Logics and complexity in finite model theory — The quest for P

I. Hodkinson

Dept. of Computing, Imperial College, London

0: Introduction

We discuss finding a logic characterising the polynomial-time-computable queries in the absence of linear order.

Themes:

1. extending fixed-point logic by generalised quantifiers to capture P on all classes of structures

2. restricting the classes until fixed-point logic itself captures P

3. detecting classes on which fixed-point logic collapses to first-order logic

Definition 0.1 A (k-ary) *query* on a class of structures is a map ρ associating with each structure A in the class a set $\rho^A \subseteq A^k$, such that if $\theta : A \cong B$ then $\theta(\rho^A) = \rho^B$.

Eg: any formula φ of a logic defines a query ($A \mapsto \varphi^A$).

Fix a class $\mathcal{C}$ of finite structures.

- A query is in P on $\mathcal{C}$ if, for any $A \in \mathcal{C}$ and $\bar{a} \in A^k$, it is decidable in P-time in $|A|$ whether $\bar{a} \in \rho^A$.

- A logic $\mathcal{L}$ is said to be in P ('$\mathcal{L} \leq P$') on the class if every formula defines a query that is in P on the class.

- A logic $\mathcal{L}$ *captures* P on the class ('$\mathcal{L} = P$') if the logic is in P on the class, and every query in P on the class is definable in the logic.

We are by now warmly familiar with:

Theorem 0.2 (Immerman,Vardi) *On any class of linearly ordered structures, LFP captures P.*

Aim: dispense with the order.

In general, LFP $<$ P.

So we may try to:

1. extend LFP to get P in all cases

2. find restricted classes of finite structures on which LFP captures P

3. (subsidiary) find when LFP expresses no more queries than first-order logic on a class.

We will need. . .

- $LFP = IFP \leq L^{\omega}_{\infty\omega}$ on finite structures.

- first-order types, L^k-types, $L^k_{\infty\omega}$-types. . .

- For any k there is an LFP-formula linearly ordering the $L^k_{\infty\omega}$-types in all finite structures. [Abiteboul–Vianu]

- Generalised quantifiers ('GQs'). Uniform sequences of GQs (to recognise interpretations of arbitrary arity). Counting quantifiers.

  As GQs may not be monotone, positive $\Rightarrow$ monotone may fail.

  So we use "FP": read 'FP+Q' as IFP+Q where Q is a set of GQs. Read 'FP' as either LFP or IFP.

1: A logic for P?

Would like to capture P on any signature.

Problem: there are limitative results.

Theorem 1.1 (Hella, 1992) *For any set Q of GQs of bounded arity, FP + Q does not capture P in all signatures.*

Is this a problem? Can interpret any M in a graph G(M), and G(M) in M. So add to FP a GQ for every P-time graph property (arity 2). Do we get all P-properties in all signatures?

No, by theorem 1.1.

The trouble is. . . arity of interpretation of G(M) in M = arity of M.

Need uniform sequences of GQs.

So it is a problem!

Now what?

A. Try to capture P on fixed signature — say graphs.

B. Go for it: try to capture P on all signatures.

Plan A: P on fixed signatures?

Take graphs.

Wide open to find a logic for P on graphs.

Theorem 1.2 (Cai–Fürer–Immerman) *FP + all counting quantifiers does not capture P on graphs.*

Proof generalises to give Hella's theorem (1.1).

Proof.

- Take a graph G and duplicate its edges:

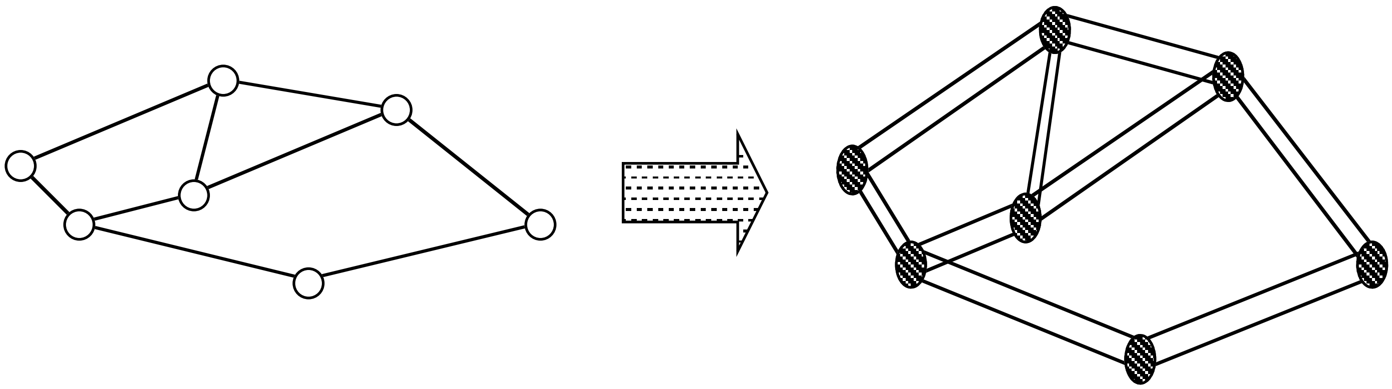

- nodes become 'gadgets' X — 'synchromesh' devices with automorphisms flipping even number of incoming edge-pairs:

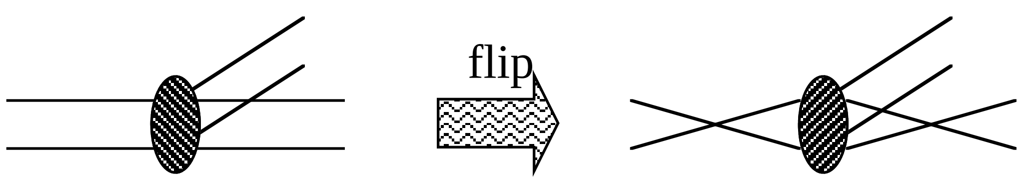

In more detail, showing the inside of X:

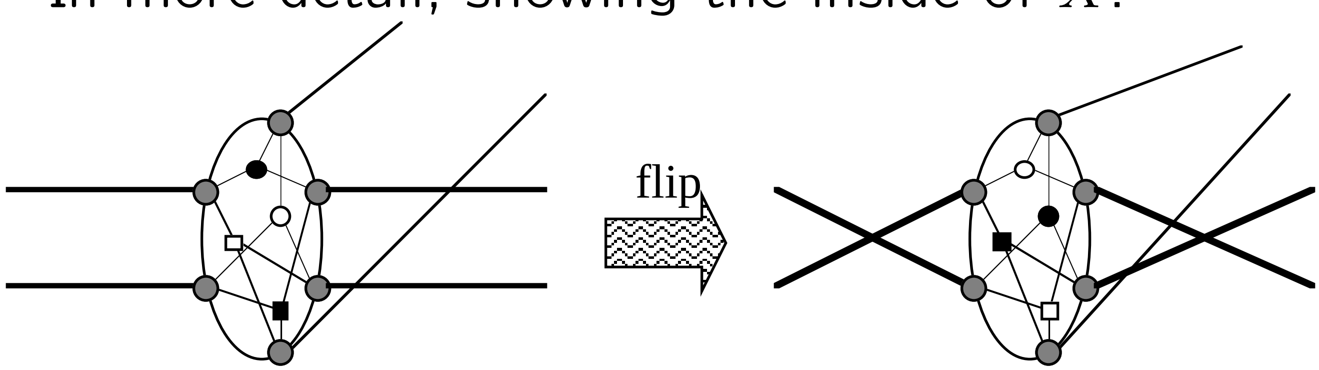

So, if G is connected, flipping an *even* number of edge pairs makes no difference to the isomorphism type:

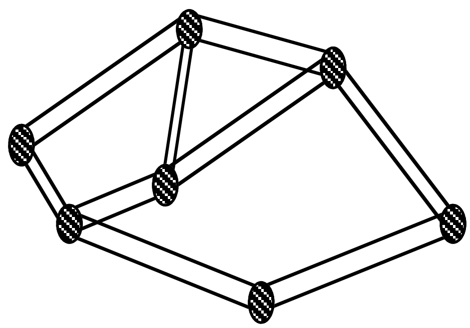 $\cong$ 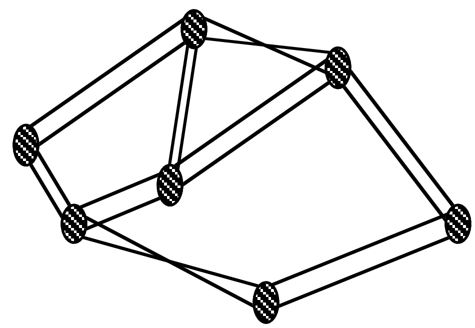

But flipping an *odd* number of edges gives a different graph up to isomorphism:

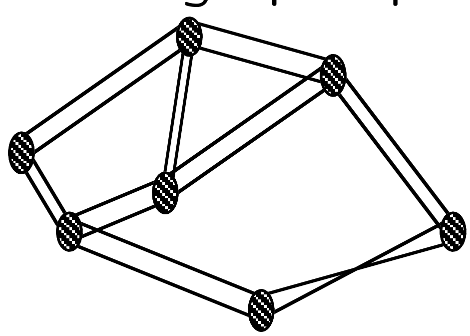 $\cong$ 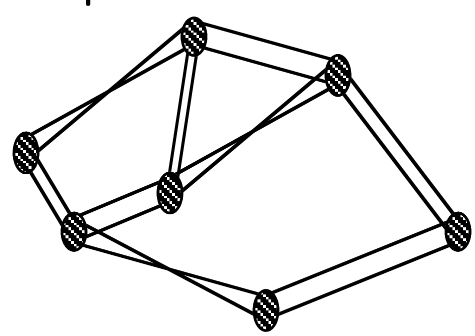

Don't care which edges are twisted. They can be moved around. Only parity of no. of twisted edges matters.

Fix k. Play k-pebble Ehrenfeucht–Fraïssé game (of length ω) between even- and odd-flipped varieties $X(G), \tilde{X}(G)$.

Ajtai: can choose G so that deleting the nodes of G corresponding to the k pebbled points in $X(G)/\tilde{X}(G)$ always leaves a connected component of size $> |G|/2$.

Duplicator wins by keeping the twisted edge (in $\tilde{X}(G)$) in this component — can do this, as successive cpts. must overlap!

Variation of game does counting quantifiers. But $X(G), \tilde{X}(G)$ differ so minutely that the *number* of points with a given $L^k_{\infty\omega}$-property is same in both.

Thus, for any k there is G s.t. $X(G)$, $\tilde{X}(G)$ are indistinguishable in $L^k_{\infty\omega}$ + counting.

But $X(G)$ and $\tilde{X}(G)$ can be distinguished in P-time.

So there are P-time graph properties not expressible in FP + counting. ♠

How about higher-arity GQs?

Theorem 1.3 (Dawar–Hella)
FP + finitely many GQs (of any arity) does not capture P on graphs.

Proof. Only need consider complete graphs! Show: $\forall Q$ (finite set of GQs): $L^\omega_{\infty\omega} + Q < P$.

Write A, B for complete graphs. Fix $k < \omega$.

Any $L^k + Q$-formula is equivalent in a given A to a disjunction of equality types (as these are orbits of $Aut(A)$ = automorphism group of A).

Let η_A be a ('Scott') sentence stating, for each formula of quantifier depth 1, exactly which disjunction it's equivalent to, in A.

Claim. For all B, if $B \models \eta_A$ then A, B agree on all $L^k_{\infty\omega} + Q$-sentences.

Proof of claim. Show by induction on $\varphi \in L^k_{\infty\omega} + Q$ that η_A forces φ to be equivalent to the same disjunction of equality types in B as in A. **Claim proved.**

So any sentence σ of $L^k_{\infty\omega} + Q$ is equivalent over complete graphs to $\bigvee_{A \models \sigma} \eta_A$.

As Q, L are finite, $\exists$ finitely many sentences of the form η_A.

Hence $\bigvee_{A \models \sigma} \eta_A$ reduces to a finite disjunction — a 'first-order' sentence of $L^k + Q$.

So on complete graphs, $L^k_{\infty\omega} + Q$ 'collapses' to first-order logic + Q.

So we can evaluate σ in any B by evaluating a finite (bounded) number of η_A in B.

Now show that if all quantifiers in Q are in P, this can be done in time of a fixed polynomial order, independent of k.

Time to evaluate η_A *for fixed* k?

- To evaluate η_A in B, enough to determine,

  – for each $q \in Q$ of the form $q\langle R_1, \ldots, R_n\rangle$ where R_i has arity m_i,

  – for each $E_i \subseteq B^{m_i}$, a disjunction of equality types over k parameters in B,

  whether $M = (B, E_1, \ldots, E_n)$ is in q.

- No. of such M is indep. of B. Each such M can be built in time $O(|B|^m)$, where m bounds the arities of the quantifiers in Q.

- If all quantifiers in Q are evaluable in P-time, there's a fixed l such that '$M \in q$' can be evaluated in time $O(|M|^l) = O(|B|^l)$.

Hence any η_A can be evaluated in time $O(|B|^r)$ for $r = \max(m, l)$.

We know:

- any $L^k_{\infty\omega} + Q$-sentence σ is equivalent to a bounded disjunction of η_As

- any η_A can be evaluated in any B in time $O(|B|^r)$.

So σ can also be evaluated in B in time $O(|B|^r)$.

But r does not depend on k.

Hence any $L^\omega_{\infty\omega} + Q$-sentence can be evaluated in a complete graph B in time $O(|B|^r)$.

But by the polynomial hierarchy theorem, there are P-time 0-ary queries on complete graphs not evaluable in time $O(|B|^r)$.

Hence $FP + Q \leq L^\omega_{\infty\omega} + Q < P$. ♠

Generalisation of Claim to bounded classes

We'll need this later.

Fix k and a finite set Q of GQs.

Definition 1.4 *A class of structures is said to be* $\boldsymbol{k, Q}$**-bounded** *if there's a finite bound on the number of* $L^k + Q$*-types realised in any of its structures.*

We write $\boldsymbol{k}$**-bounded** *where* Q *is empty.*

The argument up to the claim extends to these:

Theorem 1.5 (mostly Weinstein)
A class is k, Q*-bounded iff* $L^k_{\infty\omega} + Q$ *collapses to* $L^k + Q$ *on it.*

Proof. $\Rightarrow$: Let d^* bound the number of L^k+Q-types realised in the class. For A in the class, there is $d \leq d^*$ such that any L^k+Q-formula of quantifier depth $d+1$ is equivalent in A to a formula of quantifier depth $\leq d$.

Write a L^k+Q-sentence η_A expressing *which* formula each such formula is equivalent to, in A, and the inclusion relations between formulas of quantifier depth $\leq d$ in A.

Claim. Any $L^k_{\infty\omega}+Q$-formula φ is equivalent to a particular L^k+Q-formula φ_A of quantifier depth $\leq d$ in any model of η_A. **As before.**

η_A is of quantifier depth $\leq d^*+1$. There are finitely many such sentences up to logical equivalence. Now any $L^k_{\infty\omega}+Q$-formula φ is equivalent over the class to $\bigvee_A \eta_A \wedge \varphi_A$ — essentially a L^k+Q-formula.

$\Leftarrow$: If infinitely many types get realised in the class, say $p_0, p_1, \ldots$, then the $L^k_{\infty\omega}+Q$-formulas $\bigvee_{i\in I} \bigwedge p_i$ $(I \subseteq \omega)$ are pairwise inequivalent over the class. As L^k+Q is countable, there can be no collapse. ♠

Plan B: P in any signature?

Idea: factor out arity problem. We get a normal form for a logic for P (if one exists).

Theorem 1.6 (Dawar) *TFAE:*

1. *There's a problem complete for P w.r.t. first-order reductions (the reductions must not use a linear order!!)*

2. *Can capture P (on all structures, in any signature) using FP + a uniform seq. of GQs (generated by a single GQ).*

3. *Can capture P using FP + r.e. set of GQs.*

4. *P is recursively indexable.*

All are open. But if they fail then P $\neq$ NP.

Proof (sketch). $1 \Rightarrow 2$ (There's a problem complete for P w.r.t. first-order reductions $\Rightarrow$ can capture P using FP + a uniform seq. of GQs):

— take the class defining the GQ to be the complete problem.

$2 \Rightarrow 3 \Rightarrow 4$: obvious.

$4 \Rightarrow 1$ (P is recursively indexable $\Rightarrow \exists$ problem complete for P w.r.t. first-order reductions):

'P is recursively indexable' says:

- there's a Turing machine K that, on input $n < \omega$, outputs the code of a P-time Turing machine T_n accepting a class in P

- every class of structures in P is recognised by some T_n.

The complete problem is essentially the class $\mathcal{P}$ of structures of the form

$$A = (M, n, \text{padding}), \quad \text{where}$$

1. M: a structure that T_n accepts in time $|A|$

2. K runs on n to give the code of T_n also in time $|A|$.

Then $\mathcal{P}$ is P-complete:

- $\mathcal{P} \in$ P (use UTM to run K on n, then T_n on M; reject if don't finish in time $|A|$).

- If $\mathcal{C} \in$ P is recognised by T_n, the map $M \mapsto (M, n, \sharp)$, where $\sharp$ is enough padding to ensure (1–2), can be done by a first-order interpretation.

  This map gives a first-order reduction of $\mathcal{C}$ to $\mathcal{P}$.

♠

Conclusion of part 1

- FP + counting doesn't capture P on graphs.
- Neither does FP + any finite set of GQs.

Question: what does?

We found a normal form for any logic capturing P on all structures.

To come: what if we restrict the classes considered?

2: When does FP capture P?

When (on which classes) does FP remain as strong as P?

If there's an order around, there's hope of getting the fixed point operation to work really well.

But the order may not be accessible to FP, so we have to be careful.

Theorem 2.1 (Dawar, Hella–Kolaitis–Luosto) *TFAE in any class of finite structures:*

1. *the class is of* **bounded rigidity** *(it's k-rigid for some k — i.e., any two distinct elements of any structure have different L^k-types).*

2. *there's an FP-definable linear order*

3. *there's an $L^\omega_{\infty\omega}$-definable linear order.*

Proof. 1 $\Rightarrow$ 2: Abiteboul–Vianu $\Rightarrow$ $\exists$ FP-definable linear order on the $L^k_{\infty\omega}$-types. (1) implies this linearly orders the structures.

2 $\Rightarrow$ 3 is clear.

3 $\Rightarrow$ 1: suppose the order is definable in $L^k_{\infty\omega}$ ($k \geq 2$). Clearly, any two distinct elements of a structure will have different types in $L^k_{\infty\omega}$. But on a single finite structure, $L^k_{\infty\omega}$ collapses to L^k (apply theorem 1.5 [Weinstein] to a singleton class). ♠

So on classes of bounded rigidity, there's an FP-definable order and FP captures P.

Clearly, the structures in such classes are all **rigid** (no nontrivial automorphisms).

Is there an FP-order on rigid structures?

Theorem 2.2 (mainly Gurevich–Shelah) *There is a class of rigid structures (even one defined by a first-order sentence) that's not of bounded rigidity.*

Proof. Almost surety of k-extension axioms and of rigidity gives an example without the first-order restriction [Hella–Kolaitis–Luosto]. The full result [GS] uses *'multipedes'* and a probabilistic argument. ♠

Theorem 2.3 (Hella–Kolaitis–Luosto) *TFAE*

1. *the structures are rigid*

2. *there's a linear order definable implicitly in* $L^{\omega}_{\infty\omega}$.

Proof. $1 \Rightarrow 2$: If they're all rigid, choose an arbitrary linear order $<_A$ on each A in the class. Using $<_A$, we regard dom(A) as an initial segment of the natural numbers.

Let $k \geq 3$ bound the arity of the signature L of the structures. Write a first-order sentence χ_A of $(L \cup \{<\})^k$ with a conjunct for each $R(x_1, \ldots, x_{k'}) \in L$ and each k-tuple $n_1, \ldots, n_{k'}$ of numbers $< |A|$, saying whether R holds or fails at this k'-tuple in A.

Add another conjunct to χ_A, saying that $<$ is a linear order.

By rigidity, $<_A$ is the unique interpretation of $<$ as a linear order in A making χ_A true. So χ_A implicitly defines a linear order on A.

Let $\chi = \bigvee_A \chi_A$, a sentence of $(L \cup \{<\})^k_{\infty\omega}$. Then χ implicitly defines a linear order on every A in the class.

$2 \Rightarrow 1$: obvious. ♠

Conjecture 2.4 (Stolboushkin, 1992) *For first-order-definable classes of rigid structures, there's a linear order implicitly definable in first-order logic.*

Note: Kolaitis showed FP $\leq$ IMP(FO), so this is a natural weakening of the conjecture (of Dawar) refuted by theorem 2.2. If this one fails then P $\neq$ NP [HKL], so it seems fairly safe!

3: When does FP collapse to FO?

This is interesting subsidiary information: if FP collapses, we're not likely to capture P with it.

Historically its study has been influential, and nowadays a great deal is known about some of the classes on which FP collapses.

However, there are grave limitations:

- even if FP doesn't collapse, it's not clear that FP=P: e.g., we may not have an order

- even if it does collapse, it's not certain that FP $<$ P.

Five examples of collapse

Classes where FP collapses to FO include:

1. complete structures (boring)

2. homogeneous structures

   A is homogeneous if any partially-defined isomorphism : $A \to A$ extends to an automorphism.

3. More generally, classes that are **k-trivial** for all k [Dawar] — the automorphism group of each structure in the class has a bounded number of orbits on k-tuples for all k.

   Much is known about these: e.g., $k = 4$ is enough. This follows from a structure theory for them [Cherlin–Hrushovski], using some of the deepest ideas from classical model theory.

Theorem 3.1 (Dawar–Hella) *A class is k-trivial iff it is k, Q-bounded for any finite set Q of GQs.*

Proof. $\Rightarrow$: clear.

$\Leftarrow$: the same technique used to prove the Abiteboul–Vianu theorem on ordering the $L^k_{\infty\omega}$-types can be used to order the full first-order k-types.

Let q be a GQ expressing this order. In any finite structure A, two k-tuples have the same $L^k + q$-type iff they have the same full first-order type, iff they lie in the same orbit of $Aut(A)$. So for any class of structures, k, q-bounded implies k-trivial. ♠

By theorem 1.5 (mostly Weinstein), we get:

Corollary 3.2 *On k-trivial classes, $FP + Q$ collapses to $FO + Q$ for any finite set Q of GQs.*

4. More generally still, there are classes where we can't get a recursion going, so obviously FP collapses to first-order logic. McColm called these classes **non-proficient**.

Theorem 3.3 *TFAE*

(a) the class does not support an unbounded first-order induction (it's not 'proficient')

(b) it's k-bounded for all k (definition 1.4)

(c) $L^{\omega}_{\infty\omega}$ collapses to first-order logic on it

$a \Rightarrow c$: Kolaitis–Vardi. $c \Rightarrow a$: McColm. Rest: Dawar.

Proof. $a \Rightarrow b$: Assume not, and take a counterexample class. Then FP collapses to FO on the class; yet for some k, infinitely many L^k-types are realised in it.

We can order these types by a single FP-formula [Abiteboul–Vianu]. So by collapse, there's a first-order formula that defines arbitrarily long linear pre-orders in structures in the class. We can now do an unbounded first-order induction along the order, a contradiction.

$b \Rightarrow c$ follows from theorem 1.5 [Weinstein].

$\neg a \Rightarrow \neg c$: In proficient classes we can write infinitely many pairwise inconsistent $L^k_{\infty\omega}$-formulas for some k, picking out the stages of an unbounded induction. There are 2^ω disjunctions of these, all in $L^k_{\infty\omega}$ and all inequivalent over the class. As FO is countable, collapse is impossible. ♠

Hence FP = FO on non-proficient classes; and they include the classes that are k-trivial for all k.

5. Certain proficient classes! (McColm had conjectured this couldn't happen.)

Theorem 3.4 (Gurevich–Shelah, Immerman) *There is a first-order-definable* proficient *class of finite structures on which FP = FO.*

> **Proof.** Gurevich–Shelah used a probabilistic argument.
>
> Independently, Immerman used a deterministic constructon involving attaching cliques to a graph to record the effects of each fixed point operation. The cliques are first-order-detectable, giving FP = FO. ♠

Conjecture 3.5 (Kolaitis–Vardi) *On infinite ordered classes ('highly proficient'), FP>FO.*

True for otherwise unary signatures [Dawar–Lindell-Weinstein].

Question: find natural necessary and sufficient conditions for collapse of FP to FO.

Conclusion

We can say when there's a FP-definable order (need bounded rigidity; rigidity is not enough).

But maybe FP = P in other cases too. . . ?

Saying when FP collapses to FO is a subtle business.

Lots of open questions, lots of ideas needed.