

Seamless distributed computing from the Geometry of Interaction

Olle Fredriksson and Dan R. Ghica

The University of Birmingham, U.K.

Abstract. In this paper we present a seamless approach to writing and compiling distributed code. By “seamless” we mean that the syntax and semantics of the distributed program remain the same as if it was executed on one node only, except for label annotations indicating on what node sub-terms of the program are to be executed. There are no restrictions on how node labels are to be assigned to sub-terms. We show how the paradigmatic (higher-order functional recursive) programming language PCF, extended with node annotations, can be used for this purpose. The compilation technique is directly inspired by game semantics and the Geometry of Interaction.

1 Introduction

The conventional view of a program running in a distributed system, commonly called a *distributed program*, is of processes running on the nodes of a network, exchanging information by passing messages. This is the view supported by the ubiquitous *Message Passing Interface* system [15]. Although an effective method for writing distributed programs, MPI-style programming is far more laborious than programming a stand-alone computer. In order to create a distributed version of a program developed to run on a single node the code needs to be partitioned into modules that will run on each node in the distributed system, and the communication between the modules needs to be reduced to data communication and handled explicitly by the programmer.

A more abstract, and therefore more convenient, way of programming the interaction between process running on different nodes is through *Remote Procedure Call* [16]. In RPC, *communication* is subsumed by *function call*, which entails a certain protocol between the caller and the callee, i.e. first the arguments are sent by the caller along with the function name, then the result is returned by the callee. The success of RPC is reflected in the large number of variations it spawned, such as Java’s Remote Method Invocation, SOAP, CORBA’s object request broker, etc. With the advent of cloud computing more comprehensive frameworks were developed based on this idea, such as Microsoft’s Windows Communication Framework, Google’s Web Toolkit or Facebook’s Thrift framework, and aimed specifically at web-delivered computational services. Although significantly more convenient than MPI-style programming, RPC-style programming still requires a certain amount of boilerplate code and, more importantly, it is not compatible with higher-order functions and functional programming.

A higher-order version of RPC seems to require the sending of functions (code) from the caller to the callee and vice versa, because functions can now be both arguments and results. One solution is that all nodes have access to a local instance of each function in the program and can send references to such functions (possibly paired with some free variables, forming a closure). Erlang [1] and Cloud Haskell [4] take this approach. Erlang, which runs in a virtual machine, even allows the sending of syntax trees for terms that do not exist on the remote node. However, both these approaches have the disadvantage that a program running on a single node needs to be “ported” to the distributed setting by including significant amounts of non-trivial boilerplate code.

What we will describe in this paper is a *seamless* approach to distributed programming: the distributed program is syntactically and semantically identical to the same program running on the single node, except for annotations (labels) indicating the names of the nodes where particular terms are to be executed. There is no language-induced restriction regarding the way locations are assigned: any syntactic sub-term of the program can be given an arbitrary node label, which will mean that it will be executed on the node of that label. There is no explicit communication between nodes, all the interaction being automatically handled “under the hood” by the generated code.

Example To illustrate this point consider the same program written in Erlang versus PCF annotated with location information. Consider the PCF program $\mathbf{let } f = \lambda x. x * x \mathbf{ in } f 3 + f 4$, and suppose that we want to delegate the execution of the multiplication operation on a node C while the rest of the program executes on the (main) node A . The annotated PCF code is simply: $(\mathbf{let } f = (\lambda x. x * x) @ C \mathbf{ in } f 3 + f 4) @ A$

In Erlang, things are much more complicated. The f function can be set up as a server which receives request messages:

```
c(A_pid) -> receive X -> A_pid ! X * X end, c(A_pid).
main() ->
  C_pid = spawn(f, c, [self()]), C_pid ! 3,
  receive X -> C_pid ! 4, receive Y -> X + Y end
end.
```

Arguably, the logical structure of the program is lost in the detail. Moreover, if we want to further delegate the addition to a server B , the annotated PCF is $(\mathbf{let } f = (\lambda x. x * x) @ C \mathbf{ in } (f 3 + f 4) @ B) @ A$ whereas the Erlang version of the three-server distribution is even more complicated than the two-server version, which further obscures its logical structure:

```
c() -> receive {Pid, X} -> Pid ! X * X end, c().
b(A_pid, C_pid) ->
  receive
    request0 -> C_pid ! {self(), 3}, receive X -> A_pid ! X end;
    request1 -> C_pid ! {self(), 4}, receive X -> A_pid ! X end
  end,
```

```

    b(A_pid, C_pid).
main() ->
    C_pid = spawn(f2, c, []),
    B_pid = spawn(f2, b, [self(), C_pid]),
    B_pid ! request0,
    receive X -> B_pid ! request1, receive Y -> X + Y end
end.

```

Contribution The main technical challenge we address in this paper is handling higher-order and recursive computation. To be able to give a focussed and technically thorough treatment we will handle the paradigmatic functional programming language PCF [13], a language which is well understood semantically and which lies at the foundation of practical programming languages such as Haskell.

Conceptually, what makes the seamless distribution of PCF programs possible is an interpretation inspired by game semantics [6] and the Geometry of Interaction (GoI) [10]. These models are built on the principle of reducing function call to communication and computation to interaction.

Note that the idea of reducing computation to interaction also appeared in the context of process calculi [11]. This was a development independent of game semantics and GoI which happened around the same time. Compilation of distributed languages via process calculi is also possible, in languages such as Pict [12], but the methodology is different. Whereas we aim to hide all the communication by making it implicit in function call, Pict and related languages embed process calculus syntax to develop communication protocols. A further significant development in compilation based on process calculi was the idea of *mobile code*, where software agents are explicitly sent across the network between running processes [3, 14]. By contrast, in our methodology no code needs to be transmitted. Also note that GoI itself has been used before to compile PCF-like programs to unconventional architectures, namely reconfigurable digital circuits [7].¹

The GoI model reduces a program to a static network of elementary nodes. Although this is eminently suitable for hardware synthesis, where the elementary nodes become elementary circuits and the network becomes the circuit interconnect, it is too fine grained for distributed compilation. We address this technical challenge by introducing a new style of *abstract machine* which has elementary instructions for both control (jumps) and communication. These machines can be (almost) arbitrarily *combined* by replacing communication with jumps, which gives a high degree of control over the granularity of the network.

Our compiler works in several stages. First it creates the fine-grained network of elementary communicating abstract machines. Then, using node annotations (labels), it combines all machines which arise out of the compilation of terms using the same label. The final step is to compile the abstract machines down to executable code using C for local execution and MPI for inter-machine communication.

¹ Tool available at <http://veritygos.org>.

2 PCF and its GoI model

We use conventional PCF with natural numbers as the only base type, extended with an annotation $t @ A$, for specifying the locus of a term's computation. This annotation is to be thought of as a compiler directive: the operational semantics and typing rules simply ignore the locus specifier and are otherwise standard [13].

Girard's Geometry of Interaction [8] is a model for linear logic used for the study of the dynamics of computation, seeing a proof in the logic as a net, executed through the passing of a token along its edges. It is well known that GoI can be extended to also interpret programming languages and that it is useful in compiling programs to low-level machine code [10]. In this paper, we will use it as an interpretation for terms in our language, but we will use the notion of a proof net quite literally in that our programs will be compiled into networks of communicating nodes that operate on and send a token to each other.

We give an interpretation of terms in our language similar to the GoI interpretations by Hoshino [9] and Mackie [10], where terms are coded into linear logic proof nets. The only difference is that the interfaces of our nets are determined by type instead of being homogeneous. In this regard, the interpretation is more similar to the hardware circuits presented by the second author [5]. The reason why this works is because our language is not polymorphic, which otherwise necessitates that the interfaces are more homogeneous, since polymorphic terms may be instantiated at different types.

Term interpretations are built by connecting the ports of graphical components, that we think of as the nodes in our network. Two connected components can communicate bidirectionally through data tokens, defined by the grammar:

$$e ::= \bullet \mid 0 \mid \mathbf{S} e \mid \mathbf{inl} e \mid \mathbf{inr} e \mid \langle e, e \rangle.$$

We first give a reading of these components as partial maps between data tokens. In the next section we will give a low-level description of their inner workings as communicating abstract machines.

The standard GoI components are given in Figure 1: dereliction (d), promotion (δ) and contraction (c). The components are bidirectional and their behaviour is given by a function mapping the values of a port at a given moment to their values at the next moment. We denote the value on a port which sends/receives no data as \perp . Two well-formedness conditions of GoI nets are that at most one port is not \perp (i.e. a single-token is received at any moment) and $\bar{\perp} = (\perp, \dots, \perp)$ is a fixed-point for any net (i.e. no spontaneous output is created).

Let π_1, π_2 be the first and second projections. Components are connected by functional composition (in both directions) on the shared port, represented graphically as:

$$(t; t')(e, \perp) = t'(\pi_2 \circ t(e, \perp), \perp) \quad p_0 \text{ --- } \boxed{t} \begin{matrix} p_1 \\ p_2 \end{matrix} \text{ --- } \boxed{t'} \text{ --- } p_3$$

$$(t; t')(\perp, e) = t(\perp, \pi_1 \circ t'(\perp, e)).$$

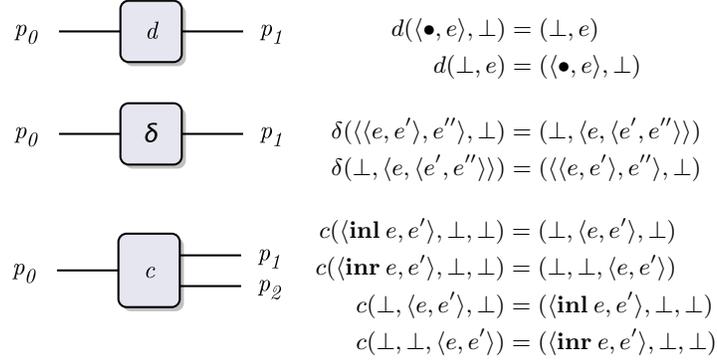
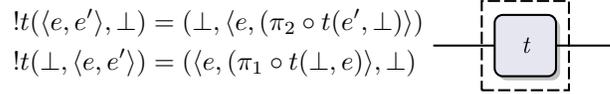
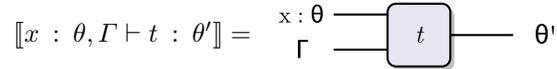


Fig. 1. Dereliction, promotion, and contraction

Exponentials Tokens need to carry both data and ‘routing’ information, but we want the basic components to have no access to the routing information but to act on data only. The role of the exponential functor (!) is to remove this routing information from the enclosed component, pass the data to the component, then restore the routing information. Diagrammatically this is represented as a dotted box around a network, defined formally below:

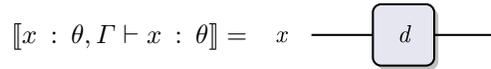


Types as interfaces The interface of a net is determined by the typing judgement of the term it interprets. The **nat** type corresponds to one port; the function type, $\theta \rightarrow \theta'$, induces an interface which is the disjoint union of those for θ, θ' . A typing environment $\Gamma = x_1 : \theta_1, \dots, x_n : \theta_n$ induces an interface which is the disjoint union of the interfaces for each θ_i . A term with typing judgement $\Gamma, x : \theta \vdash t : \theta'$ interface given by the environment on the left and its type on the right. Diagrammatically this is:

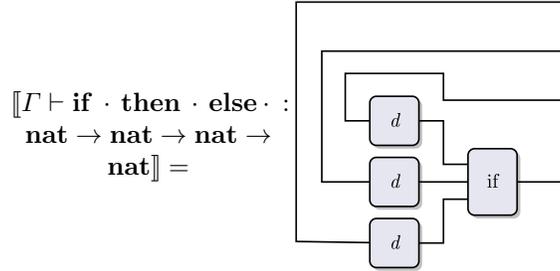


Note that type $\theta \times \theta'$ can be interpreted, as convenient, either as a pair of ports or as a single port sending or receiving pair-tokens. Because of the well-formedness conditions the pairs are always of shape (e, \perp) or (\perp, e) .

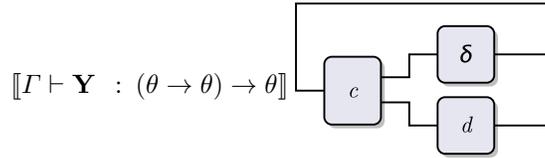
Terms as networks As the variables in the context correspond to the linear ! type, the GoI interpretation requires the use of dereliction:



For the final interpretation of conditionals in the language, derelictions are added as the arguments are of exponential type:



Recursion Recursion is interpreted as a component that is connected to itself as done by Mackie [10].



The abstract token machine interpretation given in this section is known to be sound [9, 10].

Theorem 1 (GoI Soundness). *Let $\vdash t : \mathbf{nat}$ be a closed PCF program (closed ground-type term) and $\llbracket t \rrbracket$ its GoI abstract-token machine representation. If t evaluates to n ($t \Downarrow n$) then $\llbracket t \rrbracket(\bullet) = n$.*

3 The SIC machine

To be able to describe the inner workings of the components and see how they can be compiled to executables we construct an abstract machine, the Stack-Interaction-Control (SIC) machine, which has a small instruction set tailor made for that purpose. The SIC machine works similar to Mackie's [10] but with the important distinction that it also allows sending and receiving messages to and from other machines, to model networked distribution. The machine descriptions and configurations are specified in Figure 2. An initial configuration for a machine description $\langle P, L \rangle$ is given by the function $initial(\langle P, L \rangle) = \langle \mathbf{passive}, \bullet, P, L \rangle$.

3.1 SIC semantics

The semantics of the machines are given as a transition relation in Figure 3. The machine instructions make it possible to manipulate the data token of an active machine, using the stack for state and intermediate results. The last two rules handle sending and receiving messages.

Label	l	
Port	p	
Instruction	$I ::= \mathbf{inl} \mid \mathbf{inr}$	Tags
	$\mathbf{fst} \mid \mathbf{snd}$	Projections
	$\mathbf{unfst} \mid \mathbf{unsnd}$	Reverse projections
	$\mathbf{flip} \mid \mathbf{push} \mid \mathbf{pop}$	Stack operations
	$\mathbf{zero} \mid \mathbf{suc}$	Operations on natural numbers
Code	$C ::= I; C$	Instruction cons
	$\mathbf{jump} \ l$	Jump to label l
	$\mathbf{match} \ l_1 \ l_2$	Conditional jump
	$\mathbf{if} \ l_1 \ l_2$	Conditional jump (nat)
	$\mathbf{send} \ p$	Send on port p
Stack	$S ::= \bullet$	Empty
	$d :: S$	Cons
Port map	$P : p \rightarrow l$	
Label map	$L : l \rightarrow C$	
Machine description	$::= \langle P, L \rangle$	
Machine state	$M_{\text{state}} ::= \mathbf{active} \ C \ d$	
	$\mathbf{passive}$	
Machine configuration	$M ::= \langle M_{\text{state}}, S, P, L \rangle$	

Fig. 2. SIC specification

Now it is possible to define a machine network where many of these machines can operate together. This machine network is described in the style of the Chemical Abstract Machine [2]. A machine network is simply a multiset (\mathcal{M}) of messages (port-data pairs (p, d)) that are “in the air” and a list of machine configurations (M). The transition relation for the network is:

$$\frac{M \rightarrow M'}{\langle \mathcal{M} \mid N_1, M, N_2 \rangle \rightarrow \langle \mathcal{M} \mid N_1, M', N_2 \rangle} \text{SILENT}$$

$$\frac{M \xrightarrow{\text{send } \langle p, d \rangle} M'}{\langle \mathcal{M} \mid N_1, M, N_2 \rangle \rightarrow \langle \mathcal{M} \uplus \{ \langle p, d \rangle \} \mid N_1, M', N_2 \rangle} \text{SEND}$$

$$\frac{M \xrightarrow{\text{send } \langle p, d \rangle} M'}{\langle \mathcal{M} \uplus \{ \langle p, d \rangle \} \mid N_1, M, N_2 \rangle \rightarrow \langle \mathcal{M} \mid N_1, M', N_2 \rangle} \text{RECEIVE}$$

Let $\text{active}(N)$ be a function that returns all the machines in N that are in **active** state.

Proposition 1. *For any sets of machine configurations N and N' and multisets of messages \mathcal{M} and \mathcal{M}' , if $\langle \mathcal{M} \mid N \rangle \rightarrow \langle \mathcal{M}' \mid N' \rangle$, then $|\mathcal{M}'| + |\text{active}(N')| = |\mathcal{M}| + |\text{active}(N)|$.*

Proof. By cases on the machine network step relation and construction of the SIC machine’s step relation.

$\langle \mathbf{active}(\mathbf{inl}; C) d, S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C(\mathbf{inl} d), S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{inr}; C) d, S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C(\mathbf{inr} d), S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{fst}; C) \langle d_1, d_2 \rangle, S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C d_1, d_2 :: S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{snd}; C) \langle d_1, d_2 \rangle, S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C d_2, d_1 :: S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{unfst}; C) d_1, d_2 :: S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C \langle d_1, d_2 \rangle, S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{unsnd}; C) d_2, d_1 :: S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C \langle d_1, d_2 \rangle, S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{flip}; C) d, d_1 :: d_2 :: S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C d, d_2 :: d_1 :: S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{push}; C) d, S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C \bullet, d :: S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{pop}; C) d_1, d_2 :: S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C d_1, S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{zero}; C) d, S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C 0, S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{suc}; C) n, S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} C \mathbf{S} n, S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{jump} l) d, S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} L(l) d, S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{match} l_1 l_2) (\mathbf{inl} d), S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} L(l_1) d, S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{match} l_1 l_2) (\mathbf{inr} d), S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} L(l_2) d, S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{if} l_1 l_2) 0, S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} L(l_2) \bullet, S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{if} l_1 l_2) \mathbf{S} n, S, P, L \rangle$	\rightarrow	$\langle \mathbf{active} L(l_1) \bullet, S, P, L \rangle$
$\langle \mathbf{active}(\mathbf{send} p) d, S, P, L \rangle$	$\xrightarrow{\mathbf{send} \langle p, d \rangle}$	$\langle \mathbf{passive}, S, P, L \rangle$
$\langle \mathbf{passive}, S, P, L \rangle$	$\xrightarrow{\mathbf{send} \langle p, d \rangle, p \in \mathit{dom}(P)}$	$\langle \mathbf{active} L(P(p)) d, S, P, L \rangle$

Fig. 3. SIC step relation

Case silent : In this rule, $\mathcal{M}' = \mathcal{M}$, so $|\mathcal{M}'| = |\mathcal{M}|$. The SIC machine M that takes a step goes from **active** to **active** since all SIC machine step rules that do not send or receive messages are on that form, which also implies that $|\mathit{active}(N)| = |\mathit{active}(N')|$.

Case send : Here $\mathcal{M}' = \mathcal{M} \uplus \{ \langle p, d \rangle \}$ so $|\mathcal{M}'| = |\mathcal{M}| + 1$. The only SIC machine rule that applies here is the send rule, which takes the machine M from state **active** to **passive**. So $|\mathit{active}(N')| = |\mathit{active}(N)| - 1$, and thus $|\mathcal{M}'| + |\mathit{active}(N')| = |\mathcal{M}| + 1 + |\mathit{active}(N)| - 1 = |\mathcal{M}| + |\mathit{active}(N)|$

Case receive : In this case, $\mathcal{M} = \mathcal{M}' \uplus \{ \langle p, d \rangle \}$ so $|\mathcal{M}'| = |\mathcal{M}| - 1$. The only SIC machine rule that applies is the receive rule, which takes the machine M from state **passive** to **active**, meaning that $|\mathit{active}(N')| = |\mathit{active}(N)| + 1$. Thus $|\mathcal{M}'| + |\mathit{active}(N')| = |\mathcal{M}| - 1 + |\mathit{active}(N)| + 1 = |\mathcal{M}| + |\mathit{active}(N)|$.

In particular, this proof means that if we start out with one message and no active machines, there can be at most one active machine at any point in the network's execution – the execution is *single-token*.

3.2 Components as SIC code

In each of the components we take all ports p_x and labels l_x to be distinct. Components are connected by giving an output port the same name as the input port of the component that it is connected to. To emphasise the input/output role of a port we sometimes write them as p_x^i when serving as input and p_x^o when serving as output. The machine descriptions for the different components are described by giving their port mappings P and label mappings L as a tuple.

The following three machines are stateless. They use the stack internally for intermediate results, but ultimately return the stack to the initial empty state.

Dereliction, d , removes the first component of the token tuple when going from left to right, and adds it when going in the other direction:

$$\text{dereliction} = \left\langle \begin{array}{l} p_0^i \mapsto l_0 \\ p_1^i \mapsto l_1 \end{array} \quad , \quad \begin{array}{l} l_0 \mapsto \mathbf{snd}; \mathbf{pop}; \mathbf{send} p_1^o \\ l_1 \mapsto \mathbf{push}; \mathbf{unfst}; \mathbf{send} p_0^o \end{array} \right\rangle$$

Promotion, δ , reassociates the data token to go from $\langle\langle e, e' \rangle, e''\rangle$ to $\langle e, \langle e', e'' \rangle \rangle$ and back:

$$\text{promotion} = \left\langle \begin{array}{l} p_0^i \mapsto l_0 \\ p_1^i \mapsto l_1 \end{array} \quad , \quad \begin{array}{l} l_0 \mapsto \mathbf{fst}; \mathbf{snd}; \mathbf{flip}; \mathbf{unfst}; \mathbf{unsnd}; \mathbf{send} p_1^o \\ l_1 \mapsto \mathbf{snd}; \mathbf{fst}; \mathbf{flip}; \mathbf{unsnd}; \mathbf{unfst}; \mathbf{send} p_0^o \end{array} \right\rangle$$

Contraction, c , uses matching on the first component of the token to choose the right port to send on when going from left to right.

$$\text{contraction} = \left\langle \begin{array}{l} p_0^i \mapsto l_0 \\ p_1^i \mapsto l_1 \\ p_2^i \mapsto l_2 \end{array} \quad , \quad \begin{array}{l} l_0 \mapsto \mathbf{fst}; \mathbf{match} l_3 l_4 \\ l_1 \mapsto \mathbf{fst}; \mathbf{inl}; \mathbf{unfst}; \mathbf{send} p_0^o \\ l_2 \mapsto \mathbf{fst}; \mathbf{inr}; \mathbf{unfst}; \mathbf{send} p_0^o \\ l_3 \mapsto \mathbf{unfst}; \mathbf{send} p_1^o \\ l_4 \mapsto \mathbf{unfst}; \mathbf{send} p_2^o \end{array} \right\rangle$$

Dotted boxes with k inputs are defined as follows: Let $P_{\text{in}}^i = \{p_0^i, \dots, p_{k-1}^i\}$, $P_{\text{out}}^i = \{p_k^i, \dots, p_{2k-1}^i\}$, Let $P_{\text{in}}^o = \{p_0^o, \dots, p_{k-1}^o\}$, $P_{\text{out}}^o = \{p_k^o, \dots, p_{2k-1}^o\}$, and $L_{\text{out}} = \{l_k, \dots, l_{2k-1}\}$. The box for these sets of ports and labels is then the following (note that boxes use the stack for storing their state):

$$\text{box} = \left\langle \begin{array}{l} p_i^i \mapsto l_i \mid p_i^i \in P_{\text{in}}^i \cup P_{\text{out}}^i \end{array} \quad , \quad \begin{array}{l} l_i \mapsto \mathbf{snd}; \mathbf{send} p_{i+k}^o \mid l_i \in L_{\text{in}} \\ l_i \mapsto \mathbf{unsnd}; \mathbf{send} p_{i-k}^o \mid l_i \in L_{\text{out}} \end{array} \right\rangle$$

For a constant, the component's abstract machine is defined as follows:

$$\text{constant} = \left\langle p_0^i \mapsto l_0 \quad , \quad l_0 \mapsto \mathbf{zero}; \mathbf{send} p_0^o \right\rangle$$

The successor machine first asks for its argument, then runs the successor instruction on that.

$$\text{suc} = \left\langle \begin{array}{l} p_0^i \mapsto l_0 \\ p_1^i \mapsto l_1 \end{array} \quad , \quad \begin{array}{l} l_0 \mapsto \mathbf{send} p_1^o \\ l_1 \mapsto \mathbf{suc}; \mathbf{send} p_0^o \end{array} \right\rangle$$

The conditional uses the matching instruction \mathbf{if} to choose the right branch depending on the natural number of the token.

$$\text{if} = \left\langle \begin{array}{l} p_0^i \mapsto l_0 \\ p_1^i \mapsto l_1 \\ p_2^i \mapsto l_2 \\ p_3^i \mapsto l_2 \end{array} \quad , \quad \begin{array}{l} l_0 \mapsto \mathbf{send} p_1^o \\ l_1 \mapsto \mathbf{if} l_{\text{true}} l_{\text{false}} \\ l_{\text{true}} \mapsto \mathbf{send} p_2^o \\ l_{\text{false}} \mapsto \mathbf{send} p_3^o \\ l_2 \mapsto \mathbf{send} p_0^o \end{array} \right\rangle$$

To connect the port p_0 of machine M to p_1 of M' we rename them in the machine definitions to the same port name p . A port must be connected to at most one other port; in this case the net is said to be *deterministic*, as each message will be received by at most one other machine. A port of a machine which is not connected to a port of another machine is said to be a port *of the network*. By $inputs(M)$ ($outputs(M)$) we mean the inputs (outputs) of a machine, whereas by $inputs(N)$ ($outputs(N)$) we mean the inputs (outputs) of a net. Similarly, by $\pi(M)$ ($\pi(N)$) we mean the ports of a machine (network).

Theorem 2 (Soundness). *Let $\vdash t : \mathbf{nat}$ be a closed PCF program (closed ground-type term), $\llbracket t \rrbracket$ its GoI abstract-token machine representation and N its SIC-net implementation. If t evaluates to n ($t \Downarrow n$) then $\llbracket t \rrbracket(\bullet) = n$ and $\langle \{p^i, \bullet\} \mid N \rangle \rightarrow^* \langle \{p^o, n\} \mid N \rangle$.*

4 Combining machines

When writing distributed applications, the location at which a computation is performed is vital. Traditional approaches are sometimes explicit about that, for instance by using message passing. Using our current term interpretation, and thinking of each abstract machine as running on a different node in a network, we get the communication in the network handled automatically, but will have one abstract machine for each (very small) component. The interpretation produces extremely fine grained networks where each node does very little work before passing the token along to another node. It is expected that the communication is one of the most performance critical parts in a distributed network, which is why it would be better if bigger chunks of computations happened on the same node before the token was passed along.

To make this possible, we devise a way to combine the descriptions of two abstract machines in a deterministic network to get a larger abstract machine with the same behaviour as the two original machines. Informally, the way to combine two machines is to remove ports that are used internally between the two machines (if any) and replace sends on those ports with jumps. The algorithm for combining components $M_1 = \langle P_1, L_1 \rangle$ and $M_2 = \langle P_2, L_2 \rangle$ is described formally below.

We use Δ for the symmetric difference of two sets. If $f : A \rightarrow B$ is a function we write as $f \upharpoonright A'$ the restriction of f to the domain $A' \subseteq A$ and we extend it in the obvious way to relations. We use the standard notation $C[s/s']$ to denote the replacing of all occurrences of a string s by s' in C . We write $C[s(x)/s'(x) \mid x \in A]$ to denote the substitution of all strings of shape $s(x)$ by strings of shape $s'(x)$ with x in a list A , defined inductively as

$$\begin{aligned} C[s(x)/s'(x) \mid x \in \emptyset] &= C \\ C[s(x)/s'(x) \mid x \in a :: A] &= (C[s(a)/s'(a)])[s(x)/s'(x) \mid x \in A] \end{aligned}$$

The combination of two machines is defined by keeping the ports which are not shared and by replacing in the code the send operations to shared ports by

jumps to labels given by the port mappings.

$$\begin{aligned} \text{combine}(M_1, M_2) = & \langle (P_1 \cup P_2) \upharpoonright (\pi(M_1) \Delta \pi(M_2)), \\ & (L_1 \cup L_2)[\text{send } p/\text{jump } P(p) \mid p \in \pi(M_1) \cap \pi(M_2)] \rangle. \end{aligned}$$

There are two abuses of notation above. First, the union $P_1 \cup P_2$ above is on functions taken as sets of pairs and it may not result in a proper function. However, the restriction to $\pi(M_1) \Delta \pi(M_2)$ always produces a proper function. Second, $\pi(M_1) \cap \pi(M_2)$ is a set and not a list. However, the result of this substitution is independent of the order in which the elements of this set are taken from any of its possible list representations.

A network is said to be combinable if combining any of its components does not change the overall network behaviour

Definition 1. *A deterministic network of machines $N = M_1, \dots, M_k$ is combinable if whenever*

$$\langle \{ \langle p, d \rangle \} \mid N \rangle \rightarrow^* \langle \{ \langle p', d' \rangle \} \mid N' \rangle$$

for some N' , p in $\text{inputs}(N)$, p' in $\text{outputs}(N)$, then for any N_{combined} obtained from N by replacing M_i, M_j with $\text{combine}(M_i, M_j)$ for some $i \neq j$ we have that

$$\langle \{ \langle p, d \rangle \} \mid N_{\text{combined}} \rangle \rightarrow^* \langle \{ \langle p', d' \rangle \} \mid N'_{\text{combined}} \rangle$$

for some N'_{combined} .

Note that the combined net is not equivalent to the original net (for a suitable notion of equivalence such as bisimilarity) because it will have fewer observable messages being exchanged.

Lemma 1. *If a net N is combinable then N_{combined} is also combinable.*

The set of combinable machines is hard to define exactly, so we would just like to find a sound characterisation of such machines which covers all the basic components we used and their combinations.

Definition 2. *A machine description $M = \langle P, L \rangle$ is stack-neutral if for all stacks S and S' , p in $\text{inputs}(M)$, p' in $\text{outputs}(M)$, if*

$$\langle \{ \langle p, d \rangle \} \mid \langle \text{passive}, S, P, L \rangle \rangle \rightarrow^* \langle \{ \langle p', d' \rangle \} \mid \langle \text{passive}, S', P, L \rangle \rangle$$

then $S = S'$.

Definition 3. *A machine network N of k machines described by port mappings P_i and label mappings L_i is stack-neutral, if for all stacks S_i and S'_i , p in $\text{inputs}(N)$, p' in $\text{outputs}(N)$, if*

$$\begin{aligned} \langle \{ \langle p, d \rangle \} \mid [\langle \text{passive}, S_i, P_i, L_i \rangle \mid i \in \{1, \dots, k\}] \rangle \rightarrow^* \\ \langle \{ \langle p', d' \rangle \} \mid [\langle \text{passive}, S'_i, P_i, L_i \rangle \mid i \in \{1, \dots, k\}] \rangle \end{aligned}$$

then all $S_i = S'_i$.

Note that this definition is more general than having a list of stack-neutral machines, as a stack-neutral network's machines may use the stack for state after it has been exited as long as it's cleared before an output on a network port.

Proposition 2. *If two machine networks N_1 and N_2 (of initially passive machines) are stack-neutral, combinable and N_1, N_2 is deterministic, then N_1, N_2 is stack-neutral and combinable.*

For any SIC-net N let $\text{box}(N)$ be N with an additional box machine M with input ports $\text{outputs}(N)$ and output ports $\text{inputs}(N)$, defined as in Sec. 3.2.

Proposition 3. *If a machine network N is stack-neutral and combinable then $\text{box}(N)$ is stack-neutral and combinable,*

From the following two results it follows by induction on the structure of the generated nets that

Theorem 3. *If $\Gamma \vdash t : \theta$ is a PCF term, $\llbracket t \rrbracket$ its GoI abstract-token machine representation and N the implementation of $\llbracket t \rrbracket$ as a SIC-net then N is combinable.*

With the ability to combine components, we can now exploit the $t @ A$ annotations in the language. They make it possible to specify where a piece of code should be located (A is a node identifier). When this construct is encountered in compilation, the components generated in compiling t are tagged with A (possibly overwriting older tags).

Next, the components with the same tag are combined using the algorithm above and their combined machine placed on the node identified by the tag. This allows the programmer to arbitrarily choose where the compiled representation of a part of a term is placed. Soundness (Thm. 2) along with the freedom to combine nets (Thm. 3) ensures that the resulting network is a correct implementation of any (terminating) PCF program.

5 Compiling PCF

We developed an experimental compiler that compiles to C, using MPI for communication, using SIC abstract machines as an intermediate formalism.² Each machine description in a network is mapped to a C source file in a fairly straight-forward manner, using a function for each machine instruction and global variables for the data token and the stack. An example of a predefined instruction is that for the `flip` instruction:

```
inline void flip() {
    Data d1 = pop_stack();
    Data d2 = pop_stack();
    push_stack(d1);
    push_stack(d2);}

```

² Download from <http://veritygos.org/dpcf>

An abstract machine's label l corresponds to a C function `void l()` whose definition is a list of calls to the predefined machine instruction functions. In this representation, jumps are function calls. All functions are small and not used recursively so can be efficiently inlined.

Each process in MPI has a unique identifier called its *rank*, and messages can also be assigned a *tag*. A port in a SIC machine is uniquely determined by its tag, but also has to be assigned a rank so that the message can be sent to the correct node. This is resolved at compile-time. The main loop for a machine listening on ports corresponding to tags 0 and 1 looks like this:

```
while(1) {
    int port = receive();
    switch (port) {
        case 0: l0(); break;
        case 1: l1(); break;
        default: break;}}}
```

Here `l0` and `l1` are functions corresponding to the labels associated with the ports. The predefined function `receive` calls `MPI_Recv`, which is an MPI function that blocks until a message is received. A process in this state thus corresponds to a machine in **passive** state. Upon receiving a message, the `receive` function de-serialises the message and assigns it to the global data token variable before returning the message's tag. The predefined function for the `send` instruction now has to take two parameters: the destination node's rank and the port's tag:

```
inline void send(int node, int port);
```

The function takes care of serialising the data token sending it to the correct node using `MPI_Send`.

When all machines have been compiled to C, these can in turn be compiled to executables and run on different machines in a network where they use message passing for communication.

6 Conclusion and future work

We have shown a programming language and compilation model for seamless distributed computing, that provides freedom in choosing the location at which a computation takes place with implicitly handled communication. This was achieved by basing the model on the Geometry of Interaction and constructing a way to produce nodes that are more coarse grained than the standard elementary nodes, and showing that this is still correct.

So far semantics of a term are sequential – there is nothing taking place in parallel. The next step will be to investigate how to extend the system with a safe and flexible parallelisation mechanism. A start is to identify internally sequential partitions of the network that can safely be run in parallel. Then the evaluation of a function's arguments can be performed in parallel with the function as long

as they are in different partitions of the network. Another idea is to add a more specific construct for parallelisation, e.g. one for map-reduce.

An important point that has not been discussed in this paper is fault-tolerance: A distributed system needs to be able withstand nodes crashing, becoming unavailable or being added by dynamically reassigning the locus of execution. This is also something that we would like to look into in the future.

References

1. J. Armstrong. A history of Erlang. In *HOPL III: Proceedings of the third ACM SIGPLAN conference on History of programming languages*, pages 6–1–6–26, New York, NY, USA, 2007. ACM.
2. G. Berry and G. Boudol. The chemical abstract machine. *Theor. Comput. Sci.*, 96(1):217–248, 1992.
3. A. Carzaniga, G. P. Picco, and G. Vigna. Designing distributed applications with mobile code paradigms. In W. R. Adrion, A. Fuggetta, R. N. Taylor, and A. I. Wasserman, editors, *ICSE*, pages 22–32. ACM, 1997.
4. J. Epstein, A. P. Black, and S. L. P. Jones. Towards Haskell in the cloud. In K. Claessen, editor, *Haskell*, pages 118–129. ACM, 2011.
5. D. R. Ghica. Geometry of Synthesis: a structured approach to VLSI design. In M. Hofmann and M. Felleisen, editors, *POPL*, pages 363–375. ACM, 2007.
6. D. R. Ghica. Applications of game semantics: From program analysis to hardware synthesis. In *LICS*, pages 17–26. IEEE Computer Society, 2009.
7. D. R. Ghica. Function interface models for hardware compilation. In *Formal Methods and Models for Codesign (MEMOCODE), 2011 9th IEEE/ACM International Conference on*, pages 131–142. IEEE, 2011.
8. J. Girard. Geometry of interaction 1: Interpretation of system F. *Studies in Logic and the Foundations of Mathematics*, 127:221–260, 1989.
9. N. Hoshino. A modified GoI interpretation for a linear functional programming language and its adequacy. In M. Hofmann, editor, *FOSSACS*, volume 6604 of *Lecture Notes in Computer Science*, pages 320–334. Springer, 2011.
10. I. Mackie. The geometry of interaction machine. In *POPL*, pages 198–208, 1995.
11. R. Milner. Functions as processes. *Automata, Languages and Programming*, pages 167–180, 1990.
12. B. C. Pierce and D. N. Turner. Pict: A programming language based on the pi-calculus. In G. Plotkin, C. Stirling, and M. Tofte, editors, *Proof, Language and Interaction: Essays in Honour of Robin Milner*, pages 455–494. MIT Press, 2000.
13. G. D. Plotkin. LCF considered as a programming language. *Theor. Comput. Sci.*, 5(3):223–255, 1977.
14. P. Sewell, P. T. Wojciechowski, and A. Unyapoth. Nomadic Pict: Programming languages, communication infrastructure overlays, and semantics for mobile computation. *ACM Trans. Program. Lang. Syst.*, 32(4), 2010.
15. M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra. *MPI-The Complete Reference, Volume 1: The MPI Core*. MIT Press, Cambridge, MA, USA, 2nd. (revised) edition, 1998.
16. A. M. White. A high-level framework for network-based resource sharing. Augmentation Research Centre, 1975. RFC 707.