

Instant Polymorphic Type Systems for Mobile Process Calculi: Just Add Reduction Rules and Close*

Henning Makhholm and J. B. Wells

Heriot-Watt University

Abstract. Many different *mobile process calculi* have been invented, and for each some number of type systems has been developed. Soundness and other properties must be proved separately for each calculus and type system. We present the *generic* polymorphic type system Poly★ which works for a wide range of mobile process calculi, including the π -calculus and Mobile Ambients. For any calculus satisfying some general syntactic conditions, well-formedness rules for types are derived automatically from the reduction rules and Poly★ works otherwise unchanged. The derived type system is automatically sound (i.e., has subject reduction) and often more precise than previous type systems for the calculus, due to Poly★’s *spatial polymorphism*. We present an implemented type inference algorithm for Poly★ which automatically constructs a typing given a set of reduction rules and a term to be typed. The generated typings are *principal* with respect to certain natural type shape constraints.

1 Introduction

Many calculi that intend to capture the essence of *mobile* and *distributed* computing have been invented. The most well-known of these are probably the π -calculus [18] and the *Mobile Ambients* calculus (MA) by Cardelli and Gordon [8], but they have inspired the subsequent development of a wide variety of variants and alternatives, which are variously argued to be easier to program in or reason about, and/or closer to some operational intuition about how programs in a mobile, distributed setting can be implemented. The field stays productive; new calculi are still being proposed and there is not a clear consensus about what should be considered *the* fundamental mobility calculus.

The majority of these calculi share the basic architecture of MA: They borrow from the π -calculus the syntactic machinery for talking about sets of parallel, communicating processes, plus its primitive operator ν for generating unique names. To this they add some kind of *spatial structure*, usually in the form of a tree of locations where processes can reside. The tree can generally evolve under program control as the processes in it execute; the different calculi provide quite different primitives for mutating it. Mobility calculi also provide for *communication* between processes that are near each other, usually modelled on the communication primitive of the π -calculus, but again with variations and often extended with the possibility to communicate “capabilities”, “paths”, or other restricted pieces of process syntax, rather than just names.

Most process calculi have an associated *type system*, either one that was designed

* Supported by EC FP5/IST/FET grant IST-2001-33477 “DART”, and Sun Microsystems equipment grant EDUD-7826-990410-US.

with the calculus from the beginning, or one that was retrofitted later. These type systems are closely tied to a specific calculus and its particular primitives. Once a type system has been designed and its properties (such as soundness or the applicability of a particular type inference algorithm) have been proved, it is in general not trivial to see whether these properties will survive changes to the calculus.

1.1 A generic type system. In contrast, this paper presents the *generic* type system Poly★ which works for a wide range of mobile process calculi. To use Poly★, one simply instantiates it with the reduction rules that specify the semantics of the target calculus’s primitives. From this, a set of provably sound well-formedness rules for types can be *mechanically* produced, guaranteeing that types that satisfy the rules are sound with respect to the reduction rules, i.e., subject reduction holds. The reduction rules can also be used to guide an automatic *type inference* algorithm for the instantiated type system. The inference algorithm produces a type which is *principal* with respect to certain natural constraints on the shape of types. Our implementation offers several possibilities for tuning the *precision* of the type system it implements, but the use of these is optional — it will always produce a typing even when given only the raw reduction rules of the target calculus.

For this to work, the target calculus must make one small concession to Poly★, namely that its *syntax* is sufficiently regular that the implementation can make sense of its terms and reduction rules. We define a **metacalculus** Meta★ which gives a syntax that is easy to parse and manipulate, while flexible enough that many calculi can be viewed as instances of it without deviating much from their native notations. Meta★ does not include any fixed *semantics* except for the usual semantics of parallelism and name restriction, but instead provides a common notion of substitution and a notation for rewriting rules that fits how semantics for process calculi are usually defined.

1.2 Poly★’s relation to other reasoning principles. A long-term goal of Poly★ is to make it possible to view many previously existing mobility calculi type systems as instances of Poly★, at least with regards to using the type system to statically verify that certain bad behaviours do not occur. The design we present here does not quite reach that point; there are features of existing type systems that we have not yet incorporated in Poly★. We believe it will be particularly important to express some form of the *single-threaded* locations introduced by the original type system for Safe Ambients [16].

We do not expect actual programming environments based on mobility calculi to use the fully general Poly★ formalism as their type discipline. Considerations of performance and integration will generally dictate that production environments instead use hand-crafted specialised type systems for the language they support, though *ideas* from Poly★ may well be employed.

A generic implementation of Poly★, such as the one we present here, should be a valuable tool for *exploring the design space* for mobility calculi in general. It will make it easy to change some aspect of one’s rewriting rules, try to analyse some terms, and see which effect the new rules have on, for example, the interference-control properties of one’s calculus. At the same time, our Poly★ implementation makes it easy to experiment with exactly how strong a type system one wants to use in practice, because our

implementation supports tuning the precision of types in very small steps.¹

Like every nontrivial type system with an inference algorithm, Poly \star can be used as a *control/data flow analysis* to provide the substratum for more specialised automatic *program analyses*.² (Readers who are uncomfortable about applying the term “type system” to Poly \star are invited to think “versatile program analysis framework” each time we write “type system”.) However, we have no pretension of subsuming all other analysis techniques for mobility or process calculi in general. Process calculi have provided the setting for many advanced techniques for reasoning about, for example, behavioural equivalence of processes. Poly \star does not claim to compete with these.

1.3 Spatial polymorphism. The Poly \star type system descends from (but significantly generalises and enhances) our earlier work [2] on PolyA, a polymorphic type system specific to Mobile Ambients. It inherits from PolyA the difference from most other type systems for mobility calculi that the emphasis is on types for *processes* rather than types for (ambient or channel) *names*.³ In fact, types for names have completely vanished: A name has no intrinsic type of its own, but is distinguished solely by the way it can be used to form processes.

Poly \star works by approximating the set of terms a given term can possibly evolve to using the given reduction rules. Its central concept is that of a **shape predicate** which is an automaton that describes a set of process terms. Shape predicates that satisfy certain well-formedness rules are **types**. These rules are derived from the reduction rules of the target calculus and guarantee that the set of terms denoted by a type is closed under the reduction relation, i.e., *subject reduction* holds.

This design gives rise to a new (with PolyA) form of polymorphism that we call **spatial polymorphism**. The type of a process may depend on where in the spatial structure it is found. When the process moves, it may come under influence of another part of the type which allows more reductions. For example, consider a calculus which has the single reduction rule $a[\text{eat } b \mid P] \mid b[Q] \leftrightarrow a[P \mid b[Q]]$. In this calculus, the term $x[\text{eat } z1 \mid \text{eat } z2] \mid y1[\text{eat } x \mid z1[0]] \mid y2[\text{eat } x \mid z2[0]]$ has a Poly \star type, shown in Figure 1, that says that $x[\]$ may contain $z1[\]$ when it is inside $y1[\]$, or $z2[\]$ when it is inside $y2[\]$, but can contain neither when it is found at the top level of the term. Thus Poly \star can prove that the term satisfies the safety policy that $z1$ and $z2$ may never be found side by side. To our knowledge, type systems based on earlier paradigms cannot do this.

With spatial polymorphism, *movement* is what triggers the generation of a polymorphic variant of the original analysis for a piece of code. This is different from, and orthogonal to, the more conventional form of name-parametric polymorphism in the polymorphic π -calculus [24], where it is *communication* that causes polymorphic variants to be created. Poly \star does not support the latter form of polymorphism (and neither

¹ These fine tuning options are omitted from this paper due to lack of space, but they are described in detail in the implementation’s documentation.

² Indeed it is well known [20, 3] that the difference between an advanced flow analysis and an advanced type system is often just a question of different perspectives on the same underlying machinery. The presentation of Poly \star is closer to the data-flow viewpoint than is common for type systems, though this of course does not make Poly \star any less a type system.

³ There are a number of type systems for process calculi *without* an explicit notion of locations which assign types to processes rather than names, for example [4, 14, 27, 11].

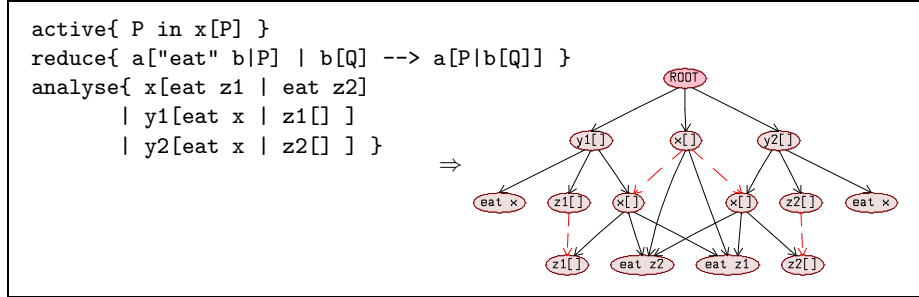


Fig. 1. Input to our type inference implementation for analysing a term in the fictional “eat calculus”, and the inferred type graph as rendered by the VCG graph layout tool [23] (the dashed lines represent subtyping edges).

does any type system for a mobility calculus with explicit locations that we are aware of); we leave it to future work to try to combine the strengths of these two principles.

1.4 Notation and preliminaries. \boxed{X} , where X is any metavariable symbol, stands for the set that X ranges over. $\mathcal{P}_{\text{fin}}(A)$ is the set of finite subsets of the set A . $A \xrightarrow{\text{fin}} B$ is the set of finite partial maps from A to B . $\text{Dom } f$ is the set of x ’s such that $f(x)$ is defined. In contexts where a sequence of similar objects are indexed with indexes up to k , it is to be understood that k can be any integer ≥ 0 . Thus, if the first index is 0, the sequence must have at least one element; sequences indexed from 1 to k may be empty.

2 Meta \star : A metacalculus of concurrent processes

The metacalculus Meta \star defined in this section is the *syntactic* setting for Poly \star . Its role is to let us present the generic properties of Poly \star without resorting to handwaving. Though we define a reduction relation and some other formal properties for Meta \star , these exist solely as support for making formal statements about Poly \star . We do not intend Meta \star to take the place of any existing calculi or frameworks.

As a first approximation, Meta \star is a “syntax without a semantics” except that it does give semantics to a few basic constructs, e.g., process replication and substitution.

2.1 Terms. Figure 2 shows the syntax of process terms in Meta \star . The trivial process 0, parallel composition of processes $P \mid Q$, process replication $!P$, and name restriction $\nu(x).P$ are all well-known from most process calculi, including π -calculus and MA. They are given their usual behaviour by the structural congruence relation \equiv .

Meta \star amalgamates all other process constructors into the general concept of a **form**. Forms have no intrinsic meaning until a set of reduction rules give them one. Examples of forms include the communication actions “ $x\langle y \rangle$ ” and “ $x(y)$ ” from the π -calculus, the movement capabilities “in x ”, “out x ”, and “open x ” from Mobile Ambients, and even ambient boundaries themselves, which we write as “ $x []$ ”. We support the traditional syntax “ $x[P]$ ” for ambients by interpreting “ $E_1 \dots E_k[P]E'_1 \dots E'_k$ ” as syntactic sugar for “ $E_1 \dots E_k [] E'_1 \dots E'_k.P$ ”. Except for this syntactic convention, the symbol $[]$ has no special interpretation in Meta \star and it is a (single) name just like in and

Names: $x, y ::= a \mid b \dots z \mid aa \mid ab \dots eas \mid eat \mid eau \dots \mid [] \mid \wedge \mid : \mid * \mid / \dots \mid \bullet$ Sub-forms: $f ::= x_0 x_1 \dots x_k$ Messages: $M, N ::= f \mid 0 \mid M.N$ Elements: $E ::= x \mid (x_1, x_1, \dots, x_k) \mid \langle M_1, M_2, \dots, M_k \rangle$ Forms: $F ::= E_0 E_1 \dots E_k$ Processes: $P, Q, R ::= F.P \mid !P \mid v(x).P \mid 0 \mid (P \mid Q)$
Free and bound names in terms are defined thus (the omitted cases being purely structural): $\begin{array}{ll} \text{FN}(x) = \{x\} & \text{BN}(x) = \emptyset \\ \text{FN}(x_1, \dots, x_k) = \emptyset & \text{BN}(x_1, \dots, x_k) = \{x_1, \dots, x_k\} \\ \text{FN}(F.P) = \text{FN}(F) \cup (\text{FN}(P) \setminus \text{BN}(F)) & \text{BN}(F.P) = \text{BN}(F) \cup \text{BN}(P) \\ \text{FN}(v(x).P) = \text{FN}(P) \setminus \{x\} & \text{BN}(v(x).P) = \text{BN}(P) \end{array}$
$\begin{array}{llll} P \equiv P & P \equiv Q \implies Q \equiv P & P \equiv Q \wedge Q \equiv R \implies P \equiv R & P \mid Q \equiv Q \mid P \\ P \mid (Q \mid R) \equiv (P \mid Q) \mid R & P \mid 0 \equiv P & !P \equiv P \mid !P & !0 \equiv 0 \\ P \equiv Q \implies F.P \equiv F.Q & P \equiv Q \implies !P \equiv !Q & P \equiv Q \implies v(x).P \equiv v(x).Q & \\ P \equiv Q \implies P \mid R \equiv Q \mid R & x \notin \text{FN}(F) \wedge x \notin \text{BN}(F) \implies F.v(x).P \equiv v(x).F.P & & \\ & x \notin \text{FN}(P) \implies P \mid v(x).Q \equiv v(x).(P \mid Q) & & \\ & y \notin \text{FN}(P) \implies v(x).P \equiv v(y).[x := y]P & & v(x).v(y).P \equiv v(y).v(x).P \end{array}$

Fig. 2. Syntax of Meta plus its structural congruence relation*

out. The process $F.0$ can be abbreviated as F .

A form consists of a nonempty sequence of **elements**, each of which is either a **name**, a **binding** element, or a **message** element. Names are used to name channels, ambients, and so on, but also work as keywords that distinguish forms with different roles in the calculus. A keyword is simply a free name that is matched explicitly by some reduction rule. Most non-alphanumeric ASCII characters that do not have any special meaning (\wedge , $:$, $*$, $/$, etc.) are also names and so can be used as keywords. With these we can encode, e.g., annotated communication actions like “ $\langle M \rangle^*$ ” or “ $(x)^y$ ” from Boxed Ambients [5] using pseudo- \TeX notation as the forms “ $\langle M \rangle \wedge^*$ ” and “ $(x) \wedge^y$ ”.

Binding elements (x_1, \dots, x_k) are used to create forms that bind names in the process they are applied to. The canonical use of this is for constructing receive actions, but again the meaning of the form is specified only by the reduction rules. Message elements $\langle \dots \rangle$ allow a form to contain other forms, which — given appropriate reduction rules for communication — can later be substituted into processes. For technical reasons we have to restrict the forms contained in message elements in that they cannot contain message or binding elements themselves. We refer to such restricted forms and their elements as **sub-forms** and **sub-elements**. In future work we hope to be able to handle calculi such as the spi-calculus [1] which communicate structured messages.

It is not uncommon for calculi to prefer using an explicit recursion construction “ $P ::= \mathbf{rec} X.P$ ” to express infinite behaviour rather than the process replication operator “ $!$ ”. There are certain technical problems with supporting this directly in Meta* (which may however be approachable by novel techniques involving regular grammars developed by Nielson et al. [19]). In the common case where the target calculus does not allow location boundaries to come between the $\mathbf{rec} X$ binder and the bound X , it can

$$\begin{array}{l}
\mathcal{S}^E x = \begin{cases} x & \text{when } x \notin \text{Dom } \mathcal{S} \\ y & \text{when } \mathcal{S}(x) = y \text{ for some } y \\ \bullet & \text{otherwise} \end{cases} & \mathcal{S}^M x = \begin{cases} \mathcal{S}(x) & \text{when } x \in \text{Dom } \mathcal{S} \\ x & \text{otherwise} \end{cases} \\
\mathcal{S}^P (v(x).P) = v(x).\mathcal{S}^P P & \mathcal{S}^P (x.P) = \begin{cases} \mathcal{S}(x) * (\mathcal{S}^P P) & \text{when } x \in \text{Dom } \mathcal{S} \\ x.(\mathcal{S}^P P) & \text{otherwise} \end{cases} \\
(M.N)_* P = M_*(N_* P) & 0_* P = P \quad f_* P = f.P
\end{array}$$

Fig. 3. The actions of term substitution. \mathcal{S}^M is the action on messages, \mathcal{S}^E the action on elements, \mathcal{S}^F on (sub)forms, and \mathcal{S}^P on processes. The omitted cases (including the one for \mathcal{S}^F) simply substitute componentwise into the syntactic element in question. The M_*P helper operator serves to linearise messages once we do not need to keep track of whether they are composite or not. (In other systems, this is often done by the structural congruence relation instead.)

easily be simulated in Meta \star by adding the reduction rule $\text{spawn } a \mid \text{rec } a.P \leftrightarrow P$ and then representing $\text{rec } X.\dots.X$ as $v(x).(\text{spawn } x \mid ! \text{rec } x.\dots.\text{spawn } x)$.

2.2 Well-scoped terms. The process term P is **well scoped** iff it contains no nested binding of the same name and none of its free names also appear bound in the term. Formally, it is required that (1) $\text{BN}(P)$ and $\text{FN}(P)$ are disjoint, (2) whenever P contains $F.Q$, $\text{BN}(F)$ and $\text{BN}(Q)$ are disjoint, and (3) whenever P contains $v(x).Q$, $x \notin \text{BN}(Q)$.

We generally require that terms are always well scoped. The reduction rules in an instantiation of Meta \star must preserve well-scopedness. This simplifies the type analysis because then we do not have to support α -conversion of ordinary binding elements.

We must still handle α -conversion of private names, which is built into the \equiv relation, but we will assume that it is not used to create terms that are not well scoped.

2.3 Substitutions. Substitutions in Meta \star substitute *messages* for *names*. The fact that entire *processes* cannot be substituted is an important technical premise of Poly \star ; it means that substitution can preserve well-scopedness. It is remarkable that mobility calculi in general refrain from substituting processes; calculi such as Seal [25] and \mathbf{M}^3 [12] which allow exchange of entire processes do it by local *movement* rather than *substitution*. This probably reflects the intuition that a running process is harder to distribute across a recipient process than a mere name or code sequence.

A (term) **substitution** \mathcal{S} is a finite map from names to messages. Figure 3 defines the action of \mathcal{S} on the various syntactic classes of Meta \star . In Mobile Ambients and its descendant calculi, the value exchanged in a communication operation can be either a name or a (sequence of) capabilities. The former is the case in reduction $\langle b \rangle \mid (a).\text{out } a.0 \leftrightarrow \text{out } b.0$ and the latter in $\langle \text{in } b \rangle \mid (a).x[\text{in } c.0] \leftrightarrow x[\text{in } b.\text{in } c.0]$. To support this, Fig. 3 contains special cases for the syntactic cases $M ::= F$ and $P ::= F.P$ when the form F is a lone name. In that case the substitution for the name is inserted directly into the message (or process structure).

In cases like $\{a \mapsto b\}^P x[\text{out } a.0]$ where the substituted name occurs properly inside a form, the substitution is carried out componentwise for each form element, and the name is replaced in the rule for $\mathcal{S}^E x$. In this context the replacement must be a name

too. This will be false if the term tries to reduce as $\langle \text{in } b \rangle \mid (a).\text{out } a.0 \hookrightarrow \text{out } (\text{in } b).0$. The published formalisms of most ambient-inspired calculi usually regard “out (in b)” as *syntactically* possible but *semantically* meaningless. That this configuration cannot occur is often the most basic soundness property of type systems for such calculi.

In Meta★ such a semantic error becomes a syntactic one: It is simply not possible to use an entire form as an element (except indirectly through a message element). If, at runtime, a substitution nevertheless tries to do so, we substitute the special name “•”, which is to be interpreted as, “an erroneous substitution happened here”. Thus, with the MA communication rule, Meta★ reduces $\langle \text{in } b \rangle.0 \mid (a).\text{out } a.0 \hookrightarrow \text{out } \bullet.0$. This convention is technically convenient because it allows us to bound the nesting depth of forms (using the sub-form restriction). Because most published calculi attach no semantics to forms like “out (in b)”, we do not lose any real expressiveness.

Forms that contain • are inert in Meta★ unless there are reduction rules that explicitly match on •. The calculus designer can also define reduction rules that create •’s in other situations to mark reduction results as “erroneous”. For example, in the polyadic π -calculus, it is usually considered a run-time error if someone tries to send an m -tuple on a channel where another process is listening for an n -tuple, with $n \neq m$. By writing explicit rules⁴ for such situations, they can be handled in parallel with malformed substitutions. (One cannot straightforwardly write patterns to test for malformed substitutions, which is one reason for building the generation of • into Meta★).

In either case, the Poly★ type system will conservatively estimate *whether* (and where) a • can occur. Which conclusions to draw from this (e.g., rejecting the input program due to “type error”) is up to the designer of the calculus.

The definitions in Figure 3 do not worry about name capture. In general, therefore, $\mathcal{S}^X X$ is only intuitively correct if $\text{BN}(X)$ is disjoint from the names mentioned in \mathcal{S} . In practise, this will always follow from the assumption that all terms are well scoped.

2.4 Reduction rules. Figure 4 defines most of the syntax and semantics of reduction rules for Meta★. Our full theory (and implementation) allows a slightly more expressive template language to the right of the arrow in **reduce** rules, but the subset we present here is sufficient to express the calculi listed in Sect. 2.5.

As an example, with this syntax we can describe Mobile Ambients by the ruleset

$$\mathcal{R}_{\text{MA}} = \{ \text{active}\{P \text{ in } a[P]\}, \\ \text{reduce}\{a[\text{in } b.P \mid Q] \mid b[S] \hookrightarrow b[a[P \mid Q] \mid S]\}, \\ \text{reduce}\{a[b[\text{out } a.P \mid Q] \mid S] \hookrightarrow a[S] \mid b[P \mid Q]\}, \\ \text{reduce}\{\text{open } a.P \mid a[R] \hookrightarrow P \mid R\}, \\ \text{reduce}\{\langle M \rangle.P \mid (a).Q \hookrightarrow P \mid \{a := M\}Q\} \}$$

These five rules are all that is necessary to instantiate Meta★ to be Mobile Ambients.⁵ The four **reduce** rules directly correspond to the reduction axioms of the target calculus. The rule **active**{P in a[P]} is the Meta★ notation for the “evaluation context” rule

⁴ E.g., **reduce**{ $\langle M1, M2 \rangle.P \mid (x1, x2, x3).Q \hookrightarrow \bullet.0$ } for $(m, n) = (2, 3)$. Our implementation provides an extension for writing a single rule that catches all pairs (m, n) at once.

⁵ The rules are not sufficient to get communication reduction with arbitrary arity. Our implementation provides a syntax for defining arbitrary-arity communication rules, but for reasons of space and clarity we omit it in our formal development.

<p>Name variables: $\hat{x} ::= a \mid b \mid c \mid \dots$ Message variables: $\hat{m} ::= M \mid N \mid \dots$ Process variables: $\hat{p} ::= P \mid Q \mid R \mid \dots$</p> <p>Substitutes: $\underline{s} ::= \hat{x} \mid \hat{m}$ Element templates: $\underline{E} ::= \hat{x} \mid x \mid (\hat{x}_1, \dots, \hat{x}_k) \mid \langle \hat{m}_1, \dots, \hat{m}_k \rangle$ Forms templates: $\underline{F} ::= \underline{E}_0 \underline{E}_1 \dots \underline{E}_k$ Process templates: $\underline{P} ::= \hat{p} \mid \underline{E}.P \mid 0 \mid (P_1 \mid P_2)$ $\quad \quad \quad \mid \{ \hat{x}_0 := \underline{s}_0, \dots, \hat{x}_k := \underline{s}_k \} \hat{p}$ (R)</p> <p>Rules: $\mathcal{R}^1 ::= \mathbf{reduce}\{P_1 \hookrightarrow P_2\} \mid \mathbf{active}\{\hat{p} \text{ in } P\}$ Rulesets: $\mathcal{R} \in \mathcal{P}_{\text{fin}}(\boxed{\mathcal{R}^1})$</p>				
<p>The syntactic choice marked (R) is allowed only in a reduce rule to the <i>right</i> of the arrow.</p>				
<p>Let an term instantiation \mathcal{V} map $\boxed{\hat{x}}$ to $\boxed{x} \setminus \{\bullet\}$, $\boxed{\hat{m}}$ to \boxed{M}, and $\boxed{\hat{p}}$ to \boxed{P}. It applies to templates strictly componentwise, except for the case that fills in and applies a substitution:</p> $\mathcal{V}^{\mathbf{P}}(\{\dots, \hat{x}_i := \underline{s}_i, \dots\} \hat{p}) = \{\dots, \mathcal{V}(\hat{x}_i) \mapsto \mathcal{V}(\underline{s}_i), \dots\}^{\mathbf{P}}(\mathcal{V}(\hat{p}))$ <p>As a special exception, $\mathcal{V}^{\mathbf{P}}P$ is considered <i>undefined</i> if $\mathcal{V}(\hat{x}_1) = \mathcal{V}(\hat{x}_2)$ for $\hat{x}_1 \neq \hat{x}_2$ such that \hat{x}_1 occurs in P below a form template containing a binding element $(\dots, \hat{x}_2, \dots)$. For example, $\{a \mapsto x, b \mapsto x, c \mapsto x\}$ cannot be applied to $(a).c.0 \mid (b).c.0$, which would otherwise capture names and produce $(x).x.0 \mid (x).x.0$.</p>				
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: center; padding: 5px;">$\frac{\mathbf{reduce}\{P_1 \hookrightarrow P_2\} \in \mathcal{R}}{\mathcal{R} \vdash \mathcal{V}^{\mathbf{P}}P_1 \hookrightarrow \mathcal{V}^{\mathbf{P}}P_2}$</td> <td style="text-align: center; padding: 5px;">$\frac{\mathcal{R} \vdash P \hookrightarrow Q}{\mathcal{R} \vdash \mathcal{V}(x).P \hookrightarrow \mathcal{V}(x).Q}$</td> </tr> <tr> <td style="text-align: center; padding: 5px;">$\frac{\mathbf{active}\{\hat{p} \text{ in } P\} \in \mathcal{R} \quad \mathcal{R} \vdash P \hookrightarrow Q}{\mathcal{R} \vdash (\mathcal{V}[\hat{p} \mapsto P])^{\mathbf{P}}P \hookrightarrow (\mathcal{V}[\hat{p} \mapsto Q])^{\mathbf{P}}P}$</td> <td style="text-align: center; padding: 5px;">$\frac{\mathcal{R} \vdash P \hookrightarrow Q}{\mathcal{R} \vdash P \mid R \hookrightarrow Q \mid R} \quad \frac{P \equiv Q \quad \mathcal{R} \vdash Q \hookrightarrow R}{\mathcal{R} \vdash P \hookrightarrow R}$</td> </tr> </table>	$\frac{\mathbf{reduce}\{P_1 \hookrightarrow P_2\} \in \mathcal{R}}{\mathcal{R} \vdash \mathcal{V}^{\mathbf{P}}P_1 \hookrightarrow \mathcal{V}^{\mathbf{P}}P_2}$	$\frac{\mathcal{R} \vdash P \hookrightarrow Q}{\mathcal{R} \vdash \mathcal{V}(x).P \hookrightarrow \mathcal{V}(x).Q}$	$\frac{\mathbf{active}\{\hat{p} \text{ in } P\} \in \mathcal{R} \quad \mathcal{R} \vdash P \hookrightarrow Q}{\mathcal{R} \vdash (\mathcal{V}[\hat{p} \mapsto P])^{\mathbf{P}}P \hookrightarrow (\mathcal{V}[\hat{p} \mapsto Q])^{\mathbf{P}}P}$	$\frac{\mathcal{R} \vdash P \hookrightarrow Q}{\mathcal{R} \vdash P \mid R \hookrightarrow Q \mid R} \quad \frac{P \equiv Q \quad \mathcal{R} \vdash Q \hookrightarrow R}{\mathcal{R} \vdash P \hookrightarrow R}$
$\frac{\mathbf{reduce}\{P_1 \hookrightarrow P_2\} \in \mathcal{R}}{\mathcal{R} \vdash \mathcal{V}^{\mathbf{P}}P_1 \hookrightarrow \mathcal{V}^{\mathbf{P}}P_2}$	$\frac{\mathcal{R} \vdash P \hookrightarrow Q}{\mathcal{R} \vdash \mathcal{V}(x).P \hookrightarrow \mathcal{V}(x).Q}$			
$\frac{\mathbf{active}\{\hat{p} \text{ in } P\} \in \mathcal{R} \quad \mathcal{R} \vdash P \hookrightarrow Q}{\mathcal{R} \vdash (\mathcal{V}[\hat{p} \mapsto P])^{\mathbf{P}}P \hookrightarrow (\mathcal{V}[\hat{p} \mapsto Q])^{\mathbf{P}}P}$	$\frac{\mathcal{R} \vdash P \hookrightarrow Q}{\mathcal{R} \vdash P \mid R \hookrightarrow Q \mid R} \quad \frac{P \equiv Q \quad \mathcal{R} \vdash Q \hookrightarrow R}{\mathcal{R} \vdash P \hookrightarrow R}$			

Fig. 4. Syntax and semantics of reduction rules.

$P \hookrightarrow P' \implies a[P] \hookrightarrow a[P']$. This is, in fact, the only concrete **active** rule that we have so far needed for encoding existing calculi. We might just have hard-coded something like this rule into **Meta***, but we find it cleaner not to have any built-in distinction between “action” forms and “process container” forms in the theory.

The lower half of Figure 4 defines how to derive a reduction relation between process terms from a ruleset. For example, let \mathcal{R}_{eat} be the ruleset for the fictional calculus from Fig. 1: $\mathcal{R}_{\text{eat}} = \{ \mathbf{active}\{P \text{ in } a[P]\}, \mathbf{reduce}\{a[\text{eat } b \mid P] \mid b[Q] \hookrightarrow a[P \mid b[Q]]\} \}$. We can then instantiate the first inference rule in the bottom third of Fig. 4 to obtain

$$\mathcal{R}_{\text{eat}} \vdash y1[\text{eat } x \mid z1[0]] \mid x[\text{eat } z1 \mid \text{eat } z2] \hookrightarrow y1[x[\text{eat } z1 \mid \text{eat } z2] \mid z1[0]]$$

by choosing \mathcal{V} to be $\{a \mapsto y1, b \mapsto x, P \mapsto z1[0], Q \mapsto (\text{eat } z1 \mid \text{eat } z2)\}$.

A reduction rule must not allow a well-scoped term to reduce to a non-well-scoped one. In order to guarantee this, the process templates in them must satisfy some scoping restrictions that are not apparent from the syntax. The restrictions will be satisfied by most rules that are intuitively sensible; because a precise understanding of how the restrictions work is not important for a high-level understanding of **Meta***, we refer to this paper’s long version [17] for a precise definition.

2.5 Example instantiations. We have checked (using machine-readable rulesets for our type inference implementation for Meta★/Poly★) that Meta★ can handle π -calculus [18]; Mobile Ambients [8]; Safe Ambients [16] and various variants regarding where the $\overline{\text{out}}$ capability must be found and which name co-capabilities must refer to (variants with anonymous co-capabilities also exist [15]); the Seal calculus [25] in the non-duplicating variation of [10]; Boxed Ambients [5], as well as its “asynchronous” and “Seal-like” variants (the latter being what later papers most often refer to as BA); Channelled Ambients [21]; NBA [6]; Push and Pull Ambient Calculus [22]; and \mathbf{M}^3 [12].

In many of these cases, Meta★ supports the straightforward way to notate process terms as flat ASCII text, but in some cases the native punctuation of the target calculus must be changed superficially to conform to Meta★ conventions about how a form looks. For example, the original send action $\overline{y}x$ from [18] is represented as “y<x>” (but “/y x” would also have worked), and “ $\overline{\text{enter}}(x,y)$ ” from [6] becomes “co-enter(x)y”, because it binds x in its continuation but uses y to handshake with the entering ambient. The “ $n[c_1, \dots, c_k; P]$ ” construction in Channelled Ambients [21] can be represented as “n[cs.(c₁.0 | \dots | c_k.0) | ps.P]”. In our ruleset for Mobile Ambients with Objective Moves [7], the fact that reduction rules cannot inspect the structure of messages forces us to represent the original “go $M.m[P]$ ” as “go.M.m[P]”.

3 Poly★: Types for Meta★

3.1 Shape predicates. As described in the introduction, *shape predicates* are the central concept in Poly★. A shape predicate denotes a set of process terms; certain shape predicates that are provably closed under reduction are *types*. The full language of shape predicates is somewhat complex, so let us introduce it piecewise. The basic idea of shape predicates can be explained simply: *A shape predicate looks like a process term. It matches any process term that can arise by repeatedly duplicating and/or removing sub-terms of the shape predicate.* Here, “duplicating” and “removing” sub-terms means applying the rewriting rules $\pi \rightsquigarrow \pi | \pi$ and $\pi \rightsquigarrow 0$ to any syntactic subterm of the shape predicate, in addition to using the structural congruence relation for terms.

For example, a shape predicate written $a[\text{in } b \mid \text{in } c] \mid c[0]$ would match the terms $a[\text{in } b \mid \text{in } c] \mid c[0]$ (which is identical to the shape predicate) and $a[\text{in } b] \mid a[\text{in } c] \mid c[0]$ (which arises by duplicating $a[\dots]$ and then removing one of the in subterms in each of the copies). But $a[\text{in } b] \mid c[a[0]]$ does not match, because duplicating subterms cannot make $a[]$ appear below a $c[]$. Neither is $\text{in } b \mid \text{in } c \mid c[0]$ allowed — when removing the $a[]$ form, the entire subterm below it must be removed.

The type in Fig. 1 can be written in term shape as $y1[\text{eat } x \mid z1[0] \mid x[z1[0] \mid \text{eat } z1 \mid \text{eat } z2]] \mid x[\text{eat } z1 \mid \text{eat } z2] \mid y2[x[\text{eat } z1 \mid \text{eat } z2 \mid z2[0]] \mid \text{eat } x \mid z2[0]]$.

In practice shape predicates cannot be exactly term-shaped, but it pays to keep this naive idea in mind as an intuition about what shape predicates are. When we introduce complications in the rest of this subsection, they should all be understood as “whatever is necessary to make the naive idea work in practice”.

Replication ($!P$) is ignored when matching shape predicates. This is sensible because $!P$ behaves like an infinite number of P ’s running in parallel, and any *finite* number of P ’s in parallel match a shape predicate exactly if a single P does.

Message types: $\mu ::= \{f_1, \dots, f_k\}^* \mid \{x\}$ Element types: $\varepsilon ::= x \mid (x_1, \dots, x_k) \mid \langle \mu_1, \dots, \mu_k \rangle$ Form types: $\varphi ::= \varepsilon_0 \varepsilon_1 \dots \varepsilon_k$ Node names: $X, Y, Z ::= X \mid Y \mid Z \mid \dots$ Type substitutions: $\mathcal{T} \in \boxed{x} \xrightarrow{\text{fin}} \boxed{\mu}$ Edges: $\eta ::= X \xrightarrow{\varphi} Y \mid X \xrightarrow{\mathcal{T}} Y$ Shape graphs: $G \in \mathcal{P}_{\text{fin}}(\boxed{\eta})$ Shape predicates: $\pi ::= \langle G \mid X \rangle$
$\frac{M \notin \boxed{x} \quad M * 0 = f_1.f_2 \dots f_k.0 \quad \{f_1, \dots, f_k\} \subseteq \{f'_1, \dots, f'_k\}}{\vdash M : \{f'_1, \dots, f'_k\}^*} \quad \frac{}{\vdash x : \{x\}}$ $\frac{}{\vdash x : x} \quad \frac{}{\vdash (x_1, \dots, x_k) : (x_1, \dots, x_k)} \quad \frac{\vdash M_1 : \mu_1 \quad \dots \quad \vdash M_k : \mu_k}{\vdash \langle M_1, \dots, M_k \rangle : \langle \mu_1, \dots, \mu_k \rangle}$ $\frac{\vdash E_0 : \varepsilon_0 \quad \dots \quad \vdash E_k : \varepsilon_k}{\vdash E_0 \dots E_k : \varepsilon_0 \dots \varepsilon_k} \quad \frac{(X \xrightarrow{\varphi} Y) \in G \quad \vdash F : \varphi \quad \vdash P : \langle G \mid Y \rangle}{\vdash F.P : \langle G \mid X \rangle}$ $\frac{\vdash P : \pi}{\vdash !P : \pi} \quad \frac{\vdash P : \pi \quad \vdash Q : \pi}{\vdash P \mid Q : \pi} \quad \frac{}{\vdash 0 : \pi}$

Fig. 5. The syntax and semantics of shape predicates. Edges of the form $X \xrightarrow{\mathcal{T}} Y$ do not influence the semantics of the shape predicate; Sect. 3.2 explains what they are for.

We want to represent all possible computational future of each term smashed together in a single shape predicate. This creates problems for the naive idea, because terms such as $!x[\text{eat } x]$ can evolve to arbitrary deep nestings of $x[\dots]$. Therefore we need shape predicates to be *infinitely* deep trees. We restrict ourselves to infinite shape predicates with finite *representations* — in other words, regular trees.

There are several known ways of representing regular trees as linear syntax, but we have found it easier to work directly with *graphs*. A shape predicate now has the form $\langle G \mid X \rangle$, where G is a directed (possibly cyclic) graph where each edge is labelled with a form, and X is a designated root node in the graph. A term matches the shape predicate if its syntax tree can be “bent into shape” to match a subgraph such that each form in the term lies atop a corresponding edge in the graph (edges may be used more than once), and groups of parallel composition, $!$, and 0 lie within a single node in the graph.

The formal structure of Poly^* uses graphs where node names are just opaque identifiers and the meaning is given by edge labels. When *displaying* the graphs (as in Fig. 1) we have found it useful to put each edge label inside the edge’s target node. Of course this can’t be done in the rare cases when two edges that share a target disagree.

Graphs alone are not enough to guarantee a finite type for every term. For example, the term $\langle x \rangle \mid ! (y). \langle y.y \rangle$ can (given the reduction rules of MA) evolve into terms with messages that contain arbitrarily long chains of x ’s *within* a single form. We need to abstract over messages such that an infinity of forms that look alike except having messages of different length within them can be described by the same shape graph label. This is the job of **message types** μ , which are defined in Figure 5.

The message type $\{f_1, \dots, f_k\}^*$ describes any message built from the any of forms f_i — *except* messages that are single names; such a message is matched by the message type $\{x\}$ instead. When $\{x\}$ is the *only* message type that matches x , we can see unambiguously from a message type whether \bullet will result from trying to substitute a message it matches into an element position. We use **element types** ε and **form types** φ to build form-like structures out of message types and non-message elements.

The syntax and semantics of shape predicates is defined in Figure 5. To save space and present the basic theory more clearly we do not handle *name restriction*; how to treat it is described in [17]. We have also omitted a third form of message types, sequenced message types, which allow more precise types and are defined in [17, sec. 5.2].

Define the **meaning** of message/element/form types and of shape predicates by

$$\llbracket \mu \rrbracket = \{ M \mid \vdash M : \mu \} \quad \llbracket \varepsilon \rrbracket = \{ E \mid \vdash E : \varepsilon \} \quad \llbracket \varphi \rrbracket = \{ F \mid \vdash F : \varphi \} \quad \llbracket \pi \rrbracket = \{ P \mid \vdash P : \pi \}$$

Proposition 3.1. *The meanings of shape predicates respect the structural congruence: If $P \equiv Q$ then $\vdash P : \pi \iff \vdash Q : \pi$ for all π .* \square

Let $\mu_1 * \mu_2$ be the least message type whose meaning includes $M.N$ for all $M \in \llbracket \mu_1 \rrbracket, N \in \llbracket \mu_2 \rrbracket$. With the language of message types presented here (omitting sequenced message types from [17]), $\mu_1 * \mu_2$ always has the form $\{f_1, \dots, f_k\}^*$, where the f_i 's are all the sub-forms that appear in either μ_1 or μ_2 in some canonical order (for this purpose the sub-form x is considered to appear in $\mu = \{x\}$). The $\mu_1 * \mu_2$ operation is associative.

3.2 Flow edges and subtyping. The only part of the shape predicate syntax of Figure 5 that has yet not been explained is the **flow edges** $X \dashrightarrow_{\mathcal{T}} Y$. They are not used at all in the above definition of the *meaning* of the shape graph, but they will be important for distinguishing between types and non-types. In brief, the flow edge $X \dashrightarrow_{\mathcal{T}} Y$ asserts that there may be a reduction where a process described by X is moved to Y and in the process incurs a substitution described by \mathcal{T} .

Alternatively, $X \dashrightarrow_{\mathcal{T}} Y$ can be viewed as a *demand* that whenever $P \in \llbracket \langle G \mid X \rangle \rrbracket$ and Q arises by applying a substitution described by \mathcal{T} to P , it must hold that $Q \in \llbracket \langle G \mid Y \rangle \rrbracket$. Because flow edges do not contribute to the meaning of shape predicates, there is no guarantee that this demand is actually satisfied for a shape predicate that contains the flow edge. This is a global property of the shape graph, and we will shortly define a class of **flow closed** shape graphs where the interpretation of flow edges is always true.

An important special case is when $\mathcal{T} = \emptyset$, where the process moves *without* any substitution. Then $X \dashrightarrow_{\emptyset} Y$ can also be viewed as an assertion that $\langle G \mid X \rangle$ is a *subtype* of $\langle G \mid Y \rangle$, or, symbolically, that $\llbracket \langle G \mid X \rangle \rrbracket \subseteq \llbracket \langle G \mid Y \rangle \rrbracket$. We therefore also speak of $X \dashrightarrow_{\emptyset} Y$ as a **subtyping edge**.

Write $\vdash \mathcal{S} : \mathcal{T}$ iff $\text{Dom } \mathcal{S} = \text{Dom } \mathcal{T}$ and $\vdash \mathcal{S}(x) : \mathcal{T}(x)$ for all $x \in \text{Dom } \mathcal{S}$.

Define the action of type substitution on subforms and message/element types by the rules in Figure 6. This definition ensures that $\llbracket \mathcal{T}^{\mathbf{M}} \mu \rrbracket$ contains the result of every *term* substitution $\mathcal{S}^{\mathbf{M}} M$ where $\vdash \mathcal{S} : \mathcal{T}$ and $\vdash M : \mu$, and likewise for elements.

Definition 3.2. *The shape graph G is **flow closed** iff whenever G contains $X \xrightarrow{\mathcal{Q}} Y$ and $X \dashrightarrow_{\mathcal{T}} Z$ such that $\text{BN}(\varphi) \cap \text{Dom } \mathcal{T} = \emptyset$, then there is a W such that G contains $Y \dashrightarrow_{\mathcal{T}} W$ and additionally it holds that*

$$\begin{aligned}
\mathcal{T}^f(x_0 \dots x_k) &= \begin{cases} \mathcal{T}(x_0) & \text{when } k = 0 \text{ and } x_0 \in \text{Dom } \mathcal{T} \\ \{\mathcal{T}^E_{x_0} \dots \mathcal{T}^E_{x_k}\}^* & \text{otherwise} \end{cases} \\
\mathcal{T}^M\{f_1, \dots, f_k\}^* &= \{\mathcal{T}^M_{f_1} \dots \mathcal{T}^M_{f_k}\}^* & \mathcal{T}^M\{x\} &= \begin{cases} \mathcal{T}(x) & \text{if } x \in \text{Dom } \mathcal{T} \\ \{x\} & \text{otherwise} \end{cases} \\
\mathcal{T}^E x &= \begin{cases} y & \text{when } \mathcal{T}^M\{x\} = \{y\} \text{ for some } y \\ \bullet & \text{otherwise} \end{cases} \\
\mathcal{T}^E \langle \mu_1, \dots, \mu_k \rangle &= \langle \mathcal{T}^M \mu_1, \dots, \mathcal{T}^M \mu_k \rangle \\
\mathcal{T}^E (x_1, \dots, x_k) &= (x_1, \dots, x_k)
\end{aligned}$$

Fig. 6. The action of type substitutions. $\mathcal{T}^f f$ is the action on subforms, $\mathcal{T}^M \mu$ is the action on message types, and $\mathcal{T}^E \varepsilon$ is the action on element types.

Let a **type instantiation** \mathcal{U} map $\boxed{\hat{x}}$ to $\boxed{x} \setminus \{\bullet\}$, $\boxed{\hat{m}}$ to \boxed{m} , and $\boxed{\hat{p}}$ to \boxed{X} . It applies to element and form templates completely componentwise (giving element and form types), just like term instantiations do. The relation between type instantiations and process templates are given by this inference system:

$$\begin{array}{c}
\frac{X = \mathcal{U}(\hat{p})}{\mathcal{U} \vDash_{\perp} \hat{p} : \langle G | X \rangle} \quad \frac{(\mathcal{U}(\hat{p}) \dashv \boxed{\bullet} \rightarrow X) \in G}{\mathcal{U} \vDash_{\text{R}} \hat{p} : \langle G | X \rangle} \quad \frac{(\mathcal{U}(\hat{p}) \dashv \dots, \mathcal{U}(\hat{x}_i) \dashv \mathcal{U}^S \hat{y}_i, \dots \rightarrow X) \in G}{\mathcal{U} \vDash_{\text{R}} \{\dots, \hat{x}_i := \hat{y}_i, \dots\} \hat{p} : \langle G | X \rangle} \\
\frac{}{\mathcal{U} \vDash_s 0 : \pi} \quad \frac{\mathcal{U} \vDash_s P_1 : \pi \quad \mathcal{U} \vDash_s P_2 : \pi}{\mathcal{U} \vDash_s P_1 \mid P_2 : \pi} \quad \frac{(X \xrightarrow{\mathcal{U}^E E} Y) \in G \quad \mathcal{U} \vDash_s P : \langle G | Y \rangle}{\mathcal{U} \vDash_s E.P : \langle G | X \rangle}
\end{array}$$

where $\mathcal{U}^S \hat{x} = \{\mathcal{U}(\hat{x})\}$ and $\mathcal{U}^S \hat{m} = \mathcal{U}(\hat{m})$. The rules for template processes have an L variant and an R variant; the variable letter s ranges over L and R.

As a special exception, $\mathcal{U} \vDash_s P : \pi$ is *not* considered to hold if $\mathcal{U}(\hat{x}_1) = \mathcal{U}(\hat{x}_2)$ for $\hat{x}_1 \neq \hat{x}_2$ such that \hat{x}_1 occurs in P below a form template containing a binding element $(\dots, \hat{x}_2, \dots)$.

Fig. 7. Matching of reduction rules to shape graphs.

1. If $\varphi = x$ and $\mathcal{T}(x) = \{y\}$, then G contains $Z \xrightarrow{y} W$.
2. If $\varphi = x$ and $\mathcal{T}(x) = \{f_1, \dots, f_k\}^*$, then $W = Z$ and G contains $Z \xrightarrow{f_i} Z$ for $1 \leq i \leq k$.
3. If none of the above apply and $\varphi = \varepsilon_0 \dots \varepsilon_k$, then G contains $Z \xrightarrow{\mathcal{T}^E \varepsilon_0 \dots \mathcal{T}^E \varepsilon_k} W$.

We call a shape predicate $\langle G | X \rangle$ flow closed iff its G component is. \square

Intuitively, a flow-closed graph is one where the intuitive meanings of flow edges are true. That is the content of the following theorem:

Proposition 3.3. *Let G be flow closed and contain $X \dashv \boxed{\tau} \rightarrow Y$. Assume that $\vdash S : \mathcal{T}$ and that $\text{BN}(P) \cap \text{Dom } \mathcal{T} = \emptyset$. Then $\vdash P : \langle G | X \rangle$ implies $\vdash S^P P : \langle G | Y \rangle$* \square

The assumption that $\text{BN}(P) \cap \text{Dom } \mathcal{T} = \emptyset$ will be true in our uses because we are assuming that all terms are well-scoped.

Definition 3.5. G is *locally closed* at X with respect to \mathcal{R} iff whenever \mathcal{R} contains $\text{reduce}\{P_1 \hookrightarrow P_2\}$ it holds that $\mathcal{U} \vdash_{\mathcal{L}} P_1 : \langle G|X \rangle$ implies $\mathcal{U} \vdash_{\mathcal{R}} P_2 : \langle G|X \rangle$. \square

Definition 3.6. The shape predicate $\pi = \langle G|X \rangle$ is *syntactically closed* with respect to \mathcal{R} iff G is flow closed and also locally closed at every $X \in \text{active}_{\mathcal{R}}(\pi)$. When this holds, we call π an (\mathcal{R} -)type. \square

Checking that a purported type is really syntactically closed is algorithmically easy; see [17] for details.

Theorem 3.7 (Subject reduction). If π is syntactically closed with respect to \mathcal{R} , then it is also semantically closed with respect to \mathcal{R} . \square

3.4 What to do with types. Once we have a type, what can we use it for? An obvious possibility is to check whether the term may “go wrong”. The user (or, more likely, the designer of a programming environment that uses Poly \star) specifies what “going wrong” means. It is a clear indication of error if \bullet turns up in an active position, but opinions may differ about if a \bullet is produced at a place in process tree that never becomes active.

One can also imagine very application-specific properties to check for, for example “this process can never evolve to a configuration where an ambient named a is inside one named b ”. This is easy to check for in the shape graph. Alternatively, one may want to write this as a rule, to have Poly \star do the checking: $\text{reduce}\{b[a[P] \mid Q] \hookrightarrow \bullet.0\}$. The ability to write such rules is one of the reasons why Meta \star does not distinguish strictly between “names” and “keywords”.

Poly \star makes it fairly easy to check *safety properties* like “no unauthorised ambients (e.g., a) inside secure ambients (e.g., b)”, but there are also questions of safety that Poly \star cannot help determine. This includes properties that depend on the *order* in which things happen, such as the “correspondence assertions” often used to specify properties of communication protocols. There are type systems for process calculi that can reason about such temporal properties (for example, [13] for the π -calculus), but we are aware of none that also handle locations and ambient-style mobility.

4 Type inference for Poly \star

Assume now that we are given a process term P and a ruleset \mathcal{R} ; we want to produce an \mathcal{R} -type for P . It is trivial to construct *some* type for P – one with a single node and a lot of edges in the shape graph. However, such a type may need to contain \bullet ’s and thus not prove that P “cannot go wrong”. In this section we discuss how to automatically infer more informative types.

We do not know how to do type inference that is complete for the full Poly \star type system; it allows too many types. Therefore we begin by defining a set of *restricted* types, for which we *can* have complete type inference.

Definition 4.1. Write $\varphi_1 \approx \varphi_2$ iff $\llbracket \varphi_1 \rrbracket \cap \llbracket \varphi_2 \rrbracket \neq \emptyset$. \square

The \approx relation is close to being equality. The only way for two non-identical φ ’s to be related by \approx is if they contain message types of the shape $\{\dots\}\star$. It is relatively safe to imagine the \approx is just a fancy way to write $=$, at least to a first approximation.

Definition 4.2. G satisfies the **width restriction** iff whenever it contains $X \xrightarrow{\varphi} Y$ and $X \xrightarrow{\varphi'} Y'$ with $\varphi \approx \varphi'$, it holds that $Y = Y'$. \square

Definition 4.3. G satisfies the **depth restriction** iff whenever it contains a chain $X_0 \xrightarrow{\varphi_1} X_1 \xrightarrow{\varphi_2} \dots \xrightarrow{\varphi_k} X_k$ with $\varphi_1 \approx \varphi_k$, it holds that $X_1 = X_k$. \square

Our type inference algorithm only produces types that satisfy both restrictions. The width restriction means that when type inference needs to add an outgoing edge from a node in the graph, it never has to choose between reusing an existing edge starting there and creating a new edge with a fresh target node, because the latter is forbidden by the restriction when there is any reusable edge. The depth restriction bounds the length of a simple path in a shape graph that can be constructed with a given set of names and a given maximal form arity, and so also bounds the total number of nodes in the graph. Therefore the closing operation described below cannot keep adding edges and nodes to the graph indefinitely and will eventually stop. (These two restrictions replace the notions of “discrete” and “modest” types in [2], which sometimes admitted slightly more precise types, but were very complex and hard to understand).

In [17] we describe a feature of our implementation which allows it to loosen the two restrictions by tracking the origin of each φ in the type graph.

The type inference proceeds in two phases. First we construct a minimal shape predicate which the term matches. Then we **close** the shape predicate — that is, rewrite its shape graph as necessary to make it syntactically closed.

The initial phase is simple. Because shape predicates “look like terms”, we can just convert the abstract syntax tree of the term to a tree-shaped shape graph. This graph may or may not satisfy the width and depth restrictions. If it does not, unify the nodes that must be equal.⁷ That may cause further violations of the two restrictions; continue unifying nodes as necessary until the restrictions are satisfied.

The closing of the shape graph is where the real work of type inference happens. It happens in a series of steps. In each step, check whether the shape graph is syntactically closed. If it is, the algorithm ends. Otherwise, the lack of closure can only be because edges already in the graph imply that some other edges *ought* to exist (by Definitions 3.2 or 3.5) but do not. In that case, add the new nodes and edges required by the failing rule, and do another round of unifications to enforce the width and depth restrictions.

The width and depth restriction together guarantee that the closure phase terminates. We do not have any good worst-case complexity bounds for the closure phase; instead our implementation allows further restrictions on types to be applied in order to quench blow-ups one might observe with particular calculi and example terms. The tightest restrictions will enforce polynomial complexity, at the cost of losing the possibility of spatial polymorphism. Thus restricted, Poly \star has a strength roughly comparable to current non-polymorphic type systems for ambient-derived calculi.

Theorem 4.4 (Principal typings). *A result of type inference π is a principal typing [26] for the input term P : For every π' such that $\vdash P : \pi'$ it holds that $\llbracket \pi' \rrbracket \supseteq \llbracket \pi \rrbracket$. \square*

⁷ This is the only way to reach a graph that satisfies the restrictions. If the width and depth restrictions had used $=$ instead of \approx , there might also have been the option of rewriting a φ to something larger but different, but there would not be a unique “best way” of doing that.

4.1 Implementation. We have implemented our type inference algorithm. Our implementation is available at the URL (<http://www.macs.hw.ac.uk/DART/software/PolyStar/>), as both a source download and an interactive web interface.

Beyond the features in this paper, our implementation allows fine-tuning of the analysis precision, which influences inference speed as well as inferred type size.

5 Conclusion

Poly \star extends basic properties of our previous system PolyA to a more general setting:

1. Poly \star has *subject reduction*. Also, given a process term P and a shape predicate π , one can decide by checking purely local conditions whether π is a *type*, and it is similarly decidable whether P *matches* π . Thus, it is decidable whether a process belongs to a specific type.
2. Poly \star supports a notion of *spatial polymorphism* that achieves what Cardelli and Wegner [9] called “the purest form of polymorphism: the same object or function can be used uniformly in different type context without changes, coercions or any kind of run-time tests or special encodings of representations”.
3. The types of Poly \star are sufficiently precise that many interesting *safety/security properties* can be checked, especially those that can be formulated as questions on the possible configurations that can arise at run-time.

In addition, this paper makes these completely novel contributions:

4. Meta \star is a syntactic framework that can be instantiated into a large family of mobile process calculi by supplying reduction rules.
5. The *generic type system* Poly \star works for any instantiation of Meta \star . We have checked that it works for π -calculus, a large number of ambient calculi, and a version of the Seal calculus. In [2] we claimed PolyA would be easy to extend to ambient-like calculi by hand, but extending the proofs for PolyA manually would be tedious. With Meta \star we have developed the theory to do such extensions fully automatically.
6. For the subsystem of Poly \star satisfying the *width* and *depth* restrictions, there is a type inference algorithm (which we have implemented) that always successfully infers a *principal type* for any process term. This means that Poly \star has the potential for *compositional analysis*.
7. The width and depth restriction are more natural and intuitive than the “discreteness” and “modesty” properties with respect to which we showed existence of principal types for PolyA.
8. Poly \star ’s handling of *communication* and *substitution* has been redesigned to be more direct and intuitive than in PolyA.

5.1 Related work. Another generic type system for process calculi was constructed by Igarashi and Kobayashi [14]. Like the shape predicates in Poly \star , their types look like process terms and stand for sets of structurally similar processes. Beyond that, however, their focus is different from ours. Their system is specific to the π -calculus and does not handle locations or ambient-style mobility. On the other hand, it is considerably

more flexible than Poly \star within its domain and can be instantiated to do such things as deadlock and race detection which are beyond the capabilities of Poly \star .

Yoshida [27] used graph types much like our shape predicates to reason about the order of messages exchanged on each channel in the π -calculus. Since this type system reasoned about *time* rather than *location*, it is not directly comparable to Poly \star , despite the rather similar type structure.

The spatial analysis of Nielson et al. [19] produces results that somewhat resemble our shape graphs, but does not have spatial polymorphism.

References

- [1] M. Abadi, A. D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Inform. & Comput.*, 148(1), 1999.
- [2] T. Amtoft, H. Makhholm, J. B. Wells. PolyA: True type polymorphism for Mobile Ambients. In *IFIP TC1 3rd Int'l Conf. Theoret. Comput. Sci. (TCS '04)*. Kluwer Academic Publishers, 2004.
- [3] T. Amtoft, F. Turbak. Faithful translations between polyvariant flows and polymorphic types. In *Programming Languages & Systems, 9th European Symp. Programming*, vol. 1782 of LNCS. Springer-Verlag, 2000.
- [4] G. Boudol. The π -calculus in direct style. In *Conf. Rec. POPL '97: 24th ACM Symp. Princ. of Prog. Langs.*, 1997.
- [5] M. Bugliesi, G. Castagna, S. Crafa. Boxed ambients. In *4th International Conference on Theoretical Aspects of Computer Science (TACS '01)*, vol. 2215 of LNCS. Springer-Verlag, 2001.
- [6] M. Bugliesi, S. Crafa, M. Merro, V. Sassone. Communication interference in mobile boxed ambients. In *FST & TCS 2002*, 2002.
- [7] L. Cardelli, G. Ghelli, A. D. Gordon. Mobility types for mobile ambients. In J. Wiedermann, P. van Emde Boas, M. Nielsen, eds., *ICALP'99*, vol. 1644 of LNCS. Springer-Verlag, 1999. Extended version appears as Microsoft Research Technical Report MSR-TR-99-32, 1999.
- [8] L. Cardelli, A. D. Gordon. Mobile ambients. In M. Nivat, ed., *FoSSaCS'98*, vol. 1378 of LNCS. Springer-Verlag, 1998.
- [9] L. Cardelli, P. Wegner. On understanding types, data abstraction, and polymorphism. *Computing Surveys*, 17(4), 1985.
- [10] G. Castagna, G. Ghelli, F. Z. Nardelli. Typing mobility in the Seal calculus. In K. G. Larsen, M. Nielsen, eds., *CONCUR*, vol. 2154 of LNCS. Springer-Verlag, 2001.
- [11] S. Chaki, S. K. Rajamani, J. Rehof. Types as models: Model checking message-passing programs. In *Conf. Rec. POPL '02: 29th ACM Symp. Princ. of Prog. Langs.*, 2002.
- [12] M. Coppo, M. Dezani-Ciancaglini, E. Giovannetti, I. Salvo. M3: Mobility types for mobile processes in mobile ambients. In *CATS 2003*, vol. 78 of ENTCS, 2003.
- [13] A. D. Gordon, A. S. A. Jeffrey. Typing correspondence assertions for communication protocols. *Theoret. Comput. Sci.*, 300(1-3), 2003.
- [14] A. Igarashi, N. Kobayashi. A generic type system for the pi-calculus. In *Conf. Rec. POPL '01: 28th ACM Symp. Princ. of Prog. Langs.*, 2001.
- [15] F. Levi, C. Bodei. A control flow analysis for safe and boxed ambients. In *Programming Languages & Systems, 13th European Symp. Programming*, vol. 2986 of LNCS. Springer-Verlag, 2004.
- [16] F. Levi, D. Sangiorgi. Controlling interference in ambients. In *POPL'00, Boston, Massachusetts*. ACM Press, 2000.
- [17] H. Makhholm, J. B. Wells. Instant polymorphic type systems for mobile process calculi: Just add reduction rules and close. Technical Report HW-MACS-TR-0022, Heriot-Watt Univ., School of Math. & Comput. Sci., 2004.
- [18] R. Milner, J. Parrow, D. Walker. A calculus of mobile processes. *Inform. & Comput.*, 100(1), 1992.
- [19] H. R. Nielson, F. Nielson, H. Pilegaard. Spatial analysis of BioAmbients. In R. Giacobazzi, ed., *Static Analysis: 11th Int'l Symp.*, vol. 3148 of LNCS, Verona, Italy, 2004. Springer-Verlag.
- [20] J. Palsberg, C. Pavlopoulou. From polyvariant flow information to intersection and union types. *J. Funct. Programming*, 11(3), 2001.
- [21] S. M. Pericas-Geertsen. *XML-Fluent Mobile Ambients*. PhD thesis, Boston University, 2001.
- [22] I. Phillips, M. G. Vigliotti. On reduction semantics for the push and pull ambient calculus. In *Theoretical Computer Science: 2nd IFIP Int'l Conf.*, vol. 223 of *IFIP Conference Proceedings*. Kluwer, 2002.
- [23] G. Sander. Graph layout through the VCG tool. In R. Tamassia, I. G. Tollis, eds., *Graph Drawing: DIMACS International Workshop, GD '94*, vol. 894 of LNCS. Springer-Verlag, 1994.
- [24] D. N. Turner. *The Polymorphic Pi-Calculus: Theory and Implementation*. PhD thesis, University of Edinburgh, 1995. Report no ECS-LFCS-96-345.
- [25] J. Vitek, G. Castagna. Seal: A framework for secure mobile computations. In *Internet Programming Languages*, vol. 1686 of LNCS. Springer-Verlag, 1999.
- [26] J. B. Wells. The essence of principal typings. In *Proc. 29th Int'l Coll. Automata, Languages, and Programming*, vol. 2380 of LNCS. Springer-Verlag, 2002.
- [27] N. Yoshida. Graph types for monadic mobile processes. In *Foundations of Software Technology and Theoret. Comput. Sci.*, 16th Conf., vol. 1180 of LNCS. Springer-Verlag, 1996.