Symmetric Computation: Part 2

Anuj Dawar

Department of Computer Science and Technology, University of Cambridge

PhD Open, Warsaw, 18 October 2019

Fixed-Point Logic with Counting

FPC is the class of *decision problems* definable in *fixed-point logic with counting*.

> *The decision problems are (isomorphism-closed) classes (or properties) of finite structures (such as graphs, Boolean formulas, systems of equations).*

Every problem in FPC is in P;

Symmetric Circuits

Say C_n is *symmetric* if any permutation of $[n]$ applied to its inputs can be extended to an automorphism of C_n.

> *i.e., for each $\pi \in S_n$, there is an automorphism of C_n that takes input (i, j) to $(\pi i, \pi j)$.*

Any symmetric circuit is invariant, but *not* conversely.

> *Consider the natural circuit for deciding whether the number of edges in an n-vertex graph is even.*

Any invariant circuit can be converted to a symmetric circuit, but with potentially *exponential blow-up*.

Logic and Circuits

Any formula of φ *first-order logic* translates into a uniform family of circuits C_n

> *For each subformula $\psi(\overline{x})$ and each assignment $\overline{a}$ of values to the free variables, we have a gate.*
> *Existential quantifiers translate to big disjunctions, etc.*

The circuit C_n is:

- of *constant* depth (given by the depth of φ);
- of size at mose $c \cdot n^k$ where c is the number of subformulas of φ and k is the *maximum number of free variables* in any subformula of φ.
- *symmetric* by the action of $\pi \in S_n$ that takes $\psi[\overline{a}]$ to $\psi[\pi(\overline{a})]$.

FP and Circuits

For every sentence φ of FP there is a k such that for every n, there is a formula φ_n of L^k that is equivalent to φ on all graphs with at most n vertices.

The formula φ_n has

- *depth* n^c for some constant c;
- at most k free variables in each sub-formula for some constant k.

It follows that every graph property definable in FP is given by a family of *polynomial-size, symmetric* circuits.

FPC and Counting

For every sentence φ of FP there is a k such that for every n, there is a formula φ_n of C^k that is equivalent to φ on all graphs with at most n vertices.

The formula φ_n has

- *depth* n^c for some constant c;
- at most k free variables in each sub-formula for some constant k.

It follows that every graph property definable in FP is given by a family of *polynomial-size, symmetric* circuits in a basis with *threshold gates*.

Note: we could also alternatively take a basis with *majority* gates.

Relating Circuits and Logic

The following are established in **(Anderson, D. 2017)**:

Theorem

A class of graphs is accepted by a P-uniform, polynomial-size, symmetric family of Boolean circuits if, and only if, it is definable by an FP formula interpreted in $G \uplus ([n], <)$.

Theorem

A class of graphs is accepted by a P-uniform, polynomial-size, symmetric family of threshold circuits if, and only if, it is definable in FPC.

Expressive Power of FPC

Most *"obviously"* polynomial-time algorithms can be expressed in FPC.

This includes P-complete problems such as

> CVP—*the Circuit Value Problem*
> *Input: a circuit, i.e. a labelled DAG with source labels from* $\{0, 1\}$, *internal node labels from* $\{\vee, \wedge, \neg\}$.
> *Decide: what is the value at the output gate.*

CVP is expressible in FPC.

It is expressible in FPC also for circuits that may include *threshold or counting gates*.

Expressive Power of FPC

Many non-trivial polynomial-time algorithms can be expressed in FPC:

FPC captures all of P over any *proper minor-closed class of graphs*

(Grohe 2010)

But some cannot be expressed:

- There are polynomial-time decidable properties of graphs that are not definable in FPC. **(Cai, Fürer, Immerman, 1992)**
- *XOR-Sat*, or more generally, solvability of a system of linear equations over a finite field cannot be expressed in FPC. **(Atserias, Bulatov, D. 2009)**

Some NP-complete problems are *provably* not in FPC, including *Sat*, *Hamiltonicity* and *3-colouraiblity*.

Counting Quantifiers

C^k is the logic obtained from *first-order logic* by allowing:

- allowing *counting quantifiers*: $\exists^i x\, \varphi$; and
- only the variables $x_1, \ldots .x_k$.

Every formula of C^k is equivalent to a formula of first-order logic, albeit one with more variables.

For every sentence φ of FPC, there is a k such that if $\mathbb{A} \equiv^{C^k} \mathbb{B}$, then

$$\mathbb{A} \models \varphi \quad \text{if, and only if,} \quad \mathbb{B} \models \varphi.$$

Counting Game

Immerman and Lander (1990) defined a *pebble game* for C^k.
This is again played by *Spoiler* and *Duplicator* using k pairs of pebbles $\{(a_1, b_1), \ldots, (a_k, b_k)\}$.

> *At each move, Spoiler picks i and a set of vertices of one structure (say $X \subseteq B$)*
>
> *Duplicator responds with a set of vertices of the other structure (say $Y \subseteq A$) of the same size.*
>
> *Spoiler then places a_i on an element of Y and Duplicator must place b_i on an element of X.*
>
> *Spoiler wins at any stage if the partial map from $\mathbb{A}$ to $\mathbb{B}$ defined by the pebble pairs is not a partial isomorphism*
>
> *If Duplicator has a winning strategy for p moves, then $\mathbb{A}$ and $\mathbb{B}$ agree on all sentences of C^k of quantifier rank at most p.*

Bijection Games

$\equiv^{C^k}$ is also characterised by a k-pebble *bijection game*. **(Hella 96)**.
The game is played on structures $\mathbb{A}$ and $\mathbb{B}$ with pebbles $a_1, \ldots, a_k$ on $\mathbb{A}$ and $b_1, \ldots, b_k$ on $\mathbb{B}$.

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

Duplicator loses if the partial map $a_i \mapsto b_i$ is not a partial isomorphism.
Duplicator has a strategy to play forever if, and only if, $\mathbb{A} \equiv^{C^k} \mathbb{B}$.

Equivalence of Games

It is easy to see that a winning strategy for *Duplicator* in the bijection game yields a winning strategy in the counting game:

> *Respond to a set* $X \subseteq A$ *(or* $Y \subseteq B$*) with* $h(X)$ *(*$h^{-1}(Y)$*, respectively).*

For the other direction, consider the partition induced by the equivalence relation

$$\{(a, a') \mid (\mathbb{A}, \mathbf{a}[a/a_i]) \equiv^{C^k} (\mathbb{A}, \mathbf{a}[a'/a_i])\}$$

and for each of the parts X, take the response Y of *Duplicator* to a move where *Spoiler* would choose X.

Stitch these together to give the bijection h.

Cai-Fürer-Immerman Graphs

Cai-Fürer-Immerman show that there is a polynomial-time graph property that is not in FPC by constructing a sequence of pairs of graphs $G_k, H_k (k \in \omega)$ such that:

- $G_k \equiv^{C^k} H_k$ for all k.
- There is a polynomial time decidable class of graphs that includes all G_k and excludes all H_k.

In particular, the first point shows that $\equiv^{C^k}$ (for any fixed k) does not capture isomorphism everywhere

Constructing G_k and H_k

Given any graph G, we can define a graph X_G by replacing every edge with a pair of edges, and every vertex with a gadget.

The picture shows the gadget for a vertex v that is adjacent in G to vertices w_1, w_2 and w_3.
The vertex v^S is adjacent to $a_{vw_i} (i \in S)$ and $b_{vw_i} (i \notin S)$ and there is one vertex for all *even size* S.
The graph $\tilde{X}_G$ is like X_G except that at *one vertex* v, we include v^S for *odd size* S.

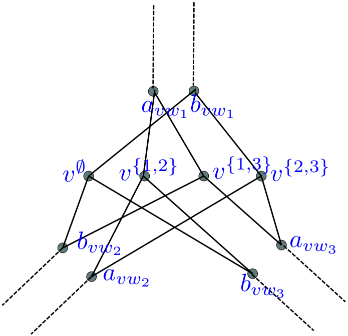

Properties

If G is *connected* and has *treewidth* at least k, then:

1. $X_G \not\cong \tilde{X}_G$; and
2. $X_G \equiv^{C^k} \tilde{X}_G$.

(1) allows us to construct a polynomial time property separating X_G and $\tilde{X}_G$.
(2) is proved by a game argument.

> *The original proof of* **(Cai, Fürer, Immerman)** *relied on the existence of balanced separators in* G*. The characterisation in terms of treewidth is from* **(D., Richerby 07)***.*

TreeWidth

The *treewidth* of a graph is a measure of how tree-like the graph is. A graph has treewidth k if it can be covered by subgraphs of at most $k+1$ nodes in a tree-like fashion.

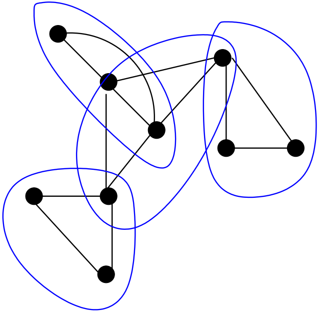

TreeWidth

Formal Definition:
For a graph $G = (V, E)$, a *tree decomposition* of G is a relation $D \subset V \times T$ with a tree T such that:

- for each $v \in V$, the set $\{t \mid (v,t) \in D\}$ forms a connected subtree of T; and
- for each edge $(u,v) \in E$, there is a $t \in T$ such that $(u,t),(v,t) \in D$.

We call $\beta(t) := \{v \mid (v,t) \in D\}$ the *bag* at t.

The *treewidth* of G is the least k such that there is a tree T and a tree-decomposition $D \subset V \times T$ such that for each $t \in T$,

$$|\{v \in V \mid (v,t) \in D\}| \leq k+1.$$

Cops and Robbers

A game played on an undirected graph $G = (V, E)$ between a player controlling k cops and another player in charge of a robber.

At any point, the cops are sitting on a set $X \subseteq V$ of the nodes and the robber on a node $r \in V$.

A move consists in the cop player removing some cops from $X' \subseteq X$ nodes and announcing a new position Y for them. The robber responds by moving along a path from r to some node s such that the path does not go through $X \setminus X'$.

The new position is $(X \setminus X') \cup Y$ and s. If a cop and the robber are on the same node, the robber is caught and the game ends.

Strategies and Decompositions

Theorem (Seymour and Thomas 93):
There is a winning strategy for the *cop player* with k cops on a graph G if, and only if, the tree-width of G is at most $k - 1$.

It is not difficult to construct, from a tree decomposition of width k, a winning strategy for $k + 1$ cops.

Somewhat more involved to show that a winning strategy yields a decomposition.

Cops and Robbers on the Grid

If G is the $k \times k$ toroidal grid, than the *robber* has a winning strategy in the k-*cops and robbers* game played on G.

To show this, we note that for any set X of at most k vertices, the graph $G \setminus X$ contains a connected component with at least half the vertices of G.

If all vertices in X are in distinct rows then $G \setminus X$ is connected. Otherwise, $G \setminus X$ contains an entire row and in its connected component there are at least $k - 1$ vertices from at least $k/2$ columns.

Robber's strategy is to stay in the large component.

Cops, Robbers and Bijections

We use this to construct a winning strategy for Duplicator in the k-pebble bijection game on X_G and $\tilde{X}_G$.

- A bijection $h : X_G \to \tilde{X}_G$ is *good bar* v if it is an isomorphism everywhere except at the vertices v^S.
- If h is good bar v and there is a path from v to u, then there is a bijection h' that is good bar u such that h and h' differ only at vertices corresponding to the path from v to u.
- Duplicator plays bijections that are good bar v, where v is the *robber position* in G when the cop position is given by the currently pebbled elements.

Counting Width

For any class of structures $\mathcal{C}$, we define its *counting width* $\nu_{\mathcal{C}} : \mathbb{N} \to \mathbb{N}$ so that

> $\nu_{\mathcal{C}}(n)$ *is the least* k *such that* $\mathcal{C}$ *restricted to structures with at most* n *elements is closed under* $\equiv^{C^k}$.

Every class in FPC has counting width bounded by a *constant*.

The *CFI* construction based on *grids* gives a class of graphs in P that has counting width $\Omega(\sqrt{n})$.

This can be improved to $\Omega(n)$ by taking, instead of grids, *expander graphs*.

Interpretations

Given two relational signatures σ and τ, where $\tau = \langle R_1, \ldots, R_r \rangle$, and arity of R_i is n_i

A *first-order interpretation of τ in σ* is a sequence:

$$\langle \pi_U, \pi_1, \ldots, \pi_r \rangle$$

of first-order σ-formulas, such that, for some d:

- the free variables of π_U are among $x_1, \ldots, x_d$,
- and the free variables of π_i (for each i) are among $x_1, \ldots, x_{d \cdot n_i}$.

d is the dimension of the interpretation.

Interpretations II

An interpretation of τ in σ maps σ-structures to τ-structures.

If $\mathbb{A}$ is a σ-structure with universe A, then
$\pi(\mathbb{A})$ is a structure $(B, R_1, \ldots, R_r)$ with

- $B \subseteq A^d$ is the relation defined by π_U.
- for each i, R_i is the relation on B defined by π_i.

An FO *reduction* of a class of structures $\mathcal{C}$ to a class $\mathcal{D}$ is a single FO interpretation θ such that $\mathbb{A} \in \mathcal{C}$ if, and only if, $\theta(\mathbb{A}) \in \mathcal{D}$.
We write $\mathcal{C} \leq_{\mathsf{FO}} \mathcal{D}$.

FPC-Reductions

More generally, we can defined reductions in any logic, *e.g.* FPC.

If $\mathcal{C} \leq_{\mathsf{FPC}} \mathcal{D}$ then

$$\nu_{\mathcal{D}} = \Omega(\nu_{\mathcal{C}}^{1/d}).$$

If the reduction takes $\mathcal{C}$-instances to $\mathcal{D}$-instances of *linear size*, then

$$\nu_{\mathcal{D}} = \Omega(\nu_{\mathcal{C}}).$$

By means of reductions, we can estalish *3-Sat*, *XOR-Sat*, *3-Colourability*, *Hamiltonicity* all have counting width $\Omega(n)$.

Relating Circuits and Logic

The following are established in **(Anderson, D. 2017)**:

Theorem

A class of graphs is accepted by a P-uniform, polynomial-size, symmetric family of Boolean circuits if, and only if, it is definable by an FP formula interpreted in $G \uplus ([n], <)$.

Theorem

A class of graphs is accepted by a P-uniform, polynomial-size, symmetric family of threshold circuits if, and only if, it is definable in FPC.

Some Consequences

We get a natural and purely circuit-based characterisation of FPC definability.

Inexpressibility results for FP and FPC yield lower bound results against natural circuit classes.

- There is no polynomial-size family of symmetric Boolean circuits deciding if an n vertex graph has an even number of edges.
- Polynomial-size families of uniform symmetric *threshold circuits* are more powerful than Boolean circuits.
- Invariant circuits *cannot* be translated into equivalent symmetric threshold circuits, with only polynomial blow-up.

Symmetric Circuits for non-Boolean Queries

Instead of circuits computing *Boolean* (i.e. $0/1$) queries, we can consider circuits C that compute an m-ary relation on an input graph.

The output gate is not unique. Instead, we have an *injective function* $\Omega : [n]^m \to C$.
The range of Ω forms the *output gates*.

The requirement that $\pi \in S_n$ extends to an automorphism $\hat{\pi}$ of C includes the condition:

$$\hat{\pi}(\Omega(x)) = \Omega(\pi(x))$$

Automorphisms of Symmetric Circuits

For a symmetric circuit C_n we can assume *w.l.o.g.* that the automorphism group is the symmetric group S_n acting in the natural way.

That is:

- Each $\pi \in S_n$ gives rise to a *non-trivial* automorphism of C_n (otherwise C_n would compute a constant function).
- There are no *non-trivial* automorphisms of C_n that fix all the inputs (otherwise there is redundancy in C_n that can be eliminated).

We call a circuit satisfying these conditions *rigid*.

By abuse of notation, we use $\pi \in S_n$ both for permutations of $[n]$ and automorphisms of C_n.

Stabilizers

For a gate g in C_n, $\mathrm{Stab}(g)$ denotes the *stabilizer group of g*, i.e. the *subgroup* of S_n consisting:

$$\mathrm{Stab}(g) = \{\pi \in S_n \mid \pi(g) = g\}.$$

The *orbit* of g is the set of gates $\{h \mid \pi(g) = h \text{ for some } \pi \in S_n\}$

By the *orbit-stabilizer* theorem, there is one gate in the orbit of g for each *co-set* of $\mathrm{Stab}(g)$ in S_n.
Thus the size of the *orbit* of g in C_n is $[S_n : \mathrm{Stab}(g)] = \frac{n!}{|\mathrm{Stab}(g)|}$.

So, an upper bound on $\mathrm{Stab}(g)$ gives us a lower bound on the orbit of g.

Conversely, knowing that the orbit of g is at most polynomial in n tells us that $\mathrm{Stab}(g)$ is *big*.

Supports

For a group $G \subseteq S_n$, we say that a set $X \subseteq [n]$ is a *support* of G if

For every $\pi \in S_n$, if $\pi(x) = x$ for all $x \in X$, then $\pi \in G$.

In other words, G contains all permutations of $[n] \setminus X$.
So, if $|X| = k$, $[S_n : G]$ is at most $\frac{n!}{(n-k)!} \leq n^k$.

Groups with small support are *big*.

The converse is clearly false since $[S_n : A_n] = 2$, but A_n has no support of size less than $n - 1$.

Note: For the family of circuits $(C_n)_{n \in \omega}$ obtained from an FPC formula there is a constant k such that all gates in each C_n have a support of size at most k.

Support Theorem

In *polynomial size* symmetric circuits, all gates have (stabilizer groups with) *small* support:

Theorem

For any polynomial p, there is a k such that for all sufficiently large n, if C is a symmetric circuit on $[n]$ of size at most $p(n)$, then every gate in C has a support of size at most k.

The general form of the support theorem in **(Anderson, D. 2017)** gives bounds on the size of supports in *sub-exponential* circuits.

Alternating Supports

Groups with small support are *big*.

The converse is clearly false since $[S_n : A_n] = 2$, but A_n has no support of size less than $n - 1$.

In a sense, the alternating group is the *only* exception, due to a standard result from permutation group theory.

Theorem

If $n > 8$, $1 \leq k \leq n/4$, and G is a subgroup of S_n with $[S_n : G] < \binom{n}{k}$, then there is a set $X \subseteq [n]$ with $|X| < k$ such that $A_{(X)} \leq G$.

where $A_{(X)}$ denotes the group $\{\pi \in A_n : \pi(i) = i \text{ for all } i \in X\}$

Supports of Gates

Theorem

If $n > 8$ and $1 \leq k \leq n/4$, and G is a subgroup of S_n with $[S_n : G] < \binom{n}{k}$, then there is a set $X \subseteq [n]$ with $|X| < k$ such that $A_{(X)} \leq G$.

If $(C_n)_{n\in\omega}$ is a family of *symmetric* circuits of size n^k, then for all sufficiently large n and gates g in C_n, there is a set $X \subseteq [n]$ with $|X| \leq k$ such that $A_{(X)} \leq \mathrm{Stab}(g)$.

It follows that if *any odd* permutation of $[n]$ that fixes X pointwise, also fixes g, then $S_{(X)} \leq \mathrm{Stab}(g)$, so X is a support of g.

where $S_{(X)}$ denotes the group $\{\pi \in S_n : \pi(i) = i \text{ for all } i \in X\}$

Supports of Gates

Some odd permutation of $[n]$ that fixes X pointwise, also fixes g. $(*)$

We can prove, by induction on the depth of g in the circuit C_n that this must be the case.

It is clearly true for input gates $R(\overline{a})$, as any permutation that fixes $\overline{a}$ fixes the gate.

Let g be a gate such that $(*)$ is true for all gates that are inputs to g.

Since g computes a *symmetric* Boolean function, and C_n is *rigid*, any $\pi \in S_n$ that fixes the inputs to g *setwise*, fixes g.
Let H be the set of inputs to g. By induction hypothesis, they all have a support of size at most k

Supports of Gates

Some odd permutation of $[n]$ that fixes X pointwise, also fixes g. $(*)$

Let g be a gate such that $(*)$ is true for all gates in H, but false for g

For any $i, j \in [n] \setminus X$, the permutation $(i\ j)$ moves g, so moves some $h \in H$.

$$[n] \setminus X \subseteq \bigcup_{h \in H} \mathrm{sp}(h)$$

We can then find $\frac{n-k}{k}$ elements of H with *pairwise disjoint* support.

This gives us $\frac{n-k}{k}$ distinct permutations $(i\ j)$ which we can *independently* combine to show that the orbit of g has size at least $2^{(n-k)/k}$.

Support Theorem

In *polynomial size* symmetric circuits, all gates have (stabilizer groups with) *small* support:

Theorem

For any $1 > \epsilon \geq \frac{2}{3}$*, let* C *be a symmetric* s*-gate circuit over* $[n]$ *with* $n \geq 2^{\frac{56}{\epsilon^2}}$*, and* $s \leq 2^{n^{1-\epsilon}}$*. Then every gate* g *of* C *has a support of size at most* $\frac{33}{\epsilon}\frac{\log s}{\log n}$*.*

We write $\mathrm{sp}(g)$ for the small support of g given by this theorem and note that it can be computed in polynomial time from a symmetric circuit C.

Translating Symmetric Circuits to Formulas

Given a polynomial-time function $n \mapsto C_n$ that generates symmetric circuits:

1. There are formulas of FP interpreted on $([n], <)$ that define the structure C_n.
2. We can also compute in polynomial time (and therefore in FP on $([n], <)$) $\mathrm{sp}(g)$ for each gate g.
3. For an input structure $\mathbb{A}$ and an assignment $\gamma : [n] \to \mathbb{A}$ of the inputs of C_n to elements of $\mathbb{A}$, whether g is made true depends only on $\gamma(\mathrm{sp}(g))$.
4. We define, by induction on the structure of C_n, the set of tuples $\Gamma(g) \subseteq \mathbb{A}^{\mathrm{sp}(g)}$ that represent assignments γ making g true.
5. This inductive definition can be turned into a formula (of FP for a Boolean circuit, of FPC for one with threshold gates.)

Circuits and Pebble Games

If C is a symmetric circuit on n-vertex graphs such that every gate of C has a support of size at most k, and $\mathbb{A}$ and $\mathbb{B}$ are graphs such that $\mathbb{A} \equiv^{C^{2k}} \mathbb{B}$ then:

C accepts $\mathbb{A}$ if, and only if, C accepts $\mathbb{B}$.

As a consequence, if $\mathcal{C}$ is a class of structures of *counting width* $k : \mathbb{N} \to \mathbb{N}$, then any family of symmetric circuits accepting $\mathcal{C}$ has size $\Omega(n^k)$.

at least for $k \leq \frac{n}{\log n}$.