

Abstract machines for game semantics, revisited

Olle Fredriksson
University of Birmingham, UK

Dan R. Ghica
University of Birmingham, UK

Abstract—We define new abstract machines for game semantics which correspond to networks of conventional computers, and can be used as an intermediate representation for compilation targeting distributed systems. This is achieved in two steps. First we introduce the HRAM, a *Heap and Register Abstract Machine*, an abstraction of a conventional computer, which can be structured into HRAM nets, an abstract point-to-point network model. HRAMs are multi-threaded and subsume communication by tokens (cf. IAM) or jumps. Game Abstract Machines (GAM), are HRAMs with additional structure at the interface level, but no special operational capabilities. We show that GAMs cannot be naively composed, but composition must be mediated using appropriate HRAM combinators. HRAMs are flexible enough to allow the representation of game models for languages with state (non-innocent games) or concurrency (non-alternating games). We illustrate the potential of this technique by implementing a toy distributed compiler for ICA, a higher-order programming language with shared state concurrency, thus significantly extending our previous distributed PCF compiler. We show that compilation is sound and memory-safe, i.e. no (distributed or local) garbage collection is necessary.

I. INTRODUCTION

One of the most profound discoveries in theoretical computer science is the fact that logical and computational phenomena can be subsumed by relatively simple communication protocols. This understanding came independently from Girard’s work on the Geometry of Interaction (GOI) [16] and Milner’s work on process calculi [21], and had a profound influence on the subsequent development of game semantics (see [12] for a historical survey). Of the three, game semantics proved to be particularly effective at producing precise mathematical models for a large variety of programming languages, solving a long-standing open problem concerning higher-order sequential computation [1, 18].

One of the most appealing features of game semantics is that it has a dual denotational and operational character. By *denotational* we mean that it is compositionally defined on the syntax and by *operational* we mean that it can be effectively presented and can form a basis for compilation [13]. This feature was apparent from the earliest presentations of game semantics [17] and is not very surprising, although the operational aspects are less perspicuous than in interpretations based on process calculi or GOI, which quickly found applications in compiler [20] or interpreter [2] development and optimisation.

An important development, which provided essential inspiration for this work, was the introduction of the *Pointer Abstract Machine* (PAM) and the *Interaction Abstract Machine* (IAM), which sought to fully restore the operational intuitions

of game semantics [5] by relating them to two kinds of abstract machines, one based on term rewriting (PAM) and one based on networks of automata (IAM) profoundly inspired by GOI. A further optimisation of IAM, the *Jumping Abstract Machine* (JAM) was introduced subsequently to avoid the overheads of the IAM [6].

Contribution: In this paper we are developing the line of work on PAM/IAM/JAM in order to define new abstract machines which correspond more closely to *networks of conventional computers* and can be used as an intermediate representation for compilation targeting distributed systems. This is achieved in two steps. First we introduce the HRAM, a *Heap and Register Abstract Machine*, an abstraction of a conventional computer, which can be structured into HRAM nets, an abstract point-to-point network model. HRAMs are multi-threaded and subsume communication by tokens (cf. IAM) or jumps. GAMs, *Game Abstract Machines*, are HRAMs with additional structure at the interface level, but no special operational capabilities. We show that GAMs cannot be naively composed, but composition must be mediated using appropriate HRAM combinators. Starting from a formulation of game semantics in the nominal model [9] has two benefits. First, pointer manipulation requires no encoding or decoding, as in integer-based representations, but exploits the HRAM ability to create locally fresh *names*. Second, token size is constant as only names are passed around; the computational history of a token is stored by the HRAM rather than passing it around (cf. IAM). HRAMs are also flexible enough to allow the representation of game models for languages with state (*non-innocent* games) or concurrency (*non-alternating* games). We illustrate the potential of this technique by implementing a compiler targeting distributed systems for ICA, a higher-order programming language with shared state concurrency [14], thus significantly extending our previous distributed PCF compiler [8]. We show that compilation is sound and memory-safe, i.e. no (distributed or local) garbage collection is necessary.¹

Other related and relevant work: The operational intuitions of GOI were originally confined to the sequential setting, but more recent work on Ludics showed how they can be applied to concurrency [7] through an abstract treatment not immediately applicable to our needs. Whereas our work takes the IAM/JAM as the starting point, developing abstract machines akin to the PAM revealed interesting syntactic and operational connections between game semantics and Böhm trees [4]. The connection between game semantics, syntactic

A full version of this paper is in arXiv:1304.4159 [cs.LO].

¹Available from <http://veritygos.org/gams>.

recursion schemes and automata also had several interesting applications to verifying higher-order computation (see e.g. [22]). Finally, the connection between game semantics and operational semantics can be made more directly by eliding all the semantic structure in the game and reducing them to a very simple communication mechanism between a program and its environment, which is useful in understanding hostile opponents and verifying security properties [15].

II. SIMPLE NETS

In this section we introduce a class of basic abstract machines for manipulating heap structures, which also have primitives for communications and control. They represent a natural intermediate stage for compilation to machine language, and will be used as such in Sec. IV. The machines can naturally be organised into communication networks which give an abstract representation of distributed systems. We find it formally convenient to work in a nominal model in order to avoid the difficulties caused by concrete encoding of game structures, especially *justification pointers*, as integers. We assume a certain familiarity from the reader with basic nominal concepts. The interested reader is referred to the literature ([10] is a starting point).

A. Heap and register abstract machines (HRAM)

We fix a set of *port names* (\mathbb{A}) and a set of *pointer names* (\mathbb{P}) as disjoint sets of atoms. Let $L \triangleq \{\mathbf{O}, \mathbf{P}\}$ be the set of polarities of a port. To maintain an analogy with game semantics from the beginning, port names correspond to game-semantic *moves* and input/output polarities correspond to opponent/proponent. A *port structure* is a tuple $(l, a) \in \mathit{Port} = L \times \mathbb{A}$. An *interface* $A \in \mathcal{P}_{\text{fin}}(\mathit{Port})$ is a set of port structures such that all port names are unique, i.e. $\forall p = (l, a), p' = (l', a') \in A$, if $a = a'$ then $p = p'$. Let the support of an interface be $\text{sup}(A) \triangleq \{a \mid (l, a) \in A\}$, its set of port names.

The *tensor* of two interfaces is defined as $A \otimes B \triangleq A \cup B$, where $\text{sup}(A) \cap \text{sup}(B) = \emptyset$. The dual of an interface is defined as $A^* \triangleq \{p^* \mid p \in A\}$ where $(l, a)^* \triangleq (l^*, a)$, $\mathbf{O}^* \triangleq \mathbf{P}$ and $\mathbf{P}^* \triangleq \mathbf{O}$. An arrow interface is defined in terms of tensor and dual, $A \Rightarrow B \triangleq A^* \otimes B$.

We introduce notation for opponent ports of an interface $A^{(\mathbf{O})} \triangleq \{(\mathbf{O}, a) \in A\}$. The player ports of an interface $A^{(\mathbf{P})}$ is defined analogously. The set of all interfaces is denoted by \mathcal{I} . We say that two interfaces *have the same shape* if they are equivariant, i.e. there is a permutation $\pi : \mathbb{A} \rightarrow \mathbb{A}$ such that $\{\pi \cdot p \mid p \in A_1\} = A_2$, and we write $\pi \vdash A_1 =_{\mathbb{A}} A_2$, where $\pi \cdot (l, a) \triangleq (l, \pi(a))$ is the permutation action of π . We may only write $A_1 =_{\mathbb{A}} A_2$ if π is obvious or unimportant.

Let the set of data \mathcal{D} be $\emptyset \in \mathbb{1}$, pointer names $a \in \mathbb{P}$ or integers $n \in \mathbb{Z}$. Let the set of instructions *Instr* be as below, where $i, j, k \in \mathbb{N} + \mathbb{1}$ (which permits ignoring results and allocating “null” data).

- $i \leftarrow \text{new } j, k$ allocates a new pointer in the heap and populates it with the values stored in registers j and k , storing the pointer in register i .
- $i, j \leftarrow \text{get } k$ reads the tuple pointed at by the name in the register k and stores it in registers i and j .
- $\text{update } i, j$ writes the value stored in register j to the second component of the value pointed to by the name in register i .
- $\text{free } i$ releases the memory pointed to by the name in the register i and resets the register.
- $\text{flip } i, j$ flips the values of registers i and j .
- $i \leftarrow \text{set } j$ sets register i to value j .

Let code fragments \mathcal{C} be $\mathcal{C} ::= \text{Instr}; \mathcal{C} \mid \text{ifzero } \mathbb{N} \mathcal{C} \mathcal{C} \mid \text{spark } a \mid \text{end}$. The port names occurring in the code fragment are $\text{sup} \in \mathcal{C} \rightarrow \mathcal{P}_{\text{fin}}(\mathbb{A})$, defined in the obvious way (only the `spark a` instruction can contribute names). An `ifzero i` instruction will branch according to the value stored in register i . A `spark a` will either jump to a or send a message to a , depending on whether a is a local port or not.

An *engine* is an interface together with a port map, $E = (A, P) \in \mathcal{I} \times (\text{sup}(A^{(\mathbf{O})}) \rightarrow \mathcal{C})$ such that for each code fragment $c \in \text{cod } P$ and each port name $a \in \text{sup}(c)$, $(\mathbf{P}, a) \in A$, meaning that ports that are “sparked” must be output ports of the interface A . The set of all engines is \mathcal{E} .

Engines have threads and shared heap. All threads have a fixed number of registers r , which is a global constant. For the language ICA we will need four registers, but languages with more kinds of pointers in the game model, e.g. control pointers [19], may need and use more registers.

A *thread* is a tuple $t = (c, \vec{d}) \in T = \mathcal{C} \times \mathcal{D}^r$: a code fragment and an r -tuple of data register values.

An *engine configuration* is a tuple $k = (\vec{t}, h) \in \mathcal{K} = \mathcal{P}_{\text{fin}}(T) \times (\mathbb{P} \rightarrow \mathbb{P} \times \mathcal{D})$: a set of threads and a heap that maps pointer names to pairs of pointer names and data items.

A pair consisting of an engine configuration and an engine will be written using the notation $k : E \in \mathcal{K} \times \mathcal{E}$. Define the function *initial* $\in \mathcal{E} \rightarrow \mathcal{K} \times \mathcal{E}$ as *initial*(E) $\triangleq (\emptyset, \emptyset) : E$ for an engine E . This function pairs the engine up with an engine configuration consisting of no threads and an empty heap.

HRAMs communicate using *messages*, each consisting of a port name and a vector of data items of size r_m : $m = (x, \vec{d}) \in \mathcal{M} = \mathbb{A} \times \mathcal{D}^{r_m}$. The constant r_m specifies the size of the messages in the network, and has to fulfil $r_m \leq r$. For a set $X \subseteq \mathbb{A}$, define $\mathcal{M}_X = X \times \mathcal{D}^{r_m}$, the subset of \mathcal{M} whose port names are limited to those of X .

We specify the operational semantics of an engine $E = (A, P)$ as a transition relation $\xrightarrow[E, \chi]{-} \subseteq \mathcal{K} \times (\{\bullet\} \cup (L \times \mathcal{M})) \times \mathcal{K}$. The relation is either labelled with \bullet — a silent transition — or a polarised message — an observable transition. The messages will be constructed simply from the first r_m registers of a thread, meaning that on certain actions part of the register contents become observable in the transition relation.

When a transition is tagged by \mathbf{P} we elide the label (output) and when it is tagged by \mathbf{O} (input) we tag it with $-\bullet$.




Fig. 1. Example HRAM net

The network connectivity is specified by the function χ , which will be described in more detail in the next sub-section. For a port name a , $\chi(a)$ can be read as “the port that a is connected to”. The interesting rule is that for `spark` because it depends on whether the port where the next computation is “sparked” is local or not. If the port is local then `spark` makes a jump, and if the port is non-local then it produces an output token and the current thread of execution is terminated, similar to the IAM. The set of threads \bar{t} and the heap h which frame these rules are unchanged by the rules for `spark` so we elide them for readability. Since messages may have fewer components than we have registers, we use the functions $msg \in \mathcal{D}^r \rightarrow \mathcal{D}^{r_m}$, which takes the first r_m components of its input, and $regs \in \mathcal{D}^{r_m} \rightarrow \mathcal{D}^r$, which pads its input with \emptyset at the end (i.e. $regs(\bar{d}) \triangleq (d_0, \dots, d_{r_m-1}, \emptyset, \dots)$).

$$\begin{aligned}
 (\text{spark } a, \bar{d}) &\xrightarrow[E, \chi]{(\chi(a), msg(\bar{d}))} \emptyset, & (\mathbf{O}, \chi(a)) \notin A. \\
 (\text{spark } a, \bar{d}) &\xrightarrow[E, \chi]{} (P(\chi(a)), regs(msg(\bar{d}))), & (\mathbf{O}, \chi(a)) \in A
 \end{aligned}$$

An input token sparks a new thread. The framing thread set and heap are again elided for simplicity:

$$\emptyset \xrightarrow[E, \chi]{(a, \bar{d})^\bullet} (P(a), regs(\bar{d})), \quad (\mathbf{O}, a) \in A.$$

B. HRAM nets

A well-formed *HRAM net* $S \in \mathcal{S}$ is a set of engines, a function over port names specifying what ports are connected, and an external interface, $S = (\bar{E}, \chi, A)$, where $E \in \mathcal{E}$, $A \in \mathcal{I}$, and χ is a permutation of port names.

Fig. 1 shows a diagram of an HRAM net with two HRAMs (interfaces A, A' , two ports each), each with two running threads (t_i, t'_i) with local registers (d_i, d'_i) and shared heaps (h, h'). Two of the HRAM ports are connected and two are part of the global interface B .

The function χ gives the network connectivity, mapping each input port name of the net’s interface and output port name of the net’s engines to either an output port name of the net’s interface or an input port name of one of its engines. Since it is a bijection, each port name (and thus port) is connected to exactly one other port name, so the abstract network model we are using is point-to-point.

For an engine $e = (A, P)$, we define a *singleton* net with e as its sole engine as $singleton(e) = (\{e\}, \chi, A')$, where A' is an interface such that $\chi \vdash A =_{\mathbb{A}} A'$.

A *net configuration* is a set of tuples of engine configurations and engines and a multiset of pending messages: $n = (\bar{e} : \bar{E}, \bar{m}) \in \mathcal{N} = \mathcal{P}_{fin}(\mathcal{K} \times \mathcal{E}) \times \mathbf{Mset}_{fin}(\mathcal{M})$. Define the function $initial \in \mathcal{S} \rightarrow \mathcal{N}$ as $initial(\bar{E}, \chi, A) \triangleq (\{initial(E) \mid E \in \bar{E}\}, \emptyset)$, a net configuration with only initial engines and no pending messages.

The operational semantics of a net $S = (\bar{E}, \chi, A)$ is specified as a transition relation $- \xrightarrow{\bullet} - \subseteq \mathcal{N} \times (\{\bullet\} \cup (L \times \mathcal{M}_{sup(A)})) \times \mathcal{N}$. The semantics is given in the style of the Chemical Abstract Machine (CHAM) [3], where HRAMs are “molecules” and the pending messages of the HRAM net is the “solution”. HRAM inputs (outputs) are to (from) the set of pending messages. Silent transitions of any HRAM are silent transitions of the net.

C. Semantics of HRAM nets

We define $\mathbf{List}[A]$ for a set A to be finite sequences of elements from A , and use $s::s'$ for concatenation. A *trace* for a net (\bar{E}, χ, A) is a *finite sequence* of messages with polarity: $s \in \mathbf{List}[L \times \mathcal{M}_{sup(A)}]$. Write $\alpha \in L \times \mathcal{M}_{sup(A)}$ for single polarised messages. We use the same notational convention as before to identify inputs ($-\bullet$).

For a trace $s = \alpha_1::\alpha_2::\dots::\alpha_n$, define \xrightarrow{s} to be the following composition of relations on net configurations: $\xrightarrow{\alpha_1} \xrightarrow{*} \xrightarrow{\alpha_2} \xrightarrow{*} \dots \xrightarrow{\alpha_n}$, where $\xrightarrow{*}$ is the reflexive transitive closure of \rightarrow , i.e. any number of silent steps are allowed in between those that are observable.

Write $traces_A$ for the set $\mathbf{List}[L \times \mathcal{M}_{sup(A)}]$. The denotation $\llbracket S \rrbracket \subseteq traces_A$ of a net $S = (\bar{E}, \chi, A)$ is the set of traces of observable transitions reachable from the initial net configuration $initial(S)$ using the transition relation:

$$\llbracket S \rrbracket \triangleq \{s \in traces_A \mid \exists n. initial(S) \xrightarrow{s} n\}$$

The denotation of a net includes the empty trace and is prefix-closed by construction.

As with interfaces, we are not interested in the actual port names occurring in a trace, so we define *equivariance* for sets of traces. Let $S_1 \subseteq traces_{A_1}$ and $S_2 \subseteq traces_{A_2}$ for $A_1, A_2 \in \mathcal{I}$. $S_1 =_{\mathbb{A}} S_2$ if and only if there is a permutation $\pi \in \mathbb{A} \rightarrow \mathbb{A}$ such that $\{\pi \cdot s \mid s \in S_1\} = S_2$, where $\pi \cdot \epsilon \triangleq \epsilon$ and $\pi \cdot (s::(l, (a, \bar{d}))) \triangleq (\pi \cdot s)::(l, (\pi(x), \bar{d}))$.

Define the *deletion* operation $s-A$ which removes from a trace all elements $(l, (x, \bar{d}))$ if $x \in sup(A)$ and define the interleaving of sets of traces $S_1 \subseteq traces_A$ and $S_2 \subseteq traces_B$ as $S_1 \otimes S_2 \triangleq \{s \mid s \in traces_{A \otimes B} \wedge s-B \in S_1 \wedge s-A \in S_2\}$.

Define the composition of the sets of traces $S_1 \subseteq traces_{A \Rightarrow B}$ and $S_2 \subseteq traces_{B' \Rightarrow C}$ with $\pi \vdash B =_{\mathbb{A}} B'$ as the usual *synchronisation and hiding* in trace semantics:

$$\begin{aligned}
 S_1; S_2 &\triangleq \{s-B \mid s \in traces_{A \otimes B \otimes C} \wedge s-C \in S_1 \\
 &\quad \wedge \pi \cdot s^*B-A \in S_2\}
 \end{aligned}$$

(where s^{*B} is s where the messages from B have reversed polarity.)

Two nets, $f = (\overline{E}_f, \chi_f, I_f)$ and $g = (\overline{E}_g, \chi_g, I_g)$ are said to be *structurally equivalent* if they are graph-isomorphic, i.e. $\pi \cdot \overline{E}_f = \overline{E}_g$, $\pi \vdash I_f =_{\mathbb{A}} I_g$ and $\chi_g \circ \pi = \pi \circ \chi_f$.

Theorem II.1. *If S_1 and S_2 are structurally equivalent nets, then $\llbracket S_1 \rrbracket =_{\mathbb{A}} \llbracket S_2 \rrbracket$.*

HRAM nets form a category where objects are interfaces $A \in \mathcal{P}_{fin}(\text{Port})$ identified up to equivariance, and morphisms $f : A \rightarrow B$ are well-formed nets of the form $(\overline{E}, \chi, A \Rightarrow B)$, for some \overline{E} and χ . We identify morphisms that have the same denotation, up to equivariance, i.e. if $\pi \vdash A =_{\mathbb{A}} B$ and $\pi \vdash \llbracket f \rrbracket =_{\mathbb{A}} \llbracket g \rrbracket$ then $f = g$ (in the category). The identity morphism for an object A is $id_A \triangleq (\emptyset, \chi, A \Rightarrow A')$, a net without HRAMs, for an A' such that $\pi \vdash A =_{\mathbb{A}} A'$ and a χ that takes \mathbf{O} -labelled port names $a \in A^{*(\mathbf{O})}$ to $\pi(a)$ and similarly for $A'^{(\mathbf{O})}$. This means that the identity is pure connectivity.

Composition of two morphisms $f = (\overline{E}_f, \chi_f, A \Rightarrow B) : A \rightarrow B$ and $g = (\overline{E}_g, \chi_g, B' \Rightarrow C) : B' \rightarrow C$, such that $\pi \vdash B =_{\mathbb{A}} B'$, is $f;g = (\overline{E}_f \cup \overline{E}_g, \chi_{f;g}, A \Rightarrow C) : A \rightarrow C$ where $\chi_{f;g}$ connects the ports of B and B' .

The following results establishes that our definitions are sensible and that HRAM nets are equal up to topological isomorphisms. This result also shows that the structure of HRAM nets is very loose.

Theorem II.2. *HRAM nets form a symmetric compact-closed category, called **HRAMnet**.*

Note: We identify HRAMs with interfaces of the same shape in the category, which means that our objects and morphisms are in reality unions of equivariant sets. In defining the operations of our category we use *representatives* for these sets, and require that the representatives are chosen such that their sets of port names are disjoint (but same-shaped when the operation calls for it). The composition operation may appear to be partial because of this requirement, but we can always find equivariant representatives that fulfil it.

The tensor product of two objects A, B , is the interface tensor, $A \otimes B$, as defined before. The tensor of two morphisms $f = (\overline{E}_f, \chi_f, A \Rightarrow B), g = (\overline{E}_g, \chi_g, C \Rightarrow D)$ as $f \otimes g = (\overline{E}_f \cup \overline{E}_g, \chi_f \otimes \chi_g, A \otimes C \Rightarrow B \otimes D)$. The requisite isomorphisms of the category are all the obvious identity-like nets of pure connectivity.

The following result explicates how communicating HRAMs can be combined into a single machine, where the intercommunication is done with jumping rather than message passing, in a sound way:

Theorem II.3. *If $E_1 = (A_1, P_1)$ and $E_2 = (A_2, P_2)$ are engines and $S = (\{E_1, E_2\}, \chi, A)$ is a net, then $E_{12} = (A_1 \otimes A_2, P_1 \cup P_2)$ is an engine, $S' = (\{E_{12}\}, \chi, A)$ is a net and $\llbracket S \rrbracket \subseteq \llbracket S' \rrbracket$.*

We define a family of projection HRAM nets $\Pi_{i, A_1 \otimes \dots \otimes A_n} : A_1 \otimes \dots \otimes A_n \rightarrow A_i$ by first constructing a family of “sinks” $!_A : A \rightarrow I \triangleq \text{singleton}((A \Rightarrow I, P))$ where $I = \emptyset$ and

$P(a) = \text{end}$ for each a in its domain and then defining e.g. $\Pi_{1, A \otimes B} : A \otimes B \rightarrow A \triangleq id_A \otimes !_B$.

III. GAME NETS FOR ICA

The structure of a **HRAMnet** token is determined by the number of registers r and the message size r_m , which are globally fixed. To implement game-semantic machines we require four message components: a port name, two pointer names, and a data fragment, meaning that $r_m = 3$. We choose $r = 4$, to get an additional register for temporary thread values to work with. The message structure is intended to capture the structure of a move when game semantics is expressed in the nominal model. The port name is the move, the first name is the “point” whereas the second name is the “butt” of a justification arrow, and the data is the value of the move. This direct and abstract encoding of the justification pointer as names is quite different to that used in PAM and in other GOI-based token machines. In PAM the pointer is represented by a sequence of integers encoding the hereditary justification of the move, which is a snap-shot of the computational causal history of the move, just like in GOI-based machines. Such encodings have an immediate negative consequence, as tokens can become impractically large in complex computations, especially involving recursion. Large tokens entail not only significant communication overheads but also the computational overheads of decoding their structure. A subtler negative consequence of such an encoding is that it makes supporting the semantic structures required to interpret state and concurrency needlessly complicated and inefficient. The nominal representation is simple and compact, and efficiently exploits local machine memory (heap) in a way that previous abstract machines, of a “functional” nature, do not.

The price that we pay is a failure of compositionality, which we will illustrate shortly. The rest of the section will show how compositionality can be restored without substantially changing the HRAM framework. If in HRAM nets compositionality is “plug-and-play”, as apparent from its compact-closed structure, *Game Abstract Machine* (GAM) composition must be mediated by a family of operators which are themselves HRAMs.

In this simple motivating example it is assumed that the reader is familiar with game semantics, and several of the notions to be introduced formally in the next sub-sections are anticipated. We trust that this will be not confusing.

Let S be a HRAM representing the game semantic model for the *successor* operation $S : \text{int} \rightarrow \text{int}$. The HRAM net in Fig. 2 represents a (failed) attempt to construct an interpretation for the term $x : \text{int} \vdash S(S(x)) : \text{int}$ in a context $C[-\text{int}] : \text{int}$. This is the standard way of composing GOI-like machines.

The labels along the edges of the HRAM net trace a token (a, p_0, p_1, d) sent by the context $C[-]$ in order to evaluate the term. We elide a and d , which are irrelevant, to keep the diagram uncluttered. The token is received by S and propagated to the other S HRAM, this time with tokens (p_1, p_2) . This trace of events $(p_0, p_1)::(p_1, p_2)$ corresponds




Fig. 2. Non-locality of names in HRAM composition

to the existence of a justification pointer from the second action to the first in the game model. The essential correctness invariant for a well-formed trace representing a game-semantic play is that each token consists of a *known* name and a *fresh* name (if locally created, or *unknown* if externally created). However, the second S machine will respond with (p_2, p_3) to (p_1, p_2) , leading to a situation where $C[-]$ receives a token formed from two unknown tokens.

In game semantics, the composition of $(p_0, p_1)::(p_1, p_2)$ with $(p_1, p_2)::(p_2, p_3)$ should lead to $(p_0, p_1)::(p_1, p_3)$, as justification pointers are “extended” so that they never point into a move hidden through composition. This is precisely what the composition operator, a specialised HRAM, will be designed to achieve.

A. Game abstract machines (GAM) and nets

Definition III.1. We define a game interface (cf. arena) as a tuple $\mathfrak{A} = (A, \text{qst}_{\mathfrak{A}}, \text{ini}_{\mathfrak{A}}, \vdash_{\mathfrak{A}})$ where

- $A \in \mathcal{I}$ is an interface. For game interfaces $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}$ we will write A, B, C and so on for their underlying interfaces.
- The set of ports is partitioned into a subset of question port names $\text{qst}_{\mathfrak{A}}$ and one of answer port names $\text{ans}_{\mathfrak{A}}$, $\text{qst}_{\mathfrak{A}} \uplus \text{ans}_{\mathfrak{A}} = \text{sup}(A)$.
- The set of initial port names $\text{ini}_{\mathfrak{A}}$ is a subset of the \mathbf{O} -labelled question ports.
- The enabling relation $\vdash_{\mathfrak{A}}$ relates question port names to non-initial port names such that if $a \vdash_{\mathfrak{A}} a'$ for port names $a \in \text{qst}_{\mathfrak{A}}$ with $(l, a) \in A$ and $a' \in \text{sup}(A) \setminus \text{ini}_{\mathfrak{A}}$ with $(l', a') \in A$, then $l \neq l'$.

For notational consistency, write $\text{opp}_{\mathfrak{A}} \triangleq \text{sup}(A^{(\mathbf{O})})$ and $\text{prop}_{\mathfrak{A}} \triangleq \text{sup}(A^{(\mathbf{P})})$. Call the set of all game interfaces $\mathcal{I}_{\mathfrak{G}}$. Game interfaces are equivariant, $\pi \vdash \mathfrak{A} =_{\mathbb{A}} \mathfrak{B}$, if and only if $\pi \vdash A =_{\mathbb{A}} B$, $\{\pi(a) \mid a \in \text{qst}_{\mathfrak{A}}\} = \text{qst}_{\mathfrak{B}}$, $\{\pi(a) \mid a \in \text{ini}_{\mathfrak{A}}\} = \text{ini}_{\mathfrak{B}}$ and $\{(\pi(a), \pi(a')) \mid a \vdash_{\mathfrak{A}} a'\} = \vdash_{\mathfrak{B}}$.

Definition III.2. For game interfaces (with disjoint sets of port names) \mathfrak{A} and \mathfrak{B} , we define:

$$\mathfrak{A} \otimes \mathfrak{B} \triangleq (A \otimes B, \text{qst}_{\mathfrak{A}} \cup \text{qst}_{\mathfrak{B}}, \text{ini}_{\mathfrak{A}} \cup \text{ini}_{\mathfrak{B}}, \vdash_{\mathfrak{A}} \cup \vdash_{\mathfrak{B}})$$

$$\mathfrak{A} \Rightarrow \mathfrak{B} \triangleq (A \Rightarrow B, \text{qst}_{\mathfrak{A}} \cup \text{qst}_{\mathfrak{B}}, \text{ini}_{\mathfrak{B}}, \vdash_{\mathfrak{A}} \cup \vdash_{\mathfrak{B}} \cup (\text{ini}_{\mathfrak{B}} \times \text{ini}_{\mathfrak{A}})).$$

A GAM net is a tuple $G = (S, \mathfrak{A}) \in \mathcal{S} \times \mathcal{I}_{\mathfrak{G}}$ consisting of a net and a game interface such that $S = (\bar{E}, \chi, A)$, i.e. the interface of the game net is the same as that of the game interface. The denotational semantics of a GAM




Fig. 3. A typical play for copycat

net $G = (S, \mathfrak{A})$ is just that of the underlying HRAM net: $\llbracket G \rrbracket \triangleq \llbracket S \rrbracket$.

B. Copycat

The quintessential game-semantic behaviour is that of the *copy-cat strategy*, as it appears in various guises in the representation of all structural morphisms of any category of strategies. A copy-cat not only replicates the behaviour of its Opponent in terms of moves, but also in terms of justification structures. Because of this, the copy-cat strategy needs to be either history-sensitive (stateful) or the justification information needs to be carried along with the token. We take the former approach, in contrast to IAM and other GOI-inspired machines.

Consider the identity (or copycat) strategy on $\text{com} \Rightarrow \text{com}$, where com is a two-move arena (one question, one answer). A typical play may look as in Fig. 3. The full lines represent justification pointers, and the trace (play) is represented nominally as

$$(r_4, p_0, p_1)::(r_2, p_1, p_2)::(r_1, p_2, p_3)::(r_3, p_1, p_4)::(d_3, p_4) \cdots$$

To preserve the justification structure, a copycat engine only needs to store “copycat links”, which are shown as dashed lines in the diagram between question moves. In this instance, for an input on r_4 , a heap value mapping a freshly created p_2 (the pointer to r_2) to p_1 (the pointer from r_4) is added.

The reason for mapping p_2 to p_1 becomes clear when the engine later gets an input on r_1 with pointers p_2 and p_3 . It can then replicate the move to r_3 , but using p_1 as a justifier. By following the p_2 pointer in the heap it gets p_1 so it can produce (r_3, p_1, p_4) , where p_4 is a fresh heap value mapping to p_3 . When receiving an answer, i.e. a d move, the copycat link can be dereferenced and then *discarded* from the heap.

The following HRAM macro-instructions are useful in defining copy-cat machines to, respectively, handle the point-

ers in an initial question, a non-initial question and an answer:

$$\begin{aligned} \text{cci} &\triangleq \text{flip } 0, 1; 1 \leftarrow \text{new } 0, 3 \\ \text{ccq} &\triangleq 1 \leftarrow \text{new } 1, 3; 0, 3 \leftarrow \text{get } 0 \\ \text{cca} &\triangleq \text{flip } 0, 1; 0, 3 \leftarrow \text{get } 1; \text{free } 1 \end{aligned}$$

For game interfaces \mathfrak{A} and \mathfrak{A}' such that $\pi \vdash \mathfrak{A} =_{\mathbb{A}} \mathfrak{A}'$, we define a generalised copycat engine as $\mathcal{C}_{C,\pi,\mathfrak{A}} = (A \Rightarrow A', P)$, where:

$$\begin{aligned} P &\triangleq \{q_2 \mapsto C; \text{spark } q_1 \mid q_2 \in \text{ini}_{\mathfrak{A}'} \wedge q_1 = \pi^{-1}(q_2)\} \\ &\cup \{q_2 \mapsto \text{ccq}; \text{spark } q_1 \\ &\quad \mid q_2 \in (\text{opp}_{\mathfrak{A}'} \cap \text{qst}_{\mathfrak{A}'}) \setminus \text{ini}_{\mathfrak{A}'} \wedge q_1 = \pi^{-1}(q_2)\} \\ &\cup \{a_2 \mapsto \text{cca}; \text{spark } a_1 \\ &\quad \mid a_2 \in \text{opp}_{\mathfrak{A}'} \cap \text{ans}_{\mathfrak{A}'} \wedge a_1 = \pi^{-1}(a_2)\} \\ &\cup \{q_1 \mapsto \text{ccq}; \text{spark } q_2 \mid q_1 \in \text{opp}_{\mathfrak{A}} \cap \text{qst}_{\mathfrak{A}} \wedge q_2 = \pi(q_1)\} \\ &\cup \{a_1 \mapsto \text{cca}; \text{spark } a_2 \mid a_1 \in \text{opp}_{\mathfrak{A}} \cap \text{ans}_{\mathfrak{A}} \wedge a_2 = \pi(a_1)\} \end{aligned}$$

This copycat engine is parametrised with an initial instruction C , which is run when receiving an initial question. The engine for an ordinary copycat, i.e. the identity of games, is $\mathcal{C}_{\text{cci},\pi,\mathfrak{A}}$. By slight abuse of notation, write $\mathcal{C}_{\mathfrak{A}}$ for the singleton copycat game net ($\text{singleton}(\mathcal{C}_{\text{cci},\pi,\mathfrak{A}}), \mathfrak{A} \Rightarrow \pi \cdot \mathfrak{A}$).

Following [9], we define a partial order \leq over polarities, L , as $\mathbf{O} \leq \mathbf{O}, \mathbf{O} \leq \mathbf{P}, \mathbf{P} \leq \mathbf{P}$ and a preorder \preceq over traces from $P_{\mathfrak{A}}$ to be the least reflexive and transitive such that if $l_1 \leq l_2$ then

$$\begin{aligned} s_1 :: (l_1, (a_1, p_1, p'_1, d_1)) :: (l_2, (a_2, p_2, p'_2, d_2)) :: s_2 \\ \preceq s_1 :: (l_2, (a_2, p_2, p'_2, d_2)) :: (l_1, (a_1, p_1, p'_1, d_1)) :: s_2, \end{aligned}$$

where $p'_1 \neq p_2$. A set of traces $S \subseteq P_{\mathfrak{A}}$ is *saturated* if and only if, for $s, s' \in P_{\mathfrak{A}}$, $s' \preceq s$ and $s \in S$ implies $s' \in S$. If $S \subseteq P_{\mathfrak{A}}$ is a set of traces, let $\text{sat}(S)$ be the smallest saturated set of traces that contains S .

The usual definition of the copycat strategy (in the alternating and single-threaded setting) as a set of traces is

$$\mathcal{C}_{\mathfrak{A},\mathfrak{A}'}^{\text{st,alt}} \triangleq \{s \in P_{\mathfrak{A} \Rightarrow \mathfrak{A}'}^{\text{st,alt}} \mid \forall s' \leq_{\text{even}} s. s'^* \upharpoonright A =_{\Delta \mathbf{P}} s' \upharpoonright A'\}$$

Definition III.3. A set of traces S_1 is \mathbf{P} -closed with respect to a set of traces S_2 if and only if $s' \in S_1 \cap S_2$ and $s = s' :: (\mathbf{P}, (a, p, p', d)) \in S_1$ implies $s \in S_2$.

The intuition of \mathbf{P} -closure is that if the trace s' is “legal” according to S_2 , then any outputs that can occur after s' in S_1 are also legal.

Definition III.4. We say that a GAM net f implements a set of traces S if and only if $S \subseteq \llbracket f \rrbracket$ and $\llbracket f \rrbracket$ is \mathbf{P} -closed with respect to S .

This is the form of the statements of correctness for game nets that we want; it certifies that the net f can accommodate all traces in S and, furthermore, that it only produces legal outputs when given valid inputs.

The main result of this section establishes the correctness of the GAM for copycat.

Theorem III.5. $\mathcal{C}_{\pi,\mathfrak{A}}$ implements $\mathcal{C}_{\mathfrak{A},\pi \cdot \mathfrak{A}}$.

C. Composition

The definition of composition in Hyland-Ong games [18] is eerily similar to our definition of trace composition, so we might expect HRAM net composition to correspond to it. That is, however, only superficially true: the nominal setting that we are using [9] brings to light what happens to the justification pointers in composition.

If A is an interface, $s \in \text{traces}_A$ and $X \subseteq \text{sup}(A)$, we define the *reindexing deletion* operator $s \downarrow X$ as follows, where $(s', \rho) = s \downarrow X$ inductively:

$$\begin{aligned} \epsilon \downarrow X &\triangleq (\epsilon, \text{id}) \\ s :: (l, (a, p, p', d)) \downarrow X &\triangleq (s' :: (l, (a, \rho(p), p', d)), \rho) \quad \text{if } a \notin X \\ s :: (l, (a, p, p', d)) \downarrow X &\triangleq (s', \rho \cup \{p' \mapsto \rho(p)\}) \quad \text{if } a \in X \end{aligned}$$

We write $s \downarrow X$ for s' when $s \downarrow X = (s', \rho)$ in the following definition:

Definition III.6. The game composition of the sets of traces $S_1 \subseteq \text{traces}_{A \Rightarrow B}$ and $S_2 \subseteq \text{traces}_{B' \Rightarrow C}$ with $\pi \vdash B =_{\mathbb{A}} B'$ is

$$\begin{aligned} S_1;_{\mathfrak{B}} S_2 &\triangleq \{s \downarrow B \mid s \in \text{traces}_{A \otimes B \otimes C} \wedge s \downarrow C \in S_1 \\ &\quad \wedge \pi \cdot s^{*B} \downarrow A \in S_2\} \end{aligned}$$

Clearly we have $S_1; S_2 \neq S_1;_{\mathfrak{B}} S_2$ for sets of traces S_1 and S_2 , which reinforces the practical problem in the beginning of this section.

Composition is constructed out of three copycat-like behaviours, as sketched in Fig. 4 for a typical play at some types A, B and C . As a trace in the nominal model, this is:

$$\begin{aligned} (q_6, p_0, p_1) :: (q_4, p_1, p_2) :: (q_3, p_2, p_3) :: \\ (q_2, p_1, p_4) :: (q_1, p_4, p_5) :: (q_5, p_1, p_6) :: (a_5, p_6) :: \\ (a_1, p_5) :: (a_2, p_4) :: (a_3, p_3) :: (a_4, p_2) :: (a_6, p_1) \end{aligned}$$

We see that this *almost* corresponds to three interleaved copycats as described above; between A, B, C and A', B', C' . There is, however, a small difference: The move q_1 , if it were to blindly follow the recipe of a copycat, would dereference the pointer p_4 , yielding p_3 , and so incorrectly make the move q_5 justified by q_3 , whereas it really should be justified by q_6 as in the diagram. This is precisely the problem explained at the beginning of this section.

To make a pointer *extension*, when the B -initial move q_3 is performed, it should map p_4 not only to p_3 , but also to the pointer that p_2 points to, which is p_1 (the dotted line in the diagram). When the A -initial move q_1 is performed, it has access to both of these pointers that p_4 maps to, and can correctly make the q_5 move by associating it with pointers p_1 and a fresh p_6 .

$$(A \Rightarrow B) \otimes (B' \Rightarrow C) \rightarrow (A' \Rightarrow C')$$




Fig. 4. Composition from copycat

Let \mathfrak{A}' , \mathfrak{B}' , and \mathfrak{C}' be game interfaces such that $\pi_{\mathfrak{A}} \vdash \mathfrak{A} =_{\mathbb{A}} \mathfrak{A}'$, $\pi_{\mathfrak{B}} \vdash \mathfrak{B} =_{\mathbb{A}} \mathfrak{B}'$, $\pi_{\mathfrak{C}} \vdash \mathfrak{C} =_{\mathbb{A}} \mathfrak{C}'$, and

$$\begin{aligned} (A' \Rightarrow A, P_A) &= \mathcal{C}_{\text{exq}, \pi_{\mathfrak{A}}^{-1}, \mathfrak{A}'} \\ (B \Rightarrow B', P_B) &= \mathcal{C}_{\text{exi}, \pi_{\mathfrak{B}}, \mathfrak{B}} \\ (C \Rightarrow C', P_C) &= \mathcal{C}_{\text{cci}, \pi_{\mathfrak{C}}, \mathfrak{C}}, \text{ where} \\ \text{exi} &\triangleq 0, 3 \leftarrow \text{get } 0; 1 \leftarrow \text{new } 1, 0 \\ \text{exq} &\triangleq \emptyset, 0 \leftarrow \text{get } 0; 1 \leftarrow \text{new } 1, 3 \end{aligned}$$

Then the game composition operator $K_{\mathfrak{A}, \mathfrak{B}, \mathfrak{C}}$ is:

$$K_{\mathfrak{A}, \mathfrak{B}, \mathfrak{C}} \triangleq ((A \Rightarrow B) \otimes (B' \Rightarrow C) \Rightarrow (A' \Rightarrow C'), P_A \cup P_B \cup P_C).$$

Using the game composition operator K we can define GAM-net composition using **HRAMnet** compact closed combinators. Let $f : \mathfrak{A} \Rightarrow \mathfrak{B}, g : \mathfrak{B} \Rightarrow \mathfrak{C}$ be GAM-nets. Then their composition is defined as

$$\begin{aligned} f;_{GAM} g &\triangleq \Lambda_A^{-1}(\Lambda_A(f) \otimes \Lambda_B(g)); K_{\mathfrak{A}, \mathfrak{B}, \mathfrak{C}}, \text{ where} \\ \Lambda_A(f : A \rightarrow B) &\triangleq (\eta_A; (id_{A^*} \otimes f)) : I \rightarrow A^* \otimes B \\ \Lambda_A^{-1}(f : I \rightarrow A \otimes B) &\triangleq ((id_A \otimes f); (\varepsilon_A \otimes id_B)) : A \rightarrow B. \end{aligned}$$

Composition is represented diagrammatically as in Fig. 5. Note the comparison with the naive composition from Fig. 2. HRAMs f and g are not plugged in directly, although the interfaces match. Composition is mediated by the operator K , which preserves the locality of freshly generated names, exchanging non-local pointer names with local pointer names and storing the mapping between the two as copy-cat links, indicated diagrammatically by dotted lines in K .

Theorem III.7. *If $f : \mathfrak{A} \rightarrow \mathfrak{B}$ and $g : \mathfrak{B}' \rightarrow \mathfrak{C}$ are game nets such that $\pi_{\mathfrak{B}} \vdash \mathfrak{B} =_{\mathbb{A}} \mathfrak{B}'$, f implements $S_f \subseteq P_{\mathfrak{A} \Rightarrow \mathfrak{B}}$,*




Fig. 5. Composing GAMs using the K HRAM

and g implements $S_g \subseteq P_{\mathfrak{B}' \Rightarrow \mathfrak{C}}$, then $f;_{GAM} g$ implements $(S_f;_{\mathfrak{C}} S_g)$.

D. Diagonal

For game interfaces $\mathfrak{A}_1, \mathfrak{A}_2, \mathfrak{A}_3$ and permutations π_{ij} such that $\pi_{ij} \vdash \mathfrak{A}_i =_{\mathbb{A}} \mathfrak{A}_j$ for $i \neq j \in \{1, 2, 3\}$, we define the family of diagonal engines as:

$$\delta_{\pi_{12}, \pi_{13}, \mathfrak{A}} = (A_1 \Rightarrow A_2 \otimes A_3, P_1 \otimes P_2 \otimes P_3)$$

where, for $i \in \{2, 3\}$,

$$\begin{aligned} P_1 &\triangleq \{q_1 \mapsto \text{ccq}; \text{ifzero } 3 \text{ (spark } q_2) \text{ (spark } q_3) \\ &\quad | q_1 \in \text{opp}_{\mathfrak{A}_1} \cap \text{qst}_{\mathfrak{A}_1} \wedge q_2 = \pi_{12}(q_1) \wedge q_3 = \pi_{13}(q_1)\} \\ &\quad \cup \{a_1 \mapsto \text{cca}; \text{ifzero } 3 \text{ (spark } a_2) \text{ (spark } a_3) \\ &\quad | a_1 \in \text{opp}_{\mathfrak{A}_1} \cap \text{ans}_{\mathfrak{A}_1} \wedge a_2 = \pi_{12}(a_1) \wedge a_3 = \pi_{13}(a_1)\} \\ P_i &\triangleq \{q_i \mapsto 3 \leftarrow \text{set } (i-2); \text{cci}; \text{spark } q_1 \\ &\quad | q_i \in \text{ini}_{\mathfrak{A}_i} \wedge q_1 = \pi_{1i}^{-1}(q_i)\} \\ &\quad \cup \{q_i \mapsto \text{ccq}; \text{spark } q_1 \\ &\quad | q_i \in (\text{opp}_{\mathfrak{A}_i} \cap \text{qst}_{\mathfrak{A}_i}) \setminus \text{ini}_{\mathfrak{A}_i} \wedge q_1 = \pi_{1i}^{-1}(q_i)\} \\ &\quad \cup \{a_i \mapsto \text{cca}; \text{spark } a_1 \\ &\quad | a_i \in \text{opp}_{\mathfrak{A}_i} \cap \text{ans}_{\mathfrak{A}_i} \wedge a_1 = \pi_{1i}^{-1}(a_i)\}. \end{aligned}$$

The diagonal is almost identical to the copycat, except that an integer value of 0 or 1 is associated, in the heap, with the name of each message arriving on the A_2 and A_3 interfaces (hence the `set` statements, to be used for routing back messages arriving on A_1 using `ifzero` statements). By abuse of notation, we also write δ for the net $\text{singleton}(\delta)$.

Lemma III.8. *The δ net is the diagonal net, i.e. $\llbracket \delta_{\pi_{12}, \pi_{23}, \mathfrak{A}}; \Pi_i \rrbracket = \llbracket \mathcal{C}_{\pi_i, \mathfrak{A}_i} \rrbracket$.*

E. Fixpoint

We define a family of GAMs $\text{Fix}_{\mathfrak{A}}$ with interfaces $(\mathfrak{A}_1 \Rightarrow \mathfrak{A}_2) \Rightarrow \mathfrak{A}_3$ where there exist permutations $\pi_{i,j}$ such that $\pi_{i,j} \vdash \mathfrak{A}_i =_{\mathbb{A}} \mathfrak{A}_j$ for $i \neq j \in \{1, 2, 3\}$. The fixpoint engine is defined as $\text{Fix}_{\pi_{12}, \pi_{13}, \mathfrak{A}} = \Lambda_A^{-1}(\delta_{\pi_{12}, \pi_{13}, \mathfrak{A}})$.

Let $\text{fix}_{\pi_{12}, \pi_{13}, \mathfrak{A}} : (\mathfrak{A} \Rightarrow \pi_{12} \cdot \mathfrak{A}) \Rightarrow \pi_{13} \cdot \mathfrak{A}$ be the game-semantic strategy for fixpoint in Hyland-Ong games [18, p. 364].

Theorem III.9. $Fix_{\pi_{12}, \pi_{13}, \mathfrak{A}}$ implements $fix_{\pi_{12}, \pi_{13}, \mathfrak{A}}$.

The proof of this is immediate considering the three cases of moves from the definition of the game-semantic strategy. It is interesting to note here that we “force” a HRAM with interface $A_1 \Rightarrow A_2 \otimes A_3$ into a GAM with game interface $(\mathfrak{A}_3 \Rightarrow \mathfrak{A}_1) \Rightarrow \mathfrak{A}_2$, which has underlying interface $(A_3 \Rightarrow A_1) \Rightarrow A_2$. In the **HRAMnet** category, which is symmetric compact-closed, the two interfaces are isomorphic (with $A_1^* \otimes A_2 \otimes A_3$), but as game interfaces they are not. It is rather surprising that we can reuse our diagonal GAMs in such brutal fashion: in the game interface for fixpoint there is a reversed enabling relation between A_3 and A_1 . The reason why this still leads to legal plays only is because the onus of producing the justification pointers in the initial move for A_3 lies with the Opponent, which cannot exploit the fact that the diagonal is “wired illegally”. It only sees the fixpoint interface and must play accordingly. It is fair to say that that fixpoint interface is more restrictive to the Opponent than the diagonal interface, because the diagonal interface allows extra behaviours, e.g. sending initial messages in A_3 , which are no longer legal.

F. Other ICA constants

A GAM net for an integer literal n can be defined using the following engine (whose interface corresponds to the ICA exp type).

$$\text{lit}_n \triangleq (\{(\mathbf{O}, q), (\mathbf{P}, a)\}, P), \text{ where}$$

$$P \triangleq \{q \mapsto \text{flip } 0, 1; 1 \leftarrow \text{set } \emptyset; 2 \leftarrow \text{set } n; \text{spark } a\}$$

We see that upon getting an input question on port q , this engine will respond with a legal answer containing n as its value (register 2).

The conditional at type exp can be defined using the following engine, with the convention that $\{(\mathbf{O}, q_i), (\mathbf{P}, a_i)\} = \text{exp}_i$.

$$\text{if} \triangleq (\text{exp}_1 \Rightarrow \text{exp}_2 \Rightarrow \text{exp}_3 \Rightarrow \text{exp}_4, P), \text{ where}$$

$$P \triangleq \{q_4 \mapsto \text{cci}; \text{spark } q_1,$$

$$a_1 \mapsto \text{cca}; \text{flip } 0, 1; \text{cci};$$

$$\text{ifzero } 2 (\text{spark } q_3) (\text{spark } q_2),$$

$$a_2 \mapsto \text{cca}; \text{spark } a_4,$$

$$a_3 \mapsto \text{cca}; \text{spark } a_4\}$$

We can also define primitive operations, e.g. $+ : \text{exp} \Rightarrow \text{exp} \Rightarrow \text{exp}$, in a similar manner. An interesting engine is that for *newvar*:

$$\text{newvar} \triangleq ((\text{exp}_1 \otimes (\text{exp}_2 \Rightarrow \text{com}_3) \Rightarrow \text{exp}_4) \Rightarrow \text{exp}_5, P)$$

$$P \triangleq \{q_5 \mapsto 3 \leftarrow \text{set } 0; \text{cci}; \text{spark } q_4,$$

$$q_1 \mapsto \emptyset, 2 \leftarrow \text{get } 0; \text{flip } 0, 1; 1 \leftarrow \text{set } \emptyset;$$

$$\text{spark } a_1,$$

$$q_3 \mapsto \text{flip } 0, 1; 1 \leftarrow \text{new } 0, 1; \text{spark } q_2,$$

$$a_2 \mapsto \emptyset, 3 \leftarrow \text{get } 0; \text{update } 3 \ 2; \text{cca}; \text{spark } a_3,$$

$$a_4 \mapsto \text{cca}; \text{spark } a_5\}$$

We see that we store the variable in the second component of the justification pointer that justifies q_4 , so that it can be accessed in subsequent requests. A slight problem is that moves in exp_2 will actually not be justified by this pointer which we remedy in the q_3 case, by storing a pointer to the pointer with the variable as the second component of the justifier of q_2 , which means that we can access and update the variable in a_2 .

We can easily extend the HRAMs with new instructions to interpret parallel execution and semaphores, but we omit them from the current presentation.

IV. SEAMLESS DISTRIBUTED COMPILATION FOR ICA

A. The language ICA

ICA is PCF extended with constants to facilitate local effects. Its ground types are expressions and commands (exp, com) , with the type of assignable variables desugared as $\text{var} \triangleq \text{exp} \times (\text{exp} \rightarrow \text{com})$. Dereferencing and assignment are desugared as the first, respectively second, projections from the type of assignable variables. The local variable binder is $\text{new} : (\text{var} \rightarrow \text{com}) \rightarrow \text{com}$. ICA also has a type of split binary semaphores $\text{sem} \triangleq \text{com} \times \text{com}$, with the first and second projections corresponding to set, get , respectively (see [14] for the full definition, including the game-semantic model).

In this section we give a compilation method for ICA into GAM nets. The compilation is compositional on the syntax and it uses the constructs of the previous section. ICA types are compiled into GAM interfaces which correspond to their game-semantic arenas in the obvious way. We will use A, B, \dots to refer to an ICA type and to the GAM interface. Sec. III has already developed all the infrastructure needed to interpret the constants of ICA (Sec. III-F), including fixpoint (Sec. III-E). Given an ICA type judgment $\Gamma \vdash M : A$ with Γ a list of variable-type assignments $x_i : A_i$, M a term and A a type, a GAM implementing it G_M is defined compositionally on the syntax as follows:

$$G_{\Gamma \vdash MM' : A} = \delta_{\pi_1, \pi_2, \Gamma; \text{GAM}}$$

$$(G_{\Gamma \vdash M : A \rightarrow B} \otimes G_{\Gamma \vdash M' : B}); \text{GAM } \text{eval}_{A, B}$$

$$G_{\Gamma \vdash \lambda x : A. M : A \rightarrow B} = \Lambda_A (G_{\Gamma, x : A \vdash M : B})$$

$$G_{x : A, \Gamma \vdash x : A} = \Pi_{\emptyset A}; \mathbb{C}_{A, \pi},$$

Where $\text{eval}_{A, B} \triangleq \Lambda_B^{-1}(\mathbb{C}_{A \Rightarrow B, \pi})$ for a suitably chosen port renaming π and $\Pi_{\emptyset A}$ and $\Pi_{\emptyset 1}$ and $\Pi_{\emptyset 2}$ are HRAMs with signatures $\Pi_{\emptyset i} = (A_1 \otimes A_2 \Rightarrow A_3, P_i)$ such that they copycat between A_3 and A_i and ignore $A_{j \neq i}$. The interpretation of function application, which is the most complex, is shown diagrammatically in Fig. 6. The copycat connections are shown using dashed lines.

Theorem IV.1. *If M is an ICA term, G_M is the GAM implementing it and σ_M its game-semantic strategy then G_M implements σ_M .*




Fig. 6. GAM net for application




Fig. 7. Optimised GAM net for application

The correctness of compilation follows directly from the correctness of the individual GAM nets and the correctness of GAM composition $;_{GAM}$.

B. Prototype implementation

Following the recipe in the previous section we can produce an implementation of any ICA term as a GAM net. GAMs are just special-purpose HRAMs, with no special operations. HRAMs, in turn, can easily be implemented on any conventional computer with the usual store, control and communication facilities. A GAM net is also just a special-purpose HRAM net, which is a powerful abstraction of communication processes, as it subsumes through the `spark` instruction communication between processes (threads) on the same physical machine or located on distinct physical machines and communicating via a point-to-point network.

The actual distribution is achieved using light pragma-like code annotations. In order to execute a program at node A but delegate one computation to node B and another computation to node C we simply annotate an ICA program with node names, e.g.:

$$\{\text{new } x. x := \{f(x)\}@B + \{g(x)\}@C; !x\}@A$$

Note that this gives node B , via function f , read-write access to memory location x which is located at node A . Accessing non-local resources is possible, albeit possibly expensive.

Several facts make the compilation process quite remarkable:

- It is *seamless* (in the sense of [8]), allowing distributed compilation where communication is never explicit but always realised through function calls.
- It is *flexible*, allowing any syntactic sub-term to be located at any designated physical location, with no impact on the semantics of the program. The access of *non-local* resources is always possible, albeit possibly at a cost (latency, bandwidth, etc.).
- It is *dynamic*, allowing the relocation of GAMs to different physical nodes at run time. This can be done with extremely low overhead if the GAM heap is empty.
- It does not require any form of *garbage collection*, even on local nodes, although the language combines (ground) state, higher-order functions and concurrency. This is because a pointer associated with a pointer is not needed if and only if the question is answered; then it can be safely deallocated.

The current implementation does not perform any optimisations, and the resulting code is inefficient. Looking at the implementation of application in Fig. 6 it is quite clear that a message entering the GAM net via port A needs to undergo four pointer renamings before reaching the GAM for M . This is the cost we pay for compositionality. However, the particular configuration for application can be significantly simplified using standard peephole optimisation, and we can reach the much simpler, still correct implementation in Fig. 7. Here the functionality of the two compositions, the diagonal, and the *eval* GAMs have been combined and optimised into a single GAM, requiring only one pointer renaming before reaching M . Other optimisations can be introduced to simplify GAM nets, in particular to obviate the need for the use of composition GAMs K , for example by showing that composition of first-order closed terms (such as those used for most constants) can be done directly.

V. CONCLUSIONS, FURTHER WORK

In a previous paper we have argued that distributed and heterogeneous programming would benefit from the existence of architecture-agnostic, seamless compilation methods for conventional programming languages which can allow the programmer to focus on solving algorithmic problems without being overwhelmed by the minutiae of driving complex computational systems [8]. In *loc. cit.* we give such a compiler for PCF, based directly on the Geometry of Interaction. In this paper we show how Game Semantics can be expressed operationally using abstract machines very similar to networked conventional computers, a further development of the IAM/JAM game machines. We believe any programming language with a semantic model expressed as Hyland-Ong-style pointer games [18] can be readily represented using GAMs and then compiled to a variety of platforms such as MPI. Even more promising is the possible leveraging of more powerful infrastructure for distributed computing that can mask much of the complexities of distributed programming, such as fault-tolerance.

The compositional nature of the compiler is very important because it gives rise to a very general notion of foreign-function interface, expressible both as control and as communication, which allows a program to interface with other programs, in a syntax-independent way (see [13] for a discussion), opening the door to the seamless development of heterogeneous open systems in a distributed setting.

We believe we have established a solid foundational platform on which to build realistic seamless distributed compilers. Further work is needed in optimising the output of the compiler which is currently, as discussed, inefficient. The sources of inefficiency in this compiler are not just the generation of heavy-duty plumbing, but also the possibly unwise assignment of computation to nodes, requiring excessive network communication. Previous work in game semantics for resource usage can be naturally adapted to the operational setting of the GAMs and facilitate the automation of optimised task assignment [11].

REFERENCES

- [1] S. Abramsky, R. Jagadeesan, and P. Malacaria. Full Abstraction for PCF. *Inf. Comput.*, 163(2):409–470, 2000.
- [2] N. Benton. Embedded interpreters. *J. Funct. Program.*, 15(4):503–542, 2005.
- [3] G. Berry and G. Boudol. The Chemical Abstract Machine. In *Conference Record of the Seventeenth Annual ACM Symposium on Principles of Programming Languages, San Francisco, California, USA, January 1990*, pages 81–94. ACM Press, 1990.
- [4] P.-L. Curien and H. Herbelin. Abstract machines for dialogue games. *CoRR*, abs/0706.2544, 2007.
- [5] V. Danos, H. Herbelin, and L. Regnier. Game Semantics & Abstract Machines. In *Proceedings, 11th Annual IEEE Symposium on Logic in Computer Science, New Brunswick, New Jersey, USA, July 27-30, 1996*, pages 394–405. IEEE Computer Society, 1996.
- [6] V. Danos and L. Regnier. Reversible, Irreversible and Optimal lambda-Machines. *Theor. Comput. Sci.*, 227(1-2):79–97, 1999.
- [7] C. Faggian and F. Maurel. Ludics Nets, a game Model of Concurrent Interaction. In *20th IEEE Symposium on Logic in Computer Science (LICS 2005), 26-29 June 2005, Chicago, IL, USA, Proceedings*, pages 376–385. IEEE Computer Society, 2005.
- [8] O. Fredriksson and D. R. Ghica. Seamless distributed computing from the geometry of interaction. In *Trustworthy Global Computing*, 2012. forthcoming.
- [9] M. Gabbay and D. R. Ghica. Game Semantics in the Nominal Model. *Electr. Notes Theor. Comput. Sci.*, 286:173–189, 2012.
- [10] M. Gabbay and A. M. Pitts. A New Approach to Abstract Syntax Involving Binders. In *14th Annual IEEE Symposium on Logic in Computer Science, Trento, Italy, July 2-5, 1999*, pages 214–224. IEEE Computer Society, 1999.
- [11] D. R. Ghica. Slot games: a quantitative model of computation. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, pages 85–97. ACM, 2005.
- [12] D. R. Ghica. Applications of Game Semantics: From Program Analysis to Hardware Synthesis. In *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS 2009, 11-14 August 2009, Los Angeles, CA, USA*, pages 17–26. IEEE Computer Society, 2009.
- [13] D. R. Ghica. Function interface models for hardware compilation. In *9th IEEE/ACM International Conference on Formal Methods and Models for Codesign, MEMOCODE 2011, Cambridge, UK, 11-13 July, 2011*, pages 131–142. IEEE, 2011.
- [14] D. R. Ghica and A. S. Murawski. Angelic semantics of fine-grained concurrency. *Ann. Pure Appl. Logic*, 151(2-3):89–114, 2008.
- [15] D. R. Ghica and N. Tzevelekos. A System-Level Game Semantics. *Electr. Notes Theor. Comput. Sci.*, 286:191–211, 2012.
- [16] J.-Y. Girard. Geometry of interaction 1: Interpretation of System F. *Studies in Logic and the Foundations of Mathematics*, 127:221–260, 1989.
- [17] J. M. E. Hyland and C.-H. L. Ong. Pi-Calculus, Dialogue Games and PCF. In *FPCA*, pages 96–107, 1995.
- [18] J. M. E. Hyland and C.-H. L. Ong. On Full Abstraction for PCF: I, II, and III. *Inf. Comput.*, 163(2):285–408, 2000.
- [19] J. Laird. Exceptions, Continuations and Macro-expressiveness. In *Programming Languages and Systems, 11th European Symposium on Programming, ESOP 2002, held as Part of the Joint European Conference on Theory and Practice of Software, ETAPS 2002, Grenoble, France, April 8-12, 2002, Proceedings*, pages 133–146. Springer, 2002.
- [20] I. Mackie. The Geometry of Interaction Machine. In *Conference Record of POPL’95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Francisco, California, USA, January 23-25, 1995*, pages 198–208. ACM Press, 1995.
- [21] R. Milner. Functions as Processes. In *Automata, Languages and Programming, 17th International Colloquium, ICALP90, Warwick University, England, July 16-20, 1990, Proceedings*, pages 167–180. Springer, 1990.
- [22] C.-H. L. Ong. Verification of Higher-Order Computation: A Game-Semantic Approach. In *Programming Languages and Systems, 17th European Symposium on Programming, ESOP 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008, Proceedings*, pages 299–306. Springer, 2008.