

Full Abstraction for Reduced ML

Andrzej S. Murawski* and Nikos Tzevelekos

Oxford University Computing Laboratory
Wolfson Building, Parks Road, Oxford OX1 3QD, UK

Abstract. We present the first effectively presentable fully abstract model for Stark’s Reduced ML, the paradigmatic higher-order programming language combining call-by-value evaluation and integer-valued references. The model is constructed using techniques of nominal game semantics. Its distinctive feature is the presence of carefully restricted information about the store in plays, combined with conditions concerning the participants’ ability to distinguish reference names. This leads to an explicit characterization of program equivalence.

1 Introduction

Reduced ML is a programming language introduced by Stark [22] as part of his investigations into generative aspects of programming languages. It combines higher-order functions with integer references and is defined simply by extending the call-by-value λ -calculus with primitives for integer and reference manipulation. Despite its economy, Reduced ML manages to embody several important paradigms (imperative programming, functional programming, call-by-value evaluation, scope extrusion), which makes it an attractive object for theoretical study. On the other hand, research into it offers wide scope for applicability, as Reduced ML is intimately related to Standard ML [15] and, in fact, has been designed with faithfulness to the latter in mind.

The first steps in the semantic analysis of Reduced ML were taken by Stark, who has identified a matching categorical framework and considered example categories, albeit without a general full abstraction result¹. Further progress was possible with the arrival of game semantics [4, 9, 19]. Although the first papers concerned call-by-name computation, attention soon turned to the call-by-value framework [8, 6]. In particular, Abramsky and McCusker presented a fully abstract model for a language called RML [6], which is essentially Reduced ML extended with the “bad-variable” constructor `mkvar`. Its presence is a consequence of adopting Reynolds’s principle of modelling references as objects with read and write methods [21]. Thus, `mkvar` allows one to define terms of reference type that need not correspond to actual memory locations. Unfortunately, this affects

* Supported by an EPSRC Advanced Research Fellowship (EP/C539753/1).

¹ A denotational model of a programming language is *fully abstract* iff equality of denotations coincides with program equivalence. Programs are *equivalent* iff they can be used interchangeably without observable differences.

the induced notion of program equivalence, so the full abstraction result of [6] does not apply to Reduced ML. More precisely, it can fail at types containing occurrences of `int ref`. Typical counterexamples are the failures of equivalences between $x := !x$ and $()$ (the terminating command), or between $x := 1; x := 1$ and $x := 1$. In the former case the terms are inequivalent in RML, because x may be instantiated with a `mkvar`-object whose reading or writing method diverges, or causes side effects. Similarly, in the latter case, an assignment to a `mkvar`-object might trigger a side effect that effectively allows one to count how many assignments took place.

The “bad-variable” phenomenon, also present in the call-by-name setting, has inspired subsequent developments in game semantics. It turned out that, in the call-by-name framework, it could be circumvented by employing suitably crafted (pre)orders on plays [14, 18], but no result of this kind has been reported for call-by-value. However, an alternative and general approach to dealing with bad variables seems to have emerged in the form of *nominal game semantics* [10, 2, 23]. Nominal game semantics advocates a departure from Reynolds’s modelling rule and stipulates that reference types be modelled by names rather than objects. Using this approach, Laird showed a full abstraction result for a call-by-value language $\lambda\nu!$ with storable names rather than integers [10]. $\lambda\nu!$ turns out more expressive than Reduced ML in its ability to distinguish reference names and, consequently, the obvious adaptation of the model to Reduced ML results in a failure of full abstraction. This can be illustrated by the terms²

$$f : \text{int ref} \rightarrow \text{unit} \vdash \text{let } n_1 = \text{ref } 0 \text{ in let } n_2 = \text{ref } 0 \text{ in } ((fn_1); (n_2 := !n_1); n_2) : \text{int ref}.$$

and $f : \text{int ref} \rightarrow \text{unit} \vdash \text{let } n = \text{ref } 0 \text{ in } (fn); n : \text{int ref}$, which are equivalent in Reduced ML, but inequivalent³ in $\lambda\nu!$. This is because a $\lambda\nu!$ -context can detect the difference between n_1 from n_2 by storing the names and subsequently comparing them. In contrast, as our results confirm, the same effect cannot be achieved by a context belonging to Reduced ML.

Previous research into ML-like languages has also produced fully abstract game models for more significant extensions of Reduced ML, notably languages with higher-order references [3, 23]. The first of these models suffers from the “bad-variable” problem outlined above. The second one, while adaptable to Reduced ML, leans rather too heavily on quotienting in order to achieve full abstraction (information on local state and store update is too explicit in the intensional model and leads to substantial undesirable distinctions). Therefore, it does not lead to an explicit characterization of program equivalence, which is obtained in the present paper.

Our point of departure is the observation that, although a Reduced ML program will in general not be able to keep track of all the names it encounters during the course of interaction with another program, at any given execution point there is a subset of such names that the program may have access to. In game

² $\text{let } x = M \text{ in } N$ stands for the Reduced ML term $(\lambda x.N)M$.

³ Strictly speaking, the terms are not in $\lambda\nu!$, but the scenario can be easily recast in $\lambda\nu!$ by replacing `ref 0` with $\nu n.n$.

semantics, using the notion of P-view, we can describe this set conservatively as one consisting of names that occur in the current P-view as well as those that the program created itself. We call such names P-available. Intuitively, whenever a program returns a name, it will be P-available. A corresponding condition inside our model will be called *P-availability*.

As a consequence, since a program cannot have access to reference names that are not P-available, its immediate behaviour will be independent of the associated values kept in the store, because the program is simply unable to read them. This leads us to found our game model on justified sequences with partial information about the store, restricted to P-available names. Note that this form of representation also conveys the idea that the program might depend on former (possibly outdated) values of currently unavailable references, recorded when the references were still available.

Unfortunately, P-availability and partiality of store alone do not yet suffice to establish a definability result. As our example demonstrates, a Reduced ML context may be unable to distinguish some occurrences of names introduced by the environment. In game semantics, we can capture this oversight in concrete terms: two (occurrences of) such names are indistinguishable to the program iff they have never occurred within the same P-view. Consequently, regardless of whether such occurrences are the same or not, the program's behaviour should remain the same. We formalize this observation via a saturation condition, called *blindness*, and show that any finitary strategy subject to all the conditions discussed above is definable, i.e. is a denotation of a Reduced ML term. This naturally leads to a fully abstract model via the usual *intrinsic quotient* construction.

To obtain a more accessible account of program equivalence we next examine the structure of the quotient in more detail. Crucially, we observe that blind strategies are determined uniquely by plays in which the environment provides a fresh name each time the name cannot be related by the program to any existing names. We call such plays *strict*. Then, by symmetrizing the model, we eventually obtain an explicit characterization of equivalence: terms of Reduced ML are equivalent iff they induce the same *mutually strict complete protoplays* (complete plays where O plays only O-available names and in which stores are restricted to *mutually available* names).

For example, *each* of the two terms introduced above generates the following such plays

$$* \overbrace{n_1^{(n_1,0)} - *^{(n_1,k)} n_2^{(n_2,k)}} ,$$

where k ranges over the set of integers. Hence, the terms are indeed equivalent.

Notes. A long version with proofs is available from the authors' webpages. We are grateful to Jim Laird for email discussions.

2 Reduced ML

Reduced ML is the call-by-value λ -calculus over the ground types `unit`, `int`, `int ref` augmented with basic commands (termination, divergence), primitives for inte-

$$\begin{array}{c}
\frac{}{u, \Gamma \vdash () : \text{unit}} \quad \frac{}{u, \Gamma \vdash \Omega : \text{unit}} \quad \frac{i \in \mathbb{Z}}{u, \Gamma \vdash i : \text{int}} \quad \frac{l \in u}{u, \Gamma \vdash l : \text{int ref}} \quad \frac{(x : \theta) \in \Gamma}{u, \Gamma \vdash x : \theta} \\
\frac{u, \Gamma \vdash M_1 : \text{int} \quad u, \Gamma \vdash M_2 : \text{int}}{u, \Gamma \vdash M_1 \oplus M_2 : \text{int}} \quad \frac{u, \Gamma \vdash M : \text{int} \quad u, \Gamma \vdash N_0 : \theta \quad u, \Gamma \vdash N_1 : \theta}{u, \Gamma \vdash \text{if } M \text{ then } N_1 \text{ else } N_0 : \theta} \\
\frac{u, \Gamma \vdash M : \text{int ref}}{u, \Gamma \vdash !M : \text{int}} \quad \frac{u, \Gamma \vdash M : \text{int ref} \quad u, \Gamma \vdash N : \text{int}}{u, \Gamma \vdash M := N : \text{unit}} \quad \frac{u, \Gamma \vdash M : \text{int}}{u, \Gamma \vdash \text{ref } M : \text{int ref}} \\
\frac{u, \Gamma \vdash M : \theta \rightarrow \theta' \quad u, \Gamma \vdash N : \theta}{u, \Gamma \vdash MN : \theta'} \quad \frac{u, \Gamma \cup \{x : \theta\} \vdash M : \theta'}{u, \Gamma \vdash \lambda x^\theta. M : \theta \rightarrow \theta'}
\end{array}$$

Fig. 1. Syntax of Reduced ML.

ger arithmetic (constants, zero-test, binary integer functions) and reference manipulation (locations, dereferencing, assignment, memory allocation). The typing rules are given in Figure 1, where \mathcal{L} stands for a countable set of *locations*, u for a finite subset of \mathcal{L} , and \oplus for binary integer functions (e.g. $+$, $-$, $*$, $=$). Their precise choice is to some extent immaterial: for the full abstraction argument to hold it suffices to be able to compare integer variables with integer constants and act on the result. The same can be said about the lack of recursion, which can be added without affecting our results. Note that we did not include reference equality testing, because it is expressible [20]. For instance, one can define $\text{eq} : \text{int ref} \rightarrow \text{int ref} \rightarrow \text{int}$ as

$$\begin{array}{l}
\lambda x^{\text{int ref}}. \lambda y^{\text{int ref}}. \text{let } v_x = \text{ref } !x \text{ in} \\
\quad \text{let } v_y = \text{ref } !y \text{ in} \\
\quad \text{let } b = \text{ref } 0 \text{ in} \\
\quad (x := 0; y := 1; (\text{if } !x = 1 \text{ then } b := 1 \text{ else } ()); x := !v_x; y := !v_y; !b)
\end{array}$$

In the above and in what follows, we write $M; N$ for the term $(\lambda z^\theta. N)M$, where z does not occur in N and θ matches the type of M .

To define the operational semantics of Reduced ML, we need to introduce a notion of store. A *store* will simply be a function from a finite set of locations to \mathbb{Z} . We write $s(l \mapsto i)$ for the store obtained by updating s so that l is mapped to i (this may extend the domain of s). Given a store $s : \{l_1, \dots, l_n\} \rightarrow \mathbb{Z}$ and a term M we say that the pair (s, M) is *compatible* iff all locations occurring in M are from $\{l_1, \dots, l_n\}$. We say that a term is *canonical* if it is either $()$, an integer constant, a location, a variable or a λ -abstraction. The big-step reduction rules are given as judgements of the shape $s, M \Downarrow s', V$, where (s, M) , (s', V) are compatible, $\text{dom } s \subseteq \text{dom } s'$ and V is canonical. We present them in Figure 2, where we let l range over locations. Most rules take the form

$$\frac{M_1 \Downarrow V_1 \quad M_2 \Downarrow V_2 \quad \dots \quad M_n \Downarrow V_n}{M \Downarrow V}$$

$$\begin{array}{c}
\frac{V \text{ is canonical}}{s, V \Downarrow s, V} \quad \frac{M \Downarrow 0 \quad N_0 \Downarrow V}{\text{if } M \text{ then } N_1 \text{ else } N_0 \Downarrow V} \quad \frac{i \neq 0 \quad M \Downarrow i \quad N_1 \Downarrow V}{\text{if } M \text{ then } N_1 \text{ else } N_0 \Downarrow V} \\
\frac{M_1 \Downarrow i_1 \quad M_2 \Downarrow i_2}{M_1 \oplus M_2 \Downarrow i_1 \oplus i_2} \quad \frac{M \Downarrow \lambda x.M' \quad N \Downarrow V' \quad M'[V'/x] \Downarrow V}{MN \Downarrow V} \\
\frac{s, M \Downarrow s', i \quad l \notin \text{dom } s'}{s, \text{ref } M \Downarrow s'(l \mapsto i), l} \quad \frac{s, M \Downarrow s', l \quad s'(l) = i}{s, !M \Downarrow s', i} \quad \frac{s, M \Downarrow s', l \quad s', N \Downarrow s'', i}{s, M := N \Downarrow s''(l \mapsto i), ()}
\end{array}$$

Fig. 2. Big-step operational semantics of Reduced ML

which is meant to abbreviate

$$\frac{s_1, M_1 \Downarrow s_2, V_1 \quad s_2, M_2 \Downarrow s_3, V_2 \quad \cdots \quad s_n, M_n \Downarrow s_{n+1}, V_n}{s_1, M_1 \Downarrow s_{n+1}, V}$$

In particular, this means that the ordering of the hypotheses is significant. We shall write $\Gamma \vdash M : \theta$ iff $\emptyset, \Gamma \vdash M : \theta$ can be derived using the rules of Figure 1. Similarly, $\vdash M : \theta$ is shorthand for $\emptyset, \emptyset \vdash M : \theta$. Given $\vdash M : \text{unit}$ we write $M \Downarrow$ iff $\emptyset, M \Downarrow s, ()$ for some store s .

Definition 1. We say that the term-in-context $\Gamma \vdash M_1 : \theta$ **approximates** $\Gamma \vdash M_2 : \theta$ (written $\Gamma \vdash M_1 \sqsubseteq M_2$) iff $C[M_1] \Downarrow$ implies $C[M_2] \Downarrow$ for any context $C[-]$ such that $\vdash C[M_1], C[M_2] : \text{unit}$. Two terms-in-context are **equivalent** if one approximates the other (written $\Gamma \vdash M_1 \cong M_2$).

The only difference between our definition of Reduced ML and Stark's is the presence of Ω , the divergent constant without a reduction rule. Thanks to it, we can define \sqsubseteq and reason about \cong in a more concise way. At the same time, program equivalence of Ω -free Reduced ML terms remains unaffected, because $C[M] \Downarrow$, where M is Ω -free, is equivalent to $\emptyset, C'[M] \Downarrow s', 0$ where $C'[-] \equiv$ let $x = \text{ref } 0$ in $C[x := 1/\Omega][M]; !x$ and s' is a state.

3 Nominal game semantics

We begin this section with a brief review of the fundamentals of nominal game semantics [2, 11, 24]. Let us fix a countably infinite set \mathbb{A} , the set of *atoms*, the elements of which we denote by a, b, c, n and variants. In nominal game semantics two participants play a game by exchanging moves that might involve atoms. However, when employing such moves, we are not interested in what exactly the names are, though we would like to know how they relate to names that have already been in play. Hence, the objects of study are rather the induced equivalence classes with respect to name-invariance. Since we want all game-semantic notions and constructions to be compatible with name-invariance, their obvious adaptations would repeatedly have to include conditions that enforce

closure under name-renamings. Fortunately, this overhead can be dealt with robustly using the language of nominal set theory [7].

Definition 2. *Let us write $\text{PERM}(\mathbb{A})$ for the group of finite permutations of \mathbb{A} . A **nominal set** X is a set $|X|$ (usually written X) equipped with a group action of $\text{PERM}(\mathbb{A})$ ⁴. Moreover, each $x \in X$ must have finite support, that is, there exists a finite set $S \subseteq \mathbb{A}$ such that, for all permutations π , $(\forall a \in S. \pi(a) = a) \implies \pi \cdot x = x$.*

Finite support is closed under intersection, and hence each element x of a nominal set has a least support $\nu(x)$, which we call **the support of x** . Intuitively, $\nu(x)$ is the set of names “involved” in x . Accordingly, we say that a is *fresh* for x if $a \notin \nu(x)$. Clearly, \mathbb{A} is a nominal set by taking $\pi \cdot a = \pi(a)$, for each π and a . More interestingly, so is the set \mathbb{A}^* of finite lists of atoms with permutations acting elementwise. If X and Y are nominal sets then so is their cartesian product $X \times Y$, with permutations acting componentwise, and their disjoint union $X \uplus Y$. Moreover, $X' \subseteq X$ is a *nominal subset* of X if X' is closed under permutation actions, these acting as on X . Then we can define $R \subseteq X \times Y$ to be a *nominal relation* iff R a nominal subset of $X \times Y$. A *nominal function* is a function which is also a nominal relation.

In game semantics a particular strengthening of the notion of support, called *strong support*, has turned out necessary to guarantee correct behaviour under strategy composition (see [24] for motivation and a detailed explanation of its significance). Here we consider an even stronger notion of support, one in which the support of each element can be linearly ordered in a canonical, nominal manner. A nominal set X is called a **linear nominal set** if, for each element x of X , there exists a linear order $<_x$ on $\nu(x)$ such that, for all $a, b \in \mathbb{A}$, $a <_x b$ implies $\pi(a) <_{\pi \cdot x} \pi(b)$ for any permutation π ⁵. It is easy to check that all elements in a linear nominal set have strong support. For example, the nominal set \mathbb{A}^* is linear, whereas $\mathcal{P}_{fin}(\mathbb{A})$ is not. A nominal subset of a linear nominal set is itself linear. Moreover, by straightforward manipulations of the orderings available for linear X, Y we can render the nominal sets $X \times Y$ and $X \uplus Y$ linear.

Finally, in nominal sets we can *define* atom-abstractions. The form of abstraction we will be using is that of complete support abstraction, that is, for a nominal set X and $x \in X$, we define $[x]$ to be $\{y \in X \mid \exists \pi. y = \pi \cdot x\}$.

3.1 Nominal arenas

Here we present nominal arenas (and prearenas), which are essentially the call-by-value arenas of Honda and Yoshida [8] cast inside the theory of nominal sets.

Definition 3. *An arena $A = (M_A, I_A, \vdash_A, \lambda_A)$ is given by:*

⁴ A group action of $\text{PERM}(\mathbb{A})$ on X is a function $\cdot : \text{PERM}(\mathbb{A}) \times X \rightarrow X$ such that, for all $x \in X$ and $\pi, \pi' \in \text{PERM}(\mathbb{A})$, $\pi \cdot (\pi' \cdot x) = (\pi \circ \pi') \cdot x$ and $\text{id} \cdot x = x$, where id is the identity permutation.

⁵ Equivalently, the relation $\{(a, b, x) \mid a, b \in \mathbb{A}, x \in X, a <_x b\}$ is nominal.

- a linear nominal set M_A of moves,
- a nominal subset $I_A \subseteq M_A$ of initial moves,
- a nominal justification relation $\vdash_A \subseteq M_A \times (M_A \setminus I_A)$,
- a nominal labelling function $\lambda_A : M_A \rightarrow \{O, P\} \times \{Q, A\}$;

satisfying, for each $m, m' \in M_A$, the conditions:

- $m \in I_A \implies \lambda_A(m) = (P, A)$,
- $m \vdash_A m' \wedge \lambda_A^{QA}(m) = A \implies \lambda_A^{QA}(m') = Q$,
- $m \vdash_A m' \implies \lambda_A^{OP}(m) \neq \lambda_A^{OP}(m')$.

The role of λ_A is to label moves as *Opponent* or *Proponent* moves and as *Questions* or *Answers*. The simplest arena is $0 = (\emptyset, \emptyset, \emptyset, \emptyset)$. Other “flat” arenas are 1 , \mathbb{Z} and \mathbb{A} , defined by $M_1 = I_1 = \{*\}$, $M_{\mathbb{Z}} = I_{\mathbb{Z}} = \mathbb{Z}$, $M_{\mathbb{A}} = I_{\mathbb{A}} = \mathbb{A}$.

We take advantage of the following constructions on arenas. By \bar{I}_A we denote $M_A \setminus I_A$, by $\bar{\lambda}_A$ the *OP*-complement of λ_A ; i_A and i_B range over initial moves in the respective arenas.

$$\begin{aligned}
M_{A \otimes B} &= (I_A \times I_B) \uplus \bar{I}_A \uplus \bar{I}_B \\
I_{A \otimes B} &= I_A \times I_B \\
\lambda_{A \otimes B} &= [((i_A, i_B), PA), \lambda_A \upharpoonright \bar{I}_A, \lambda_B \upharpoonright \bar{I}_B] \\
\vdash_{A \otimes B} &= \{(i_A, i_B), m \mid i_A \vdash_A m \text{ or } i_B \vdash_B m\} \cup (\vdash_A \upharpoonright \bar{I}_A^2) \cup (\vdash_B \upharpoonright \bar{I}_B^2)
\end{aligned}$$

$$\begin{aligned}
M_{A \Rightarrow B} &= \{*\} \uplus I_A \uplus \bar{I}_A \uplus M_B \\
I_{A \Rightarrow B} &= \{*\} \\
\lambda_{A \Rightarrow B} &= [(*, PA), (i_A, OQ), \bar{\lambda}_A \upharpoonright \bar{I}_A, \lambda_B] \\
\vdash_{A \Rightarrow B} &= \{(*, i_A), (i_A, i_B)\} \cup \vdash_A \cup \vdash_B
\end{aligned}$$

The types of Reduced ML will be interpreted by arenas in the following way: $\llbracket \text{unit} \rrbracket = 1$, $\llbracket \text{int} \rrbracket = \mathbb{Z}$, $\llbracket \text{int ref} \rrbracket = \mathbb{A}$ and $\llbracket \theta_1 \rightarrow \theta_2 \rrbracket = \llbracket \theta_1 \rrbracket \Rightarrow \llbracket \theta_2 \rrbracket$. Although types are interpreted by arenas, the actual games will be played in *prearenas*, which are defined in the same way as arenas with the exception that initial moves are *O*-questions. For given arenas A, B we can construct a prearena $A \rightarrow B$ by

$$\begin{aligned}
M_{A \rightarrow B} &= M_A \uplus M_B & \lambda_{A \rightarrow B} &= [(i_A, OQ) \cup (\bar{\lambda}_A \upharpoonright \bar{I}_A), \lambda_B] \\
I_{A \rightarrow B} &= I_A & \vdash_{A \rightarrow B} &= \{(i_A, i_B)\} \cup \vdash_A \cup \vdash_B .
\end{aligned}$$

Typing judgements $\Gamma \vdash \theta$, where $\Gamma = \{x_1 : \theta_1, \dots, x_n : \theta_n\}$, will eventually be interpreted by strategies for the prearena $\llbracket \theta_1 \rrbracket \otimes \dots \otimes \llbracket \theta_n \rrbracket \rightarrow \llbracket \theta \rrbracket$ (if $n = 0$ we take the left-hand side to be 1), which we shall denote by $\llbracket \Gamma \vdash \theta \rrbracket$.

3.2 Plays

Analogously to the definition of store in Section 2, in this section a store will be a partial function $S : \mathbb{A} \rightarrow \mathbb{Z}$ such that $\text{dom } S$ is finite.

A *basic justified sequence* in a prearena A is a finite sequence s of moves of A satisfying the following conditions: the first move must be initial, but all other

moves m must be equipped with a pointer to an earlier occurrence of another move m' such that $m' \vdash_A m$ (we then say that m' justifies m ; if m is an answer, we might also say that m answers m'). A **justified sequence** in A is a basic justified sequence s in which each move is, in addition, decorated with a store to yield a *move-with-store*, typically denoted by m^S . Given a justified sequence s , we write \underline{s} for the underlying basic justified sequence. It should be clear that, similarly to the set of finite sequences of moves, the set of justified sequences can be viewed as a (not necessarily linear) nominal set with permutations preserving the pointer structure, but acting on moves as in A and on stores by permuting the domain.

Below we define the notions of O-view $\lfloor s \rfloor$ and P-view $\lceil s \rceil$ of a justified sequence, using o and p to range over O-moves and P-moves respectively. We write $s' \sqsubseteq s$ if s' is a prefix of s and use $\sqsubseteq^{\text{even}}$ if s' is of even length.

$$\begin{array}{ll} \lfloor \epsilon \rfloor = \epsilon & \lceil \epsilon \rceil = \epsilon \\ \lfloor s o^S \rfloor = \lfloor s \rfloor o^S & \lceil s p^{S\lceil} \rceil = \lceil s \rceil p^S \\ \lfloor s o^S \bar{t} p^{S'} \rfloor = \lfloor s \rfloor o^S p^{S'} & \lfloor s p^S \bar{t} o^{S'} \rfloor = \lfloor s \rfloor p^S o^{S'} \end{array}$$

A name in s is said to be **introduced by player** X ($X \in \{O, P\}$) iff its first occurrence in \underline{s} is in (the support of) an X -move. Names introduced by X in s will be referred to as **X -names** in s and denoted with $X(s)$. We define the set $\text{Av}_X(s)$ of **X -available names after** s by:

$$\text{Av}_O(s) = O(s) \cup \nu(\lfloor \underline{s} \rfloor) \quad \text{Av}_P(s) = P(s) \cup \nu(\lceil \underline{s} \rceil).$$

Definition 4. A justified sequence is **legal** iff it satisfies the following conditions.

Alternation No two adjacent moves belong to the same player.

Bracketing The justifier of each answer is the most recent unanswered question.

Visibility For any $tm^S \sqsubseteq s$, the justifier of m is in $\lfloor tm^S \rfloor$ if m is an O-move and in $\lceil tm^{S\lceil} \rceil$ otherwise.

Frugality For any $tm^S \sqsubseteq s$, $\text{dom } S \subseteq \nu(\underline{tm^S})$.

The set of legal justified sequences will be denoted by L_A .

Note that legal sequences contain those of [11]. Because of frugality, the support of a legal sequence is that of its underlying basic sequence, and therefore L_A is a linear nominal set. Our model will be based on still more restrictive plays.

Definition 5. A legal sequence s is a **play** iff it satisfies the following two conditions.

P-availability For each $s'p^S \sqsubseteq^{\text{even}} s$ and any $a \in \mathbb{A}$, if $a \in \nu(p) \cap \nu(s')$ then $a \in \text{Av}_P(s')$.

P-storage For any $s'm^S \sqsubseteq s$, $\text{dom } S = \text{Av}_P(s'm^S)$.

The set of plays over prearena A will be denoted by P_A .

Note that the two conditions are biased towards P. Equivalently, P-availability can be restated as: for any $s'p^S \sqsubseteq^{\text{even}} s$ and any $a \in \nu(p)$, if $a \in O(s')$ then $a \in \nu(\lceil s' \rceil)$. It is worth observing that, given $s \in P_A$, we have $\text{Av}_P(s) = \nu(\lceil s \rceil)$.

3.3 Strategies

Definition 6. A strategy σ on a prearena A , written $\sigma : A$, is a set of equivalence classes $[s]$ of even-length plays of A satisfying

Even-prefix closure If $[so^S p^{S'}] \in \sigma$ then $[s] \in \sigma$,

Determinacy If $[sp_1^{S_1}], [s'p_2^{S_2}] \in \sigma$ and $[s] = [s']$ then $[sp_1^{S_1}] = [s'p_2^{S_2}]$.

Next we show how strategies can be composed. First, following [11], let us define an endofunction γ on justified sequences that restricts a given justified sequence to a frugal one by removing from the stores the atoms violating it. Now, let γ' be an analogous *partial* function enforcing P-storage, i.e. γ' will remove O-names violating P-storage (it is undefined when the domain of any of the stores involved contains an insufficient supply of atoms, i.e. some of the P-available names required are missing).

Now we turn to defining a suitable notion of interaction between plays. Given arenas A, B, C , let u be a sequence $m_1^{S_1} \cdots m_k^{S_k}$ of moves from $M_A + M_B + M_C$ with store, equipped with pointers such that no C -move has a pointer to an A -move and vice versa. We define $u \upharpoonright A, B$ to be u in which all C -moves are suppressed along with associated pointers. $u \upharpoonright B, C$ is defined analogously. $u \upharpoonright A, C$ is defined similarly with the caveat that, if there was a pointer from a C -move to a B -move which in turn had a pointer to an A -move, we add a pointer from the C -move to the A -move. By $u_{\leq m_i}$ we mean the initial segment of u ending in $m_i^{S_i}$.

Definition 7. u is an interaction sequence of A, B, C iff $\gamma'(u \upharpoonright A, B) \in P_{A \rightarrow B}$, $\gamma'(u \upharpoonright B, C) \in P_{B \rightarrow C}$ and the following conditions hold:

- $P(u \upharpoonright A, B) \cap P(u \upharpoonright B, C) = \emptyset$;
- $O(u \upharpoonright A, C) \cap (P(u \upharpoonright A, B) \cup P(u \upharpoonright B, C)) = \emptyset$;
- for each $u' \sqsubseteq u$ ending in a move-with-store m^S ,
 $\text{dom } S = (O(u' \upharpoonright A, C) \cap \nu(\ulcorner u' \upharpoonright A, C \urcorner)) \cup P(u' \upharpoonright A, B) \cup P(u' \upharpoonright B, C)$;
- for each $u' \sqsubseteq u$ ending in $m^S m'^{S'}$, if m' is:
 - a P-move in $A \rightarrow B$ then $S'(a) = S(a)$ for all $a \in \text{dom } S' \setminus \text{Av}_P(u' \upharpoonright A, B)$;
 - a P-move in $B \rightarrow C$ then $S'(a) = S(a)$ for all $a \in \text{dom } S' \setminus \text{Av}_P(u' \upharpoonright B, C)$;
 - an O-move in $A \rightarrow C$ then $S'(a) = S(a)$ for all $a \in \text{dom } S' \setminus \text{Av}_P(u' \upharpoonright A, C)$.

The set of all interaction sequences of A, B, C will be denoted by $\text{Int}(A, B, C)$.

The first two conditions ensure that name-privacy is not broken under composition; the third one imposes an extended notion of P-availability for sequences; and the fourth set of conditions ensures that participants do not change parts of the store inaccessible to them. It can be shown that, if $u \in \text{Int}(A, B, C)$ then $\gamma(u \upharpoonright A, C) \in P_{A \rightarrow C}$. Two strategies $\sigma : A \rightarrow B$ and $\tau : B \rightarrow C$ can now be composed as follows

$$\sigma;\tau = \{[\gamma(u \upharpoonright A, C)] \mid u \in \text{Int}(A, B, C), [\gamma'(u \upharpoonright A, B)] \in \sigma, [\gamma'(u \upharpoonright B, C)] \in \tau\}.$$

Associativity of composition can be established using similar arguments to those

in [11]⁶. Using the standard definition of the identity strategy one can then obtain a category of arenas and games. Next we shall define its lluf subcategory that will be used to prove the full abstraction result.

Given a non-empty justified sequence s , let us write s^- for s without its last element. The following definition aims to capture plays that differ by renamings of names that O introduces in the P-view.

- Definition 8.** – Given a prearena A , $s \in P_A$, $a \in \mathbf{O}(s)$ and an O-move o in s , we say that a is **P-new at o in s** iff $a \in \nu(o)$ and $a \notin \nu(\ulcorner s_{\leq o} \urcorner^-)$.
- Given A, s, a, o as above and $b \in \mathbb{A}$, we say that a is **renameable for b at o in s** provided $b \notin \mathbf{P}(s)$ and, for any $s' \sqsubseteq s$, if o occurs in $\ulcorner s' \urcorner$ then $b \notin \nu(\ulcorner s' \urcorner)$.
 - Under the assumptions above, we define the **renaming** $(a\ b)_o \cdot s$ of s by induction on the subsequences of s^\ulcorner :

$$(a\ b)_o \cdot \epsilon = \epsilon \quad (a\ b)_o \cdot (tm^S) = \begin{cases} ((a\ b)_o \cdot t) ((a\ b) \cdot m^S) & o \in \ulcorner tm^S \urcorner \\ ((a\ b)_o \cdot t) m^S & o \notin \ulcorner tm^S \urcorner \end{cases}$$

where $(a\ b)$ is the permutation swapping a with b . We write $s \overset{\ulcorner}{\sim} s'$ iff s can be obtained from s' through a sequence of renamings.

Observe that, if a is P-new at m in s , a need not be fresh for $s_{<m}$ (the converse holds, though, as long as $a \in \nu(m)$). A play s in which every a that is P-new at m in s is also fresh at $s_{<m}$ will be called **strict**.

Example 9. Let $A = \mathbb{A} \rightarrow (\mathbb{A} \Rightarrow 1)$,

$$s_1 = n_1^{(n_1,0)} \widehat{*}^{(n_1,1)} \widehat{n_2^{(n_1,2),(n_2,3)} *^{(n_1,4),(n_2,5)}} n_3^{(n_1,6),(n_3,7)} \widehat{*}^{(n_1,8),(n_3,9)}$$

and

$$s_2 = n_1^{(n_1,0)} \widehat{*}^{(n_1,1)} \widehat{n_3^{(n_1,2),(n_3,3)} *^{(n_1,4),(n_3,5)}} n_3^{(n_1,6),(n_3,7)} \widehat{*}^{(n_1,8),(n_3,9)}.$$

Then n_2 is P-new at the third move (also n_2) in s_1 , n_2 is renameable for n_3 at that move and $(n_2\ n_3)_{n_2} \cdot s_1 = s_2$. Note also that $(n_3\ n_2)_{n_3} \cdot s_2 = s_1$, where the subscript n_3 stands for the third move of s_2 , and that s_1 is strict, whereas s_2 is not.

In general it can be shown that renamings are reversible, so $\overset{\ulcorner}{\sim}$ is an equivalence relation. Observe that, for any play s , there exists a *strict* play s' (determined uniquely up to atom-abstraction) such that $s \overset{\ulcorner}{\sim} s'$. We write \tilde{s} for $[s']$.

Definition 10. A strategy $\sigma : A$ is **blind** iff $[s] \in \sigma$ and $s \overset{\ulcorner}{\sim} s'$ imply $[s'] \in \sigma$.

Since the identity strategy is blind and blind strategies compose we obtain a category \mathcal{G} of arenas and blind strategies. Observe that blind strategies are uniquely determined by the underlying strict positions (via renamings).

⁶ In fact, strategies in our framework can be regarded as compact representations of a class of strategies from [11], namely those independent of P-unavailable names.

⁷ We write $o \in s$ to mean that the distinguished occurrence of o is present in s .

factors out non-innocence (also using new_0). Finally, we prove a direct definability result for finitarily innocent strongly deterministic strategies. We discuss the innocent factorization in more detail below, as it involves the key novelties of our framework.

Lemma 13. *Let $\sigma : 1 \rightarrow A$ be a finitary strongly deterministic blind strategy. There exists a finitarily innocent strongly deterministic strategy $\dot{\sigma} : \mathbb{A} \rightarrow A$ such that $new_0; \dot{\sigma} = \sigma$.*

The standard way [5] of proving such results is to store the history of play using the additional \mathbb{A} component. This is impossible in our case, because atoms cannot be stored. However, given a play s , we can try to store a numerical representation of $[s]$ instead. Recall that the set of moves of an arena is a linear nominal set, i.e. there is a *canonical* ordering of atoms in any move. Hence, in any legal sequence, atoms can also be ordered in a *canonical* way according to their order of appearance and, if they were introduced in the same move, using the canonical ordering associated with that move. Consequently, we can represent $[s]$ as an integer by representing atoms in s with integers that correspond to their position in the associated canonical order and by encoding such a sequence as an integer. Let us write $\#(s)$ for such an encoding. In particular we have $\#(s_1) = \#(s_2)$ iff $[s_1] = [s_2]$.

Unfortunately, this is not yet sufficient for a successful factorization through an innocent strategy because, given $\ulcorner so^S \urcorner$ and $\#(s)$, it will in general be impossible to extract so^S (or $[so^S]$) due to the fact that o might contain O-names occurring in s , but not in $\ulcorner so^S \urcorner$. As a result, given $\#(s)$ and $\ulcorner so^S \urcorner$, we may then be unable to relate some names in o with those of s , which will prevent us from reconstructing $[so^S]$. Note, however, that given $\#(s)$ and $\ulcorner so^S \urcorner$ we can still determine $\widetilde{so^S}$ in absence of P-names. Furthermore, since σ is blind and strongly deterministic (in particular plays satisfy P-availability), we can uniquely identify $p^{S'}$ such that $so^S p^{S'} \in \sigma$ (though not necessarily the whole of $so^S p^{S'}$), by referring¹⁰ to $\widetilde{so^S}$ and σ . Analogously, we can also deduce $\widetilde{so^S p^{S'}}$. Thus, the familiar factorization technique can be employed provided that, instead of $\#(s)$, the argument will rely on $\#(\widetilde{s})$, where $\#(\widetilde{s})$ stands for $\#(s')$ and s' is a strict play such that $s' \overset{\tau}{\sim} s$ (by previous remarks the code is independent of the exact choice of s').

Thanks to the definability result for finitary blind strategies, we can define a fully abstract model of Reduced ML in the usual way by quotienting \mathcal{G} by the induced *intrinsic preorder* defined below.

Definition 14. *Suppose $\sigma_1, \sigma_2 : 1 \rightarrow A$. We define $\sigma_1 \leq \sigma_2$ to hold iff, for any $\rho : A \rightarrow 1$, $\sigma_1; \rho \neq \{\{\epsilon\}\}$ implies $\sigma_2; \rho \neq \{\{\epsilon\}\}$.*

It is common to refer to the above preorder as *testing σ_i with ρ* , where $\sigma_i; \rho \neq \{\{\epsilon\}\}$ is regarded as a successful outcome.

Theorem 15. *Given Reduced ML terms $\vdash M_1 : \theta$, $\vdash M_2 : \theta$, we have $\vdash M_1 \sqsubseteq M_2$ iff $\llbracket \vdash M_1 \rrbracket \leq \llbracket \vdash M_2 \rrbracket$.*

¹⁰ Recall that blind strategies are generated by their strict plays.

5 Program equivalence explicitly

Let σ_1, σ_2, ρ be as in the definition of \leq . Note that during composition of σ_i with ρ there is a full symmetry between O-names and P-names, i.e. names which are O-names in σ_i are viewed as P-names in ρ , and vice versa. This can be contrasted with the general case of composition, where both strategies may regard a name as an O-name during composition, though not a P-name. This symmetry of roles means that, because plays of ρ satisfy P-availability, a successful outcome can only be reached by interaction with a play of σ_i that satisfies the dual condition of **O-availability**: for each $s' o^S \sqsubseteq^{\text{odd}} s$ and any $a \in \mathbb{A}$, if $a \in \nu(o) \cap \nu(s')$ then $a \in \text{Av}_O(s')$.

Similarly, whenever the play s engages with ρ successfully, the **O-passivity** condition holds: for each $s' p^S o^{S'} \sqsubseteq s$ and any $a \in \mathbb{A}$, if $a \in \text{P}(s' p^S) \setminus \text{Av}_O(s' p^S)$ then $S(a) = S'(a)$. This time this is due to the definition of composition, which stipulates that the part of store irrelevant to one of the strategies must be copied. This means that the plays of σ_i that “matter” must necessarily meet the above condition. Finally, whenever $\sigma_i; \rho \neq \{\epsilon\}$, the play witnessing this is **complete**, i.e. all of its questions are answered.

Definition 16. *A play is **relevant** iff it is complete, satisfies O-availability and O-passivity. We write $\text{rel}(\sigma)$ for the set of relevant plays of σ .*

We can represent *relevant* plays more succinctly by restricting the associated stores to mutually available names (both O- and P-available). The outcome is not a play any more, though it remains a legal justified sequence. We call such sequences **protoplays** and let γ'' be the obvious operation on justified sequences that simply erases the O-unavailable names in stores. Although some information about σ is seemingly lost by applying γ'' to $\text{rel}(\sigma)$, the missing values turn out inessential for testing. By O-passivity, the lost values of O-unavailable names can be uniquely retrieved in O-moves, by copying values from the preceding P-moves. However, more surprisingly, it does not matter what values such names have in P-moves either. This is because the names are then P-unavailable for ρ and, during composition, are dealt with *uniformly* by propagation as long as they remain unavailable.

Finally, we take advantage of the fact that the test ρ is a blind strategy. Recall that blind strategies are uniquely determined by their strict plays, i.e. plays in which O-names fresh in the P-view must be genuinely fresh at the point of introduction. Consequently, if one wants to check if σ_i passes the ρ test, we can take advantage of the fact that any contribution from ρ will originate from a strict play. Let $s'' = \gamma''(s')$ ($s' \in \text{rel}(\sigma)$) be a protoplay generated by σ_i . To test whether s'' represents a renaming of a strict play from ρ , it suffices to “refresh” P-names in s'' and try to match it with that the strict play. The desired refreshing operation (for P-names using O-views) is entirely dual to renamings introduced in Definition 8, though it needs to be defined on protoplays to be correct. For two protoplays s, s' , we write $s \underset{\tau}{\sim} s'$ iff s can be obtained from s' by a series of dual renamings. A protoplay is **dually strict** iff any P-name fresh in the O-view is fresh at the point of introduction. Given $\sigma : A$, let $\hat{\sigma}$ be the

following set of equivalence classes of protoplays $\widehat{\sigma} = \{[s] \mid s \text{ is dually strict, } s' \in \text{rel}(\sigma), s \underset{r}{\sim} \gamma''(s')\}$. We can then show the following result.

Lemma 17. *Given $\sigma_1, \sigma_2 : 1 \rightarrow A$, $\sigma_1 \leq \sigma_2$ iff $\widehat{\sigma}_1 \subseteq \widehat{\sigma}_2$.*

Observe that $\widehat{\sigma}$, like σ , is saturated under renamings (extended to act on protoplays). This makes it possible to simplify the above result along the following lines. We call a protoplay *mutually strict* iff it is both strict and dually strict. Note that by using $\underset{r}{\sim}$ and $\underset{r}{\sim}$ (in any order) we can convert a protoplay to a mutually strict protoplay, unique up to atom-abstraction. Given $\sigma : A$, let $\widehat{\widehat{\sigma}}$ be $\{[s] \mid s \text{ is mutually strict, } s' \in \text{rel}(\sigma), s (\underset{r}{\sim}; \underset{r}{\sim}) \gamma''(s')\}$.

Theorem 18. *Given $\sigma_1, \sigma_2 : 1 \rightarrow A$, $\sigma_1 \leq \sigma_2$ iff $\widehat{\widehat{\sigma}}_1 \subseteq \widehat{\widehat{\sigma}}_2$.*

It follows that terms of Reduced ML are equivalent iff they induce the same mutually strict protoplays.

Thus we have shown that program equivalence and approximation in Reduced ML can be captured explicitly, which is the first result of this kind for Reduced ML. The characterization immediately implies that equivalence is decidable for finitary strategies and that the fully abstract model of Reduced ML is effectively presentable.

Our results identify mutually strict protoplays as an appealing object for future study and, indeed, our fully abstract model can be presented in a more direct way by founding the games on them. It has to be said, though, that composition of such plays is quite intricate, because they cannot be combined by parallel composition with hiding: although this kind of interaction is sufficient to test plays, composition in general lacks the convenient duality between O-names and P-names. Consequently, in order to compose mutually strict protoplays, one has to allow for renamings before synchronization and follow with dual renamings afterwards. We intend to present an account of this procedure in the full version of the paper.

In this submission however we have chosen to present the model gradually: starting from the intuitive framework in which full information about the store is available we successively imposed a series of restrictions. We believe this leads to a more informative presentation and decomposes the difficulties involved in dealing with mutually strict protoplays into smaller arguments. For instance, the correctness proof of compositionality of mutually strict protoplays (for the algorithm sketched above) draws on insights obtained from all of our compositionality proofs (P-availability, P-storage, blindness) as well as the argument behind the explicit characterization.

We hope to bring our results to bear on the on-going research into algorithmic aspects of game models [1] and to contribute new methods of reasoning about program equivalence in Reduced ML. This direction has been pursued using logical relations in [20]. However, there are limits to what can be achieved, as program equivalence of finitary Reduced ML (finite types) is already undecidable at second order, due to a similar result for Reduced ML with `mkvar` in [17].

References

1. Abramsky, S., Ghica, D.R., Murawski, A.S., Ong, C.-H.L.: Applying game semantics to compositional software modelling and verification. Proc. of TACAS. Vol. 2988 of LNCS (2004) 421–435
2. Abramsky, S., Ghica, D.R., Murawski, A.S., Ong, C.-H.L., Stark, I.D.B.: Nominal games and full abstraction for the nu-calculus. Proc. of LICS (2004) 150–159
3. Abramsky, S., Honda, K., McCusker, G.: Fully abstract game semantics for general references. Proc. of LICS (1998) 334–344
4. Abramsky, S., Jagadeesan, R., Malacaria, P.: Full abstraction for PCF. Information and Computation **163** (2000) 409–470
5. Abramsky, S., McCusker, G.: Linearity, sharing and state: a fully abstract game semantics for Idealized Algol with active expressions. In: Algol-like languages, Birkhäuser (1997) 297–329
6. Abramsky, S., McCusker, G.: Call-by-value games. Proc. of CSL. Vol. 1414 of LNCS (1997) 1–17
7. Gabbay, M.J., Pitts, A.M.: A new approach to abstract syntax with variable binding. Formal Aspects of Computing **13** (2002) 341–363
8. Honda, K., Yoshida, N.: Game-theoretic analysis of call-by-value computation. Proc. of ICALP. Vol. 1256 of LNCS (1997) 225–236
9. Hyland, J.M.E., Ong, C.-H.L.: On Full Abstraction for PCF. Information and Computation **163(2)** (2000) 285–408
10. Laird, J.: A game semantics of local names and good variables. Proc. of FOSSACS. Vol. 2987 of LNCS (2004) 289–303
11. Laird, J.: A game semantics of names and pointers. Annals of Pure and Applied Logic **151** (2008) 151–169
12. McCusker, G.: Games for recursive types. BCS Distinguished Dissertation. Cambridge University Press (1998)
13. McCusker, G.: Games and full abstraction for FPC. Information and Computation **160 (1-2)** (2000) 1–61
14. McCusker, G.: On the semantics of Idealized Algol without the bad-variable constructor. Proc. of MFPS. Vol. 83 of ENTCS (2003)
15. Milner, R., Tofte, M., Harper, R.: The Definition of Standard ML. The MIT Press, Cambridge, Massachusetts (1990)
16. Moggi, E.: Notions of computation and monads. Information and Computation **93** (1991) 55–92
17. Murawski, A.S.: Functions with local state: regularity and undecidability. Theoretical Computer Science **338(1/3)** (2005) 315–349
18. Murawski, A.S.: Bad variables under control. Proc. of CSL. Vol. 4646 of LNCS (2007) 558–572
19. Nickau, H.: Hereditarily sequential functionals. Proc. of LFCS. Vol. 813 of LNCS (1994) 253–264
20. Pitts, A.M., Stark, I.D.B.: Operational reasoning for functions with local state. In: Higher-Order Operational Techniques in Semantics. CUP (1998) 227–273
21. Reynolds, J.C.: The essence of Algol. In de Bakker, J.W., van Vliet, J., eds.: Algorithmic Languages. North Holland (1978) 345–372
22. Stark, I.D.B.: Names and Higher-Order Functions. PhD thesis, University of Cambridge (1995)
23. Tzevelekos, N.: Full abstraction for nominal general references. Proc. of LICS (2007) 399–410
24. Tzevelekos, N.: Nominal game semantics. D.Phil. thesis, University of Oxford (2008)