

# Integrating Multithreading into the Spineless Tagless G-machine

Manuel M.T. Chakravarty\*  
Software Engineering Research Group  
Technical University Berlin

## Abstract

To reduce the adverse effects of long latency communication operations in distributed implementations of the Spineless Tagless G-machine (STGM), a variant of the original abstract machine that contains explicit support for multithreading is introduced. In particular, source-to-source transformations can be used on the level of the abstract machine code to foster the tolerance to long latency communication.

The changes to the original STG-language include a separation of demand from case selection together with the introduction of a new construct that provides an abstract notion of thread boundaries and thread synchronization.

## 1 Introduction

A static mapping of the components of a parallel program to the physical processor elements of a parallel computer such that these components communicate in a regular fashion is only possible for restricted programs or programs that contain explicit layout information, e.g., [Kel89, DP93]. In contrast, distributed implementations of general lazy functional programs issue demands for remote data in a dynamic and unpredictable fashion; the resulting long latency communication operations have adverse effects on distributed implementations of abstract machines, such as the Spineless Tagless G-machine (STGM) [PS88, Pey92] with its implementations for GRIP [PCS89] and GUM [THJ<sup>+</sup>96].

The subject of this paper is a variant of the STGM that is designed to reduce the impact of long latency communication on the execution time; to this end, it exploits the inherent fine-grain parallelism contained in functional programs by employing multithreading and supporting stateless threads. A special feature of the new abstract machine, called STGM<sub>MT</sub>, is that source-to-source transformations can be used on the level of the abstract machine code to foster the tolerance to long latency communication.

Section 2 introduces the basic ideas underlying the development of the STGM<sub>MT</sub>. Section 3 describes the changes to the original machine language of the STGM, and Section 4 demonstrates the use of the new constructs. Then in Section 5, the operational semantics of the new machine language is presented. Related work is discussed in Section 6, and Section 7 contains the conclusions.

---

\*E-mail: [chak@cs.tu-berlin.de](mailto:chak@cs.tu-berlin.de), URL: <http://www.cs.tu-berlin.de/~chak/>

## 2 The Use of Multithreading

We assume a distributed implementation of the STGM where on each of the many processor elements (PEs) a number of *tasks* are running interleaved. When the active task of one of these processors attempts the evaluation of the expression  $(a + b) * (c + d)$ , it may start by accessing the value of the variable  $a$ . If the value of  $a$  is stored in the memory of another PE, a remote access is triggered (by entering a FETCHME closure; cf. [Pey92]). When such a long latency communication operation stalls the progress of a task, it is common practice to suspend that task and, instead, to execute some other, runnable task. Such a task switch includes an expensive context switch where the current values of some processor registers have to be saved and other values have to be loaded; furthermore, the data in the processor cache is probably useless for the computation performed by the newly selected task.

Crucial for the techniques developed in this paper is, returning to the example, the observation that the evaluation of the subexpression  $c + d$  is (i) independent from the result of the remote access induced by  $a$  and (ii) can be achieved with less context switching costs than switching to an arbitrary runnable task. The context needed to evaluate  $c + d$  is similar to the context needed to evaluate  $a + b$ , in particular, both computations are part of the same function activation.

Across function boundaries, a similar observation can be made. When the result of a function activation depends on a long latency operation and there are no more computations in this function activation that are independent from the long latency operation, then it is usually less costly to switch to the execution of some independent computation in the *calling* function than to activate a completely unrelated task—consider the processor cache.

Overall, when the delay of a long latency operation has to be covered by some independent computation, take a computation whose context is as close to the current context as possible. Such a computation can be chosen by the heuristic that computations that would have been executed shortly after the current computation are closely related in their context—an assumption justified by the locality that computations usually exhibit.

The fine-grain parallelism inherent in functional programs can be used to mechanically partition a program into *threads* realizing independent computations. Within a function activation, long latency operations can then be covered by switching to any of the ready threads. The abstract machine language described below provides a construct that explicitly displays independent computations on the level of the abstract machine language. These independent computations can be implemented by *stateless* threads [CGSE93, EAL93]. The crucial feature of stateless threads is that they never suspend, i.e., they start only when all resource requirements are satisfied and execute to completion without any further need for synchronization—in essence, they represent the smallest unit of non-synchronizing computation. There is evidence that the use of stateless threads minimizes the thread switching time while simultaneously allowing to exploit the properties of the memory hierarchy [CGSE93, H<sup>+</sup>95].

### 3 The Abstract Machine Language

Starting from the STG-language as described in [Pey92], two principal changes are required to integrate support for multithreading: demand must be separated from case selection, and an abstract notion of thread boundaries and thread synchronization is needed. In the following, a third modification will also be applied: an abstract notion of distribution is added in order to be able to observe the effects of multithreading on the abstract level of the STG-language. The variant of the STG-language that is defined in this paper is called the STG<sub>MT</sub>-language.

To focus the following presentation on issues relevant to multithreading, support for built-in data types, such as integers and floating point numbers, is omitted—they can be handled in a similar way as in the original STGM, namely by an explicit treatment of unboxed values. Furthermore, the explicit top-level, which contains the global definitions, is omitted—here also, the mechanisms of the STGM can still be applied in the STGM<sub>MT</sub>.

#### 3.1 The Grammar

A definition of the grammar of the STG<sub>MT</sub>-language can be found in Figure 1.

In comparison to the original STG-language, note the addition of **letrec** and **letpar**, and the fact that in a case expression the keywords **case** and **of** enclose an identifier and not an arbitrary expression. An intuition of the behaviour of the added or changed constructs is provided in the following subsections.

#### 3.2 An Abstract Notion of Threads

In the original STGM, there is only a single kind of bindings. It associates variable names with lambda forms. In the STGM<sub>MT</sub> these bindings are called *function bindings* and are produced by the nonterminal *fbind*. In addition, *value bindings*—nonterminal *vbind*—are introduced in the STGM<sub>MT</sub>.

Value bindings occur only in the **letpar** construct, which has the following general form:

$$\mathbf{letpar} \ v_1\# = e_1; \dots; v_n\# = e_n \ \mathbf{in} \ e$$

In contrast to the **letrec** construct (cf. [Pey92]), no closures are created, but the expressions  $e_1$  to  $e_n$  are evaluated, and only after all results are assigned to the  $v_i\#$ , evaluation proceeds with  $e$ . Furthermore, the  $v_i\#$  may not occur free in  $e_1$  to  $e_n$ .

The last restriction guarantees the independence that we required in Section 2 for computations that may be used to cover long latency operations. More precisely, it allows to evaluate  $e_1$  to  $e_n$  in an arbitrary order without any need to synchronize on the  $v_i\#$ . Should the evaluation of any  $e_i$  suspend due to a remote access, then it is still possible to continue the computation locally with any  $e_j$  where  $j \neq i$ . In short, **letpar** allows to express the independence

$exp$	$\rightarrow$	<b>letrec</b> $fbinds$ <b>in</b> $exp$	(mutually recursive closures)
		<b>letrem</b> $fbind$ <b>in</b> $exp$	(remote closure)
		<b>letpar</b> $vbinds$ <b>in</b> $exp$	(parallel demands)
		<b>case</b> $uvid$ <b>of</b> $alts$	(selection)
		$vid$ $args$	(closure evaluation)
		$cid$ $env$	(return constructor)
		$uvid$	(return unboxed value)
$fbinds$	$\rightarrow$	$fbind_1; \dots; fbind_n$	( $n \geq 1$ )
$fbind$	$\rightarrow$	$vid = env \ \backslash \pi \ args \ \rightarrow \ exp$	(function binding)
$\pi$	$\rightarrow$	<b>u</b>	(updatable)
		<b>n</b>	(not updatable)
$vbinds$	$\rightarrow$	$vbind_1; \dots; vbind_n$	( $n \geq 1$ )
$vbind$	$\rightarrow$	$uvid = exp$	(value binding)
$alts$	$\rightarrow$	$alt_1; \dots; alt_n; dft$	( $n \geq 1$ )
$alt$	$\rightarrow$	$cid \ args \ \rightarrow \ exp$	(case alternative)
$dft$	$\rightarrow$	<b>default</b> $\rightarrow \ exp$	(case default)
$args$	$\rightarrow$	$\{arg_1, \dots, arg_n\}$	(argument variables, $n \geq 0$ )
$arg$	$\rightarrow$	$vid$	(boxed value)
		$uvid$	(unboxed value)
$env$	$\rightarrow$	$\{vid_1, \dots, vid_n\}$	(environment variables, $n \geq 0$ )
$uvid$	$\rightarrow$	$vid\#$	(variable for unboxed values)
$vid$	$\rightarrow$	<b>lowercase identifier</b>	(variable)
$cid$	$\rightarrow$	<b>uppercase identifier</b>	(data constructor)

Figure 1: The grammar of the STG<sub>MT</sub>-language.

of local computations on the level of the abstract machine language. Furthermore, the fact that the evaluation of the body expression  $e$  must wait for the delivery of all  $v_i\#$  can be seen as an abstract form of synchronization barrier.

The hash marks (**#**) behind the  $v_i$  indicate that the  $v_i\#$  store unboxed values. The treatment of unboxed values in the STG<sub>MT</sub> is related to, but not identical to the use of unboxed values in the original STGM. In particular in the original STGM, which follows [JL91], all types of boxed and unboxed values have to be explicitly introduced, but in the STG<sub>MT</sub>, there is implicitly a corresponding unboxed value for each boxed value. The coercion from boxed to unboxed and from unboxed to boxed types is made explicitly by value bindings and function bindings, respectively. For example, the expression

```
letpar v# = w in ...
```

binds to  $v\#$  the unboxed value associated with the boxed value stored in  $w$ .

Conversely,

```
letrec w = {} \n {} -> v# in ...
```

boxes the unboxed value contained in  $v\#$ . The deviation from the technique used in the STGM becomes necessary due to the fact that demand for evaluation, in the  $\text{STGM}_{\text{MT}}$ , is issued when an expression occurs in the right-hand side of a value bindings while, in the STGM, the value of an expression is only demanded when it is scrutinized by a **case** expression.

Note that unboxed variables are not allowed to be in the list of free variables of a function binding, i.e., only boxed values can be stored in the environment of closures, and that it is forbidden to use them as arguments to constructors. These restrictions can be relaxed, but here they are enforced here to make the presentation simpler.

Furthermore, the expressions appearing as right-hand sides of value bindings must not be of functional type, i.e., must not be of type  $\alpha \rightarrow \beta$ . This restriction corresponds to the restriction of the original STGM that says that **case** expressions must not inspect values of functional type.

### 3.3 Selection Without Demand

In the original STGM, **case** expressions play two roles: first, they demand the evaluation of the scrutinized expression, i.e., the expression between the keywords **case** and **of**; second, they select one of several alternatives by matching the value of the scrutinized expression against the patterns of the alternatives.

It was already mentioned that value bindings issue demands for evaluation in the  $\text{STGM}_{\text{MT}}$ , and hence, the single purpose of **case** expressions is pattern matching. Overall, we have the following correspondence:

$$\begin{array}{ll}
 \text{Original STGM} & \text{STGM}_{\text{MT}} \\
 \text{case } e \text{ of } alt_1; \dots; alt_n; dft \approx & \text{letpar } v\# = e \\
 & \text{in case } v\# \text{ of } alt_1; \dots; alt_n; dft
 \end{array} \tag{1}$$

### 3.4 An Abstract Notion of Distribution

In the original presentation of the STGM [Pey92], the potential distribution of the heap of the abstract machine over multiple processing elements is left implicit. To make the need for long latency operations explicit, we expose the potential for distribution in the  $\text{STGM}_{\text{MT}}$ .

To this end, the concept of a *machine instance* is introduced. Each machine instance has a *local* heap and is able to evaluate closures in its local heap independent from the other instances. When the local evaluation depends on a closure stored within another machine instance—we call this a *remote* closure—a long latency operation is triggered.<sup>1</sup> On the level of the abstract

---

<sup>1</sup>Note that already in the original STGM, closures are the only kind of structure in the heap.

machine code, no assumptions are made about the number of machine instances available.

The **letrem** construct specifies those closures that *may* be allocated remotely. In contrast, the closures associated with the function bindings of a **letrec** are bound to be allocated locally. To simplify matters, there may only be one binding in a **letrem** and it must not be recursive—recursion can be introduced by using a **letrec** in the right-hand side of the single function binding of the **letrem**. Furthermore, the binding of a **letrem** must not have any arguments, i.e., it has to represent a nullary function.

Overall, the  $\text{STG}_{\text{MT}}$ -language allows, even requires, to explicitly specify the *partitioning* of a parallel program, but it abstracts over its *mapping* (cf. [Fos95] for a definition of these notions).

## 4 Using the New Constructs

In summary, the  $\text{STG}_{\text{MT}}$ -language modifies the original STG-language in three ways: it introduces an explicit, but abstract, notion of (i) local, independent computations (**letpar**) and (ii) closures that may be allocated on a remote instance (**letrem**); and (iii) demand and selection are separated. These modifications lead to a number of interesting properties that are discussed in the following.

### 4.1 Generating code for the $\text{STGM}_{\text{MT}}$

The translation of a functional program into the  $\text{STG}_{\text{MT}}$ -language corresponds closely to the generation of code for the original STGM. The main difference is that we have to observe the correspondence stated in Equation (1) and **case** expressions that are just used for unboxing require no **case** in the  $\text{STGM}_{\text{MT}}$ , but only a **letpar**. In contrast to **letpar** expressions, **letrem** constructs are not expected to be generated automatically; instead they are generated from explicit annotations, i.e., the programmer decides which computations are coarse-grained enough to be worth the shipping to another processor element.

### 4.2 Covering Long Latency Operations

Following the stated scheme for code generation, only **letpar** constructs containing a single value binding are generated. Such code does not exhibit any tolerance to long latency operations. An important characteristic of the  $\text{STG}_{\text{MT}}$ -language is that simple source-to-source transformations can be used to increase this tolerance. In particular, we can apply the following transformation rule:

$$\begin{array}{l} \text{letpar } v_1\# = e_1 \\ \text{in letpar } v_2\# = e_2 \text{ in } e_3 \end{array} \quad \Rightarrow \quad \begin{array}{l} \text{letpar} \\ v_1\# = e_1 \\ v_2\# = e_2 \\ \text{in } e_3 \end{array} \quad (2)$$

when  $v_1$  is not free in  $e_2$ .

In case of the example from Section 2,  $(a + b) * (c + d)$ , the transformation rule (2) has a dramatic effect on the corresponding STGM<sub>MT</sub>-code:

<pre> letpar av# = a [] in letpar bv# = b [] in letpar ab# = add# [av#, bv#] in letpar cv# = c [] in letpar dv# = d [] in letpar cd# = add# [cv#, dv#] in mul# [ab#, cd#] </pre>	$\Rightarrow$	<pre> letpar   ab# = letpar     av# = a [];     bv# = b []   in     add# [av#, bv#]; cd# = letpar   cv# = c [];   dv# = d [] in   add# [cv#, dv#]; in mul# [ab#, cd#] </pre>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

The code to the left is similar to the **case** cascade used to represent this computation in the original STGM. The transformed code, to the right, explicitly represents the independence of those subcomputations that can be used to cover long latency operations. In particular, when demand **a** triggers a remote access, the demand of **b** as well as the demands to **c** and **d** together with the evaluation of **add# [cv#, dv#]** can be done while waiting for the value **av#**. In the worst case, when all data is remote, at least the access to **a**, **b**, **c**, and **d** is overlapped. In essence, the code to the right is a textual representation of the partial ordering induced on the code by the data dependencies.

Overall the separation of demand and selection allows to move demand, i.e., value bindings, outwards in order to collect multiple value bindings in a single **letpar** as in the following code fragment (the  $y_i$  are not free in  $e_2$ ):

<pre> letpar x# = e<sub>1</sub> in case x# of   C {y<sub>1</sub>, ..., y<sub>n</sub>} -&gt;     letpar z# = e<sub>2</sub>     in case z# of ... </pre>	$\Rightarrow$	<pre> letpar   x# = e<sub>1</sub>   z# = e<sub>2</sub> in case x# of   C {y<sub>1</sub>, ..., y<sub>n</sub>} -&gt; case z# of ... </pre>
--------------------------------------------------------------------------------------------------------------------------------------------------------	---------------	------------------------------------------------------------------------------------------------------------------------------------------

Apart from data dependencies, the outwards movement of value bindings is stopped by **case** expressions with more than one alternative; moving a value binding over a **case** with multiple alternatives can change the termination behaviour of a program.

In principle, function boundaries also stop the outwards movement, but this should, for the following reason, not become a problem in practice: either the function is very simple, then it can be inlined; or it is complex and, then, it usually contains a **case** with multiple alternatives, which hinders the outward movement anyway.

### 4.3 Distribution

The following program fragment displays the essentials of a parallel program exploiting pipelined parallelism:

```

consume      :: [Value] -> Result
consume []   = <trivial result>
consume (x:xs) = let r = consume xs
                  in
                  <make result from x and r>

main = let stream = <producer of some list of Values>
        in
        stream 'par' consume stream   — consume in parallel

```

The `par` meta-function indicates that its two arguments may be evaluated (in parallel) on different machine instances; it is compiled into a `letrem` construct in the following `STGMMT`-code:

```

consume = {} \n {l} ->
  letpar lv# = l {} in case lv# of
    Nil {}      -> <trivial result>
    Cons {x, xs} -> letpar
                      xv# = x {}           — independent
                      rv# = consume {xs}  — computations
                    in
                    <make result from xv# and rv#>

main = {} \u {} ->
  letrem                — potential remote allocation
  stream = {...} \u {} -> <producer of some list of Values>
  in
  consume {stream}

```

The above code for `consume` is already transformed; the original code would place the expression `consume {xs}` into the function binding of a separate `letrec`. But the immediate following occurrence of the bound variable in the right-hand side of a value binding allows the transformation into the shown code.

If there are machine instances that need additional work, then closures created with `letrem` can be shipped to those instances; otherwise they can also be allocated and evaluated locally—the latter case corresponds to the idea of the absorption of previously sparked children [HMP94]. In the above example, let us assume that `stream` is allocated remotely. Then, the value of `stream` and `consume {stream}` are evaluated on different instances, in parallel—any closures created with `letrec` while evaluating `stream` are also allocated and, thus, evaluated on the remote instance. This implies that the access to `x` in the body of `consume` triggers a remote access, which is, at least partially, covered by the recursive call to `consume` (in the same `letpar`).

## 5 The Meaning

To formalize the operational semantics of the STG<sub>MT</sub>-language, a transition system is presented in this section; it is derived from the system in [Pey92]. It makes the effects of multithreading on the abstract level of the STG<sub>MT</sub>-language explicit. The notation used in this section is similar to that used in [Pey92]; details are provided in Appendix A.

### 5.1 Machine Configurations

A *machine configuration* is represented by a mapping  $\mathcal{I}$  from names to machine instances. Each instance consists of several components, including a code component, a task pool, an argument stack, a return stack, and a heap. A detailed description of these data structures is provided in Appendix B.

The machine instances in a configuration share a global name space, but the computations within one instance  $i$  may only access the components of  $i$ . When  $i$  needs to access a closure, named  $o$ , that is located in the heap of another instance  $j$ , it has to request  $j$  to evaluate the closure  $o$  and to return the WHNF of  $o$  back to  $i$ . This operation is the single form of long latency operation in the STG<sub>MT</sub>.

In the transition rules, we assume an unbound number of instances, and each closure allocated by a `letrem` is created on a not-yet-used instance. This exposes the maximal parallelism of the program. In a concrete implementation, the scheme outlined Section 4.3 is used, i.e., closures are only distributed upon request from processing elements with an insufficient amount of work.

### 5.2 The Transition Rules

The following transition rules affect either one or two instances at a time. To get a parallel and not only an interleaving semantics, we define a *parallel transition step* to be a set of applications of transition rules such that this set contains at least a single element and no instance is affected by more than one transition rule. A transition rule is said to affect an instance if this instance occurs in the rules pre- or postcondition.

#### 5.2.1 The Initial State

The initial machine state used to evaluate an expression  $e$  is the following:

inst.	current task	task pool	arg. frames	arg. stack	ret. stack	heap	dissem. map
$[i \mapsto$	$Eval\ e\ []$	$\emptyset$	$[]$	$[]$	$[\langle d, Next, [] \rangle]$	$[]$	$[]$

The configuration consists of a single instance named  $i$ . Its current task is to evaluate  $e$  within an empty environment. The task pool is empty (i.e., there is no further work), just as the frame map and argument stack. The single continuation on the return stack indicates that the result of  $e$  has to be delivered via the (non-existing) slot  $d$  of the dissemination map.

The machine terminates when it attempts to distribute some value over the dissemination slot  $d$ ; this is the value computed for  $e$ .

Intuitively, the roles of the components of an instance are as follows. The task pool contains the waiting (for the completion of a long latency operation) and the ready-to-run tasks that have to be executed on this instance. At this point it is important to clearly distinguish between tasks and threads. Tasks are unrelated, coarse-grain computations that are distributed over the machine instances to gain speedup by parallel evaluation; they are indirectly introduced by the `letrem` construct. Threads are clustered into closely related groups represented by the `letpar` construct and are fine-grain computations that are used to efficiently cover long latency operations. Only when a task contains no more ready-to-run threads and it is still waiting for a long latency operation, it is suspended and placed in the task pool. Every distributed implementation of the STGM uses tasks, but threads are the uncommon feature of the STGM<sub>MT</sub>.

For every `letpar` construct that is executed, a frame is created; it contains a counter storing the number of value bindings that have not been completed yet and the local environment used to store the value bindings. The argument stack has the same function as in the original STGM, but the return stack assumes the functionality of both the return and the update stack of the STGM—this is necessary to correctly deal with updates of closures whose evaluation triggered a long latency operation. The heap is used to store closures—just as in the original STGM. Finally, the dissemination map supports the dissemination of the results of long latency operations to multiple receivers.

### 5.2.2 Applications

Execution of the application of a function to some arguments, pushes the arguments on the stack and enters the closure that represents the function. In contrast, the application of a data constructor, initiates a return operation.

$$\frac{\mathcal{I}[i \mapsto \text{Eval } (f \{\overline{x_N}\}) \rho \quad \Gamma \quad fs \quad as \quad rs \quad h \quad ds]}{\stackrel{(1)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{Enter } (\rho f) \quad \Gamma \quad fs \quad [\overline{\rho x_N}] ++ as \quad rs \quad h \quad ds]}$$


---


$$\frac{\mathcal{I}[i \mapsto \text{Eval } (c \{\overline{x_N}\}) \rho \quad \Gamma \quad fs \quad as \quad rs \quad h \quad ds]}{\stackrel{(2)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{RetTerm } \langle c, [\overline{\rho x_N}] \rangle \quad \Gamma \quad fs \quad as \quad rs \quad h \quad ds]}$$

We use  $\overline{\phantom{x}}$  to indicate repetition, e.g.,  $[\overline{\rho x_N}]$  stands for  $[\rho x_1, \dots, \rho x_N]$

Evaluating an unboxed variable returns the unboxed value represented by this variable in the local environment  $\rho$ .

$$\frac{\mathcal{I}[i \mapsto \text{Eval } x \# \rho \quad \Gamma \quad fs \quad as \quad rs \quad h \quad ds]}{\stackrel{(3)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{RetTerm } (\rho x) \quad \Gamma \quad fs \quad as \quad rs \quad h \quad ds]}$$

### 5.2.3 Entering a Closure

A not updatable closure is entered by evaluating its code under an environment built from the closures arguments and the appropriate number of parameters from the argument stack. The body  $\gamma$  of the closure is a function that applied

to the environment  $\rho$  yields the code form that has to be executed. The environment  $\rho$  is constructed by taking  $length\ xs$  arguments from the stack and associating the free variables  $vs$  with the environment  $eos$  of the closure.

$$\frac{\mathcal{I}[i \mapsto \text{Enter } o \quad \Gamma \quad fs \quad as \quad rs \quad h[o \mapsto \langle (vs \setminus n\ xs \rightarrow \gamma), eos \rangle] \quad ds]}{\xrightarrow{(4)} \mathcal{I}[i \mapsto (\gamma \rho) \quad \Gamma \quad fs \quad as' \quad rs \quad h \quad ds]}$$

where

$$\begin{aligned} aos \ ++ \ as' &= as, \text{ such that } length\ aos = length\ xs \\ \rho &= [xs \mapsto aos] \ ++ [vs \mapsto eos] \end{aligned}$$

Updatable closures are always nullary (cf. [Pey92]). In the original STGM, such closures push an update frame; in the STGM<sub>MT</sub>, they create an *Upd* dissemination entry—as soon as a value is passed to this entry, the closure is updated with this value. Depending on the type of the value that is computed by the closure, we distinguish two cases: first, if the type is non-functional it is sufficient to extend the dissemination entry referenced by the topmost return continuation; second, if the type is functional, the closure has to be reentered after the update, i.e., a return continuation initiating the reentering is pushed and a new slot  $d$  is created in the dissemination map. Note that in the first case, the argument stack is guaranteed to be empty in type correct programs.

$$\frac{\mathcal{I}[i \mapsto \text{Enter } o \quad \Gamma \quad fs \quad [] \quad \langle d, cont, as_p \rangle : rs \quad h \quad ds[d \mapsto ms]]}{\text{when } h\ o = \langle (vs \setminus u \ [] \rightarrow \gamma), eos \rangle \text{ and has non-functional type}} \xrightarrow{(5)} \mathcal{I}[i \mapsto (\gamma \rho) \quad \Gamma \quad fs \quad [] \quad \langle d, cont, as_p \rangle : rs \quad h \quad ds[d \mapsto (Upd\ o) : ms]]$$

where

$$\rho = [vs \mapsto eos]$$

$$\frac{\mathcal{I}[i \mapsto \text{Enter } o \quad \Gamma \quad fs \quad as \quad rs \quad h \quad ds]}{\text{when } h\ o = \langle (vs \setminus u \ [] \rightarrow \gamma), eos \rangle \text{ and has functional type}} \xrightarrow{(6)} \mathcal{I}[i \mapsto (\gamma \rho) \quad \Gamma \quad fs \quad [] \quad \langle d, (Enter\ o), as \rangle : rs \quad h \quad ds[\underline{d} \mapsto [Upd\ o]]]$$

where

$$\rho = [vs \mapsto eos]$$

Finally, entering a (not updatable) closure needing more than the available arguments indicates that a partial application has to be passed to the topmost return continuation, i.e., the partial application must be distributed using the dissemination slot  $d$  referenced by the return continuation. This case occurs when either a thunk (cf. [Pey92]) has to be updated with a partial application or a partial application has to be communicated to a remote instance.

$$\frac{\mathcal{I}[i \mapsto \text{Enter } o \quad \Gamma \quad fs \quad as \quad \langle d, cont, as_p \rangle : rs \quad h \quad ds]}{\text{when } h\ o = \langle (vs \setminus n\ xs \rightarrow \gamma), eos \rangle \text{ and } length\ as < length\ xs} \xrightarrow{(7)} \mathcal{I}[i \mapsto (MsgPAPP\ d\ o_p\ cont) \quad \Gamma \quad fs \quad as_p \quad rs \quad h' \quad ds]$$

where

$$\begin{aligned} xs_1 \ ++ \ xs_2 &= xs, \text{ such that } length\ xs_1 = length\ as \\ h' &= h[o_p \mapsto \langle ((f.xs_1) \setminus n \ [] \rightarrow Eval\ (f\ xs_1)), o : as \rangle] \end{aligned}$$

A new closure, named  $o_p$ , is created, which implements the partial application; its structure corresponds to the representation of partial applications in the original STGM.

#### 5.2.4 Local Bindings

**letrec** expressions behave as in the original STGM, but note the partial application of the code form *Eval* in the body of the new closures.

$$\frac{\mathcal{I}[i \mapsto \text{Eval} \left( \frac{\mathbf{letrec}}{x_N = vs_N \setminus \pi_N xs_N \rightarrow e_N} \right) \rho \quad \Gamma \quad fs \quad as \quad rs \quad h \quad ds]}{\stackrel{(8)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{Eval } e \rho' \quad \Gamma \quad fs \quad as \quad rs \quad h' \quad ds]}$$

where

$$\begin{aligned} h' &= h[o_N \mapsto \langle (vs_N \setminus \pi_N xs_N \rightarrow (Eval e_N)), \rho' vs_N \rangle] \\ \rho' &= \rho[x_N \mapsto o_N] \end{aligned}$$

Evaluating a **letrem** expression allocates a closure on a new instance  $k$ . Additional forwarding closures that contain *EnterOn* code forms are used for two purposes: first, to reference the new closure  $o_k$  that is allocated on the new instance  $k$  from the current instance  $i$  and, second, to reference the closures ( $\rho v_N$ ) that are contained in the environment of the new closure but are located on the current instance  $i$ .

$$\frac{\mathcal{I}[i \mapsto \text{Eval} \left( \frac{\mathbf{letrem}}{x = \{v_N\} \setminus \pi \{ \} \rightarrow e_1} \right) \rho \quad \Gamma_i \quad fs_i \quad as_i \quad rs_i \quad h_i \quad ds_i]}{\stackrel{(9)}{\Longrightarrow} \begin{array}{l} \mathcal{I}[i \mapsto \text{Eval } e_2 \rho[x \mapsto o] \quad \Gamma_i \quad fs_i \quad as_i \quad rs_i \quad h'_i \quad ds_i, \\ \underline{k} \mapsto \text{Next} \quad \square \quad \square \quad \square \quad \square \quad h_k \quad \square \end{array}}$$

where

$$\begin{aligned} h'_i &= h_i[o \mapsto \langle ([u] \setminus u \square \square \rightarrow EnterOn k v), o_k \rangle] \\ h_k &= \left[ \frac{o_k \mapsto \langle ([v_N] \setminus \pi \square \square \rightarrow (Eval e_1)), [o_N] \rangle}{o_N \mapsto \langle ([v'] \setminus u \square \square \rightarrow EnterOn i v'), \rho v_N \rangle} \right] \end{aligned}$$

In Section 4.3, **letrem** was used to implement the meta-function **par**. Following the definition of **par** in [HMP94], the closure allocated on the remote instance  $k$  must be evaluated immediately. To achieve this behaviour, the initial code form of  $k$  must be (*Enter o*) instead of *Next*.

**letpar** constructs specify related, but independent work; furthermore, the evaluation of the body expression has to be synchronized with the delivery of the values demanded in the value bindings. To this end, the code form *Sync* together with a new frame  $f$  are employed. The first argument of *Sync* contains the still to be evaluated value bindings. The frame maintains a counter of the number of value bindings whose value was not yet added to the environment that is also held in the frame. Note that the number of still awaited values will be greater than the number of bindings in the *Sync* form when the computation of some values is hindered by long latency operations.

$$\frac{\mathcal{I}[i \mapsto \text{Eval } (\mathbf{letpar} \ b_1; \dots; b_n \ \mathbf{in} \ e) \ \rho \ \Gamma \ fs \ as \ rs \ h \ ds]}{\stackrel{(10)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{Sync } [b_1, \dots, b_n] \ e \ f \ \Gamma \ fs' \ as \ rs \ h \ ds]}$$

where

$$fs' = fs[f \mapsto \langle n, \rho \rangle]$$

In a concrete implementation, the guaranteed independence of the value bindings within one **letpar** can be used to partition the code generated from an STG<sub>MT</sub>-program into non-synchronizing threads, i.e., stateless threads.

If there are unprocessed value bindings in a *Sync*, one is selected and its right hand side  $e_1$  is evaluated. A new return continuation is pushed on the stack; it contains the remaining part of the *Sync* code form and the values currently on the argument stack. The new slot  $d$  in the dissemination map is used to eventually distribute the result of  $e_1$ . The dissemination entry (*Store f x*) indicates that the result has to be stored in the environment of frame  $f$  with the local name  $x$ .

$$\frac{\mathcal{I}[i \mapsto \text{Sync } ((x\# = e_1):bs) \ e_2 \ f \ \Gamma \ fs[f \mapsto \langle n, \rho \rangle] \ as \ rs \ h \ ds]}{\stackrel{(11)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{Eval } \ e_1 \ \rho \ \Gamma \ fs \ [] \ rs' \ h \ ds']}$$

where

$$\begin{aligned} rs' &= \langle d, (\text{Sync } bs \ e_2 \ f), as \rangle : rs \\ ds' &= ds[d \mapsto [\text{Store } f \ x]] \end{aligned}$$

If there are no more unprocessed value bindings in a *Sync* form, the behaviour depends on the value of the synchronization counter in the associated frame  $f$ . If it is zero, all values are available, and the body expression can be evaluated; otherwise, the evaluation of the body form has to wait for the delivery of the remaining values, but there is no more independent work in the **letpar** that created this *Sync* form. Nevertheless, it is usually not necessary to suspend the current task; there may be further work in textually enclosing **letpars** or in the calling function. In order to utilize such work the code form *RetDelay* is used. The evaluation of the expression  $e$  is deferred to a new task *cont* that is activated only after the long latency operation is completed.

$$\frac{\mathcal{I}[i \mapsto \text{Sync } [] \ e \ f \ \Gamma \ fs[f \mapsto \langle 0, \rho \rangle] \ as \ rs \ h \ ds]}{\stackrel{(12)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{Eval } \ e \ \rho \ \Gamma \ fs \ as \ rs \ h \ ds]}$$

$$\mathcal{I}[i \mapsto \text{Sync } [] \ e \ f \ \Gamma \ fs \ as \ rs \ h \ ds]$$

when  $fs \ f = \langle n, \rho \rangle$  and  $n > 0$

$$\stackrel{(13)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{RetDelay } \underline{x} \ [x \mapsto d] \ \Gamma \ .cont \ fs \ as \ rs \ h \ ds[d \mapsto []]]$$

where

$$cont = \langle \text{Sync } [] \ e \ f, \langle d, \text{Next}, [] \rangle \rangle$$

The form *RetDelay* informs the enclosing computation about the fact that a long latency operation delayed the delivery of the requested value and that no more local work is available. This information is propagated through all the return continuations until a *Sync* form is found that has some work. When the delayed value eventually becomes available it is distributed using the new slot  $d$  of the dissemination map.

### 5.2.5 Remote Method Invocation

In the STGM<sub>MT</sub> as in the original STGM, accessing a data structure means to evaluate a closure. Hence, the code form *EnterOn* represents a remote data access; it initiates the evaluation of the closure on the remote instance  $k$  by placing a new task *enter* into the task pool  $\Gamma_k$ . This task eventually distributes the result of the remote computation using the dissemination slot  $d_k$ , which forwards it to the slot  $d_i$  on the instance  $i$  that initiated the whole process.

Note that the time between the execution of an *EnterOn* and the delivery of its result is, in general, unbound. Hence, it is important that the instance  $i$  can do some useful work while waiting for the delivery of the result. The delay induced by the remote access is signaled with *RetDelay* to the enclosing computation. If the return continuation is a *Sync* form that still has some unevaluated value bindings left, it can continue by evaluating one of these bindings as they do not depend on the delayed value.

$$\begin{array}{c} \mathcal{I}[i \mapsto \text{EnterOn } k \ v \ \rho \quad \Gamma_i \quad fs_i \ as_i \ rs_i \ h_i \ ds_i, \\ \quad k \mapsto \text{task} \quad \Gamma_k \quad fs_k \ as_k \ rs_k \ h_k \ ds_k] \\ \xrightarrow{(14)} \mathcal{I}[i \mapsto \text{RetDelay } \underline{x} \ [x \mapsto d_i] \quad \Gamma_i \quad fs_i \ as_i \ rs_i \ h_i \ ds'_i, \\ \quad k \mapsto \text{task} \quad \Gamma_k.\text{enter} \ fs_k \ as_k \ rs_k \ h'_k \ ds'_k] \end{array}$$

where

$$\begin{array}{l} \text{enter} = \langle \text{Eval } v \ [v \mapsto \rho \ v], \langle d_k, \text{Next}, [] \rangle \rangle \\ ds'_i = ds_i[d_i \mapsto []] \\ ds'_k = ds_k[\underline{d_k} \mapsto [\text{RetTo } i \ d_i]] \end{array}$$

### 5.2.6 Selection

A **case** selects the appropriate alternative on the basis of the scrutinized value.

$$\begin{array}{c} \mathcal{I}[i \mapsto \text{Eval} \left( \begin{array}{c} \text{case } x \# \text{ of} \\ \dots ; c \{ \overline{v_N} \} \rightarrow e ; \dots \end{array} \right) \rho \quad \Gamma \quad fs \ as \ rs \ h \ ds] \\ \text{when } \rho \ x = \langle c, [\overline{w_N}] \rangle \\ \xrightarrow{(15)} \mathcal{I}[i \mapsto \text{Eval } e \ \rho[\overline{v_N} \mapsto \overline{w_N}]] \quad \Gamma \quad fs \ as \ rs \ h \ ds] \\ \mathcal{I}[i \mapsto \text{Eval} \left( \begin{array}{c} \text{case } x \# \text{ of} \\ \overline{c_N \ v_N} \rightarrow e_N ; \text{default } \rightarrow e \end{array} \right) \rho \quad \Gamma \quad fs \ as \ rs \ h \ ds] \\ \text{when } \rho \ x = \langle c, \overline{ws} \rangle \text{ and } \overline{c} \neq \overline{c_N} \\ \xrightarrow{(16)} \mathcal{I}[i \mapsto \text{Eval } e \ \rho \quad \Gamma \quad fs \ as \ rs \ h \ ds] \end{array}$$

If no alternative matches, a fatal failure occurs.

### 5.2.7 Returning a Proper Value

Returning the unboxed form of a data term implies to distribute it using the dissemination slot  $d$  referenced by the topmost return continuation; afterwards, the continuation is executed. Returning a data term, the argument stack always is empty for type correct programs.

$$\frac{\mathcal{I}[i \mapsto \text{RetTerm } \langle c, ws \rangle \quad \Gamma \quad fs \quad [] \quad \langle d, cont, as_p \rangle : rs \quad h \quad ds]}{\stackrel{(17)}{\Longrightarrow} \mathcal{I}[i \mapsto (\text{MsgTerm } d \langle c, ws \rangle \text{ cont}) \quad \Gamma \quad fs \quad as_p \quad \quad \quad rs \quad h \quad ds]}$$

### 5.2.8 Returning a Delayed Value

When the code form *RetDelay* is executed, a long latency operation delayed the delivery of some intermediate result and local, but independent, computations should be employed to cover the delay, i.e., not yet evaluated value bindings in surrounding **letpars** should be executed. For this mechanism to work properly, two jobs have to be carried out: first, some independent work has to be found and, second, when the delayed value finally arrives, it has to be introduced into the ongoing computation.

Imagine a *Sync* code form with multiple value bindings. When the first binding is evaluated, according to Rule (11), a return continuation is pushed that contains the *Sync* form with the remaining bindings. To utilize the independent work constituted by these bindings, we just have to return to this *Sync*. This is what *RetDelay* does while simultaneously taking care of the second issue, namely, preparing the asynchronous delivery of the delayed value. The latter is done by placing *Fwd* entries in the dissemination slot ( $\rho v$ ) that will eventually be used to deliver the remote value; as a result the asynchronous delivery updates closures that have to be updated with the remote value and it stores the value in frames whose associated *Sync* forms wait for that value.

When the topmost return continuation belongs to a closure that has to be updated (this is handled by the “if” in the rule below), this closure has to be overwritten with a new closure that contains a *RetDelay* code form. This ensures that repeatedly entering the closure does not cause multiple remote accesses; instead, remote accesses are shared.

When there are no arguments on the stack, the delayed value can be forwarded to the entry specified by the return continuation by using a *Fwd* entry, which sends any value delivered via this entry on to the entry given in its argument (in the following rule, to  $d$ ).

$$\frac{\mathcal{I}[i \mapsto \text{RetDelay } v \rho \quad \Gamma \quad fs \quad [] \quad rs \quad h \quad \quad \quad ds[\rho v \mapsto ms]]}{\text{when } rs = \langle d, cont, as_p \rangle : rs'}$$

$$\stackrel{(18)}{\Longrightarrow} \mathcal{I}[i \mapsto cont \quad \Gamma \quad fs \quad as_p \quad rs' \quad h' \quad ds[\rho v \mapsto (\text{Fwd } d) : ms]]$$

where

$$h' = \begin{cases} \text{if } (\text{Enter } o_u) = cont \\ \text{then } h[o_u \mapsto \langle ([v'] \setminus u \quad [] \rightarrow \text{RetDelay } v'), \rho v \rangle \\ \text{else } h \end{cases}$$

The above rule together with Rule (13) implies that the continuations on the return stack are popped in the process of exploiting work that is independent from the delayed value. Only when the stack is completely empty, it is necessary to switch to a completely unrelated task. An interesting consequence

of this property is that only one stack per machine instance is needed in the  $\text{STGM}_{\text{MT}}$ —instead of one stack per task.

If there are arguments remaining on the stack, the delayed value is a partial application, which has to be applied to the arguments on the stack when it is eventually delivered. To realize the synchronization between the delivery of the partial application via the dissemination entry  $(\rho v)$  and the code performing the application, a new frame  $f$  and a task  $enter$  are created. The form  $Store$  in the dissemination entry is used to put the partial application into the frame  $f$  when it arrives; this, in turn, enables the task  $enter$ .

$$\frac{\mathcal{I}[i \mapsto \text{RetDelay } v \ \rho \ \Gamma \quad fs \quad as \quad rs \quad h \quad ds[\rho v \mapsto ms]]}{\text{when } rs = \langle d, cont, as_p \rangle : rs'}$$

$$\xrightarrow{(19)} \mathcal{I}[i \mapsto cont \quad \Gamma. enter \quad fs' \quad as_p \quad rs' \quad h' \quad ds']$$

where

$$h' = \begin{cases} \text{if } (Enter \ o_u) = cont \\ \text{then } h[o_u \mapsto \langle ([\underline{x}] \setminus u \ \square) \rightarrow \text{RetDelay } v \rangle, \rho v] \\ \text{else } h \end{cases}$$

$$fs' = fs[f \mapsto (1, [\underline{us} \mapsto as])] \\ enter = (Sync \ \square \ (x \ vs) \ f, \langle d, Next, \square \rangle) \\ ds' = ds[\rho v \mapsto (Store \ f \ \underline{x}) : ms]$$

As mentioned above, an empty return stack indicates that the current task has no more work to offer.

$$\frac{\mathcal{I}[i \mapsto \text{RetDelay } v \ \rho \ \Gamma \quad fs \ \square \ \square \quad h \quad ds]}{\xrightarrow{(20)} \mathcal{I}[i \mapsto Next \quad \Gamma \quad fs \ \square \ \square \quad h \quad ds]}$$

### 5.2.9 Task Management

*Next* passes control to some arbitrary task from the task pool. The return continuation of the new task is placed on the stack, then the task is executed.

$$\frac{\mathcal{I}[i \mapsto Next \ \Gamma.(task, r) \quad fs \ \square \ \square \quad h \quad ds]}{\xrightarrow{(21)} \mathcal{I}[i \mapsto task \ \Gamma \quad fs \ \square \ [r] \quad h \quad ds]}$$

A concrete implementation would, of course, ensure that the selected task is ready to run, e.g., by maintaining a list of such tasks; otherwise, the selected task may just be suspended again immediately.

### 5.2.10 Dissemination of Messages

The distribution of unboxed data terms is performed by considering one entry of the dissemination slot after the other and use the value accordingly to store it into the environment of a frame (*Store*), update a closure (*Upd*), forward it to another dissemination slot (*Fwd*), or send it to another instance (*RetTo*).

$$\frac{\mathcal{I}[i \mapsto \text{MsgTerm } d \ \langle c, ws \rangle \ task \ \Gamma \quad fs \quad as \quad rs \quad h \quad ds[d \mapsto \square]]}{\mathcal{I}[i \mapsto task \quad \Gamma \quad fs \quad as \quad rs \quad h \quad ds]}$$

$$\begin{array}{c}
\mathcal{I}[i \mapsto \text{MsgTerm } d \langle c, ws \rangle \text{ task } \Gamma \text{ fs } \text{ as } \text{ rs } \text{ h} \quad ds] \\
\text{when } fs \text{ f} = \langle m, \rho \rangle \text{ and } ds \text{ d} = (\text{Store } f \text{ x}):ms \\
\stackrel{(23)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{MsgTerm } d \langle c, ws \rangle \text{ task } \Gamma \text{ fs}' \text{ as } \text{ rs } \text{ h} \quad ds[d \mapsto ms]] \\
\text{where} \\
fs' = fs[f \mapsto \langle m-1, \rho[x \mapsto \langle c, ws \rangle] \rangle] \\
\hline
\mathcal{I}[i \mapsto \text{MsgTerm } d \langle c, [\overline{w}_N] \rangle \text{ task } \Gamma \text{ fs } \text{ as } \text{ rs } \text{ h} \quad ds] \\
\text{when } ds \text{ d} = (\text{Upd } o_u):ms \\
\stackrel{(24)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{MsgTerm } d \langle c, [\overline{w}_N] \rangle \text{ task } \Gamma \text{ fs } \text{ as } \text{ rs } \text{ h}_u \quad ds[d \mapsto ms]] \\
\text{where} \\
h_u = h[o_u \mapsto \langle ([\overline{v}_N] \setminus n \ [] \rightarrow \text{Eval } (c \{ \overline{v}_N \})), [\overline{w}_N] \rangle] \\
\hline
\mathcal{I}[i \mapsto \text{MsgTerm } d \langle c, ws \rangle \text{ task } \Gamma \text{ fs } \text{ as } \text{ rs } \text{ h} \quad ds] \\
\text{when } ds \text{ d} = (\text{Fwd } d'):ms \\
\stackrel{(25)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{MsgTerm } d' \langle c, ws \rangle \text{ cont } \Gamma \text{ fs } \text{ as } \text{ rs } \text{ h} \quad ds[d \mapsto ms]] \\
\text{where} \\
\text{cont} = \text{MsgTerm } d \langle c, ws \rangle \text{ task} \\
\hline
\mathcal{I}[i \mapsto \text{MsgTerm } d_i \langle c, [\overline{w}_N] \rangle \text{ task}_i \quad \Gamma_i \text{ fs}_i \text{ as}_i \text{ rs}_i \text{ h}_i \quad ds_i, \\
k \mapsto \text{task}_k \quad \Gamma_k \text{ fs}_k \text{ as}_k \text{ rs}_k \text{ h}_k \quad ds_k] \\
\text{when } ds \text{ d}_i = (\text{RetTo } k \text{ d}_k):ms \\
\stackrel{(26)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{MsgTerm } d_i \langle c, [\overline{w}_N] \rangle \text{ task}_i \quad \Gamma_i \text{ fs}_i \text{ as}_i \text{ rs}_i \text{ h}_i \quad ds'_i, \\
k \mapsto \text{MsgTerm } d_k \langle c, [\overline{w}'_N] \rangle \text{ task}_i \quad \Gamma_k \text{ fs}_k \text{ as}_k \text{ rs}_k \text{ h}'_k \quad ds_k] \\
\text{where} \\
ds'_i = ds_i[d_i \mapsto ms] \\
h'_k = h_k[w'_N \mapsto \langle ([\overline{v}] \setminus u \ [] \rightarrow \text{EnterOn } i \text{ v}), [\overline{w}_N] \rangle] \\
\hline
\text{In the last rule, where the value is transmitted to another instance, the arguments must be replaced by forwarding closures using the } \text{EnterOn} \text{ code form (compare this to Rule (9)).} \\
\text{The distribution of partial applications is similar to that of data terms.} \\
\hline
\mathcal{I}[i \mapsto \text{MsgPAPP } d \text{ o } \text{task } \Gamma \text{ fs } \text{ as } \text{ rs } \text{ h} \quad ds[d \mapsto []]] \\
\mathcal{I}[i \mapsto \text{task} \quad \Gamma \text{ fs } \text{ as } \text{ rs } \text{ h} \quad ds] \\
\hline
\mathcal{I}[i \mapsto \text{MsgPAPP } d \text{ o } \text{task } \Gamma \text{ fs } \text{ as } \text{ rs } \text{ h} \quad ds] \\
\text{when } fs \text{ f} = \langle m, \rho \rangle \text{ and } ds \text{ d} = (\text{Store } f \text{ x}):ms \\
\stackrel{(28)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{MsgPAPP } d \text{ o } \text{task } \Gamma \text{ fs}' \text{ as } \text{ rs } \text{ h} \quad ds[d \mapsto ms]] \\
\text{where} \\
fs' = fs[f \mapsto \langle m-1, \rho[x \mapsto o] \rangle] \\
\hline
\mathcal{I}[i \mapsto \text{MsgPAPP } d \text{ o } \text{task } \Gamma \text{ fs } \text{ as } \text{ rs } \text{ h} \quad ds[d \mapsto (\text{Upd } o_u):ms]] \\
\stackrel{(29)}{\Longrightarrow} \mathcal{I}[i \mapsto \text{MsgPAPP } d \text{ o } \text{task } \Gamma \text{ fs } \text{ as } \text{ rs } \text{ h}_u \quad ds[d \mapsto ms]] \\
\text{where} \\
h_u = h[o_u \mapsto \langle ([\overline{f}] \setminus n \ [] \rightarrow \text{Eval } (f \{ \overline{f} \})), o \rangle] \\
\hline
\end{array}$$

$$\begin{array}{c} \mathcal{I}[i \mapsto \text{MsgPAPP } d \ o \ \text{task} \ \Gamma \ fs \ as \ rs \ h \ ds[d \mapsto (\text{Fwd } d'):ms]] \\ \xrightarrow{(30)} \mathcal{I}[i \mapsto \text{MsgPAPP } d' \ o \ \text{cont} \ \Gamma \ fs \ as \ rs \ h \ ds[d \mapsto ms]] \end{array}$$

where

$$\text{cont} = \text{MsgPAPP } d \ o \ \text{task}$$

$$\frac{\mathcal{I}[i \mapsto \text{MsgPAPP } d_i \ o_i \ \text{task}_i \ \Gamma_i \ fs_i \ as_i \ rs_i \ h_i \ ds_i, \quad k \mapsto \text{task}_k \ \Gamma_k \ fs_k \ as_k \ rs_k \ h_k \ ds_k]}{\text{when } \left( \begin{array}{l} ds \ d_i = (\text{RetTo } k \ d_k):ms, \\ h_i \ o_i = \langle (vs \setminus n \ [] \rightarrow d), o_f:[w_N] \rangle, \text{ and} \\ h_i \ o_f = \langle (vs_f \setminus \pi \ xs_f \rightarrow d_f), [w_{f_N}] \rangle \end{array} \right)}$$

$$\xrightarrow{(31)} \mathcal{I}[i \mapsto \text{MsgPAPP } d_i \ o_i \ \text{task}_i \ \Gamma_i \ fs_i \ as_i \ rs_i \ h_i \ ds_i[d_i \mapsto ms], \quad k \mapsto \text{MsgPAPP } d_k \ o_k \ \text{task}_k \ \Gamma_k \ fs_k \ as_k \ rs_k \ h'_k \ ds_k]$$

where

$$h'_k = h_k \left[ \begin{array}{l} o_k \mapsto \langle (vs \setminus n \ [] \rightarrow d), o'_f:[w'_N] \rangle \\ o'_f \mapsto \langle (vs_f \setminus \pi \ xs_f \rightarrow d_f), [w'_{f_N}] \rangle \\ w'_N \mapsto \langle ([v] \setminus u \ [] \rightarrow \text{EnterOn } i \ v), [w_N] \rangle \\ w'_{f_N} \mapsto \langle ([v] \setminus u \ [] \rightarrow \text{EnterOn } i \ v), [w_{f_N}] \rangle \end{array} \right]$$

The complexity of the last rule is due to the fact that forwarding *EnterOn* closures have to be created for all the objects referenced in the environments of the transmitted closures.

It is not sufficient to transmit only the partial application  $(h_i \ o_i)$ , but the closure  $(h_i \ o_f)$  referenced in the partial application's first environment argument has to be transmitted, too. This closure represents the function that was (partially) applied to the  $w_i$ . To execute the partial application on the instance  $k$ , the closure  $(h_i \ o_f)$  has obviously to be on instance  $k$  also.

## 6 Related Work

The Threaded Abstract Machine (TAM) [CGSE93] is designed to implement the dataflow language Id [Nik90]. It applies multithreading based on stateless threads to tolerate long latency operations. In this respect it is close to the work presented in this paper, but the realization of this basic idea differs considerably. Instead of source-to-source transformations, the Id compiler based on the TAM builds a structured dataflow graph as an intermediate representation and has to apply sophisticated thread partitioning schemes [SCG95]. These partitioning algorithms require graphs without cyclic dependencies, which can not be guaranteed for a lazy language. In contrast to the STGM<sub>MT</sub>, which employs asynchronous operations only when a long latency operation is actually encountered, the TAM uses asynchronous operations by default.

In comparison to the parallel implementation of the STGM, it is interesting to note that the STGM<sub>MT</sub>, only needs a single argument and return stack per machine instance, instead of one stack per task. Furthermore, on entry, closures need not be overwritten with a “queue-me” code pointer. The updating performed in the rules (18) and (19) is sufficient and happens only when a remote

access occurred. A proper comparison with a distributed implementation of the original STGM, e.g., GUM [THJ<sup>+</sup>96], has to wait until a first implementation of the STGM<sub>MT</sub> is working.

## 7 Conclusion

To decrease the impact of long latency operations on the execution time of distributed implementations of the Spineless Tagless G-machine, the STGM<sub>MT</sub> extends the abstract machine language with abstract notions of independent local computations (thread boundaries and thread synchronization) and with an abstract notion of distribution. This enables the use of source-to-source transformations to increase the tolerance of the code to long latency operations.

While the behaviour of the new abstract machine can be studied using the operational semantics provided in this paper, it remains to be shown that the proposed techniques decrease the impact of the communication overhead in an actual implementation of a lazy functional language on parallel computers with distributed memory.

**Acknowledgments.** I am grateful to Paul H. J. Kelly, Hans-Wolfgang Loidl, and Simon L. Peyton Jones for their excellent comments on a previous version of this paper.

## References

- [CGSE93] David E. Culler, Seth Copen Goldstein, Klaus Erik Schauer, and Thorsten von Eicken. TAM—a compiler controlled threaded abstract machine. *Journal of Parallel and Distributed Computing*, 18:347—370, 1993.
- [DP93] Marco Danelutto and Susanna Pelagatti. Structuring parallelism in a functional framework. Technical Report TR-29/93, Dipartimento di Informatica, Università di Pisa, 1993.
- [EAL93] Dawson R. Engler, Gregory R. Andrews, and David K. Lowenthal. Filaments: Efficient support for fine-grain parallelism. TR 93-13a, University of Arizona, 1993.
- [Fos95] Ian Foster. *Designing and Building Parallel Programs*. Addison-Wesley, 1995.  
(URL: <http://www.mcs.anl.gov/dbpp/text/book.html>).
- [H<sup>+</sup>95] Herbert H.J. Hum et al. A design study of the EARTH multiprocessor. In *Proceedings of Parallel Architectures and Compilation Techniques*, 1995.

- [HMP94] Kevin Hammond, Jim S. Mattson, and Simon L. Peyton Jones. Automatic spark strategies and granularity for a parallel functional language reducer. In *CONPAR '94*, 1994.
- [JL91] Simon L. Peyton Jones and John Launchbury. Unboxed values as first class citizens in a non-strict functional language. In J. Hughes, editor, *FPCA '91*, 1991.
- [Kel89] P. Kelly. *Functional Programming for Loosely-Coupled Multiprocessors*. Pitman, 1989.
- [Nik90] R.S. Nikhil. Id (version 90.0) reference manual. Technical report, MIT Laboratory for Computer Science, July 1990. CSG Memo 284-1.
- [PCS89] Simon L. Peyton Jones, Chris Clack, and Jon Salkild. High-performance parallel graph reduction. In E. Odijk, M. Rem, and J.-C. Syre, editors, *Proceedings of PARLE (Volume 1)*. Springer Verlag, 1989. LNCS 365.
- [Pey92] Simon L. Peyton-Jones. Implementing lazy functional languages on stock hardware: the Spineless Tagless G-machine. *Journal of Functional Programming*, 2(2), 1992.
- [PS88] Simon L. Peyton-Jones and J. Salkild. The spineless tagless G-machine. In *Workshop on Implementations of Lazy Functional Languages*, Aspensas, Schweden, 1988.
- [SCG95] Klaus E. Schauser, David E. Culler, and Seth C. Goldstein. Separation constraint partitioning - a new algorithm for partitioning non-strict programs into sequential threads. In *Proceedings of the Symposium on Principles of Programming Languages*, 1995.
- [THJ<sup>+</sup>96] P. W. Trinder, K. Hammond, J. S. Mattson Jr, A. S. Partridge, and S. L. Peyton Jones. Gum: a portable implementation of Haskell. Submitted for publication, 1996.

## A Notation

Tuples with  $n$  components are written as  $\langle x_1, \dots, x_n \rangle$  and have type  $\langle \alpha_1, \dots, \alpha_n \rangle$  when the  $x_i$  have type  $\alpha_i$ .

Sequences have the form  $[x_1, \dots, x_n]$  and type  $[\alpha]$  when  $\alpha$  is the type of all elements  $x_i$ . The functions  $(++)$ ,  $(:)$ , and *length* represent concatenation of two sequences, prepending an element at a sequence, and obtaining the length of a sequence, respectively—they have the obvious types.

Maps are sequences of associations, denoted by  $[o_1 \mapsto x_1, \dots, o_n \mapsto x_n]$ , with type  $[Name \mapsto \alpha]$  when the  $x_i$  are of type  $\alpha$ . The notation  $m[o \mapsto x]$  is, in left-hand sides of transition rules, used to indicate that the name  $o$  is mapped

to  $x$  in the map  $m$  and, in right-hand sides, to denote the map that is obtained by replacing the value associated with  $o$  in  $m$  with  $x$ . We write for  $[o_1 \mapsto x_1, \dots, o_n \mapsto x_n]$  also  $[os \mapsto xs]$  with  $os = [o_1, \dots, o_n]$  and  $xs = [x_1, \dots, x_n]$ . Sometimes maps of type  $[Name \mapsto \alpha]$  are used as representatives of functions from  $Names$  to values of type  $\alpha$ . Hence,  $(m[o \mapsto x]) o = x$ .

Multisets with elements of type  $\alpha$  are of type  $\{\alpha\}$ ; we use the notation  $\Gamma.e$  as a shorthand for  $\Gamma \cup \{e\}$ .

To support conciseness, we use  $\overline{form}$  to represent the repetition of  $form$ . Using  $N$  as an index for a subform in  $form$  denotes a family of the subform, e.g.,  $\overline{[x_N \mapsto a_N]}$  denotes  $[x_1 \mapsto a_1, \dots, x_N \mapsto a_N]$ .

When new, unique names are required, they are marked by underlining them. Underlying a variable representing a sequence means a sequence of new names.

## B The Structure of Machine Configurations

Machine configurations are maps from instance names to machine instances:

$$Config = [IName \mapsto Inst]$$

The type  $IName$  is a synonym for  $Name$ , we use it for names of machine instances. In the transition rules from Section 5.2, we have  $\mathcal{I} :: Config$ .

A machine instance is a tuple containing seven components. It is of the following type:

$$Inst = \langle Code, \quad \text{— currently executed code} \\ \{Task\}, \quad \text{— remaining tasks (ready \& waiting)} \\ [FName \mapsto Frame], \quad \text{— sync. frames for `letpars`} \\ [Value], \quad \text{— stack of pending arguments} \\ [Cont], \quad \text{— return continuations} \\ [HName \mapsto CIs], \quad \text{— heap: named closures} \\ [DName \mapsto [DEntry]] \rangle \text{— target locations for results}$$

In the transition rules, the  $\langle, \rangle$ , and comma are omitted, and the seven components are just placed side-by-side. The types  $FName$ ,  $HName$ , and  $DName$  are synonyms of the type  $Name$  and are used for names of frames, heap-allocated closures, and dissemination entries, respectively.

The types of tasks and frames are

$$Task = \langle Code, Cont \rangle \quad \text{— return to } Cont \text{ with result of } Code \\ Frame = \langle Int, [Env] \rangle \quad \text{— sync. counter \& completed bindings} \\ Env = [LName \mapsto Value] \quad \text{— local environment}$$

The type  $LName$  is a synonym for  $Name$ , it is used for names of local variables.

Values appear either boxed, i.e., they are the name of a closure in the heap, or they are unboxed, i.e., a data constructor with its arguments. In addition a value can be the name of a dissemination entry. In the following definition of

the data type *Value*, data constructors are omitted—this avoids some clutter in the transition rules and there is no danger that ambiguities arise.

$$\begin{aligned}
\textit{Value} &= \textit{HName} && \text{--- boxed: name of a closure} \\
&| \textit{UValue} && \text{--- unboxed: constructor \& arguments} \\
&| \textit{DName} && \text{--- name of a dissemination entry} \\
\textit{UValue} &= \langle \textit{CName}, [\textit{HName}] \rangle && \text{--- constructor \& arguments}
\end{aligned}$$

The type *CName* is a synonym for *Name*, it is used for names of data constructors.

The type of closures is defined as

$$\textit{Cls} = \langle \underbrace{([\textit{LName}]}_{\text{free}} \backslash \pi \underbrace{[\textit{LName}]}_{\text{args}} \rightarrow \underbrace{(\textit{Env} \rightarrow \textit{Code})}_{\text{body}}, \underbrace{[\textit{HName}]}_{\text{env}} \rangle$$

The sequences of free and argument variables correspond to the list of environment variables and the list of argument variables, respectively, in the function bindings of the STG<sub>MT</sub>-language. The body of a closure is, in contrast to the original STGM, not an expression of the STG<sub>MT</sub>-language, but a function from environments to code forms—the reason for this change is that we sometimes need to place code forms other than *Eval* (e.g., *EnterOn*) into the body of closures.

Return continuations, the elements of the return stack, are triples made of the name of a dissemination entry, a code form, and a sequence of values. The meaning is that the currently executed task distributes its result value using the dissemination entry, and then continues with the code after placing the values on the argument stack.

$$\textit{Cont} = \langle \textit{DName}, \textit{Code}, [\textit{Value}] \rangle$$

The various code forms are defined in the following data type (*Eval* and *Enter* are equal to the corresponding forms of the original STGM):

$$\begin{aligned}
\textit{Code} &= \textit{Eval Expr Env} && \text{--- evaluate expression} \\
&| \textit{Enter HName} && \text{--- enter the given closure} \\
&| \textit{EnterOn IName LName Env} && \text{--- enter on other instance} \\
&| \textit{RetTerm UValue} && \text{--- return unboxed value} \\
&| \textit{RetDelay LName Env} && \text{--- cover long latency} \\
&| \textit{Sync [VBind] Expr FName} && \text{--- wait for local bindings} \\
&| \textit{Next} && \text{--- switch to other task} \\
&| \textit{MsgTerm DName UValue Code} && \text{--- dissem. unboxed value} \\
&| \textit{MsgPAPP DName HName Code} && \text{--- dissem. partial appl.}
\end{aligned}$$

Details on the meaning of the variants are provided in Section 5.2.

Finally, elements of the lists of the dissemination map describe the locations where the code forms *MsgTerm* and *MsgPAPP* have to place disseminated values:

$$\begin{aligned}
\textit{DEntry} &= \textit{Store FName LName} && \text{--- store in the env. of a frame} \\
&| \textit{Upd HName} && \text{--- update closure} \\
&| \textit{Fwd DName} && \text{--- forward to other dissem. entry} \\
&| \textit{RetTo IName DName} && \text{--- communicate to other instance}
\end{aligned}$$