


UNIVERSITY OF SUSSEX

COMPUTER SCIENCE

UNIVERSITY OF



**Trust and Partial Typing in Open
Systems of Mobile Agents**

James Riely and Matthew Hennessy

Trust and Partial Typing in Open Systems of Mobile Agents

JAMES RIELY AND MATTHEW HENNESSY

ABSTRACT. We present a *partially-typed* semantics for $D\pi$, a distributed π -calculus. The semantics is designed for mobile agents in *open* distributed systems in which some sites may harbor malicious intentions. Nonetheless, the semantics guarantees traditional type-safety properties at *good* locations by using a mixture of static and dynamic type-checking. We show how the semantics can be extended to allow *trust* between sites, improving performance and expressiveness without compromising type-safety.

1 Introduction

In [13] we presented a type system for controlling the use of resources in a distributed system, or network. The type system guarantees two properties:

- resource access is always *safe*, *e.g.* integer resources are always accessed with integers and string resources are always accessed with strings, and
- resource access is always *authorized*, *i.e.* resources may only be accessed by agents that have been granted *permission* to do so.

While these properties are desirable, they are properties of a network *as a whole*. In *open systems* it is impossible to verify the system as a whole, *e.g.* to “type-check the web”. In this paper, we present type systems and semantics for open systems

Trust and Partial Typing in Open Systems of Mobile Agents

corrupted; an example of this phenomenon is described in Section 3.2. Technically

Table 1 Syntax of names e , values u , patterns X , threads P , and networks M .

$e ::= k$	Location	$u, v, w ::= \text{bv}$	Base Value
a	Resource	e	Name
$X, Y ::= x$	Variable	x	Variable
(X_1, \dots, X_n)	Tuple	(u_1, \dots, u_n)	Tuple
$P, Q, R ::= \text{stop}$	Termination	$M, N ::= \mathbf{0}$	Empty
$P Q$	Composition	$M N$	Composition
$(\nu e:T)P$	Restriction	$(\nu_k e:T)N$	Restriction
$\text{goto } u.P$	Movement	$k[[P]]$	Agent
$u!\langle v \rangle P$	Output		
$u?(X:T)P$	Input		
$*P$	Replication		
$\text{if } u = v \text{ then } P \text{ else } Q$	Matching		

contain the main contributions of the paper. Section 4 presents the formalization of filters and dynamic typing, showing how these are incorporated into the runtime semantics. In Section 5 this framework is extended to include *trust*. Both sections include several examples, as well as proofs of Subject Reduction and Type Safety. In Section 6 we discuss the design of the semantics and describe some of its limitations, pointing to topics for further research. The paper ends with a brief survey of related work.

2 The Language and Standard Typing

In this section we review the syntax and standard semantics of $\text{D}\pi$. For a full treatment of the language, including many examples, see [13]. Our formalization of the language differs slightly from that of [13], as discussed in the conclusion.

2.1 Syntax

The syntax is given in Table 1, although discussion of types, T , is postponed to Section 2.3. The syntax is parameterized with respect to the following syntactic sets, which we assume to be disjoint:

- *Base*, of *base values*, ranged over by bv ,
- *Loc*, of *location names*, ranged over by $k-m$,
- *Res*, of *resource names*, ranged over by $a-d$,
- *Var*, of *variables*, ranged over by $x-z$.

Names, e , include location names and resource names. *Values*, $u-w$, include base values, names, variables and tuples of values. We occasionally use the metavariables $u-w$ to range over restricted classes of values, such as $\text{Var} \cup \text{Loc}$ or $\text{Var} \cup \text{Res}$;

such cases should be clear from context. *Patterns, X–Y*, include variables and tu-

Table 2 Standard Reduction

On tuples, the definition is by homomorphic extension:

$$\begin{aligned} \tilde{S} <: \tilde{T} & \text{ if } \forall i: S_i <: T_i \\ K[\tilde{A}] <: L[\tilde{B}] & \text{ if } K <: L \text{ and } \tilde{A} <: \tilde{B} \end{aligned}$$

An important property of the subtyping preorder is that it has a partial meet operator \sqcap .

DEFINITION 2.1. A partial binary operator \sqcap on a preorder (S, \preceq) is a partial meet operator if it satisfies the following for every $r, s, t \in S$:

For networks and threads, the main rules of interest are for agents and movement. For the agent $\ell[[P]]$ to be well-typed, P must be well-typed at location ℓ ; whereas for the thread $\text{goto } u.P$ to be well-typed at some location w , P must be well-typed at location u .

With this notation the rule for restriction in networks, for example, should be easily understandable. The network $(\nu_k e:T)N$ is well-typed with respect to Γ , $\Gamma \vdash (\nu_k e:T)N$, if e is new to Γ and N is well-typed with respect to Γ extended at k by the type information in declaration $e:T$, i.e. $\Gamma \sqcap \{_k e:T\} \vdash N$.

The rule for matching allows the combination of capabilities available on different instances of a location name. Note that the rule may only be applied when $S \sqcap T$ is defined. In the case that $S = T$, the rule degenerates to the standard rule for conditionals:

$$\frac{\Gamma \vdash_w u:T, v:T, P, Q}{\Gamma \vdash_w \text{if } u = v \text{ then } P \text{ else } Q}$$

The extra generality of the rule is necessary to type threads such as the following:

$$a?(z[x] \) \ b?(w[y] \) \ \text{if } z = w \ \text{then goto } z. \ (x?(u) \ y!\langle u \rangle)$$

This thread receives two remote channels from different sources, then forwards messages from one channel to the other. Further examples are given in [13] where we argue that the more general rule is crucial for typing many practical applications.

The typing system satisfies several standard properties such as type specialization, weakening and a substitution lemma, as described in [13]. The following result establishes that well-typed terms are free of runtime errors throughout their execution.

THEOREM

Table 4 Partial Typing Relation

All rules from Table 3 but those for restriction (v)

$$\begin{array}{c}
\text{(thread-bad)} \frac{\Gamma(w) = \text{lbad}}{\Gamma \vdash_w P} \\
\text{(net-new}_b) \frac{\Gamma(k) = \text{lbad} \\ \ell \notin \text{fn}(\Gamma) \\ \Gamma \sqcap \{\ell : \text{lbad}\} \vdash Pw}{\Gamma \vdash_w P} \\
\text{(thread-new}_g) \frac{\Gamma \neq \text{lbad} \\ e \notin \text{fn}(\Gamma) \\ \Gamma, (e) \{w e : T\} \vdash_w P}{\Gamma \vdash_w (v e : T) P}
\end{array}$$

The *partial typing relation* is defined in Table 4. All of the rules of the standard type system carry over to the partial typing system but for those concerning restriction, which require an additional side condition. Most important, the introduction of the rule (thread-bad) allows untyped locations to have truly arbitrary behavior, including the ability to (attempt to) send malicious agents to good locations. Thus the partial typing relation validates the judgment $\Gamma \vdash m \llbracket \text{goto } k.a! \langle t \rangle \rrbracket$, with Γ as given in the previous paragraph.

The rule (net-new_b) says that locations created at untyped locations should themselves be untyped. This rule is required to maintain well-typing under reductions such as:

$$k \llbracket (\nu \ell : L) \text{goto } \ell.P \rrbracket \mapsto (\nu_k \ell : L) k \llbracket \text{goto } \ell.P \rrbracket \mapsto (\nu_k \ell : L) \ell \llbracket P \rrbracket$$

The rules (thread-new_g) and (net-new_g) are as in the standard type system, but require that typed locations not create untyped ones. This “reasonableness requirement” is necessary to establish Type Safety, as formulated in Theorem 4.10.

3.2 An Example

Consider the following (partial) type environment:

$$\Gamma = \left\{ \begin{array}{l} k : \text{loc} \{ a : \text{res} \langle \text{int} \rangle \} \\ \ell : \text{loc} \left\{ \begin{array}{l} b : \text{res} \langle \text{loc} [\text{res} \langle \text{bool} \rangle] \rangle \\ c : \text{res} \langle \text{loc} [\text{res} \langle \text{int} \rangle] \rangle \end{array} \right\} \\ m : \text{lbad} \end{array} \right\}$$

Here we have three locations, k , ℓ and m , the first two of which are typed, and the last untyped. Of the good (typed) sites, we know that k has an integer channel a , and ℓ has two channels: c , which communicates dependent tuples with the second element being an integer channel; and b , which communicates dependent tuples with the second element being boolean channels.

Consider a system with two agents at ℓ , waiting to receive data on channels c and b , respectively. The first agent will expect, as the second element of the tuple it receives, the name of an integer channel, whereas the second will expect the name of a boolean channel. In addition suppose that there are agents at k and m poised to send data to ℓ on channels c and b , respectively. Such a system is the following:

$$\begin{aligned} P = & \ell \llbracket c? \langle w[y] \rangle \text{goto } w.y! \langle 0 \rangle \rrbracket \\ & | \ell \llbracket b? \langle z[x] \rangle \text{goto } z.x! \langle t \rangle \rrbracket \\ & | k \llbracket \text{goto } \ell.c! \langle k[a] \rangle \rrbracket \\ & | m \llbracket \text{goto } \ell.b! \langle k[a] \rangle \rrbracket \end{aligned}$$

Here the agents at ℓ and k are all quite reasonable; they could be typed using the standard type system of Table 3. The final agent, at m , however, flagrantly violates

the types of channels a and b ; this agent intends to send an integer channel (a) where a boolean channel is expected (on b).

One can easily see that, using the standard typing system (without lbad), for no Δ do we have $\Delta \vdash P$. This is because channel a at k may be bound to either y or x , and these identifiers are subject to conflicting uses. There is no assignment of standard types to a , b and c that satisfies all of the constraints given in P . On the other hand, using the partial typing system, we have $\Gamma \vdash P$. This well typing, however, is not preserved by reduction.

First consider the agents communicating on c . Using standard reduction, as defined in Table 2, these agents reduce as follows.

$$\begin{aligned} & \ell\llbracket c?\langle w[y] \rangle \text{ goto } w.y!\langle 0 \rangle \rrbracket \mid k\llbracket \text{goto } \ell.c!\langle k[a] \rangle \rrbracket \\ \longrightarrow & \ell\llbracket c?\langle w[y] \rangle \text{ goto } w \rrbracket \end{aligned} \quad (1)$$

It is worth contrasting this approach with the “purely local” approach adopted for “anonymous networks” in [14] (and outlined in Appendix B). In anonymous networks, the authority of incoming threads is not known. The semantics of [14] uses a weaker typing system requiring consistency only of *local* resource types. Thus, in that work, (6) is taken to be well-typed, with subject reduction failing only in the move from (7) to (8). The chief advantage of the current work is that it permits the use of *trust*, which appears to be incompatible with terms such as (6).

4 Filters and Authorities

In this section we propose a semantics which recovers subject reduction for partially-typed networks. The solution assumes that the origin, or *authority*, of incoming agents can be reliably determined.

4.1 Syntax and Semantics

To accomplish dynamic typechecking, it is necessary to add type information to running networks. We do this by adding a *filter* $k\langle\Delta\rangle$ for each location k in a network. The filter includes a type environment Δ which gives k 's view of the resources in the network. Suppose that in a network N , location k knows that there is resource named a of type A at location ℓ

Table 5 Typing and reduction using filters

Static typing: all rules from Table 4

$$\begin{array}{c}
 \Gamma <: \Delta \\
 \text{(net-filter}_g\text{)} \frac{\Gamma(k) = \Delta(k)}{\Gamma \vdash k \langle\langle \Delta \rangle\rangle}
 \end{array}
 \qquad
 \begin{array}{c}
 \Gamma(k) = \text{lbad} \\
 \text{(net-filter}_b\text{)} \frac{\Gamma(k) = \text{lbad}}{\Gamma \vdash k \langle\langle \Delta \rangle\rangle}
 \end{array}$$

Reduction precongruence: (r-split), (r-eq₁) and (r-eq₂) rules for \equiv from Table 2

$$\begin{array}{l}
 \text{(r}_f\text{-move)} \quad k \llbracket \text{goto } \ell. P \rrbracket \mid \ell \langle\langle \Delta \rangle\rangle \quad \text{if } k = \ell \text{ or } \Delta \Vdash_{\ell}^k P \\
 \longmapsto \quad \ell \llbracket P \rrbracket \mid \ell \langle\langle \Delta \rangle\rangle \\
 \\
 \text{(r}_f\text{-newr)} \quad k \llbracket (\nu a:A) P \rrbracket \mid k \langle\langle \Delta \rangle\rangle \quad \text{if } a \notin \text{fn}(\Delta) \\
 \longmapsto (\nu_k a:A) (k \llbracket P \rrbracket \mid k \langle\langle \Delta \sqcap \{k a:A\} \rangle\rangle) \\
 \\
 \text{(r}_f\text{-newl)} \quad k \llbracket (\nu \ell:L) P \rrbracket \mid k \langle\langle \Delta \rangle\rangle \quad \text{if } \ell \notin \text{fn}(\Delta) \cup \{k\} \\
 \longmapsto (\nu_k \ell:L) (k \llbracket P \rrbracket \mid k \langle\langle \Delta \sqcap \{\ell:L\} \rangle\rangle \mid \ell \langle\langle \{\ell:L\} \rangle\rangle) \\
 \\
 \text{(r}_f\text{-comm)} \quad k \llbracket a! \langle v \rangle P \rrbracket \mid k \llbracket a?(X:T) Q \rrbracket \mid k \langle\langle \Delta \rangle\rangle \\
 \longmapsto k \llbracket P \rrbracket \mid k \llbracket Q \{v/X\} \rrbracket \mid k \langle\langle \Delta \sqcap \{k v:T\} \rangle\rangle
 \end{array}$$

Dynamic typing: all rules from Table 4, with ‘

Here $\Delta \Vdash_\ell^k P$ is a *dynamic typing relation*, which intuitively says that P is well-formed to move to location ℓ , if acting under *authority of k* . Agents originating locally are assumed to be well-typed and therefore need not be checked dynamically.

Dynamic Typing. One approach to dynamic typing would be to take the dynamic typing relation to be the same as the static typing relation: $(\Vdash_w^k) = (\vdash_w)$. In effect, this would limit incoming agents to include only names of resources that are known in advance. While this is certainly sound, it is much too restrictive; for example, new resources could only be used by agents that originated locally. Consider the system:

$$k \llbracket (\text{va}) \text{ goto } \ell. b! \langle k[a] \rangle \rrbracket \mid \ell \llbracket b?(z[x]) P \rrbracket \mid \ell \langle \langle \Delta \rangle \rangle \quad (*)$$

Here k creates a new resource and wishes to communicate it to ℓ . However with $(\Vdash_w^k) = (\vdash_w)$ the move from k to ℓ is refused — (r_f -move) cannot be applied — since the filter Δ at ℓ can have no knowledge of the new resource a .

At the opposite extreme, we might allow threads to include any reference to non-local resources. However, this approach is clearly unsound from the counter-example given in the last section. The difficulty is that threads from bad locations may provide incorrect information about good locations, breaking subject reduction.

To straddle the gap between sound-but-useless and unsound-but-expressive, we introduce the notion of *authority*. We say that an agent leaving a location k acts *under the authority of k* . When an agent with authority k enters another location, we say that k is the *authority* of the agent.

While it is not safe to allow incoming agents to refer to *any* non-local resources, it is safe to allow them to refer to resources located at their authority, *i.e.* at their “home” location. Intuitively this is true because, under this discipline, “bad” agents can only to “lie” about resources located at their authority, which must have been a bad location to begin with. Lies about bad locations don’t hurt well-typing, since bad locations are untyped.

Formally, the rules for runtime typing extend those of the static type system given in Tables 3 and 4 with two additional rules for values and one for threads. These rules allow references to an incoming agent’s authority to go unchecked. The rule ($\text{val}_f\text{-self}_1$) allows an incoming agent to refer to its authority k , regardless of whether the filter environment Δ contains any information about k . (Note that the condition $\text{!bad} \prec: K$ is vacuously satisfied; we include it here only for reference in the next section.) The rule ($\text{val}_f\text{-self}_2$) allows an incoming agent to refer to resources at its authority. As an example, let $\Delta_\ell = \{\ell:\text{loc}\{a:\text{res}\langle K[B] \rangle\}\}$. Although we cannot infer that $\Delta_\ell \vdash_\ell a! \langle$

The rule (thread_f-return) allows a thread to return to its home location without subjecting the returning thread to further typechecking. This rule allows some additional expressiveness and reduces the burdens of typechecking somewhat.

Note that while the static typing system interprets the rules of Tables 3 and 4 with respect to an omniscient authority (Γ

$\Delta(\ell) = \Gamma(\ell)$, and therefore $\Delta(\ell)$ must have the entry $b:\text{res}\langle\text{loc}[\text{res}\langle\text{bool}\rangle]\rangle$; therefore to type the term we must be able to deduce $\Delta \Vdash_k^m a:\text{res}\langle\text{bool}\rangle$. Next note that Δ must be consistent with reality, namely Γ . This means that if Δ has knowledge of the resource a at k then it must be at the conflicting type $\text{res}\langle\text{int}\rangle$; therefore the rules of Table 3 cannot be used to infer $\Delta \Vdash_k^m a:\text{res}\langle\text{bool}\rangle$. Finally, since k is not the authority of the thread, neither can the additional rules of Table 5 be used to justify the claim that $\Delta \Vdash_k^m a:\text{res}\langle\text{bool}\rangle$. It follows that the inference $\Delta \Vdash_\ell^m b!\langle k[a] \rangle$ is impossible. \square

EXAMPLE 4.4. Let us now modify the previous example so that m attempts to relate information about its *own* resources, rather than those of k . In such cases, movement always succeeds, whether or not the source site is bad. For example, we have the reduction:

$$m[\text{goto } \ell.b!\langle m[a] \rangle] \mid \ell\langle\Delta\rangle \longrightarrow \ell[b!\langle m[a] \rangle] \mid \ell\langle\Delta\rangle$$

This follows since $\Delta \Vdash_\ell^m m[a:\text{loc}[\text{res}\langle\text{bool}\rangle]]$ can be inferred using ($\text{val}_f\text{-self}_1$) and ($\text{val}_f\text{-self}_2$)

bioRxiv preprint doi: <https://doi.org/10.1101/2018.04.04.299874>; this version posted April 4, 2018. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

since $\Delta' \Vdash_{\ell}^m c!\langle k[a] \rangle$. In the absence of other agents, the migrations can only be executed in one order (k first). \square

EXAMPLE 4.7. As a filter is updated, contradictory evidence may be obtained about a site, in which case the site must be untyped and can safely be deemed to be bad. As an example let Γ and the filter $\Delta = \{\ell:\Gamma(\ell)\}$ be as before, and consider the network:

$$m \llbracket \text{goto } \ell. b!\langle m[d] \rangle c!\langle m[d] \rangle \rrbracket \mid \ell \llbracket b?(z[x]) c?(w[y]) P \rrbracket \mid \ell \langle \langle \Delta \rangle \rangle$$

After the migration from m to ℓ and one communication this reduces to

$$\ell \llbracket c!\langle m[d] \rangle \rrbracket \mid \ell \llbracket c?(w[y]) P' \rrbracket \mid \ell \langle \langle \Delta' \rangle \rangle$$

where $\Delta' = \Delta \sqcap \{m:\text{loc}\{d:\text{res}\langle \text{bool} \rangle\}\}$. After the second communication, the network reduces to

$$\ell \llbracket P \rrbracket$$

Table 6 Runtime Error


$\ell \llbracket a?(X:T)P \rrbracket \mid \ell \langle \Delta \rangle \xrightarrow{err \ell}$	if $\Delta(\ell) \not\prec: \text{loc}\{a:\text{res}\langle T \rangle\}$
$\ell \llbracket a!\langle v \rangle P \rrbracket \mid \ell \langle \Delta \rangle \xrightarrow{err \ell}$	if $\Delta(\ell) \not\prec: \text{loc}\{a:\text{res}\langle T \rangle\}$, all T
$\ell \llbracket a!\langle v \rangle P \rrbracket \mid \ell \langle \Delta \rangle \xrightarrow{err \ell}$	if $\Delta(\ell) <: \text{loc}\{a\}$

a. Thus we may have trusted locations with certain resources for handling trusted data and others for handling untrusted data. In a similar vein we may have untrusted locations containing resources that communicate trusted data. As we shall see, these resources at untrusted locations cannot be used to increase the level of trust in a network.

The extension of the subtyping relation to these new types is based on two ideas:

- Every trusted location is also a location.
- Every trusted location guarantees good behavior; therefore, a “bad” or untyped location can never be trusted by a good site. This means that the type lbad is no longer the minimal location type in the subtyping preorder.

The subtyping relation is therefore built up using the ordering:



The formal definition is given in Appendix A.

PROPOSITION 5.1. *The set of types, extended with lbad and ltrust , under the subtyping preorder, has a partial meet operator.*

Proof. See Appendix A. □

With the addition of ltrust , the filters in a network may contain more detailed information about remote sites. Consider a network N which contains a filter $\ell\langle\langle\Delta\rangle\rangle$. As before, if k is not mentioned in Δ , this means that ℓ has no knowledge of k . But now there are now three possibilities with respect to a remote location k mentioned in $\ell\langle\langle\Delta\rangle\rangle$:

- $\Delta(k) <: \text{lbad}$, which means that ℓ has accumulated sufficient contradictory information about k to conclude that k is untyped.
- $\Delta(k) <: \text{ltrust}$, which means that ℓ trusts k . Note that this notion of trust is asymmetric; ℓ may trust k without k trusting ℓ . Also note that in well-typed systems, the rule (net-filter) in Table 5 ensures that k , trusted by ℓ , cannot be an untyped location unless ℓ itself is untyped; this is enforced by the requirement that $\Gamma(k) <: \Delta(k)$, since $\text{lbad} \not<: \text{ltrust}$.
- $\Delta(k) <: \text{loc}$, which means that ℓ knows of k , but cannot determine whether or not k is well-typed.

As we have seen in the previous section, the information in a filter may increase as the network evolves, *i.e.* $\ell\langle\Delta\rangle$ may evolve to $\ell\langle\Delta'\rangle$, where $\Delta' \prec \Delta$. But the subtyping relation between types ensures that once a location k is deemed “bad” in $\ell\langle\Delta\rangle$ it will remain so forever, and similarly with sites that are deemed “trusted”. It is only the third category which may change. In Example 4.7 we have seen that new information may result in $\Delta(k)$ changing from `loc` to `lbad`. We shall soon see that new information can also “improve” the status of k from `loc` to `ltrust`.

With the addition of trust, we can revise the reduction relation of the previous section to eliminate dynamic typechecking of agents arriving from trusted sites. We adopt the semantics of Table 5, replacing (r_f -move) with:

$$(r_t\text{-move}) \quad k\llbracket\text{goto } \ell.P\rrbracket \mid \ell\langle\Delta\rangle \longmapsto \ell\llbracket P\rrbracket \mid \ell\langle\Delta\rangle \quad \text{if } \Delta(\ell) \prec \text{ltrust} \text{ or } \Delta \Vdash_{\ell}^k P$$

Note that the presence of `ltrust` changes the importance of the condition `lbad` \prec `K` in the dynamic typing rule ($\text{val}_f\text{-self}_1$). Whereas this condition was tautological in Section 4, here it is not. The side condition precludes the use of ($\text{val}_f\text{-self}_1$) to infer $\Delta \Vdash_{\ell}^k k:\text{ltrust}$. This is important, as it prevents bad sites from becoming trusted.

EXAMPLE 5.2. Let $\Delta = \{\ell:\text{loc}\{d:\text{res}\langle\text{ltrust}\rangle\}, k:\text{ltrust}\}$ and consider the network:

$$\ell\langle\Delta\rangle \mid \ell\llbracket d?(z).P\rrbracket \mid k\llbracket\text{goto } \ell.d!\langle m\rangle\rrbracket \mid m\llbracket\text{goto } \ell.d!\langle n\rangle\rrbracket$$

Here the locations m and n are unknown to ℓ , *i.e.* $\Delta(m)$ and $\Delta(n)$ are undefined. In addition, d is a resource at ℓ for communicating trusted locations. The migration from m to ℓ is not immediately allowed since Δ

EXAMPLE 5.3. Consider the network

$$m[[\text{goto } \ell_0. \text{goto } \ell_1. \text{goto } \ell_2. P] \mid \ell_i \langle \Delta_i \rangle]$$

rule (*r-comm*), from the standard semantics. Another possibility would be to add analysis to the filter operation. Then the move rule would become:

$$k[[\text{goto } \ell.P]] \mid \ell\langle\langle\Delta\rangle\rangle \mapsto \ell[[P]] \mid \ell\langle\langle\Delta \sqcap \Delta'\rangle\rangle \quad \text{if } \Delta \Vdash_{\ell}^k P : \Delta'$$

The idea is that while checking an incoming term, the filter could also note any new names that are received with authority. Another possibility is to abandon non-local filter updates altogether; in this case, to allow a reasonable amount of expressiveness while preserving type safety, one would have to add further constructs to the language, as outlined at the end of the next subsection.

6.3 Progress

While subject reduction is important, it is purely a safety property; it does not imply that any reductions are ever performed. The semantics of Section 5 enjoys the property that whenever an agent attempts to move from a site k to a location that trusts k , the movement is always successful. This *liveness* property relies on the fact that the target trusts k , however. It works because agents from trusted sites come in with “universal authority”, *i.e.* the authority to say whatever they like.

A stronger property, which we call *progress*, is that whenever a well-typed agent attempts to move between two good locations, the movement is successful. Suppose we add the following clause to the definition of runtime error in Table 6:

$$k[[\text{goto } \ell.P]] \mid \ell\langle\langle\Delta\rangle\rangle \xrightarrow{\text{err}\{k,\ell\}} \text{if } k[[\text{goto } \ell.P]] \mid \ell\langle\langle\Delta\rangle\rangle \not\rightarrow$$

We then say that the typing system guarantees *progress* if

$$\Gamma \vdash N \text{ and } \Gamma(\ell) \not\prec \text{lbad}, \Gamma(k) \not\prec \text{lbad} \text{ implies } \neg(N \xrightarrow{\text{err}\{k,\ell\}})$$

Note that this property is not dependent on the trust relation between sites. Unfortunately, this progress property does not hold for our semantics, as can be seen from the following example. Let Γ , Δ and N be defined as follows:

$$\begin{aligned} \Gamma &= \left\{ \begin{array}{l} k : \text{ltrust}\{a : \text{res}\langle\text{int}\rangle\} \\ \ell : \text{ltrust}\{c : \text{res}\langle\text{loc}[\text{res}\langle\text{int}\rangle]\rangle\} \\ m : \text{ltrust} \end{array} \right\} \\ \Delta &= \{ \ell : \text{ltrust}\{b : \text{res}\langle\text{loc}[\text{res}\langle\text{bool}\rangle]\rangle\} \} \\ N &= m[[\text{goto } \ell.c!\langle k[a] \rangle]] \mid \ell\langle\langle\Delta\rangle\rangle \end{aligned}$$

Then $\Gamma \vdash N$, but $N \not\rightarrow$. The problem here is that, although the agent at m is well-typed, the reference to a is made without authority.

In practice, progress may not be that important, depending upon the application and the underlying implementation. In the example above where the move from m to ℓ is unsuccessful, an implementation of the filter at ℓ might report to m the reasons for the failure. It would then be up to m to resend the agent (or some piece of it) via k .

On the other hand, one way to guarantee progress would be to allow an incoming agent to refer only to *local* values or values at its *authority*. It is straightforward to design a static type system to enforce this constraint.⁴ However such an approach is very restrictive without some addition to the language. One possibility would be to introduce the notion of *signed* values (possibly based on [1]) which would allow certain values in an agent to be received (and typed) under a *different authority* than that of the agent itself. Even without full progress, signatures could be useful. In the example sketched above, after m 's agent is refused entry to ℓ , m might itself resend the agent, rather than forwarding it to k , this time carrying a signed value to prove that $k[a]$ is of the appropriate type.

6.4 Anonymous Networks

In [14] we presented a semantics for open system in which the authority of incoming agents is not known. We call such systems *anonymous networks*. In Appendix B we recast the semantics of [14] using filters and `lbad`. An attractive property of the semantics is that filter updating is purely local, *i.e.* no non-local data need be stored in filters. However because the origin of incoming agents cannot be determined it is not possible to incorporate notions of trust into this semantics, which implies that incoming agents must always be typechecked. In addition, it is very easy for good sites to develop misconceptions about other good sites, frustrating progress.

6.5 Plugins

One quickly discovers a limitation of $D\pi$ when trying to model *mirroring* of names across a network. The idea is to create a new resource, say a class name, at one

static, the second dynamic:

$$\frac{\Gamma \vdash_u v:A \quad \Gamma \sqcap \{w:\text{loc}\{u.v:A\}\} \vdash_w P}{\Gamma \vdash_w (\text{load } u.v:\text{class } A)P} \quad \frac{}{\Delta \Vdash_w^k k.a:\text{class } A}$$

We believe that using indexed names for mirrorable values will be crucial to establish Subject Reduction for such a language under partial typing. Note that such a naming strategy has been adopted by the Java community, although perhaps for different reasons, where class names are of the form `com.ibm.aglet`.

7 Conclusions

We introduced the notion of *partial typing*, which captures the intuition that “bad” sites in a network may harbor malicious agents while “good” sites may not. We demonstrated that in the presence of partial typing, some form of dynamic typechecking is required to ensure that good sites remain uncorrupted. We presented a semantics for $D\pi$ incorporating such dynamic typechecking, showing that it prevented type violations at good sites, and discussed the extent to which it guaranteed progress. Finally, we added *webs of trust* to the language, reducing the need for dynamic typechecking while retaining type safety at good sites.

The presentation of $D\pi$ given here differs somewhat from that of [13]; for example, we have added base types and moved some of the semantic rules from the structural equivalence to the reduction relation. Most of the changes are stylistic rather than substantive. Two of the changes, however, are essential for the treatment of partial typing. First, we have moved the rule (r-new) from the structural equivalence to the reduction relation; this is necessary to allow filter updating. Second, we have split the space of names in two, syntactically distinguishing locations from resources; this is necessary to prevent the filter updating rules from producing nonsense environments such as $\{\ell:\text{loc}\{\ell:\text{res}\langle \rangle\}\}$.

Several other distributed variants of the π -calculus have been defined, and it is informative to see how partial typing might be added to these languages. Syntactically, $D\pi$ is most similar to the language of Amadio and Prasad [4, 5], which also uses a “goto” operator for thread movement, written “`spawn(ℓ, P)`”. However, in Amadio and Prasad’s language, the set of resources available to a thread does not vary as the thread moves about the network. This means that an agent at ℓ can access resources at a different location k without requiring thread movement. While this makes the language very expressive, it also frustrates the use of filters to typecheck incoming threads. To add partial typing to such a language, one would need to typecheck *messages* dynamically, rather than threads, violating the third principle given in the introduction. In addition, the fact that names are assigned unique locations in [5] appears to be incompatible with partial typing, as outlined at the end of Section 4.1.

The join calculus of Fournet, Gonthier, Levy, Marganet and Remy [11] shares many of these properties. Whereas Amadio's language adds thread movement to message movement, however, the join calculus adds location movement. Unfortu-

A Proofs

The Subject Reduction and Type Safety results for Section 4 are special cases of those of Section 5, in which no trusted types appear. We present only the more general results. First we establish Proposition 5.1.

The formal definition of subtyping with lbad and ltrust is:

$$\begin{array}{lcl} \text{ltrust}\{\tilde{u}:\tilde{S}, \tilde{v}:\tilde{T}\} <: \text{loc}\{\tilde{u}:\tilde{S}\} & \text{ltrust}\{\tilde{u}:\tilde{S}, \tilde{v}:\tilde{T}\} <: \text{ltrust}\{\tilde{u}:\tilde{S}\} \\ \text{lbad} <: \text{loc}\{\tilde{u}:\tilde{S}\} & \text{loc}\{\tilde{u}:\tilde{S}, \tilde{v}:\tilde{T}\} <: \text{loc}\{\tilde{u}:\tilde{S}\} \\ & \text{lbad} <: \text{lbad} \end{array}$$

PROPOSITION (5.1). *The set of types, extended with lbad and ltrust , under the subtyping preorder, has a partial meet operator.*

Proof. Ignoring resources, the meet operator can be defined as follows:

	lbad	loc	ltrust
lbad	lbad	lbad	<i>undef</i>
loc	lbad	loc	ltrust
ltrust	<i>undef</i>	ltrust	ltrust

Combining this with the subtyping rules already given for resources, we have (omitting symmetric cases):

$$\begin{array}{l} \text{lbad} \sqcap \text{lbad} = \text{lbad} \\ \text{lbad} \sqcap \text{loc}\{\tilde{v}:\tilde{T}\} = \text{lbad} \\ \text{lbad} \sqcap \text{ltrust}\{\tilde{v}:\tilde{T}\} = \end{array}$$

The first result is proved by induction on the definition of \equiv , the second by induction on the definition of \mapsto . The proofs of both results, and the accompanying lemmas, can easily be derived from those found in [13]; in particular see Lemmas 4.7 and A.2, Proposition 4.5 and Theorem 5.1 of that paper. The only substantial differences are in the rules (r_f -comm) and (r_f -move), which we discuss below.

For the most part, the proof for (r_f -comm) follows that given in [13]. The only additional complication is presence of filter updating. Suppose that $\Gamma(k) \neq \text{lbad}$, $\Gamma \vdash_{\bar{k}} v:T$ and $\Gamma \triangleleft: \Delta$. We must show that $\Delta' = \Delta \sqcap \{_{\bar{k}}v:T\}$ is defined and that $\Gamma \triangleleft: \Delta'$, but this follows immediately from Lemma 2.3c and Lemma 2.3a.

Now let us turn to (r_f -move). Suppose that $\Gamma \vdash k\llbracket \text{goto } \ell.P \rrbracket \mid \ell\langle\langle \Delta \rangle\rangle$ and $k\llbracket \text{goto } \ell.P \rrbracket \mid \ell\langle\langle \Delta \rangle\rangle \mapsto \ell\llbracket P \rrbracket \mid \ell\langle\langle \Delta \rangle\rangle$. To establish the result, it is sufficient to show that $\Gamma \vdash_{\bar{\ell}} P$. There are three cases to consider:

-

- Suppose that $\ell \llbracket \text{if } u = v \text{ then } P \text{ else } Q \rrbracket \xrightarrow{\text{err}^\ell}$ because for every R either $\{\ell u : R\}$ or $\{\ell v : R\}$ is undefined. By way of contradiction, suppose that $\Gamma \vdash_{\ell} \text{if } u = v \text{ then } P \text{ else } Q$ and therefore for some S, T we have:

Note that the definition of static typing here is much weaker than that presented in the body of the paper. For example the network (6) of Section 3.2 is well-typed, although (8) is not. Using these definitions, one can establish Subject Reduction and a weaker notion of Type Safety (given in [14]).

This formulation has certain advantages over that of [14], such as the stronger language of partial types. Moreover it allows self moves to go untyped; *i.e.* reductions of the form $\ell[\text{goto } \ell.P] \mapsto \ell[P]$ are always allowed.

References

- [1] M. Abadi and A. D. Gordon. A calculus for cryptographic protocols: The spi calculus. Technical Report 414, University of Cambridge Computer Laboratory, January 1997.
- [2] Martín Abadi. Secrecy by typing in security protocols. Draft, 1997. Available from http://www.research.digital.com/RC/personal/Martin_Abadi/ome.tml.
- [3] *Conference Record of the ACM Symposium on Principles of Programming Languages*, San Diego, January 1998. ACM Press.
- [4] R. Amadio and S. Prasad. Localities and failures. In *Proc. 14th Foundations of Software Technology and Theoretical Computer Science*, volume 880 of *Lecture Notes in Computer Science*. Springer-Verlag, 1994.
- [5] Roberto Amadio. An asynchronous model of locality, failure, and process mobility. In *COORDINATION '97*, volume 1282 of *Lecture Notes in Computer Science*. Springer-Verlag, 1997.
- [6] L. Cardelli and A. D. Gordon. Mobile ambients, 1997. Draft, Available from <http://www.cl.cam.ac.uk/users/adg/>.
- [7] Luca Cardelli. A language with distributed scope.

