A Hybrid Type System for Lock-Freedom of Mobile Processes

Naoki Kobayashi[1] and Davide Sangiorgi[2]

[1] Tohoku University
[2] Università di Bologna

Abstract. We propose a type system for lock-freedom in the π-calculus, which guarantees that certain communications will eventually succeed. Distinguishing features of our type system are: it can verify lock-freedom of concurrent programs that have sophisticated recursive communication structures; it can be fully automated; it is hybrid, in that it combines a type system for lock-freedom with local reasoning about deadlock-freedom, termination, and confluence analyses. Moreover, the type system is parameterized by deadlock-freedom/termination/confluence analyses, so that any methods (e.g. type systems and model checking) can be used for those analyses. A lock-freedom analysis tool has been implemented based on the proposed type system, and tested for non-trivial programs.

1 Introduction

In this paper, we attack the problem of verifying concurrent programs that create threads and communication channels dynamically. More specifically, we choose the π-calculus [16] as the target language, and consider the problem of verifying the lock-freedom property, which intuitively means that certain communications (or synchronizations) will eventually succeed (possibly under some fairness assumption). Lock-freedom is important for communication-centric computation models like the π-calculus; indeed, in the pure π-calculus, most liveness properties can be turned into the lock-freedom property. For example, the following properties can be reduced to instances of lock-freedom: Will the request of accessing a resource be eventually granted? In a client-server system, will a client request be eventually received from the server? And if so, will the server eventually send back an answer to the client? In multi-threaded programs, can a thread eventually acquire a lock? And if so, will the thread eventually release the lock? The lock-freedom property has also applications to other verification problems and program transformation, such as information flow analysis and program slicing (dependency analysis in general). Verification of liveness properties such as lock-freedom is notoriously hard in concurrency. In formalisms for mobile processes, such as the π-calculus, it is even harder, because of dynamic creation of threads and first-class channels. In these formalisms, *type systems* have emerged as a powerful means for disciplining and controlling the behaviors of the processes.

A. Gupta and S. Malik (Eds.): CAV 2008, LNCS 5123, pp. 80–93, 2008.

Type systems for lock-freedom include [1,8,9,20,21]. An automatic verification tool, TyPiCal [10], has been implemented based on Kobayashi's system [9]. The expressive power of such type systems is, however, very limited. This shows up clearly, for instance, in the treatment of recursion. For example, even primitive recursive functions cannot be expressed in Kobayashi's lock-free type system, since it ignores value-dependent behaviors completely.

In this paper, we tackle lock-freedom by pursuing a different route. We overcome limitations of previous type systems by combining the lock-freedom analysis with two other analysis: *deadlock-freedom* and *termination*. The result, therefore, is not a "pure" type system, but one that is *parametric* in the techniques employed to ensure deadlock-freedom and termination. Such techniques may themselves be based on type systems (and indeed in the paper we indicate such type systems, or develop them when needed), but could also use other methods (model checking, theorem provers, etc.). The parameterization allows us to go beyond certain limits of type systems, by appealing to other methods. For instance, a type system, as a form of static analysis, may have difficulties in handling value-dependent behaviors (even very simple ones), which are more easily dealt with by other methods such as model checking.

Roughly, we use the deadlock-freedom analysis to ensure that a system can reduce if some of its expected communications have not yet occurred. We then apply a termination analysis to discharge the possibility of divergence and guarantee lock-freedom (i.e., the expected communication will indeed occur). The reasons for appealing to deadlock-freedom are that powerful type-based analyzers exist (notably Kobayashi's systems [11]), and that deadlock-freedom is a safety property, which is easier than liveness to verify in other verification methods such as model checking.

A major challenge was to be able to apply the deadlock and termination analysis *locally*, to subsystems of larger systems. The local reasoning is particularly important for termination. A result forcing a global termination analysis would not be very useful in practice: first, valid concurrent programs may not terminate (e.g., operating systems); second, even if a program is terminating, it can be extremely hard to verify it if the program is large, particularly in languages for mobile processes such as the π-calculus that subsume higher-order formalisms such as the λ-calculus.

Very approximately, our hybrid rule for local reasoning looks as follows:

$$\frac{\models_{\mathtt{DF}} P \qquad \models_{\mathtt{Ter}} P}{\Delta \vdash_{\mathtt{LT}} P} \qquad (*)$$

where $\models_{\mathtt{DF}} P$ and $\models_{\mathtt{Ter}} P$ indicate, respectively, that P is deadlock-free and terminating, and $\Delta \vdash_{\mathtt{LT}} P$ is a typing judgment for lock-freedom. The type environment Δ captures conditions, or "contracts", on the way P interacts with its environment, of the kind "P will eventually send a message on a" and "if P receives a message on a, then P is lock-free afterwards". Such contracts are necessary for the compositionality of the type system for lock-freedom (i.e., local reasoning on lock-freedom). We use Kobayashi's lock freedom types [9], which refine those of the simply-typed π-calculus with *channel usages*, to express the

contracts. Therefore we add rule (∗), as an "axiom", to the rules of Kobayashi's lock freedom type system [9].

The contracts in Δ, however, are completely ignored—and are not guaranteed—in the premises of rule (∗). As a consequence, the resulting type system is unsound. In other words, knowing that P is deadlock-free and terminating is not sufficient to guarantee compositionality and local reasoning. As an example of missing information, P being terminating ensures that P itself has no infinite reductions; but it says nothing on the behaviour of P after it receives a message from other components in the system. (Indeed rule (∗) is only sound if applied globally, to the whole system.)

The first refinement we make for the soundness of rule (∗) is to replace deadlock-freedom and termination with more robust notions, which we call, respectively, *robust deadlock-freedom under* Δ, written $\Delta \models_{\mathtt{RD}} P$, and *robust termination*, written $\models_{\mathtt{RTer}} P$. These stronger notions approximately mean that P is deadlock-free or terminating after any substitution (P may be open, and therefore contain free variables), and any interaction with its environment; $\Delta \models_{\mathtt{RD}} P$ further ensures that P fulfills certain obligations in Δ. The problems of verifying robust deadlock-freedom and robust termination are more challenging than the ordinary ones, due to the additional requirements (e.g., quantifications over substitutions and transition sequences). Existing type systems for deadlock-freedom, notably [11], do meet however the extra conditions for robust deadlock-freedom. We also show how to tune type systems for ordinary termination in a generic manner so to guarantee the stronger property of robust termination. We should stress nevertheless that $\Delta \models_{\mathtt{RD}} P$ and $\models_{\mathtt{RTer}} P$ are semantic requirements: our type system is parametric on the verification methods that guarantee them—one need not employ type systems.

Even with the above refinement of the deadlock-freedom and termination conditions, the hybrid rule (∗) remains unsound. The reason is, roughly, the same as why assume-guarantee reasoning in concurrency often fails in the presence of circularity. In fact, the judgment $\Delta \vdash_{\mathtt{LT}} P$ can be considered a kind of assume-guarantee reasoning, where Δ expresses both assumptions on the environment and guarantees about P's behavior. To prevent circular reasoning, we add a condition $nocap(\Delta)$ that intuitively ensures us that P is independent of its environment, in the sense that P will fulfill its obligations (to perform certain input/output actions) without relying on its environment's behavior. (For example, suppose that there is an obligation to send a message on channel a. The process $\overline{a}[1]$, which sends 1 on a, is fine, since it fulfills the obligation with no assumption. On the other hand, the process $b(x).\overline{a}[x]$, which waits to receive a value on b before sending x on a, is not allowed since it fulfills the obligation only *on the assumption* that the environment will send a message on b.) This leads to the following hybrid rule:

$$\frac{\Delta \models_{\mathtt{RD}} P \qquad \models_{\mathtt{RTer}} P \qquad nocap(\Delta)}{\Delta \vdash_{\mathtt{LT}} P} \qquad \text{(LT-Hyb)}$$

The resulting type system guarantees that any well-typed process P is *weakly lock-free*, in the sense that if an input/output action is declared in P as an action

that should succeed, and if $P \longrightarrow^* Q$, then the action has already succeeded in $P \longrightarrow^* Q$ or there is a further reduction sequence from Q in which the action will succeed. This is similar to the way in which success of passing a test is defined in fair should/must testing [4] and bisimulation, (and also in accordance with other definitions of similar forms of liveness for π-calculus such as [20]).

We have also considered a stronger form of lock-freedom, guaranteeing that the marked actions will eventually succeed on the assumption that the scheduler is strongly fair. We show that our type system can be strengthened to guarantee the strong lock-freedom by adding a condition of partial confluence to rule LT-Hyb above. Again, the partial confluence is only required locally; the whole program need not be confluent.

The verification framework outlined above for lock-freedom (including an automated robust termination analysis) has been implemented as an extension of TyPiCal program analysis tool (except the extension to strong lock-freedom; adding this on top of the present implementation would be tedious but not difficult). We have succeeded in automatically verifying several non-trivial programs, such as symbol tables and binary tree search. These examples are non-trivial because lists and trees are implented as networks of processes connected by channels, and they grow dynamically (both channels and processes are dynamically created and linked). Recursive structures of the kind illustrated in these examples are common in programming languages for mobile processes (the examples in fact, were taken or inspired from Pict programs [15]).

2 Target Language

Syntax. We write $\mathcal{L}$ for the set of *links* (also called *channels*), and $\mathcal{V}$ for the (disjoint) set of *variables*. We use meta-variables $a, b, c, \ldots$ and $x, y, z, \ldots$ for links and variables, respectively. We write $\mathcal{N}$ for the set $\mathcal{L} \cup \mathcal{V} \cup \{\mathtt{true}, \mathtt{false}\}$ of *names* (sometimes called *values*), where `true` and `false` are the usual boolean values. We use meta-variables u, v, w for names. The grammar is the following:

$$P ::= \mathbf{0} \mid \overline{v}^{\chi}[\widetilde{w}].\, P \mid v^{\chi}(\widetilde{y}).\, P \mid (P \,|\, Q) \mid *P \mid (\nu a)\, P \mid \mathbf{if}\ v\ \mathbf{then}\ P\ \mathbf{else}\ Q$$

Here, χ is either $\circ$ or $\bullet$, and $\widetilde{w}$ abbreviates a possibly empty sequence $w_1, \ldots, w_n$. The constructs are the standard ones of the polyadic π-calculus: nil, output and input prefixes, parallel composition, replication ($*P$ behaves like infinitely many copies of P running in parallel), restriction, and a conditional. The only difference is the annotation χ in prefixes, which indicates whether the action is expected to succeed (symbol $\circ$) or not (symbol $\bullet$). (In the type inference of TyPiCal these annotations are actually inferred, in the sense that if the analysis succeed then a set of prefixes that will eventually succeed is marked, see Section 5.) We call a prefix *marked* if its annotation is $\circ$. We usually omit the $\bullet$ annotation, thus for example $a(x).P$ stands for $a^{\bullet}(x).\, P$. As usual, restriction and input prefix are binders. A *closed* process has no free variables. We often omit trailing $\mathbf{0}$, and write $\overline{v}^{\chi}[\widetilde{w}]$ for $\overline{v}^{\chi}[\widetilde{w}].\, \mathbf{0}$. We also write $\overline{v}^{\chi}.P$ and $v^{\chi}.P$ for $\overline{v}^{\chi}[\,].\, P$ and $v^{\chi}(\,).\, P$ respectively. In examples, we use an extension of the above language with natural numbers, list, etc. as they are straightforward to accommodate.

Typing. The type systems that we will propose are defined on top of the simply-typed π-calculus ($\mathtt{ST}$). The set of *simple types* is given by:

$$\mathtt{S} ::= \mathtt{Bool} \mid \sharp[\mathtt{S}_1, \ldots, \mathtt{S}_n]$$

$\sharp[\mathtt{S}_1, \ldots, \mathtt{S}_n]$ is the type of channels that are used for transmitting tuples consisting of values of types $\mathtt{S}_1, \ldots, \mathtt{S}_n$. A type judgment is of the form $\Gamma \vdash_{\mathtt{ST}} P$. A type environment Γ is a mapping from names to simple types, with the constraint that $\mathtt{true}$ and $\mathtt{false}$ are mapped to $\mathtt{Bool}$, and that the links are mapped to channel types. $\Gamma, \widetilde{v}:\widetilde{\mathtt{S}}$ indicates the type environment obtained by extending Γ with the type assignments $\widetilde{v}:\widetilde{\mathtt{S}}$, with the understanding that for all v_i already defined in Γ it should be $\Gamma(v_i) = \mathtt{S}_i$. The standard typing rules are omitted.

Operational Semantics. We use the standard (early) labeled transition relation $P \xrightarrow{\eta} Q$ for the π-calculus. Here, η, called a transition label, is either a silent action τ (which represents an internal communication), an output action $(\nu\widetilde{c})\,\overline{a}[\widetilde{b}]$, or an input action $a[\widetilde{b}]$. We write $\xrightarrow{\tau}{}^*$ for the reflexive and transitive closure of $\xrightarrow{\tau}$; we write $P \xrightarrow{\tau}$ and $P \xrightarrow{\tau}{}^*$ if there is P' s.t. $P \xrightarrow{\tau} P'$ and $P \xrightarrow{\tau}{}^* P'$, respectively.

We extend the above transition relation to a *typed* transition relation, to show how a type environment evolves when a process performs a transition. $\Gamma \vdash_{\mathtt{ST}} P \xrightarrow{\eta} \Gamma' \vdash_{\mathtt{ST}} P'$ holds if: (1) $P \xrightarrow{\eta} P'$; (2) $\Gamma \vdash_{\mathtt{ST}} P$; and (3) if $\eta = \tau$ then $\Gamma = \Gamma'$; otherwise if η is an output $(\nu\widetilde{c})\,\overline{a}[\widetilde{b}]$ or an input $a[\widetilde{b}]$ and $\Gamma(a) = \sharp[\widetilde{\mathtt{S}}]$, then $\Gamma' = \Gamma, \widetilde{b} : \widetilde{\mathtt{S}}$. Note that $\Gamma \vdash_{\mathtt{ST}} P \xrightarrow{\eta} \Gamma' \vdash_{\mathtt{ST}} P'$ implies $\Gamma' \vdash_{\mathtt{ST}} P'$. We write $\Gamma_0 \vdash_{\mathtt{ST}} P_0 \xrightarrow{\eta_1} \cdots \xrightarrow{\eta_k} P_k$ to mean that $\Gamma_0 \vdash_{\mathtt{ST}} P_0$, and there are $\Gamma_1, ..., \Gamma_k$ s.t. for all $i < k$ it holds that $\Gamma_i \vdash_{\mathtt{ST}} P_i \xrightarrow{\eta_{i+1}} \Gamma_{i+1} \vdash_{\mathtt{ST}} P_{i+1}$.

Deadlock-Freedom and Lock-Freedom. A prefix is *at top level* if it is not underneath another input/output prefix or underneath a replication.

Definition 1 (deadlock-freedom). *P is* deadlock-free *if, whenever $P \xrightarrow{\tau}{}^* Q$ and Q has at least one marked prefix at top level, then $Q \xrightarrow{\tau}$.*

Deadlock-freedom indicates only the possibility for the system to evolve further; on the other hand, lock-freedom indicates the eventual success of marked actions at top-level. In the definition of lock-freedom, we track the success of a specific action (as several marked actions may simultaneously appear at top-level) by tagging it. We then demand success for all possible taggings. We call *tagged* a process in which exactly one unguarded and unreplicated prefix—the prefix that we wish to track—has the special annotation $\Box$ (instead of $\circ$ as in the marked prefixes). Transitions of tagged processes are defined as for the untagged ones, except that the labels of transitions emanating from the tagged prefix are also tagged. We call a tagged τ-transition, written $P \xrightarrow{\tau^\Box} P'$, a *success.*

Definition 2 ((weak) lock-freedom). A tagged process P is *successful* if whenever $P \xrightarrow{\tau}{}^* Q$ then $Q \xrightarrow{\tau}{}^* \xrightarrow{\tau^\Box}$. Given an untagged process P, the *tagging*

of P is the set of tagged processes obtained from P by replacing the annotation of a marked prefix at top level with $\Box$. Process P is *(weakly) lock-free* if whenever $P \xrightarrow{\tau}{}^{*} Q$ then all processes in the tagging of Q are successful.

A sequence of transitions $\xrightarrow{\tau}$ or $\xrightarrow{\tau^{\Box}}$ is *full* if it is finite and ends with an irreducible process, or if it is infinite. A sequence of transitions is *strongly fair* if, intuitively, any τ-action that is enabled infinitely often will eventually succeed (see [8, 3] for a formal definition of strong fairness in the π-calculus).

Definition 3 (strong lock-freedom). *P is* strongly lock-free *if whenever $P \xrightarrow{\tau}{}^{*} Q$ then every full and strongly fair transition sequence of each process in the tagging of Q contains the success transition $\xrightarrow{\tau^{\Box}}$.*

Experts in concurrency will easily recognize the difference between weak lock-freedom and strong lock-freedom: Weak lock-freedom combines safety and liveness guarantees, by requiring that a system never reaches a state where a marked action is at top-level, but there is no sequence of τ-actions in which the marked action is consumed. On other hand, strong lock-freedom is a purely liveness property that says that if a marked action is at top-level, the action will eventually be consumed. The example below shows the difference between weak lock-freedom and strong lock-freedom.

Example 1. Consider the following process P:

$$b^{\circ}(\,) \mid \overline{a}[b] \mid *a(y).\,(\nu c)\,(\overline{c}[y] \mid c(y).\,\overline{y}[\,] \mid c(y).\,\overline{a}[y])$$

The rightmost subprocess $(*a(y).\,\cdots)$ receives b on a and either sends a message on b or forwards b to itself non-deterministically. Since c is freshly created everytime b is received from a, the strong fairness does not guarantee that a message is eventually sent on b, and P is therefore *not* strongly lock-free. On the other hand, however, after any number of forwardings, there is a chance for a message to be sent on b; hence, P is weakly lock-free. See Example 3 for another example of a process that is weakly lock-free but not strongly lock-free.

3 Type System for Lock-Freedom

We introduce the type systems for weak/strong lock-freedom. They are obtained by augmenting Kobayashi's type system [9] with hybrid rules appealing to deadlock/termination/confluence analyses. For lack of space, precise definitions are often omitted; see the extended version [13].

3.1 Review of Previous Type System for Lock-Freedom

As mentioned in Section 1, to enable local reasoning about lock-freedom in terms of deadlock and termination analyses, we need to express some contracts between a process and its environment. We reuse the type judgments of Kobayashi's lock-freedom type system [9] to express the contracts. A type judgment is of the form

$\Delta \vdash_{\mathtt{LT}} P$, where Δ is a type environment, which expresses both requirements on the behavior of P, and assumptions on its environment. Ordinary channel types are extended with *usages*, which express how each communication channel is used. For example, $\sharp_{?.!}[\mathtt{Bool}]$ describes a channel that should be first used for receiving a boolean once, and then for sending a boolean once. A channel of type $\sharp_{?}[\sharp_{!}[\mathtt{Bool}]]$ should be first used for receiving a channel once, and then the received channel should be used once for sending a boolean. (! and ? express an output and an input respectively, and "." denotes the sequential composition;)

In order to express both assumptions on the environment (like, "a process can eventually receive a message from its environment") and guarantees by the process (like, "a process will certainly send a message"), each action (! or ?) in a usage is further annotated with *capability levels* and *obligation levels*, which range over the set of natural numbers extended with ∞. If a capability level of an action is finite, then that action is guaranteed to succeed (in other words, its co-action will be provided by the environment) if it becomes ready for execution (i.e., it is at top-level). If an obligation level of an action is finite, then that action must become ready for execution, only by relying on capabilities of smaller levels. For example, the type judgment $a : \sharp_{?_0^\infty}[\mathtt{Bool}], b : \sharp_{!_\infty^1}[\mathtt{Bool}] \vdash_{\mathtt{LT}} P$ means that P has a capability of level 0 to receive a boolean on channel a (but not an obligation to receive it) , and P has an obligation of level 1 to send a boolean on b. (Here, the superscript of ! or ? is the obligation level, and the subscript is the capability level.) Thus, P can be $\overline{b}[\mathtt{true}]$ or $a(x).\overline{b}[x]$, but not $a(x).\mathbf{0}$. Thanks to the abstraction of process behavior by usages, the problem of checking lock-freedom of a process is reduced to that of checking whether the usage of each channel is consistent in the sense that, for each capability of level t, there is a corresponding obligation of level less than or equal to t.

To understand how this kind of judgment can be used for compositional reasoning about lock-freedom, consider the (deadlocked) process $a^{\circ}(x).\overline{b}[x] \mid b^{\circ}(x).\overline{a}[x]$. We have the following judgment for subprocesses:

$$a : \sharp_{?_0^0}[\mathtt{Bool}], b : \sharp_{!_\infty^1}[\mathtt{Bool}] \vdash_{\mathtt{LT}} a^{\circ}(x).\overline{b}[x]$$
$$a : \sharp_{!_\infty^1}[\mathtt{Bool}], b : \sharp_{?_0^0}[\mathtt{Bool}] \vdash_{\mathtt{LT}} b^{\circ}(x).\overline{a}[x]$$

For the entire process, we can simply combine both type environments by combining usages pointwise:

$$a : \sharp_{?_0^0 \mid !_\infty^1}[\mathtt{Bool}], b : \sharp_{!_\infty^1 \mid ?_0^0}[\mathtt{Bool}] \vdash_{\mathtt{LT}} a^{\circ}(x).\overline{b}[x] \mid b^{\circ}(x).\overline{a}[x]$$

Now, the capability level of the input on a (which is 0) is smaller than the obligation level of the corresponding output on a (which is 1); this indicates a failure of assume-guarantee reasoning (the assumption made by the left subprocess is not met by the guarantee by the right subprocess). Thus, we know the process may not be lock-free. On the other hand, if we replace the subprocess in the righthand side with $\overline{a}[\mathtt{true}].b(x)$, then we get:

$$a : \sharp_{?_0^0 \mid !_0^0}[\mathtt{Bool}], b : \sharp_{!_1^1 \mid ?_1^1}[\mathtt{Bool}] \vdash_{\mathtt{LT}} a^{\circ}(x).\overline{b}[x] \mid \overline{a}[\mathtt{true}].b^{\circ}(x)$$

The capability of each action is matched by the obligation of its co-action, which implies that the process is lock-free.

$$U \text{ (usages)} ::= \mathbf{0} \mid ?^{t_1}_{t_2}.U \mid !^{t_1}_{t_2}.U \mid (U_1 \mid U_2) \mid *U \qquad t \text{ (levels)} \in \mathbf{Nat} \cup \{\infty\}$$
$$\mathtt{L} \text{ (usage types)} ::= \mathtt{Bool} \mid \sharp_U[\widetilde{\mathtt{L}}] \qquad \Delta \text{ (type environments)} ::= v_1 : \mathtt{L}_1, \ldots, v_n : \mathtt{L}_n$$

$$\frac{\Delta_1 \vdash_{\mathtt{LT}} P \qquad t_c = \infty \Rightarrow \chi = \bullet}{v : \sharp^{0}_{!^{t_c}}[\widetilde{\mathtt{L}}]; (\Delta_1 \mid \widetilde{w} : \uparrow\widetilde{\mathtt{L}}) \vdash_{\mathtt{LT}} \overline{v}^{\chi}[\widetilde{w}].\, P} \qquad \frac{\Delta, \widetilde{y} : \widetilde{\mathtt{L}} \vdash_{\mathtt{LT}} P \qquad t_c = \infty \Rightarrow \chi = \bullet}{v : \sharp^{0}_{?^{t_c}}[\widetilde{\mathtt{L}}]; \Delta \vdash_{\mathtt{LT}} v^{\chi}(\widetilde{y}).\, P}$$

$$\frac{}{\emptyset \vdash_{\mathtt{LT}} \mathbf{0}} \qquad \frac{\Delta_1 \vdash_{\mathtt{LT}} P_1 \qquad \Delta_2 \vdash_{\mathtt{LT}} P_2}{\Delta_1 \mid \Delta_2 \vdash_{\mathtt{LT}} P_1 \mid P_2} \qquad \frac{\Delta' \vdash_{\mathtt{LT}} P \qquad \Delta \leq \Delta'}{\Delta \vdash_{\mathtt{LT}} P} \qquad \frac{\Delta \vdash_{\mathtt{LT}} P}{*\Delta \vdash_{\mathtt{LT}} *P}$$

$$\frac{\Delta, a : \sharp_U[\widetilde{\mathtt{L}}] \vdash_{\mathtt{LT}} P \qquad rel(U)}{\Delta \vdash_{\mathtt{LT}} (\nu a)\, P} \qquad \frac{\Delta \vdash_{\mathtt{LT}} P \qquad \Delta \vdash_{\mathtt{LT}} Q}{\Delta \mid (v : \mathtt{Bool}) \vdash_{\mathtt{LT}} \mathbf{if}\ v\ \mathbf{then}\ P\ \mathbf{else}\ Q}$$

Fig. 1. Kobayashi's type system for lock-freedom [9]

Figure 1 summarizes the syntax of types, and typing rules of Kobayashi's lock-freedom type system [9]. A type inference algorithm for the type system, which serves as a lock-freedom verification algorithm, is discussed in [9].

3.2 Robust Deadlock-Freedom/Termination/Confluence

To enable local reasoning in the new type system for lock-freedom that we will present, we introduce a strengthening of the notions of deadlock-freedom, termination, and confluence.

A substitution $\sigma = [\widetilde{w}/\widetilde{x}]$ *respects* $\Gamma = \widetilde{v} : \widetilde{\mathtt{S}}$ if $\sigma\Gamma = \widetilde{\sigma v} : \widetilde{\mathtt{S}}$ is well-defined. A substitution σ is *closing for* Γ if σ respects Γ and $\sigma\Gamma$ has no variables. A process is robustly terminating if it cannot diverge, after any sequence of transition that conforms to the base type system ST.

Definition 4 (robust termination). *A process P is* terminating *if there is no infinite internal transition sequence $P \xrightarrow{\tau} P_1 \xrightarrow{\tau} P_2 \xrightarrow{\tau} \cdots$. An (open) process P is* robustly terminating under Γ, *written $\Gamma \models_{\mathtt{RTer}} P$, if $\Gamma \vdash_{ST} P$, and for every closing substitution σ for Γ and for any Q, k, and $\eta_1, \cdots \eta_k$ such that $\sigma\Gamma \vdash_{ST} \sigma P \xrightarrow{\eta_1} \cdots \xrightarrow{\eta_k} Q$, the derivative Q is terminating.*

We say that Δ is closed if $dom(\Delta) \cap \mathcal{V} = \emptyset$. We write $rel(\Delta)$ intuitively to mean that each capability in Δ is guaranteed by a corresponding obligation; and $ob_!(\mathtt{L})$ for the level of the obligation to send a message: again, precise definitions are in the extended version [13].

In the definition of robust deadlock-freedom below, the first condition say that P is deadlock-free when it is executed by itself, and that P either fulfills its obligations or reduces further. The other conditions say that the robust deadlock-freedom is preserved by substitutions and transitions. The relation $\Delta \xrightarrow{\eta} \Delta'$ (see [13] for the definition) expresses the increase/decrease of capabilities/obligations in Δ by the transition η. For example, $a : \sharp^{0}_{?^{\infty}}[\sharp^{1}_{!^{\infty}}[\mathtt{Bool}]] \xrightarrow{a\,[b]} a : \sharp_{\mathbf{0}}[\sharp^{1}_{!^{\infty}}[\mathtt{Bool}]], b : \sharp^{1}_{!^{\infty}}[\mathtt{Bool}]$ holds (where the usage $\mathbf{0}$ indicates that the channel

cannot be used at all). Thus, $a : \sharp^{0}_{?\infty}[\sharp^{1}_{!\infty}[\mathtt{Bool}]] \models_{\mathtt{RD}} P$ means that P will eventually perform an input on a, and then send a boolean on the received channel, unless P at some point diverges.

Definition 5 (robust deadlock-freedom). *The relation $\Delta \models_{\mathtt{RD}} P$ is the largest relation such that $\Delta \models_{\mathtt{RD}} P$ implies all of the following conditions.*

1. *If Δ is closed and $rel(\Delta)$, then: (i) P is deadlock-free; (ii) If $ob_!(\Delta(a)) \neq \infty$, then either $P \xrightarrow{(\nu \widetilde{c})\,\overline{a}[\widetilde{b}]}$ or $P \xrightarrow{\tau}$; and (iii) If $ob_?(\Delta(a)) \neq \infty$ then either $P \xrightarrow{a[\widetilde{b}]}$ or $P \xrightarrow{\tau}$.*
2. *If $[v \mapsto a]\Delta$ is well-defined, then $[v \mapsto a]\Delta \models_{\mathtt{RD}} [v \mapsto a]P$.*
3. *If $P \xrightarrow{\eta} P'$ and, furthermore, when η is an input, all names received are fresh, then $\Delta \xrightarrow{\eta} \Delta'$ and $\Delta' \models_{\mathtt{RD}} P'$ for some Δ'.*

We say that P is robustly deadlock-free *under Δ if $\Delta \models_{\mathtt{RD}} P$ holds.*

Partial confluence means that any τ-transition commutes with any other transitions. To define the partial confluence, we assume that each prefix is uniquely labeled (as in [3]), and extend the transition relation to $\xrightarrow{\eta,S}$ where S is the set of the labels of the prefixes involved in the transition: see [13]. Robust confluence indicates partial confluence after any sequence of transition that conforms to the base type system ST.

Definition 6 (robust confluence). *A process P is* partially confluent*, if whenever $P_1 \xleftarrow{\tau,S_1} P \xrightarrow{\eta,S_2} P_2$, either $\eta = \tau \wedge S_1 = S_2$, or $P_1 \xrightarrow{\eta,S_2} \equiv \xleftarrow{\tau,S_1} P_2$. A process P is* robustly confluent *under Γ, written $\Gamma \models_{\mathtt{RConf}} P$, if $\Gamma \vdash_{ST} P$ and for any closing substitution σ that respects Γ and for any Q, k, and $\eta_1, \cdots \eta_k$ such that $\sigma\Gamma \vdash_{ST} \sigma P \xrightarrow{\eta_1} \cdots \xrightarrow{\eta_k} Q$, the derivative Q is partially confluent.*

While termination, deadlock-freedom, and confluence are frequently discussed in the literature, we are not aware of previous work that defines the robust counterparts above and verification methods for them.

We have proved that robust deadlock-freedom is guaranteed by Kobayashi's type system for deadlock-freedom [11]. In applications of robust deadlock-freedom, it is often the case that the environment Δ needed is of a restricted form, so that $\Delta \models_{\mathtt{RD}} P$ then boils down to the verification of a few simple behavioral properties for which other type systems and model checkers can also be used. For example, if Δ is $a : \sharp^{!0}_{\infty}[\mathtt{Bool}]$, then $\Delta \models_{\mathtt{RD}} P$ only means that P is deadlock-free and P will eventually send a boolean on a unless it diverges. Robust confluence is guaranteed, for instance, by types systems for linear channels [12] and race-freedom [18]; other static analysis methods such as model checking could also be used. Verification of robust termination is discussed in Section 4.

3.3 Hybrid Typing Rules

We now introduce the new rules LT-Hyb (for weak lock-freedom), and SLT-Hyb (for strong lock-freedom).

$$\frac{\Delta \models_{\mathtt{RD}} P \qquad Er(\Delta) \models_{\mathtt{RTer}} P \qquad nocap(\Delta)}{\Delta \vdash_{\mathtt{LT}} P} \quad \text{(LT-Hyb)}$$

$$\frac{\Delta \models_{\mathtt{RD}} P \qquad Er(\Delta) \models_{\mathtt{RTer}} P \qquad Er(\Delta) \models_{\mathtt{RConf}} P \qquad nocap(\Delta)}{\Delta \vdash_{\mathtt{SLT}} P} \quad \text{(SLT-Hyb)}$$

Here, $Er(\Delta)$ is the simple type environment obtained from Δ by removing all usage annotations. The condition $nocap(\Delta)$ holds if, intuitively, Δ describes a process that fulfills its obligations without relying on the environment. As mentioned in Section 1, this is used to avoid circular, unsound, assume-guarantee reasoning. The precise definition of $nocap(\Delta)$, given in [13], is subtle; for nested channel types, the nocap condition depends on whether a channel is used for input or output. For example, $nocap(\sharp^{?0}_{\infty}[\sharp^{!0}_{\infty}[\,]])$ holds but $nocap(\sharp^{!0}_{\infty}[\sharp^{!0}_{\infty}[\,]])$ does not. In the rule for strong lock-freedom, the robust confluence ensures that once a marked prefix is enabled, it cannot be disabled by any other transitions. See Example 3 for a non-trivial example, for which the rule LT-Hyb fails to guarantee strong lock-freedom.

We write $\Delta \vdash_{\mathtt{LT}} P$ if it is derivable by using the typing rules in Section 3.1 and LT-Hyb, and write $\Delta \vdash_{\mathtt{SLT}} P$ if it is derivable by using SLT-Hyb instead of LT-Hyb. The theorem below states the soundness of the type systems. Its proof is non-trivial because of the presence of the hybrid rules; for instance, conditions such as $nocap(\Delta)$ are not preserved by transitions, so in the proof we had to refine and extend the type systems. See the extended version [13].

Theorem 1 (lock-freedom). *If $\emptyset \vdash_{\mathtt{LT}} P$, then P is (weakly) lock-free. If $\emptyset \vdash_{\mathtt{SLT}} P$, then P is strongly lock-free.*

Example 2. Consider the following processes.

$$\begin{aligned} \mathit{Clients} &\stackrel{\text{def}}{=} *(\nu r_1)\,(\overline{\mathit{fact}}^{\circ}[\mathtt{rnd}(), r_1] \mid r_1{}^{\circ}(x).\,\mathbf{0}) \\ \mathit{Server} &\stackrel{\text{def}}{=} (\nu \mathit{fact\_it})\,(*\mathit{fact}(n,r).\,\overline{\mathit{fact\_it}}[n,1,r] \\ &\qquad \mid *\mathit{fact\_it}(n,x,r).\,\mathbf{if}\ n=0\ \mathbf{then}\ \overline{r}[x]\ \mathbf{else}\ \overline{\mathit{fact\_it}}[n-1, x\times n, r]) \end{aligned}$$

The process *Server* creates an internal communication channel *fact_it* (used for computing factorial numbers in a tail-recursive manner), and waits on *fact* for a request $[n, r]$ on computing the factorial of n. Upon receiving a request, it returns the result on r. *Client* consists of infinitely many copies of the process that creates a fresh channel r_1 for receiving a reply, sends a request $[\mathtt{rnd}(), r_1]$ (where $\mathtt{rnd}()$ creates a random number) and then waits for the result on r_1.

Let Δ be $\mathit{fact} : \sharp_{*?0}^{\infty}[\mathtt{Nat}, \sharp_{!1}^{\infty}[\mathtt{Nat}]]$. Then, we have $\Delta \models_{\mathtt{RD}} \mathit{Server}$, $Er(\Delta) \models_{\mathtt{RTer}} \mathit{Server}$, and $Er(\Delta) \models_{\mathtt{RConf}} \mathit{Server}$ with $nocap(\Delta)$. Thus, by using SLT-Hyb, we obtain $\Delta \vdash_{\mathtt{SLT}} \mathit{Server}$. From this judgment and $\mathit{fact} : \sharp_{*!0}^{\infty}[\mathtt{Nat}, \sharp_{!1}^{\infty}[\mathtt{Nat}]] \vdash_{\mathtt{SLT}} \mathit{Clients}$, we obtain: $\emptyset \vdash_{\mathtt{SLT}} (\nu \mathit{fact})\,(\mathit{Server} \mid \mathit{Clients})$. This means that all the clients can eventually receive replies. Note that the whole process diverges (since

there are infinitely many clients), but we can derive strong lock-freedom by local reasoning based on SLT-Hyb.

Example 3. This example shows a binary tree data structure, offering services for inserting and searching natural numbers. Each node of the tree is implemented as a process that has: a state, given by the integer stored in the node and pointers to the left and right subtree and that contain, respectively, smaller and greater integers; channels for the insert and search operations. In Figure 2, G is a generator of new nodes, which can then grow and originate a tree, and where: i and s will be the insertion and search channels; `state` stores the state of the node. Initially the node is a leaf. `TInit` is the initial tree, with an empty state and public channels `insert` and `search` to communicate with the environment. Once received a query for an integer n, the tree lets the request ripple down the nodes, following the order on the integers to find the right path, until either t is found in a node, or the end of the tree is reached. There is parallelism in the system: many requests can be rippling down the tree at the same time; in doing so, requests can even overtake each other.

Let Δ be $\texttt{insert}:\sharp_{*?0}^{\infty}[\texttt{Nat},\sharp_{!1}^{\infty}[\,]],\texttt{search}:\sharp_{*?0}^{\infty}[\texttt{Nat},\sharp_{!1}^{\infty}[\texttt{Bool}]]$. Then, we have:

$$\Delta \models_{\texttt{RD}} \texttt{TInit} \quad Er(\Delta) \models_{\texttt{RTer}} \texttt{TInit} \quad nocap(\Delta)$$

Thus, by using LT-Hyb, we obtain $\Delta \vdash_{\texttt{LT}} \texttt{TInit}$. By applying rules for `LT` to the rest of the system, we get $\Delta \vdash_{\texttt{LT}} \texttt{Sys}$.

Note that SLT-Hyb is not applicable since `TInit` is not robustly confluent (because, when multiple requests arrive simultaneously, there can be a race on the channel `state`). Indeed, the example is NOT strongly lock-free! A search request may never be replied if the request is overtaken by insertion requests so often that the tree grows faster than the search request goes down the tree. See [13] for a strongly lock-free version of binary trees.

4 Types for Robust Termination

For our analysis we need a refinement of the standard termination property, that we call robust termination. *Termination* of a term means that all its reduction sequences are of finite length. *Robust termination* guarantees that termination is maintained when the process interacts with its environment. Termination is strictly weaker than robust termination. Consider for instance the term $P \stackrel{\text{def}}{=} \overline{c}[b] \mid c(x).(\overline{x} \mid *a.\overline{x})$. The process P has one reduction only, and therefore it is terminating. It is indeed typable in the simplest of the type systems in [7]. However, P is not robustly terminating. It can interact with other processes via the input at c and, in doing so, it may receive a resulting in the non-terminating derivative $\overline{c}[b] \mid \overline{a} \mid *a.\overline{a}$.

A number of type systems for termination of mobile processes have appeared in the literature [6, 7, 17, 21]. We have isolated some some abstract conditions which allows us to turn a type system for termination into one for robust termination. For lack of space we refer the reader to [13] for the details.

$$
\begin{aligned}
G \stackrel{\text{def}}{=}\ & *\mathtt{newtree}(i,s).(\nu \mathtt{state})\ \Big(\overline{\mathtt{state}}[\mathbf{leaf}] \\
& \mid *i(n,r).\mathtt{state}(x). \quad \texttt{/*** insertion ***/} \\
& \quad \mathbf{match}\ x\ \mathbf{with}\ \mathbf{leaf}\ \rightarrow \\
& \qquad (\nu \mathtt{left\_i}, \mathtt{left\_s}, \mathtt{right\_i}, \mathtt{right\_s}) \\
& \qquad \Big(\overline{\mathtt{newtree}}[\mathtt{left\_i}, \mathtt{left\_s}] \mid \overline{\mathtt{newtree}}[\mathtt{right\_i}, \mathtt{right\_s}] \\
& \qquad \mid \overline{\mathtt{state}}[\mathbf{node}(n, \mathtt{left\_i}, \mathtt{left\_s}, \mathtt{right\_i}, \mathtt{right\_s})] \mid \overline{r}\Big) \\
& \quad \| \mathbf{node}(n_1, i_l, s_l, i_r, s_r)\ \rightarrow \\
& \qquad \Big(\overline{\mathtt{state}}[x] \mid \mathbf{if}\ n = n_1\ \mathbf{then}\ \overline{r}[]\ \mathbf{else\ if}\ n < n_1\ \mathbf{then}\ \overline{i_l}[n,r]\ \mathbf{else}\ \overline{i_r}[n,r]\Big) \\
& \mid *s(n,r).\mathtt{state}(x).\Big(\overline{\mathtt{state}}[x] \quad \texttt{/*** search ***/} \\
& \quad \mid \mathbf{match}\ x\ \mathbf{with}\ \mathbf{leaf}\ \rightarrow\ \overline{r}[\mathtt{false}] \\
& \qquad \| \ \mathbf{node}(n_1, i_l, s_l, i_r, s_r) \rightarrow \\
& \qquad\ \mathbf{if}\ n_1 = n\ \mathbf{then}\ \overline{r}[\mathtt{true}]\ \mathbf{else\ if}\ n < n_1\ \mathbf{then}\ \overline{s_l}[n,r]\ \mathbf{else}\ \overline{s_r}[n,r]\Big)\Big) \\
\mathtt{TInit} \stackrel{\text{def}}{=}\ & (\nu \mathtt{newtree})\ (G \mid \overline{\mathtt{newtree}}[\mathtt{insert}, \mathtt{search}]) \\
\mathtt{Sys} \stackrel{\text{def}}{=}\ & (\nu \mathtt{insert}, \mathtt{search}) \\
& (\mathtt{TInit} \mid *(\nu r_1)\ (\overline{\mathtt{insert}}^{\circ}[\mathtt{rnd}(), r_1] \mid {r_1}^{\circ}) \mid *(\nu r_2)\ (\overline{\mathtt{search}}^{\circ}[\mathtt{rnd}(), r_2] \mid {r_2}^{\circ}(x)))
\end{aligned}
$$

Fig. 2. A binary tree

5 Implementation

We have implemented the new weak lock-freedom analysis as a feature of TyPiCal Version 1.6.0 [10]. TyPiCal takes as an input a program written in the π-calculus, and marks all input/output prefixes that are guaranteed to succeed.

The original type system for lock-freedom (reviewed in Section 3.1) had been implemented already [11, 9]. A major challenge in the implementation of the new system was to automate verification of the robust termination property. We have modified the type systems of Deng and Sangiorgi [7], so that the resulting systems can guarantee robust termination, and also so to make them more suited for automatic verification (e.g., using heuristic and incomplete algorithms when the original ones were NP-complete). We also integrated them with a termination analysis based on size-change graphs [2]. See the extended version for details.

We have applied the implementation to non-trivial programs (including the examples in Section 3), and verified them fully automatically (without any type annotations). According to benchmark results (shown in [13]), the new components (dealing with termination) run fast; most of the analysis time is spent by the other components (dealing with deadlock- and lock-freedom). For the binary tree (Example 3), the verification time was 5.47 sec., of which the time for robust termination analysis was only 0.02 sec.

6 Related Work

Several type systems for lock-freedom (sometimes referred to by different names) have been already proposed [8, 9, 20, 1, 19, 21]. Our type system substantially

improves the expressiveness of previous type systems; for instance, it can handle non-trivial recursive structures (e.g., the binary trees as in Example 3), and value-dependent behaviors. This is possible through a parameterization that appeals to other analyzers, in particular those for deadlock freedom (so that more powerful analyzers make the lock-freedom type system more powerful too). Another important point is that none of the previous type systems for lock-freedom, except Kobayashi's one [9], has been implemented. In fact, most of the type systems classify channels into a few usage patterns, and prepare separate typing rules for each of the usage patterns. Thus, verification based on those type systems would not be possible without heavy program annotations.

Type systems for deadlock-freedom have been studied extensively. As already mentioned, deadlock-freedom is weaker than lock-freedom, so that those type systems alone cannot be used for lock-freedom analysis. For example, the divergent process obtained by replacing $\overline{fact\_it}[n-1, x \times n, r]$ in Example 2 with $\overline{fact\_it}[n, x \times n, r]$ is deadlock-free.

The idea of reducing verification of lock-freedom to verification of robust termination is a reminiscence of Cook et al.'s work on reducing verification of liveness properties to that of fair termination [5]. The target language of their work is a sequential, imperative language and is quite different from our language, which is concurrent and allows dynamic creation of communication channels and threads. The used techniques are also quite different; they use model checking while we use types.

There are a number of methods for proving termination of programs, and they have been extensively studied in the context of term rewriting systems and sequential programs. The point of parameterizing our type system for lock-freedom by the robust termination property was to reuse those techniques for termination verification, instead of developing a sophisticated type system that can reason about both termination and deadlock within the single type system.

Parameterized, or hybrid, type systems of this kind presented in this paper are fairly rare in the literature, mainly due to the difficulties in combining the analyses. For instance, in Leroy's modular module system [14] a type system for module is presented that is parametric on the type system used for the core language. This is quite different from ours, as the world on which the two type systems operate—modules and core languages—are stratified, hence clearly separated.

References

1. Acciai, L., Boreale, M.: Responsiveness in process calculi. In: Proc. of 11th Annual Asian Computing Science Conference (ASIAN 2006). LNCS, vol. 4435, pp. 136–150. Springer, Heidelberg (2006)
2. Ben-Amram, A.M., Lee, C.S.: Program termination analysis in polynomial time. ACM Trans. Prog. Lang. Syst. 29(1 (Article 5)) (2007)
3. Bidinger, P., Compagnoni, A.B.: Pict correctness revisited. In: Bonsangue, M.M., Johnsen, E.B. (eds.) FMOODS 2007. LNCS, vol. 4468, pp. 206–220. Springer, Heidelberg (2007)

4. Brinksma, E., Rensink, A., Volger, W.: Fair testing. In: Lee, I., Smolka, S.A. (eds.) CONCUR 1995. LNCS, vol. 962, pp. 313–327. Springer, Heidelberg (1995)
5. Cook, B., Gotsman, A., Podelski, A., Rybalchenko, A., Vardi, M.Y.: Proving that programs eventually do something good. In: Proc. of POPL, pp. 265–276 (2007)
6. Demangeon, R., Hirschkoff, D., Kobayashi, N., Sangiorgi, D.: On the complexity of termination inference for processes. In: Barthe, G., Fournet, C. (eds.) Proceedings of TGC 2007. LNCS, vol. 4912, pp. 140–155. Springer, Heidelberg (2008)
7. Deng, Y., Sangiorgi, D.: Ensuring termination by typability. Info. Comput. 204(7), 1045–1082 (2006)
8. Kobayashi, N.: A type system for lock-free processes. Info. Comput. 177, 122–159 (2002)
9. Kobayashi, N.: Type-based information flow analysis for the pi-calculus. Acta Informatica 42(4-5), 291–347 (2005)
10. Kobayashi, N.: TyPiCal: A type-based static analyzer for the pi-calculus, `http://www.kb.ecei.tohoku.ac.jp/~koba/typical/`
11. Kobayashi, N.: A new type system for deadlock-free processes. In: Baier, C., Hermanns, H. (eds.) CONCUR 2006. LNCS, vol. 4137, pp. 233–247. Springer, Heidelberg (2006)
12. Kobayashi, N., Pierce, B.C., Turner, D.N.: Linearity and the pi-calculus. ACM Trans. Prog. Lang. Syst. 21(5), 914–947 (1999)
13. Kobayashi, N., Sangiorgi, D.: A hybrid type system for lock-freedom of mobile processes. An extended version (2008), `http://www.kb.ecei.tohoku.ac.jp/~koba/papers/hybrid.pdf`
14. Leroy, X.: A modular module system. J. Funct. Program. 10(3), 269–303 (2000)
15. Pierce, B.C., Turner, D.N.: Pict: A programming language based on the pi-calculus. In: Plotkin, G., Stirling, C., Tofte, M. (eds.) Proof, Language and Interaction: Essays in Honour of Robin Milner, pp. 455–494. MIT Press, Cambridge (2000)
16. Sangiorgi, D., Walker, D.: The Pi-Calculus: A Theory of Mobile Processes. Cambridge University Press, Cambridge (2001)
17. Sangiorgi, D.: Termination of processes. Math. Struct. Comput. Sci. 16(1), 1–39 (2006)
18. Terauchi, T., Aiken, A.: A Capability Calculus for Concurrency and Determinism. In: Baier, C., Hermanns, H. (eds.) CONCUR 2006. LNCS, vol. 4137, pp. 218–232. Springer, Heidelberg (2006)
19. Sangiorgi, D.: The name discipline of uniform receptiveness. Theor. Comput. Sci. 221(1-2), 457–493 (1999)
20. Yoshida, N.: Type-based liveness guarantee in the presence of nontermination and nondeterminism. Technical Report 2002-20, MSC Technical Report, University of Leicester (April 2002)
21. Yoshida, N., Berger, M., Honda, K.: Strong normalisation in the pi-calculus. Info. Comput. 191(2), 145–202 (2004)