A Logical Characterisation of the Computational Complexity Class BPP

Phillip Kaye *

September 14, 2001

Abstract

We give a characterisation of the class of languages decided by a probabilistic Turing machine in polynomial time (BPP) in terms of formulas in first-order logic with a least fixed point operator, and a special relation-counting operator which we define. We show that BPP is equivalent (modulo a suitable encoding of logical structures) to the set of finite structures which satisfy some such formula having a specific semantic property.

1 The complexity class BPP

A probabilistic Turing machine is a Turing machine with the addition of a read-only random oracle tape. Before the start of the computation, the oracle tape contains an (infinite) random binary string. At each step of the computation the oracle tape head reads a random bit, and moves one position to the right. Formally, such a machine is defined as follows:

Definition 1 *A* Probabilistic Turing Machine (PTM) *is a quintuple* $\langle Q, \Sigma, q_0, F, \delta \rangle$ *where*

Q *is a finite set of* states.

Σ *is a finite set of input symbols, and* $\{0,1\} \in \Sigma$.

$q_0 \in Q$ *is the initial state.*

$F \subseteq Q$ *is a set of accepting states.*

δ *is a transition function,* $\delta : Q \times \Sigma \times \{0,1\} \mapsto Q \times \Sigma \times \{L, R\}$.

A PTM M is said to *decide* a language $L \subseteq \Sigma^*$ with error probability ϵ if, for any input x:

$x \in L \Longrightarrow M$ accepts with probability $\geq 1 - \epsilon$,

*prkaye@cacr.math.uwaterloo.ca, Department of Combinatorics & Optimization, University of Waterloo, Waterloo, ON, N2T 2L1, Canada.

$x \notin L \Longrightarrow M$ rejects with probability $\geq 1 - \epsilon$.

Definition 2 *The class* BPP *is the class of languages decided by a probabilistic Turing machine with error probability* $\epsilon = \frac{1}{2^n}$.

2 FO(LFP)

A formula in first order logic (FO) contains *nonlogical symbols* as well as *logical symbols*. The nonlogical symbols are $\{=, \wedge, \vee, \neg, \rightarrow, \forall, \exists, V\}$ where V is an infinite set of variables, and are common to every first order language. Parentheses are often written in first order formulas for clarity but are not, strictly speaking, symbols of first order logic. A particular first order language $\mathcal{L}$ is determined by specifying a *vocabulary* of nonlogical symbols:

$\mathcal{C}$ a set of *constant* symbols.

$\mathcal{R}$ a set of *relation* symbols (each having an associated arity).

$\mathcal{F}$ a set of *function* symbols (each having an associated arity).

For our purposes it is sufficient to restrict attention to vocabularies with $\mathcal{F} = \varnothing$, sometimes called *relational* vocabularies. A first order language $\mathcal{L}$ is the set of formulas built from the nonlogical symbols and the vocabulary for $\mathcal{L}$ in a manner syntactically consistent with the inductive definition of first order formulas. This definition can be found in any introductory logic text (see for example ??). For convenience we will use $\mathcal{L}$ to denote both a vocabulary, and the first-order language determined by the vocabulary.

For a given relational vocabulary $\mathcal{L}$, a *structure appropriate for* $\mathcal{L}$ (sometimes called an $\mathcal{L}$-structure) consists of

a. A nonempty set called the *universe* of $\mathcal{A}$. We will denote structures by calligraphic letters such as $\mathcal{A}$, and the universe of $\mathcal{A}$ by $|\mathcal{A}|$.

b. For each constant symbol $c \in \mathcal{C}$, an specific element $c^{\mathcal{A}} \in |\mathcal{A}|$.

c. For each relation symbol $R \in \mathcal{R}$, an actual relation $\mathcal{R}^{\mathcal{A}}$ over $|\mathcal{A}|$, of the appropriate arity. If R has arity a, then $R^{\mathcal{A}} \subseteq |\mathcal{A}|^a$.

A first order formula $\varphi \in \mathcal{L}$ can be interpreted as expressing some property of $\mathcal{L}$-structures. For a specific $\mathcal{L}$-structure $\mathcal{A}$, we can ask whether $\mathcal{A}$ possesses the property expressed by φ. When it does, we say that $\mathcal{A}$ *satisfies* φ, or equivalently that φ is *true* for $\mathcal{A}$. We denote this situation by $\mathcal{A} \models \varphi$. The inductive definition of the semantics of first order logic can be found in a logic text.

Sometimes it is convenient to consider a first order formula $\varphi \in \mathcal{L}$, having free variables $\overline{x} = x_1, \ldots x_k$, *evaluated for* a particular k-tuple of elements $\overline{a} = a_1 \ldots a_k$ from an $\mathcal{L}$-structure. The notation $\varphi(\overline{a})$ is often used to denote the formula φ with all occurrences of $x_1, \ldots x_k$ replaced by $a_1 \ldots a_k$.

Suppose $\varphi(R, x_, \ldots, x_k)$ is a first order formula in $\mathcal{L}$ with free (unquantified) variables $x_1, \ldots, x_k$ and having one or more occurrences of the new k-ary relation symbol $R \notin \mathcal{L}$. For any $\mathcal{L}$-structure $\mathcal{A}$, φ defines a map from relations on $|\mathcal{A}|$ to relations on $|\mathcal{A}|$:

$$\varphi^{\mathcal{A}}(R) = \{\langle a_1, a_2, \ldots a_k\rangle : \mathcal{A} \models \varphi(R, a_1, a_2, \ldots a_k)\},$$

where $a_1, a_2, \ldots a_k$ range over $|\mathcal{A}|$. The *least fixed point* of $\varphi^{\mathcal{A}}$ is the minimal relation T such that $\varphi^{\mathcal{A}}(T) = T$. We can define an operator on first order formulas, known as the *least fixed point* (LFP) operator. For $\varphi \in \mathcal{L}$ with free variables $x_1, \ldots, x_k$ and occurrences of a k-ary relation symbol $R \notin \mathcal{L}$, $(\mathrm{LFP}_{R,x_1,\ldots,x_k}\varphi)$ is a k-ary relation symbol, which when evaluated for a structure $\mathcal{A}$ is the least fixed point relation of $\varphi^{\mathcal{A}}$, with respect to $R, x_1, \ldots, x_k$.

Define FO(LFP) be the set of all first order formulas from any vocabulary $\mathcal{L}$, having any number of occurrences of the LFP operator.

3 FO(LFP,FRAC)

We begin by giving a syntactic definition of expressions in FO(LFP,FRAC).

Definition 3 *A sentence in* FO(LFP,FRAC) *is a logical expression of the form*

$$\mathrm{FRAC}(R)\varphi$$

where φ is a sentence in FO(LFP) from a first order language $\mathcal{L}$, and $R \notin \mathcal{L}$ is a new relation symbol of arity a.

For a finite $\mathcal{L}-$structure $\mathcal{A}$, and an actual a-ary relation $R^{\mathcal{A}}$ over the universe of $\mathcal{A}$, let $\mathcal{A} \cup \{R^{\mathcal{A}}\}$ be the $\mathcal{L} \cup \{R\}$-structure obtained by augmenting $\mathcal{A}$ with $R^{\mathcal{A}}$. Let $n = ||\mathcal{A}||$ be the cardinality of the universe of $\mathcal{A}$. Then the *value* of $\Psi = \mathrm{FRAC}(R)\varphi$ *on* structure $\mathcal{A}$ is given by

$$\Psi^{\mathcal{A}} = \frac{|\{R^{\mathcal{A}} : (\mathcal{A} \cup \{R^{\mathcal{A}}\}) \models \varphi\}|}{2^{n^a}}.$$

Noting that there are exactly 2^{n^a} distinct a-ary relations $R^{\mathcal{A}}$, we see that $\Psi^{\mathcal{A}}$ evaluates to the fraction of all a-ary relations over the universe of $\mathcal{A}$ which make φ true. The following definition completes the description of the semantics of FO(LFP,FRAC).

Definition 4 *Let Ψ be a sentence in* FO(LFP,FRAC) *from a first order language $\mathcal{L}$. For any finite $\mathcal{L}-$structure $\mathcal{A}$, we say that $\mathcal{A}$* satisfies Ψ *(or equivalently that Ψ is* true *for $\mathcal{A}$), written $\mathcal{A} \models \Psi$, if $\Psi^{\mathcal{A}} \geq 1 - \frac{1}{2^n}$, where $n = ||\mathcal{A}||$.*

Definition 5 *Let $\mathcal{L}$ be a relational vocabulary. Define*

$$\mathrm{FIN}(\mathcal{L}) = \{\text{finite } \mathcal{L}\text{-structures } \mathcal{A} \text{ having an ordering on the universe}\}$$

Definition 6 *The* spectrum *of a sentence* $\Psi \in$ FO(LFP,FRAC) *is defined as follows:*

$$\mathrm{SPEC}(\Psi) = \{\mathcal{A} \in \mathrm{FIN}(\mathcal{L}) : \mathcal{A} \models \Psi\}$$

Note that any such spectrum will be closed under isomorphism of $\mathcal{L}$-structures.

Definition 7 *We say that a sentence* Ψ *in* FO(LFP,FRAC) *from a first order language* $\mathcal{L}$ *is* decisive *provided that there do not exist any* $\mathcal{A} \in \mathrm{FIN}(\mathcal{L})$ *for which* $\frac{1}{2^n} < \Psi^{\mathcal{A}} < 1 - \frac{1}{2^n}$, *where* $n = ||\mathcal{A}||$.

4 A Logical Characterisation of BPP

We begin by defining an encoding that allows us to relate sets of structures to languages decided by a Turing machine.

Definition 8 *Suppose* $\mathcal{A} \in \mathrm{FIN}(\mathcal{L})$, *for some finite vocabulary* $\mathcal{L}$, *and let* $n = ||\mathcal{A}||$. *Let* $e(n)$ *be the binary representation of* n. *For each constant* $c_i^{\mathcal{A}} \in \mathcal{C}$, *let* $e(c_i^{\mathcal{A}})$ *be the index of the corresponding element of* $|\mathcal{A}|$ *in binary (by definition of* $\mathrm{FIN}(\mathcal{L})$, $|\mathcal{A}|$ *is ordered; assume the elements of* $|\mathcal{A}|$ *are indexed by the integers 0 through* $n-1$*). For each relation* $R_i^{\mathcal{A}} \in \mathcal{R}$ *of arity* a_i *let* $e(R_i^{\mathcal{A}})$ *be a binary string of length* $||\mathcal{A}||^{a_i}$ *with the* i^{th} *bit equal to* 1 *if the corresponding element of* $|\mathcal{A}|^{a_i}$ *is in the relation, and equal to zero otherwise. Then define the* encoding *of* $\mathcal{A}$ *by*

$$e(\mathcal{A}) = e(n).e(c_1^{\mathcal{A}}) \dots e(c_{|\mathcal{C}|}^{\mathcal{A}}).e(R_1^{\mathcal{A}}) \dots e(c_{|\mathcal{R}|}^{\mathcal{A}})$$

where $b_i.b_j$ *denotes the concatenation of binary strings* b_i *and* b_j. *Note that* $|e(\mathcal{A})| = p(n)$ *for some polynomial* p. *Also note that, up to a re-ordering of the elements of the universe,* $e(\mathcal{A}) = e(\mathcal{B})$ *iff* $\mathcal{A}$ *and* $\mathcal{B}$ *are isomorphic* $\mathcal{L}$ *structures. For* $S \subseteq \mathrm{FIN}(\mathcal{L})$ *define* $E(S) = \{e(\mathcal{A}) : \mathcal{A} \in S\}$.

Theorem 1 *Assume* $S \subseteq \mathrm{FIN}(\mathcal{L})$ *and that* S *is closed under isomorphism. Then* $S = \mathrm{SPEC}(\Psi)$ *for some decisive sentence* $\Psi \in$ FO(LFP,FRAC) *iff* $E(S) \in$ BPP.

The remainder of this section is dedicated to proving Theorem 1. We have the following lemma due to Immerman [I87].

Lemma 1 *Let* $\mathcal{L}$ *be a vocabulary, and let* $S \subseteq \mathrm{FIN}(\mathcal{L})$*be closed under isomorphism. We have that*

$$S = \{\mathcal{A} \in \mathrm{FIN}(\mathcal{L}) : \mathcal{A} \models \varphi\}$$

for some $\varphi \in$ FO(LFP) *from* $\mathcal{L}$, *iff*

$$\mathrm{E(S)} \in P.$$

We now prove the first direction of theorem 1. Suppose $S \subseteq \mathrm{FIN}(\mathcal{L}) = \mathrm{SPEC}(\Psi)$ for some $\Psi \in$ FO(LFP,FRAC). Let $\Psi = \mathrm{FRAC}(R)\varphi$, where R is a-ary. We construct a PTM M with $\Sigma = \{0, 1\}$ which, on input $x = e(\mathcal{A})$ for some

$\mathcal{A} \in \text{FIN}(\mathcal{L})$ decides whether $e(\mathcal{A}) \in E(S)$. As before, let $n = ||\mathcal{A}||$. The machine M can generate a random a-ary relation $\hat{R}^{\mathcal{A}}$ over the universe of $\mathcal{A}$. This is accomplished by reading a series of n^a random oracle tape cells, where the i^{th} element of $|\mathcal{A}|^a$ is determined to be in the relation iff the i^{th} random oracle bit is 1. The PTM can then test whether $\mathcal{A} \cup \hat{R}^{\mathcal{A}} \vDash \varphi$, accepting if the test succeeds, and rejecting if it fails. By Lemma 1, this test can be performed in polynomial time. If $\mathcal{A} \in S$ then $\mathcal{A} \vDash \Psi$, and by definition of satisfaction for $\Psi \in$ FO(LFP,FRAC), the above test will succeed for at least $\frac{1}{2^n}$ of the possible random relations generated. Furthermore, since Ψ is decisive by assumption, if $\mathcal{A} \notin S$, then the test fails with probability at least $\frac{1}{2^n}$. So the error probability for the PTM is $\frac{1}{2^n}$, and we have $E(S) \in$ BPP. This proves one direction of theorem 1.

Now assume a language $E(S)$ (for some $S \in \text{FIN}(\mathcal{L})$) is in BPP. Note that any $L \subseteq \{0,1\}^*$ is $E(S)$ for some S in a vocabulary whose sructures are isomorphic to binary strings. Let M be a PTM deciding $E(S)$ with error probability $\frac{1}{2^n}$. On input $x = e(\mathcal{A})$ (where $\mathcal{A}$ is some structure in $\text{FIN}(\mathcal{L})$, suppose M halts in time bounded by n^a. Without loss of generality, we can assume that M halts after exactly n^a steps for all possible random sequences of bits in the first n^a cells of the oracle tape. We briefly justify this assumption. The machine M may be modified to include a new special holding state corresponding to each halting state. When M is in such a holding state, it will continue to read the oracle tape, doing nothing else. The transition function is modified so that if a particular random sequence would cause M to halt in fewer than n^a steps, M is put into the holding state until time n^a. At time n^a, M moves from any holding state to the corresponding halting state.

Define a new vocabulary $\mathcal{L}' = \mathcal{L} \cup \{P\}$ where P is a relation symbol of arity a. For a given $\mathcal{L}'$-structure $\mathcal{A}'$ with $||\mathcal{A}'|| = n$, $e(P^{\mathcal{A}})$ is a binary string of length n^a (see definition 8). The oracle string for a computation of M on input $e(\mathcal{A})$ is a binary string of length n^a. A computation of M on input $e(\mathcal{A})$ with oracle string equal to p can thus be represented as an $\mathcal{L}'$-structure $\mathcal{A}' = \mathcal{A} \cup P^{\mathcal{A}'}$, where $p = e(P^{\mathcal{A}'})$. Then M has an equivalent polynomial-time deterministic Turing machine M' which takes inputs $e(\mathcal{A}') = \langle e(\mathcal{A}), e(P^{\mathcal{A}'})\rangle$. Machine M' will accept on input $e(\mathcal{A}')$ iff M accepts on input $e(\mathcal{A})$ when $e(P^{\mathcal{A}'})$ appears on it's oracle tape. Let X' be the set of inputs $e(\mathcal{A}')$ for which M' accepts. Since M' runs in polynomial time, by Lemma 1 there exists a $\varphi \in$ FO(LFP) from first order language $\mathcal{L}'$, for which $\mathcal{A}' \vDash \varphi$ iff $e(\mathcal{A}') = \langle e(\mathcal{A}), e(P^{\mathcal{A}'})\rangle \in X'$. Suppose $\mathcal{A} \in S$. Then M accepts for at least $1 - \frac{1}{2^n}$ of the possible oracle strings, and so the deterministic machine M' accepts for at least $1 - \frac{1}{2^n}$ of its 2^a possible inputs $\langle e(\mathcal{A}), e(P^{\mathcal{A}'})\rangle$. Consider

$$\Psi = \text{FRAC}(P)\varphi$$

where $\varphi \in$ FO(LFP) is as above. By definition, we have that $\Psi^{\mathcal{A}} \geq 1 - \frac{1}{2^n}$, and $\mathcal{A} \vDash \Psi$. Conversely, suppose $A \notin S$. Then M' accepts for less than $\frac{1}{2^n}$ of the 2^a possible inputs. Then by definition $\Psi^{\mathcal{A}} < \frac{1}{2^n}$. So $\mathcal{A} \nvDash \Psi$ and furthermore Ψ is decisive. $\square$

5 A Proof that BPP $\subseteq$ PH

Definition 9 *A formula of* second order logic *SO is a logical formula of the form*

$$Q_1 R_1 \ldots Q_s R_s \varphi$$

where s is any integer, $Q_i \in \{\forall, \exists\}$, and $R_i \notin \mathcal{R}$ are new relation symbols.

The semantics of second order logic are defined in the obvious way. Each existentially quantified relation symbol is interpreted as asserting the existence of an actual relation in the structure making φ true, and each universally quantified relation symbol is interpreted as asserting that every possible relation in the structure makes φ true.

The complexity class PH is defined in terms of *oracle Turing machines.*

Definition 10 *An* oracle Turing machine *is a Turing machine equipped with an* oracle *for a given language B. At each step an oracle Turing Machine M can query the oracle, by writing the string in question on an oracle tape. In one time step, the oracle will put M into an* oracle yes *or an* oracle no *state, giving the answer to the oracle query. A machine M with an oracle for a language B is denoted by $M(B)$. For a complexity class $\mathcal{C}$, let $\mathrm{NP}(\mathcal{C})$ be the class of languages decided by nondeterministic oracle Turing machines using an oracle $B \in \mathcal{C}$.*

Definition 11 *For integers n, the computational complexity classes Σ_n^{P} and Π_n^{P} are defined as follows:*

1. $\Sigma_0^{\mathrm{P}} = \Pi_0^{P} = \mathrm{P}$
2. $\Sigma_{n+1}^{P} = \mathrm{NP}(\Sigma_n^{\mathrm{P}})$
3. $\Pi_{n+1}^{\mathrm{P}} = co - \Sigma_{n+1}^{\mathrm{P}}$

The polynomial-time Hierarchy (PH) *is defined as*

$$\mathrm{PH} = \cup_{n \geq 0} \Sigma_n^{\mathrm{P}}$$

The following logical characterisation of PH is due to Immerman [ ].

Lemma 2 *Let $\mathcal{L}$ be a vocabulary, and let $S \subseteq \mathrm{FIN}(\mathcal{L})$be closed under isomorphism. We have that*

$$S = \{\mathcal{A} \in \mathrm{FIN}(\mathcal{L}) : \mathcal{A} \vDash \Phi\}$$

for some $\Phi \in \mathrm{SO}$ from $\mathcal{L}$, iff

$$\mathrm{E(S)} \in PH.$$

Two further of definitions will be useful.

Definition 12 *Two logical formulas Φ and Ψ from a vocabulary $\mathcal{L}$ are said to be* logically equivalent*, written $\Phi \equiv \Psi$, if for all $\mathcal{L}$-structures $\mathcal{A}$, $\mathcal{A} \vDash \Phi \iff \mathcal{A} \vDash \Psi$.*

Definition 13 *Let $\mathcal{A} \in \mathrm{FIN}(\mathcal{L})$. Let $R_1^{\mathcal{A}}$, and $R_2^{\mathcal{A}}$ be two a-ary relations over $|\mathcal{A}|$, and let $\overline{x}$ be an a-tuple of variables. Consider the new a-ary relation over $|\mathcal{A}|$ defined by $\varphi_{\oplus} \equiv (R_1\overline{x} \wedge \neg R_2\overline{x}) \vee (\neg R_1\overline{x} \wedge R_2\overline{x})$ (i.e. for an a-tuple $\overline{a}$ of elements from $|\mathcal{A}|$, $\overline{a}$ is taken to be in the new relation iff $\mathcal{A} \models phi_{\oplus}(\overline{a})$). For any pair $R_1^{\mathcal{A}}$, and $R_2^{\mathcal{A}}$, $R_1^{\mathcal{A}} \oplus R_2^{\mathcal{A}}$ denotes the new relation defined by $\varphi_{\oplus}$ in this way. For any pair of a-ary relation symbols R_1, R_2, $R_1 \oplus R_2$ can be used as a new a-ary relation symbol in logical formulas. For an a-tuple $\overline{x}$, $(R_1 \oplus R_2)\overline{x}$ is logically equivalent to $(R_1\overline{x} \wedge \neg R_2\overline{x}) \vee (\neg R_1\overline{x} \wedge R_2\overline{x})$.*

Definition 14 *For a set $\overline{R}^{\mathcal{A}}$ of a-ary relations and another fixed a-ary relation $T^{\mathcal{A}}$, define $\overline{R}^{\mathcal{A}} \oplus T^{\mathcal{A}} = \{R^{\mathcal{A}} \oplus T^{\mathcal{A}} : R^{\mathcal{A}} \in \overline{R}^{\mathcal{A}}\}$, and call $\overline{R}^{\mathcal{A}} \oplus T^{\mathcal{A}}$ the* translation *of $\overline{R}^{\mathcal{A}}$ by $T^{\mathcal{A}}$.*

For a fixed a-ary relation $T^{\mathcal{A}}$, $\oplus T^{\mathcal{A}}$ can be thought of as a one-to-one function on relations over $\mathcal{A}$. The following properties of this function are easily checked. Firstly, for any set $\overline{R}^{\mathcal{A}}$ of a-ary relations, $|\overline{R}^{\mathcal{A}}| = |\overline{R}^{\mathcal{A}} \oplus T^{\mathcal{A}}|$. Secondly, for any a-ary relation $R^{\mathcal{A}}$, $R^{\mathcal{A}} \oplus T^{\mathcal{A}} = S^{\mathcal{A}}$ iff $S^{\mathcal{A}} \oplus T^{\mathcal{A}} = R^{\mathcal{A}}$. Finally, if $R^{\mathcal{A}}$ is a random a-ary relation over $\mathcal{A}$, then $R^{\mathcal{A}} \oplus T^{\mathcal{A}}$ is also random.

Definition 15 *Let $\mathcal{A} \in \mathrm{FIN}(\mathcal{L})$, and let $\overline{R}^{\mathcal{A}}$ be a set of a-ary relations over $|\mathcal{A}|$. Let $\overline{T} = \{T_1^{\mathcal{A}}, \ldots T_m^{\mathcal{A}}\}$ be a set of translations of $\overline{R}^{\mathcal{A}}$ (i.e. another set of a-ary relations over $|\mathcal{A}|$). Let B be a fixed a-ary relation over $|\mathcal{A}|$. We say $\overline{T}$* covers *B if $B \in \overline{R} \oplus T_i^{\mathcal{A}}$ for some $1 \leq i \leq m$.*

The following theorem can now be proven using Theorem 1 and Lemma 2.

Theorem 2

$$\mathrm{BPP} \subseteq \mathrm{PH}.$$

Proof: It suffices to show that for every decisive formula $\Psi \in \mathrm{FO(LFP,FRAC)}$ there exists a logically equivalent $\Phi \in \mathrm{SO}$. Let $\Psi = \mathrm{FRAC}(R)\varphi$ be an arbitrary formula in FO(FRAC,LFP) from a vocabulary $\mathcal{L}$, where $\mathrm{arity}(R) = a$. For any finite $\mathcal{L}$-structure $\mathcal{A}$, define $\mathrm{GOOD}(\mathcal{A}) = \{R^{\mathcal{A}} : \mathcal{A} \cup \{R^{\mathcal{A}}\} \models \varphi\}$. Letting $n = ||\mathcal{A}||$, note that $|\mathrm{GOOD}(\mathcal{A})| = 2^{n^a}\Psi^{\mathcal{A}}$. Note also that the set of all a-ary relations over a universe $|\mathcal{A}|$ is the power set of all a-tuples over $|\mathcal{A}|$, $\mathcal{P}(|\mathcal{A}|^a)$.

Claim: There exists a set of translations $\overline{T}^{\mathcal{A}} \subseteq \mathrm{GOOD}(\mathcal{A})$, of $\mathrm{GOOD}(\mathcal{A})$ with $|\overline{T}^{\mathcal{A}}| = n^a$, that covers all of $\mathcal{P}(|\mathcal{A}|^a)$ iff $\mathcal{A} \models \Psi$.

Proof of claim: Suppose $\mathcal{A} \models \Psi$. Consider a random set of n^a translations $\overline{T}^{\mathcal{A}} = \{T_1^{\mathcal{A}}, T_2^{\mathcal{A}}, \ldots T_{n^a}^{\mathcal{A}}\} \subseteq \mathrm{GOOD}(\mathcal{A})$. Fix a relation $B^{\mathcal{A}} \in \mathcal{P}(|\mathcal{A}|^a)$. We want to show that $B^{\mathcal{A}} \in \mathrm{GOOD}(\mathcal{A}) \oplus T_j^{\mathcal{A}}$, for some $j \leq n^a$. Note that $B^{\mathcal{A}} \in \mathrm{GOOD}(\mathcal{A}) \oplus T_j^{\mathcal{A}}$ iff $B^{\mathcal{A}} \oplus T_j^{\mathcal{A}} \in \mathrm{GOOD}(\mathcal{A})$. Since $B^{\mathcal{A}} \oplus T_j^{\mathcal{A}}$ is as random as $T_j^{\mathcal{A}}$, and since $\mathcal{A} \models \Psi$ by assumption, we have $\mathrm{prob}[B^{\mathcal{A}} \oplus T_j^{\mathcal{A}} \notin \mathrm{GOOD}(\mathcal{A})] = \mathrm{prob}[B^{\mathcal{A}} \notin \mathrm{GOOD}(\mathcal{A}) \oplus T_j^{\mathcal{A}}] \leq \frac{1}{2^n}$. So the probability that some relation in $\mathcal{P}(|\mathcal{A}|^a)$ is not covered is at most $2^{-(n-1)n^a} < 1$. So a random sequence of n^a translations has a positive probability that it covers all of $\mathcal{P}(|\mathcal{A}|^a)$, and therefore such a set of translations must exist. Conversely, suppose $\mathcal{A} \not\models \Psi$. Then $|\mathrm{GOOD}(\mathcal{A})| \leq \frac{1}{2^n} 2^{n^a} = 2^{n^a - n}$. If $|\overline{T}^{\mathcal{A}}| = n^a$, then

$$|\overline{T}^{\mathcal{A}} \oplus \mathrm{GOOD}(\mathcal{A})| \leq |\overline{T}| \cdot |\mathrm{GOOD}(\mathcal{A})| \leq n^a \cdot 2^{n^a - n} < 2^{n^a} = |\mathcal{P}(|\mathcal{A}|^a)|$$

and so translations of $\mathrm{GOOD}(\mathcal{A})$ exists that covers all of $\mathcal{P}(|\mathcal{A}|^a)$.

By definition, $B^{\mathcal{A}} \oplus T_j \in \mathrm{GOOD}(\mathcal{A})$ iff $\mathcal{A} \models \varphi(B \oplus T_j)$. By the claim, for each n, Ψ is equivalent to the following logical formula which asserts "there exist a set of n^a relations $T_1^{\mathcal{A}}, \ldots T_{n^a}^{\mathcal{A}}$ that covers $|\mathcal{A}|^a$":

$$\exists T_1 \exists T_2 \ldots \exists T_{n^a} \forall B \left[\varphi(B \oplus T_1) \vee \varphi(B \oplus T_2) \vee \cdots \vee \varphi(B \oplus T_{n^a})\right].$$

What is required is a *single* second-order sentence Φ which is logically equivalent to Ψ *for all* n. A set of n^a a-ary relations $T_1^{\mathcal{A}}, T_2^{\mathcal{A}} \ldots T_{n^a}^{\mathcal{A}}$ has an equivalent $2a$-ary relation $T^{\mathcal{A}}$, defined as follows. Recall that an a-tuple $\overline{y}^{\mathcal{A}}$ over $|\mathcal{A}|$ can encode any number $1 \leq y \leq n^a$. Let $\overline{x}^{\mathcal{A}}$ and $\overline{y}^{\mathcal{A}}$ be two a-tuples over $|\mathcal{A}|$. The $2a$-ary relation $T^{\mathcal{A}}$ is defined as: $\langle \overline{x}^{\mathcal{A}}, \overline{y}^{\mathcal{A}} \rangle \in T^{\mathcal{A}}$ iff the a-tuple $\overline{x}^{\mathcal{A}}$ is in the relation $T_y^{\mathcal{A}}$ above.

The following second-order formula Φ is logically equivalent to Ψ, where T is a $2a$-ary relation symbol, S and B are a-ary relation symbols and $\overline{a}$ and $\overline{b}$ are a-tuples of variables:

$$\Phi \equiv \exists T \forall B \exists S \exists \overline{a} \forall \overline{b} \left((T\overline{a}\overline{b} \longleftrightarrow S\overline{b}) \wedge \varphi(B \oplus S) \right)$$

6 Acknowledgements

I would like to gratefully acknowledge the support and guidance of my supervisors, Professors Ross Willard and Michele Mosca at the University of Waterloo, as well as the helpful suggestions I received from Professor John Watrous at the University of Calgary.