similar to that used by Ambos-Spies, Nies, and Shore [3], for the r.e. *wtt*-degrees. Cholak and Downey have proven the undecidability of the recursive sets under any of the m-reductions of section 2. The proofs are technically rather more difficult since the orderings are Σ_4 in the non-strongly uniform case. These proofs use $0^{(4)}$ priority techniques together with the speedup technique of [6]. For the reader familiar with [3] or [2] things are also hampered by the lack of an exact pair theorem for the relevant structures. In [5] it is also shown that one cannot just hope to lift the P-time result by, say, taking as a representative of a set A for this purpose the set which is empty on all rows except the first where is copies A. For instance, it is shown that this process can take a minimal pair for $\leq_T^p$ and turn the pair into a pair without infimum for $\leq_T^u$. This uses new techniques to analyse the normal polynomial degrees and shows how the study of these new structures and reducibilities can shed real light on the structure of the classic Karp and Cook reducibilities.

Many problems remain to be resolved. Aside from the general density question, one of great interest is whether collapse at the k-th level propogates upward in the W hierarchy. We conclude this section by mentioning some related oracle results. In [11], we proved:

Theorem 18. *There exist recursive oracles A, B, and C such that $P^A = NP^A$, $P^B \neq NP^B$, yet $W[P]^A = W[P]^B = FPT^A = FPT^B$, and $W[1]^C \neq W[P]^C$.*

These results would seem to support the thesis that the W hierarchy is distinct from NP.

References

1. K. Abrahamson, R. Downey, and M. Fellows, "Fixed Parameter Tractability and Completeness IV: $W[P]$ and *PSPACE*," to appear.
2. K. Ambos-Spies and A. Nies, "The Theory of The Polynomial Time Many-One Degrees is Undecidable," to appear.
3. K. Ambos-Spies, A. Nies, and R.A. Shore, "The Theory of the Recursively Enumerable Weak Truth Table Degrees is Undecidable," to appear.
4. K. Aoki, J. Shinoda, and T. Tsuda, "On Π_2 Theories of hp-T Degrees of Low Sets," to appear.
5. P. Cholak and R. Downey, "Undecidability and Definability for Parameterized Polynomial Time Reducibilities," in preparation.
6. R. Downey, "Nondiamond Theorems for Polymomial Time Reducibility," to appear, J.C.S.S.
7. R. Downey and M. Fellows, "Fixed Parameter Tractability and Completeness," *Congr. Num.*, 87 (1992) 161-187.
8. R. Downey and M. Fellows, "Fixed Parameter Tractability and Completeness I: Basic Results," to appear.
9. R. Downey and M. Fellows, "Fixed Parameter Tractability and Completeness II: On Completeness for $W[1]$," to appear.
10. R. Downey and M. Fellows, "Fixed Parameter Intractability (Extended Abstract)," *Proc. 7th Conf. on Structure in Complexity Theory* (1992), 36-49.
11. R. Downey and M. Fellows, "Fixed Parameter Tractability and Completeness III: Some Structural Aspects of the W-Hierarchy," in preparation.
12. R. Downey and M. Fellows, "Fixed Parameter Tractability," monograph in preparation.
13. M. R. Fellows and M. A. Langston, "On Search, Decision and the Efficiency of Polynomial-Time Algorithms." In *Proc. Symp. on Theory of Computing (STOC)* (1989), 501-512.
14. M. R. Fellows and M. A. Langston, "An Analogue of the Myhill-Nerode Theorem and Its Use in Computing Finite Basis Characterizations." In *Proc. Symp. Foundations of Comp. Sci. (FOCS)* (1989), 520-525.
15. M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness* (Freeman, San Francisco, 1979).
16. R. Ladner, "On the Structure of Polynomial Time Reducibility," *J.A.C.M.* 22 (1975), 155-171.
17. N. Robertson and P. D. Seymour, "Graph Minors XIII. The Disjoint Paths Problem," to appear.
18. N. Robertson and P. D. Seymour, "Graph Minors XV. Wagner's Conjecture," to appear.
19. T. J. Schaefer, "Complexity of Some Two-person Perfect Information Games," *J. Comput. Sys. Sci.* 16 (1978), 185-225.
20. J. Shinoda, Personal Communication, 1991.