

Analysing Input/Output-Capabilities of Mobile Processes with a Generic Type System^{*}

Barbara König (koenigb@in.tum.de)

Fakultät für Informatik, Technische Universität München

Abstract. We introduce a generic type system (based on Milner's sort system) for the synchronous polyadic π -calculus, allowing us to mechanise the analysis of input/output capabilities of mobile processes. The parameter of the generic type system is a lattice-ordered monoid, the elements of which are used to describe the capabilities of channels with respect to their input/output-capabilities. The type system can be instantiated in order to check process properties such as upper and lower bounds on the number of active channels, confluence and absence of blocked processes.

1 Introduction

For the analysis and verification of processes there are basically two approaches: methods that are complete, but cannot be fully mechanised, and fully automatic methods which are consequently not complete, i.e. not all processes satisfying the property to be checked are recognised.

One promising direction for the latter approach is to use type or sort systems and type inference with rather complex types abstracting from process behaviour. In the last few years there have been several papers presenting such type systems for the polyadic π -calculus and other process calculi, checking e.g. input/output behaviour [15], absence of deadlocks [7], security properties [1, 4], allocation of permissions to names [16] and many others. Types are compositional and thus allow reuse of information obtained in the analysis of smaller subsystems.

One drawback of the type systems mentioned above is the fact that they are specialised to check very specific properties. A much more general approach is a theory of types by Honda [6] which is based on typed algebras and gives a classification of type systems. This theory is very general and it is thus necessary to prove the subject reduction property and the correctness of a type system for every instance. Our paper attempts to fill the gap between the two extremes. We present a generic type system where we can show the subject reduction property for the general case, and by instantiating the type system we are able to analyse specific properties of processes. Despite its generality, our type system can be used to generate existing type systems, or at least subsets of them. With the introduction of residuation (explained below) we can even type some processes which are not typable by comparable type systems.

We concentrate on properties connected to input/output capabilities of processes in the synchronous polyadic π -calculus. In our examples (see section 5) we check properties such as upper and lower bounds on the number of active channels, confluence,

^{*} Research supported by SFB 342 (subproject A3) of the DFG.

absence of blocked input or output prefixes. Determining these capabilities of a process involves counting and we attempt to keep this concept as general as possible by basing the generic type system on commutative monoids. Instantiating a type system mainly involves choosing an appropriate monoid, and monoid elements associated with input and output prefixes (e.g. for counting the number of prefixes with a certain subject).

Instead of giving the precise answer to every question, our type system uses over-approximation (e.g. we can expect results of the form “there are at most two active channels with subject x at any given time”). Hence plain monoids are not sufficient, but we need ordered monoids (so-called lattice-ordered monoids or l-monoids), equipped with a partial order compatible with summation.

There is a huge class of lattice-ordered monoids which are residuated, i.e. some limited form of subtraction can be defined. Residuation can be put to good use in process analysis. Consider, e.g. the process $P = \bar{x}.x.\mathbf{0}$. While P increases the number of occurrences of the output prefix \bar{x} by one, it does not do so for the input prefix x , since we are interested exclusively in the number of prefixes on the outer level (i.e. in prefixes which are currently active) and x can only be reached by a communication with \bar{x} which decreases the number of input prefixes in the environment by one. This decrease can be anticipated when typing P , and is taken into consideration by subtracting one from the number of input prefixes.

The type of a process contains an assignment of names to sorts and a mapping of sorts to strings of sorts (as in [13]), keeping track of channel arities, i.e. if channel x has sort s , and n -ary tuples are communicated via x , then s will be mapped to a string of sorts having length n , being the sorts of the respective channels. Thus, successful typing also guarantees the absence of runtime errors produced by mismatching arities. Furthermore a monoid element is assigned to each sort s . The monoid element is expected to be an upper bound for the capabilities of all channels having sort s .

2 Preliminaries

2.1 The π -Calculus

The π -calculus [12, 13] is an influential paradigm describing communication and mobility of processes. In this paper we will consider the synchronous polyadic π -calculus without choice and matching, and replication is only defined for input prefixes. Its syntax is as follows:

$$P ::= \mathbf{0} \mid (\nu x : s)P \mid P_1 \mid P_2 \mid \bar{x}(\tilde{z}).P \mid x(\tilde{y}).P \mid !x(\tilde{y}).P$$

where s is an element from a fixed set of sorts S and x is taken from a fixed set of names \mathcal{N} . $\tilde{y} = y_1 \dots y_n$ and $\tilde{z} = z_1 \dots z_n$ are abbreviations for sequences with elements from \mathcal{N} . We call $\bar{x}(\tilde{z})$ *output prefix* and $x(\tilde{y})$ *input prefix*.

The set of all free names (i.e. names not bound by either ν or by an input prefix) of a process P is denoted by $fn(P)$. The process obtained by replacing the free names y_i by x_i in P (and avoiding capture) is called $P\{\tilde{x}/\tilde{y}\}$.

Structural congruence is the smallest congruence obeying the rules in the upper part of table 1, and equating processes that can be converted into one another by consistent

renaming of bound names (α -conversion). We use a reduction semantics as for the chemical abstract machine [2] instead of a labelled transition semantics.

| | | |
|-----------------------|--|---|
| Structural Congruence | (C-COM) $P_1 P_2 \equiv P_2 P_1$ | (C-0) $P 0 \equiv P$ |
| | (C-ASS) $P_1 (P_2 P_3) \equiv (P_1 P_2) P_3$ | |
| | (C-RESTR1) $(\nu x:s)(\nu y:t)P \equiv (\nu y:t)(\nu x:s)P$ if $x \neq y$ | |
| | (C-RESTR2) $((\nu x:s)P_1) P_2 \equiv (\nu x:s)(P_1 P_2)$ if $x \notin \text{fn}(P_2)$ | |
| Reduction Rules | (R-COMM) $\bar{x}\langle\bar{z}\rangle.Q \mid x(\bar{y}).P \rightarrow Q \mid P\{\bar{z}/\bar{y}\}$ | |
| | (R-REP) $\bar{x}\langle\bar{z}\rangle.Q \mid !x(\bar{y}).P \rightarrow Q \mid P\{\bar{z}/\bar{y}\} \mid !x(\bar{y}).P$ | |
| | (R-PAR) $\frac{P \rightarrow P'}{P Q \rightarrow P' Q}$ | (R-RESTR) $\frac{P \rightarrow P'}{(\nu x:s)P \rightarrow (\nu x:s)P'}$ |
| | (R-EQU) $\frac{Q \equiv P, P \rightarrow P', P' \equiv Q'}{Q \rightarrow Q'}$ | |

Table 1. operational semantics of the π -calculus

Consider the following processes which we will use as an example in this paper (we omit the final 0):

$$F = c(r).\bar{d}\langle r \rangle.d(a).\bar{c}\langle a \rangle \quad S = d(s).s(h_1, h_2).\bar{d}\langle h_1 \rangle \quad T = \bar{c}\langle h \rangle.c(x) \quad H = \bar{h}\langle i_1, i_2 \rangle$$

There is a forwarder F which receives requests on a channel c , forwards them on a channel d to a server, receives the answer and sends it back on c . The server S receives requests on d , and we assume that these requests come with a name s where the server can get further information. The server obtains this information, processes it and sends the answer back on d (in our example we keep the “processing part” very simple, the server just sends back the first component). Furthermore T is a trigger process, starting the execution of F and receiving the result in the end, and H delivers information to the server.

We can combine the processes F, S, T, H to obtain P as the entire system. If we want F and S to be persistent, we regard P' .

$$P = T \mid H \mid (\nu d: s_d)(F \mid S) \quad P' = T \mid H \mid (\nu d: s_d)(!F \mid !S)$$

A programmer analysing this piece of code might be interested in the following properties: input/output behaviour, upper and lower bound on the number of channels being active, confluence properties and absence of blocked prefixes that never find a communication partner. E.g., examining P will reveal that at any given time every name is used for input and output at most once and that P is therefore confluent.

2.2 Residuated Lattice-ordered Monoids

Lattice-ordered monoids are a well-developed mathematical concept (see e.g. [3]). We are interested in commutative residuated l-monoids in order to represent input/output capabilities.

Definition 1. (Lattice-ordered Monoid)

A commutative lattice-ordered monoid (l-monoid) is a tuple $(I, +, \leq)$ where I is a set, $+$: $I \times I \rightarrow I$ is a binary operation and \leq is a partial order which satisfy:

- $(I, +)$ is a commutative monoid, i.e. $+$ is associative and commutative, and there is a unit 0 with $0 + a = a$ for every monoid element $a \in I$.
- (I, \leq) is a lattice, i.e. \leq is a partial order, where two elements $a, b \in I$ have a join (or least upper bound) $a \vee b$ and a meet (or greatest lower bound) $a \wedge b$.
- I contains a bottom element \perp , the smallest element in I , and a top element \top , the greatest element in I .
- For $a, b, c \in I$: $a + (b \vee c) = (a + b) \vee (a + c)$ and $a + (b \wedge c) = (a + b) \wedge (a + c)$

Any l-monoid $(I, +, \leq)$ is associated with an l-monoid (I, \oplus, \leq) where $a \oplus b = (a + b) \vee a \vee b$ and \perp is the unit. The significance of \oplus can be made clear with the following consideration: monoid elements will be used to label sorts, being an upper bound for the capabilities of channels having this sort. E.g., we assume that a free name x and a bound name y have sort s , indicating that, during reduction, x might replace y . The capabilities of x and y are a respectively b . What capability should be associated with s ? In the presence of positive monoid elements only, $a + b$ is the correct answer. If, however, a is negative, $a + b$ is actually smaller than b and if x has not yet replaced y , the monoid element associated with s underestimates the capabilities of y . Since we use over-approximation the correct sort label is $a \oplus b$.

Definition 2. (Residuated l-monoid) Let $(I, +, \leq)$ be an l-monoid and let $a, b \in I$. The residual $a - b$ is the smallest x (if it exists) such that $a \leq x + b$. I is called residuated if all residuals $a - b$ exist in I for $a, b \in I$.

Example: one residuated l-monoid which we will later use for the analysis of processes is $IO = (\{none, I, O, both\}, \vee, \leq)$ where $none \leq I \leq both$, $none \leq O \leq both$ and the monoid operation is the join, i.e. the l-monoid degenerates to a lattice. A channel name has for example capability O if it is used at most for output and capability $both$ if it may be used for both output and input.

In order to count the number of inputs or outputs we use the residuated l-monoid $\mathbb{Z}^\infty = (\mathbb{Z} \cup \{\infty, -\infty\}, +, \leq)$ with all integers including ∞ and $-\infty$ ($\infty + (-\infty) = -\infty$). Residuation is subtraction for all monoid elements different from ∞ and $-\infty$.

The cartesian product of two l-monoids, e.g. $\mathbb{Z}^\infty \times \mathbb{Z}^\infty$, is also an l-monoid.

We use the following inequations concerning residuated l-monoids: for all elements a, b, c of a residuated l-monoid it holds that

$$\begin{array}{l} a \leq (a - b) + b \quad (a + b) - b \leq a \quad (a + b) - c \leq (a - c) + b \\ (a + b) \vee 0 \leq (a \vee 0) + (b \vee 0) \quad a + b \leq a \oplus b \quad \perp + \perp = \perp \quad \top + \top = \top \end{array}$$

$$\text{And we define: } sig(a) = \begin{cases} \perp & \text{if } a < 0 \\ 0 & \text{if } a = 0 \\ \top & \text{if } a > 0 \\ \text{undefined} & \text{otherwise} \end{cases}$$

3 The Type System and its Properties

We define the notion of types and type assignments which have already been informally introduced in section 1.

Definition 3. (Type Assignment) *Let S be a fixed set of sorts and let $(I, +, \leq)$ be a fixed l-monoid. A type assignment $\Gamma = ob_\Gamma; m_\Gamma; x_1 : \langle s_1/a_1 \rangle, \dots, x_n : \langle s_n/a_n \rangle$ (abbreviated by $ob_\Gamma; m_\Gamma; \tilde{x} : \langle \tilde{s}/\tilde{a} \rangle$) consists of a sort mapping $ob_\Gamma : S \rightarrow S^*$ (mapping sorts to object sorts), a mapping $m_\Gamma : S \rightarrow I$ (assigning a monoid element to every sort) and an assignment of channel names x_i to tuples consisting of a sort s_i and a monoid element a_i .*

We define $sort_\Gamma(x_i) = s_i$ and $\Gamma, y : \langle t/b \rangle$ denotes $ob_\Gamma; m_\Gamma; \tilde{x} : \langle \tilde{s}/\tilde{a} \rangle, y : \langle t/b \rangle$.

Sorts are used to control the mobility of names. That is if $ob_\Gamma(s) = s_1 \dots s_n$, we know that only n -tuples of channel names with sorts s_i are sent or received via a channel with sort s . If a free name x and a bound name y have the same sort, we have to take into account that x may replace y during the reduction. We also use sorts as an intermediate level between names and monoid elements, since with α -conversion it is problematic to assign monoid elements directly to names.

Monoid elements appear in two places: in the range of m_Γ and in the tuples $x : \langle s/a \rangle$. The idea is to sum up the capabilities of x with $+$ in a while x is still free and add a to $m_\Gamma(s)$ with \oplus as soon as x is hidden. We have to use \oplus according to the explanation given in section 2.2. The other possibility would be to immediately add the capabilities to $m_\Gamma(s)$ with \oplus (without storing them in a first), but since $a + b \leq a \oplus b$, this would lead to looser bounds. (It would, however, be possible in the case where we only consider monoid elements greater than or equal to 0, since in this case $+$ and \oplus always coincide.)

In the rest of this paper we use the operations on type assignments given in table 2 (all operations on sequences are conducted pointwise): in (ADD-MON) we add a monoid element a to a type assignment Γ by adding a to $m_\Gamma(s)$ (with \oplus) and leaving everything else unchanged. And $\Gamma^{\langle \tilde{s}, \tilde{a} \rangle}$ denotes $(\dots (\Gamma^{\langle s_1, a_1 \rangle})^{\langle s_2, a_2 \rangle} \dots)^{\langle s_n, a_n \rangle}$.

| |
|---|
| $\text{(ADD-MON)} \quad (ob; m; \tilde{x} : \langle \tilde{s}/\tilde{a} \rangle)^{\langle s, a \rangle} = ob; m'; \tilde{x} : \langle \tilde{s}/\tilde{a} \rangle$ $\text{where } m'(s') = \begin{cases} m(s) \oplus a & \text{if } s = s' \\ m(s) & \text{otherwise} \end{cases}$ $\text{(SUM)} \quad (ob; m; \tilde{x} : \langle \tilde{s}/\tilde{a} \rangle) \oplus (ob; m'; \tilde{x} : \langle \tilde{s}/\tilde{b} \rangle) = ob; m \oplus m'; \tilde{x} : \langle \tilde{s}/\tilde{a} + \tilde{b} \rangle$ $\text{(JOIN)} \quad (ob; m; \tilde{x} : \langle \tilde{s}/\tilde{a} \rangle) \vee b = ob; m; \tilde{x} : \langle \tilde{s}/\tilde{a} \vee b \rangle$ $\text{(ORD)} \quad (ob; m; \tilde{x} : \langle \tilde{s}/\tilde{a} \rangle) \leq (ob; m'; \tilde{x} : \langle \tilde{s}/\tilde{b} \rangle) \iff m \leq m' \text{ and } \tilde{a} \leq \tilde{b}$ $\text{(REMOVE)} \quad \text{If } \Gamma = \Delta, x : \langle s/a \rangle, \text{ then } \Gamma \setminus \{x\} = \Delta$ |
|---|

Table 2. Operations on type assignments

Summation $\Gamma \otimes \Gamma'$ (SUM) is defined for type assignments which contain the same names (having identical sorts) and which satisfy $ob_\Gamma = ob_{\Gamma'}$ (i.e. they have the same sort structure). In this case $m \oplus m'$ and $\tilde{a} + \tilde{b}$ denote the pointwise summation. The summation on type assignments has a counterpart (denoted by the same symbol) in Honda's work [6].

In (JOIN) a pointwise join with every monoid element assigned to a channel name and the monoid element b is defined. And finally we need a partial order on type assignments (ORD) and an operation removing an assumption on a name x from a type assignment (REMOVE).

We are now ready to define the rules of the type system (see table 3). *out* and *in* are fixed monoid elements (where *in* must be comparable to 0) representing the capabilities of output respectively input prefixes.

| | | |
|--|---|---|
| $\frac{\Gamma \vdash P, \Gamma \leq \Delta}{\Delta \vdash P} \text{(T-}\leq\text{)}$ | $\Gamma \vee 0 \vdash \mathbf{0} \text{(T-NIL)}$ | $\frac{\Gamma_1 \vdash P_1, \Gamma_2 \vdash P_2}{\Gamma_1 \otimes \Gamma_2 \vdash P_1 \mid P_2} \text{(T-PAR)}$ |
| $\frac{\Gamma, x: \langle s/a \rangle, \tilde{z}: \langle \tilde{t}, \tilde{b} \rangle \vdash P}{(\Gamma, \tilde{z}: \langle \tilde{t}, \tilde{b} \rangle) \vee 0, x: \langle s/(a - in) \vee 0 + out \rangle \vdash \bar{x}(\tilde{z}).P} \text{(T-OUT)} \quad \text{if } ob_\Gamma(s) = \tilde{t}$ | | |
| $\frac{\Gamma, x: \langle s/a \rangle, \tilde{y}: \langle \tilde{t}/\tilde{b} \rangle \vdash P}{(\Gamma \vee 0, x: \langle s/(a - out) \vee 0 + in \rangle)^{\langle \tilde{t}, \tilde{b} \rangle} \vdash x(\tilde{y}).P} \text{(T-IN)} \quad \text{if } ob_\Gamma(s) = \tilde{t}$ | | |
| $\frac{\Gamma \setminus \{x\}, x: \langle s/a \rangle \vdash P}{\Gamma^{(x,a)} \vdash (\nu x: s)P} \text{(T-RESTR)}$ | $\frac{\Gamma \vdash x(\tilde{y}).P}{\Delta, x: \langle s/a + sig(in) \rangle \vdash !x(\tilde{y}).P} \text{(T-REP)}$ if $\Gamma \otimes \Gamma \leq \Gamma$ and $\Gamma = \Delta, x: \langle t/a \rangle$ | |

Table 3. Rules of the type system

The intuitive meaning of the rules is as follows:

- (T- \leq) We can always over-approximate the capabilities of a process.
 - (T-NIL) The nil process can have an arbitrary type assignment, provided the monoid elements of the free names are greater than 0.
 - (T-PAR) The parallel composition of two processes can be typed by adding their respective type assignments.
 - (T-OUT) First we subtract *in* from the monoid element a associated with the subject x of the output prefix, since the emergence of P means the removal of an input prefix with subject x somewhere else in the environment. We then take the join of all monoid elements and 0, since we only consider capabilities on the outer level of processes and thus we only consider future influence by positive capabilities, but not by negative ones (since we are doing over-approximation). In the end *out* is added to the monoid element associated with x .
- Furthermore we have to check that the sort structure is correct, i.e. since z_1, \dots, z_n are communicated via x , the string of their sorts must be the object sort of the sort of x .

- (T-IN) As described for (T-OUT), we subtract *out*, take the join with 0 and then add *in*. Furthermore we check the correctness of the sort structure as above.
 Since, in this case, y_1, \dots, y_n are bound by the input prefix, all assumptions on these names are removed from the type assignment, and their monoid elements are added to the rest of the type assignment with \oplus .
- (T-RESTR) If a name is hidden, we remove the assumption on it, but retain information on its capabilities by adding its monoid element to the type assignment and by keeping the sort.
- (T-REP) In this rule we have to make sure that a replicated process has a type assignment which is either idempotent or gets smaller when added to itself. This can be achieved if Γ contains only negative or idempotent monoid elements.
 And furthermore, since we know that infinitely many copies of the input prefix with subject x are available, we add \perp , \top or 0, according to the value of *in*.

The type system satisfies the following substitution lemma, which is central for proving the subject reduction property:

Lemma 1. (Substitution) *Let $x \neq y$ be two names.*

If $\Gamma, x: \langle s/a \rangle, y: \langle s/b \rangle \vdash P$, then $\Gamma, x: \langle s/a + b \rangle \vdash P\{x/y\}$.

Remark: the proofs can be found in the extended version of this paper [11].

The types defined in table 3 are not yet invariant under reduction: rather than Γ , a modified type assignment $\bar{\Gamma}$ satisfies the subject reduction property.

Let $\Gamma = ob; m; \tilde{x}: \langle \tilde{s}/\tilde{a} \rangle$ and define $\bar{\Gamma} = (ob; m; \tilde{x}: \langle \tilde{s}/\tilde{0} \rangle)^{\langle \tilde{s}, \tilde{a} \rangle}$. That is we add all monoid elements of the remaining free names to m with \oplus . Constructing $\bar{\Gamma}$ directly during the typing process does not seem to be possible, since we first have to sum up monoid elements with $+$ and then add them to the type tree with \oplus the moment they are hidden.

Proposition 1. (Subject Reduction Property) *If $P \equiv Q$ and $\Gamma \vdash P$, then it holds also that $\Gamma \vdash Q$. And if $P \rightarrow P'$ and $\Gamma \vdash P$ then there exists a type assignment Γ' such that $\Gamma' \vdash P'$ and $\bar{\Gamma}' \leq \bar{\Gamma}$.*

4 Using the Type System for Process Analysis

As in other type systems for mobile processes, a type guarantees absence of runtime errors which may appear in the form of arity mismatches in the communication rules (R-COMM) and (R-REP), but it also enables us to perform more detailed process analysis.

4.1 Process Capabilities

The aim of this paper is to construct type systems yielding useful results for the analysis and verification of parallel processes. In our case the generic type system gives information concerning structural properties of a process, especially concerning its input and

output capabilities. We will now formally define the connection between the type of a process and its capabilities.

Let P be a process and let x be a free name occurring in P . We define P 's *capability* wrt. x by adding the following monoid elements: for every use of x as an output port we add *out* and for every use of x as an input port we add *in*. Notice that we do not continue summation after prefixes (see table 4).

| | |
|---|--|
| $C_x(\mathbf{0}) = 0 \quad C_x(P \mid Q) = C_x(P) + C_x(Q)$ | |
| $C_x(\bar{z}\langle\tilde{y}\rangle.P) = \begin{cases} \textit{out} & \text{if } x = z \\ 0 & \text{otherwise} \end{cases}$ | $C_x(z\langle\tilde{y}\rangle.P) = \begin{cases} \textit{in} & \text{if } x = z \\ 0 & \text{otherwise} \end{cases}$ |
| $C_x(!z\langle\tilde{z}\rangle.P) = \begin{cases} \textit{sig(in)} & \text{if } x = z \\ 0 & \text{otherwise} \end{cases}$ | $C_x((\nu y: s)P) = \begin{cases} C_x(P) & \text{if } x \neq y \\ 0 & \text{otherwise} \end{cases}$ |

Table 4. Determining the capabilities of a process

Proposition 2. *If $\Gamma \vdash P$, $P \rightarrow P'$ and x is a free name of P it follows that $C_x(P') \leq m_{\overline{\Gamma}}(\textit{sort}_{\Gamma}(x))$, i.e. we determine the sort of x and look up the corresponding monoid element in $\overline{\Gamma}$. And if P contains a subexpression $(\nu y: s')Q$ it follows that the capabilities of y will never exceed $m_{\overline{\Gamma}}(s')$.*

4.2 Type Inference

In order to support our claim that the type system is useful for the automated analysis of processes, we roughly sketch a type inference algorithm, determining the smallest type (in the \leq relation defined in section 3) of a process P , provided P has a type. In order to make sure that a smallest type exists, we impose the following condition on the l-monoid: for every monoid element $a \in I$ there is a smallest element a' such that $a \leq a'$ and $a' + a' \leq a'$ (the same must be true for the operation \oplus)¹.

The algorithm proceeds in two steps:


- In the first step we determine the assignment of sorts to names and the mapping ob_{Γ} . This may be done by representing ob as a graph and refining ob step by step by collapsing graph nodes every time we encounter a constraint of the form $ob(s) = \bar{s}$. Or we can use the sort inference algorithm by Simon Gay [5].
- In the second step we compute the monoid elements by induction on the structure of P . In this case the typing rules are already very constructive, the main complication arises from typing rule (T-REP). Here we require that the monoid I satisfies the condition stated above. So (because of rule (T- \leq)) we may replace every monoid element a with its corresponding a' in the type assignment that we have derived so far.

A straightforward implementation of the algorithm has a runtime complexity quadratic in the size of P . Ameliorations are certainly possible by using efficient algorithms for unification and by finding an intelligent strategy for computing the monoid elements.

¹ Every l-monoid useful for process analysis that we have come across so far satisfies this condition. In the case of \mathbb{Z}^{∞} , a' is ∞ for positive a and a itself for all other elements.

5 Examples

We now get back to the two example processes P and P' introduced in section 2.1 and type them with several instantiations of our type system, and thereby show how to mechanise process analysis in these cases.



We use the algorithm presented in section 4.2 to derive a type assignment Γ for P and P' and in the first step obtain a sort structure ob_Γ as shown in the figure to the left (ob_Γ is the same for P and P'). If there is an arrow labelled ob_i from sort s to sort t , then t is the i -th element of the sequence $ob_\Gamma(s)$. The assignment of names (in brackets we give the bound names) to sorts is:

$$c: s_c \quad (d: s_d) \quad h, i_1(, r, a, s, h_1): s_1 \quad i_2(, h_2): s_2$$

In the second step the monoid elements $m_{\overline{\Gamma}}(s)$ are computed (see below) in order to give an upper bound for all names having sort s .

5.1 Input/Output Behaviour of Channels

One simple application of our type system is to check whether channels are used for input, output or for both. We use the monoid IO (with elements $none$, O —“output only”, I —“input only” and $both$) introduced in section 2.2. We set $in = I$, $out = O$.

For both processes P and P' we obtain the same type assignments with monoid elements shown in table 5 (row 1), i.e. i_2, h_2 are used neither for input nor output while all other names may be used for both. Note that, because of residuation, typing F alone yields capability I for name c , but no output capability. c acquires output capability only if communication with the environment is taking place.

This type system is similar to the one in [15] (apart from the fact that we consider types as a representation of process capabilities, rather than constraints on the environment), our type system however lacks a concept of co- and contravariance and thus our bounds are less tight.

5.2 Upper Bounds on the Number of Active Channels

We attempt to define a type system, similar to the one presented in [8] for our framework, i.e. we want to check how often a channel is used either for input or output.

We use the l-monoid $\mathbb{Z}^\infty \times \mathbb{Z}^\infty$ (cartesian product of the set of integers with ∞ and $-\infty$) introduced in section 2.2. The first component represents the number of active output prefixes (with a fixed subject) and the second component represents the number of active input prefixes.

We set $out = (1, 0)$, $in = (0, 1)$, and typing the processes P and P' yields the results given in table 5 (rows 2 & 3). Since for P the upper bound is always $(1, 1)$ or smaller we can conclude that there is at most one active input port and one active output port for any given subject at a time. For P' we can guarantee that, e.g. \bar{c} always occurs at most once as an output prefix, although it occurs under a replication (see monoid element $m_{\overline{\Gamma}}(s_c)$).

| Property to be checked | $m_{\overline{F}}(s_d)$ | $m_{\overline{F}}(s_c)$ | $m_{\overline{F}}(s_1)$ | $m_{\overline{F}}(s_2)$ |
|--|--------------------------|--------------------------|-------------------------|-------------------------|
| 1 Input/Output behaviour of P and P' | <i>both</i> | <i>both</i> | <i>both</i> | <i>none</i> |
| 2 Upper bounds on active channels in P | (1, 1) | (1, 1) | (1, 1) | (0, 0) |
| 3 Upper bounds on active channels in P' | (∞ , ∞) | (1, ∞) | (1, ∞) | (0, 0) |
| 4 Lower bounds on active channels in P | (-1, 0) | (-1, 0) | (-1, -1) | (0, 0) |
| 5 Lower bounds on active channels in P' | ($-\infty$, ∞) | ($-\infty$, ∞) | ($-\infty$, -1) | (0, 0) |
| 6 Avoiding blocked output prefixes in P' | (∞ , ∞) | (1, ∞) | (1, -1) | (0, 0) |

Table 5. Resulting monoid elements for different instantiations of the generic type system

5.3 Confluence

As in [8] we can use upper bounds on the number of active channels to guarantee confluence for π -calculus processes (see also [14]). Let Q be a process, and for every name x in Q which is either free or bound by the scope operator ν it holds that its capabilities never exceed (1, 1). Then we can guarantee that every channel (also bound channels) occurs at most once at any given time as active input and output prefix, and we have non-overlapping redexes in (R-COMM). Thus we can conclude that if $Q \rightarrow^* Q'$, $Q' \rightarrow Q_1$ and $Q' \rightarrow Q_2$, then either $Q_1 \equiv Q_2$ or there is a process Q_3 such that $Q_1 \rightarrow Q_3$ and $Q_2 \rightarrow Q_3$.

Row 2 in table 5 provides upper bound (1, 1) for all capabilities in P . So we can state that P is confluent. Note that the same process would not be recognised as confluent by the type system in [8].

5.4 Lower Bounds on the Number of Active Channels

The type system is not limited to statements of the form: “there *at most* n active channels”, we can also guarantee that there are *at least* m active channels. In order to achieve this, we use the type system above and just invert the partial order, i.e. we take \geq instead of \leq , *out* and *in* remain unchanged. This means also that the join \vee in the new partial order is now the meet \wedge of the original partial order. Typing P does not give us much information, since we cannot guarantee that there are at least $m > 0$ prefixes active at any given time (see table 5, row 4) for any channel. In fact, some lower bounds are even (-1) stating that the respective channel removes input (or output) prefixes instead of making them available. In this case $P \rightarrow^* \mathbf{0}$ which means that no lower bounds can be guaranteed.

Typing P' yields the monoid elements given in table 5 (row 5) which states that input prefixes with subjects c, d are available infinitely often.

5.5 Avoiding Blocked Prefixes

Another interesting feature is to avoid blocked prefixes, i.e. prefixes which are waiting for a non-existing communication partner. We will first define—with the help of a lattice-ordered monoid—what it means for an output prefix to be blocked.

We take $\mathbb{Z}^\infty \times \mathbb{Z}^\infty$ as an l-monoid and define a new partial order: $(i, j) \sqsubseteq (i', j')$ iff $i \leq i'$ and $j \geq j'$. The first component represents the number of output prefixes and the

second the number of input prefixes of the same subject. $out = (1, 0)$ and $in = (0, 1)$. We say a name x is *blocking* in P , if $P \rightarrow^* P'$, $C_x(P') \sqsupseteq (1, 0)$ (i.e. there is at least one output prefix with subject x and no corresponding input prefixes) and for all P'' with $P' \rightarrow^* P''$ it follows that $C_x(P'') \sqsupseteq (1, 0)$ (no communication with x will ever take place).

We can, e.g., avoid this situation, by demanding that it is always the case that $C_x(P') = (a, b)$ and either $a \leq 0$ or $b \geq 1$ (i.e. $(a, b) \not\sqsupseteq (1, 0)$). We take the l-monoid and out, in introduced above. This type system can be obtained by composing a type system establishing upper bounds for input prefixes and one establishing lower bounds for output prefixes. In this way we find out that all output prefixes with subjects c and d are non-blocking in P' .

6 Conclusion and Future Work

This work has a similar aim as that of Honda [6], in that it attempts to describe a general framework for process analysis using type systems. We concentrate on a more specialised but still generic type system, which enables us to prove the subject reduction property for the general case. We have shown that, despite its generality, the type system can be instantiated in order to yield type systems related to existing ones. We have also shown how to parameterise type systems and what kind of parameters are feasible (in our case an l-monoid).

Another type system that has close connections to ours is the linear type system by Kobayashi, Pierce and Turner [8], since it also involves the typing of input/output capabilities of processes. Apart from the more general approach, one new feature of our type system is the introduction of residuation which allows us to recognise the process P in section 5 as confluent, in contrast to the type system in [8]. In some other cases however, our bounds are less tight. The central aim of [8] is to introduce a new notion of barbed congruence by reducing the possible contexts of a process. This question has not been addressed in this paper, it is an interesting direction for future work. For a more detailed discussion of the relation between the two type systems see the full version of this paper [11].

Our type system was derived from a type system for a graph-based process calculus with graphs as types, which make it easier to add additional behaviour information and which have a clear correspondence to associated monoid elements (via morphisms and categorical functors) [9]. A graph-based type system with lattices instead of monoids was presented in [10]. For lattices or positive cones of l-monoids, generic type systems are much easier to present. The main complication arises from non-positive elements and residuation.

Inspiration for this work came from papers deriving information on the behaviour of a process by inspecting its input/output capabilities, such as [15, 14, 8]. In order to conduct process analysis concerning more complex properties (as was done e.g. in [7, 4]) it is necessary to use type systems assigning behaviour information (i.e. monoid elements in our case) not only to single channels, but rather to tuples of channels or other more complex structures. This normally results in a semi-additive type system, in the terminology of Honda [6], while our present type system is strictly additive. In order

to extend this type system, a first solution would be to allow monoid labels for n -ary tuples of names. Another idea is to integrate it into the categorical framework presented in [10], which would allow us to specify very general behaviour descriptions.

We believe that generic type systems can be developed into tools suitable for fast debugging and the analysis of concurrent programs. The next step is to apply the type system presented here to “real-life examples” and to more realistic programming languages.

Acknowledgements: I would like to thank the anonymous referees for their helpful comments, especially for the suggestion to use a sort system instead of type trees.

References

1. Martín Abadi. Secrecy by typing in security protocols. In *Theoretical Aspects of Computer Software*, pages 611–638. Springer-Verlag, 1997.
2. Gérard Berry and Gérard Boudol. The chemical abstract machine. *Theoretical Computer Science*, 96:217–248, 1992.
3. G. Birkhoff. *Lattice Theory*. American Mathematical Society, third edition, 1967.
4. Chiara Bodei, Pierpaolo Degano, Flemming Nielson, and Hanne Riis Nielson. Control flow analysis for the pi-calculus. In *Proc. of CONCUR '98*, pages 84–98. Springer-Verlag, 1998. LNCS 1466.
5. Simon J. Gay. A sort inference algorithm for the polyadic π -calculus. In *Proc. of POPL '93*. ACM, 1993.
6. Kohei Honda. Composing processes. In *Proc. of POPL'96*, pages 344–357. ACM, 1996.
7. Naoki Kobayashi. A partially deadlock-free typed process calculus. In *Proc. of LICS '97*, pages 128–139. IEEE, Computer Society Press, 1997.
8. Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. Linearity and the pi-calculus. In *Proc. of POPL'96*, pages 358–371. ACM, 1996.
9. Barbara König. *Description and Verification of Mobile Processes with Graph Rewriting Techniques*. PhD thesis, Technische Universität München, 1999.
10. Barbara König. Generating type systems for process graphs. In *Proc. of CONCUR '99*, pages 352–367. Springer-Verlag, 1999. LNCS 1664.
11. Barbara König. Analysing input/output-capabilities of mobile processes with a generic type system (extended version). Technical Report TUM-I0009, Technische Universität München, 2000.
12. R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes. *Information and Computation*, 100(1):1–77, 1992.
13. Robin Milner. The polyadic π -calculus: a tutorial. In F. L. Hamer, W. Brauer, and H. Schwichtenberg, editors, *Logic and Algebra of Specification*. Springer-Verlag, Heidelberg, 1993.
14. Uwe Nestmann and Martin Steffen. Typing confluence. In *Second International ERCIM Workshop on Formal Methods in Industrial Critical Systems (Cesena, Italy, July 4–5, 1997)*, pages 77–101, 1997.
15. Benjamin Pierce and Davide Sangiorgi. Typing and subtyping for mobile processes. In *Proc. of LICS '93*, pages 376–385, 1993.
16. James Riely and Matthew Hennessy. Distributed processes and location failures. In *Proc. of ICALP'97*, pages 471–481. Springer-Verlag, 1997. LNCS 1256.