

# Encoding $\pi$ -calculus

David Guaspari

January 28, 2010

## 1 Introduction

This note describes an encoding of  $\pi$ -calculus in the process model of [1], hereafter called “our model.” It assumes familiarity with that paper and with  $\pi$ -calculus. We consider the formulation of monadic  $\pi$ -calculus as given in, for example, [2]. The syntax of a  $\pi$ -calculus process term is:

$$P ::= \mathbf{0} \mid \Sigma_{i=1}^n \pi_i.P_i \mid P|Q \mid !P \mid (\nu x)P$$

where  $n > 0$  and a prefix  $\pi_i$  is either the “get” operation  $c(x)$  or the “put” operation  $\bar{c}x$ . We will model the basic  $\pi$ -calculus semantics that does not require communications to be chosen “fairly.” Section 5 discusses this further.

Several basic differences between  $\pi$ -calculus and our model must be negotiated. The notion of location, fundamental in our model, doesn’t exist in  $\pi$ -calculus. In our model, a process sends a message to a known recipient by labeling the message with the recipient’s location, whereas a  $\pi$ -calculus process sends a message to a channel without knowing which other processes might be reading from it. Our model is asynchronous and processes can act only on local knowledge; whereas in  $\pi$ -calculus, communication is synchronous and rendezvous are effected by the environment, using global information.

**Terminology and notation** Unless otherwise indicated, “process” refers to a process in the sense of our model. “Target processes” will encode  $\pi$ -calculus process terms and “bookkeeping processes” will manage and constrain their interactions. A “Target component” is a component, in the sense of our model, that encapsulates a target process. Unexplained notation comes from [1].

## 2 Communication

It seems appealing, at first glance, to encode a  $\pi$ -calculus channel as a process, but that runs into problems because  $\pi$ -calculus semantics can require global decisions to choose which messages to transmit (and which processes will receive them); our model has no built-in magic to make such global decisions.

Instead, we introduce a central bookkeeping process  $Comm$ , at location  $l_{comm}$ , that manages all communication. A target process at location  $l$  says

it’s available to communicate on channel  $b$  by sending a message to  $Comm$  that says, roughly, “location  $l$  is willing to get from [put to] channel  $b$ .”  $Comm$  decides which communications will occur and carries them out by sending a value to the process doing the “get” and an acknowledgement to the process doing the “put”—both of which block until they receive these messages. In general, target processes communicate directly only with bookkeeping processes, not with one another.<sup>1</sup>

$Comm$ ’s decisions about which communications occur must be made nondeterministically, without ruling out any possibilities. In our model all nondeterminism comes from the environment, and processes—in particular,  $Comm$ —must be deterministic, but a simple trick can be used to achieve the desired effect.

$\pi$ -calculus terms that can immediately engage in communications have the form

$$\pi_1.P_1 + \dots + \pi_n.P_n$$

A process at location  $l$  that encodes this term will ask  $Comm$  to perform one of those communications by placing on its external part<sup>2</sup>

$$[l_{comm} : \langle l, [\pi_1, \dots, \pi_n] \rangle]$$

a message of type

$$Reqs = Loc \times (Prefix List)$$

We abuse notation by using expressions of the form  $c(x)$  and  $\bar{c}v$  to denote both  $\pi$ -calculus prefixes and NuPl terms of type  $Prefix$  that represent them. The message includes  $l$  in order to provide  $Comm$  with a return address.

Note that the list of  $\pi_i$  could contain repetitions—e.g., the process could be

$$c(x).P + c(x).Q$$

so that the corresponding  $Reqs$  is  $\langle l, [c(x), c(x)] \rangle$ . Accordingly,  $Comm$ ’s replies will have type

$$Value \times \mathbb{N}$$

A reply will contain both a return value (or acknowledgement) and the index of the request chosen. In this case it would send to location  $l$  a message of the form  $\langle v, 1 \rangle$  or  $\langle v, 2 \rangle$ . The requesting process will use the index to determine whether to continue as  $P$  or as  $Q$ .

In making communication synchronous, we will ensure

**Invariant 1** *The external part of a target component contains at most one element (representing the list of communications for one guarded choice).*

<sup>1</sup>There’s a tiny exception: one target process may send a special **fire** message to another in order to activate it; but such messages do not correspond to communications in the  $\pi$ -calculus execution being simulated.

<sup>2</sup>Strictly speaking, a process does not have an external part; only a component does. The official definitions will not be so loose.

## 2.1 A nondeterministic version of *Comm*

We will first describe *Comm* by allowing it to invoke a nondeterministic choice operator. As noted, our model has no such thing. We then show how to implement that choice operator with a trick that pushes the nondeterminism into the environment.

We define *Comm* with the operator *RecPr*, defined in [1], that encapsulates a recursive definition of a process in terms of an implicit state  $\Sigma$  and a *next* function

$$\text{next} : \text{Msg} \times \Sigma \rightarrow \Sigma \times \text{Ext}$$

Given an input from *Msg* and a state, *next* computes the next state and a list of messages to send.

In this nondeterministic version, the state of *Comm* is simply a finite partial function from locations to lists of prefixes:

$$\text{state}_C : \text{Loc} \xrightarrow{fp} \text{Prefix List}$$

$\text{state}_C[l]$  is defined when *Comm* has a pending request from  $l$ —the request to perform one of the communications in  $\text{state}_C[l]$ .

The only inputs to *Comm* are of type *Reqs*. Consider an input  $\text{loc} : \vec{\pi}$ , where  $\vec{\pi} = [\pi_1, \dots, \pi_n]$ . The definition of  $\text{next}(\text{loc} : \vec{\pi}, \text{state}_C)$  splits into cases, according to whether there exists a match—that is, whether there exist  $i, j \in \mathbb{N}$ ,  $l \in \text{Loc}$  such that  $\pi_i$  and  $\text{state}_C[l][j]$  are matching prefixes (a put and a get on the same channel). Because a  $\pi$ -calculus process can't rendezvous with itself we should also insist that  $l \neq \text{loc}$ ; invariant 2 guarantees that any match found will satisfy that constraint.

### No match exists

$$\text{next}(\text{loc} : \vec{\pi}, \text{state}_C) = \langle \text{state}_C[\text{loc} := \vec{\pi}], \text{nil} \rangle$$

If there is no match, add  $\vec{\pi}$  to  $\text{state}_C$  and send no messages.

**A match exists** If there is a match, nondeterministically choose such a triple,  $\langle i', j', l' \rangle$ .

$$\text{next}(\text{loc} : \vec{\pi}, \text{state}_C) = \langle \text{state}_C[l' := \text{nil}], \text{msgs} \rangle$$

where the list *msgs* consists of the two messages that carry out the rendezvous determined by  $\langle i', j', l' \rangle$ . The pending request  $\text{state}_C[l']$  is deleted.

The definition of *msgs* splits into two cases (depending on who's sending and who's receiving):

$$\text{If } \pi_{i'} = c(x) \text{ and } \text{state}_C[l'][j'] = \bar{c}v, \text{ msgs} = [\text{loc} : \langle v, i' \rangle, l' : \langle \text{ack}, j' \rangle]$$

$$\text{If } \pi_{i'} = \bar{c}v \text{ and } \text{state}_C[l'][j'] = c(x), \text{ msgs} = [\text{loc} : \langle \text{ack}, j' \rangle, l' : \langle v, i' \rangle]$$

As noted, the qualification  $l \neq \text{loc}$  is redundant because of invariant 2.

**Invariant 2** *Comm* cannot receive a message from location  $l$  unless  $\text{state}_C[l] = \text{nil}$ .

## 2.2 Encoding nondeterministic choice

The choice of the triple  $\langle i', j', k' \rangle$  has to range nondeterministically over all possible matches. It suffices to implement the operation “choose nondeterministically from a list” because *Comm* can compute the list of all possible matching triples and request a choice from that. The (trivial) trick is to have the environment make the decision.

Making *one* such choice is easy. Introduce a new process, *Choose*, at a fresh location  $l_C$ . *Choose* communicates only with *Comm*. When *Comm* wants *Choose* to pick an element from the list  $[x_1, \dots, x_n]$  it places the  $n$  separate messages  $l_C : x_1, \dots, l_C : x_n$  on its external part. The action of *Choose* is simply to send back to *Comm* the first message it receives and to ignore subsequent messages.

If we wish *Choose* to make choices repeatedly, it must know when it's seeing a stale message (from a choice that has already been made, and which should be ignored) and when it's seeing a new message (which should be echoed). It suffices for *Comm* to keep an integer counter, *id*, and proceed as follows: when it wants a choice from  $[x_1, \dots, x_n]$  it increments *id* and then place the messages  $l_C : \langle x_1, id \rangle, \dots, l_C : \langle x_n, id \rangle$  on its external part. *Choose* keeps track of the counter values that it has seen; it echoes  $\langle m, k \rangle$  if it has not seen  $k$ , and ignores  $\langle m, k \rangle$  if it has. (An equally good alternative would be to fork off a new version of a one-time-only *Choose* every time a choice is wanted.)

This is not quite good enough, because the decisions that *Comm* makes are not independent of one another. Consider this scenario:  $state_C$  is defined only at location  $l$  and  $state_C[l] = [\bar{c}v]$ . *Comm* receives the message  $\langle l_1, [c(x)] \rangle$ , determines the list of possibilities,  $[\langle 1, 1, l \rangle]$  and asks *Choose* to pick one. Before *Choose* returns a result, *Comm* receives the message  $\langle l_2, [c(y)] \rangle$ , determines *the same* list of possibilities, because  $state_C$  has not changed, and asks *Choose* to choose from it. Then *Choose* will return  $\langle 1, 1, l \rangle$  to *Comm* in response to both requests and *Comm* will execute both rendezvous, sending the value  $v$  both to  $l_1$  and  $l_2$ —which a mistake. We therefore construct *Comm* so that after submitting one request for a choice, it waits for an answer before it submits the next one; and it queues any messages received from target processes while it's waiting. To provide one more illustration of programming in our model, section 6 shows how to do the (straightforward) bookkeeping.

## 3 Encoding $\pi$ -calculus terms and programs

The preceding section has factored the model of communication from the rest of the semantics: to send or messages a target process sends an appropriate vector to *Comm* and waits for a response.

If  $P$  is a  $\pi$ -calculus process term, then its denotation will be a map

$$\llbracket P \rrbracket : Loc \rightarrow Process$$

Intuitively, for any location  $l$ ,  $\llbracket P \rrbracket(l)$  represents the behavior of  $P$  if it is “installed” at location  $l$ .

**Notation:** For readability, we'll commonly write  $\llbracket P \rrbracket(l)$  as  $\llbracket P, l \rrbracket$

We also define  $\mathcal{M}(P) \in \text{System}$ , a system (in the sense of [1]) whose runs simulate the executions of  $P$  as a stand-alone  $\pi$ -calculus program. One component of  $\mathcal{M}(P)$  will be constructed from  $\llbracket P \rrbracket$  and the others will be constructed from bookkeeping processes.

Processes of the form  $\llbracket P, l \rrbracket$  will have certain features in common:

- Each responds to a special message, called *fire*, that activates it.
- Each communicates only with three bookkeeping processes: *Comm*; the *location server LServer*; and the *name server NServer*.

Strictly speaking, *LServer* is a map whose inputs are finite sets of locations and whose outputs are processes. If  $L$  is a finite set of locations, then  $LServer(L)$  is a process that, in response to an input returns a message with a name that differs from any location in  $L$  and any location it has previously returned. *NServer* does the corresponding thing for names. The description of  $\llbracket P \rrbracket$  will treat “get a new location” and “get a new name” as basic operations, though they really involve sending a message to the appropriate server and waiting for a reply—analogous to, but simpler than, the interaction between *Comm* and *Choose*.

**Note:** If locations are modeled as atoms then, for technical reasons, *LServer* is limited to serving up an arbitrarily large but finite number of new locations. However, there is no reason why locations have to be atoms. Alternatively, one may adopt the less elegant solution of simulating the encoding described here at a single location. (Each process has a unique id that simulates its location; a message addressed to a process contain that id in its header; etc.)

$\llbracket P, l \rrbracket$  may create new processes. Our model requires that we know the level of each process being created. Accordingly, we introduce a map  $\Lambda$  from  $\pi$ -calculus terms to  $\mathbb{N}$  and will guarantee invariant 3.

**Invariant 3** For any  $l$ ,  $\llbracket P, l \rrbracket$  is a process of level  $\Lambda(P)$ .

The definition of  $\Lambda$  is simple:

$$\begin{aligned} \Lambda(\mathbf{0}) &= 0 \\ \Lambda(\sum_{i=1}^n \pi_i.P_i) &= \max\{\Lambda(P_1), \dots, \Lambda(P_n)\} \\ \Lambda(P|Q) &= \max\{\Lambda(P), \Lambda(Q)\} + 1 \\ \Lambda(!P) &= \Lambda(P) + 1 \\ \Lambda((\nu x)P) &= \Lambda(P) \end{aligned}$$

**Notation:** This section will contain process definitions in the form “if the input satisfies condition  $\phi$ , the transition is ... elseif  $\psi$ , the transition is ...” They should be understood as saying that, in all other cases, the transition is a no-op. Formally, this is a recursive definition:

$$P = \lambda m. \text{if } \phi(m) \text{ then } \dots \text{elseif } \psi(m) \dots \text{else } \langle P, nil \rangle$$

We will use  $P$  and  $Q$ , with decorations, to vary over  $\pi$ -calculus process terms and boldface  $\mathbf{P}$  and  $\mathbf{Q}$ , decorated, to vary over processes.

**Note:** The definitions in this section make use of primitive operations “get new name” and “get new location.” Accordingly, the only messages a target process receives are `fire` and messages from *Comm*. Explicitly implementing the “new name” and “new location” operations by interaction with *NServer* and *LServer* would add trivial complications.

**The base case** The base case in the definition of  $\llbracket P \rrbracket$  is trivial:

$$\llbracket \mathbf{0}, l \rrbracket = \mathit{null}$$

where *null* is the process that does nothing:

$$\mathit{null} = \lambda m. \langle \mathit{null}, \mathit{nil} \rangle$$

**Guarded choice**

$$\llbracket \pi_1.P_1 + \dots + \pi_n.P_n, l \rrbracket = \lambda m. \text{if } m = \text{fire then } \langle \mathbf{Q}, [l_{\text{comm}} : \langle l, [\pi_1 \dots, \pi_n] \rangle] \rangle$$

where

$$\mathbf{Q} = \lambda m. \text{let } \langle v, i \rangle = m \text{ in } \langle \mathbf{P}_i, [l : \text{fire}] \rangle$$

and the definition of  $\mathbf{P}_i$  depends on  $\pi_i$ :

$$\mathbf{P}_i = \text{if } \pi_i = c(x) \text{ then } \llbracket P_i[x := v], l \rrbracket \text{ else } \llbracket P_i, l \rrbracket$$

That is, when it receives the `fire` message, the process sends its list of possible communications to *Comm* and becomes process  $\mathbf{Q}$ .  $\mathbf{Q}$  waits for a reply from *Comm*, becomes the appropriate continuation process, and then activates that process by sending it the `fire` message.

What is the point of “fire”? A process is purely reactive; it acts only in response to an input. But the components that represent  $\pi$ -calculus terms cannot be purely reactive: they must initiate communications by sending requests to *Comm*. We can make a component (initially) “active” by creating it with a nonempty external part; the environment will deliver those messages. To do this, we could generalize the notion of “process message” to include both a process and an initial external part, generalize the notion of boot process correspondingly, etc. Instead, we construct each target process so that it responds appropriately to the `fire` message, and we have the boot process install the result of firing the target process.

**Digression: a fancy boot process** The obvious strategy for encoding parallel composition or replication is to create new processes at new locations. Thus,  $\llbracket P|Q, l \rrbracket$  will proceed by getting two new locations  $l_1$  and  $l_2$  and sending them, respectively, messages containing processes that encode  $P$  and  $Q$ . The effect of such “process messages” depends on the boot process to be applied. If we

use the default boot process the newly installed processes will not be “active” because they’re waiting for the fire message. Thus we must both install them and send the activating message. Without adding assumptions about the environment it will not suffice to have  $\llbracket P|Q, l \rrbracket$  send both the processes and the fire messages—for there is no guarantee that either process will be installed before the fire message arrives.

The solution we choose is to introduce a fancier boot process that both installs a process and activates it by sending it the message fire. It’s a simple modification of the boot process defined in [1]. (To improve readability the definitions write the constructor  $pmsg$  as  $\mu$ .)

$$boot^+(n) = RecPr(next, \perp)$$

where

$$\begin{aligned} next(\perp) &= \lambda m. \quad \text{if } m = pmsg(k, \mathbf{Q}) \text{ and } k < n \\ &\quad \text{then let } \langle \mathbf{P}, e \rangle = \mathbf{Q}(\text{fire}) \text{ in } \langle \text{inl}(\mathbf{P}), e \rangle \\ &\quad \text{else } \langle \perp, nil \rangle \\ next(\text{inl}(\mathbf{Q})) &= \lambda m. \quad \text{let } \langle \mathbf{Q}', e' \rangle = \mathbf{Q}(m) \text{ in } \langle \text{inl}(\mathbf{Q}'), e' \rangle \end{aligned}$$

**Parallel composition** The process encoding  $P|Q$  will create processes for  $P$  and  $Q$  at new locations, activate them, and then terminate. This definition clearly preserves invariant 3.

$$\begin{aligned} \llbracket P|Q, l \rrbracket &= \lambda m. \text{ if } m = \text{fire then} \\ &\quad \text{let } l_1, l_2 \text{ be new locations in} \\ &\quad \langle null, [l_1 : pmsg(\Lambda(P), \llbracket P, l_1 \rrbracket), l_2 : pmsg(\Lambda(Q), \llbracket Q, l_2 \rrbracket)] \rangle \end{aligned}$$

One subtlety: we want  $l_1$  and  $l_2$  to be, in particular, different from  $l$ . We will guarantee that by insuring that the locations of all the target processes are either in the finite set with which  $LServer$  is seeded or are themselves returned from a call on  $LServer$ .

**Replication** Replication is just a variant of parallel composition. The process  $\llbracket !P, l \rrbracket$  will get a new location  $l_1$ , install  $\llbracket P, l_1 \rrbracket$  there, and resend the fire message to itself (enabling it, when the message is delivered, to install further models of  $P$  at other locations). The definition clearly preserves invariant 3.

$$\begin{aligned} \llbracket !P, l \rrbracket &= \lambda m. \text{ if } m = \text{fire then} \\ &\quad \text{let } l_1 \text{ be a new location in} \\ &\quad \langle \llbracket !P, l \rrbracket, [l : \text{fire}, l_1 : pmsg(\Lambda(P), \llbracket P, l_1 \rrbracket)] \rangle \end{aligned}$$

**The  $\nu$  operator** The  $\nu$  operator is encoded straightforwardly by calling on  $NServer$ .

$$\llbracket (\nu x)P, l \rrbracket = \lambda m. \text{ if } m = \text{fire then let } n \text{ be a new name in } \langle \llbracket P[x := n], l \rrbracket, nil \rangle$$

## 4 Simulating a $\pi$ -calculus program

In our model, what executes is not a process but a system. Thus, if  $P$  is a  $\pi$ -calculus process term, we define the system  $\mathcal{M}(P)$  to represent execution of  $P$  as a stand-alone program, its elements interacting only with one another. We may assume that all the bound names of  $P$  are distinct.

Let  $N$  be the set of all names occurring in  $P$  and choose five distinct locations:  $l_{comm}$ ,  $l_C$ ,  $l_N$ ,  $l_L$ , and  $l$ .  $\mathcal{M}(P)$  will consist of the following five components, together with the boot process  $boot^+(\Lambda(P) + 1)$ .

- $\langle l_{comm}, Comm, nil \rangle$
- $\langle l_C, Choose, nil \rangle$
- $\langle l_N, NServer(N), nil \rangle$
- $\langle l_L, LServer(\{l_C, l_{comm}, l_N, l_L, l_0\}), nil \rangle$
- $\langle l, \llbracket P, l \rrbracket(\text{fire}) \rangle$

As discussed in the next section, our semantic model of  $P$  is the collection of *fair* runs of  $\mathcal{M}(P)$ .

## 5 Fairness, etc.

A correctness proof for this encoding would show that the collection of fair runs of  $\mathcal{M}(P)$  simulates execution of  $P$ . Instead of proof, we offer hand-waving. However, one possible subtlety should be addressed—why we insist on the fair runs of  $\mathcal{M}(P)$ , those in which every message placed in an external part is eventually delivered. Without that assumption, the simulation is not guaranteed to make progress. With it—the hands wave— $\mathcal{M}(P)$  simulates the “standard”  $\pi$ -calculus semantics of  $P$ , not any kind of “fair” semantics.

To see why we need guaranteed message delivery, consider  $P = !c(x).\mathbf{0}! \bar{c}v.\mathbf{0}$ . Execution of the  $\pi$ -calculus term  $P$  produces only one trace, an infinite sequence that repeatedly sends  $v$  on channel  $c$ . However, there is an unfair run of  $\mathcal{M}(P)$  that does nothing but create new copies of  $Q = ! \bar{c}v.\mathbf{0}$  and never performs any communication at all: At each step the environment has the option to deliver a fire message that forks off another copy of  $Q$ , and chooses to do so instead of delivering any “get” or “put” message to  $Comm$ .

To see why fair runs of  $\mathcal{M}(R)$  need not correspond to “fair” executions of  $R$ , consider  $R = !c(x).P_1!c(x).P_2! \bar{c}v.Q$ . We are guaranteed that all requests for communications will be delivered to  $Comm$  but not guaranteed that  $Comm$  responds to every request it receives. For example, there are fair runs of  $\mathcal{M}(R)$  in which  $Comm$  is repeatedly asked to choose between replying to the “get” message that continues as  $P_1$  and the “get” message that continues as  $P_2$ —and always chooses the first of these.

## 6 Appendix: *Comm*

Here is a definition of a deterministic version of *Comm* that makes choices by communicating with *Choose*. A state is now a tuple  $\langle reqArray, q, waiting \rangle$  where

$$\begin{aligned} reqArray & : Loc \xrightarrow{fp} PrefixList \\ q & : Queue(Reqs) \\ waiting & : \mathbb{B} \end{aligned}$$

The variable *reqArray* is what used to be called *state<sub>C</sub>*, and plays the same role; *q* keeps track of the communication requests that have been received by *Comm* but not yet examined; and *waiting* is true when *Comm* is waiting for an answer from *Choose*.

The construction will guarantee invariant 4:

**Invariant 4** *In any reachable state  $\langle reqArray, q, waiting \rangle$ :*

- *If  $l_1 : p_1$  and  $l_2 : p_2$  are distinct entries in  $q$ ,  $l_1 \neq l_2$ .*
- *If  $l : p$  is an entry in  $q$ ,  $reqArray[l] = nil$ .*

Let  $nil_Q$  be the empty queue and define

$$\begin{aligned} add(nil_Q, reqArray) & = reqArray \\ add(enqueue(l : \vec{\pi}, q), reqArray) & = add(q, reqArray[l := \vec{\pi}]) \end{aligned}$$

Note that if  $q' \subseteq q$  and  $q$  satisfies invariant 4, then so does  $q'$ . If  $q'$  and  $reqArray$  satisfy invariant 4 the result of  $add(q, reqArray)$  is to add all the requests in  $q'$  to  $reqArray$ .

*Comm* responds to three kinds of input messages: *Reqs* from target processes, triples from *Choose*, and *continue* messages that it sends to itself. So if state  $\sigma = \langle reqArray, q, waiting \rangle$ , the definition of  $next(m, \sigma)$  splits into three cases:

**Case  $m = reqs \in Reqs$**

$$next(m, \sigma) = (\langle reqArray, enqueue(reqs, q), waiting \rangle, [l_{comm} : continue])$$

When it receives *reqs* from a target task, *Comm* enqueues *reqs* on *q* and sends itself a *continue* message.

**Case  $m = \langle i', j', l' \rangle$**

$$next(m, \sigma) = (\langle reqArray[l' := nil], q, false \rangle, [l_{comm} : continue]@msgs)$$

where *msgs* is as defined as in section 2.1.

When it receives a triple from *Choose*, *Comm* sends *msgs*, the value and acknowledgment determined by that triple, and deletes from *reqArray* the request that has been matched; it is no longer *waiting*, and sends itself a *continue* message (so that it will process unexamined messages in *q*).

**Case  $m = \text{continue}$**

There are two subcases.

If  $\text{waiting} = \text{true}$ ,

$$\text{next}(m, \sigma) = (\sigma, \text{nil})$$

A `continue` message that arrives while *Comm* is awaiting a reply from *Choice* is ignored.

If  $\text{waiting} = \text{false}$ , let  $q = q_1 \oplus q_2$ , where

- no entry in  $q_2$  matches any other entry in  $q_2$  or any entry in  $\text{reqArray}$
- either  $q_1 = \text{nil}_Q$  or  $\text{head}(q_1)$  has a match in  $\text{add}(q_2, \text{reqArray})$ .

It's easy to see that there is a unique pair  $(q_1, q_2)$  satisfying these conditions. Split once again, on the value of  $q_1$ .

If  $q_1 = \text{nil}_Q$

$$\text{next}(m, \sigma) = (\langle \text{add}(q, \text{reqArray}), \text{nil}_Q, \text{false} \rangle, \text{nil})$$

In this case, there are no matches among the unexamined requests, all are added to the pending requests, emptying the queue. *Comm* is not waiting and sends no messages.

If  $q_1 \neq \text{nil}_Q$

$$\text{next}(m, \sigma) = (\langle \text{add}(q_2, \text{reqArray}), \text{tail}(q_1), \text{true} \rangle, [l_{\text{comm}} : \text{continue}, l_C : \text{matches}])$$

where *matches* is the list of triples representing all possible matches of elements of  $\text{head}(q_2)$  with elements of  $\text{add}(q_2, \text{reqArray})$ .

In this case, no matches can be found in  $q_2$ , so all of its elements are added to the pending requests in  $\text{reqArray}$ ;  $\text{head}(q_1)$  can be matched to something in  $\text{add}(q_2, \text{reqArray})$ , so *Choice* is asked to pick one and *Comm* is now waiting; the unexamined requests are those remaining in  $q_2$ ; finally, since  $q_2$  may not be empty, *Comm* tells itself to `continue`.

## References

- [1] Mark Bickford and Robert Constable. Generating Event Logics with Higher-Order Processes as Realizers. *To appear*.
- [2] Robin Milner. The Polyadic  $\pi$ -Calculus: A Tutorial. LFCS report ECS-LFCS-91-180. <http://www.lfcs.inf.ed.ac.uk/reports/91/ECS-LFCS-91-180/>