

Secure Composition of Untrusted Code: Wrappers and Causality Types

Peter Sewell
Computer Laboratory
University of Cambridge
Peter.Sewell@cl.cam.ac.uk

Jan Vitek
Department of Computer Sciences
Purdue University
jv@cs.purdue.edu

November 21, 1999

Abstract

We consider the problem of assembling concurrent software systems from untrusted or partially trusted off-the-shelf components, using *wrapper programs* to encapsulate components and enforce security policies. In previous work we introduced the *box- π* process calculus with constrained interaction to express wrappers and discussed the rigorous formulation of their security properties. This paper addresses the verification of wrapper information flow properties. We present a novel *causal type system* that statically captures the allowed flows between wrapped possibly-badly-typed components; we use it to prove that a unidirectional-flow wrapper enforces a causal flow property.

Contents

1	Introduction	2
2	A Boxed π Calculus	3
2.1	Syntax	4
2.2	Semantics	5
2.2.1	Reduction Semantics	5
2.2.2	Labelled Transitions	7
3	A Filtering Example	8
4	The Unidirectional-flow Wrapper	8
5	Colouring and Causal Flow	10
5.1	Colouring the Box- π Calculus	10
5.2	The Causal Flow Property	11
6	Causality Types	12
6.1	Soundness	15
6.2	Typing the Unidirectional-flow Wrapper	17
7	Discussion	17
8	Conclusion	19
A	Coloured LTS and Typing Rules	20
B	Proofs	22
B.1	Soundness	22
B.2	Proving Causal Flow for \mathcal{F}	33
	References	34

1 Introduction

A typical desktop software environment nowadays contains components – whole programs, plug-ins, or smaller code fragments – obtained from different untrusted or partially-trusted sources; they interact in intricate ways. Components may be faulty or malicious, or designed with a weaker security policy than the user requires – what is legitimate marketing data to a vendor may be considered sensitive by a user. It is difficult for a user to gain assurance that the composed system is secure, particularly because many off-the-shelf components are only available as object code. Furthermore current operating systems fail to provide support for the kind of fine-grained policies that could control the execution of such components [12, 9].

Recent practical work advocates interposing security code at the operating system boundary to observe and modify the data passing through [36, 15, 10, 12, 9]. Interposition techniques effectively encapsulate untrusted components in *wrapper programs* that have full control over the interactions between encapsulated components and the OS and over the interactions among components. The code of a wrapper can, for instance, perform access control checks, audit, attempt to detect intruders, and even monitor covert channels. In [9] Fraser, Badger and Feldman presented a system that splits the task of writing a wrapper into two parts. The wrapper’s body is written in a variant of C called the Wrapper Definition Language. The dynamic aspects of creating wrappers and instantiating concurrently executing components are specified in the Wrapper Life Cycle framework. While quite expressive, their approach does not provide guarantees that the wrappers actually enforce the desired security policies. The powerful wrapper language, the fact that all wrappers execute in kernel mode, and the fact that components are concurrent combine to make it difficult to understand precisely what properties a wrapper enforces.

Our work is exploring secure composition using wrappers, focussing on the rigorous statement and proof of their security properties. To this end, we have abstracted the essential characteristics of the problem in a process calculus – powerful enough to express the code of non-trivial wrappers and to express the concurrent composition of components, but small enough to be amenable to formal proof. In this paper we study information flow properties of wrappers. To express clear statements of such properties we equip our calculus with an annotated operational semantics, regarding a wrapper and each wrapped component as a different principal and colouring processes with the sets of principals that have causally affected them. This allows a direct statement of the property that one component cannot causally affect another. Verifying such a causal flow property directly can be laborious, requiring a characterisation of the state space of a wrapper containing arbitrary components. We therefore introduce a type system that statically captures causal flows. Since components are often provided as object code, which is impractical for the user to typecheck, our type system must admit programs with badly-typed subcomponents.

Expressing wrappers requires a language for composing concurrently-executing components, including primitives for encapsulating components and controlling their interactions. We use the *box- π* calculus of [28], recapitulated in Sections 2 and 3. *Box- π* is a minimal extension of the π -calculus with encapsulation; it is sufficiently expressive for components and wrappers while retaining the simplicity and tractable semantics needed for proving properties. Moreover *Pict* [22] demonstrates how to build a real programming language above a π -calculus core, a similar approach could be used for *box- π* .

Our main example, in Section 4, is a unidirectional-flow wrapper that encapsulates two components, allowing messages to be sent only in one direction between them and both components to interact with the environment. The following *box- π* program is a

simplified version of this example.

$$(\nu a, b) (a[P] \quad | \quad !c^a x. \bar{c}^b x \quad | \quad b[Q])$$

Processes P and Q are arbitrary, possibly malicious, components. They are encapsulated in named boxes, with private names a and b , and placed in parallel with a forwarder process on channel c from box a to box b . The term $\bar{c}^b x$ is an output to channel c in box b of value x . The term $c^a x. \bar{c}^b x$ prefixes this with an input on channel c from box a ; here the first x is a formal parameter that binds the second. The $!$ operator indicates a replicated input, so the forwarder persists after use. The boxes restrict communication of the encapsulated processes and ensure that P and Q cannot interact with each other directly; the private names ensure that they cannot interact with their environment in any other way. This simplified forwarder sends only unordered asynchronous messages; our main example provides FIFO communication (this is related to the NRL pump [16], as discussed in Section 4).

Intuitively the system enforces an information flow policy that prevents Q from leaking secrets to P . When one attempts to make such properties precise, however, there are many choices. A body of model-theoretic work on *non-interference* uses delicate extensional properties of the trace sets of systems. In our programming language setting a more intensional approach allows what we believe to be clearer statements. We start with a labelled transition semantics that specifies the input/output behaviour of programs and extend it to represent and propagate causal dependencies explicitly. In terms of this, one can state the desired property as ‘no visible action of P is causally dependent on any action of Q ’. The causal semantics and property are defined in Section 5.

The causal type system, given in Section 6, allows us to prove information flow properties of box- π programs. For the example above, to statically allow the flow from a to b but disallow the converse we can associate the components with principals p and q , then take a to be a box name whose contents may be affected by p , written $a : \mathbf{box}_{\{p\}}$, b to be a box name whose contents may be affected by p or q , written $b : \mathbf{box}_{\{p,q\}}$, and c to be a channel, carrying values of a top type \top , which can be affected only by p , so $c : \mathbf{chan}_{\{p\}} \top$. The fragment is then typable, whereas the converse forwarder $c^b x. \bar{c}^a x$ is not. The type system also deals with tracking causes through computation within a wrapper, including communication of channel names, and with interaction between a wrapper and badly-typed components. All boxes are assumed to contain untyped processes; wrapper code is statically typed; run-time type checking is required only when receiving from a component.

Further discussion of related work is given in Section 7; Section 8 concludes with future work. Proofs are deferred to an Appendix.

2 A Boxed π Calculus

The language – known as the *box- π calculus* – that we use for studying encapsulation properties must allow interacting components to be composed. The components will typically be executing concurrently, introducing nondeterminism. It is therefore natural to base the language on a process calculus. The box- π calculus lies in a large design space of distributed calculi that build on the π -calculus of Milner, Parrow and Walker [19], including among others the related calculi [2, 6, 8, 23, 26, 30, 32]. A brief overview of the design space can be found in [27]; here we highlight the main design choices for box- π .

The calculus is based on asynchronous message passing, with components interacting only by the exchange of unordered asynchronous messages. Box- π has an asynchronous π -calculus as a subcalculus – we build on a large body of work studying such calculi, notably

[14, 4]. They are known to be very expressive, supporting many programming idioms including functions and objects, and are Turing-complete; a box- π process may therefore perform arbitrary internal computation. The choice of asynchronous communication is important as it allows two components to interact without creating causal connections in both directions between them.

Box- π requires facilities for constraining communication – in standard π -calculi, if one process can send a message to another then the only way to prevent information flowing in the reverse direction is to impose a type system on components, which (as observed above) is not appropriate here. We therefore add a boxing primitive – boxes may be nested, giving hierarchical protection domains; communication across box boundaries is strictly limited. Underlying the calculus design is the principle that each box should be able to control all interactions of its children, both with the outside world and with each other. Boxes can be viewed as protection domains, akin to operating system-enforced address spaces. Direct communication is therefore allowed only between a box and its parent, or within the process running in a particular box. All other communication, in particular that between two sibling boxes, must be mediated by code running in the parent. This code can enforce an arbitrary security policy, even supporting dynamically-changing policies and interfaces (in contrast to static restriction or blocking operators [5, 33]).

Turning to the values that may be communicated, it is convenient to allow arbitrary tuples of names (or other tuples). Note that we do *not* allow communication of process terms. Moreover, no primitives for movement of boxes are provided, in contrast to most work cited above. The calculus is therefore entirely first order, which is important for the tractable theory of behaviour (the labelled transition semantics) that we require to state and prove security properties. The calculus is also untyped – we wish to consider the wrapping of ill-understood, probably buggy and possibly malicious programs.

2.1 Syntax

The syntax of the calculus is as follows:

Names We take an infinite set \mathcal{N} of *names*, ranged over by a, b, c etc. (except i, j, k, o, p, u, v). Both boxes and communication channels are named; names also play the role of variables, as in the π -calculus.

Values and Patterns Processes will interact by communicating values which are deconstructed by pattern-matching by the receiver. Values u, v can be names or tuples, with patterns p correspondingly tuple-structured.

$u, v ::=$	x	name
	$\langle v_1 .. v_k \rangle$	tuple ($k \geq 0$)
$p ::=$	$-$	wildcard
	x	name pattern
	$(p_1 .. p_k)$	tuple pattern ($k \geq 0$, no repeated names)

Processes The main syntactic category is that of *processes*, ranged over by P, Q . We introduce the primitives in three groups.

Boxes A box $n[P]$ has a name n , it can contain an arbitrary process P . Box names are not necessarily unique – the process $n[0] \mid n[0]$ consists of two distinct boxes named n ,

both containing an empty process, in parallel.

$P ::= n[P]$	box named n containing P
$P \mid P'$	P and P' in parallel
0	the nil process
\dots	

Communication The standard asynchronous π -calculus communication primitives are $\bar{x}v$, indicating an output of value v on the channel named x , and $xp.P$, a process that will receive a value output on channel x , binding it to p in P . Here we refine these with a tag indicating the direction of the communication in the box hierarchy. An *input tag* ι can be either \star , for input within a box, \uparrow , for input from the parent box, or a name n , for input from a sub-box named n . An *output tag* o can be any of these, similarly. For technical reasons we must also allow an output tag to be $\bar{\uparrow}$, indicating an output received from the parent that has not yet interacted with an input, or \bar{n} , indicating an output received from child n that has not yet interacted. The communication primitives are then

$P ::= \dots$	
$\bar{x}^o v$	output v on channel x to o
$x^{\iota} p.P$	input on channel x from ι
$!x^{\iota} p.P$	replicated input
\dots	

The replicated input $!x^{\iota} p.P$ behaves essentially as infinitely many copies of $x^{\iota} p.P$ in parallel. This gives computational power, allowing e.g. recursion to be encoded simply, while keeping the theory simple. In $x^{\iota} p.P$ and $!x^{\iota} p.P$ the names occurring in the pattern p bind in P . Empty patterns and tuples will often be elided.

New name creation Both box and channel names can be created fresh, with the standard π -calculus $(\nu x)P$ operator. This declares any free instances of x within P to be instances of a globally fresh name.

$P ::= \dots$	
$(\nu x)P$	new name creation

In $(\nu x)P$ the x binds in P . We work up to alpha conversion of bound names throughout, writing the free name function, defined in the obvious way for values, tags and processes, as $\text{fn}(_)$.

2.2 Semantics

This subsection defines the operational semantics of Box- π . The reader unfamiliar with process calculi may wish to skim to the start of Section 3 on a first reading.

2.2.1 Reduction Semantics

The simplest semantic definition of the calculus is a *reduction semantics*, a one-step reduction relation $P \rightarrow P'$ indicating that P can perform one step of internal computation to become P' . We first define the complement $\bar{\iota}$ of a tag ι in the obvious way, with $\bar{\star} = \star$ and $\bar{\iota} = \iota$. We define a partial function $\{-/_ \}$, taking a pattern and a value and giving,

where it is defined, a partial function from names to values.

$$\begin{aligned} \{v/-\} &= \{\} \\ \{v/x\} &= \{x \mapsto v\} \\ \{(v_1 \dots v_{k'})/(p_1 \dots p_k)\} &= \{v_1/p_1\} \cup \dots \cup \{v_k/p_k\} \quad \text{if } k = k' \\ &\text{undefined, otherwise} \end{aligned}$$

The natural definition of the application of a substitution σ (from names to values) to a process term P , written σP , is also a partial operation, as the syntax does not allow arbitrary values in all the places where free names can occur. We write $\{v/p\}P$ for the result of applying the substitution $\{v/p\}$ to P . This may be undefined either because $\{v/p\}$ is undefined, or because $\{v/p\}$ is a substitution but the application of that substitution to P is undefined. For example, $\{(z z)/x\}\bar{x}^*(\cdot)$ is not defined as $\overline{(z z)}^*(\cdot)$ is not in the syntax. Note that the result $\{y/x\}P$ of applying a name-for-name substitution is always defined. This definition of substitution leads to a lightweight notion of runtime error¹.

The definition of reduction involves an auxiliary *structural congruence* \equiv , defined as the least congruence relation such that the axioms below hold. This allows the parts of a redex to be brought syntactically adjacent.

$$\begin{aligned} P \mid 0 &\equiv P & (\nu x)(\nu y)P &\equiv (\nu y)(\nu x)P \\ P \mid Q &\equiv Q \mid P & (\nu x)(P \mid Q) &\equiv P \mid (\nu x)Q \quad x \notin \text{fn}(P) \\ (P \mid Q) \mid R &\equiv P \mid (Q \mid R) & (\nu x)n[P] &\equiv n[(\nu x)P] \quad x \neq n \end{aligned}$$

The reduction relation is now the least relation over processes satisfying the axioms and rules below. The (Red Comm) and (Red Repl) axioms are subject to the condition that $\{v/p\}P$ is well-defined.

$$\begin{aligned} n[\bar{x}^\dagger v \mid Q] &\longrightarrow \bar{x}^n v \mid n[Q] && \text{(Red Up)} \\ \bar{x}^n v \mid n[Q] &\longrightarrow n[\bar{x}^\dagger v \mid Q] && \text{(Red Down)} \\ \bar{x}^\dagger v \mid x^t p.P &\longrightarrow \{v/p\}P && \text{(Red Comm)} \\ \bar{x}^\dagger v \mid !x^t p.P &\longrightarrow !x^t p.P \mid \{v/p\}P && \text{(Red Repl)} \\ P \longrightarrow Q &\Rightarrow P \mid R \longrightarrow Q \mid R && \text{(Red Par)} \\ P \longrightarrow Q &\Rightarrow (\nu x)P \longrightarrow (\nu x)Q && \text{(Red Res)} \\ P \longrightarrow Q &\Rightarrow n[P] \longrightarrow n[Q] && \text{(Red Box)} \\ P \equiv P' &\longrightarrow Q' \equiv Q \Rightarrow P \longrightarrow Q && \text{(Red Struct)} \end{aligned}$$

The (Red Up) axiom allows an output to the parent of a box to cross the enclosing box boundary. Similarly, the (Red Down) axiom allows an output to a child box n to cross the boundary of n . The (Red Comm) axiom then allows synchronisation between a complementary output and input within the same box. The (Red Repl) axiom is similar, but preserves the replicated input in the resulting state.

Communications across box boundaries take two reduction steps, as in the following upwards and downwards communications.

$$\begin{aligned} n[\bar{x}^\dagger v \mid x^n p.P] &\longrightarrow n[0] \mid \bar{x}^n v \mid x^n p.P \\ &\longrightarrow n[0] \mid \{v/p\}P \\ \bar{x}^n v \mid n[x^\dagger p.P] &\longrightarrow n[\bar{x}^\dagger v \mid x^\dagger p.P] \\ &\longrightarrow n[\{v/p\}P] \end{aligned}$$

¹A more conventional notion of runtime error would give errors only when a tuple is used as a name, e.g. for output. The substitution-based notion is forced by our choice of syntax, which disallows values in various places where names may appear. In general it will report errors sooner than the conventional notion.

This removes the need for 3-way synchronisations between a box, an output and an input (as in [32]), simplifying both the semantics and the implementation model.

2.2.2 Labelled Transitions

The reduction semantics defines only the internal computation of processes. The statements of our security properties must involve the interactions of processes with their environments, requiring more structure: a *labelled transition relation* characterising the potential inputs and outputs of a process. We give a labelled semantics for box- π in an explicitly-indexed early style, defined inductively on process structure by an SOS. The *labels* are

$$\begin{aligned} \ell ::= & \tau && \text{internal action} \\ & \bar{x}^o v && \text{output action} \\ & x^\gamma v && \text{input action} \end{aligned}$$

where γ ranges over tags \star , n , \uparrow and \bar{n} . The labelled transitions can be divided into those involved in moving messages across box boundaries and those involved in communications between outputs and inputs. The movement labels are

$$\begin{array}{ll} \bar{x}^n v \text{ (sending to child } n) & x^{\bar{n}} v \text{ (box } n \text{ receiving from its parent)} \\ \bar{x}^\uparrow v \text{ (sending to the parent)} & \end{array}$$

Say $\text{mv}(o)$ is true if o is of the form n or \uparrow . The communication labels are

$$\begin{array}{ll} \bar{x}^\star v \text{ (local output)} & x^\star v \text{ (local input)} \\ \bar{x}^{\bar{n}} v \text{ (output received from child } n) & x^n v \text{ (input a message received from child } n) \\ \bar{x}^\uparrow v \text{ (output received from parent)} & x^\uparrow v \text{ (input a message received from parent)} \end{array}$$

Labels synchronise in the pairs $\bar{x}^\gamma v$ and $x^\gamma v$. The labelled transition relation has the form

$$A \vdash P \xrightarrow{\ell} Q$$

where A is a finite set of names and $\text{fn}(P) \subseteq A$; it should be read as ‘in a state where the names A may be known to P and its environment, process P can do ℓ to become Q ’. The relation is defined as the smallest relation satisfying the rules in Figure 2 omitting all transition subscripts, occurrences of $\mathbf{C} :$ and occurrences of $\mathbf{C} \bullet$. We write A, x for $A \cup \{x\}$ where x is assumed not to be in A , and A, p for the union of A and the names occurring in the pattern p , where these are assumed disjoint.

The labelled semantics is explained further in [28]. It is similar to a standard π semantics but must also deal with boxes and with reductions such as

$$((\nu x) \bar{x}^n z) \mid n[0] \longrightarrow (\nu x) n[\bar{x}^\uparrow z]$$

in which a new-bound name enters a box boundary. The two semantics coincide in the following sense.

Theorem 1 *If $\text{fn}(P) \subseteq A$ then $A \vdash P \xrightarrow{\tau} Q$ iff $P \longrightarrow Q$.*

This give confidence that the labelled semantics carries enough information. The proof is somewhat delicate; it is sketched in [29] and given in detail in [28].

3 A Filtering Example

To demonstrate the use of $\text{box-}\pi$ we give the definition of a wrapper that restricts the interface for user programs. In most operating systems, programs installed and run by a user enjoy the same access rights as the user, so if the user is allowed to open a socket and send data out on the network then so can any component. We idealize this scenario with the configuration below – an idealized single-user OS in which user Alice is executing a program P . Here the box around P stands for the operating system enforced user protection domain.

$$\begin{array}{l} \text{alice}[P] \mid \\ \dots \overline{\text{in}}^{\text{alice}} x \dots \mid \quad \text{OS write on Alice's } \textit{in} \text{ port} \\ \text{!out}^{\text{alice}} x \dots \mid \quad \text{OS read from Alice's } \textit{out} \text{ port} \\ \text{!net}^{\text{alice}} x \dots \quad \text{OS read from Alice's } \textit{net} \text{ port} \end{array}$$

The OS provides three channels \textit{in} , \textit{out} and \textit{net} , to respectively allow the user's program to read from and write to the terminal and to send data out on a network connection. The program P is executing within a box and so interacts with the OS using the \uparrow tag – for example $P = \text{in}^\uparrow x. \overline{\text{out}}^\uparrow \langle x x \rangle$ receives a value from the terminal and then sends a pair of copies of the value back to the terminal.

To execute some untrusted code fragment Q , Alice may run the code in parallel with her other applications, perhaps as $\text{alice}[P \mid Q]$. But, this grants too much privilege to Q . In particular, if $Q = \text{!in}^\uparrow x. \overline{\text{net}}^\uparrow x$ then any terminal input may be redirected to the net. A wrapper is a $\text{box-}\pi$ context which can provide fine-grain control of the behaviour of Q . For example, the filtering wrapper \mathcal{W}_1 of [28] prevents Q from accessing the network:

$$\mathcal{W}_1(-) \stackrel{\text{def}}{=} (\nu a) (a[-] \mid \text{!in}^\uparrow x. \overline{\text{in}}^a x \mid \text{!out}^a x. \overline{\text{out}}^\uparrow x)$$

The system becomes $\text{alice}[P \mid \mathcal{W}_1(Q)]$. The untrusted code is placed in a box with a fresh name a , so $a \notin \text{fn}(Q)$. In parallel with the box are two forwarders for \textit{in} and \textit{out} messages. The first, $\text{!in}^\uparrow x. \overline{\text{in}}^a x$, is a replicated input receiving values from the OS and sending them to a ; the second is dual. Only these two processes can interact with a due to the scope of the restriction, so even when put in parallel with other code the wrapper guarantees that Q will not be able to send on \textit{net} .

We show a small reduction sequence where $P = 0$ and $Q = \text{in}^\uparrow x. \overline{\text{net}}^\uparrow x$. Here B is the forwarders $\text{!in}^\uparrow x. \overline{\text{in}}^a x \mid \text{!out}^a x. \overline{\text{out}}^\uparrow x$.

$$\begin{aligned} \overline{\text{in}}^{\text{alice}} y \mid \text{alice}[P \mid \mathcal{W}_1(Q)] &\equiv \overline{\text{in}}^{\text{alice}} y \mid \text{alice}[(\nu a)(a[Q] \mid B)] \\ \rightarrow \text{alice}[\overline{\text{in}}^\uparrow y \mid (\nu a)(a[Q] \mid B)] &\equiv \text{alice}[(\nu a)(\overline{\text{in}}^\uparrow y \mid a[Q] \mid B)] \\ \rightarrow \text{alice}[(\nu a)(\overline{\text{in}}^a y \mid a[Q] \mid B)] & \\ \rightarrow \text{alice}[(\nu a)(a[\overline{\text{in}}^\uparrow y \mid Q] \mid B)] & \\ \rightarrow \text{alice}[(\nu a)(a[\overline{\text{net}}^\uparrow y] \mid B)] & \\ \rightarrow \text{alice}[(\nu a)(\overline{\text{net}}^a y \mid a[0] \mid B)] & \end{aligned}$$

At the final step the output from Q is prevented from leaving the alice box directly as B does not contain a forwarder for \textit{net} . It is prevented from interaction with any P (although here P was empty) by the restriction on a .

4 The Unidirectional-flow Wrapper

There is a tension between the strength of communication primitive supported by a wrapper and the strength of the security property it can guarantee. The examples of

the introduction and [28] provide only asynchronous unordered communication between components, which would be awkward to use in most real systems. At the other extreme, synchronous communication introduces causal flows in both directions (the causal flow property we state in Section 5 would not hold in a synchronous calculus, so a more delicate property would be required – perhaps stating that there are only data-less acks from one component to another). There are two intermediate points – one can provide asynchronous ordered communication, as we do below, or use some form of weak acknowledgments, as in the NRL pump [16]. The former still guarantees an absence of information flow (albeit at the cost of maintaining an unbounded buffer) while the latter limits bandwidth of covert channels. In both cases, it is essential to be able to guarantee that the implementation of the communication primitives does actually have the desired flow property, this is what we set to do here.

In Figure 1 we give a wrapper \mathcal{F} that takes two components and allows the first to communicate with the second by a first-in, first-out buffer. The wrapper has been written with care to avoid any information leak from the second component to the first. For simplicity both components have simple unordered input and output ports in_i and out_i to the environment; it would be routine to make these FIFO also.

$$\begin{array}{l}
\mathcal{F}_{(-1, -2)} = (\nu a, b) (a[-1] \mid b[-2] \mid \\
\quad (\nu buff, full) (\\
\quad \quad (\nu front, back) (\\
\text{(create FIFO buffer)} \quad \quad \quad \overline{buff}^* \langle front \ back \rangle \mid \\
\text{(connect } from^a \text{ to buffer)} \quad \quad \quad ! from^a \langle v \ r \rangle . (\nu r') (\overline{front}^* \langle v \ r' \rangle \mid r'^* . \overline{r}^a) \mid \\
\text{(connect buffer to } to^b \text{)} \quad \quad \quad ! back^* \langle v \ r \rangle . (\nu r') (\overline{to}^b \langle v \ r' \rangle \mid r'^b . \overline{r}^*) \mid \\
\text{(buffer code)} \quad \quad \quad ! buff^* \langle front \ back \rangle . front^* \langle v \ r \rangle . (\overline{r}^* \mid (\nu back') (\overline{buff}^* \langle front \ back' \rangle \\
\quad \quad \quad \mid \overline{full}^* \langle back' \ back \ v \rangle)) \mid \\
\quad \quad \quad ! full^* \langle back' \ back \ v \rangle . (\nu r) (\overline{back}^* \langle v \ r \rangle \mid r^* . back'^* \langle v' \ r' \rangle . (\overline{r}^{\prime*} \\
\quad \quad \quad \mid \overline{full}^* \langle back' , \ back \ v' \rangle)) \mid \\
\text{(I/O forwarders)} \quad \quad \quad ! in_1^\uparrow x . \overline{in}_1^a x \mid ! out_1^a x . \overline{out}_1^\uparrow x \mid \\
\quad \quad \quad ! in_2^\uparrow x . \overline{in}_2^b x \mid ! out_2^b x . \overline{out}_2^\uparrow x)
\end{array}$$

Figure 1: FIFO Pipeline Wrapper \mathcal{F} .

The interface to the wrapper is as follows. To write to the buffer a producer sends a value together with an acknowledgment channel to the wrapper (using a standard asynchronous π -calculus idiom). The wrapper inserts the value in a queue and acknowledges reception. For value v the producer may contain

$$(\nu ack) (\overline{from}^\uparrow \langle v \ ack \rangle \mid ack^\uparrow \dots),$$

sending the value and a new acknowledgement channel ack to the wrapper and, in parallel, waiting for a reply before proceeding with its computation. On the receiver side, we may have a process that waits for a pair of a value and an ack channel:

$$to^\uparrow \langle z \ r \rangle . (\overline{r}^\uparrow \mid \dots)$$

The name of the receiving channel is to ; channel r is used to send the acknowledgement back to the wrapper. Thus a configuration where B stands for the body of the wrapper could be:

$$(\nu a, b) (a [(\nu ack) (\overline{from}^\uparrow \langle v \ ack \rangle \mid ack^\uparrow . 0)] \mid b [to^\uparrow \langle z \ r \rangle . \overline{r}^\uparrow] \mid B)$$

The implementation of the wrapper is somewhat tricky, as we have to be careful not to introduce covert channels between the components. Within the wrapper there is a replicated input on *buff* that creates a new empty FIFO buffer and a replicated input on *full* that creates a new buffer cell containing a value. The key is to ensure that the acknowledgment to the first component not be dependent on any action performed by the second component. The glue process that connects the *from*^{*a*} channel to the buffer has a subprocess, $r'^*.\bar{r}^a$, to send the ack to *a*. This small process itself expects an ack from the head of the buffer saying that the message was inserted in the queue. The buffer code $front^*(vr).\bar{r}^* \dots$ acks on *r* immediately, in parallel with placing the new message in a full buffer cell at the head of the queue. The asynchrony here is essential.

So far we have been vague about the statement of the properties that we expect wrappers to enforce. For \mathcal{W}_1 it may be clear from examination of the code and the semantics that the wrapper is satisfactory, but it is unclear exactly what properties are guaranteed. For \mathcal{F} the situation is worse – even this simple wrapper is complex enough that a rigorous statement and proof of its security properties is essential; the user should not be required to examine the code of a wrapper in order to understand the security that it provides. We now turn to the task of formalizing these properties and developing the tools needed to prove them.

5 Colouring and Causal Flow

The intuitive property of \mathcal{F} that we wish to express is that the second wrapped component should not be able to affect the first. In [28] we expressed the intuitive property that one wrapped component does not causally affect another using a simple *coloured reduction semantics* for box- π . Output processes were annotated with sets of colours that record their causal histories – essentially the sets of principals that have affected them in the past – and the reduction semantics propagated this causal history data. In this paper we introduce also a coloured labelled transition semantics, allowing more direct statements of security properties of wrappers that interact with their environment. The coloured calculus is a trade-off – it captures less detailed causality information than the non-interleaving models studied in concurrency theory [37, 3, 7] but is much simpler; it captures enough information to express interesting security properties.

In [28] we also expressed a number of other desirable properties of wrappers – that they *honestly* forward messages between component and environment, and that they *mediate* all communication between components. The latter, related to intransitive non-interference [24], was expressed using the coloured semantics. Two further information flow properties were expressed using the uncoloured LTS: *new name directionality* and *permutation*. They illustrate the wide range of precise properties which the intuitive statement might be thought to mean.

5.1 Colouring the Box- π Calculus

We take a set *col* of *colours* or *principals* (we use the terms interchangeably) disjoint from \mathcal{N} . Let *k, p, q* range over elements of *col* and *C, D, K* range over subsets of *col*. We define a coloured box- π calculus by annotating all outputs with sets of colours:

$$P ::= C:\bar{x}^o v \mid x^t p.P \mid !x^t p.P \mid n[P] \mid 0 \mid P \mid P' \mid (\nu x)P$$

If *P* is a coloured term we write $|P|$ for the term of the original syntax obtained by erasing all annotations. Conversely, for a term *P* of the original syntax $C \circ P$ denotes the term with every particle coloured by *C*. For a coloured *P* we write $C \bullet P$ for the

coloured term which is as P but with C unioned to every set of colours occurring in it. We sometimes confuse \mathfrak{p} and the set $\{\mathfrak{p}\}$. Let $\text{pn}(P)$ be the set of colours that occur in P . We write CD for the union $C \cup D$.

In the coloured output $C:\bar{x}^o v$ think of C as recording the causal history of the output particle – C is the set (possibly empty) of principals $\mathfrak{p} \in C$ that have affected the particle in the past. In an initial state all outputs might typically be coloured by singleton sets giving their actual principals, for example colouring the code of wrapper \mathcal{F} and two wrapped components with different colours w, p, q :

$$(w \circ \mathcal{F}) (p \circ P \mid q \circ Q)$$

The coloured reduction semantics is obtained by replacing the first four axioms of the uncoloured semantics by the rules

$$\begin{aligned} n[C:\bar{x}^\dagger v \mid Q] &\longrightarrow C:\bar{x}^n v \mid n[Q] && \text{(C Red Up)} \\ C:\bar{x}^n v \mid n[Q] &\longrightarrow n[C:\bar{x}^\dagger v \mid Q] && \text{(C Red Down)} \\ C:\bar{x}^v v \mid x^t p.P &\longrightarrow C \bullet (\{v/p\}P) && \text{(C Red Comm)} \\ C:\bar{x}^v v \mid !x^t p.P &\longrightarrow !x^t p.P \mid C \bullet (\{v/p\}P) && \text{(C Red Repl)} \end{aligned}$$

that propagate colour sets. The coloured calculus has essentially the same reduction behaviour as the original calculus:

Proposition 2 *For any coloured P we have $|P| \longrightarrow Q$ iff $\exists P' . P \longrightarrow P' \wedge |P'| = Q$.*

The coloured labelled transitions have labels ℓ exactly as before. The coloured labelled transition relation has the form

$$A \vdash P \xrightarrow{\ell}_C Q$$

where A is a finite set of names and $\text{fn}(P) \subseteq A$; it should be read as ‘in a state where the names A may be known to P and its environment, process P can do ℓ , coloured C , to become Q ’. Again C records causal history, giving all the principals which have directly or indirectly contributed to this action. The relation is defined as the smallest relation satisfying the rules in Figure 2. It coincides with the previous LTS and with the coloured reduction semantics in the following senses.

Proposition 3 *For any coloured P we have $A \vdash |P| \xrightarrow{\ell} Q$ iff $\exists C, P' . A \vdash P \xrightarrow{\ell}_C P' \wedge |P'| = Q$.*

Proposition 4 *For coloured P and Q , if $\text{fn}(P) \subseteq A$ then $A \vdash P \xrightarrow{\tau}_{\emptyset} Q$ iff $P \longrightarrow Q$.*

5.2 The Causal Flow Property

The property can now be stated. Say an *instantiation* of some binary wrapper \mathcal{W} is an uncoloured process $\mathcal{W}(P, Q)$ where P and Q are uncoloured processes not containing the new-bound names scoping the holes of \mathcal{W} . Say \mathcal{W} is a *pure binary wrapper* if for any instantiation and any transition sequence

$$A \vdash \mathcal{W}(P, Q) \xrightarrow{\ell_1} \dots \xrightarrow{\ell_k} R$$

the labels ℓ_j have the form τ , $in_i^\dagger v$, or $\overline{out}_i^\dagger v$, for $i \in \{1, 2\}$. It is easy to see that \mathcal{F} is pure. Purity simply means that the wrapper has a fixed interface and thus simplifies the statement of the causal flow property.

Definition 1 (Causal flow property) *A pure binary wrapper \mathcal{W} has the causal flow property if for any instantiation $\mathcal{W}(P, Q)$ and any coloured trace*

$$A \vdash \emptyset \circ \mathcal{W}(P, Q) \xrightarrow{\ell_1}_{C_1} \dots \xrightarrow{\ell_k}_{C_k},$$

such that all input transitions $in_1^\uparrow v$ and $in_2^\uparrow v$ in $\ell_1.. \ell_k$ are coloured with principal sets $\{p\}$ and $\{q\}$ respectively, we have $\ell_j = \overline{out_1}^\uparrow v$ implies that $q \notin C_j$.

This property forbids any causal flow from an input on in_2 to an output on out_1 .

Different variants of the flow property, with different characteristics, can be stated – for example, to also prevent information in the initial state of Q affecting outputs on out_1 we could consider coloured traces

$$A \vdash (\emptyset \circ \mathcal{W})(p \circ P, q \circ Q) \xrightarrow{\ell_1}_{C_1} \dots \xrightarrow{\ell_k}_{C_k}$$

This second definition still allows the Q to communicate with P but only on the condition that P does not perform any further output dependent on the communicated values. Forbidding Q affecting P at all (even if there are no inputs or outputs of either component) can be done with a slightly more intricate coloured semantics. There is no clear cut ‘best’ solution, yet the use of causal semantics allows succinct statement of the alternatives and eases the comparison of these different properties.

6 Causality Types

Verifying a causal flow property directly can be laborious, requiring a characterisation of the state space of a wrapper containing arbitrary components. We therefore introduce a type system that statically captures causal flows; a wrapper can be shown to satisfy the causal flow property simply by checking that it is well-typed. This section introduces the type system, gives its soundness theorems, and applies it to \mathcal{F} .

A simple type system for $\text{Box-}\pi$ would have types

$$T ::= \mathbf{chan} T \mid \mathbf{box} \mid \langle T .. T \rangle$$

for the types of channel names carrying T , box names, and tuples. We annotate the first two by sets K of principals and add a type **name**, of arbitrary names, and \top , of arbitrary values, giving the *value types*

$$T ::= \mathbf{chan}_K T \mid \mathbf{box}_K \mid \langle T .. T \rangle \mid \mathbf{name} \mid \top$$


If $x : \mathbf{chan}_K T$ then x is the name of a channel carrying T ; moreover, in an output process $C : \overline{x} v$ on x the typing rules will require $C \subseteq K$ – intuitively, such an output may have been causally affected only by the principals $k \in K$. In an input $x^! p.P$ on x the continuation P must therefore be allowed to be affected by any $k \in K$, so any output within P must be on a channel of type $\mathbf{chan}_{K'} T$ with $K \subseteq K'$.

We are concerned with the encapsulation of possibly badly-typed components, so must allow a box $a[P]$ in a well-typed term to contain an untyped process P . The type system cannot be sensitive to the causal flows within such a box; it can only enforce an upper bound on the set of principals that can affect any part of the contents. If $a : \mathbf{box}_K$ then a is a box name; the contents may have been causally affected only by $k \in K$.

We take *type environments* Γ to be finite partial functions from names to value types. The type system has two main judgments, $\Gamma \vdash v : T$ for values and $\Gamma \vdash P : \mathbf{proc}_K$ for processes. The typing for processes records just enough information to determine when

prefixing a process with an input is legitimate – if $P : \mathbf{proc}_K$ then P can be prefixed by an input on a channel $x : \mathbf{chan}_{K'} \langle \rangle$, to give $x^*.P$, iff $K' \subseteq K$. Note, however, that a $P : \mathbf{proc}_K$ may have been affected by (and so syntactically contain) $k \notin K$.

To type interactions between well-typed wrapper code and a badly-typed boxed component some simple subtyping is useful. We take the subtype order $T \leq T'$ as below, and write $\bigwedge \{ T_i \mid i \in 1..k \}$ for the greatest lower bound of T_1, \dots, T_k , where this exists.



The complete type system is given in Figure 3; we now explain the key aspects by giving some admissible typing rules.

Basic Flow Typing Consider $x : \mathbf{chan}_K \langle \rangle$, $y : \mathbf{chan}_L \langle \rangle$ and the reduction

$$C : \bar{x}^* \mid x^*.D : \bar{y}^* \longrightarrow (C \cup D) : \bar{y}^*$$

During the reduction the output \bar{y}^* on y is causally affected by the output on x – the right-hand process term $(C \cup D) : \bar{y}^*$ records that the output on y has been (indirectly) affected by all the principals that had affected the output on x . For the left process to be well-typed we must clearly require $C \subseteq K$ and $D \subseteq L$; for the right process to be well-typed we need also $C \subseteq K$, to guarantee this the typing rules require $K \subseteq L$. The relevant admissible rules are below.

$$\frac{\Gamma \vdash x : \mathbf{chan}_K T \quad \Gamma \vdash v : T \quad C \subseteq K}{\Gamma \vdash C : \bar{x}^* v : \mathbf{proc}_K} \quad \frac{\Gamma \vdash x : \mathbf{chan}_K T \quad \Gamma, y : T \vdash P : \mathbf{proc}_{K''} \quad K \subseteq K''}{\Gamma \vdash x^* y.P : \mathbf{proc}_K}$$

Now consider also $y : \mathbf{chan}_L \langle \rangle$ and the process

$$C : \bar{x}^* \mid x^*. (D : \bar{y}^* \mid D' : \bar{y}'^*)$$

Here both the output on y and that on y' must be affectable by C , so the typing rule for parallel must take the intersection of allowed-cause sets:

$$\frac{\Gamma \vdash P : \mathbf{proc}_K \quad \Gamma \vdash Q : \mathbf{proc}_{K'}}{\Gamma \vdash P \mid Q : \mathbf{proc}_{K \cap K'}}$$

The examples above involve only communication within a wrapper, with tag \star . Communication between a wrapper and its parent, with tag \uparrow , has the same typing rules, as the parent is presumed well-typed.

Channel Passing Channel passing involves no additional complication. Consider the type environment $\Gamma = z : \mathbf{chan}_{K''} \langle \rangle$, $x : \mathbf{chan}_K \mathbf{chan}_{K''} \langle \rangle$, and the reduction

$$C : \bar{x}^* z \mid x^* y.D : \bar{y}^* \longrightarrow (C \cup D) : \bar{z}^*$$

The left-hand process is typable using the rules above if $C \subseteq K$ for the x output, $D \subseteq K''$ for the y output, and $K \subseteq K''$ for the input, using $\Gamma, y : \mathbf{chan}_{K''} \langle \rangle \vdash D : \bar{y}^* : \mathbf{proc}_{K''}$. Together these imply $(C \cup D) \subseteq K''$, so the right-hand process is well-typed.

Interacting with a box (at \top) As discussed above, the contents of a box may be badly-typed, yet a wrapper must still be able to interact with them. The simplest case is that in which a wrapper sends and receives values that it considers to be of type \top ; we consider more general communication in the next paragraph. The typing rule for boxes requires only that the principals $\text{pn}(P)$ syntactically occurring within the contents P of a box are contained in the permitted set and that P 's free names are all declared in the type environment.

$$\frac{\begin{array}{l} \Gamma \vdash a : \mathbf{box}_K \\ \text{pn}(P) \subseteq K \\ \text{fn}(P) \subseteq \text{dom}(\Gamma) \end{array}}{\Gamma \vdash a[P] : \mathbf{proc}_K}$$

Consider sending to and receiving from a box $a : \mathbf{box}_K$.

$$C : \bar{x}^a v \mid a[P] \mid z^a y.Q$$

For the output to be well-typed we must insist only that $C \subseteq K$; for the input to be well-typed Q must be allowed to be affected by any principal that might have affected the contents P .

$$\frac{\begin{array}{l} \Gamma \vdash a : \mathbf{box}_K \\ \Gamma \vdash x : \mathbf{name} \\ \Gamma \vdash v : \top \\ C \subseteq K \end{array}}{\Gamma \vdash C : \bar{x}^a v : \mathbf{proc}_K} \quad \frac{\begin{array}{l} \Gamma \vdash a : \mathbf{box}_K \\ \Gamma \vdash x : \mathbf{chan}_{K'} \top \\ \Gamma, y : \top \vdash P : \mathbf{proc}_{K''} \\ K \subseteq K' \subseteq K'' \end{array}}{\Gamma \vdash x^a p.P : \mathbf{proc}_{K'}}$$

Interacting with a box (at any transmissible S) More generally, a wrapper may receive from a box tuples containing names which are to be used *for communicating with the box* as channel names, for example

$$x^a (v r). (C : \bar{r}^a \mid \dots)$$

receives a value v and name r from box a and uses r to send an ack back into a . This necessarily involves some run-time typechecking, as the box may send a tuple instead of a name for r . There is a design choice here: how strong should this typechecking be? Requiring an implementation to maintain a run-time record of the types of all names would be costly, so we check only the structure of values received from boxes. We suppose the run-time representations of values allow names (bit-patterns of some fixed length) and tuples to be distinguished, and the number of items in a tuple to be determined, but no more (so e.g. $x : \mathbf{chan}_K T$ and $y : \mathbf{box}_L$ will both be represented as bit patterns of the same length). We introduce the supertype **name** of $\mathbf{chan}_K T$ and \mathbf{box}_L , and allow a wrapper to receive only values of the *transmissible types*

$$S ::= \top \mid \mathbf{name} \mid \langle S .. S \rangle$$

To send a value to a box by $C : \bar{x}^a v$ it is necessary only for x to be of type **name**.

The operational semantics expresses this run-time typechecking by means of the condition that $\{v/p\}P$ is well-defined in the reduction communication rule and the labelled-transition input rules – for example, $\{z^z/x\}C : \bar{x}^a$ is not well-defined, as the syntax does not allow a tuple to occur in channel-name position of an output. We would like to

ensure that run-time typechecking is only required when receiving values from a box, i.e. that for communication within a wrapper or between a wrapper and its parent such a substitution is always well-defined. This is guaranteed by requiring a box input prefix to immediately test all parts of a received value that are assumed of type **name** – in typing an input $x^a p.P$ the type environment Δ derived from the pattern p must contain no tuples, and all $x : \mathbf{name}$ in Δ must be used within P as a channel or box. For example, if $a : \mathbf{box}_K$ and $x : \mathbf{chan}_K(\mathbf{name name})$ then

$$x^a(yz).(\mathbf{K}:\bar{y}^a \mid \mathbf{K}:\bar{z}^a)$$

is well-typed as the pattern (yz) completely decomposes values of type $\langle \mathbf{name name} \rangle$ and both y and z are used as channels in $\mathbf{K}:\bar{y}^a \mid \mathbf{K}:\bar{z}^a$. On the other hand

$$x^a w.\bar{x}^* w$$

is not, as it may receive (for example) a triple from the box, leading to a later run-time error within the wrapper. The type system is conservative in also excluding $x^a(yz).(\mathbf{K}:\bar{y}^a)$. Say a type is *atomic* if it is of the form **name**, $\mathbf{chan}_K T$ or \mathbf{box}_K and *flat* if it is of the form \top , **name**, $\mathbf{chan}_K T$, or \mathbf{box}_K . Say Γ is atomic or flat if all types in $\text{ran}(\Gamma)$ are. The atomic types are those which can be dynamically extended using restriction. We consider dynamics (reductions and labelled transitions) only for processes with respect to atomic typing contexts; the definitions ensure that an extruded name can always be taken to be of an atomic type. The calculus has no basic data types, e.g. a type of integers, that are not dynamically extensible. This makes the type system a little degenerate.

The rest The typing rules for nil and restriction are straightforward; there is also a specialisation rule allowing some permitted affectees of a process to be forgotten.

$$\frac{}{\Gamma \vdash 0 : \mathbf{proc}_K} \quad \frac{\Gamma, x : T \vdash P : \mathbf{proc}_K \quad T \text{ atomic}}{\Gamma \vdash (\nu x)P : \mathbf{proc}_K} \quad \frac{\Gamma \vdash P : \mathbf{proc}_{K'} \quad K \subseteq K'}{\Gamma \vdash P : \mathbf{proc}_K}$$

6.1 Soundness

We wish to infer properties of the coloured input/output behaviour of wrappers from the soundness of the type system, and therefore need a subject reduction result which refers not only to reductions (equivalently, τ transitions) but also to input/output transitions. Define typed labelled transitions by

$$\Gamma \vdash_K P \xrightarrow{\ell}_C Q \quad \text{iff} \quad (\Gamma \text{ atomic} \wedge \Gamma \vdash P : \mathbf{proc}_K \wedge \text{dom}(\Gamma) \vdash P \xrightarrow{\ell}_C Q)$$

The subject reduction theorem for ℓ an output $\bar{x}^o v$ should state that x, o, v and Q have suitable types; the theorem for ℓ an input should state that if ℓ can be typed then Q can. The result is complicated by the fact that $\mathbf{box}\text{-}\pi$ is a calculus with new name generation, so new names can be extruded and intruded. Type environments for these names are calculated as follows. For a type environment Γ , with Γ atomic, and a value v extruded at type T define the type environment $tc(\Gamma, v, T)$ for new names in v as follows.

$$\begin{aligned} tc(\Gamma, x, T) &= x : T && \text{if } x \notin \text{dom}(\Gamma) \\ &&& \text{and } T \text{ atomic} \\ tc(\Gamma, x, \top) &= x : \mathbf{name} && \text{if } x \notin \text{dom}(\Gamma) \\ tc(\Gamma, x, T) &= \emptyset && \text{if } \Gamma \vdash x : T \\ tc(\Gamma, \langle v_1 \dots v_k \rangle, \top) &= \bigwedge_{1..n} tc(\Gamma, v_i, \top) \\ tc(\Gamma, \langle v_1 \dots v_k \rangle, \langle T_1 \dots T_k \rangle) &= \bigwedge_{1..n} tc(\Gamma, v_i, T_i) \\ tc(\Gamma, v, T) &\text{undefined elsewhere} \end{aligned}$$

Here $\bigwedge_{i \in 1..k} \Gamma_i$ is the type environment that maps each x in some $\text{dom}(\Gamma_i)$ to $\bigwedge \{T \mid \exists i . x : T \in \Gamma_i\}$, where all of these are defined. $\bigwedge_{i \in 1..k} \Gamma_i$ is undefined otherwise. Note that in the \top case the $tc(\Gamma, v_i, \top)$ will necessarily all be well-defined and will be consistent. To see the need for \bigwedge , consider $\Gamma = c : \mathbf{chan}_K(\mathbf{box}_K \mathbf{name})$ and $P = (\nu x) \bar{c}^* \langle x x \rangle$. P has an extrusion transition with value $\langle x x \rangle$; the type context $tc(\Gamma, \langle x x \rangle, \langle \mathbf{box}_K \mathbf{name} \rangle)$ should be well-defined and equal to $x : \mathbf{box}_K$.

Further, the type system involves subtyping, so $tc(\Gamma, v, T)$ can only be used as a bound on the extruded/intruded type environments. Say $\Gamma \leq \Gamma'$ iff $\text{dom}(\Gamma) = \text{dom}(\Gamma')$ and $\forall x \in \text{dom}(\Gamma) . \Gamma(x) \leq \Gamma'(x)$.

We can now state the subject reduction result. For output tags $\{\star, \bar{\uparrow}\}$ and \uparrow the name x is guaranteed to have a channel type and v the type carried; for a and \bar{a} they are only guaranteed to be a **name** and a value of type \top . $\{\star, \bar{\uparrow}\}$ and \bar{a} are communication tags, so x cannot be extruded, whereas \uparrow and a are movement tags, so x may be extruded. By convention we elide a conjunct that $tc(\dots)$ is defined wherever it is mentioned.

Theorem 5 (Subject Reduction) *If $\Gamma \vdash_K P \xrightarrow{\bar{x}^o v}_C Q$ then*

case $o \in \{\star, \bar{\uparrow}\}$: *for some K', T we have $C \subseteq K'$, $\Gamma \vdash x : \mathbf{chan}_{K'} T$, and there exists $\Theta \leq tc(\Gamma, v, T)$ such that $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$.*

case $o = \uparrow$: *for some K', T we have $C \subseteq K'$ and there exists $\Theta \leq tc(\Gamma, \langle x v \rangle, \langle \mathbf{chan}_{K'} T T \rangle)$ such that $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$.*

case $o = a$: *for some K' we have $C \subseteq K'$, $\Gamma \vdash a : \mathbf{box}_{K'}$, and there exists a type environment $\Theta \leq tc(\Gamma, \langle x v \rangle, \langle \mathbf{name}, \top \rangle)$ such that $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$.*

case $o = \bar{a}$: *for some K' we have $C \subseteq K'$, $\Gamma \vdash a : \mathbf{box}_{K'}$, $\Gamma \vdash x : \mathbf{name}$, and there exists $\Theta \leq tc(\Gamma, v, \top)$ such that $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$.*

If $\Gamma \vdash_K P \xrightarrow{x^\gamma v}_C Q$ then

case $\gamma \in \{\star, \uparrow\}$: *for some K', T we have $\Gamma \vdash x : \mathbf{chan}_{K'} T$. If moreover $C \subseteq K'$ and $\Theta \leq tc(\Gamma, v, T)$ then $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$.*

case $\gamma = a$: *for some $K' \subseteq K''$, and S we have $\Gamma \vdash a : \mathbf{box}_{K'}$, $\Gamma \vdash x : \mathbf{chan}_{K''} S$, $tc(\Gamma, v, S)$ well-defined, and $\text{ran}(tc(\Gamma, v, S)) \subseteq \{\mathbf{name}\}$. If moreover $C \subseteq K''$ and $\Theta \leq tc(\Gamma, v, S)$ then $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$.*

case $\gamma = \bar{a}$: *for some K' we have $\Gamma \vdash a : \mathbf{box}_{K'}$. If moreover $C \subseteq K'$ and we have $\Theta \leq tc(\Gamma, \langle x v \rangle, \langle \mathbf{name}, \top \rangle)$ then $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$.*

If $\Gamma \vdash_K P \xrightarrow{\tau}_C Q$ then $C = \emptyset$ and $\Gamma \vdash Q : \mathbf{proc}_K$.

A run-time error for $\mathbf{box}\text{-}\pi$ is a process in which a potential communication fails because the associated substitution is not defined. More precisely, P contains a run-time error if it contains subterms $\bar{x}^\gamma v$ and $x^\gamma p.P$ in parallel (and not under an input prefix) and $\{v/p\}P$ is not defined. In a well-typed process run-time errors can only occur within boxes (whose contents are untyped) or at communications from a box to the wrapper. Internal transitions of the wrapper and communications between the wrapper and its parent therefore do not require dynamic typechecking.

Theorem 6 (Limited Runtime Errors) *If $\Gamma \vdash P : \mathbf{proc}_K$, $P \equiv (\nu x_1 \dots x_n)(\bar{x}^\gamma v \mid x^\gamma p.P' \mid Q)$, Γ atomic, P' does not contain a box and $\gamma \in \{\star, \uparrow\}$ then $\{v/p\}P$ is well-defined. Similarly for replicated input.*

6.2 Typing the Unidirectional-flow Wrapper

Finally, we can show that instantiations of \mathcal{F} are well-typed and use the subject reduction theorem to conclude that \mathcal{F} has the causal flow property.

Theorem 7 (*\mathcal{F} typing*) *If*

$$\begin{aligned} \Gamma = \quad & in_1 : \mathbf{chan}_{\{p\}}^\top, & out_1 : \mathbf{chan}_{\{p\}}^\top, \\ & in_2 : \mathbf{chan}_{\{q\}}^\top, & out_2 : \mathbf{chan}_{\{p,q\}}^\top, \\ & from : \mathbf{chan}_{\{p\}} \langle \top \mathbf{name} \rangle, \\ & to : \mathbf{chan}_{\{p,q\}} \langle \top \mathbf{chan}_{\{p,q\}} \rangle, \\ & \Gamma_1 \end{aligned}$$

and also $\text{fn}(P, Q) \subseteq \text{dom}(\Gamma) - \{a, b\}$ then $\Gamma \vdash \emptyset \circ \mathcal{F}(P, Q) : \mathbf{proc}_p$.

The proof of this involves type assumptions for the new-bound names of \mathcal{F} as follows.

$$\begin{aligned} & a : \mathbf{box}_{\{p\}} \\ & b : \mathbf{box}_{\{p,q\}} \\ buff : & \mathbf{chan}_{\{p\}} \langle \mathbf{chan}_{\{p\}} \langle \top \mathbf{chan}_{\{p\}} \rangle \rangle \\ & \mathbf{chan}_{\{p,q\}} \langle \top \mathbf{chan}_{\{p,q\}} \rangle \rangle \\ full : & \mathbf{chan}_{\{p,q\}} \langle \mathbf{chan}_{\{p,q\}} \langle \top \mathbf{chan}_{\{p,q\}} \rangle \rangle \\ & \mathbf{chan}_{\{p,q\}} \langle \top \mathbf{chan}_{\{p,q\}} \rangle \rangle \\ & \top \rangle \end{aligned}$$

A straightforward induction on trace lengths using the Subject Reduction theorem then proves the desired causal flow result:

Theorem 8 *Wrapper \mathcal{F} has the causal flow property.*

7 Discussion

Policy enforcement mechanisms: Wrappers impose security policies on components for which it is impractical to analyze the internal structure, e.g. where only untyped object code is available.

Several alternative approaches are possible, differing in the level of trust required, the flexibility of the security policy enforced, and their costs to component producers and users. Code signing and Java-style sandboxing have low cost but cannot enforce flexible policies – signed components may behave in arbitrary ways whereas sandboxed components should not be able to interact with each other at all. Code signing requires the user to have total trust in the component producers – not just in their intent, but also in their ability to produce bug-free components. Sandboxing requires no trust, but the lack of any interaction is often too restrictive. More delicate policies can be enforced by shipping code together with data allowing the user to type-check it in a security-sensitive type system [34, 13], or to check a proof of a security-relevant behavioural property [20]. In the long term these seem likely to be the best approaches, but they require component producers to invest effort and to conform to a common standard for types or proofs – in the short term this is prohibitive. Shifting the burden of proof to the user, by performing type inference or static analysis of downloaded code, seems impractical given only the object code, which may not have been written with security in mind and so not conform to any reasonable type system. In contrast, wrappers have been shown to have low-cost

– none to the producer and only a small run-time cost to the user [9]. They allow more flexible interaction than sandboxing, albeit coarser-grain policies than proof-carrying components or security-type-checked components.

Information flow properties: The causal flow property is related to the property, studied in many contexts, that there is no information flow from a high to a low security level (though most work addresses components, which may have the property, rather than wrappers, which may enforce it on subcomponents). The literature contains a range of definitions that aim to capture this intuition in some particular setting; the formalisations vary widely. A basic choice is whether the property is stated purely extensionally, in terms of a semantics that describes only the input/output behaviour of a system, or using a more intensional semantics. A line of work on Non-Interference, summarised in [18], takes an extensional approach, stating properties in terms of the traces of input and output events of a system. Related definitions, adapted to a programming language setting, are used in [34, 13]. In the presence of nondeterminism, however, non-interference becomes problematic – as discussed in [35], the property may only be meaningful given probabilistic scheduling, which has a high run-time cost.

We believe that the basic difficulty is that the intuitive property is an *intensional* one – the notion of one component affecting another depends on some understanding of how components interact; a precise statement requires a semantics that captures some aspects of internal execution, not just input/output behaviours. This might be denotational or operational. Intensional denotational semantics have been used in the proofs (and, in the last, statements) of non-interference properties in [13, 1, 25], which use a logical relations proof and PER-based models. [35] and [25] go on to consider probabilistic properties.

For wrappers, it is important that the end-user be able to understand the security that they provide as clearly as possible. We therefore wish to use as lightweight a semantics as possible, as this must be understood before any security property stated using it, and so adopt an annotated operational semantics (developing a satisfactory denotational semantics of $\text{box-}\pi$, dealing with name creation, boxes, and untyped components, would be a challenging research problem in its own right). In a sequential setting annotated operational semantics have been used by [38]; see also [17]. The definition of the coloured semantics for $\text{box-}\pi$ seems unproblematic, but in general one might validate an annotated semantics by relating it to a lower-level execution model (as mentioned below).

Information flow type systems: The type system differs from previous work [34, 35, 21] primarily in handling badly typed components. Necessarily, it does not provide fine-grain tracking of information flow through these components. It also handles non-determinism, new name creation and channel passing. Precise comparisons with related type systems are difficult as the languages involved differ widely. One can, however, embed fragments of these languages into $\text{box-}\pi$ (noting that this only exploits the fully-typed part of our calculus). For example, in the work of Smith and Volpano [31] an assignment to a low security variable can follow an assignment to a high variable – the program $h:=3; l:=1$ is well-typed. The natural translation of this program in $\text{box-}\pi$ would be

$$\bar{h}^*0 \mid \bar{l}^*0 \quad | \quad h^*y.(\bar{h}^*3 \mid l^*y.\bar{l}^*1)$$

with an initial store assigning 0 to h and l . This would not be well-typed in the system of this paper, taking $h : \mathbf{chan}_{\{H,L\}} \mathbf{Int}$, $l : \mathbf{chan}_{\{L\}} \mathbf{Int}$ and a new base type \mathbf{Int} . Here the low assignment is causally dependent on the high, even though no high information can leak. On the other hand a $\text{box-}\pi$ encoding of branches would not forbid high variable guards.

Causal flow is a robust and straightforward property; it can be enforced by a remarkably simple type system. But, as the example above shows, it is sometimes overconstraining. We envisage that in a large system the bulk of the code will be typeable in a secure type system, a small portion will be in clearly-identified unsafe modules that are subject only to conventional typechecking, and a small portion (any untrusted code) will be encapsulated in wrappers. Automatic type inference would be required to relieve the burden of adding security annotations to all declarations.

8 Conclusion

The issue of securely composing untrusted or partially trusted components has great practical relevance. In this paper we have studied techniques for formally proving that software wrappers – the glue between components – actually enforce user-specified information flow constraints. We have defined a coloured operational semantics for a concurrent wrapper language. By keeping track of all the principals that have affected a process in the semantics it becomes easy to formulate clear statements of information flow properties. To prove that particular wrappers are secure, we defined a causal type system and so only need show that the wrappers are well typed.

Throughout the paper we focussed on wrapper properties – the calculus, statement of security properties and type system are all designed specifically for wrappers – but we believe similar techniques are applicable to other situations in which interaction must be controlled but not completely excluded, for example in isolating a security-critical kernel of a single application, or in controlling interactions between packets in an active network. Allowing untyped code fragments in otherwise typed programs gives a way to loosen security restrictions when necessary.

In future work we intend to integrate the causal type system with a lower-level semantics for object code, such as the typed assembly language of [11]. We also intend to address the issue of type inference of security levels and the statements of properties involving dynamic changes in information flow policy.

Acknowledgements We would like to thank J. Leifer and J. Palsberg for comments. The first author was supported by a Royal Society University Research Fellowship and by EPSRC grant GR/L 62290 *Calculi for Interactive Systems: Theory and Experiment*. The second author did part of this work in the Object System Group at the University of Geneva.

A Coloured LTS and Typing Rules

$\frac{}{A \vdash C : \bar{x}^o v \xrightarrow{\bar{x}^o v}_C 0} \text{ (Out)}$	$\frac{}{A \vdash x^t p.P \xrightarrow{x^t v}_C C \bullet \{v/p\} P} \text{ (c) (In)}$
$\frac{A \vdash P \xrightarrow{\ell}_C P'}{A \vdash P \mid Q \xrightarrow{\ell}_C P' \mid Q} \text{ (Par)}$	$\frac{}{A \vdash !x^t p.P \xrightarrow{x^t v}_C !x^t p.P \mid C \bullet \{v/p\} P} \text{ (c) (Repl)}$
$\frac{A \vdash P \xrightarrow{\bar{x}^\uparrow v}_C P' \quad A \vdash Q \xrightarrow{x^\uparrow v}_C Q'}{A \vdash P \mid Q \xrightarrow{\tau}_\emptyset (\nu \text{fn}(x, v) - A)(P' \mid Q')} \text{ (Comm)}$	
$\frac{A \vdash P \xrightarrow{\bar{x}^\uparrow v}_C P'}{A \vdash n[P] \xrightarrow{\tau}_\emptyset (\nu \text{fn}(x, v) - A)(C : \bar{x}^\uparrow v \mid n[P'])} \text{ (Box-1)}$	
$\frac{}{A \vdash n[P] \xrightarrow{x^\uparrow v}_C n[C : \bar{x}^\uparrow v \mid P]} \text{ (Box-2)}$	$\frac{A \vdash P \xrightarrow{\tau}_C P'}{A \vdash n[P] \xrightarrow{\tau}_C n[P']} \text{ (Box-3)}$
$\frac{A, x \vdash P \xrightarrow{\ell}_C P'}{A \vdash (\nu x)P \xrightarrow{\ell}_C (\nu x)P'} \text{ (a) (Res-1)}$	$\frac{A, x \vdash P \xrightarrow{\bar{y}^o v}_C P'}{A \vdash (\nu x)P \xrightarrow{\bar{y}^o v}_C P'} \text{ (b) (Res-2)}$
$\frac{A \vdash P \xrightarrow{\ell}_C P' \quad P' \equiv P''}{A \vdash P \xrightarrow{\ell}_C P''} \text{ (Struct)}$	
<p>(a) The (Res-1) rule is subject to $x \notin \text{fn}(\ell)$. (b) The (Res-2) rule is subject to $x \in \text{fn}(v) - \text{fn}(y, o)$, if o is \star, \uparrow or \bar{n}, and to $x \in \text{fn}(y, v) - \text{fn}(o)$ otherwise. (c) In the (In) and (Repl) axioms there is a side condition that $\{v/p\}P$ is well-defined. In all rules with conclusion of the form $A \vdash P \xrightarrow{\ell}_C Q$ there is an implicit side condition $\text{fn}(P) \subseteq A$. Symmetric versions of (Par) and (Comm) are elided.</p>	

Figure 2: Coloured Box- π Labelled Transition Semantics

Patterns:	
$\frac{}{\vdash _ : T \triangleright \emptyset}$	$\frac{}{\vdash x : T \triangleright x : \bar{T}}$
$\frac{\vdash p_1 : T_1 \triangleright \Delta_1 \dots \vdash p_k : T_k \triangleright \Delta_k}{\vdash (p_1 \dots p_k) : \langle T_1 \dots T_k \rangle \triangleright \Delta_1, \dots, \Delta_k}$	
Values:	
$\frac{}{\Gamma, x : T \vdash x : \bar{T}}$	$\frac{\Gamma \vdash v_1 : T_1 \dots \Gamma \vdash v_k : T_k}{\Gamma \vdash \langle v_1 \dots v_k \rangle : \langle T_1 \dots T_k \rangle}$
$\frac{\text{fn}(v) \subseteq \text{dom}(\Gamma)}{\Gamma \vdash v : \top}$	$\frac{T \text{ atomic}}{\Gamma, x : T \vdash x : \mathbf{name}}$
Processes:	
$\frac{\begin{array}{l} o \in \{\star, \uparrow, \bar{\uparrow}\} \\ \Gamma \vdash x : \mathbf{chan}_K T \\ \Gamma \vdash v : T \\ C \subseteq K \end{array}}{\Gamma \vdash C : \bar{x}^o v : \mathbf{proc}_K} \quad (\text{Out-}\star, \uparrow, \bar{\uparrow})$	$\frac{\begin{array}{l} \iota \in \{\star, \uparrow\} \\ \Gamma \vdash x : \mathbf{chan}_K T \\ \vdash p : T \triangleright \Delta \\ \Gamma, \Delta \vdash P : \mathbf{proc}_K \end{array}}{\Gamma \vdash x' p.P : \mathbf{proc}_K} \quad (\text{In-}\star, \uparrow)$
$\frac{\begin{array}{l} o \in \{a, \bar{a}\} \\ \Gamma \vdash a : \mathbf{box}_K \\ \Gamma \vdash x : \mathbf{name} \\ \Gamma \vdash v : \top \\ C \subseteq K \end{array}}{\Gamma \vdash C : \bar{x}^o v : \mathbf{proc}_K} \quad (\text{Out-}a, \bar{a})$	$\frac{\begin{array}{l} \Gamma \vdash a : \mathbf{box}_{K'} \\ \Gamma \vdash x : \mathbf{chan}_K S \\ \vdash p : S \triangleright \Delta \\ \Gamma, \Delta \vdash P : \mathbf{proc}_K \\ K' \subseteq K \\ \Delta \text{ flat} \\ P \text{ tests all names of type } \mathbf{name} \text{ in } \Delta \\ p \text{ contains no wildcards} \end{array}}{\Gamma \vdash x^a p.P : \mathbf{proc}_K} \quad (\text{In-}a)$
$\frac{\begin{array}{l} \Gamma \vdash P : \mathbf{proc}_K \\ \Gamma \vdash Q : \mathbf{proc}_{K'} \end{array}}{\Gamma \vdash P \mid Q : \mathbf{proc}_{K \cap K'}} \quad (\text{Par})$	$\frac{\begin{array}{l} \Gamma \vdash n : \mathbf{box}_K \\ \text{pn}(P) \subseteq K \\ \text{fn}(P) \subseteq \text{dom}(\Gamma) \end{array}}{\Gamma \vdash n[P] : \mathbf{proc}_K} \quad (\text{Box})$
$\frac{}{\Gamma \vdash 0 : \mathbf{proc}_K} \quad (\text{Nil})$	$\frac{\begin{array}{l} \Gamma, x : T \vdash P : \mathbf{proc}_K \\ T \text{ atomic} \end{array}}{\Gamma \vdash (\nu x) P : \mathbf{proc}_K} \quad (\text{Res})$
$\frac{\begin{array}{l} \Gamma \vdash P : \mathbf{proc}_{K'} \\ K \subseteq K' \end{array}}{\Gamma \vdash P : \mathbf{proc}_K} \quad (\text{Spec})$	
<p>The replicated input rules are similar to the input rules. The predicate ‘P tests all names of type \mathbf{name} in Δ’ is defined to be true iff for all $y : \mathbf{name}$ in Δ, y occurs free in channel or box position within P.</p>	

Figure 3: Coloured Box- π Typing

B Proofs

This appendix gives the soundness proofs for the type system (of the Subject Reduction and Limited Runtime Error theorems) and the proof that \mathcal{F} has the causal flow property.

B.1 Soundness

The proof of Subject Reduction is divided into three main parts. First we require lemmas giving conditions under which a substitution is well-defined and well-typed (here ‘good’). We then prove substitution lemmas for values and processes by induction on typing derivations, and finally the Subject Reduction result by induction on pairs of transition and typing derivations. The Limited Runtime Error result is almost an immediate consequence of these lemmas.

Say $\Gamma; \Delta \vdash \{u/p\}$ good iff $\{u/p\}$ is well-defined, $\text{dom}(\{u/p\}) = \text{dom}(\Delta)$, and $\forall x: T \in \Delta. \Gamma \vdash \{u/p\}x: T$. We adopt the convention below that wherever $tc(\Gamma, v, T)$ is mentioned it is also assumed well-defined.

Lemma 9 *If $\Gamma = \bigwedge_{i \in 1..n} \Gamma_i$ and for some $j \in 1..n$ $\Gamma_j \vdash v: T$ then $\Gamma \vdash v: T$.*

Proof Induction on derivation of $\Gamma_j \vdash v: T$.

(Var) $v = x$. We have $\Gamma(x) \leq \Gamma_j(x) = T$ so by examination of the value typing rules $\Gamma \vdash x: T$.

(Tuple) By the induction hypothesis.

(Top) We have $\text{fn}(v) \subseteq \text{dom}(\Gamma_j) \subseteq \text{dom}(\Gamma)$ so $\Gamma \vdash v: \top$.

(Name) $v = x$. We have $\Gamma(x) \leq \Gamma_j(x)$ so as the atomic types are down-closed $\Gamma \vdash x: \mathbf{name}$.

□

Lemma 10 (tc) *If $\Theta \stackrel{\text{def}}{=} tc(\Gamma, v, T)$ then Θ atomic and $\Gamma, \Theta \vdash v: T$.*

Proof The first part is by induction on v , noting that the set of atomic types is closed under defined glbs. The second part is also by induction on v .

Case x . Immediate from the definition of $tc(\Gamma, x, _)$.

Case $\langle v_1 .. v_n \rangle$. For $tc(\Gamma, v, T)$ to be defined one of the following two cases must hold.

Case $T = \top$. We have that $\Theta_i \stackrel{\text{def}}{=} tc(\Gamma, v_i, \top)$ are well-defined and $\Theta = \bigwedge_{i \in 1..n} \Theta_i$. By the induction hypothesis $\Gamma, \Theta_i \vdash v_i: \top$, so $\text{fn}(v_i) \subseteq \text{dom}(\Gamma, \Theta_i)$, so $\text{fn}(\langle v_1 .. v_n \rangle) \subseteq \text{dom}(\Gamma, \Theta)$, so $\Gamma, \Theta \vdash \langle v_1 .. v_n \rangle: \top$.

Case $T = \langle T_1 .. T_n \rangle$. We have that $\Theta_i \stackrel{\text{def}}{=} tc(\Gamma, v_i, T_i)$ are well-defined and $\Theta = \bigwedge_{i \in 1..n} \Theta_i$. By the induction hypothesis $\Gamma, \Theta_i \vdash v_i: T_i$. Now for each i Γ and Θ_i have disjoint domains, so $\Gamma, \bigwedge_{i \in 1..n} \Theta_i = \bigwedge_{i \in 1..n} \Gamma, \Theta_i$. By Lemma 9 $\Gamma, \Theta \vdash v_i: T_i$. By the (Tuple) rule $\Gamma, \Theta \vdash \langle v_1 .. v_n \rangle: \langle T_1 .. T_n \rangle$.

□

Lemma 11 *If Γ atomic then $tc(\Gamma, v, \top)$ is well-defined and is equal to the type context mapping each $x \in \text{fn}(v) - \text{dom}(\Gamma)$ to \mathbf{name} .*

Proof Induction on v . □

Lemma 12 (Goodness - Standard - Preliminary) *If*

$$\begin{aligned} & \Gamma \text{ atomic} \\ & \Gamma \vdash u : U \\ & \vdash p : U \triangleright \Delta \end{aligned}$$

then $\Gamma; \Delta \vdash \{u/p\}$ good.

Proof By induction on the two typing derivations. Consider the last rule of the pattern judgement.

(Wild) $\vdash _ : U \triangleright \emptyset$. Clearly $\{u/_ \} = \{\}$; the second clause of the definition of good is trivial; the third is vacuous.

(Var) $\vdash y : U \triangleright y : U$. Clearly $\{u/y\} = \{y \mapsto u\}$, $\text{dom}(\{y \mapsto u\}) = \text{dom}(y : U)$, and $\Gamma \vdash \{u/y\}y : U$.

(Tuple) $\vdash (p_1 \dots p_k) : \langle U_1 \dots U_k \rangle \triangleright \Delta_1, \dots, \Delta_k$. As Γ atomic the last rule of the value judgement must be the value (Tuple) rule, with conclusion $\Gamma \vdash \langle u_1 \dots u_k \rangle : \langle U_1 \dots U_k \rangle$. By the induction hypothesis $\Gamma; \Delta_i \vdash \{u_i/p_i\}$ good for each i . By the definition of substitution $\{\langle u_1 \dots u_k \rangle / (p_1 \dots p_k)\}$ is defined and equal to $\bigcup_{i \in 1..k} \{u_i/p_i\}$. As for each i $\text{dom}(\{u_i/p_i\}) = \text{dom}(\Delta_i)$, we have $\text{dom}(\bigcup_{i \in 1..k} \{u_i/p_i\}) = \text{dom}(\Delta_1, \dots, \Delta_k)$. To check $\Gamma; \Delta_1, \dots, \Delta_k \vdash \{\langle u_1 \dots u_k \rangle / (p_1 \dots p_k)\}$ good it remains only to observe that for each $z : T \in \Delta_1, \dots, \Delta_k$ there is an i with $z : T \in \Delta_i$, hence $\Gamma \vdash \{u_i/p_i\}z : T$, hence $\Gamma \vdash (\bigcup_{i \in 1..k} \{u_i/p_i\})z : T$. □

Note that this result requires that the range of Γ contains no tuple types. Consider $u = x$, $U = \langle \mathbf{box}_K \mathbf{box}_K \rangle$, $\Gamma = x : U$ and $p = (y z)$. We have $\vdash p : U \triangleright y : \mathbf{box}_K, z : \mathbf{box}_K$ but $\{x/(y z)\}$ is not well-defined.

Lemma 13 (Goodness - Dynamic) *If*

$$\begin{aligned} & \Gamma \text{ atomic} \\ & \vdash p : S \triangleright \Delta \\ & \text{dom}(\Gamma) \text{ and } \text{dom}(\Delta) \text{ disjoint} \\ & \{u/p\} \text{ well defined} \\ & \Delta \text{ flat (so } \text{ran}(\Delta) \subseteq \{\top, \mathbf{name}\}) \\ & \forall y : \mathbf{name} \in \Delta . \{u/p\}y \text{ is a name} \\ & p \text{ contains no wildcards} \end{aligned}$$

then $\Theta \stackrel{\text{def}}{=} tc(\Gamma, u, S)$ is well-defined, $\text{ran}(\Theta) \subseteq \{\mathbf{name}\}$, and $\Gamma, \Theta; \Delta \vdash \{u/p\}$ good.

Proof By induction on the pattern p .

Case $_$. Contradicts the premises.

Case y . We have $\Delta = y : S$. Trivially $\{u/y\}$ has domain $\text{dom}(\Delta)$. As S transmissible and Δ flat either $S = \mathbf{name}$ or $S = \top$.

Case $S = \mathbf{name}$. We have $\{u/y\}y = x$ for some name x and $u = x$.

If $x \notin \text{dom}(\Gamma)$ then $tc(\Gamma, x, \mathbf{name}) = x : \mathbf{name}$.

If $x \in \text{dom}(\Gamma)$ then by Γ atomic we have $\Gamma \vdash x : \mathbf{name}$, so $tc(\Gamma, x, \mathbf{name}) = \emptyset$.

In either case we have $\Gamma, \Theta \vdash \{x/y\}y : S$, so $\Gamma, \Theta; \Delta \vdash \{x/y\}$ good.

Case $S = \top$. By Lemma 11 Θ is well-defined and maps each $x \in \text{fn}(u) - \text{dom}(\Gamma)$ to **name**, so $\Gamma, \Theta \vdash \{u/y\}y : \top$, so $\Gamma, \Theta; \Delta \vdash \{u/y\}$ good.

Case $(p_1 \dots p_k)$. We have $S = \langle S_1 \dots S_k \rangle$, $\Delta = \Delta_1, \dots, \Delta_k$, and $\vdash p_i : S_i \triangleright \Delta_i$ for $i \in 1..k$. By $\{u/(p_1 \dots p_k)\}$ well-defined we know there exist u_1, \dots, u_k such that $u = \langle u_1 \dots u_k \rangle$, each $\{u_i/p_i\}$ is well-defined and $\{u/(p_1 \dots p_k)\} = \bigcup_i \{u_i/p_i\}$.

By the induction hypothesis each $\Theta_i \stackrel{\text{def}}{=} \text{tc}(\Gamma, u_i, S_i)$ is well-defined, $\text{ran}(\Theta_i) \subseteq \{\mathbf{name}\}$, and $\Gamma, \Theta_i; \Delta_i \vdash \{u_i/p_i\}$ good. $\Theta \stackrel{\text{def}}{=} \bigwedge_i \Theta_i$ is then well-defined, with domain $\bigcup_i \text{dom}(\Theta_i)$ and maps all names in its domain to **name**.

For $\Gamma, \Theta; \Delta \vdash \{\langle u_1 \dots u_k \rangle / (p_1 \dots p_k)\}$ good the well-definedness of the substitution is a premise; the fact that it has domain $\text{dom}(\Delta)$ follows from $\Gamma, \Theta_i; \Delta_i \vdash \{u_i/p_i\}$ good. Consider $x : T \in \Delta$. For some $i \in 1..k$ we have $x : T \in \Delta_i$, so $\Gamma, \Theta_i \vdash \{u_i/p_i\}x : T$. By Lemma 9 $\Gamma, \Theta \vdash \{u_i/p_i\}x : T$. We have $x \in \text{dom}(\{u_i/p_i\})$ so $\{\langle u_1 \dots u_k \rangle / (p_1 \dots p_k)\}x = \{u_i/p_i\}x$, hence $\Gamma, \Theta \vdash \{u/p\}x : T$.

□

Lemma 14 (Substitution - values) *If*

$$\begin{array}{l} \Gamma, \Delta \vdash v : T \\ \Gamma, \Theta; \Delta \vdash \{u/p\} \text{ good} \end{array}$$

then $\{u/p\}v$ *is well-defined and* $\Gamma, \Theta \vdash \{u/p\}v : T$

Proof By $\Gamma, \Theta; \Delta \vdash \{u/p\}$ good we have that $\{u/p\}$ is well-defined, so $\{u/p\}v$ is well-defined. The second part is proved by induction on the value typing derivation.

(Var) $\Gamma, \Delta \vdash x : T$.

Case $x : T \in \Gamma$. By goodness the domain of $\{u/p\}$ coincides with that of Δ and hence is disjoint from Γ , so $\{u/p\}x = x$, so $\Gamma, \Theta \vdash \{u/p\}x : T$.

Case $x : T \in \Delta$. By goodness $\Gamma, \Theta \vdash \{u/p\}x : T$.

(Tuple) $\Gamma, \Delta \vdash \langle v_1 \dots v_k \rangle : \langle T_1 \dots T_k \rangle$. By the induction hypothesis for each i we have $\Gamma, \Theta \vdash \{u/p\}v_i : T_i$ hence $\Gamma, \Theta \vdash \{u/p\}\langle v_1 \dots v_k \rangle : \langle T_1 \dots T_k \rangle$

(Top) $\Gamma, \Delta \vdash v : \top$. We have $\text{fn}(v) \subseteq \text{dom}(\Gamma, \Delta)$. By goodness $\text{dom}(\{u/p\}) = \text{dom}(\Delta)$ and $\text{fn}(\text{ran}(\{u/p\})) \subseteq \text{dom}(\Gamma, \Theta)$, so $\text{fn}(\{u/p\}v) \subseteq \text{dom}(\Gamma, \Theta)$, hence $\Gamma, \Theta \vdash \{u/p\}v : \top$.

(Name) $\Gamma, \Delta \vdash x : \mathbf{name}$. We have $x : T \in \Gamma, \Delta$ for some T atomic.

Case $x : T \in \Gamma$. By goodness the domain of $\{u/p\}$ coincides with that of Δ and hence is disjoint from Γ , so $\{u/p\}x = x$, so $\Gamma, \Theta \vdash \{u/p\}x : \mathbf{name}$.

Case $x : T \in \Delta$. By goodness $\Gamma, \Theta \vdash \{u/p\}x : T$. As T atomic this must be derivable using the (Var) or (Name) rules. In the (Var) case $\{u/p\}x = y$ for some name y with $y : T \in \Gamma, \Theta$, so $\Gamma, \Theta \vdash y : \mathbf{name}$. In the (Name) case $T = \mathbf{name}$ so we are done.

□

Note that for this to hold the typing rules must ensure that names of tuple types do not have type **name**. Note also that this lemma does not require any atomicity, and that that is important in the first input clause of the process substitution lemma.

Lemma 15 (Substitution - processes) *If*

$$\begin{array}{l} \Gamma, \Delta \vdash P : \mathbf{proc}_K \\ \Gamma, \Theta; \Delta \vdash \{^u/p\} \text{ good} \end{array}$$

then (1) if $\{^u/p\}P$ is well-defined then $\Gamma, \Theta \vdash \{^u/p\}P : \mathbf{proc}_K$ and (2) if P contains no subterm $n[Q]$ then $\{^u/p\}P$ is well-defined.

Proof We prove both parts simultaneously by induction on the size of type derivation for P . For (1) we give two instances of each typing rule; in each case showing that the premises of the right-hand instance follow from those of the left-hand instance.

(Out- $\star, \uparrow, \bar{\uparrow}$)

$$\frac{\begin{array}{l} o \in \{\star, \uparrow, \bar{\uparrow}\} \\ \Gamma, \Delta \vdash x : \mathbf{chan}_K T \\ \Gamma, \Delta \vdash v : T \\ C \subseteq K \end{array}}{\Gamma, \Delta \vdash C : \bar{x}^o v : \mathbf{proc}_K} \quad (\text{Out-}\star, \uparrow, \bar{\uparrow}) \quad \frac{\begin{array}{l} o \in \{\star, \uparrow, \bar{\uparrow}\} \\ \Gamma, \Theta \vdash \{^u/p\}x : \mathbf{chan}_K T \\ \Gamma, \Theta \vdash \{^u/p\}v : T \\ C \subseteq K \end{array}}{\Gamma, \Theta \vdash C : \overline{\{^u/p\}x}^o \{^u/p\}v : \mathbf{proc}_K} \quad (\text{Out-}\star, \uparrow, \bar{\uparrow})$$

For (1) we use Lemma 14 for x and v , so all the premises of the right-hand instance hold. For (2) by $\Gamma, \Theta \vdash \{^u/p\}x : \mathbf{chan}_K T$ we have that $\{^u/p\}x$ is a name, so $\{^u/p\}C : \bar{x}^o v$ is defined and is equal to the conclusion of the right-hand instance.

(Out- a, \bar{a})

$$\frac{\begin{array}{l} o \in \{a, \bar{a}\} \\ \Gamma, \Delta \vdash a : \mathbf{box}_K \\ \Gamma, \Delta \vdash x : \mathbf{name} \\ \Gamma, \Delta \vdash v : T \\ C \subseteq K \end{array}}{\Gamma, \Delta \vdash C : \bar{x}^o v : \mathbf{proc}_K} \quad (\text{Out-}a, \bar{a}) \quad \frac{\begin{array}{l} \{^u/p\}o \in \{\{^u/p\}a, \overline{\{^u/p\}a}\} \\ \Gamma, \Theta \vdash \{^u/p\}a : \mathbf{box}_K \\ \Gamma, \Theta \vdash \{^u/p\}x : \mathbf{name} \\ \Gamma, \Theta \vdash \{^u/p\}v : T \\ C \subseteq K \end{array}}{\Gamma, \Theta \vdash C : \overline{\{^u/p\}x}^{\{^u/p\}o} \{^u/p\}v : \mathbf{proc}_K} \quad (\text{Out-}a, \bar{a})$$

For (1) we use Lemma 14 for a , x and v , so the middle 3 premises of the right-hand instance hold. By $\Gamma, \Theta \vdash \{^u/p\}a : \mathbf{box}_K$ we have that $\{^u/p\}a$ is a name, so the first premise holds. For (2) note also that by $\Gamma, \Theta \vdash \{^u/p\}x : \mathbf{chan}_K T$ we have that $\{^u/p\}x$ is a name, so $\{^u/p\}C : \bar{x}^o v$ is defined and is equal to the conclusion of the right-hand instance.

(In- \star, \uparrow)

$$\frac{\begin{array}{l} \iota \in \{\star, \uparrow\} \\ \Gamma, \Delta \vdash x : \mathbf{chan}_K T \\ \vdash q : T \triangleright \Phi \\ \Gamma, \Delta, \Phi \vdash P : \mathbf{proc}_K \end{array}}{\Gamma, \Delta \vdash x^\iota q.P : \mathbf{proc}_K} \quad (\text{In-}\star, \uparrow) \quad \frac{\begin{array}{l} \iota \in \{\star, \uparrow\} \\ \Gamma, \Theta \vdash (\{^u/p\}x) : \mathbf{chan}_K T \\ \vdash q : T \triangleright \Phi \\ \Gamma, \Theta, \Phi \vdash \{^u/p\}P : \mathbf{proc}_K \end{array}}{\Gamma, \Theta \vdash (\{^u/p\}x)^\iota q.\{^u/p\}P : \mathbf{proc}_K} \quad (\text{In-}\star, \uparrow)$$

We assume without loss of generality that $\text{dom}(\Phi)$ and $\text{dom}(\Theta)$ are disjoint (but elide the result on renaming of typing derivations that justifies this).

For (1) we use Lemma 14 for x , so the second premises of the right-hand instance holds, and further $\{^u/p\}x$ is a name. By the definition of Good we have $\Gamma, \Theta, \Phi; \Delta \vdash$

$\{^u/p\}$ good, so by the induction hypothesis the fourth premise of the right-hand instance holds.

For (2) we noted above that $\{^u/p\}x$ is a name, and by the induction hypothesis $\{^u/p\}P$ is well-defined, so $\{^u/p\}(x^i q.P)$ is well-defined. By the disjointness assumption it is equal to the conclusion of the right-hand side.

(In-a)

$$\begin{array}{c}
\Gamma, \Delta \vdash a : \mathbf{box}_{K'} \\
\Gamma, \Delta \vdash x : \mathbf{chan}_{K'} S \\
\vdash q : S \triangleright \Phi \\
\Gamma, \Delta, \Phi \vdash P : \mathbf{proc}_K \\
K' \subseteq K \\
\Phi \text{ flat} \\
P \text{ tests all } y : \mathbf{name} \text{ in } \Phi \\
p \text{ contains no wildcards} \\
\hline
\Gamma, \Delta \vdash x^a q.P : \mathbf{proc}_K
\end{array}
\quad (\text{In-a}) \quad
\begin{array}{c}
\Gamma, \Theta \vdash \{^u/p\}a : \mathbf{box}_{K'} \\
\Gamma, \Theta \vdash \{^u/p\}x : \mathbf{chan}_{K'} S \\
\vdash q : S \triangleright \Phi \\
\Gamma, \Theta, \Phi \vdash \{^u/p\}P : \mathbf{proc}_K \\
K' \subseteq K \\
\Phi \text{ flat} \\
\{^u/p\}P \text{ tests all } y : \mathbf{name} \text{ in } \Phi \\
p \text{ contains no wildcards} \\
\hline
\Gamma, \Theta \vdash \{^u/p\}x^{\{^u/p\}a} q. \{^u/p\}P : \mathbf{proc}_K
\end{array}
\quad (\text{In-a})$$

For (1), as in the previous case we have the right-hand side requirements on a , x , and P . The last right-hand premise follows from the disjointness of $\text{dom}(\Phi)$ and $\text{dom}(\{^u/p\})$.

For (2), by the typing of $\{^u/p\}a$ and $\{^u/p\}x$ both must be names; it remains only to use the induction hypothesis.

(Par)

$$\frac{\Gamma, \Delta \vdash P : \mathbf{proc}_K \quad \Gamma, \Delta \vdash Q : \mathbf{proc}_{K'}}{\Gamma, \Delta \vdash P \mid Q : \mathbf{proc}_{K \cap K'}} \quad (\text{Par}) \quad
\frac{\Gamma, \Theta \vdash \{^u/p\}P : \mathbf{proc}_K \quad \Gamma, \Theta \vdash \{^u/p\}Q : \mathbf{proc}_{K'}}{\Gamma, \Theta \vdash \{^u/p\}P \mid \{^u/p\}Q : \mathbf{proc}_{K \cap K'}} \quad (\text{Par})$$

Straightforward use of the induction hypothesis.

(Box)

$$\frac{\Gamma, \Delta \vdash n : \mathbf{box}_K \quad \text{pn}(P) \subseteq K \quad \text{fn}(P) \subseteq \text{dom}(\Gamma, \Delta)}{\Gamma, \Delta \vdash n[P] : \mathbf{proc}_K} \quad (\text{Box}) \quad
\frac{\Gamma, \Theta \vdash (\{^u/p\}n) : \mathbf{box}_K \quad \text{pn}(\{^u/p\}P) \subseteq K \quad \text{fn}(\{^u/p\}P) \subseteq \text{dom}(\Gamma, \Theta)}{\Gamma, \Theta \vdash (\{^u/p\}n)[\{^u/p\}P] : \mathbf{proc}_K} \quad (\text{Box})$$

For (1) we use Lemma 14 for n , so the first premise of the right-hand instance holds. The second is immediate as value-for-pattern substitutions do not affect principal sets. The third follows from the definition of goodness. Clause (2) is vacuously true.

(Nil)

$$\frac{}{\Gamma, \Delta \vdash 0 : \mathbf{proc}_K} \quad (\text{Nil}) \quad
\frac{}{\Gamma, \Theta \vdash 0 : \mathbf{proc}_K} \quad (\text{Nil})$$

Immediate.

(Res)

$$\frac{\Gamma, \Delta, x : T \vdash P : \mathbf{proc}_K \quad T \text{ atomic}}{\Gamma, \Delta \vdash (\nu x)P : \mathbf{proc}_K} \quad (\text{Res}) \quad
\frac{\Gamma, \Theta, x : T \vdash \{^u/p\}P : \mathbf{proc}_K \quad T \text{ atomic}}{\Gamma, \Theta \vdash (\nu x)\{^u/p\}P : \mathbf{proc}_K} \quad (\text{Res})$$

We assume without loss of generality that $x \notin \text{dom}(\Theta)$. By the definition of Good we have $\Gamma, x : T, \Theta; \Delta \vdash \{^u/p\}$ good, so by the induction hypothesis the first premise of the right-hand instance holds.

(Spec)

$$\frac{\Gamma, \Delta \vdash P : \mathbf{proc}_{K'} \quad K \subseteq K'}{\Gamma, \Delta \vdash P : \mathbf{proc}_K} \text{ (Spec)} \quad \frac{\Gamma, \Theta \vdash \{^u/p\}P : \mathbf{proc}_{K'} \quad K \subseteq K'}{\Gamma, \Theta \vdash \{^u/p\}P : \mathbf{proc}_K} \text{ (Spec)}$$

Straightforward use of the induction hypothesis. □

To see the need for the condition that P is box-free, consider $\Gamma = \emptyset, \Delta = x : \langle \mathbf{name name} \rangle, \Theta = z : \mathbf{name}, P = (\nu n)n[\bar{x}^*(\cdot)]$, and $\{^u/p\} = \{^z/z\}/x$. The premises of the Lemma hold, but $\{^u/p\}P$ is not well-defined.

Lemma 16 (Painting – Jackson Pollack style) *If*

$$\begin{array}{l} \Gamma \vdash P : \mathbf{proc}_K \\ C \subseteq K \end{array}$$

then $\Gamma \vdash C \circ P : \mathbf{proc}_K$

Proof Routine induction on typing derivations. □

Lemma 17 *If* $\Gamma \vdash v : T$ *then* $tc(\Gamma, v, T) = \emptyset$.

Proof Induction on v . □

Lemma 18 *If* $\text{dom}(\Delta)$ *is disjoint from* $\text{dom}(\Gamma)$, $\text{fn}(v)$ *and* $\text{fn}(P)$ *then*

1. $\Gamma \vdash v : T \iff \Gamma, \Delta \vdash v : T$.
2. $tc(\Gamma, v, T) = tc((\Gamma, \Delta), v, T)$.
3. $\Gamma \vdash P : \mathbf{proc}_K \iff \Gamma, \Delta \vdash P : \mathbf{proc}_K$

Proof Routine inductions. □

Lemma 19 *If* $tc((\Gamma, y : U), v, T)$ *well-defined and* $y \in \text{fn}(v)$ *then there exists some* V *with* $U \leq V$ *and* $tc(\Gamma, v, T) = tc((\Gamma, y : U), v, T), y : V$.

Proof Induction on v .

Case x . We must have $x = y$. From the tc definition $tc((\Gamma, y : U), y, T) = \emptyset$ and $\Gamma, y : U \vdash y : T$, hence $U \leq T$. From the tc definition $tc(\Gamma, y, T) = y : T$, so take $V = T$.

Case $\langle v_1 \dots v_n \rangle$.

Case $T = \langle T_1, \dots, T_n \rangle$. We have $tc((\Gamma, y : U), v, T) = \bigwedge_i tc((\Gamma, y : U), v_i, T_i)$.

Let $J = \{j \mid y \in \text{fn}(v_j)\}$.

For all $j \in J$ by the induction hypothesis we have V_j with $U \leq V_j$ and $tc(\Gamma, v_j, T_j) = tc((\Gamma, y : U), v_j, T_j), y : V_j$.

For all $i \notin J$ by Lemma 18.2 we have $tc(\Gamma, v_i, T_i) = tc((\Gamma, y : U), v_i, T_i)$.

The types $\{V_j \mid j \in J\}$ have a lower bound U so a $\text{glb } V \stackrel{\text{def}}{=} \bigwedge_{j \in J} V_j$ exists and $U \leq V$.

It follows that $\bigwedge_i tc(\Gamma, v_i, T_i)$ exists and moreover is equal to $(\bigwedge_i tc((\Gamma, y : U), v_i, T_i)), y : V$, so $tc(\Gamma, \langle v_1 \dots v_n \rangle, \langle T_1 \dots T_n \rangle) = tc((\Gamma, y : U), \langle v_1 \dots v_n \rangle, \langle T_1 \dots T_n \rangle), y : V$.

Case $T = \top$. This is similar to the previous case. We have $tc((\Gamma, y : U), v, \top) = \bigwedge_i tc((\Gamma, y : U), v_i, \top)$.

Let $J = \{j \mid y \in \text{fn}(v_j)\}$.

For all $j \in J$ by the induction hypothesis we have V_j with $U \leq V_j$ and $tc(\Gamma, v_j, \top) = tc((\Gamma, y : U), v_j, \top), y : V_j$.

For all $i \notin J$ by Lemma 18.2 we have $tc(\Gamma, v_i, \top) = tc((\Gamma, y : U), v_i, \top)$.

The types $\{V_j \mid j \in J\}$ have a lower bound U so a $\text{glb } V \stackrel{\text{def}}{=} \bigwedge_{j \in J} V_j$ exists and $U \leq V$.

It follows that $\bigwedge_i tc(\Gamma, v_i, \top)$ exists and moreover is equal to $(\bigwedge_i tc((\Gamma, y : U), v_i, \top)), y : V$, so $tc(\Gamma, \langle v_1 \dots v_n \rangle, \top) = tc((\Gamma, y : U), \langle v_1 \dots v_n \rangle, \top), y : V$.

□

Lemma 20 *If $\Gamma, y : U$ atomic, $\Theta \leq tc((\Gamma, y : U), v, T)$ and $y \in \text{fn}(v)$ then $\Theta, y : U \leq tc(\Gamma, v, T)$.*

Proof An immediate corollary of Lemma 19, which gives that there exists some V with $U \leq V$ and $tc(\Gamma, v, T) = tc((\Gamma, y : U), v, T), y : V$. □

Lemma 21 (Structural Congruence) *If $\Gamma \vdash P : \text{proc}_\kappa$ and $P \equiv Q$ then $\Gamma \vdash Q : \text{proc}_\kappa$.*

Proof Induction on derivations of $P \equiv Q$. □

Say $\Gamma \leq \Gamma'$ iff $\text{dom}(\Gamma) = \text{dom}(\Gamma')$ and $\forall x \in \text{dom}(\Gamma) . \Gamma(x) \leq \Gamma'(x)$.

Lemma 22 *If $\Gamma \leq \Gamma'$ and $\Gamma' \vdash v : T$ then $\Gamma \vdash v : T$.*

Proof Induction on typing derivation of v . □

Lemma 23 *If $\Gamma \leq \Gamma'$ and $\Gamma'; \Delta \vdash \{u/p\}$ good then $\Gamma; \Delta \vdash \{u/p\}$ good.*

Proof By the definition of Good and Lemma 22. □

Lemma 24 *If $A \vdash P \xrightarrow{\bar{x}^u}_C Q$ then $C \subseteq \text{pn}(P)$ and $\text{pn}(Q) \subseteq \text{pn}(P)$.*

Proof Routine induction on transition derivations. □

We can now restate and prove Theorem 5.

Theorem 25 (subject reduction)

1. If $\Gamma \vdash_K P \xrightarrow{\bar{x}^o v}_C Q$ and $o \in \{\star, \bar{\uparrow}\}$ then for some K', T

$$\Gamma \vdash x : \mathbf{chan}_{K'} T$$

there exists $\Theta \leq tc(\Gamma, v, T)$ such that $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$
 $C \subseteq K'$
2. If $\Gamma \vdash_K P \xrightarrow{\bar{x}^i v}_C Q$ then for some K', T

there exists $\Theta \leq tc(\Gamma, \langle x v \rangle, \langle \mathbf{chan}_{K'} T T \rangle)$ such that $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$
 $C \subseteq K'$
3. If $\Gamma \vdash_K P \xrightarrow{\bar{x}^a v}_C Q$ then for some K'

$$\Gamma \vdash a : \mathbf{box}_{K'}$$

there exists $\Theta \leq tc(\Gamma, \langle x v \rangle, \langle \mathbf{name}, \top \rangle)$ such that $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$
 $C \subseteq K'$
4. If $\Gamma \vdash_K P \xrightarrow{\bar{x}^{\bar{v}} v}_C Q$ then for some K'

$$\Gamma \vdash a : \mathbf{box}_{K'}$$

$$\Gamma \vdash x : \mathbf{name}$$

there exists $\Theta \leq tc(\Gamma, v, \top)$ such that $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$
 $C \subseteq K'$
5. If $\Gamma \vdash_K P \xrightarrow{\bar{x}^\gamma v}_C Q$ and $\gamma \in \{\star, \bar{\uparrow}\}$ then for some K', T $\Gamma \vdash x : \mathbf{chan}_{K'} T$. If moreover
$$\Theta \leq tc(\Gamma, v, T)$$

$$C \subseteq K'$$

then $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$.
6. If $\Gamma \vdash_K P \xrightarrow{\bar{x}^a v}_C Q$ then for some $K' \subseteq K''$, and S we have $\Gamma \vdash a : \mathbf{box}_{K'}$, $\Gamma \vdash x : \mathbf{chan}_{K''} S$, $tc(\Gamma, v, S)$ well-defined, and $\text{ran}(tc(\Gamma, v, S)) \subseteq \{\mathbf{name}\}$. If moreover
$$\Theta \leq tc(\Gamma, v, S)$$

$$C \subseteq K''$$

then $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$.
7. If $\Gamma \vdash_K P \xrightarrow{\bar{x}^{\bar{v}} v}_C Q$ then for some K' we have $\Gamma \vdash a : \mathbf{box}_{K'}$. If moreover
$$\Theta \leq tc(\Gamma, \langle x v \rangle, \langle \mathbf{name}, \top \rangle)$$

$$C \subseteq K'$$

then $\Gamma, \Theta \vdash Q : \mathbf{proc}_K$.
8. If $\Gamma \vdash_K P \xrightarrow{\bar{\tau}}_C Q$ then $C = \emptyset$ and $\Gamma \vdash Q : \mathbf{proc}_K$.

Proof We give first the output part, for clauses 1–4, then the input part, for 5–7, then the tau part, for 8. Each is by induction on pairs of transition and typing derivations.

Output Consider the last pair of rules used:

(Out), (Out- \star , \uparrow , $\bar{\uparrow}$) Take $K' = K$ and consider cases of o :

1: $\star, \bar{\uparrow}$ By Lemma 17 we can take $\Theta = tc(\Gamma, v, T) = \emptyset$.

2: \uparrow By Lemma 17 we can take $\Theta = tc(\Gamma, \langle x v \rangle, \langle \mathbf{chan}_{K'} T T \rangle) = \emptyset$.

(Out), (Out- a , \bar{a}) Take $K' = K$ and consider cases of o :

3: a By Lemma 17 we can take $\Theta = tc(\Gamma, \langle x v \rangle, \langle \mathbf{name}, \top \rangle) = \emptyset$.

4: \bar{a} By Lemma 17 we can take $\Theta = tc(\Gamma, v, \top) = \emptyset$.

(Struct), (*) 1–4 follow from the same clauses of the induction hypothesis and Lemma 21.

(*), (Spec) 1–4 follow from the same clauses of the induction hypothesis and a use of (Spec) for Q .

(Par), (Par) Consider $\Gamma \vdash_K P \mid P' \xrightarrow{\ell}_C Q \mid P'$ with $\Gamma \vdash_K P \xrightarrow{\ell}_C Q$. 1–4 follow from the same clauses of the induction hypothesis and a use of Lemma 18.3 for P' .

(Res-1), (Res) Consider $\Gamma \vdash_K (\nu y)P \xrightarrow{\ell}_C (\nu y)Q$ with $\Gamma, y:U \vdash_K P \xrightarrow{\ell}_C Q$, ℓ an output $\bar{x}^o v$ and $y \notin \text{fn}(\ell)$. Suppose $o \in \{\star, \bar{\uparrow}\}$. By clause 1 of the induction hypothesis for some K', T

$$\begin{array}{l} \Gamma, y:U \vdash x: \mathbf{chan}_{K'} T \\ \text{there exists } \Theta \leq tc((\Gamma, y:U), v, T) \text{ such that } \Gamma, y:U, \Theta \vdash Q: \mathbf{proc}_K \\ C \subseteq K' \end{array}$$

By Lemma 18.1 $\Gamma \vdash x: \mathbf{chan}_{K'} T$. By Lemma 18.2 $tc((\Gamma, y:U), v, T) = tc(\Gamma, v, T)$, so taking the same Θ and using the (Res) typing rule we have $\Gamma, \Theta \vdash (\nu y)Q: \mathbf{proc}_K$ as required. The other cases of o are similar.

(Res-2), (Res) Consider $\Gamma \vdash_K (\nu y)P \xrightarrow{\ell}_C Q$ with $\Gamma, y:U \vdash_K P \xrightarrow{\ell}_C Q$, ℓ an output $\bar{x}^o v$ and $y \in \text{fn}(\ell)$.

Case $\neg \text{mv}(o)$. We have $o \in \{\star, \bar{\uparrow}, \bar{a}\}$ and $y \in \text{fn}(v) - \text{fn}(x, o)$.

Suppose $o \in \{\star, \bar{\uparrow}\}$. By clause 1 of the induction hypothesis for some K', T

$$\begin{array}{l} \Gamma, y:U \vdash x: \mathbf{chan}_{K'} T \\ \text{there exists } \Theta \leq tc((\Gamma, y:U), v, T) \text{ such that } \Gamma, y:U, \Theta \vdash Q: \mathbf{proc}_K \\ C \subseteq K' \end{array}$$

By Lemma 18.1 $\Gamma \vdash x: \mathbf{chan}_{K'} T$. By Lemma 20 $\Theta, y:U \leq tc(\Gamma, v, T)$. The case $o = \bar{a}$, for clause 4, is similar.

Case $\text{mv}(o)$. We have $o \in \{\uparrow, a\}$ and $y \in \text{fn}(x, v) - \text{fn}(o)$.

Suppose $o = \uparrow$. By clause 2 of the induction hypothesis for some K', T

$$\begin{array}{l} \text{there exists } \Theta \leq tc((\Gamma, y:U), \langle x v \rangle, \langle \mathbf{chan}_{K'} T T \rangle) \text{ such that } \Gamma, y:U, \Theta \vdash Q: \mathbf{proc}_K \\ C \subseteq K' \end{array}$$

By Lemma 20 $\Theta, y:U \leq tc(\Gamma, \langle x v \rangle, \langle \mathbf{chan}_{K'} T T \rangle)$. The case $o = a$, for clause 3, is similar.

Input Consider the last pair of rules used:

(In)(In-*, ↑) Clause 5. Take $K' = K$. By $tc(\Gamma, v, T)$ defined and Lemma 10 we have $tc(\Gamma, v, T)$ atomic and $\Gamma, tc(\Gamma, v, T) \vdash v : T$. It follows that Θ atomic and by Lemma 22 $\Gamma, \Theta \vdash v : T$. By Lemma 12 $\Gamma, \Theta; \Delta \vdash \{v/p\}$ good. By the definition of labelled transitions $\{v/p\}P$ is well-defined so by Lemma 15 $\Gamma, \Theta \vdash \{v/p\}P : \mathbf{proc}_K$. By Lemma 16 $\Gamma, \Theta \vdash C \circ \{v/p\}P : \mathbf{proc}_K$.

(In)(In-a) Clause 6. Take $K'' = K$. By the definition of labelled transitions $\{v/p\}P$ is well-defined so $\{v/p\}$ is well-defined. As P tests all $y : \mathbf{name} \in \Delta$ and $\{v/p\}P$ is defined $\forall y : \mathbf{name} \in \Delta . \{v/p\}y$ is a name. By Lemma 13 $tc(\Gamma, v, S)$ is well-defined, $\text{ran}(tc(\Gamma, v, S)) \subseteq \{\mathbf{name}\}$, and $\Gamma, tc(\Gamma, v, S); \Delta \vdash \{v/p\}$ good. By Lemma 23 $\Gamma, \Theta; \Delta \vdash \{v/p\}$ good. By Lemma 15 $\Gamma, \Theta \vdash \{v/p\}P : \mathbf{proc}_K$. By Lemma 16 $\Gamma, \Theta \vdash C \circ \{v/p\}P : \mathbf{proc}_K$.

(Repl)(Repl-*, ↑) and (Repl)(Repl-a) Similar to the two cases above.

(Box-2)(Box) Clause 7. Take $K' = K$. To check $\Gamma, \Theta \vdash n[C : \bar{x}^\top v \mid P] : \mathbf{proc}_K$ observe that $\Gamma, \Theta \vdash n : \mathbf{box}_K$ by weakening, $\text{pn}(C : \bar{x}^\top v \mid P) \subseteq C \cup \text{pn}(P) \subseteq K$, and $\text{fn}(C : \bar{x}^\top v \mid P) \subseteq \text{fn}(x, v) \cup \text{fn}(P) \subseteq \text{dom}(\Gamma, \Theta)$.

(Struct), (*) 5–7 follow from the same clauses of the induction hypothesis and Lemma 21.

(*), (Spec) 5–7 follow from the same clauses of the induction hypothesis and a use of (Spec) for Q .

(Par), (Par) Consider $\Gamma \vdash_K P \mid P' \xrightarrow{\ell}_C Q \mid P'$ with $\Gamma \vdash_K P \xrightarrow{\ell}_C Q$. 5–7 follow from the same clauses of the induction hypothesis and a use of Lemma 18.3 for P' .

(Res-1)(Res) Consider $\Gamma \vdash_K (\nu y)P \xrightarrow{\ell}_C (\nu y)Q$ with $\Gamma, y : U \vdash_K P \xrightarrow{\ell}_C Q$, ℓ an input $x^\gamma v$ and $y \notin \text{fn}(\ell)$. Suppose $\gamma \in \{\star, \uparrow\}$. By clause 5 of the induction hypothesis for some K', T $\Gamma, y : U \vdash x : \mathbf{chan}_{K'} T$ and

$$\begin{aligned} \Theta &\leq tc((\Gamma, y : U), v, T) \\ C &\subseteq K' \end{aligned}$$

implies $\Gamma, y : U, \Theta \vdash Q : \mathbf{proc}_K$. By $y \notin \text{fn}(\ell)$ we have $\Gamma \vdash x : \mathbf{chan}_{K'} T$. Now suppose

$$\begin{aligned} \Theta &\leq tc(\Gamma, v, T) \\ C &\subseteq K' \end{aligned}$$

By Lemma 18.2 $tc(\Gamma, v, T) = tc((\Gamma, y : U), v, T)$ so by the implication in the induction hypothesis $\Gamma, y : U, \Theta \vdash Q : \mathbf{proc}_K$, hence $\Gamma, \Theta \vdash (\nu y)Q : \mathbf{proc}_K$. The case of $\gamma = \bar{a}$, for 7, is similar.

Now suppose $\gamma = a$. By clause 6 of the induction hypothesis for some $K' \subseteq K''$, and S we have $\Gamma, y : U \vdash a : \mathbf{box}_{K'}$, $\Gamma, y : U \vdash x : \mathbf{chan}_{K''} S$, $tc((\Gamma, y : U), v, S)$ well-defined, and $\text{ran}(tc((\Gamma, y : U), v, S)) \subseteq \{\mathbf{name}\}$. Moreover

$$\begin{aligned} \Theta &\leq tc((\Gamma, y : U), v, S) \\ C &\subseteq K'' \end{aligned}$$

implies $\Gamma, y : U, \Theta \vdash Q : \mathbf{proc}_K$.

As $y \notin \text{fn}(\ell)$ the various strengthening results suffice to show clause 6.

Tau Consider the last pair of rules used:

(Comm)(Par) We have

$$\frac{\Gamma \vdash_{K_1} P_1 \xrightarrow{\bar{x}^\top v}_C P'_1 \quad \Gamma \vdash_{K_2} P_2 \xrightarrow{x^\top v}_C P'_2}{\Gamma \vdash_{K_1 \cap K_2} P_1 \mid P_2 \xrightarrow{\tau}_\emptyset (\nu \text{fn}(x, v) - \text{dom}(\Gamma))(P'_1 \mid P'_2)} \text{ (Comm)}$$

Consider cases of γ and the corresponding output and input clauses:

Case \star, \uparrow 1,5. By the induction hypotheses there exists $\Theta \leq tc(\Gamma, v, T)$ such that $\Gamma, \Theta \vdash P'_1 : \mathbf{proc}_{K_1}$ and $\Gamma, \Theta \vdash P'_2 : \mathbf{proc}_{K_2}$. By the (Par) and (Res) typing rules $\Gamma \vdash (\nu \text{fn}(x, v) - \text{dom}(\Gamma))(P'_1 \mid P'_2) : \mathbf{proc}_{K_1 \cap K_2}$.

Case a 4,6 By clause 4 of the induction hypothesis there exists $\Theta \leq tc(\Gamma, v, \top)$ such that $\Gamma, \Theta \vdash P'_1 : \mathbf{proc}_{K_1}$. By Lemma 11 $tc(\Gamma, v, \top)$ is the type context mapping each $x \in \text{fn}(v) - \text{dom}(\Gamma)$ to **name**.

By clause 4 $C \subseteq K'$ and by clause 6 $K' \subseteq K''$ so $C \subseteq K''$.

By clause 6 of the induction hypothesis $tc(\Gamma, v, S)$ is well-defined and has range contained in $\{\mathbf{name}\}$, so $tc(\Gamma, v, S) = tc(\Gamma, v, \top)$, so $\Theta \leq tc(\Gamma, v, S)$, so $\Gamma, \Theta \vdash P'_2 : \mathbf{proc}_{K_2}$.

By the (Par) and (Res) typing rules $\Gamma \vdash (\nu \text{fn}(x, v) - \text{dom}(\Gamma))(P'_1 \mid P'_2) : \mathbf{proc}_{K_1 \cap K_2}$.

Case \bar{a} 3,7 Similar to case \star, \uparrow above.

(Par)(Par) By the induction hypothesis.

(Box-1)(Box) We have

$$\frac{\text{dom}(\Gamma) \vdash P \xrightarrow{\bar{x}^\top v}_C Q}{\text{dom}(\Gamma) \vdash n[P] \xrightarrow{\tau}_\emptyset (\nu \text{fn}(x, v) - \text{dom}(\Gamma))(C : \bar{x}^\top v \mid n[Q])} \text{ (Box-1)}$$

and

$$\frac{\begin{array}{l} \Gamma \vdash n : \mathbf{box}_K \\ \text{pn}(P) \subseteq K \\ \text{fn}(P) \subseteq \text{dom}(\Gamma) \end{array}}{\Gamma \vdash n[P] : \mathbf{proc}_K} \text{ (Box)}$$

Note that we do not have $\Gamma \vdash P : \mathbf{proc}_K$, so the induction hypothesis is not applicable.

Take $\Theta = tc(\Gamma, \langle xv \rangle, \top)$.

By weakening $\Gamma, \Theta \vdash n : \mathbf{box}_K$. By Lemma 24 $\text{pn}(Q) \subseteq K$. In addition we have $\text{fn}(Q) \subseteq \text{dom}(\Gamma, \Theta)$, so $\Gamma, \Theta \vdash n[Q] : \mathbf{proc}_K$.

We have also $\Gamma, \Theta \vdash x : \mathbf{name}$, $\Gamma, \Theta \vdash v : \top$ and (again by Lemma 24) $C \subseteq K$, so $\Gamma, \Theta \vdash C : \bar{x}^\top v : \mathbf{proc}_K$.

By the (Par) and (Res) typing rules $\Gamma \vdash (\nu \text{fn}(x, v) - \text{dom}(\Gamma))(C : \bar{x}^\top v \mid n[Q]) : \mathbf{proc}_K$.

(Box-3)(Box) As a τ transitions cannot increase the principal set or free name set of a process.

(Res-1)(Res) By the induction hypothesis.

(Struct)(*) Follows from the induction hypothesis and Lemma 21.

(*)(Spec) Follows from the induction hypothesis and a use of (Spec) for Q .

□

Proof (of Theorem 6) By $\Gamma \vdash P : \mathbf{proc}_K$ and for some T and K , we get $x : \mathbf{chan}_K T \in \Gamma$. Furthermore, we have $\Gamma \vdash v : T$ and since $\gamma \in \{\star, \uparrow\}$ we also have $\Gamma \vdash p : T$. By Lemma 12 and the facts that Γ atomic, $\Gamma \vdash v : T$ and $\vdash p : T \triangleright \Delta$, we have $\Gamma; \Delta \vdash \{v/p\}$ good. By Lemma 15 and the facts that P does not contain a box, $\Gamma, \Delta \vdash P' : \mathbf{proc}_K$ and $\Gamma; \Delta \vdash \{v/p\}$ good, we have $\{v/p\}P$ is well-defined. □

B.2 Proving Causal Flow for \mathcal{F}

The proof that \mathcal{F} has the causal flow property is a straightforward induction on the traces of $\mathcal{F}(P, Q)$ using the Subject Reduction theorem.

Proof (of Theorem 8) Consider an instantiation $\mathcal{F}(P, Q)$ and coloured trace

$$A \vdash \emptyset \circ \mathcal{F}(P, Q) \xrightarrow{\ell_1}_{C_1} R_1 \dots \xrightarrow{\ell_k}_{C_k} R_k,$$

such that all inputs on in_1 in $\ell_1.. \ell_k$ are coloured with p and all inputs on in_2 are coloured with q .

By the definition of transitions (if $k \geq 1$) we have $\text{fn}(\mathcal{F}(P, Q)) \subseteq A$.

Let Γ_0 be the type environment Γ above with Γ_1 removed.

Let Γ_1 be the type environment mapping $\text{fn}(P, Q) - \text{dom}(\Gamma_0)$ to **name** and $\Gamma = \Gamma_0, \Gamma_1$.

Clearly Γ atomic.

By the definition of instantiation we have $\text{fn}(P, Q) \subseteq \text{dom}(\Gamma) - \{a, b\}$.

By Theorem 7 $\Gamma \vdash \emptyset \circ \mathcal{F}(P, Q) : \mathbf{proc}_p$.

By \mathcal{F} pure we know the ℓ_j have the form τ , $in_i \uparrow v$, or $\overline{out}_i \uparrow v$, for $i \in \{1, 2\}$.

Take $R_0 = \emptyset \circ \mathcal{F}(P, Q)$ and $\Theta_0 = \emptyset$. We now show by induction on k that for all $j \in 1..k$ $\ell_j = \overline{out}_1 \uparrow v \implies q \notin C_j$ and there exists Θ_j atomic such that $\Gamma, \Theta_j \vdash R_j : \mathbf{proc}_p$.

Consider the transition $R_{k-1} \xrightarrow{\ell_k}_{C_k} R_k$. We have Γ, Θ_{k-1} atomic, $\Gamma, \Theta_{k-1} \vdash R_{k-1} : \mathbf{proc}_p$, and $\text{dom}(\Gamma, \Theta_{k-1}) \vdash R_{k-1} \xrightarrow{\ell_k}_{C_k} R_k$, so

$$\Gamma, \Theta_{k-1} \vdash_p R_{k-1} \xrightarrow{\ell_k}_{C_k} R_k$$

Consider cases of ℓ_j .

Case $\overline{out}_1 \uparrow v$. By Theorem 5 for some K', T we have $C_k \subseteq K'$ and there exists $\Theta \leq tc(\Gamma, \Theta_{k-1}, \langle out_1 v \rangle, \langle \mathbf{chan}_{K'} T T \rangle)$ such that $\Gamma, \Theta_{k-1}, \Theta \vdash R_k : \mathbf{proc}_p$.

As $tc(\dots)$ is defined and $out_1 : \mathbf{chan}_p \top \in \Gamma$ we have $K' = \{p\}$ and $T = \top$, so $C_k \subseteq \{p\}$, so $q \notin C_k$.

Take $\Theta_k = \Theta_{k-1}, \Theta$; it is clearly atomic.

Case $in_1 \uparrow v$. By Theorem 5 for some K', T we have $\Gamma, \Theta_{k-1} \vdash in_1 : \mathbf{chan}_{K'} T$. If moreover $C_k \subseteq K'$ and $\Theta \leq tc(\Gamma, \Theta_{k-1}, v, T)$ then $\Gamma, \Theta_{k-1}, \Theta \vdash R_k : \mathbf{proc}_p$.

As $in_1 : \mathbf{chan}_p \top \in \Gamma$ we have $K' = \{p\}$ and $T = \top$. By the premises $C_k \subseteq \{p\}$.

As $T = \top$ we have $tc(\Gamma, \Theta_{k-1}, v, T)$ defined and atomic; take Θ equal to this and $\Theta_k = \Theta_{k-1}, \Theta$.

The other cases are similar. □

References

- [1] Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon G. Riecke. A core calculus of dependency. In ACM, editor, *POPL '99. Proceedings of the 26th ACM SIGPLAN-SIGACT on Principles of programming languages, January 20–22, 1999, San Antonio, TX*, pages 147–160, New York, NY, USA, 1999. ACM Press.
- [2] Martín Abadi, Cédric Fournet, and Georges Gonthier. Secure implementation of channel abstractions. In *LICS 98 (Indiana)*, pages 105–116. IEEE, Computer Society Press, July 1998.
- [3] Michele Boreale and Davide Sangiorgi. A fully abstract semantics for causality in the pi-calculus. In E. W. Mayr and C. Puech, editors, *Proceedings of STACS'95*, volume 900 of *Lecture Notes in Computer Science*, pages 243–254. Springer-Verlag, 1995.
- [4] Gérard Boudol. Asynchrony and the π -calculus (note). Rapport de Recherche 1702, INRIA Sofia-Antipolis, May 1992.
- [5] S.D. Brookes, C.A.R. Hoare, and A.W. Roscoe. A theory of communicating sequential processes. *Journal of the ACM*, 31(3):560–599, 1984.
- [6] Luca Cardelli and Andrew D. Gordon. Mobile ambients. In *Proc. of Foundations of Software Science and Computation Structures (FoSSaCS), ETAPS'98, LNCS 1378*, pages 140–155, March 1998.
- [7] Pierpaolo Degano and Corrado Priami. Causality for mobile processes. In Zoltán Fülöp and Ferenc Gécseg, editors, *Proceedings of ICALP '95*, volume 944 of *Lecture Notes in Computer Science*, pages 660–671. Springer-Verlag, 1995.
- [8] Cédric Fournet, Georges Gonthier, Jean-Jacques Lévy, Luc Maranget, and Didier Rémy. A calculus of mobile agents. In *Proceedings of CONCUR '96. LNCS 1119*, pages 406–421. Springer-Verlag, August 1996.
- [9] Timothy Fraser, Lee Badger, and Mark Feldman. Hardening COTS software with generic software wrappers. In *IEEE Symposium on Security and Privacy*, Berkeley, California, May 1999.
- [10] Douglas P. Ghormley, Steven H. Rodrigues, David Petrou, and Thomas E. Anderson. Interposition as an operating system extension mechanism. Technical Report CSD-96-920, University of California, Berkeley, April 9, 1997.
- [11] Neal Glew and Greg Morrisett. Type-safe linking and modular assembly language. In ACM, editor, *POPL '99. Proceedings of the 26th ACM SIGPLAN-SIGACT on Principles of programming languages, January 20–22, 1999, San Antonio, TX*, pages 250–261, New York, NY, USA, 1999. ACM Press.
- [12] Ian Goldberg, David Wagner, Randi Thomas, and Eric A. Brewer. A secure environment for untrusted helper applications. In *Sixth USENIX Security Symposium*, San Jose, California, July 1996.
- [13] Nevin Heintze and Jon G. Riecke. The SLam calculus: Programming with secrecy and integrity. In *Proceedings of the 25th POPL*, January 1998.
- [14] Kohei Honda and Mario Tokoro. An object calculus for asynchronous communication. In Pierre America, editor, *Proceedings of ECOOP '91, LNCS 512*, pages 133–147, July 1991.
- [15] Michael B. Jones. Interposition agents: Transparently interposing user code at the system interface. In Jan Vitek and Christian Jensen, editors, *Secure Internet Programming: Security Issues for Mobile and Distributed Objects*. Springer Verlag, 1999.
- [16] Myong H. Kang, Ira S. Moskowitz, and Daniel C. Lee. A network pump. *IEEE Transactions on Software Engineering*, 22(5):329–338, May 1996.
- [17] Xavier Leroy and François Rouaix. Security properties of typed applets. In *Conference Record of POPL '98: The 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 391–403, San Diego, California, 19–21 January 1998.

- [18] J. McLean. Security models. In J. Marciniak, editor, *Encyclopedia of Software Engineering*, Wiley & Sons, 1994.
- [19] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, Parts I + II. *Information and Computation*, 100(1):1–77, 1992.
- [20] G. C. Necula and P. Lee. Safe, untrusted agents using proof-carrying code. In G. Vigna, editor, *Mobile Agents and Security*, volume 1419 of *LNCS*, pages 61–91. SV, 1998.
- [21] Jens Palsberg and Peter Ørbæk. Trust in the lambda-calculus. *Journal of Functional Programming*, 7(6):557–591, November 1997.
- [22] Benjamin C. Pierce and David N. Turner. Pict: A programming language based on the pi-calculus. In Gordon Plotkin, Colin Stirling, and Mads Tofte, editors, *Proof, Language and Interaction: Essays in Honour of Robin Milner*. MIT Press, 1999.
- [23] James Riely and Matthew Hennessy. A typed language for distributed mobile processes. In *Proceedings of the 25th POPL*, January 1998.
- [24] A.W. Roscoe and M.H. Goldsmith. What is intransitive noninterference? In *Proceedings of the 12th IEEE Computer Security Foundations Workshop (CSFW-12)*, Mordano, Italy, June 1999.
- [25] Andrei Sabelfeld and David Sands. A PER model of secure information flow in sequential programs. In *Proceedings of European Symposium on Programming*, Amsterdam, Netherlands, March 1999.
- [26] Peter Sewell. Global/local subtyping and capability inference for a distributed π -calculus. In *Proceedings of ICALP '98, LNCS 1443*, pages 695–706, 1998.
- [27] Peter Sewell. A brief introduction to applied π , January 1999. Lecture notes for the Mathfit Instructional Meeting on Recent Advances in Semantics and Types for Concurrency: Theory and Practice, July 1998. Available from <http://www.cl.cam.ac.uk/users/pes20/>.
- [28] Peter Sewell and Jan Vitek. Secure composition of insecure components. In *Proceedings of the 12th IEEE Computer Security Foundations Workshop (CSFW-12)*, Mordano, Italy, June 1999.
- [29] Peter Sewell and Jan Vitek. Secure composition of insecure components. Trusted objects, Centre Universitaire d'Informatique, University of Geneva, July 1999. Also available as University of Cambridge TR 463.
- [30] Peter Sewell, Pawel T. Wojciechowski, and Benjamin C. Pierce. Location-independent communication for mobile agents: a two-level architecture. In *Internet Programming Languages, LNCS 1686*. Springer-Verlag, October 1999.
- [31] Geoffrey Smith and Dennis Volpano. Secure information flow in a multi-threaded imperative language. In *Conference Record of POPL '98: The 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 355–364, San Diego, California, 19–21 January 1998.
- [32] Jan Vitek and Giuseppe Castagna. Towards a calculus of mobile computations. In *Workshop on Internet Programming Languages, Chicago*, May 1998.
- [33] Jose-Luis Vivas and Mads Dam. From higher-order pi-calculus to pi-calculus in the presence of static operators. In Davide Sangiorgi and Robert de Simone, editors, *CONCUR '98: Concurrency Theory (9th International Conference, Nice, France)*, volume 1466 of *lncs*, pages 115–130. sv, September 1998.
- [34] D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):1–21, 1996.
- [35] Dennis Volpano and Geoffrey Smith. Confinement properties for programming languages. *SIGACT News*, 29(3):33–42, September 1998.
- [36] Dan S. Wallach, Dirk Balfanz, Drew Dean, and Edward W. Felten. Extensible security architectures for Java. In *Proceedings of the 16th Symposium on Operating System Principles*, 1997.

- [37] G. Winskel and M. Nielsen. Models for concurrency. In Abramsky, Gabbay, and Maibaum, editors, *Handbook of Logic in Computer Science*, volume IV, pages 1–148. Oxford University Press, 1995.
- [38] Steve Zdancewic, Dan Grossman, and Greg Morrisett. Principals in programming languages: A syntactic proof technique. In *International Conference on Functional Programming*, Paris, France, September 1999.