

Relational Analysis of Correlation ^{*}

Jörg Bauer¹, Flemming Nielson², Hanne Riis Nielson², and Henrik Pilegaard²

¹ Institut für Informatik, Technische Universität München, Germany.

² DTU Informatics, Technical University of Denmark
Kongens Lyngby, Denmark.

`joba@model.in.tum.de` {`nielson,riis,hepi`}@imm.dtu.dk

Abstract. In service-oriented computing, correlations are used to determine links between service providers and users. A correlation contains values for some variables received in a communication. Subsequent messages will only be received when they match the values of the correlation. Correlations allow for the implementation of sessions, local shared memory, gradually provided input, or input provided in arbitrary order – thus presenting a challenge to static analysis.

In this work, we present a static analysis in relational form of correlations. It is defined in terms of a fragment of the process calculus COWS that itself builds on the Fusion Calculus. The analysis is implemented and practical experiments allow us to automatically establish properties of the flow of information between services.

1 Introduction

Process calculi have proved their usefulness in describing and analysing distributed systems. It is therefore natural that they are also applied to model and analyse services made available over loosely coupled networks. *Correlation* has been identified as a useful feature to model services and hence appears in emerging service-oriented process calculi [6, 5].

Correlation is an idea that stems from executable business process languages such as BPEL. In the presence of multiple concurrent instances of the same service, messages sent from a user to the provider must be delivered to the correct instance of the service. This is achieved by associating specific, already available data in the messages to maintain a unique reference to a specific service instance. For example, such data may be derived from personal information like a social security number.

Technically, in process calculi, correlation may be realised by the *decoupling of name binding and input actions* in combination with *pattern matching* in input prefixes. Decoupling was first formulated in the Fusion Calculus [14], where inputs are *not* binders. Rather, a *scope construct* $(x)P$ is the only binder binding x in P . The effect of a communication inside P may then be to *fuse* x with another name, e.g., y . The names x and y will then be considered identical. The scope of the subsequent substitution of y for x will then be all of P . This allows

^{*} This work has been partially sponsored by the project SENSORIA, IST-2005-016004.

to model local, shared state, which is essential for correlations. As argued in [3], this decoupling is impossible to encode naturally in the π -calculus.

Based on these considerations we have delineated the *Calculus for Web Services*, CWS. It is mainly a fragment of the Calculus of Orchestration of Web Services (COWS, [6]) developed in the EU project SENSORIA. In contrast to the Fusion Calculus (or D-Fusion developed in [3]), COWS is clearly service-centred and focuses on, among others, correlation. In CWS, however, we discard some of the technicalities associated with session management.

The goal of this work is not the design of a new language but the development of a relational analysis for correlation. The insights gained here should be easily transferable to any other correlation-based language for services.

Contribution. The study of process calculi is a challenging avenue for the development and application of static analysis. Such analyses often tend to fall between two extremes: rather simple 0-CFA analyses (e.g., [2]) expressing only rudimentary properties, or extremely powerful relational and polyhedral analyses (e.g., [16, 4]). The latter technology seems to be mastered by only a few, and hence might not obtain widespread use, and furthermore seems to require that the calculus in question be presented in a *non-trivial normalised form* in order to aid the analysis thereby making it a major effort to apply the techniques. To be specific, the work of [16] on developing relational and polyhedral analyses for the π -calculus depends heavily on the PhD thesis of [15] that develops the non-trivial normalised form used. We therefore consider it an important achievement of our work that we populate the middle ground between the two extremes – this approach was used successfully in [8] dealing with the π -calculus. Here we apply this approach to develop a relational static analysis for correlation, which is a challenging and highly relevant aspect of process calculi for service-oriented computing. To the best of our knowledge [1], we are the first to develop a static analysis for a descendant of the Fusion Calculus.

Outline. In Section 2, we shall introduce the CWS calculus and illustrate it by giving an example of an accident service. In Section 3, we develop a static analysis in relational form for CWS. Before we report on our implementation and experimental results, we establish the correctness of our analysis expressed by two major theorems, subject reduction and adequacy. It is noteworthy that our subject reduction result does not claim that analysability is preserved under reduction, but rather that a stronger notion of “analysability of all permitted instances” is preserved under reduction. This provides a new insight in the development of static analyses of process calculi. Section 5 reports on related work and Section 6 concludes.

2 Calculus for Web Services

Table 1 defines the syntactic domains and the syntax of CWS, the Calculus for Web Services. We have three syntactic domains, *names*, *variables*, and *labels*.

Services		Name	Domain	Symbols
$s ::=$	$u \bullet u'!\bar{v}^\ell$ (invoke)			
	g (input-guarded choice)			
	$s \mid s$ (parallel composition)	Variables	Var	x, y
	$(n)s$ (name binding box)	Names	Name	n, m
	$[x]^\ell s$ (variable scope)	Partners	Name	p, p'
	$*s$ (replication)	Operations	Name	o, o'
$g ::=$	(input-guarded choice)		Name \cup Var	u, v
	$\mathbf{0}$ (nil)	Labels	Lab	ℓ
	$p \bullet o?v^\ell.s$ (request processing)			
	$g + g$ (choice)			

Table 1. Syntax of CWS and its syntactic domains. Service invocation, request processing, as well as the variable scope are decorated with labels, $\ell \in \mathbf{Lab}$. An overline denotes tuples, for instance, \bar{x} denotes tuples of variables.

While the latter merely serve as pointers into the syntax guiding the analysis specification of Section 3, names denote computational values and variables are used to bind names. In the sequel we assume that services are consistently labelled: two equally labelled actions are either both inputs – they coincide on partner, operation, and on the input length – or both outputs – they coincide on the output length.

The computational entities of CWS are called *services*. In contrast to, e.g., the π -calculus, communication endpoints in CWS are pairs $p \bullet o$ of a *partner* and an *operation* modeling several services available at the same site. Communication endpoints at reception sites are *statically determined*, because a service is supposed to know itself. In contrast, received partner and operation names may be used for subsequent service invocation. Additionally, CWS has output – rendered *asynchronous* by the missing continuation after the service invocation – input actions, input-guarded choice, parallel composition, name binding, variable scope declaration, and replication; each with the expected meaning. Names in input prefixes are used for pattern matching.

Variable binding and global scope are defined similarly to the Fusion Calculus and differently from π -like calculi: Variables are *not* bound at input actions, the only variable binder is the *scope construct* $[x]^\ell s$. If variable x occurs in an input action a , at which name n may be received, then the scope of the induced substitution, $[x \mapsto n]$, is the whole of s ; not just the continuation of a .

2.1 An Accident Service

Our running example is given in Table 2. It describes an arbitrary number of cars (1-3) that have subscribed to an *accident service* (4-10). Each car is equipped with a GPS device tracking its position. In case an on-board sensor detects any abnormal behaviour, the alarm operation of the service centre is invoked (1) by

	$*(self) *(gps) *(sid)$	$p_s \bullet o_{alarm}!(acc, self, sid)^1$	(1)		
		$ (self \bullet o_{confirm}?(sid)^2.p_s \bullet o_{confirm}!(ok, self, sid)^3 +$	(2)		
		$self \bullet o_{confirm}?(sid)^4.p_s \bullet o_{confirm}!(ko, gps, sid)^5)$	(3)		
$*[x_{info}]^6$	$[x_{id}]^7$	$[x_{reply}]^8$	$[x_{sid}]^9$	$p_s \bullet o_{alarm}?(acc, x_{id}, x_{sid})^{10}.$	(4)
				$x_{id} \bullet o_{confirm}!(x_{sid})^{11}$	(5)
				$ p_s \bullet o_{confirm}?(x_{reply}, x_{info}, x_{sid})^{12}.$	(6)
				$p_s \bullet o_{Sos}!(x_{reply}, x_{info}, x_{sid})^{13}$	(7)
				$ p_s \bullet o_{Sos}?(ko, x_{info}, x_{sid})^{14}.$	(8)
				$p_{amb} \bullet o_{Sos}!(x_{info}, x_{sid})^{15}$	(9)
				$ p_s \bullet o_{Sos}?(ok, x_{info}, x_{sid})^{16}.\mathbf{0}$	(10)

Table 2. An accident service.

sending an accident message with the identity $self$ of the driver and the nonce sid , which prevents malign session interference. The centre processes this request (4) and asks for confirmation (5), whereupon the alarm is either revoked (ok ,2) or confirmed (ko ,3) by the driver. The driver identity is attached to the revocation message, while the current GPS position is attached to the confirmation message. Then the centre processes the answer and invokes another internal service (6,7). Depending on the driver's answer either an ambulance is called and informed about the location of the accident (8,9), or a false alarm is detected (10).

Note that there may be many cars around having a number of false alarms involving different locations. Due to the safety-critical nature of this example, the service centre must be able to handle many of these service calls concurrently without mixing up information from different sessions. In particular, we would like to guarantee (and shall indeed show) that:

1. The ambulance is not called, when an ok was received.
2. Messages sent to the ambulance always contain GPS information. Note, that the validity of this property is not obvious, since x_{info} may be bound to both driver identities and positions at run-time.
3. If several cars employ the accident service at once, then only the positions of the ones confirming the accident are reported to the ambulance.

2.2 Labelled Transition System for CWS

To define the semantics of CWS we employ a notion of structural congruence defined as the least congruence that incorporates the axioms of Table 3, contains disciplined α -renaming of names³, and asserts that choice and parallel are associative, commutative, and have $\mathbf{0}$ as neutral element. We do *not* allow the extrusion of *variable scopes*. Rather, [BINDER₄] allows the binding boxes of *names* to migrate freely in and out of variable scopes. *Substitutions*, σ , are mappings

³ For disciplined α -renaming, we identify each name with its defining syntactic occurrence, its canonical name, thus partitioning the name space into finitely many equivalence classes rendering canonical names stable under evaluation.

$*\mathbf{0} \equiv \mathbf{0}$	[REPL ₁]
$*s \equiv s \mid *s$	[REPL ₂]
$(n)\mathbf{0} \equiv \mathbf{0}$	[BINDER ₁]
$(n)(m)s \equiv (m)(n)s$	[BINDER ₂]
$s_1 \mid (n)s_2 \equiv (n)(s_1 \mid s_2)$ if $n \notin fn(s_1)$	[BINDER ₃]
$[x](n)s \equiv (n)[x]s$	[BINDER ₄]

Table 3. An excerpt of the CWS structural congruence rules. For a given service expression s , $fn(s)$ denotes the set of free names.

with domain $\mathbf{Var} \rightarrow \mathbf{Name}$. The application of a substitution, $[x \mapsto n]$, to a service s is written $s \cdot [x \mapsto n]$ and replaces free occurrences of x in s by n . The disjoint union of two substitutions σ_1 and σ_2 is written $\sigma_1 \uplus \sigma_2$ and the substitution with empty domain is written \emptyset . A tuple \bar{u} of names and variables is matched against a tuple \bar{n} of names as follows (yielding a substitution):

$$\mathcal{M}(x, n) = x \mapsto n \quad \mathcal{M}(n, n) = \emptyset \quad \frac{\mathcal{M}(u_1, n_1) = \sigma_1 \quad \mathcal{M}(\bar{u}_2, \bar{n}_2) = \sigma_2}{\mathcal{M}((u_1, \bar{u}_2), (n_1, \bar{n}_2)) = \sigma_1 \uplus \sigma_2}$$

The semantics of CWS (Table 4) is given as a labelled transition system using transition labels α :

$$\alpha ::= (p \bullet o) \triangleleft \bar{n}^\ell \mid (p \bullet o) \triangleright \bar{u}^\ell \mid p \bullet o[\sigma] \bar{u}^{\ell_i} \bar{n}^{\ell_o}$$

Transition labels $(p \bullet o) \triangleleft \bar{n}^\ell$ and $(p \bullet o) \triangleright \bar{u}^\ell$ result from applying rules for invocation and reception of names, respectively, that is, rules [INV] and [REC]. They can engage in a communication ([COM]) giving rise to a transition label $p \bullet o[\sigma] \bar{u}^{\ell_i} \bar{n}^{\ell_o}$, where σ is the resulting substitution. Note that the only transition labels being *observable* are those of the form $p \bullet o[\emptyset] \bar{u}^{\ell_i} \bar{n}^{\ell_o}$. We use $d(\alpha)$ to denote the set of names and variables occurring in α ; if $\alpha = p \bullet o[\sigma] \bar{u}^{\ell_i} \bar{n}^{\ell_o}$ then $d(\alpha)$ contains only the names and variables in the domain and codomain of σ .

At top level, the semantics is only defined for *closed services*, that is, no free variables may occur when a computation step is applied. Deeper in the inference tree there may of course be free occurrences of variables. A substitution $[x \mapsto n]$ is only applied, when the enclosing scope of x is met ([DEL_{sub}]). This ensures that the effect of a communication is visible *globally* within a scope and allows to model shared memory of a session instance. Services can proceed normally under name bindings or variable scopes, unless the enclosing entity is mentioned by the transition label ([NAME] and [SCOPE]). Note that for an output to occur all variables in it must have been substituted by names earlier ([INV]).

Comparison with COWS. CWS lacks the orchestration constructs (kill and protect) found in COWS, because we focus on correlation rather than on orchestration. The semantics of CWS and COWS are comparable up to rule [COM]: Let two different communications both yield a valid match. While our rule picks non-deterministically from the set of choices, the COWS semantics [6] imposes a

[INV] $p \bullet o! \bar{n}^\ell \xrightarrow{(p \bullet o) < \bar{n}^\ell} \mathbf{0}$	[REC] $p \bullet o? \bar{u}^\ell . s \xrightarrow{(p \bullet o) > \bar{u}^\ell} s$
[CHOICE] $\frac{g_1 \xrightarrow{\alpha} s}{g_1 + g_2 \xrightarrow{\alpha} s}$	[DEL _{sub}] $\frac{s \xrightarrow{p \bullet o[\sigma \uplus [x \mapsto n]] \bar{u}^{\ell_i} \bar{n}^{\ell_o}} s'}{[x]^\ell s \xrightarrow{p \bullet o[\sigma] \bar{u}^{\ell_i} \bar{n}^{\ell_o}} s' \cdot [x \mapsto n]}$
[NAME] $\frac{s \xrightarrow{\alpha} s' \quad n \notin d(\alpha)}{(n)s \xrightarrow{\alpha} (n)s'}$	[SCOPE] $\frac{s \xrightarrow{\alpha} s' \quad x \notin d(\alpha)}{[x]s \xrightarrow{\alpha} [x]s'}$
[COM] $\frac{s_1 \xrightarrow{(p \bullet o) > \bar{u}^{\ell_i}} s'_1 \quad s_2 \xrightarrow{(p \bullet o) < \bar{n}^{\ell_o}} s'_2 \quad \mathcal{M}(\bar{u}, \bar{n}) = \sigma}{s_1 \mid s_2 \xrightarrow{p \bullet o[\sigma] \bar{u}^{\ell_i} \bar{n}^{\ell_o}} s'_1 \mid s'_2}$	
[PAR] $\frac{s_1 \xrightarrow{\alpha} s'_1}{s_1 \mid s_2 \xrightarrow{\alpha} s'_1 \mid s_2}$	[CONG] $\frac{s \equiv s_1 \quad s_1 \xrightarrow{\alpha} s_2 \quad s_2 \equiv s'}{s \xrightarrow{\alpha} s'}$

Table 4. CWS operational semantics.

constraint, *noc*, selecting the match that gives rise to the smallest substitution. In case of equally long substitutions, the choice is random. In [6], this feature is used to distinguish (less instantiated) service definitions from (more instantiated) service instances. However, we omit this feature from our semantics, because it is not analysed, and our analysis would still produce a valid over-approximation if we used it. In order to prevent malign session interferences we adapt what is considered good style in communication protocols, where one relies on the use of nonces or shared secrets (or indeed shared keys) to ensure that sessions do not interfere. Finally, note that it is indeed possible to implement the Fusion-like scoping by using explicit environments in the semantics – however, as anywhere else, we decided to follow COWS as closely as possible.

3 A Relational Analysis for CWS

In this section, we develop a relational static analysis [9] for CWS: For each program label ℓ – invocation, processing, and variable scope – we compute sets of tuples of names to which the variables that are in scope at ℓ may be bound at runtime. Tracking *sets* of tuples makes the analysis relational, while an independent analysis would track tuples of sets and loose track of which variables are bound at the same time.

Auxiliary Information. A *label environment*, L , is defined as a mapping

$$L : \mathbf{Lab} \rightarrow (\mathbf{Var} \times \mathbf{Lab})^*$$

$\mathcal{L}_{\bar{\chi}}[\mathbf{0}] = []$ $\mathcal{L}_{\bar{\chi}}[u \bullet u'!v^\ell] = [\ell \mapsto \bar{\chi}]$ $\mathcal{L}_{\bar{\chi}}[p \bullet o?u^\ell.s] = \mathcal{L}_{\bar{\chi}}[s] \uplus [\ell \mapsto \bar{\chi}]$ $\mathcal{L}_{\bar{\chi}}[(n)s] = \mathcal{L}_{\bar{\chi}}[s]$ $\mathcal{L}_{\bar{\chi}}[*s] = \mathcal{L}_{\bar{\chi}}[s]$ $\mathcal{L}_{\bar{\chi}}[[x]^\ell s] = \mathcal{L}_{\bar{\chi}x^\ell}[s] \uplus [\ell \mapsto \bar{\chi}]$ $\mathcal{L}_{\bar{\chi}}[s_1 \mid s_2] = \mathcal{L}_{\bar{\chi}}[s_1] \uplus \mathcal{L}_{\bar{\chi}}[s_2]$ $\mathcal{L}_{\bar{\chi}}[s_1 + s_2] = \mathcal{L}_{\bar{\chi}}[s_1] \uplus \mathcal{L}_{\bar{\chi}}[s_2]$	$\mathcal{F}[\mathbf{0}] = ([], \emptyset)$ $\mathcal{F}[u \bullet u'!v^\ell] = ([], \{\ell\})$ $\mathcal{F}[p \bullet o?u^\ell.s] = \text{let } (F, E) = \mathcal{F}[s] \\ \text{in } (F \uplus [\ell \mapsto E], \{\ell\})$ $\mathcal{F}[(n)s] = \mathcal{F}[s]$ $\mathcal{F}[*s] = \mathcal{F}[s]$ $\mathcal{F}[[x]^\ell s] = \text{let } (F, E) = \mathcal{F}[s] \\ \text{in } (F \uplus [\ell \mapsto E], \{\ell\})$ $\mathcal{F}[s_1 \mid s_2] = \mathcal{F}[s_1 + s_2]$ $= \text{let } (F_1, E_1) = \mathcal{F}[s_1] \\ (F_2, E_2) = \mathcal{F}[s_2] \\ \text{in } (F_1 \uplus F_2, E_1 \cup E_2)$
--	--

Table 5. Label environment $\mathcal{L}_{\bar{\chi}}[s]$ and flow information $\mathcal{F}[s]$

that to each label ℓ associates a sequence $\bar{\chi} \in (\mathbf{Var} \times \mathbf{Lab})^*$ of pairs of variables and labels that have been introduced *before* this point in the process; the label indicates the exact scope where the variable was introduced. Formally, we shall take $\mathbf{L} = \mathcal{L}_\epsilon[s]$, where $\mathcal{L}_{\bar{\chi}}$ is defined in Table 5. Here we write \uplus for joining two mappings with *disjoint* domains and $[]$ for the mapping with empty domain. The notation $\bar{\chi}_1\bar{\chi}_2$ stands for the concatenation of $\bar{\chi}_1$ and $\bar{\chi}_2$. Furthermore, we write $x\#\mathbf{L}.\ell$ for the label, at which x at position ℓ was defined, taking into account that the most recently introduced definition of x is the rightmost:

$$x\#\langle x_1^{\ell_1}, \dots, x_k^{\ell_k} \rangle = \ell_t, \text{ if } t = \max\{i \mid x_i = x\}$$

This notion is only defined, when x is among the x_i , which will always be the case, when we use it.

Flow of control in a service s is represented a *flow mapping* \mathbf{F} that to each label ℓ associates the set of labels that will become visible once the action labelled ℓ has been executed; thus

$$\mathbf{F} : \mathbf{Lab} \hookrightarrow \mathcal{P}(\mathbf{Lab})$$

Function \mathcal{F} in Table 5 computes the flow mapping, \mathbf{F} , together with the set \mathbf{E} of *visible labels* of a service s , where $(\mathbf{F}, \mathbf{E}) = \mathcal{F}[s]$. When applying \mathbf{F} and \mathbf{L} to labels, we shall write $\mathbf{L}.\ell$ and $\mathbf{F}.\ell$ instead of $\mathbf{L}(\ell)$ and $\mathbf{F}(\ell)$.

Example 1. Label environment computed from the running example in Table 2:

$$\begin{aligned} \mathbf{L}.1 = \mathbf{L}.2 = \mathbf{L}.3 = \mathbf{L}.4 = \mathbf{L}.5 = \mathbf{L}.6 = \epsilon & \quad \mathbf{L}.7 = \langle (x_{info}, 6) \rangle \\ \mathbf{L}.8 = \langle (x_{info}, 6), (x_{id}, 7) \rangle & \quad \mathbf{L}.9 = \langle (x_{info}, 6), (x_{id}, 7), (x_{reply}, 8) \rangle \\ \mathbf{L}.10 = \mathbf{L}.11 = \dots = \mathbf{L}.16 = \langle (x_{info}, 6), (x_{id}, 7), (x_{reply}, 8), (x_{sid}, 9) \rangle \end{aligned}$$

In the sequel, we shall often write pairs like $(x_{info}, 6)$ using superscript as in x_{info}^6 . Two examples of flow information are $\mathbf{F}.9 = \{10, 12, 14, 16\}$ and $\mathbf{F}.10 = \{11\}$. Finally, the set of visible labels for the running example is $\{1, 2, 4, 6\}$. \square

Analysis Domain. The *abstract environments* \hat{R} of the analysis will, given a label ℓ , return a set of sequences. Each element of such a sequence can either be a name or the *undefined* symbol \perp . The latter denotes cases, where a variable has not been assigned a value yet. This helps to track gradually provided inputs. The length of these sequences will equal that of $L.\ell$. We will thus determine the potential values of the variables at the point determined by ℓ and take:

$$\hat{R} : \mathbf{Lab} \rightarrow \mathcal{P}(\mathbf{Name}_{\perp}^*)$$

where $\mathbf{Name}_{\perp} = \mathbf{Name} \cup \{\perp\}$.⁴ We shall use w to denote elements of \mathbf{Name}_{\perp} . If $\hat{R}.\ell = \emptyset$ it means that the program point ℓ is not reachable; if $\hat{R}.\ell = \{\epsilon\}$ then also $L.\ell = \epsilon$ and it means that no variable has been introduced at that program point.

The potential values of a variable x at the label ℓ are computed by $\Pi_{x@L.\ell}(\bar{w})$, where \bar{w} has the same length as $L.\ell$ and x occurs in $L.\ell$. In cases with more than one occurrence of x in $L.\ell$ we select the rightmost, i.e., the most recently declared one. Formally:

$$\Pi_{x@L.\ell}(\bar{w}) = w_t$$

where $\bar{w} = \langle w_1, \dots, w_k \rangle$, $L.\ell = \langle x_1^{\ell_1}, \dots, x_k^{\ell_k} \rangle$, and $t = \max\{i \mid x_i = x\}$. Note, that the result may be \perp . The operation is trivially extended to names, n : $\Pi_{n@L.\ell}(\bar{w}) = n$. Also, it is extended to sequences, \bar{u} , of variables and names, as in $\Pi_{\bar{u}@L.\ell}(\bar{w})$, and to sets, R , of such sequences, i.e., $\Pi_{\bar{u}@L.\ell}(R)$. The abstract communication cache

$$\hat{K} \subseteq \mathbf{Name} \times \mathbf{Name} \times \mathbf{Name}^*$$

records the tuples of names that potentially are communicated over the channels. Elements of the domain are triples representing a partner, an operation, and the tuple of names communicated. In \hat{K} , we keep track of names only, there is no \perp involved.

Analysis Specification. The *judgements* of the analysis have the form

$$\hat{R}, \hat{K} \vdash_{L,F} s$$

where L, F, \hat{R} , and \hat{K} are as above. Intuitively, the judgements defined in Table 6 determine whether a given pair (\hat{R}, \hat{K}) is a valid analysis result. The computation of such a pair is described in Section 4.2.

The first five rules are recursive cases. The rule [RINV] deals with invocations. It looks up the available bindings of the involved variables in \hat{R} and records the resulting tuples in the abstract communication cache, \hat{K} . Since the lookup may yield undefined values, we need to intersect with \mathbf{Name}^* .

Rule [RSCOPE] extends all variable bindings available at scope definition ℓ with a single \perp and passes the information on to the program points following ℓ . This reflects that a newly introduced variable is not yet bound.

⁴ Again, all the names tracked in the analysis are canonical.

[RNIL] $\hat{R}, \hat{K} \vdash_{L,F} \mathbf{0}$	[RREP] $\frac{\hat{R}, \hat{K} \vdash_{L,F} s}{\hat{R}, \hat{K} \vdash_{L,F} *s}$	[RNAME] $\frac{\hat{R}, \hat{K} \vdash_{L,F} s}{\hat{R}, \hat{K} \vdash_{L,F} (n)s}$
[RPAR] $\frac{\hat{R}, \hat{K} \vdash_{L,F} s_1 \quad \hat{R}, \hat{K} \vdash_{L,F} s_2}{\hat{R}, \hat{K} \vdash_{L,F} s_1 \mid s_2}$	[RCHOICE] $\frac{\hat{R}, \hat{K} \vdash_{L,F} s_1 \quad \hat{R}, \hat{K} \vdash_{L,F} s_2}{\hat{R}, \hat{K} \vdash_{L,F} s_1 + s_2}$	
[RINV] $\hat{R}, \hat{K} \vdash_{L,F} u \bullet u'! \bar{v}^\ell \quad \text{if } \Pi_{uu'\bar{v}@\mathbb{L}.l}(\hat{R}.l) \cap \mathbf{Name}^* \subseteq \hat{K}$		
[RSCOPE] $\frac{\hat{R}, \hat{K} \vdash_{L,F} s}{\hat{R}, \hat{K} \vdash_{L,F} [x]^\ell s} \quad \text{if } \forall \ell' \in F.l : \hat{R}.l \times \{\perp\} \subseteq \hat{R}.\ell'$		
[RREC] $\frac{X \neq \emptyset \Rightarrow \hat{R}, \hat{K} \vdash_{L,F} s}{\hat{R}, \hat{K} \vdash_{L,F} p \bullet o? \bar{u}^\ell . s} \quad \text{if } \forall \ell' \in F.l : X \subseteq \hat{R}.\ell'$ $\forall x \in \{\bar{u}\} \forall \ell_x \in (F.(x\#\mathbb{L}.l)) : Y_x \subseteq \hat{R}.\ell_x$ where $X = \{\bar{w} \in \hat{R}.l \mid \Pi_{po\bar{u}@\mathbb{L}.l}(\bar{w}) \in \hat{K}\}$ and $Y_x = \{\bar{w}n \in \mathbf{Name}^*_\perp \mid \bar{w} \in \hat{R}.(x\#\mathbb{L}.l) \wedge$ $\exists \bar{w}' \in \mathbf{Name}^*_\perp : \Pi_{po\bar{u}@\mathbb{L}.l}(\bar{w}n\bar{w}') \in \hat{K} \}$		

Table 6. Specification of the analysis judgement $\hat{R}, \hat{K} \vdash_{L,F} s$.

Finally, the rule [RREC] describes two different flows. First, it ensures that all possible bindings at ℓ that may lead to a communicated tuple will indeed flow to the subsequent visible labels (expressed in the set X); only if this set is non-empty is it possible to perform the input, hence the test $X \neq \emptyset$ is a reachability condition for the continuation s . Second, for each variable x of the input pattern, we record in the set Y_x which names may be bound to x by any communication that matches this input pattern. This information flows to the labels just after the scope at which x was introduced (denoted by $F.(x\#\mathbb{L}.l)$).

In addition to the satisfaction of the analysis judgement, we require some initialisation information, stating that ϵ is available at all globally visible labels:

Definition 1. *Let s be a service, $(F, E) = \mathcal{F}[[s]]$ its flow information, and $L = \mathcal{L}_\epsilon[[s]]$ its label environment. A pair (\hat{R}, \hat{K}) is an acceptable analysis estimate for s , if and only if $\hat{R}, \hat{K} \vdash_{L,F} s$ and $\epsilon \in \hat{R}.l$ for all $l \in E$.*

Example 2. An acceptable analysis estimate – and in fact the least, that is, most precise one – of the running example of Table 2 comprises:

$$\hat{K} = \{ \langle p_s, o_{alarm}, acc, self, sid \rangle, \langle p_s, o_{confirm}, ok, self, sid \rangle, \\ \langle p_s, o_{confirm}, ko, gps, sid \rangle, \langle self, o_{confirm}, sid \rangle, \langle p_s, o_{Sos}, ok, self, sid \rangle, \\ \langle p_s, o_{Sos}, ko, gps, sid \rangle, \langle p_{amb}, o_{Sos}, gps, sid \rangle \}$$

Regarding the abstract environments, we state only $\hat{R}.13$ and $\hat{R}.15$ explicitly, because they are most relevant with respect to the properties we are interested in. Recall that both L.13 and L.15 amount to $\langle (x_{info}, 6), (x_{id}, 7), (x_{reply}, 8), (x_{sid}, 9) \rangle$.

$$\hat{R}.13 = \{ \langle self, self, ok, sid \rangle, \langle self, \perp, ok, sid \rangle, \langle gps, self, ko, sid \rangle, \langle gps, \perp, ko, sid \rangle \}$$

$$\hat{R}.15 = \{ \langle gps, self, ko, sid \rangle, \langle gps, self, \perp, sid \rangle, \langle gps, \perp, ko, sid \rangle, \langle gps, \perp, \perp, sid \rangle \}$$

Reconsider the three properties stated in Section 2.1. Property 1 states that the ambulance is not called, when an `ok` was received. The ambulance is called at label 15. In the analysis result, we can see that the value of x_{reply} (the third position in the tuples) cannot be `ok` proving property 1. Note that at label 13, `ok` may still occur.

Property 2 requires that an ambulance is only called with location information. It is shown by inspecting \hat{K} , which over-approximates all messages sent: All message involving p_{amb} in \hat{K} contain `gps` only.

Property 3 is not so easy to show for the analysis as is. We are able to establish it, if we unfold the definition of cars a finite number of times, though. The analysis will then show, that malign interferences are prevented by the session id `sid`. The formal result justifying our reasoning about properties in this example is stated in Theorem 2 below. \square

4 Properties of the Analysis

In this section, we establish the formal correctness of our analysis. We start by defining a *correctness predicate*, which states the analysability of *all permitted substitutions* of an analysable service, and show that its validity is preserved under observable computation steps. This constitutes our subject reduction result (Theorem 1). It is shown in Appendix A that mere analysability is not preserved under reduction. Theorem 2 states that all actually sent messages and all potential variable bindings are correctly recorded in the analysis. We conclude this section by reporting on experiments using the implementation of our analysis. The feasibility of this implementation relies on the Moore family property of the set of all acceptable analyses guaranteeing the existence of least (most precise) solutions (Theorem 3).

4.1 Correctness

In the following we assume an arbitrary but fixed given program s_* , as well as its flow information $F = \mathcal{F}[[s_*]]$, its label and its visible labels $(L, E) = \mathcal{L}_\epsilon[[s_*]]$. Moreover, we define the notion $\mathcal{E}(s)$ of the *exposed actions* of a service s , which is a set of input and output prefixes, such that $\mathcal{E}(u \bullet u'! \bar{v}^\ell) = \{u \bullet u'! \bar{v}^\ell\}$, $\mathcal{E}(p \bullet o? \bar{v}^\ell.s) = \{p \bullet o? \bar{v}^\ell.s\}$, $\mathcal{E}(*s) = \mathcal{E}((n)s) = \mathcal{E}([x]s) = \mathcal{E}(s)$, $\mathcal{E}(\mathbf{0}) = \emptyset$, $\mathcal{E}(s_1 \mid s_2) = \mathcal{E}(s_1 + s_2) = \mathcal{E}(s_1) \cup \mathcal{E}(s_2)$. Finally, we define an extended version of substitution. Let $s = p \bullet o? \bar{v}^\ell.s'$ or $s = u \bullet u'! \bar{u}^\ell$ and let $\bar{w} \in \hat{R}.\ell$ where $L.\ell = \langle x_1, \dots, x_k \rangle$ and $\bar{w} = \langle w_1, \dots, w_k \rangle$. Then we define

$$s[\bar{w}/L.\ell] = (\dots (s \cdot [x_k \mapsto w_k]) \cdot \dots) \cdot [x_1 \mapsto w_1]$$

where $s \cdot [x \mapsto \perp] = s$.

Intuitively, the correctness predicate of Definition 2 holds of a service s obtained by semantic reduction from s_* , if all exposed actions s' of s have a counterpart s'' in s_* , such that s'' is correctly analysed and such that s'' is congruent

to s' when instantiating it with information computed by the analysis. In other words, the correctness predicate describes the analysability of permitted substitution instances.

Definition 2 (Correctness Predicate). *A service s satisfies the correctness predicate with respect to s_* , written $\hat{R}, \hat{K} \models^{s_*} s$ if and only if (1-4) hold.*

1. $\hat{R}, \hat{K} \vdash_{L,F} s_*$
2. $\forall \ell \in E : \epsilon \in \hat{R}.\ell$
3. For all $p \bullet o? \bar{u}^\ell . s' \in \mathcal{E}(s)$ there exists a subexpression $p \bullet o? \bar{v}^\ell . s''$ of s_* s.t.
 - $\hat{R}, \hat{K} \vdash_{L,F} p \bullet o? \bar{v}^\ell . s''$ and
 - $\exists \bar{w} \in \hat{R}.\ell : p \bullet o? \bar{u}^\ell . s' \equiv p \bullet o? \bar{v}^\ell . s''[\bar{w}/L.\ell]$
4. For all $v \bullet v! \bar{v}^\ell \in \mathcal{E}(s)$ there exists a subexpression $u \bullet u! \bar{u}^\ell$ of s_* s.t.
 - $\hat{R}, \hat{K} \vdash_{L,F} u \bullet u! \bar{u}^\ell$ and
 - $\exists \bar{w} \in \hat{R}.\ell : v \bullet v! \bar{v}^\ell \equiv u \bullet u! \bar{u}^\ell[\bar{w}/L.\ell]$

Lemma 1 states some obvious compositionality properties of the correctness predicate. The proof is straightforward from Definition 2.

Lemma 1 (Compositionality). *Let s, s_1, s_2 be services, x a variable and n a name. It holds:*

- $\hat{R}, \hat{K} \models^{s_*} s_1 \mid s_2$ if and only if $\hat{R}, \hat{K} \models^{s_*} s_1$ and $\hat{R}, \hat{K} \models^{s_*} s_2$.
- $\hat{R}, \hat{K} \models^{s_*} s_1 + s_2$ if and only if $\hat{R}, \hat{K} \models^{s_*} s_1$ and $\hat{R}, \hat{K} \models^{s_*} s_2$.
- $\hat{R}, \hat{K} \models^{s_*} (n)s$ if and only if $\hat{R}, \hat{K} \models^{s_*} s$.
- $\hat{R}, \hat{K} \models^{s_*} [x]^\ell s$ if and only if $\hat{R}, \hat{K} \models^{s_*} s$.

Moreover, if $s_1 \equiv s_2$ then $\hat{R}, \hat{K} \models^{s_*} s_1$ if and only if $\hat{R}, \hat{K} \models^{s_*} s_2$.

We are now able to state our subject reduction result. Note that the correctness predicate is only preserved at top-level, that is, when talking about observable computation steps. A computation step is observable, when the substitution of the transition label is empty, that is, when all substitutions induced by a communication have happened.

Theorem 1 (Subject Reduction). *Let s_1 and s_2 be services. If $\hat{R}, \hat{K} \models^{s_*} s_1$ and $s_1 \xrightarrow{p \bullet o[\emptyset] \bar{u}^{\ell_i} \bar{n}^{\ell_o}} s_2$ then $\hat{R}, \hat{K} \models^{s_*} s_2$.*

The proof uses Lemma 1 and is given in Appendix B. It requires an even stronger induction hypothesis than provided by the correctness predicate, because we have to deal with all possible transition labels, in particular with transition labels carrying non-empty substitutions. The stronger induction hypothesis then makes a connection between these substitutions and the analysis result.

The following theorem constitutes the adequacy of our analysis and is proven in Appendix B. It states that every communication triggering an observed substitution and the substitution itself are in fact recorded in the analysis information.

Theorem 2 (Adequacy). *If (\hat{R}, \hat{K}) is an acceptable analysis of s_* and if $s_* \xrightarrow{*} s \xrightarrow{p \bullet o[\emptyset] \bar{u}^{\ell_o} \bar{n}^{\ell_i}} s'$ then $\langle p\bar{o}\bar{n} \rangle \in \hat{K}$ and $\bar{n} \in \Pi_{\bar{u} \otimes L.\ell_i}(\hat{R}.\ell_i)$.*

4.2 Implementation

The basis of our implementation relies on the following Moore family result. It ensures the existence and the uniqueness of a least acceptable analysis result. Its proof is obvious from the syntax directed analysis specification.

Theorem 3 (Moore Family). *For any service s the set of acceptable analysis estimates under $\vdash_{L,F}$ constitutes a Moore family, i.e.,*

$$\forall A \subseteq \{\hat{R}, \hat{K} \mid \hat{R}, \hat{K} \vdash_{L,F} s\} : \sqcap A \in \{\hat{R}, \hat{K} \mid \hat{R}, \hat{K} \vdash_{L,F} s\}.$$

If we, by a change of perspective, view the analysis specification as logical formulas and acceptable results as models of these formulas, then the Moore family result turns into a model intersection property ensuring a least model corresponding to the least acceptable analysis result.

We have implemented a fully functional prototype in Standard ML generating clauses that lie within the Alternation-free Least Fixed Point (ALFP) fragment of first order logic. Least models of such formulas always exist and can be computed efficiently by, e.g., the Succinct Solver [12, 10].

Example 3. The input communication of line (6) of Table 2 gives rise to the following clause:

$$\begin{aligned} & \forall x_{info}, x_{id}, x_{reply}, x_{sid} : \\ & \hat{R}.12(x_{info}, x_{id}, x_{reply}, x_{sid}) \wedge \hat{K}(p_s, o_{confirm}, x_{reply}, x_{info}, x_{sid}) \Rightarrow \\ & \quad [\hat{R}.13(x_{info}, x_{id}, x_{reply}, x_{sid}) \wedge \phi \wedge \psi] \end{aligned}$$

This specifies the flow into $\hat{R}.13$: All variable bindings available at $\hat{R}.12$ leading to a tuple that may be communicated flow to the continuation at line (7). This corresponds to the X set in rule [RREC]. Formula ϕ in the clause above corresponds to the Y_x set and is left out for brevity. Formula ψ in the clause above corresponds to line (7) of the example representing the analysis of an output according to [RINV] and specifying a flow into \hat{K}

$$\begin{aligned} \psi = & \forall x_{info}, x_{id}, x_{reply}, x_{sid} : \\ & \hat{R}.13(x_{info}, x_{id}, x_{reply}, x_{sid}) \wedge x_{sid} \neq \perp \wedge x_{info} \neq \perp \wedge x_{reply} \neq \perp \Rightarrow \\ & \quad \hat{K}(p_s, o_{Sos}, x_{reply}, x_{info}, x_{sid}) \end{aligned}$$

□

Complexity. Three quantities determine the complexity of solving the derived ALFP clause:

- the number n of names used in the program,
- the maximal nesting depth of variables, bounded by $m = \max_{\ell \in \mathbf{Lab}} |\mathbf{L}.\ell|$, and
- the maximal length of any sent message, bounded by $k = \max_{u \bullet u' \bar{u} \in s} |\bar{u}|$.

The size of the logical universe is bounded by n , while m and n decide the maximal arity of relations. Furthermore, the maximal nesting depth is decided by m , which, by Proposition 1 of [11], results in a complexity bound of $\mathcal{O}(n^{3+k+m})$. This is exponential in the worst case, which is only realised by pathological service specifications, s , where the number of sequenced inputs (m) and/or the arity of sent messages (k) are linear in the size (n) of s . For realistic CWS services, the complexity will be polynomially bound by constants m and k . For the accident service the solution was found in less than a second.

5 Related Work

The separation of scope and binding occurrence of a variable is originally due to Parrow and Victor, who used it to model locally shared memory in the Fusion Calculus [14]. In contrast to CWS, there is only one syntactic category, names, rendering input and output symmetric, i.e., input and output designation can be exchanged while yielding the same fusion.

In [3], the calculus D-Fusion is proposed. It extends the Fusion Calculus with another binder similar to restriction in π -calculus. It retains the symmetry of actions, in fact, it does not even distinguish input and output at all. Apart from this symmetry, D-Fusion is quite close to CWS. However, we prefer to consider CWS as a subset of the COWS [6] calculus, because the latter is more focused on correlations and their use in service-oriented computing.

COWS retains a variant of separation of input and name binding in order to faithfully model correlation. In contrast to Fusion and D-Fusion, inputs and outputs are clearly distinguished, in particular by using two syntactic categories, names and variables, where only substitutions of names for variables are possible. COWS also features pattern matching in input prefixes facilitating correlations. While we, too, embrace these concepts we discard some of the technicalities associated with session management and do not consider the fault and compensation facilities provided by COWS.

To the best of our knowledge [1] the present static analysis is the first static analysis, which is not a type system, developed for the Fusion Calculus or its descendants. Also, it seems to be the first static analysis of correlation. Our analysis is relational in form, which has previously been investigated in the simpler context of the π -calculus [8] based on a reaction semantics. The present development retains the simplicity of the former, even in the context of a more complicated calculus with a labelled structural operational semantics, and testifies to the flexibility of the Flow Logic specification style [13]. In contrast, previous relational approaches [16, 4], fashioned within the framework of Abstract Interpretation, have relied on highly customised syntax and semantics and are not easily extended beyond the original context of the π -calculus.

In the context of COWS, Lapadula et. al. use type systems in order to enforce distribution policy annotations [7]. For this approach to work the user has to annotate each piece of data with a region of maximal dissemination (a set of service principals). Static inference combined with a typed semantics, performing

appropriate run-time checks, then ensures that the policy is never violated. In contrast, our approach relies neither on user-provided annotations, nor dynamic type-checking. It is fully static and automatically computes a very precise estimate of the data-sets that may reach every single program point. The specified information is very general; hence a region based policy can easily be tested against the computed result. In [7], the complexity of the type system is not discussed, but the general tendency is that checking is polynomial, whereas inference is exponential. In the case of the Succinct Solver, however, there is no complexity gap between checking and inference [12] – both are polynomial for non-pathological specifications.

6 Conclusion

In this paper we have delineated the process calculus CWS for implementing correlation based services in order to focus on the essentials of correlation, e.g., separation of binding and input and pattern matching. We expect that our analysis of CWS transfers easily to other correlation based process calculi. CWS is formulated in a form resembling COWS but lacking its orchestration constructs. In future extensions, we aim at adding these features.

We developed a relational analysis for the CWS process calculus and showed its usefulness for ensuring that service invocations do not interfere in malign ways. We have based our work on a recent relational analysis developed for the π -calculus [8] thereby supporting the claim that the Flow Logic framework facilitates transferring analysis insights between languages (being programming languages or process calculi). While more powerful approaches to relational analysis exist, in particular the work of polyhedral analysis of certain π -processes [16, 4], they are substantially harder to transfer to other language because they rely on a special “normal form” for processes established a priori (in case of the π -calculus in the PhD-thesis of [15]).

Despite the guidance offered by the Flow Logic framework and the relational analysis developed for the π -calculus in [8], correlations, i.e., the separation of scope from binding, have presented profound obstacles that we have managed to solve. A key ingredient is the use of \perp to denote the “presence” of a variable that has not yet received its value; this technique is being used in rather deep ways to ensure the semantic correctness of the analysis. The correctness result follows the approach pioneered in [8] in making use of a subject reduction result where “analysability” is not preserved under reduction, whereas the more complex notion of “analysability of all permitted substitution instances” is.

References

1. Personal communication with B. Victor, October 2007.
2. C. Bodei, P. Degano, F. Nielson, and H. R. Nielson. Static analysis for the π -calculus with applications to security. *Information and Computation*, 168:68–92, 2001.

3. M. Boreale, M. G. Buscemi, and U. Montanari. D-fusion: A distinctive fusion calculus. In W.-N. Chin, editor, *APLAS*, volume 3302 of *LNCS*, pages 296–310. Springer, 2004.
4. J. Feret. Dependency analysis of mobile systems. In D. L. Métayer, editor, *ESOP*, volume 2305 of *LNCS*, pages 314–330. Springer, 2002.
5. C. Guidi, R. Lucchi, R. Gorrieri, N. Busi, and G. Zavattaro. SOCK: A calculus for service oriented computing. In A. Dan and W. Lamersdorf, editors, *ICSOC*, volume 4294 of *LNCS*, pages 327–338. Springer, 2006.
6. A. Lapadula, R. Pugliese, and F. Tiezzi. A calculus for orchestration of web services. In R. De Nicola, editor, *Proc. of 16th European Symposium on Programming (ESOP'07)*, LNCS. Springer, 2007.
7. A. Lapadula, R. Pugliese, and F. Tiezzi. Regulating data exchange in service oriented applications. In *Proc. of IPM International Symposium on Fundamentals of Software Engineering (FSEN'07)*, volume 4767 of *LNCS*, pages 223–239. Springer, 2007.
8. F. Nielson, H. R. Nielson, J. Bauer, C. R. Nielsen, and H. Pilegaard. Relational analysis for delivery of services. In *Trustworthy Global Computing*, 2007. To appear.
9. F. Nielson, H. R. Nielson, and C. Hankin. *Principles of Program Analysis*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1999.
10. F. Nielson, H. R. Nielson, H. Sun, M. Buchholtz, R. R. Hansen, H. Pilegaard, and H. Seidl. The succinct solver suite. In *TACAS*, pages 251–265, 2004.
11. F. Nielson and H. Seidl. Control-flow analysis in cubic time. In *Proc. ESOP'01*, number 2028 in LNCS, pages 252–268. Springer, 2001.
12. F. Nielson, H. Seidl, and H. R. Nielson. A succinct solver for ALFP. *Nord. J. Comput.*, 9(4):335–372, 2002.
13. H. R. Nielson and F. Nielson. Flow Logic: a multi-paradigmatic approach to static analysis. In T. Mogensen, D. Schmidt, and I. H. Sudborough, editors, *The Essence of Computation: Complexity, Analysis, Transformation. Essays dedicated to Neil D. Jones*, volume 2566 of *LNCS*, pages 223–244. Springer Verlag, 2002.
14. J. Parrow and B. Victor. The fusion calculus: Expressiveness and symmetry in mobile processes. In *LICS*, pages 176–185, 1998.
15. D. N. Turner. *The Polymorphic Pi-calculus: Theory and Implementation*. PhD thesis, University of Edinburgh, 1996.
16. A. Venet. Automatic determination of communication topologies in mobile systems. In G. Levi, editor, *SAS*, volume 1503 of *LNCS*, pages 152–167. Springer, 1998.

A Subject Reduction Revisited

We have established a substitution based subject reduction result in Theorem 1 based upon a correctness predicate relating correct analyses results to substitution instances of the programme originally analysed. Now, we shall present a counterexample illustrating that general subject reduction with respect to correct analyses fails. Consider the following service and a computation step:

$$\begin{array}{c} [x]^1 \\ \mid p \bullet o?x^2.\mathbf{0} \\ \mid p \bullet o!a^3 \\ \mid p \bullet o?x^4.p \bullet o!\langle \rangle^5 \end{array} \xrightarrow{p \bullet o[\emptyset]x^2a^3} p \bullet o?a^4.p \bullet o!\langle \rangle^5$$

It is obvious that a valid analysis result of the left hand side may comprise $\hat{R}.4 = \{\mathbf{a}, \perp\}$ and $\hat{R}.5 = \{\mathbf{a}\}$, because the X in [REQ] filters \perp away. However, analysing the right-hand side in isolation requires $X = \{\mathbf{a}, \perp\} \subseteq \hat{R}.5$ which is clearly not satisfied by the result given. The filter does not work in this case constituting an example of standard subject reduction not holding. A key insight is that correct analyses are not preserved whenever we have a match involving more than constants, e.g. variables. More generally speaking, the problem occurs when we have a static analysis with localised environments, where we incorporate the results of matchings into these environments. This observation becomes apparent in [8], too.

B Proof of Theorem 1

In order for the proof to work, we need to define a stronger induction hypothesis in form of a correctness predicate parameterised by a substitution σ . Definition 2 remains unchanged except for the fact that we require the substitution to be compatible with the element \bar{w} of the analysis result that we use to create a substitution instance with. We shall write $\overline{\{x_1, \dots, x_k\}}$ to denote the vector $\langle x_1 \dots x_k \rangle$ (the concrete order is irrelevant) and $\text{var}(\bar{u})$ to denote the vector that remains of \bar{u} after removing all non-variables.

Definition 3 (Parameterised Correctness Predicate). *A service s satisfies the parameterised correctness predicate with respect to s_\star and σ , written $\hat{R}, \hat{K} \models_\sigma^{s_\star} s$ if and only if (1-4) hold.*

1. $\hat{R}, \hat{K} \vdash_{L,F} s_\star$
2. $\forall \ell \in E : \epsilon \in \hat{R}.\ell$
3. For all $p \bullet o?\bar{u}^\ell.s' \in \mathcal{E}(s)$ there exists a subexpression $p \bullet o?\bar{v}^\ell.s''$ of s_\star s.t.
 - $\hat{R}, \hat{K} \vdash_{L,F} p \bullet o?\bar{v}^\ell.s''$ and
 - $\exists \bar{w} \in \hat{R}.\ell : p \bullet o?\bar{u}^\ell.s' \equiv p \bullet o?\bar{v}^\ell.s''[\bar{w}/L.\ell]$ and $\Pi_{\text{dom}(\sigma) \otimes L.\ell}(\bar{w}) = \sigma(\text{dom}(\sigma))$.
4. For all $v \bullet v!\bar{u}^\ell \in \mathcal{E}(s)$ there exists a subexpression $u \bullet u!\bar{u}^\ell$ of s_\star s.t.
 - $\hat{R}, \hat{K} \vdash_{L,F} u \bullet u!\bar{u}^\ell$ and

$$- \exists \bar{w} \in \hat{R}. \ell : v \bullet v' ! \bar{v}^\ell \equiv u \bullet u' ! \bar{u}^\ell [\bar{w}/L.\ell] \text{ and} \\ \Pi_{\overline{dom(\sigma)} @ L.\ell}(\bar{w}) = \sigma(\overline{dom(\sigma)}).$$

Notice that Definition 3 and Definition 2 coincide for empty substitutions. Therefore, the proof of Theorem 1 follows directly from case 3 of Lemma 2 below, when we choose $\sigma = \emptyset$.

Lemma 2. *Let s_1 and s_2 be services. The following statements hold.*

1. *For all $\sigma : \mathbf{Var} \rightarrow \mathbf{Name}$ holds: If $\hat{R}, \hat{K} \models_{\sigma}^{s_*} s_1$ and $s_1 \xrightarrow{(p \bullet o) < \bar{n}^\ell} s_2$ then $\hat{R}, \hat{K} \models_{\sigma}^{s_*} s_2$ and $\langle po\bar{n} \rangle \in \hat{K}$.*
2. *For all $\sigma : \mathbf{Var} \rightarrow \mathbf{Name}$ such that $dom(\sigma) = \{x \mid x \in \{\bar{u}\}\}$ holds: If $\hat{R}, \hat{K} \models_{\sigma}^{s_*} s_1$ and $s_1 \xrightarrow{(p \bullet o) > \bar{u}^\ell} s_2$ and $po\bar{u} \cdot \sigma \in \hat{K}$ then $\hat{R}, \hat{K} \models_{\sigma}^{s_*} s_2$.*
3. *If $\hat{R}, \hat{K} \models_{\sigma}^{s_*} s_1$ and $s_1 \xrightarrow{p \bullet o [\sigma] \bar{u}^{\ell_i} \bar{n}^{\ell_o}} s_2$ then $\hat{R}, \hat{K} \models_{\sigma}^{s_*} s_2$.*

When proving Lemma 2 we make use of the fact, that Lemma 1 holds for the parameterised correctness predicate and all substitutions. All three parts of Lemma 2 are proven by induction on the inference tree used to establish the semantic reduction. We start with the base case of case 1, that is, rule [INV]. Let σ be an arbitrary substitution and let

$$p \bullet o ! \bar{n}^\ell \xrightarrow{(p \bullet o) < \bar{n}^\ell} \mathbf{0}$$

such that $\hat{R}, \hat{K} \models_{\sigma}^{s_*} p \bullet o ! \bar{n}^\ell$. It is obvious that $\hat{R}, \hat{K} \models_{\sigma}^{s_*} \mathbf{0}$. We know that there exists $u \bullet u' ! \bar{u}^\ell$ in s_* and $\bar{w} \in \hat{R}. \ell$ such that $uu' \bar{u} [\bar{w}/L.\ell] = po\bar{n}$. Hence by rule [RINV], we obtain $po\bar{n} \in \hat{K}$. The inductive cases follow easily by (the variation of) Lemma 1.

When we prove part 2 of Lemma 2, it suffices to concentrate on the base case, that is, rule [REC]. So assume

$$p \bullet o ? \bar{u}^\ell . s \xrightarrow{(p \bullet o) > \bar{u}^\ell} s$$

and let σ be an appropriate substitution such that $\hat{R}, \hat{K} \models_{\sigma}^{s_*} p \bullet o ? \bar{u}^\ell . s$. We know that there exists a $p \bullet o ? v^\ell . s'$ in s_* and $\bar{w} \in \hat{R}. \ell$ such that

$$\hat{R}, \hat{K} \vdash_{L,F} p \bullet o ? v^\ell . s' \tag{1}$$

$$p \bullet o ? u^\ell . s \equiv p \bullet o ? v^\ell . s' [\bar{w}/L.\ell] \tag{2}$$

$$\Pi_{\overline{dom(\sigma)} @ L.\ell}(\bar{w}) = \sigma(\overline{dom(\sigma)}) \tag{3}$$

Effectively, according to analysis rule [RREC], it suffices to show, that $\bar{w} \in X$ (referring to the X in [RREC]), that is, $\Pi_{po\bar{v} @ L.\ell}(\bar{w}) \in \hat{K}$. Since we know $po\bar{u} \cdot \sigma \in \hat{K}$ by the assumption of part 2 of Lemma 2, and the combination of (2) and (3)

yields $(p \bullet o?u^\ell.s) \cdot \sigma \equiv p \bullet o?v^\ell.s'[\bar{w}/L.\ell]$ we can conclude this requirement and thus the proof of this part.

Proving part 3 essentially amounts to looking at two rules, [COM] and [DEL_{sub}]. We start with the first one assuming

$$\frac{s_1 \xrightarrow{(p \bullet o) \triangleright \bar{u}^{\ell_i}} s'_1 \quad s_2 \xrightarrow{(p \bullet o) \triangleleft \bar{n}^{\ell_o}} s'_2 \quad \mathcal{M}(\bar{u}, \bar{n}) = \sigma}{s_1 \mid s_2 \xrightarrow{p \bullet o[\sigma] \bar{u}^{\ell_i} \bar{n}^{\ell_o}} s'_1 \mid s'_2}$$

and $\hat{R}, \hat{K} \models_{\sigma}^{s^*} s_1 \mid s_2$, hence $\hat{R}, \hat{K} \models_{\sigma}^{s^*} s_1$ and $\hat{R}, \hat{K} \models_{\sigma}^{s^*} s_2$. We can thus apply the induction hypothesis on s_2 obtaining

$$\hat{R}, \hat{K} \models_{\sigma}^{s^*} s'_2 \tag{4}$$

$$po\bar{n} \in \hat{K} \tag{5}$$

From (5) and $\mathcal{M}(\bar{u}, \bar{n}) = \sigma$ we deduce $po\bar{u} \cdot \sigma \in \hat{K}$, which allows us to apply the induction hypothesis for s_1 yielding $\hat{R}, \hat{K} \models_{\sigma}^{s^*} s'_1$. Together with (4), this completes the proof of part 3 for rule [COM].

We are left to show part 3 for rule [DEL_{sub}] that looks as follows:

$$\frac{s \xrightarrow{p \bullet o[\sigma \uplus [x \mapsto n]] \bar{u}^{\ell_i} \bar{n}^{\ell_o}} s'}{[x]^\ell s \xrightarrow{p \bullet o[\sigma] \bar{u}^{\ell_i} \bar{n}^{\ell_o}} s' \cdot [x \mapsto n]}$$

The proof of this part amount to showing the following:

$$\hat{R}, \hat{K} \models_{\sigma}^{s^*} [x]^\ell s \text{ implies } \hat{R}, \hat{K} \models_{\sigma \uplus [x \mapsto n]}^{s^*} s \tag{6}$$

$$\hat{R}, \hat{K} \models_{\sigma \uplus [x \mapsto n]}^{s^*} s' \text{ implies } \hat{R}, \hat{K} \models_{\sigma}^{s^*} s' \cdot [x \mapsto n] \tag{7}$$

Fact 6 allows us to apply the induction hypothesis, which yields exactly the assumption of (7). The application of (7) then concludes this case and the overall proof (since again the inductive cases are straightforward). We are left to show (6) since (7) is obvious from Definition 3. Observe that $x \# F.\ell_i = \ell$. We can thus deduce (from the definition of Y_x in analysis rule [RREC] that $n \in \Pi_{x \text{ @ } L.\ell'}(\hat{R}.\ell')$ for all $\ell' \in F.\ell$. This suffices to prove (6) and Lemma 2.