DOI 10.1007/s00165-013-0284-5
BCS © 2013
Formal Aspects of Computing (2014) 26: 865–918

Formal Aspects of Computing

The stream-based service-centred calculus: a foundation for service-oriented programming

Luís Cruz-Filipe[1], Ivan Lanese[2], Francisco Martins[3], António Ravara[4] and Vasco Thudichum Vasconcelos[3]

[1] Escola Superior Náutica Infante Dom Henrique, Oeiras, Portugal
[2] Focus Team, University of Bologna/INRIA, Bologna, Italy
[3] Universidade de Lisboa, Faculdade de Ciências & LaSIGE, Lisbon, Portugal
[4] Universidade Nova de Lisboa, Faculdade de Ciências e Tecnologia & CITI, Caparica, Portugal

Abstract. We give a formal account of stream-based, service-centered calculus (SSCC), a calculus for modelling service-based systems, suitable to describe both service composition (orchestration) and the protocols that services follow when invoked (conversation). The calculus includes primitives for defining and invoking services, for isolating conversations (called sessions) among clients and servers, and for orchestrating services. The calculus is equipped with a reduction and a labelled transition semantics related by an equivalence result. SSCC provides a good trade-off between expressive power for modelling and simplicity for analysis. We assess the expressive power by modelling van der Aalst workflow patterns and an automotive case study from the European project Sensoria. For analysis, we present a simple type system ensuring compatibility of client and service protocols. We also study the behavioural theory of the calculus, highlighting some axioms that capture the behaviour of the different primitives. As a final application of the theory, we define and prove correct some program transformations. These allow to start modelling a system from a typical UML Sequence Diagram, and then transform the specification to match the service-oriented programming style, thus simplifying its implementation using web services technology.

1. Introduction

Enterprise application integration, either to reuse legacy code, or to combine third-party software modules, has long been tackled by several middleware proposals, namely using message brokers or workflow management systems. As the popularity of using the Web increased, traditional middleware was forced to provide integration across companies. The technologies developed lay in the concept of *Web service*: a way of exposing (to the Web) the functionalities performed by internal systems and making them dynamically searchable (discoverable and accessible through the Web) and composable, allowing adaptability and reusability [ACKM03].

Correspondence and offprint requests to: A. Ravara, E-mail: aravara@fct.unl.pt

1.1. A formal approach to service-oriented programming

Nowadays, Web Services are one of the main technologies used to deploy applications that coordinate behaviours over the Web. Contributing to the success of the technology are the currently available standards [GHM+07, CCMW01, BCC+04, AAA+07] that, in order to achieve business goals, allow easy orchestration of different services (distributed and belonging to different organizations), while maximizing interoperability. While standards and programming tools are continuously improving, the formal bases of Service-Oriented Computing (SOC) are still uncertain: there is an urgent need for models and techniques allowing the development of applications in a safe manner, while checking that systems provide the required functionalities. These techniques should be able to deal with the different aspects of services, including their dynamic behaviour.

To model systems, and, in particular, to be able to reason and ensure properties about specifications of service-oriented systems, one needs mathematical tools. Process calculi are one suitable tool, providing not only a description language, but a rigourous semantics as well, allowing the proof of relevant properties. Process calculi give precise semantics to system specifications, and they come equipped with a rich toolbox of analysis techniques, including type systems and contextual equivalences. When defining a calculus for SOC, different aspects influence the choice of primitives and of their behaviour, and a careful trade-off between expressiveness and suitability to analysis should be found. Our main concerns have been threefold.

Expressiveness of the language The calculus should be able to express in a direct way the different kinds of interactions that characterize SOC: invocations of services, client-server conversations, and interactions among different client-server pairs. We use three different classes of operators to this end: services, sessions, and streams. We show, via examples, that these are enough to model various kinds of SOC scenarios. We stress in particular the importance of interactions using streams, which is the heart of SSCC orchestration. Other constructs such as tuple spaces or shared memory would be as expressive as streams, but would be more difficult to analyze. Furthermore, primitives similar to our streams are commonly used in practice, for instance when programming Service Oriented Architectures.

Expressiveness of the analysis The elements to be analyzed should correspond to explicit elements in the calculus. Concerning the three classes of operators in the previous paragraph, service definition is fundamental to speak about service availability. It also allows easy extensions for service discovery based on quality of service. Sessions, instead, allow to analyze client-server compatibility and to study behavioural-based service discovery. Other mechanisms, such as BPEL correlation sets [AAA+07], would make these analyses more complex, since they rely on run-time values for determining the communication patterns, spreading the protocol code throughout the whole program. Streams are, on the one hand, used for modelling service coordination, and, on the other hand, needed to study global properties of systems such as deadlock-freedom and the fact that the system satisfies a more abstract specification.

Computability of the analysis Static analysis should be decidable, possibly also efficient to compute. Thus, the allowed communication patterns should be constrained whenever this does not destroy expressiveness. In our calculus, streams and sessions are static, and the dynamism is concentrated in service invocation. To stress the effect of these considerations on the design decisions, we give some "proof of concept" analyses to illustrate how to exploit the features of the calculus.

1.2. Our proposal

We present stream-based, service-centered calculus (SSCC), a calculus for modelling service-based systems, inspired by SCC [BBC+06] and Orc [MC07, KCM06], and developed with the above considerations in mind. A prominent feature is that it captures in a direct way the main activities in service-oriented computations: *definition and invocation of services*, *long-running interactions* between the service invoker and the service provider, and *orchestration of complex computations* exploiting services as building blocks.

SSCC builds on SCC [BBC+06], and tries to improve, in particular, the suitability for modelling orchestration. To this end, it introduces a new construct, called *stream*, with the aim of collecting the results from some ongoing computations and make them available for new ones. This is the main aspect that differentiates SSCC from CaSPiS [BBNL08, BBNL14], the most direct evolution of SCC. While proposing interesting concepts, like sessions, and featuring services as first class entities, SCC looks not fully adequate for service composition. In fact, the only way for a session to interact with other client-server pairs is the **return** primitive, and the functional style of invocation is not adequate for modelling complex patterns of interactions, such as van der Aalst workflow patterns [vdAtHKB03]. To overcome these problems we introduce streams and we allow non-persistent service invocations, thus enhancing the expressiveness of the calculus, while making it easier to program with. This design choice has been taken in order to simplify static analysis techniques, trying to find a suitable trade-off between expressiveness of the chosen primitives and suitability to analysis. In particular, since stream names cannot be communicated, their scope is known statically.

Another source of inspiration was Orc [MC07], a basic programming model for orchestration of Web services. Here a few coordination constructs are used to model the most common patterns, and a satisfying expressiveness is claimed by presenting a formalization of all van der Aalst workflow patterns [CPM06]. However, in order to model the more challenging patterns, special sites (the basic computation entity in Orc) are required, acting, e.g., as semaphores. This is a coordination concern, and in our opinion should be addressed within the language. Thus we introduced more basic mechanisms to tackle all the coordination concerns inside the calculus (most of Orc operators can be expressed as macros in our model). Also, we introduced conversations, which are absent in Orc, to model service behaviour (Orc leaves this unspecified).

1.3. Contributions

The main contribution of this paper is SSCC, a language for modelling service-oriented systems based on the session communication paradigm and equipped with a formal semantics. This paper aggregates two conference papers [LVMR07a, CFLM$^+$08] and two technical reports [LVMR07b, CFLM$^+$07]. The preliminary results in the conference papers are extended and fully proved. Thus, much of the material here presented is new.

A clear separation of concerns: conversation and orchestration. As discussed above, the main concerns leading to this language were expressiveness and intuitiveness for modelling and suitability for analysis. Concerning the modelling part, we want to emphasize here some aspects. The calculus allows for the description of service interaction and of service orchestration using distinct mechanisms; the conversation between parties engaged in a service interaction is described by a series of value send/receive, isolated inside a session, while the orchestration of services is performed using the stream operations. The two communication methods are orthogonal, as witnessed for instance by the fact that their scopes give rise to separate nesting hierarchies. This choice simplifies both the type system and the observational semantics, as discussed below. Notice that Orc [MC07] lacks the conversation primitives, and that both Orc and SCC [BBC$^+$06] feature insufficient orchestration. Furthermore, in both SCC [BBC$^+$06] and CaSPiS [BBNL08, BBNL14] orchestration is not orthogonal w.r.t. sessions, thus making it difficult to program cross-sessions orchestration patterns.

A flexible programming style. Service orchestration and service conversation are both easily structured in SSCC. (Keep in mind that SSCC is a process calculus, not a full-fledged programming language.) We were in fact able to encode all van der Aalst workflow patterns [vdAtHKB03] (apart from the ones that require termination), in intelligible code. We tried the same exercise using SCC and the results for some patterns were not satisfying at all. As a more challenging test, we put SSCC at work by modelling the Automotive case study [BFF$^+$06] of the European project Sensoria [pro]. Again, SSCC passed the test. Actually, it was able to model different programming styles, from object-oriented to session-oriented to request-response oriented.
Concerning the analysis part, we concentrated on a type system and on proofs of equivalence based on a contextual equivalence.

A type discipline. We provide a simple static analysis system to check the compatibility between service definition and service invocation, as well as protocol sequentiality. Our type system is inspired in the results achieved for session types [GH05, HVK98, YV06, THK94]. Interestingly, it exploits the separation of concerns between conversation and orchestration. In fact, the type of a process has two components, one concerning its conversation behaviour and another concerning its orchestration behaviour. This allows to simplify the reasoning in some cases: for instance, the simplified type system for session sequentiality has no need to consider the orchestration part. The same simplification is not possible in other calculi where the two aspects are intertwined.

A coinductive contextual congruence. The contextual congruence in SSCC coincides with a bisimilarity that is also a (non-input) congruence and gives rise to some axiomatic laws. These laws not only clarify the relationship among language constructs, but they also allow for meaningful program transformations.

Program transformations. We exploited the theory for defining and proving correct some program transformations. Those transformations allow to start designing a system from a standard UML Sequence Diagram [Amb04] (based on the object-oriented paradigm). Such a diagram has a direct translation into SSCC, and it induces an object-oriented style. This style does not match the session-oriented style normally used for service-oriented systems. Our first program transformation allows to transform it to fit this second style. However, most implementation technologies based on the service-oriented paradigm do not fully support general sessions, but only request-response communications, which can be seen as a particular case. Our second transformation allows to transform the program to fit this implementation requirement. Interestingly, this transformation exploits the type system for session sequentiality, since it is correct only for sequential sessions.

1.4. Related work

Among process calculi, the π-calculus (and its variants) has been frequently used in SOC. We claim that general purpose concurrent calculi are not suitable for our aims, since the different communication patterns are mixed, and most of the interesting properties are not directly reflected in text of programs. Thus, these calculi do not satisfy the requirements above. Different proposals use types, e.g., session types [GH05, HVK98, YV06, THK94], to solve this problem, but since they allow free π-calculus communications the analysis becomes difficult. We consider our proposal as some kind of tamed π-calculus, with a good trade-off between expressiveness for SOC systems and suitability to analyze SOC-related properties.

This approach was born inside the EU project Sensoria [pro] on Software Engineering for Service-Oriented Overlay Computers. The first outcome of this line of research has been SCC [BBC+06]. SCC had the merit of introducing services and sessions as first-class entities in the calculus, but lacked good primitives for orchestrating groups of services. The only orchestration-oriented mechanism was the **return** primitive, allowing a session to return a value to its enclosing session, but implementing complex interaction patterns using it was an hard task. We show examples supporting this claim in Sect. 3.1 and in Appendix B, when modelling workflow patterns [vdAtHKB03].

SSCC has been the first proposal [LVMR07a] extending SCC to provide easier orchestration, followed by the Conversation Calculus [VCS08] and CasPiS [BBNL08, BBNL14]. CasPiS approach is the closest to ours. The main difference is that CasPiS uses pipelines for orchestration instead of our streams. Pipelines allow to redirect session communications to another session. The difficulty in using pipelines is that the same messages can be used both for session communication and for orchestration. Instead we use two orthogonal features. The separation of concerns is one of the cornerstones of SSCC approach, since it allows one to use different analysis techniques on messages used for different purposes. For instance, in the type system in Sect. 2.4, session communications are modelled using behavioural types, while streams have simpler static types. A similar approach would not be as effective in CasPiS. See, e.g., [BM08] for a type system for CasPiS.

The Conversation Calculus introduces the notion of conversation, a medium where different processes can interact. Conversations are more expressive and more complex than SSCC sessions or streams, since they allow for direct multiparty interactions, while such a kind of interaction has to be programmed in some well-structured way in SSCC. As a consequence, static analysis in the Conversation Calculus is much more difficult, as can be seen by comparing its type system ensuring well-formed communications in [CV10] and our type system in Sect. 2.4, which has a similar aim.

Multiparty sessions, similar to conversations, have been studied also in use [BLMT08]. There, two running multiparty sessions can dynamically merge. This makes static analysis even more difficult.

Two other calculi have been considered inside the Sensoria project: COWS [LPT07] and SOCK [BGG+06]. Communication in these two calculi is not based on private sessions: they use instead the correlation set mechanism. Using this mechanism the recipient of a message is found based on (part of) the message content. For instance, a message including a specific user name is matched by the session created for managing all the communications for that specific user name. Correlation sets are used by the main technology for service-oriented orchestration: WS-BPEL [AAA+07]. However, since communication patterns depend on runtime values, they are more difficult to analyze statically. Only basic properties can be statically proved, as in [MC11]. In fact, private sessions can be seen as a constrained use of correlation sets ensuring more behaviour predictability.

While both COWS and SOCK are based on correlation sets, they differ in many respects. In fact, COWS is more abstract, in the spirit of name passing calculi, while SOCK is essentially an imperative language extended with primitives for correlation-based communication. This makes SOCK more complex, but closer to real service-oriented languages. In fact, Jolie [Jol, MGZ07], a full-fledged language for programming service-oriented applications, is based on the semantics of SOCK. We refer to [CNP+11] for a comparison of the calculi proposed inside the Sensoria project, and to [LRV11] for a comparison of their behavioural theories.

Orc [MC07] is a language that focuses on orchestrating the concurrent invocation of services to achieve a given goal. The theory is built upon three composition operators—parallel composition, sequencing and selective pruning—while relying on a number of primitive services (sites, in the Orc's terminology) to perform basic computations. The symmetric parallel composition of f and g, written as f | g, permits independent computations and publications (i.e., outputs) from f and g. Sequencing of f and g, written as f >x >g, conveys *all* values x published by f to g. Selective pruning of f and g, written as g **where** x :∈ f conveys *some* value x published by f to g. We find that Orc does not provide enough support for orchestrating computations and does not support separation of concerns between orchestration and computation. In particular, and with respect to SSCC, Orc does not provide primitives for expressing conversations among services. This paper shows that flexible programming together with proofs of correctness for programming transformations can only be achieved in a language that takes conversation as a primitive concept.

Another thread of research [CHY07, HYC08, BHTY10, LGMZ08, CM13] aims at capturing the principles behind Web service based business processes. A global description of communication behaviour allows one to generate automatically an "endpoint-based" description of each participant to the protocol, a projection of the global scenario. We are at the same abstraction level of the endpoint calculus, but this one relies on multiparty session communication, and the considerations above for the Conversation Calculus apply.

1.5. The rest of this paper

We start by presenting the main ideas of the process language through two basic examples (Sect. 2.1). Then, Sects. 2.2 and 2.3 present, respectively, the syntax and the operational semantics (both a reduction-based system and a labelled transition system, shown to coincide on silent transitions) of the process description language.

Section 2.4 presents a first theoretical mechanism to study the (static) behaviour of processes, a type system for SSCC. Equipped with this mechanism, we present in Sect. 3 more elaborate examples of systems specification in SSCC: Sect. 3.1 presents the encoding of common workflow patterns, and Sect. 3.2 presents a non-trivial programming exercise—the automotive case study.

Section 4 presents another theoretical mechanism to study the (dynamic) behaviour of processes, giving notions of behavioural equivalences and proving some of their properties.

As an application of these theoretical results, Sect. 5 discusses how different programming styles applied to the same concrete problem may lead to seemingly different implementations, and shows, using the properties from the previous sections, that these implementations are indistinguishable if one abstracts from internal details. Furthermore, this section introduces a generic set of rules allowing to program SSCC processes according to different programming styles, and presents a set of (behaviour-preserving) transformation rules allowing one to move from one style to another.

The last section summarizes the results obtained, and explores possible directions for future work.

2. SSCC: examples, syntax, and operational semantics

To motivate the basic constructs of the calculus, we present them one by one, incrementally building a simple example. A more elaborate example follows, to illustrate the expressive power of the calculus. Afterwards, we define rigorously the syntax, the operational semantics, and the type system of SSCC.

2.1. Modelling: basic examples

A first example. We start by defining a simple process to deliver the price of a room for a given date at a given hotel.

```
(date) ⟨query−the−hotel−db⟩.price
```

Here, the parentheses in (date) indicate the reception of a value, and an identifier alone, as in price, means publishing a value. We use the notation ⟨ activity ⟩ to refer to code for some activity which is not (yet) specified. For instance, ⟨query−the−hotel−db⟩ stands for code which looks, in the hotel database, for the price of the stay; when running, this code uses an actual value for the parameter (date, in this example). Hotel bologna may turn the process above into a service definition, by writing:

```
bologna ⇒ (date) ⟨query−the−hotel−db⟩.price
```

Here bologna is the name of a service that can be invoked to interact with the hotel, and whose behaviour is as above. A client is supposed to meet the expectations of the service by providing a date and requesting a price:

```
bologna ⇐ 31Dec2012.(price) ⟨use−price⟩
```

Here we have an invocation (⇐) of service bologna, whose behaviour first sends 31Dec2012 as date, then waits for an answer, which is stored in price. The code continues with some unspecified use of the received price.

When the service provider (⇒) and the service client (⇐) get together, by means, e.g., of parallel composition, a *conversation* takes place, and values are exchanged in both directions.

Now suppose that a broker comes to the market trying to provide better deals for its clients. The behaviour of the broker is as follows: it asks prices to three hotels that it knows of (bologna, lagoa, lisbon), waits for two results, and publishes the best offer of the two. Calling the three services for the same given date is as above:

```
bologna ⇐ date.(price1) ... |
lagoa ⇐ date.(price2) ... |
lisbon ⇐ date.(price3) ...
```

Note the use of parallel composition | to perform the three invocations concurrently.

In order to collect the prices for further processing, we introduce a *stream* constructor, playing the role of a *service orchestrator*. The various prices are fed into the stream; another process reads the stream. We write it as follows.

```
stream
  bologna ⇐ date.(price1).feed price1 |
  lagoa ⇐ date.(price2).feed price2 |
  lisbon ⇐ date.(price3).feed price3
as f in
  f(x).f(y).⟨publish−the−min−of−x−and−y⟩
```

To write price1 into a stream we use the syntax **feed** price1. To read a value from stream f we use f(x).⟨use−x⟩, where ⟨use−x⟩ is the code that uses the parameter x. Writing is an anonymous operation (feeds to the nearest enclosing stream), whereas reading is named. The above pattern is so common that we provide a special syntax for it. Inspired by Orc, **call** stands for an invocation of a (parametric) service, whereas $\mathsf{P}>^{\mathbf{n}} x_1 \ldots x_n>\mathsf{Q}$ models the flow of n values from process P to process Q, via a stream. We refer to Fig. 9 for a formal definition of the derived syntactic constructs.

```
(call bologna(date) |
 call lagoa(date) |
 call lisbon(date)) >^2
   x y > ⟨publish−the−min−of−x−and−y⟩
```

To complete the example we rely on a min service, chaining the first two answers, and publishing the result. We also turn the code into a service named broker.

```
broker ⇒ (date).(
  (call bologna(date) |
   call lagoa(date) |
   call lisbon(date)) >^2
     x y > call min(x,y) >^1 m > m)
```

In detail, the first two values returned by the services, x and y, are used to invoke service min (note again the use of special syntax **call**). The value returned by service min is bound to m and returned to the client of service broker (the last occurrence of m is the return value).

Note that a client interacts with the broker using the same protocol that it would use to interact with a particular hotel named broker. The downside of this approach is that the client does not know which hotel offers the best price; it is not difficult to adapt the example to overcome this limitation.

Using **call** and $\mathsf{P}>^{\mathbf{n}} x_1 \ldots x_n>\mathsf{Q}$ we have avoided explicitly mentioning streams altogether. Direct stream manipulation can however be quite handy. The following example shows a broker that logs all three answers by sending them to a specific service log, while publishing the best price of the first two (cf. the Discriminator Pattern [vdAtHKB03] in Sect. 3.1).

```
stream ... as f in
  f(x).f(y).call min(x,y) >^1
    m > (m | f(z).log ⇐ x.y.z)
```

Our language is equipped with a notion of types,[1] allowing to statically filter programs that may incur in *conversation errors*, such as both the service provider and the client expecting a value at the same time (while if one is expecting a value the other one should be sending it), or the service expecting a value while the client has already terminated its execution. Returning to the hotel example, we can easily see that the conversation between the service provider (⇒) and the client (⇐) is, from the point of view of the provider, as follows: expect a date; send a price; terminate. The whole process of querying the hotel database to obtain the price is opaque to the client, and does not show up in the type. We write the type for an hotel as:

```
bologna :: [?Date.!Price.end]
```

The protocol with the broker is somewhat more complex, yet its interface with the client is exactly the same.

```
broker :: [?Date.!Price.end]
```

All values in a stream are required to be of the same type. The type of a process is a pair describing the conversation it engages into and the values it writes into its stream. Considering the part **stream** P **as** f **in** Q of the broker example, we have that P is of type (**end**, Price), meaning that P does not engage in any interaction with the client, and that it feeds Price values into the stream. On the other hand, Q is of type (!Price.**end**, T), since it communicates a price to the client (the type of the stream is arbitrary, given that Q does not feed into its stream).

A memory cell. Even if a memory cell is not a common scenario in SOC, stateful services are. Examples abound in the literature, from data-structures to weblog updates. Contrary to SCC [BBC+06], our language allows writing stateful services without exploiting service termination. Inspired by the encoding of objects in the π-calculus [PT95], we set up a simple, ephemeral, service to produce a value: buffer ⇒ v. The persistent service get below (persistent services are identified by the extra $*$ in the service definition) calls the buffer service to obtain its value (thus consuming the service provider), replies the value to the client, and replaces the buffer service.

```
get *⇒ call buffer >^1 v > (v | buffer ⇒ v)
```

The service set calls the buffer service (in order to consume the service provider), then gets the new value from the client and replaces the buffer with this value.

```
set *⇒ call buffer >^1 > (w)(buffer ⇒ w)
```

Note that above, in the operator $P{>^{1}}{>}Q$, the name of the parameter is not written: we omit it since it is never used in the continuation.

Finally, the cell service sets up three services—get, set, and buffer—sends the first two to the client, and keeps buffer locally with initial value 0.

```
cell *⇒ (νbuffer,get,set).get.set.
 (buffer ⇒ 0 |
  get *⇒ call buffer >^1 v >
      (v | buffer ⇒ v) |
  set *⇒ call buffer >^1 > (w)(buffer ⇒ w))
```

2.2. Syntax

The syntax for SSCC processes is inspired by that of the π-calculus, but it includes several additional features. A process may post/read a message to/from its environment; more complex processes may be built by joining two processes in parallel, restricting a name within a process or using recursion (guarded by prefixes).

Furthermore, a process may be a service definition a ⇒ P or service invocation a ⇐ P, containing the name a of the service and a protocol P. When a service invocation meets its corresponding definition, the corresponding protocols execute within a session whose name is freshly generated, and messages posted by one endpoint are read by the other endpoint of that session.

[1] Types and the type system are presented in detail in Sect. 2.4.

$$\begin{array}{llllll}
P, Q & ::= & & \textit{Processes} & & \\
 & \mid & \mathbf{0} & \text{Terminated process} & & \\
 & \mid & P|Q & \text{Parallel composition} & \mid (\nu\, a)P & \text{Name restriction} \\
 & \mid & X & \text{Process variable} & \mid \textsf{rec}\; X.P & \text{Recursive process definition} \\
 & \mid & a \Rightarrow P & \text{Service definition} & \mid a \Leftarrow P & \text{Service invocation} \\
 & \mid & v.P & \text{Value sending} & \mid (x)P & \text{Value reception} \\
 & \mid & \textsf{stream}\; P \;\textsf{as}\; f \;\textsf{in}\; Q & \text{Stream} & \mid \textsf{feed}\; v.P & \text{Feed the process' stream} \\
 & \mid & f(x).P & \text{Read from a stream} & & \\
 & & & & & \\
v & ::= & & \textit{Values} & & \\
 & & a & \text{Service name} & \mid \textsf{unit} & \text{Unit value}
\end{array}$$

Fig. 1. The syntax of SSCC

Service orchestration is achieved by means of stream composition. A stream is a one-way static channel whose name cannot be communicated; the process "inside" a stream can only post values to that stream, whereas the "outer" process can only read from that stream. It is important to point out that any process may (directly) write to exactly one stream, but might be able to read from several; hence the need to name the stream in the reading, but not in the writing, context. This choice simplifies the type system and the behavioural analysis, not compromising expressiveness, as far as we noticed with the (large amount of) examples developed. Intuitively, communication via streams corresponds to local interaction, whereas service interaction is meant for remote communication.

Notation. Processes are built using three kinds of identifiers: *service names*, *stream names*, and *process variables*. Service names are ranged over by a and b; values, ranged over by u and v, can be either service names or the **unit** value. Basic values such as integers and strings can be easily added, and will be used in examples. Names for values can also be used as variables (bound by value reception or read from stream), and we use x and y in this case. Stream names are ranged over by f and g. Process variables are ranged over by X and Y, and are used to define recursive processes. All process variables are required to be bound, i.e., they must appear only within a recursive process definition of which they are the recursion variable. Thus, we consider a process well-formed only if it contains no free process variables.

Definition 1 (Syntax). The grammar in Fig. 1 defines the *syntax of* SSCC *processes.*

The first five cases of the grammar introduce standard process calculi operators: the terminated process **0**, parallel composition P|Q, name restriction (notice that only service names can be restricted), and recursion (to define recursion we need both process variables X and recursive process definitions **rec** X.P).

We then have two constructs to *build services*: definition (or provider) and invocation (or client). They are both defined by their name a and protocol P. Service definition and service invocation are symmetric, differently from SCC [BBC+06]. *Service protocols* are built using value sending and receiving, allowing bidirectional communication between clients and servers.

Finally, there are the three constructs for *service orchestration*, which constitute the main novelty of our calculus. The stream construct declares a stream f for communication from P to Q. P can insert a value v into the stream f using **feed** v.P', and Q can read from there using f(x).Q'. Streams can be considered either ordered or unordered. An unordered stream is a multiset, while an ordered one is a queue. In most cases the difference is not important. Herein, streams are considered ordered, acting as queues. Note that stream communication is inherently asynchronous. We write w :: $\vec{v}$ for the stream obtained by adding w to stream $\vec{v}$, and $\vec{v}$ :: w for a stream from which w can be removed. In the latter case $\vec{v}$ is what we get after removing w. We denote the empty stream by $\langle\rangle$.

Observe that these processes do not yet contain sessions. Indeed, sessions can only arise as a result of the interaction of a service definition and a service invocation, which produces an active session. Therefore, an extended run-time syntax is needed. This syntax distinguishes a fourth set of identifiers: *session names*, ranged over by r and s. The letters n and m are used to range over both session names and service names. Moreover, at runtime, values in the stream are stored together with the stream definition: the static construct **stream** P **as** f **in** Q can be seen as an abbreviation of **stream** P **as** f=$\langle\rangle$ **in** Q.

Definition 2 (Runtime syntax). The grammar in Fig. 2 defines the *syntax of runtime* SSCC *processes.*

Henceforth, we only consider processes derived from the ones in the static syntax.

$$
\begin{array}{rlllll}
P, Q & ::= & \textit{Runtime processes} & & & \\
& \;\;\ldots & \text{as in Figure 1} & & & \\
& \mid\; r \rhd P & \text{Server session} & \mid & r \lhd P & \text{Client session} \\
& \mid\; (\nu\, r)P & \text{Session restriction} & \mid & \textbf{stream}\; P \;\textbf{as}\; f = \vec{v} \;\textbf{in}\; Q & \text{Stream with values}
\end{array}
$$

Fig. 2. The runtime syntax of SSCC

$$
\begin{array}{l}
P|\mathbf{0} \;\equiv\; P \qquad P|Q \;\equiv\; Q|P \qquad (P|Q)|R \;\equiv\; P|(Q|R) \qquad \text{(S-NIL, S-COMM, S-ASSOC)} \\
(\nu\, n)P|Q \;\equiv\; (\nu\, n)(P|Q) \quad \text{if } n \notin \text{fn}(Q) \qquad r \bowtie (\nu\, a)P \;\equiv\; (\nu\, a)(r \bowtie P) \\
\hfill \text{(S-EXTR-PAR, S-EXTR-SESS)} \\
\textbf{stream}\; (\nu\, a)P \;\textbf{as}\; f = \vec{v} \;\textbf{in}\; Q \;\equiv\; (\nu\, a)(\textbf{stream}\; P \;\textbf{as}\; f = \vec{v} \;\textbf{in}\; Q) \quad \text{if } a \notin \text{fn}(Q) \cup \text{Set}(\vec{v}) \quad \text{(S-EXTR-STREAML)} \\
\textbf{stream}\; P \;\textbf{as}\; f = \vec{v} \;\textbf{in}\; (\nu\, a)Q \;\equiv\; (\nu\, a)(\textbf{stream}\; P \;\textbf{as}\; f = \vec{v} \;\textbf{in}\; Q) \quad \text{if } a \notin \text{fn}(P) \cup \text{Set}(\vec{v}) \quad \text{(S-EXTR-STREAMR)} \\
\textbf{stream}\; P \;\textbf{as}\; f \;\textbf{in}\; Q \;\equiv\; \textbf{stream}\; P \;\textbf{as}\; f = \langle\rangle \;\textbf{in}\; Q \hfill \text{(S-RUN-TIME-STREAM)} \\
(\nu\, n)(\nu\, m)P \equiv (\nu\, m)(\nu\, n)P \qquad (\nu\, a)\mathbf{0} \;\equiv\; \mathbf{0} \qquad \textbf{rec}\; X.P \;\equiv\; P[{}^{\textbf{rec}\; X.P}/_{X}] \\
\hfill \text{(S-SWAP, S-COLLECT, S-REC)}
\end{array}
$$

Fig. 3. The structural congruence relation

2.3. Operational semantics

Since the aim of SSCC is to model systems that change over time, it is essential to describe how they evolve. This is done by giving an operational semantics, which is a reduction relation on runtime processes.

Notation. Some constructs act as binders. Name x is bound in $(x)P$ and in $f(x).P$; name n is bound in $(\nu\, n)P$; stream f is bound in $\textbf{stream}\; P \;\textbf{as}\; f = \vec{v} \;\textbf{in}\; Q$ with scope Q; and process variable X is bound in $\textbf{rec}\; X.P$. The sets of free and bound names in P are denoted, respectively, by $\text{fn}(P)$ and $\text{bn}(P)$. The set of names in P is $\text{n}(P) = \text{fn}(P) \cup \text{bn}(P)$. If $\vec{w} = w_1 \cdots w_n$ with $n \geq 0$, then $\text{Set}(\vec{w}) = \{w_1, \ldots, w_n\}$. We work up to α-conversion and follow Barendregt's variable convention, whereby all variables in binding occurrences in any mathematical context are pairwise distinct and distinct from the free variables [Bar84]. We use capture avoiding substitutions (thanks to Barendregt convention) to replace names for names, as in $[{}^{v}/_{x}]$, and processes for process variables, as in $[{}^{\textbf{rec}\; X.P}/_{X}]$.

Structural congruence. Before proceeding further, it is appropriate to introduce a structural congruence relation for SSCC processes. In the rules in Fig. 3 and in the following, in processes containing sessions, $r \bowtie P$ stands for either $r \lhd P$ or $r \rhd P$; multiple occurrences of $r \bowtie P$ within the same rule are instantiated in the same way (i.e., all $r \lhd P$ or all $r \rhd P$), while occurrences of $r \,\overline{\bowtie}\, P$ stand for the opposite instantiation (i.e., where $r \bowtie P$ stands for $r \lhd P$, $r \,\overline{\bowtie}\, P$ stands for $r \rhd P$, and vice-versa).

Definition 3 (Structural congruence relation). The *structural congruence relation* $\equiv$ *on runtime* SSCC *processes* is the smallest congruence closed under the rules in Fig. 3.

Reduction relation. Interactions can happen in different active contexts. Since our interactions are binary, we introduce also two-holed contexts, which we call double contexts.

Definition 4 (Active contexts). The grammar in Fig. 4 generates *active contexts* $\mathcal{C}\llbracket\rrbracket$ and *double contexts* $\mathcal{D}\llbracket,\rrbracket$.

$$
\begin{array}{rcl}
\mathcal{C}\llbracket\rrbracket & ::= & \bullet \;\mid\; \mathcal{C}\llbracket\rrbracket|Q \;\mid\; P|\mathcal{C}\llbracket\rrbracket \;\mid\; (\nu\, n)\mathcal{C}\llbracket\rrbracket \\
& & \mid\; \textbf{stream}\; \mathcal{C}\llbracket\rrbracket \;\textbf{as}\; f = \vec{v} \;\textbf{in}\; Q \;\mid\; \textbf{stream}\; P \;\textbf{as}\; f = \vec{v} \;\textbf{in}\; \mathcal{C}\llbracket\rrbracket \;\mid\; r \bowtie \mathcal{C}\llbracket\rrbracket \\
\mathcal{D}\llbracket,\rrbracket & ::= & \mathcal{C}'\llbracket\rrbracket|\mathcal{C}''\llbracket\rrbracket \;\mid\; \textbf{stream}\; \mathcal{C}'\llbracket\rrbracket \;\textbf{as}\; f = \vec{v} \;\textbf{in}\; \mathcal{C}''\llbracket\rrbracket
\end{array}
$$

Fig. 4. Active and double contexts

$$\frac{\mathcal{D}\llbracket\,,\,\rrbracket \text{ does not bind } r \text{ or } a \qquad r \notin \mathrm{fn}(P) \cup \mathrm{fn}(Q) \cup \mathrm{fn}(\mathcal{D}\llbracket\,,\,\rrbracket)}{\mathcal{D}\llbracket a \Rightarrow P, a \Leftarrow Q\rrbracket \rightarrow (\nu\, r)\mathcal{D}\llbracket r \rhd P, r \lhd Q\rrbracket} \qquad (\text{R-SYNC})$$

$$\frac{\begin{array}{c}\mathcal{D}\llbracket\,,\,\rrbracket, \mathcal{C}\llbracket\rrbracket, \text{ and } \mathcal{C}'\llbracket\rrbracket \text{ do not bind } r \text{ or } v\\ \mathcal{C}\llbracket\rrbracket \text{ and } \mathcal{C}'\llbracket\rrbracket \text{ do not contain sessions around the hole}\end{array}}{(\nu\, r)\mathcal{D}\llbracket r \bowtie \mathcal{C}\llbracket v.P\rrbracket, r \overline{\bowtie} \mathcal{C}'\llbracket (x)Q\rrbracket\rrbracket \rightarrow (\nu\, r)\mathcal{D}\llbracket r \bowtie \mathcal{C}\llbracket P\rrbracket, r \overline{\bowtie} \mathcal{C}'\llbracket Q[{}^{v}\!/\!_{x}]\rrbracket\rrbracket} \qquad (\text{R-COMM})$$

$$\frac{\mathcal{C}\llbracket\rrbracket \text{ does not bind } w \qquad \text{the hole does not occur in } \mathcal{C}\llbracket\rrbracket \text{ in the left part of a stream context}}{\textsf{stream}\, \mathcal{C}\llbracket \textsf{feed}\, w.P\rrbracket\, \textsf{as}\, f = \vec{v}\, \textsf{in}\, Q \rightarrow \textsf{stream}\, \mathcal{C}\llbracket P\rrbracket\, \textsf{as}\, f = w :: \vec{v}\, \textsf{in}\, Q} \qquad (\text{R-FEED})$$

$$\frac{\mathcal{C}\llbracket\rrbracket \text{ does not bind } w \text{ or } f}{\textsf{stream}\, P\, \textsf{as}\, f = \vec{v} :: w\, \textsf{in}\, \mathcal{C}\llbracket f(x).Q\rrbracket \rightarrow \textsf{stream}\, P\, \textsf{as}\, f = \vec{v}\, \textsf{in}\, \mathcal{C}\llbracket Q[{}^{w}\!/\!_{x}]\rrbracket} \qquad (\text{R-READ})$$

$$\frac{P \rightarrow P'}{\mathcal{C}\llbracket P\rrbracket \rightarrow \mathcal{C}\llbracket P'\rrbracket} \qquad \frac{Q \equiv P \rightarrow P' \equiv Q'}{Q \rightarrow Q'} \qquad (\text{R-CONG, R-STR})$$

Fig. 5. The reduction relation

Applying a double context to two processes P_1 and P_2 produces the process obtained by replacing the first (in the prefix visit of the syntax tree) hole • with P_1 and the second hole • with P_2.

Definition 5 **(Reduction semantics).** The rules in Fig. 5, together with symmetric rules of R-COMMand of R-SYNC (swapping the processes in the two holes of $\mathcal{D}\llbracket\,,\,\rrbracket$), inductively define the *reduction relation* on SSCC processes.

Rule R-SYNCallows a service invocation and a service definition to interact. This interaction produces a pair of complementary sessions, distinguished by a fresh restricted name r. Notice that both the service invocation and the service definition disappear. Rule R-COMMallows communication between corresponding sessions. Since value sending and receiving refer to the innermost session, we have to ensure that r is such a session. This is done by requiring that $\mathcal{C}\llbracket\rrbracket$ and $\mathcal{C}'\llbracket\rrbracket$ do not contain sessions around the hole. Then there are the two rules dealing with streams: rule R-FEEDputs a value in the stream while rule R-READtakes a value from the stream. In R-FEED, similarly to what happens in R-COMM, the right premise ensures that the value sent by the feed operation is not captured by an inner stream construct in context $\mathcal{C}\llbracket\rrbracket$. Finally rule R-CONGallows reductions to happen inside arbitrary active contexts, and rule R-STRexploits structural congruence.

2.4. Type system

This section deals with the static behaviour of SSCC processes. Type systems provide protocol information, abstracting from the actual data being sent, and are powerful enough to prove a number of interesting properties about termination.

Typing for SSCC processes is a way of providing information about their protocol, telling the world about the data they are expecting and how they communicate with the outside. This simple type system, built along the lines of session types [GH05, HVK98, YV06, THK94], is strong enough to ensure protocol compatibility among clients and servers. It is also able to deal with many different interacting services simultaneously.

Definition 6 **(Types).** The grammar in Fig. 6 defines the *syntax of types*.

Types are divided into three classes. Types for *values* T are either **Unit**, which denotes the only basic type,[2] or $[U]$, which is the type of a service or session with protocol U. This protocol is always seen from the server point of view, regardless of the role of the process being typed. Types for *streams* are of the form $\langle T\rangle$ where T is the type of the values the stream carries. Finally, types for *processes* are of the form (U, T) where U is the protocol that the process follows, and T is the type of the values the process feeds into its stream.

The **rec** operator for types is a binder, giving rise, in the standard way, to notions of bound and free variables and α-equivalence. As was the case for processes, α-convertible types are considered the same. Furthermore, this paper takes an *equi-recursive* view of types, not distinguishing between a type $\textsf{rec}\ \alpha.U$ and its unfolding $T[{}^{\textsf{rec}\ \alpha.U}\!/\!_{\alpha}]$. We are interested on *contractive* (not including subterms of the form $\textsf{rec}\,\alpha.\textsf{rec}\,\alpha_1 \ldots \textsf{rec}\,\alpha_n.\alpha$) types only; see [Pie02].

[2] To be possibly extended with, say, integers and strings.

T	$::=$	*Types*
	$\textsf{Unit}$	unit type
	$\mid\ [U]$	service type
U	$::=$	*Conversation types*
	$?T.U$	input
	$\mid\ !T.U$	output
	$\mid\ \textsf{end}$	end of conversation
	$\mid\ \alpha$	type variable
	$\mid\ \textsf{rec}\,\alpha.U$	recursive type

Fig. 6. The syntax of value types

$$\overline{?T.U} \triangleq !T.\overline{U} \qquad \overline{!T.U} \triangleq ?T.\overline{U} \qquad \overline{\textsf{end}} \triangleq \textsf{end} \qquad \overline{\alpha} \triangleq \alpha \qquad \overline{\textsf{rec}\ \alpha.U} \triangleq \textsf{rec}\ \alpha.\overline{U}$$

Fig. 7. The complement of a conversation type

In order to establish communication, two processes must have complementary protocols, in the sense that one is ready to send values of the type that the other is expecting, and vice-versa. This is captured via the operation of complementation on conversation types, defined in Fig. 7. Intuitively, if a client executes protocol U and a server implements protocol $\overline{U}$, then the conversation between them can proceed without errors.

Typing judgments are as follows:

$\Gamma \vdash P : (U, T)$	*Processes*
$\Gamma \vdash v : T$	*Values*

where Γ is a map with entries $a : T$, $r : T$, $f : \langle T \rangle$, and $X : (U, T)$.

Definition 7 (Type system). The rules in Fig. 8 inductively define the type system for SSCC processes.

The type of a process is an abstraction of its behaviour: its first component shows the protocol of the process, while the second component traces the type of the values fed to its stream. Notice that the properties of internal sessions and streams are guaranteed by the typing derivation and the typing assumptions in Γ, and they do not influence the type of the process. For instance, if the process is a session $\mathsf{r} \rhd \mathsf{P}$, then its protocol is $\textsf{end}$, but the protocol followed by P is traced by an assumption $r : [U]$ in Γ. When the complementary session is found, the compatibility check is performed.

These types force protocols to be sequential, which the authors consider to be a good programming style. Suppose for instance that the protocol contained two parallel outputs: then there should be two inputs in the complementary protocol, and we cannot know which output is matched with each input. Either this is irrelevant for the process, and the outputs can be sorted out in an arbitrary way, or it is relevant, and it should generate an error.

Another aspect is that parallel protocols are more complex to check for compatibility. Notice that this does not forbid having, e.g., two concurrent service invocations, since sequentiality is only enforced inside protocols.

SSCC equipped with this type system is type safe: typable processes are not errors, nor can they generate errors. The proof of this result requires, as usual, a type preservation property—subject reduction—and a definition of erroneous processes. The proofs of the theorems below can be found in Appendix A.

Theorem 1 (Subject reduction). *Let $\Gamma \vdash P : (U, T)$ and $P \to P'$. Then $\Gamma \vdash P' : (U, T)$.*

Theorem 2 (Type Safety). *Let P be a typable process. Then P has no subterm of the following forms.*

Protocol:

$\mathcal{D}[\![r \bowtie \mathcal{C}[\![v.P]\!], r \,\overline{\bowtie}\, \mathcal{C}'[\![u.Q]\!]]\!]$	Two outputs
$\mathcal{D}[\![r \bowtie \mathcal{C}[\![v.P]\!], r \,\overline{\bowtie}\, \mathbf{0}]\!]$	Output and $\mathbf{0}$
$\mathcal{D}[\![r \bowtie \mathcal{C}[\![(x)P]\!], r \,\overline{\bowtie}\, \mathcal{C}'[\![(y)Q]\!]]\!]$	Two inputs
$\mathcal{D}[\![r \bowtie \mathcal{C}[\![(x)P]\!], r \,\overline{\bowtie}\, \mathbf{0}]\!]$	Input and $\mathbf{0}$

where in all the cases $\mathcal{D}[\![,]\!]$ does not bind r, and $\mathcal{C}[\![]\!]$ and $\mathcal{C}'[\![]\!]$ do not contain sessions around the $\bullet$.

$$\Gamma, n\colon T \vdash n\colon T \qquad \Gamma, f\colon \langle T\rangle \vdash f\colon \langle T\rangle \qquad \Gamma \vdash \textsf{unit}\colon \textsf{Unit} \qquad (\text{T-NAME, T-SNAME, T-UNIT})$$

$$\frac{\Gamma \vdash P\colon (U,T) \quad \Gamma \vdash v\colon T'}{\Gamma \vdash v.P\colon (!T'.U,T)} \qquad \frac{\Gamma, x\colon T' \vdash P\colon (U,T)}{\Gamma \vdash (x)P\colon (?T'.U,T)} \qquad (\text{T-SEND, T-RECEIVE})$$

$$\frac{\Gamma \vdash P\colon (U,T) \quad \Gamma \vdash a\colon [U]}{\Gamma \vdash a \Rightarrow P\colon (\textsf{end},T)} \qquad \frac{\Gamma \vdash P\colon (U,T) \quad \Gamma \vdash a\colon [\overline{U}]}{\Gamma \vdash a \Leftarrow P\colon (\textsf{end},T)} \qquad (\text{T-DEF, T-CALL})$$

$$\frac{\Gamma \vdash P\colon (U,T) \quad \Gamma \vdash r\colon [U]}{\Gamma \vdash r \rhd P\colon (\textsf{end},T)} \qquad \frac{\Gamma \vdash P\colon (U,T) \quad \Gamma \vdash r\colon [\overline{U}]}{\Gamma \vdash r \lhd P\colon (\textsf{end},T)} \qquad (\text{T-SESS-S, T-SESS-C})$$

$$\frac{\Gamma \vdash P\colon (U,T) \quad \Gamma \vdash v\colon T}{\Gamma \vdash \textsf{feed}\, v.P\colon (U,T)} \qquad \frac{\Gamma, x\colon T \vdash P\colon (U,T') \quad \Gamma \vdash f\colon \langle T\rangle}{\Gamma \vdash f(x).P\colon (U,T')} \qquad (\text{T-FEED, T-READ})$$

$$\frac{\Gamma \vdash P\colon (U,T) \quad \Gamma \vdash Q\colon (\textsf{end},T)}{\Gamma \vdash P|Q\colon (U,T)} \qquad \frac{\Gamma \vdash P\colon (\textsf{end},T) \quad \Gamma \vdash Q\colon (U,T)}{\Gamma \vdash P|Q\colon (U,T)} \qquad (\text{T-PAR-L, T-PAR-R})$$

$$\frac{\Gamma \vdash P\colon (U,T) \quad \Gamma, f\colon \langle T\rangle \vdash Q\colon (\textsf{end},T') \quad w \in \mathrm{Set}(\vec{v}) \Rightarrow \Gamma \vdash w\colon T}{\Gamma \vdash \textsf{stream}\, P\, \textsf{as}\, f = \vec{v}\, \textsf{in}\, Q\colon (U,T')} \qquad (\text{T-STREAM-L})$$

$$\frac{\Gamma \vdash P\colon (\textsf{end},T) \quad \Gamma, f\colon \langle T\rangle \vdash Q\colon (U,T') \quad w \in \mathrm{Set}(\vec{v}) \Rightarrow \Gamma \vdash w\colon T}{\Gamma \vdash \textsf{stream}\, P\, \textsf{as}\, f = \vec{v}\, \textsf{in}\, Q\colon (U,T')} \qquad (\text{T-STREAM-R})$$

$$\frac{\Gamma, X\colon (U,T) \vdash P\colon (U,T)}{\Gamma \vdash \textsf{rec}\, X.P\colon (U,T)} \qquad \frac{\Gamma, n\colon T_1 \vdash P\colon (U,T)}{\Gamma \vdash (\nu\, n)P\colon (U,T)} \qquad (\text{T-REC, T-RES})$$

$$\Gamma, X\colon (U,T) \vdash X\colon (U,T) \qquad \Gamma \vdash \mathbf{0}\colon (\textsf{end},T) \qquad (\text{T-VAR, T-NIL})$$

Fig. 8. The type system

Sequentiality:

$\mathcal{D}[\![v.P, u.Q]\!]$	Parallel outputs
$\mathcal{D}[\![(x)P, u.Q]\!]$	Parallel input and output
$\mathcal{D}[\![v.P, (y)Q]\!]$	Parallel output and input
$\mathcal{D}[\![(x)P, (y)Q]\!]$	Parallel inputs

where in all cases $\mathcal{D}[\![,]\!]$ does not contain sessions around the •.

An example of *protocol failure* is illustrated by process r ▷ v.P | r ◁ nil. This process cannot be typed, since the two parallel components require different assumptions for r, namely r: [!T.U'] where T is the type of v, and r: **end**, respectively. Similarly, a *non-sequential conversation* is r ▷ (v.P | u.Q), which also cannot be typed, since both v.P and u.Q have non-**end** protocols, thus forbidding the application of rules for parallel composition. Techniques used for session types can be adapted to *type check* SSCC processes [VGR06].

As an example we show the typing judgment for the protocol of the memory cell from Sect. 2.1. Services in SSCC are ephemeral: they do not survive invocation. Recursion can be used to provide persistent services: a service a ⇒ P can be made persistent by writing instead **rec** X.a ⇒ (P|X), which we abbreviate as a ∗⇒ P (see Fig. 9).

```
get :: [!Int.end]
get ∗⇒ call buffer >^1 v > (v | buffer ⇒ v)
```

The type captures the fact that the get service publishes the integer stored in the memory cell.

```
set :: [?Int.end]
set ∗⇒ call buffer >^1 > (w)(buffer ⇒ w)
```

The type indicates that a value is sent to the service (it will be used for updating the memory cell).

```
cell :: [![!Int.end].![?Int.end].end]
cell ∗⇒ (ν buffer, get,set).get.set.(buffer ⇒ 0 |
  get ∗⇒ call buffer >^1 v > (v | buffer ⇒ v) |
  set ∗⇒ call buffer >^1 > (w)(buffer ⇒ w))
```

The type for the Cell service is more interesting. It makes apparent the fact that, upon instantiation, the Cell service will publish two services, the first able to send a value (for getting the value of the memory cell) and the second able to receive a value (for setting the memory cell).

$$
\begin{array}{rcl}
\mathbf{call}\ \mathsf{a}(\mathsf{x}_1\ ,\ldots,\ \mathsf{x}_n) & \triangleq & \mathsf{a} \Leftarrow \mathsf{x}_1\ldots\mathsf{x}_n.(\mathsf{y})\ \mathbf{feed}\ \mathsf{y} \\
\mathsf{P} >^{\mathbf{n}} \mathsf{x}_1\ldots\ \mathsf{x}_n > \mathsf{Q} & \triangleq & \mathbf{stream}\ \mathsf{P}\ \mathbf{as}\ \mathsf{f}\ \mathbf{in}\ \mathsf{f}(\mathsf{x}_1).\ \ \ldots\ \ \mathsf{f}(\mathsf{x}_n).\mathsf{Q} \\
\mathsf{P} > \mathsf{x} > \mathsf{Q} & \triangleq & \mathbf{stream}\ \mathsf{P}\ \mathbf{as}\ \mathsf{f}\ \mathbf{in}\ \mathbf{rec}\ \mathsf{X}.\mathsf{f}(\mathsf{x}).(\mathsf{Q} \mid \mathsf{X}) \\
\mathsf{a} *\!\!\Rightarrow \mathsf{P} & \triangleq & \mathbf{rec}\ \mathsf{X}.\ \mathsf{a} \Rightarrow (\mathsf{P} \mid \mathsf{X}) \\
\mathbf{if}\ \ \mathsf{b}\ \mathbf{then}\ \mathsf{P} & \triangleq & \mathsf{b} \Leftarrow (\mathsf{x})(\mathsf{y})\ \mathsf{x} \Leftarrow \mathbf{feed\ unit} >^{\mathbf{1}} > \mathsf{P} \\
\mathbf{if}\ \ \neg\mathsf{b}\ \mathbf{then}\ \mathsf{P} & \triangleq & \mathsf{b} \Leftarrow (\mathsf{x})(\mathsf{y})\ \mathsf{y} \Leftarrow \mathbf{feed\ unit} >^{\mathbf{1}} > \mathsf{P} \\
\mathbf{if}\ \ \mathsf{b}\ \mathbf{then}\ \mathsf{P}\ \mathbf{else}\ \mathsf{Q} & \triangleq & \mathbf{if}\ \ \mathsf{b}\ \mathbf{then}\ \mathsf{P} \mid\ \mathbf{if}\ \ \neg\mathsf{b}\ \mathbf{then}\ \mathsf{Q} \\
\mathsf{T}_1 \rightarrow \ldots\ \rightarrow \mathsf{T}_n \rightarrow \mathsf{T} & \triangleq & [?\mathsf{T}_1\ldots?\mathsf{T}_n.!\mathsf{T}.\mathbf{end}] \\
\varepsilon\ \rightarrow \mathsf{T} & \triangleq & [!\,\mathsf{T}.\mathbf{end}] \\
\mathsf{Bool} & \triangleq & [![\ \mathbf{end}\,].![\mathbf{end}].\mathbf{end}]
\end{array}
$$

Assume omnipresent the following services,

```
tt  *⇒ 0
true  *⇒ tt.ff.0
false  *⇒ ff.tt.0
```

and omniabsent the service for ff.

Fig. 9. Derived constructs

3. Specifying in SSCC

This section explores examples that highlight the versatility of SSCC, while comparing to solutions written in SCC and Orc. These two languages are briefly reviewed in Sect. 1.4. For modelling we use a few useful abbreviations, which are gathered in Fig. 9. We will continue to use them in the rest of the paper.

The first example shows that naming streams can be handy. *Fork-join* is a pattern that spawns two threads, and resumes computation after receiving a value from each thread. In the example below, services a and b are run in parallel; **call** a feeds the first result produced by the service into stream f, and similarly for **call** b and stream g.

```
fork-and-join :: [?[!T1.end].?[!T2.end].!T1.!T2.end]
fork-and-join *⇒ (a)(b)(
  stream call a as f in
    stream call b as g in
      f(x).g(y).x.y)
```

The example is inspired in Orc [MC07, KCM06], but Orc, when one of the service invocations at a or b completes, kills the other. Instead, we let them run to completion. Orc is not able to match our semantics: reading a single value from an expression can only be performed via the **where** construct, and that necessarily means terminating the evaluation of the expression. We feel that termination should be distinct from normal orchestration; we leave for further work termination (and the corresponding compensation). Notice, however, that the declared type makes sure that services a and b produce each a single value.

The second example describes an idiom where for each value x produced by a process P, a second process Q is started. If process P produces its values by feeding into its stream, then, in the process below, a new copy of process Q is spawned for each value read from the stream. Process

```
  stream P as f in rec X.f(x).(Q | X)
```

can be abbreviated to P >x >Q (x can be dropped if it does not occur in Q), so that a service that reads news from sites CNN and BBC and e-mails each to a given address can be written as:

```
  email-news :: [?Address.end]
  email-news *⇒ (a)((call CNN | call BBC) > x > email ⇐ a.x)
```

The example and the short syntax is again from Orc. In this case we are faithful to the Orc semantics.

3.1. Workflow patterns

In this section we illustrate the expressiveness of SSCC by implementing the Workflow Patterns (WP) from van der Aalst et al. [vdAtHKB03]. This permits contrasting again our approach with that of SCC and of Orc [CPM06], which have similar aims. While WPs are an interesting benchmark, they are aimed at workflow description languages, not at calculi for SOC. For this reason some of the patterns are not meaningful (e.g., WP11: Implicit Termination) in our context, while others are redundant (e.g., WP12: Multiple Instances without Synchronization is similar to WP2: Parallel Split, since process calculi can obviously handle multiple instances). Also, some patterns require the ability to kill processes, which has not yet been introduced in SSCC, and thus are out of our possibilities. On the contrary WPs consider only "activities", i.e., services that receive one value and give back one result, while our calculus can also model more complex protocols.

We describe all patterns (in reference [vdAtHKB03]) as services; we also present their types. Those that have multiple entry points (the various merges, for example) are modelled with a vector of boolean values, describing which services should be invoked.

For us, an *activity* is a service that writes at most one value on the client side (replies at most one value). The simplest activity is the null service.

$$
\begin{array}{l}
\texttt{nullService} :: \varepsilon \to \textbf{Unit} \\
\texttt{nullService} \ast\!\!\Rightarrow \textbf{unit}
\end{array}
$$

Most of the patterns below allow definitions in SSCC that do not directly use either the stream operations (**stream**, **feed**, and $\texttt{f(x).P}$) or recursion.

In what follows we give a brief description of each workflow pattern and present an illustrative example, both taken from [vdAtHKB03]. To allow a comparison we also show how the patterns can be implemented in SCC. Services in SCC have always one parameter: we exploit it as first input for the server if the server protocol should start with an input, otherwise we assume it is unused and use **unit** as invocation value. We concentrate on WP1, WP3, WP4, WP9, WP10, and WP17, the patterns we find more significant and which better highlight the expressive power of SSCC. The reader should refer to Appendix B for the modelling of all the workflow patterns from [vdAtHKB03].

WP1: Sequence

"An activity in a workflow process is enabled after the completion of another activity in the same process. Example: an insurance claim is evaluated after the client's file is retrieved."

$$
\begin{array}{l}
\texttt{seq} :: (\varepsilon \to \texttt{T}_1) \to (\varepsilon \to \texttt{T}_2) \to \texttt{T}_2 \\
\texttt{seq} \ast\!\!\Rightarrow (\texttt{a}_1)(\texttt{a}_2)\,\textbf{call}\ \texttt{a}_1 >^{\mathbf{1}}>\ \textbf{call}\ \texttt{a}_2 >^{\mathbf{1}} \texttt{x} > \texttt{x}
\end{array}
$$

In Orc the implementation is similar. In SCC the most direct implementation is:

$$\texttt{seq} \Rightarrow (\texttt{a}_1)(\texttt{a}_2)\ \texttt{a}_2 \Leftarrow \texttt{a}_1 \Leftarrow \textbf{unit}$$

This implementation is fine for activities (actually here $\texttt{a}_2$ is invoked with the value from $\texttt{a}_1$ rather than with **unit**: this problem is solved in SSCC), but if $\texttt{a}_1$ is not an activity then $\texttt{a}_2$ is called for each value returned by $\texttt{a}_1$, and this is not the expected semantics. In SSCC this cannot happen, since the remaining values returned by $\texttt{a}_1$ stay forever in the stream. To enforce correct behaviour in SCC one should write:[3]

$$
\begin{array}{l}
\texttt{seq} \Rightarrow (\texttt{a}_1)(\texttt{a}_2)(\nu\texttt{r})(\texttt{r} \rhd (\texttt{a}_1 \Leftarrow \textbf{unit}\ \mid\ (\texttt{res})\ \textbf{return}\ \texttt{res})\ \mid \\
\qquad\qquad\qquad\qquad\quad \bar{\texttt{r}} \rhd (\texttt{v})\ \texttt{a}_2 \Leftarrow \textbf{unit})
\end{array}
$$

Notice that the SCC encoding above uses a conversation (a process of the form $\texttt{r} \rhd \texttt{P} \mid \bar{\texttt{r}} \lhd \texttt{Q}$), which we view as runtime syntax in SSCC. However, sessions can be also avoided in SCC using "fake" service invocations and definitions (nevertheless, service definitions stay there afterward, since they are persistent).

[3] This can be done also by type checking the protocol for $\texttt{a}_1$: this feature is not yet available in SCC, but it can be easily transferred there.

WP3: Synchronization

"A point in the workflow process where multiple parallel subprocesses/activities converge into one single thread of control, thus synchronizing multiple threads. It is an assumption of this pattern that each incoming branch of a synchronizer is executed only once. Example: insurance claims are evaluated after the policy has been checked and the actual damage has been assessed."

$$\begin{array}{l}
\mathsf{sync}\ ::\ (\varepsilon \to \mathsf{T}) \to\ \ldots \to (\varepsilon \to \mathsf{T}) \to \mathbf{Unit}\\
\mathsf{sync}\ {*}{\Rightarrow}\ (\mathsf{a}_1)\ldots(\mathsf{a}_n)(\mathbf{call}\ \mathsf{a}_1 \mid \ldots \mid \mathbf{call}\ \mathsf{a}_n) >^{\mathbf{n}}> \mathbf{unit}
\end{array}$$

Orc uses the **where** operator, and SCC uses sessions (or "fake services"):

$$\begin{array}{l}
\mathsf{sync} \Rightarrow (\mathsf{a}_1)\ldots(\mathsf{a}_n)(\nu \mathsf{r})(\mathsf{r} \rhd (\mathsf{a}_1 \Leftarrow \mathbf{unit} \mid \ldots \mid \mathsf{a}_n \Leftarrow \mathbf{unit}) \mid\\
\qquad\qquad\qquad\qquad\quad \bar{\mathsf{r}} \rhd (\mathsf{x}_1)\ldots(\mathsf{x}_n)\ \mathbf{return}\ \mathbf{unit})
\end{array}$$

WP4: Exclusive choice

"A point in the workflow process where, based on a decision or workflow control data, one of several branches is chosen. Example: based on the workload, a processed tax declaration is either checked using a simple administrative procedure or is thoroughly evaluated by a senior employee."

$$\begin{array}{l}
\mathsf{xor}\ ::\ \mathsf{Bool} \to (\varepsilon \to \mathsf{T}) \to (\varepsilon \to \mathsf{T}) \to \mathsf{T}\\
\mathsf{xor}\ {*}{\Rightarrow}\ (\mathsf{b})(\mathsf{a}_1)(\mathsf{a}_2)\ \mathbf{if}\ \mathsf{b}\ \mathbf{then}\ \mathbf{call}\ \mathsf{a}_1 >^{\mathbf{1}} \mathsf{x} > \mathsf{x}\ \mathbf{else}\ \mathbf{call}\ \mathsf{a}_2 >^{\mathbf{1}} \mathsf{x} > \mathsf{x}
\end{array}$$

Notice that **if**−**then**−**else** cannot be typed with the current type system unless both branches have the empty protocol (since they occur in parallel). This problem can be solved by adding a dedicated typing rule to exploit the knowledge that only one of the branches is actually executed.

In SCC one can implement true and false in a similar way. Then we have:

$$\begin{array}{l}
\mathbf{if}\ \mathsf{b}\ \mathbf{then}\ \mathsf{P} = (\nu \mathsf{s})\ \mathsf{s} \rhd \mathsf{b}\ \{(-)\ (\mathsf{x})(\mathsf{y})\ \mathbf{return}\ \mathsf{x}\} \Leftarrow \mathbf{unit}\\
\qquad\qquad\qquad \mid \bar{\mathsf{s}} \rhd (\mathsf{x}_1)\ (\nu \mathsf{r})\ \mathsf{r} \rhd \mathsf{x}_1\ \{(-)\mathbf{return}\ \mathbf{unit}\} \Leftarrow \mathbf{unit}\\
\qquad\qquad\qquad\qquad\qquad\qquad\quad \mid \bar{\mathsf{r}} \rhd (\mathsf{z})\ \mathsf{P}\\
\mathbf{if}\ \neg\mathsf{b}\ \mathbf{then}\ \mathsf{P} = (\nu \mathsf{s})\ \mathsf{s} \rhd \mathsf{b}\ \{(-)\ (\mathsf{x})(\mathsf{y})\ \mathbf{return}\ \mathsf{y}\} \Leftarrow \mathbf{unit}\\
\qquad\qquad\qquad \mid \bar{\mathsf{s}} \rhd (\mathsf{x}_1)\ (\nu \mathsf{r})\ \mathsf{r} \rhd \mathsf{x}_1\ \{(-)\mathbf{return}\ \mathbf{unit}\} \Leftarrow \mathbf{unit}\\
\qquad\qquad\qquad\qquad\qquad\qquad\quad \mid \bar{\mathsf{r}} \rhd (\mathsf{z})\ \mathsf{P}
\end{array}$$

Notice that while in SSCC feeds from P are not intercepted by the if context, in SCC the returns are lost, since P is executed inside a subsession. To forward the results to the caller, extra programming effort is required. Actually, since P cannot be executed at top level (since in order to start it when a trigger coming from a subsession is received, the trigger should be transmitted using a return, which either goes to the other side, or executes P inside a session) a forward of values is needed, but in SCC one is able to specify only a finite amount of forwarding. Thus, if P can give back an unbounded number of replies, this cannot be programmed. Nevertheless, in the following example we always suppose to have the **if** with forwarding of the results. Since workflow patterns deal only with activities (one result), then the forwarding can be implemented.

The **if**−**then**−**else** is as in SSCC.

$$\mathsf{xor} \Rightarrow (\mathsf{b})(\mathsf{a}_1)(\mathsf{a}_2)\ \mathbf{if}\ \mathsf{b}\ \mathbf{then}\ \mathsf{a}_1 \Leftarrow \mathbf{unit}\ \mathbf{else}\ \mathsf{a}_2 \Leftarrow \mathbf{unit}$$

Because of the above observation the xor in SCC (with the above implementation of **if**) gives back no value.

WP9: Discriminator

"The discriminator is a point in a workflow process that waits for one of the incoming branches to complete before activating the subsequent activity. From that moment on it waits for all remaining branches to complete and "ignores" them. Once all incoming branches have been triggered, it resets itself so that it can be triggered again (which is important, otherwise it could not really be used in the context of a loop). Example: to improve query response time, a complex search is sent to two different databases over the Internet. The first one that comes up with the result should proceed the flow. The second result is ignored."

$$\begin{array}{l}
\mathsf{discriminator}\ ::\ (\varepsilon \to \mathsf{T}) \to\ldots\to (\varepsilon \to \mathsf{T}) \to \mathbf{Unit}\\
\mathbf{rec}\ \mathsf{X}.\mathsf{discriminator} \Rightarrow (\mathsf{a}_1)\ldots(\mathsf{a}_n)\\
\qquad\qquad\qquad\qquad\quad \mathbf{stream}\ \mathbf{call}\ \mathsf{a}_1 \mid \ldots \mid \mathbf{call}\ \mathsf{a}_n\ \mathbf{as}\ \mathsf{f}\ \mathbf{in}\\
\qquad\qquad\qquad\qquad\quad \mathsf{f}(\mathsf{x}_1).\mathbf{unit}.\mathsf{f}(\mathsf{x}_2)\ldots\mathsf{f}(\mathsf{x}_n).\mathsf{X}
\end{array}$$

In SCC, we cannot control the point where the service discriminator becomes available again.

```
discriminator ⇒ (a_1)...(a_n)
                         (νr)  r ▷ a_1 ⇐ unit  |  ...  |a_n ⇐ unit
                               r̄ ▷ (x_1).return  unit.(x_2)...(x_n)
```

Here, the Orc implementation supposes the existence of a basic site S, with methods put and get, acting as a buffer. This site cannot be described in Orc (Orc does not deal with site programming). We think that sites should deal only with computation, while all the coordination should be done at the coordination language level. This implementation fails to satisfy this separation of concerns principle. We are not aware of better implementations in Orc.

WP10: Arbitrary cycles

"A point in a workflow process where one or more activities can be done repeatedly." Arbitrary cycles can be obtained via mutual invocations among services. We show here how an example of a structured cycle can be programmed:

```
while (service c returns true) do {call service a}.
while :: (ε → Bool) → (ε → T) → Unit
while *⇒ (c)(a) call c >1 b >
                     IfSignal(b,call a >1> call while (c,a)) >1 x > x
```

where

```
IfSignal(b,P) = if b then P else unit
```

The pattern is programmed as in Orc.

WP17: Interleaved parallel routing

"A set of activities is executed in an arbitrary order: each activity in the set is executed, the order is decided at run-time, and no two activities are executed at the same moment (i.e., no two activities are active for the same workflow instance at the same time). Example: the Navy requires every job applicant to take two tests: *physical_test* and *mental_test*. These tests can be conducted in any order, but not at the same time."

We assume that each service (a_1 to a_n) signals termination by sending a value, as witnessed by their types. Contrary to Orc, SSCC is expressive enough to describe the pattern within the language.

```
interleave:: [?[!T_1.end]...?[!T_n.end].end]
interleave ⇒ (a_1)...(a_n)(ν back)(
  stream
    back *⇒ (x)feed x
  as lock in
    back ⇐ unit |
    lock(_).a_1 ⇐ (x)(back ⇐ unit) | ... |
    lock(_).a_n ⇐ (x)(back ⇐ unit))
```

Essentially, the different activities execute in parallel, and a lock ensures that two of them are never enabled together. Note the use of auxiliary service back to relay values from the right to the left part of a stream construct, where they are fed into the stream.

In principle, the implementation in SCC can follow the same idea, replacing the stream with a session (possibly obtained using fake services, cf. WP1 above), and the **feed** x with a **return** x. However, the results produced by services a_1 ,..., a_n would naturally go into the same session, and should be forwarded to the top-level. Orc, here, exploits a basic site M implementing a lock with methods acquire and release. This site cannot be programmed inside Orc (see discussion in WP9).

3.2. The automotive scenario

In this section we show how two case studies from an automotive scenario can be represented in SSCC in a satisfactory way. The scenario and the case studies were developed in the EU project Sensoria [pro] and the project deliverable D8.0 [GtBB+06] presents them in detail. The first case study, a sight service that dynamically shows sights according to the driver's preferences, is straightforward to model. The second case study, a dinner

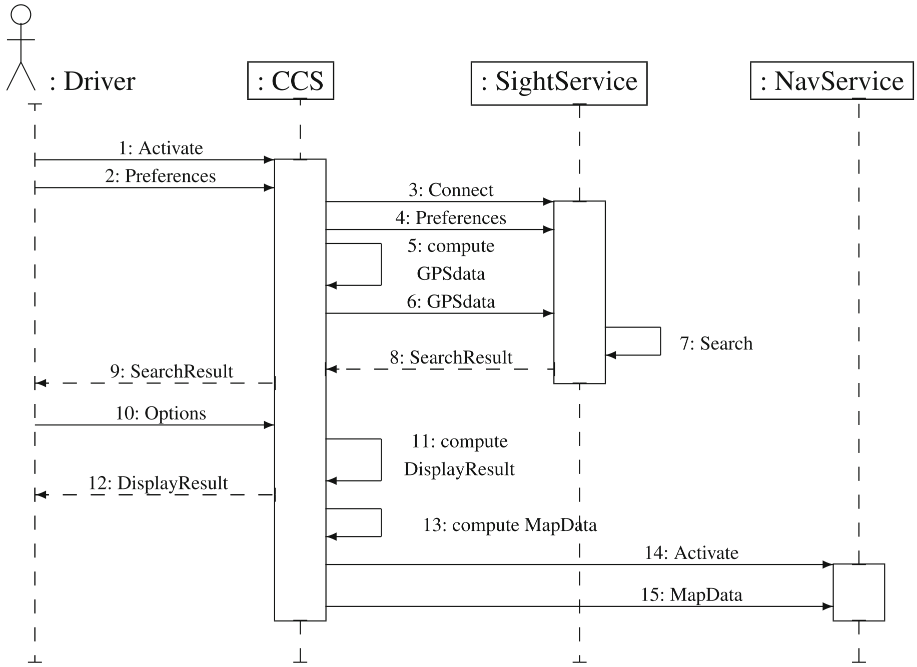

Fig. 10. SD for sight service scenario

service that allows the driver to make a reservation at a restaurant with some interaction, raises some problems with communication and typing of the resulting processes, which we then show that can be solved without compromising the motivation behind SSCC.

Road sights scenario. Here is the scenario as described in [GtBB+06].

The driver has subscribed to the dynamic sight service offered by the car company. The vehicle's GPS coordinates are automatically sent to the dynamic sight server at regular intervals, so the vehicle's location is known within a specified radius. Based on the driver's preferences that were given at the beginning of the trip, the dynamic sight server searches a sightseeing database for appropriate sights and displays them on the in-car map of the vehicle's navigation system. The driver clicks on sights he would like to visit, which results in the display of more detailed information about this specific sight (e.g., opening times, guidance to parking, etc.).

As suggested by Banci et al. [BFF+06], there are four actors in this scenario. The driver enables the sight service and sets his preferences via the Car Communication System (CCS). The CCS manages the communication to, and from, the sight service. The sight service has access to a sightseeing database where it gathers information from. Finally, the vehicle's navigation system displays the results to the driver in a graphical way.

UML Sequence Diagrams [Amb04] describe the exchange of messages among the components in a complex system; we use a variant of these diagrams (henceforth referred as SD). The dialogue between the actors is represented in Fig. 10. Service CCS can be implemented by the following process in SSCC. The notation ⟨action⟩ stands for an internal action of the system; the numbers on the right correspond to the numbers of the actions in the sequence diagram.

```
CCS ⇒ (Preferences).                           // 1:,2:
      (SightService ⇐ Preferences.             // 3:,4:
                      ⟨compute GPSdata⟩.       // 5:
                      GPSdata.                 // 6:
                      (SearchResult).          // 8:
                      feed SearchResult)
      >$^{1}$ SearchResult >
      SearchResult.                            // 9:
      (Options).                               // 10:
      ⟨compute DisplayResult⟩.                 // 11:
      DisplayResult.                           // 12:
      ⟨compute MapData⟩.                       // 13:
      (NavSystem ⇐ MapData)                    // 14:,15:
```

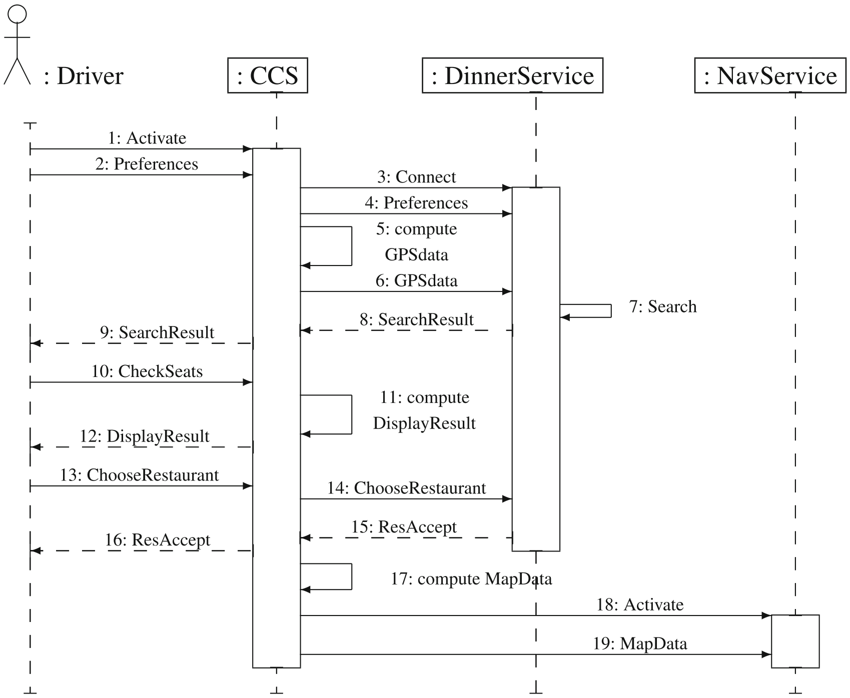

Fig. 11. SD for dinner service scenario

This implementation follows the SD diagram in Fig. 10 closely. It is easy to prove that

$\Gamma \vdash \mathsf{CCS} \Rightarrow \ldots : (\mathbf{end}, \mathsf{T})$

for any T (since the process performs no feed to its stream) whenever Γ is such that

$\Gamma \vdash \mathsf{CCS}: [?\mathsf{Preferences}.!\mathsf{SearchResult}.?\mathsf{Options}.!\mathsf{DisplayResult}.\mathbf{end}]$

$\Gamma \vdash \mathsf{SightService}: [?\mathsf{Preferences}.?\mathsf{GPSdata}.!\mathsf{SearchResult}.\mathbf{end}]$

$\Gamma \vdash \mathsf{NavSystem}: [?\mathsf{MapData}.\mathbf{end}]$

The (anonymous) stream has type $\langle \mathsf{SearchResult} \rangle$.

Observe that the types of all these services (CCS, SightService, and NavSystem) correspond precisely to their part of the conversation in the sequence diagram presented above from the point of view of the party who invokes them. For example, the type of SightService composes the arrows labelled 4:, 6:, and 8:, which is the trace of its conversation with CCS.

Instead, the communication along the stream does not correspond to any action in the SD. As we will discuss in more detail in the next section, it corresponds to the communication between two sessions within the CCS (the session CCS–SightService and the session CCS–Driver), and not to communication between different parties.

Dinner service. Once again we present the scenario as described in [GtBB+06].

Paul is very hungry, since he is driving without any food for 5 hours, so he activates the dinner service and enters a pizzeria as desired restaurant type and a price range between five and ten euros per meal. The navigation system displays a collection of nearby restaurants that match the preferred settings. Paul chooses the option to check for available seats in the participating local restaurants, and as a result the map displays only restaurants with available tables. Paul chooses "Tony's Pizza" and gets his reservation acknowledged. The way to the restaurant's parking lot is now displayed on the navigation system map.

Again, following the description in Banci et al. [BFF+06], we identify four actors in this scenario. The driver enables the dinner service and sets his preferences via the Car Communication System (CCS). The CCS manages the communication to, and from, the dinner service. The dinner service has access to a database of restaurants that it can contact in order to acknowledge the reservation. Finally, the vehicle's navigation system displays the results to the driver in a graphical way.

The dialogue between the actors is represented by the SD in Fig. 11, which is again essentially as in [BFF+06]. However, it is not so straightforward as before to implement the CCS in SSCC. The problem arises from the need to interact with the driver *after* receiving information from the dinner service, and then give feedback to the latter. We present some alternatives and discuss the drawbacks of each of them.

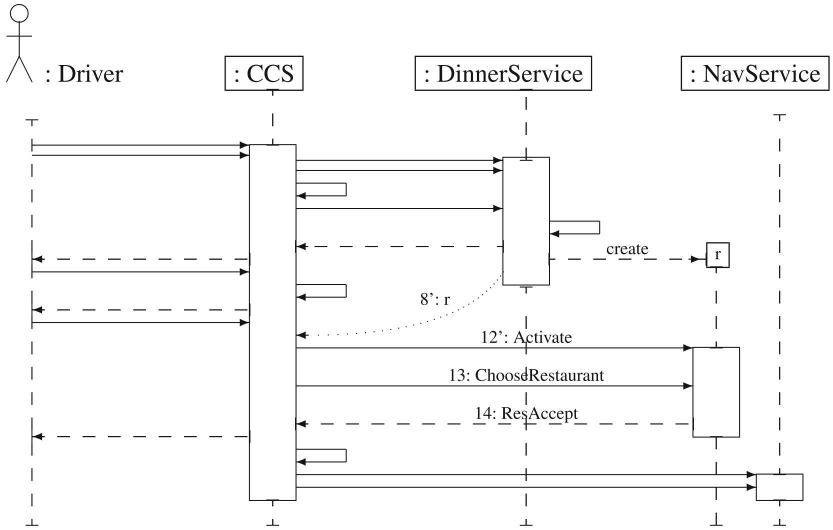

Fig. 12. SD for the dinner service scenario with creation of a continuation. Only the continuation part is detailed (cf. Fig. 11)

Implementation using a continuation. In this approach, the dinner service, when invoked for the first time, creates a (customized) new service whose (unique) name is sent back to the CCS. Afterwards, the CCS invokes *that* new service, which contains some persistent information. This solution deviates slightly from the sequence diagram above, since the session between the CCS and the dinner service is actually split into two sessions, one between the CCS and the dinner service (actions 3: to 8:), another between the CCS and the new service (actions 14: and 15:); see Fig. 12.

```
CCS ⇒ (Preferences).                              // 1:,2:
      stream (DinnerService ⇐ Preferences.        // 3:,4:
                              ⟨compute GPSdata⟩.  // 5:
                              GPSdata.            // 6:
                              (SearchResult).     // 8:
                              feed SearchResult.
                              (NewService).       // 8':
                              feed NewService)
       as f in (f(SearchResult).
                SearchResult.                     // 9:
                (CheckSeats).                     // 10:
                ⟨compute DisplayResult⟩           // 11:
                DisplayResult.                    // 12:
                (ChooseRestaurant).               // 13:
                f(NewService).
                (NewService ⇐ ChooseRestaurant.   // 13':,14:
                              (ResAccept).        // 15:
                              feed ResAccept)
                >^1 ResAccept >
                ResAccept.                        // 16:
                ⟨compute MapData⟩.                // 17:
                (NavSystem ⇐ MapData)             // 18:,19:
      )
```

This approach is interesting because it shows how continuations can be easily passed as (specialized) services, yielding some form of persistency. However, the CCS as given will not be typable because stream f cannot be typed consistently (since it is used twice to pass two bits of information of different types).

Implementation using an auxiliary service. An alternative approach is to feed new data from the communication system into the dinner service. However, this is not immediately possible using the syntax of SSCC. In order to achieve this "backward" communication, a new (linear) service b is created (whose name is private to the communication system). Whenever the communication system needs to send more data to the dinner service, it does so via b. Notice that one service is created for each message that needs to be sent back.

This solution follows the original SD (Fig. 11) faithfully.

```
CCS ⇒ (Preferences).                                        // 1:,2:
      (νb)(stream (DinnerService ⇐ Preferences.             // 3:,4:
                                   ⟨compute GPSdata⟩.       // 5:
                                   GPSdata.                 // 6:
                                   (SearchResult).          // 8:
                                   feed SearchResult.
                                   b ⇓ (ChooseRestaurant).
                                   ChooseRestaurant.        // 14:
                                   (ResAccept).             // 15:
                                   feed ResAccept)
       as f in (f(SearchResult).
                SearchResult.                               // 9:
                (CheckSeats).                               // 10:
                ⟨compute DisplayResult⟩.                    // 11:
                DisplayResult.                              // 12:
                (ChooseRestaurant).                         // 13:
                b ⇑ ChooseRestaurant.
                f(ResAccept).
                ResAccept.                                  // 16:
                ⟨compute MapData⟩.                          // 17:
                (NavSystem ⇐ MapData)                       // 18:,19:
       ))
```

where:

- b ⇑ v.P stands for (b ⇐ v.**feed unit**)$>^{\mathbf{1}}$ x >P,
  which in turn unfolds to **stream** (b ⇐ v.**feed unit**) **as** g **in** (g(x).P).
- b ⇓ (x).P stands for (b ⇒ (z)**feed** z)$>^{\mathbf{1}}$ x >P,
  which in turn unfolds to **stream** (b ⇒ (z)**feed** z) **as** g **in** (g(x).P).

Observe that both b ⇑ v.P and b ⇓ (x).P have the same type as P (but in the second case P might not be typable without the extra information of the type of x); furthermore, v:T ⊢ b:?T, hence the use of these abbreviations always fits well with the type system. Indeed, the only typing problem is the one above —namely stream f cannot be adequately given a type.

Implementation with communication via services. The constructions b ⇑ v.P and b ⇓ (x).P may also be used to communicate in the same direction as the stream, hereby avoiding the typing problems both previous solutions suffered from. Thus, we arrive at a third proposal for modelling the dinner service scenario within SSCC, which again follows the proposed SD faithfully.

```
CCS ⇒ (Preferences).                                        // 1:,2:
      (νa1,a2,b)(
                  (DinnerService ⇐ Preferences.             // 3:,4:
                                   ⟨compute GPSdata⟩.       // 5:
                                   GPSdata.                 // 6:
                                   (SearchResult).          // 8:
                                   a1 ⇑ SearchResult.
                                   b ⇓ (ChooseRestaurant).
                                   ChooseRestaurant.        // 14:
                                   (ResAccept).             // 15:
                                   a2 ⇑ ResAccept)
                  |
                  (a1 ⇓ (SearchResult).
                   SearchResult.                            // 9:
                   (CheckSeats).                            // 10:
                   ⟨compute DisplayResult⟩.                 // 11:
                   DisplayResult.                           // 12:
                   (ChooseRestaurant).                      // 13:
                   b ⇑ ChooseRestaurant.
                   a2 ⇓ (ResAccept).
                   ResAccept.                               // 16:
                   ⟨compute MapData⟩.                       // 17:
                   (NavSystem ⇐ MapData)                    // 18:,19:
                  )
      )
```

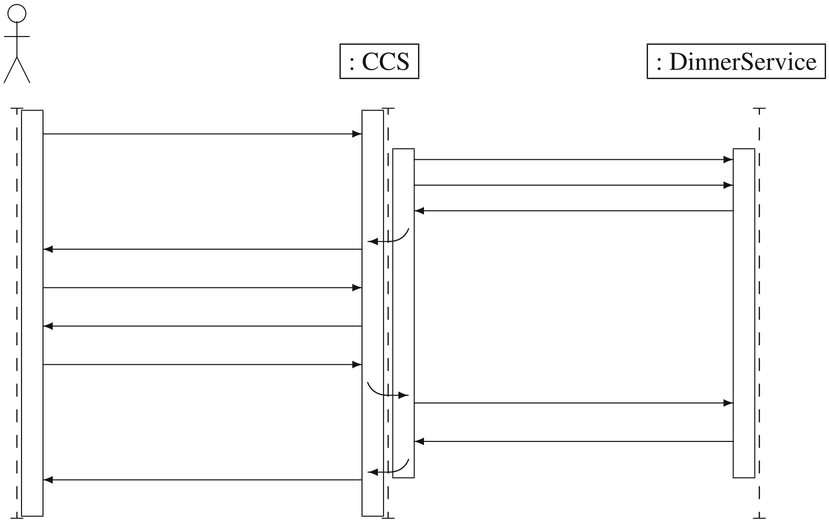

Fig. 13. Two services running in parallel (the curved arrows indicate where information needs to be transmitted between sessions)

With this solution, it can be easily verified that

$\Gamma \vdash \mathsf{CCS} \Rightarrow \ldots : (\mathbf{end}, \mathsf{T})$

for any T (since the process performs no feed to its stream) in any context Γ such that:

$$\begin{aligned}\Gamma \vdash \mathsf{CCS} : {} & [?\mathsf{Preferences}.!\mathsf{SearchResult}.?\mathsf{CheckSeats}.\\ & !\mathsf{DisplayResult}.?\mathsf{ChooseRestaurant}.!\mathsf{ResAccept}.\mathbf{end}]\\ \Gamma \vdash \mathsf{DinnerService} : {} & [?\mathsf{Preferences}.?\mathsf{GPSdata}.!\mathsf{SearchResult}.\\ & ?\mathsf{ChooseRestaurant}.!\mathsf{ResAccept}.\mathbf{end}]\\ \Gamma \vdash \mathsf{NavSystem} : {} & [?\mathsf{MapData}.\mathbf{end}].\end{aligned}$$

The types of services a_1, a_2, and b are derived as in the previous example. Again, the types of the services reflect their communication in the SD with the party who invokes them: the type of CCS is the concatenation of actions 2:, 9:, 10:, 12:, 13:, and 16:; the type of DinnerService is the concatenation of actions 4:, 6:, and 8:; and the type of NavSystem is simply action 19:.

Observe that all internal communication via a_1, a_2, and b is hidden in the type of the services and is not represented in the SD.

Remarks. The Dinner Service scenario exposes an issue with streams: they render communication asymmetric, since a running instance of a service is able to feed information into the process that invoked it, but the latter process has no way to interact back (directly) with that instance of the service (this is actually the intended motivation behind streams, since when a service is invoked it should go on running on its own).

However, in this scenario, the communication system has to synchronize two sessions running concurrently (that with the dinner service and another one with the driver) and information has to run back and forth between them—which is *a priori* not possible due to who invoked whom.

A simple way to work around this problem is using continuations: whenever a session needs extra information from the context to proceed, it saves its state in a new service, feeds the name of that service into the context and dies; later on, the context can invoke the continuation of that session with any extra information that became available in the meantime. Another possibility is to use ephemerous services to communicate in the "opposite" direction.

Figure 13 depicts the sequence of messages exchanged between the intervenients (forgetting activation commands). Observe that CCS is managing two sessions; the curved arrows denote information that must be transmitted from one into the other (which CCS *should* be able to do). The arrows going left can be implemented by **feed**ing into the appropriate stream; the arrow going right must be dealt with using one of the two mechanisms detailed earlier.

This diagram also explains why typing fails for the first and the second solution. Except for their direction, there is no essential distinction between the arrows connecting both sessions CCS is managing. However, streams are specially tailored to capture the right-to-left arrows, while the left-to-right have to be implemented by service invocation. But the purpose of stream communication is quite different: allowing a process invoking several concurrent services to receive answers from them all regardless the order.

Thus, the search for a symmetric manner of modelling all the curled arrows in Fig. 13 is motivated not only because of technical problems, but also by desire for coherence. The third solution proposed to the dinner problem satisfies both requirements: CCS manages communication between the sessions it is involved in a uniform way, and the resulting process is typable.

This example reveals the need for program transformations that make possible the analysis and verification of programs with such characteristics. The following sections address the behavioural theory of processes that allow us to reason about such program transformations.

4. Behavioural theory

This section deals with the dynamic behaviour of SSCC processes. We define first a labelled transition system (LTS) in the early style, and then notions of bisimilarity known in the literature as *strong ground bisimilarity* and *weak ground bisimilarity*. The reason for choosing them is simple: we are interested in capturing contextual equivalence, so bisimilarity should be a congruence. Therefore, we choose the simplest possible setting where this may happen. It is well-known already from the π-calculus that ground bisimilarity over a late LTS is not preserved by parallel composition, requiring the more demanding notions of late and early bisimilarity (which in turn are not preserved by input prefix, since they are not closed under general substitutions). Not surprisingly, this fact also occurs in SSCC: ground bisimilarity is a non-input congruence. Although the general strategy is the same as for π-calculus, the proof techniques themselves differ significantly.

4.1. Labelled transition system

The LTS we define herein adapts (albeit in an insignificant way in terms of expressive power, as discussed below) the original one, presented in [LVMR07a], to allow for the results we are seeking. There are in total roughly thirty transition rules, which can be grouped into four classes. However, only a few rules correspond to a process actually *doing* something (reading or writing a message); the remaining ones either just propagate an action within a component of a process to the whole system, or model interaction between two separate sub-processes.

Labels and rules. We define first the *actions* performed by a process.

Definition 8 (**Labels**). The grammar in Fig. 14 defines the syntax of *labels*.

We define now the LTS and give a short explanation of the rules.

Definition 9 (**Labelled transition system**). The labelled transition system for SSCC is inductively defined by the rules in Fig. 15.

The action of sending a value v is represented by the label $\uparrow v$, while receiving v corresponds to the label $\downarrow v$. Feeding a value into a stream is written $\Uparrow v$, reading from stream f is written $f \Downarrow v$. These labels can only arise by application of rules L-SEND, L-RECEIVE, L-FEED, and L-READ.

The request of a service is represented by the label $a \Leftarrow (r)$, and activation by $a \Rightarrow (r)$. These labels arise by means of rules L-CALL and L-DEF. In the labels a is the name of the service and r the fresh name of the session to be created. To ensure freshness, name r is bound in the label. This is obtained thanks to side conditions of propagation rules below, which guarantee that bound names in the labels are different from free names in the process.

A third group of rules is concerned with propagation of labels, defining the circumstances when a process executes an action because one of its components executes that action. Labels propagate across parallel composition (rules L-PAR and L-PAR') or session endpoints (as long as they do not bind the session names); in this case, the labels for message sending/receiving become prefixed with the session name (rule L-SESS-VAL), while other labels remain unchanged (rule L-SESS-PASS). Here, $\updownarrow v$ stands for one of $\downarrow v$ or $\uparrow v$, and multiple occurrences of $\updownarrow v$ are assumed to be instantiated in the same way, and likewise $a \Leftrightarrow (r)$ stands for one of $a \Rightarrow (r)$ and $a \Leftarrow (r)$.

μ ::=	*Labels*
$\uparrow v$	Value output
$\mid\ \downarrow v$	Value input
$\mid\ a \Rightarrow (r)$	Service definition activation
$\mid\ a \Leftarrow (r)$	Service invocation
$\mid\ \Uparrow v$	Stream feed
$\mid\ f \Downarrow v$	Stream read
$\mid\ r \rhd \uparrow v$	Server session output
$\mid\ r \rhd \downarrow v$	Server session input
$\mid\ r \lhd \uparrow v$	Client session output
$\mid\ r \lhd \downarrow v$	Client session input
$\mid\ r\tau$	Conversation step
$\mid\ \tau$	Internal step
$\mid\ (a) \uparrow a$	Value extrusion
$\mid\ (a) r \rhd \uparrow a$	Server session extrusion
$\mid\ (a) r \lhd \uparrow a$	Client session extrusion
$\mid\ (a) \Uparrow a$	Stream feed extrusion

Fig. 14. The syntax of labels

Every label is also propagated across streams (rules L-STREAM-PASS-P and L-STREAM-PASS-Q) except for feeding or reading by the process at the stream's appropriate end. In the latter cases, the value is fed into/read from the stream, and the process executes an internal action τ (rules L-STREAM-FEED and L-STREAM-CONS); if the value fed is restricted (see below), rule L-FEED-CLOSE should be applied instead. Recursion and restriction do not affect labels (rules L-REC and L-RES), in the latter case as long as these do not refer to the name being restricted. The exception occurs when an action with label $r\tau$ (see below) occurs within a process where r is restricted: it becomes simply an internal action τ. In the remaining cases, name extrusion occurs (rule L-EXTR). Name extrusion may occur whenever a value is communicated, i.e. in value sending (this also includes labels prefixed by a session name) and in value feed. As a notation, the bound name is written in parentheses.

Internal actions, labelled τ, and actions internal to session r, labelled $r\tau$, arise when two subprocesses interact. This can happen when a service invocation meets a service activation (rules L-SERV-COM-PAR and L-SERV-COM-STREAM) or when two endpoints of a session r communicate (rules L-SESS-COM-PAR and L-SERV-COM-PAR); here name extrusion may occur, in which case rules L-PAR-CLOSE or L-SESS-COM-CLOSE apply. In these rules, $\overline{\updownarrow}\ v$ stands for the opposite instantiation of $\updownarrow\ v$ and $a \overline{\Leftrightarrow}(r)$ for the opposite instantiation of $a \Leftrightarrow (r)$.

A few comments are in place at this point. The original LTS for SSCC, presented in [LVMR07a], had fewer rules (in particular, all the propagation rules were absent) and included the following additional rule.

$$\frac{P \xrightarrow{\mu} P' \quad P \equiv Q \quad P' \equiv Q'}{Q \xrightarrow{\mu} Q'} \qquad \text{(L-STRUCT)}$$

Unfortunately, allowing this rule removes the capacity to use structural induction over processes to reason over their actions. Furthermore, although all the new rules are admissible in the original presentation of SSCC, rule L-STRUCT is not admissible in the new calculus, so the latter is strictly weaker. To see why, note that in the original LTS, using L-STRUCT, L-SEND and L-PAR, one can infer that $a \mid b \xrightarrow{\uparrow a} b$. However, without L-STRUCT, the best one can show is that $a \mid b \xrightarrow{\uparrow a} \mathbf{0} \mid b$. Notice that, in the example, one still has $\mathbf{0} \mid b \equiv b$, which for practical purposes suffices. The following result shows that this is not a coincidence.

$$v.P \xrightarrow{\uparrow v} P \qquad (x)P \xrightarrow{\downarrow v} P[{}^{v}\!/\!_{x}] \qquad \mathbf{feed}\, v.P \xrightarrow{\Uparrow v} P \qquad f(x).P \xrightarrow{f \Downarrow v} P[{}^{v}\!/\!_{x}]$$

(L-SEND, L-RECEIVE,L-FEED, L-READ)

$$\frac{r \notin \mathrm{fn}(P)}{a \Leftarrow P \xrightarrow{a \Leftarrow (r)} r \lhd P} \qquad \frac{r \notin \mathrm{fn}(P)}{a \Rightarrow P \xrightarrow{a \Rightarrow (r)} r \rhd P}$$

(L-CALL, L-DEF)

$$\frac{P \xrightarrow{\mu} P' \quad \mathrm{bn}(\mu) \cap \mathrm{fn}(Q) = \emptyset}{P|Q \xrightarrow{\mu} P'|Q} \qquad \frac{Q \xrightarrow{\mu} Q' \quad \mathrm{bn}(\mu) \cap \mathrm{fn}(P) = \emptyset}{P|Q \xrightarrow{\mu} P|Q'}$$

(L-PAR, L-PAR')

$$\frac{P \xrightarrow{\updownarrow v} P'}{r \bowtie P \xrightarrow{r \bowtie \updownarrow v} r \bowtie P'} \qquad \frac{P \xrightarrow{\mu} P' \quad \mu \neq \updownarrow v \quad r \notin \mathrm{bn}(\mu)}{r \bowtie P \xrightarrow{\mu} r \bowtie P'}$$

(L-SESS-VAL, L-SESS-PASS)

$$\frac{P \xrightarrow{\mu} P' \quad \mu \neq \Uparrow v \quad \mathrm{bn}(\mu) \cap (\mathrm{fn}(Q) \cup \mathrm{Set}(\vec{w})) = \emptyset}{\mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q \xrightarrow{\mu} \mathbf{stream}\, P'\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q}$$

(L-STREAM-PASS-P)

$$\frac{Q \xrightarrow{\mu} Q' \quad \mu \neq f \Downarrow v \quad \mathrm{bn}(\mu) \cap (\mathrm{fn}(P) \cup \mathrm{Set}(\vec{w})) = \emptyset}{\mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q \xrightarrow{\mu} \mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q'}$$

(L-STREAM-PASS-Q)

$$\frac{P \xrightarrow{\Uparrow v} P'}{\mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q \xrightarrow{\tau} \mathbf{stream}\, P'\, \mathbf{as}\, f = v :: \vec{w}\, \mathbf{in}\, Q}$$

(L-STREAM-FEED)

$$\frac{P \xrightarrow{(v)\Uparrow v} P' \quad v \notin \mathrm{fn}(Q) \cup \mathrm{Set}(\vec{w})}{\mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q \xrightarrow{\tau} (\nu\, v)\mathbf{stream}\, P'\, \mathbf{as}\, f = v :: \vec{w}\, \mathbf{in}\, Q}$$

(L-FEED-CLOSE)

$$\frac{Q \xrightarrow{f \Downarrow v} Q'}{\mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w} :: v\, \mathbf{in}\, Q \xrightarrow{\tau} \mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q'}$$

(L-STREAM-CONS)

$$\frac{P[{}^{\mathbf{rec}\, X.P}\!/\!_{X}] \xrightarrow{\mu} P'}{\mathbf{rec}\, X.P \xrightarrow{\mu} P'} \qquad \frac{P \xrightarrow{\mu} P' \quad n \notin \mathrm{n}(\mu)}{(\nu\, n)P \xrightarrow{\mu} (\nu\, n)P'}$$

(L-REC, L-RES)

$$\frac{P \xrightarrow{r\tau} P'}{(\nu\, r)P \xrightarrow{\tau} (\nu\, r)P'} \qquad \frac{P \xrightarrow{\mu} P' \quad \mu \in \{\uparrow a, r \bowtie \uparrow a, \Uparrow a\}}{(\nu\, a)P \xrightarrow{(a)\mu} P'}$$

(L-SESS-RES, L-EXTR)

$$\frac{P \xrightarrow{a \Leftrightarrow (r)} P' \quad Q \xrightarrow{a\, \overline{\Leftrightarrow} (r)} Q'}{P|Q \xrightarrow{\tau} (\nu\, r)(P'|Q')} \qquad \frac{P \xrightarrow{a \Leftrightarrow (r)} P' \quad Q \xrightarrow{a\, \overline{\Leftrightarrow} (r)} Q'}{\mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q \xrightarrow{\tau} (\nu\, r)\mathbf{stream}\, P'\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q'}$$

(L-SERV-COM-PAR, L-SERV-COM-STREAM)

$$\frac{P \xrightarrow{r \bowtie \uparrow v} P' \quad Q \xrightarrow{r\, \overline{\bowtie} \downarrow v} Q'}{P|Q \xrightarrow{r\tau} P'|Q'} \qquad \frac{P \xrightarrow{r \bowtie \uparrow v} P' \quad Q \xrightarrow{r\, \overline{\bowtie} \downarrow v} Q'}{\mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q \xrightarrow{r\tau} \mathbf{stream}\, P'\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q'}$$

(L-SESS-COM-PAR, L-SESS-COM-STREAM)

$$\frac{P \xrightarrow{r \bowtie (v) \uparrow v} P' \quad Q \xrightarrow{r\, \overline{\bowtie} \downarrow v} Q' \quad v \notin \mathrm{fn}(Q)}{P|Q \xrightarrow{r\tau} (\nu\, v)(P'|Q')} \qquad \frac{P \xrightarrow{r \bowtie \downarrow v} P' \quad Q \xrightarrow{r\, \overline{\bowtie} (v) \uparrow v} Q' \quad v \notin \mathrm{fn}(P)}{P|Q \xrightarrow{r\tau} (\nu\, v)(P'|Q')}$$

(L-PAR-CLOSE, L-PAR-CLOSE')

$$\frac{P \xrightarrow{r \bowtie (v) \uparrow v} P' \quad Q \xrightarrow{r\, \overline{\bowtie} \downarrow v} Q' \quad v \notin \mathrm{fn}(Q) \cup \mathrm{Set}(\vec{w})}{\mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q \xrightarrow{r\tau} (\nu\, v)\mathbf{stream}\, P'\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q'}$$

(L-SESS-COM-CLOSE)

$$\frac{P \xrightarrow{r \bowtie \downarrow v} P' \quad Q \xrightarrow{r\, \overline{\bowtie} (v) \uparrow v} Q' \quad v \notin \mathrm{fn}(P) \cup \mathrm{Set}(\vec{w})}{\mathbf{stream}\, P\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q \xrightarrow{r\tau} (\nu\, v)\mathbf{stream}\, P'\, \mathbf{as}\, f = \vec{w}\, \mathbf{in}\, Q'}$$

(L-SESS-COM-CLOSE')

Fig. 15. The labelled transition system

Lemma 1 **(Harmony Lemma).** *Let P and Q be processes with $P \equiv Q$. If $P \xrightarrow{\alpha} P'$, then $Q \xrightarrow{\alpha} Q'$ for some Q' with $P' \equiv Q'$, and vice-versa.*

The proof of this result is in Appendix C, together with lemmas showing the derivability of the new transition rules in the original system. The new LTS is strictly weaker than the original one, since it is not always the case that two structurally congruent processes can evolve via the same action to the same process, as the example above shows. However, it is always the case that structurally congruent processes can evolve via the same action to structurally congruent processes: this is precisely the statement of the Harmony Lemma. Observe that this is sufficient for both LTSs to yield the same notion of bisimilarity: the Harmony Lemma implies that structural equivalence is a bisimulation, while all other transition rules are preserved.

The reduction relation on processes (Definition 5) coincides with the restriction of the LTS to τ-transitions.

Lemma 2 **(Correspondence Lemma).** *Let P and Q be processes. It is the case that $P \xrightarrow{\tau} Q$ if, and only if, $P \to Q$.*

4.2. Bisimilarity notions

The first notion of equivalence between processes that we present is strong bisimilarity, hereafter referred to simply as "bisimilarity". Two bisimilar processes behave in the same way, in the sense that they mimic each other's actions (even those that are not visible to the outside) perfectly.

Definition 10

- A symmetric binary relation $\mathcal{R}$ on processes is a *(strong) bisimulation* if, for any processes P, Q such that $P \mathcal{R} Q$, if $P \xrightarrow{\alpha} P'$ for some process P' and action α such that no bound name in α is free in P or Q, then there exists a process Q' such that $Q \xrightarrow{\alpha} Q'$ with $P' \mathcal{R} Q'$.
- (Strong) bisimilarity $\sim$ is the largest bisimulation.
- Two processes P and Q are said to be *(strongly) bisimilar* if $P \sim Q$.

Moreover, a *full* bisimulation is a bisimulation closed under service name substitutions, and we call full bisimilarity $\sim_{\mathrm{f}}$ the largest full bisimulation.

Since bisimilarity is not closed under service name substitutions, $\sim_{\mathrm{f}} \subsetneq \sim$. Bisimilarity (respectively full bisimilarity) can be obtained both as the union of all bisimulations (respectively full bisimulations) or as a fixed-point of a suitable monotonic operator; both characterizations are useful. Notice that "equal" (structurally congruent) processes are bisimilar, since by the Harmony Lemma, $\equiv$ is a bisimulation. Moreover, structurally congruent processes are also fully bisimilar. Bisimilarity as just defined is a non-input congruence, as it is the case in the π-calculus.

Definition 11

- A *context* $C[\![\,]\!]$ is a process where exactly one occurrence of $\mathbf{0}$ has been replaced by a hole $\bullet$. Given a process P, $C[\![P]\!]$ is the process obtained by replacing the hole in $C[\![\,]\!]$ by P.
- A context is said to be *non-input* if no hole occurs under an input prefix (x) or $f(x)$.

Note that here we consider all possible contexts, not only active contexts as defined in Definition 4.

Proposition 1 Bisimilarity is a non-input congruence: if $P \sim Q$ and $C[\![\,]\!]$ is a non-input context, then $C[\![P]\!] \sim C[\![Q]\!]$.

The proof is detailed in Appendix D. The strategy is the same as in [SW01], based on the notion and properties of a relation *progressing* to another relation; however, several major details of the proof are different due to the presence of recursion in the syntax of processes.

Corollary 1 Full bisimilarity is a congruence.

When discussing the behaviour of a system, one often wants to abstract from implementation details and ignore internal actions. In particular, processes are deemed to be equivalent if their *visible* behaviours coincide. This is the goal of weak bisimilarity. Write $P \overset{\tau}{\Longrightarrow} Q$ whenever $P \xrightarrow{\tau} \cdots \xrightarrow{\tau} Q$, and $P \overset{\alpha}{\Longrightarrow} Q$ whenever $P \overset{\tau}{\Longrightarrow}\xrightarrow{\alpha}\overset{\tau}{\Longrightarrow} Q$ for $\alpha \neq \tau$. Notice that, in particular, $P \overset{\tau}{\Longrightarrow} P$ for every process P.

Definition 12

- A symmetric binary relation $\mathcal{R}$ on processes is a *weak bisimulation* if, for any processes P, Q such that $P \mathrel{\mathcal{R}} Q$, if $P \overset{\alpha}{\Longrightarrow} P'$ for some process P' and action α such that no bound name in α is free in P or Q, there exists a process Q' such that $Q \overset{\alpha}{\Longrightarrow} Q'$ with $P' \mathrel{\mathcal{R}} Q'$.
- Weak bisimilarity $\approx$ is the largest weak bisimulation.
- Two processes P and Q are said to be *weakly bisimilar* if $P \approx Q$.

Moreover, a *full* weak bisimulation is a weak bisimulation closed under service name substitutions, and we call full weak bisimilarity $\approx_{\mathrm{f}}$ the largest full weak bisimulation.

Again, (full) weak bisimilarity can be obtained as the union of all (full) weak bisimulations or as a fixed-point of a suitable monotonic operator. The main property of weak bisimilarity is, as before, the following.

Proposition 2 Weak bisimilarity is a non-input congruence.

The proof is detailed in Appendix D. Note that the typical π-calculus examples showing that (weak) bisimulation is not a congruence can be adapted in a straightforward way to SSCC.

Corollary 2 Full weak bisimilarity is a congruence.

4.3. Useful axioms

Even if presenting a complete axiomatization for a calculus as complex as SSCC is out of the scope of this paper, it is interesting to present some axioms (equational laws correct with respect to full strong bisimilarity) that capture key facts about the behaviour of processes. Some of them will also prove to be useful in the next session.

On streams and sessions. In all cases but Axiom 8 (whose proof is described below), the correctness of each axiom below is straightforwardly proved by considering a full bisimulation including all the instances of the axiom together with the identity. The notation $\{\textsf{feed}\ v.Q \to Q\}$ in the Unused Stream law (Axiom 7) denotes a transformation on processes, defined by induction on the syntax, that is an homomorphism for all process constructors, except for **feed**, where it transforms $\textsf{feed}\ v.Q$ into Q, for all v and Q.

Proposition 3

Session Garbage Collection

$$(\nu\, \mathsf{r})\mathcal{D}[\![\mathsf{r} \rhd \mathbf{0}, \mathsf{r} \lhd \mathbf{0}]\!] \sim_{\mathrm{f}} \mathcal{D}[\![\mathbf{0}, \mathbf{0}]\!] \qquad \text{if } \mathcal{D} \text{ does not bind } r \tag{1}$$

Stream Garbage Collection

$$\textsf{stream}\ 0\ \textsf{as}\ f\ \textsf{in}\ P \sim_{\mathrm{f}} P \qquad \text{if } f \text{ does not occur in } P \tag{2}$$

Session Independence

$$r \bowtie \mathsf{Q} \mid s \bowtie \mathsf{P} \sim_{\mathrm{f}} r \bowtie (s \bowtie \mathsf{Q} \mid \mathsf{P}) \qquad \text{if } s \neq r \tag{3}$$

The same holds if the sessions have opposite polarities:

$$r \bowtie \mathsf{Q} \mid s\, \overline{\bowtie}\, \mathsf{P} \sim_{\mathrm{f}} r \bowtie (s\, \overline{\bowtie}\, \mathsf{Q} \mid \mathsf{P}) \qquad \text{if } s \neq r$$

Stream Independence

$$\textsf{stream}\ P\ \textsf{as}\ f\ \textsf{in stream}\ P'\ \textsf{as}\ g\ \textsf{in}\ Q \sim_{\mathrm{f}} \textsf{stream}\ P'\ \textsf{as}\ g\ \textsf{in stream}\ P\ \textsf{as}\ f\ \textsf{in}\ Q \qquad \text{if } f \neq g \tag{4}$$

Streams are Orthogonal to Sessions

$$r \bowtie (\textsf{feed}\ v \mid P) \sim_{\mathrm{f}} \textsf{feed}\ v \mid r \bowtie P \tag{5}$$

Stream Locality

$$\textsf{stream}\ P\ \textsf{as}\ f\ \textsf{in}\ (Q \mid Q') \sim_{\mathrm{f}} (\textsf{stream}\ P\ \textsf{as}\ f\ \textsf{in}\ Q) \mid Q' \qquad \text{if } f \notin \mathrm{fn}(Q') \tag{6}$$

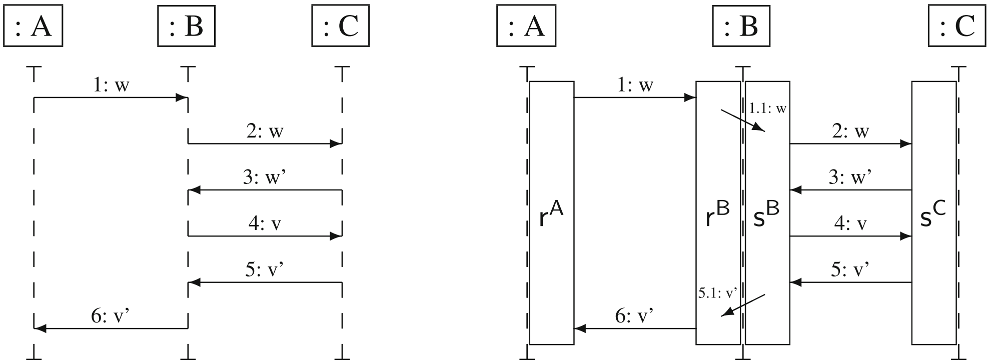

Fig. 16. SD for communication pattern: object-centred and session-centred view

Unused Stream

$$\textbf{stream}\ P\ \textbf{as}\ f\ \textbf{in}\ 0 \approx_{\mathrm{f}} P\{\textbf{feed}\ v.Q \to Q\} \tag{7}$$

Parallel Composition Versus Streams

$$\textbf{stream}\ P\ \textbf{as}\ f\ \textbf{in}\ Q \sim_{\mathrm{f}} P \mid Q \quad \text{if } f \notin \mathrm{fn}(Q) \text{ and } P \text{ does not contain } \textbf{feed} \tag{8}$$

The Session Independence law shows that different sessions are independent. Interestingly this property is strongly connected to the operators available in SSCC, failing in similar calculi such as SCC [BBC+06]. Note that Axiom 7 is correct only with respect to full weak bisimilarity. It becomes correct with respect to full strong bisimilarity if and only if P does not contain any feed which is not inside another stream. We still have to prove Axiom 8.

Proof (*Soundness of Equation* 8) Since $Q \equiv \mathbf{0} \mid Q$ we can apply Axiom 6 to obtain

$$\textbf{stream}\ P\ \textbf{as}\ f\ \textbf{in}\ Q \sim_{\mathrm{f}} (\textbf{stream}\ P\ \textbf{as}\ f\ \textbf{in}\ \mathbf{0}) \mid Q$$

The thesis then follows from Axiom 7. □

5. Program transformations

This section presents an application of the results discussed earlier by means of a more complex example. The first subsection discusses different approaches to implement services as SSCC processes; then a sequence of transformation rules is presented mapping processes obtained by one approach into processes obtained by the other. The results in Sect. 4, together with general results to break sequential sessions into smaller sessions developed in Sect. 5.3, are used to establish soundness of the transformation.

5.1. Program design

We describe how the same behaviour, initially specified in an object-oriented style, can be modelled first in a more session-centered style, and then in a request-response style suitable for implementation. Diagrams for the initial sequence diagram and those describing the result of each transformation are in Figs. 16–18.

The object-oriented view. The SD on the left of Fig. 16 describes a very common pattern appearing in scenarios involving (at least) three partners. The description of the communication pattern is as follows.

Object B receives from object A the value w and forwards it to object C. After receiving the value, object C answers with a value w'. Object B replies with v and, finally, object C replies with v'. Now object B forwards it to object A.

Notice that by "Object B receives from object A the value w" we mean that object A invokes a method in object B passing the value w.

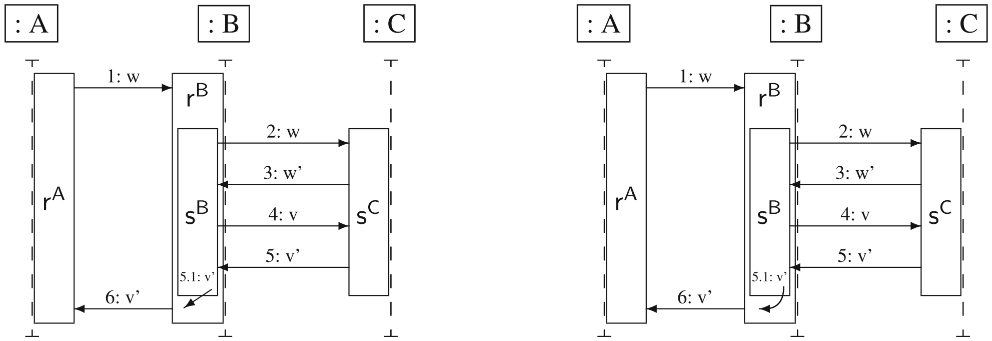

Fig. 17. SD for communication pattern: using a subsession

The session-centered view. Assume that the components of this abstract communication scenario are clients and servers of a service-oriented architecture, and further assume that communication happens via sessions. We refine the diagram by incorporating information about the running sessions, in the diagram on the right of Fig. 16, where the slanted arrows mean message passing between sessions.[4] An instance of service B (let us call *participant* such an instance) has a session r running with participant A and another session s running with participant C. Since sessions involve two partners, a session r between participants A and B has two sides—called endpoints, r^{A} at participant A and r^{B} at participant B. The communication pattern is now like this:

Participant B receives in session r^{B} the value w, passes it to its part of the session with participant C (s^{B}), and then forwards the value through this session to C. Inside the same session, C sends w' to B, B sends v to C and C sends v' to B. Participant B now forwards the value v' back to A, passing it from session s to session r.

In addition to the normal constructs in the calculus, to model object-oriented systems (that do not follow the laws of session communication), it is useful to have two constructs enabling arbitrary message passing. These can be obtained by generalizing auxiliary services to polyadic communications:

$$\begin{aligned} b \Uparrow \langle v_1, \ldots, v_n\rangle.P &\triangleq \textbf{stream}\ b{\Leftarrow}\ v_1 \ldots v_n.\textbf{feed unit as}\ f\ \textbf{in}\ f(v).P \\ b \Downarrow (x_1, \ldots, x_n)P &\triangleq \textbf{stream}\ b{\Rightarrow}\ (z_1)\ldots(z_n).\textbf{feed}\ z_1 \ldots \textbf{feed}\ z_n\ \textbf{as}\ f \\ &\qquad\qquad \textbf{in}\ f(x_1)\ldots f(x_n).P \end{aligned}$$

where name v and stream f are not used in P.

The diagram on the right of Fig. 16 is directly implemented in SSCC as

$$\mathsf{SC} \triangleq (\nu\, b, c)\,(\,\mathsf{A} \mid \mathsf{B} \mid \mathsf{C}\,),$$

where

$$\mathsf{A} \triangleq b{\Leftarrow}\, w.(y)P,\ \ \mathsf{B} \triangleq (\nu\, b_1, b_2)\,(\,\mathsf{B}_1 \mid \mathsf{B}_2\,),\ \text{ and } \mathsf{C} \triangleq c{\Rightarrow}\,(x)w'.(y)v'.S,$$

with

$$\mathsf{B}_1 \triangleq b{\Rightarrow}\,(x)b_1 \Uparrow x.b_2 \Downarrow (y)y.Q,\ \text{ and } \mathsf{B}_2 \triangleq c{\Leftarrow}\, b_1 \Downarrow (x)x.(z)v.(y)b_2 \Uparrow y.R.$$

The process above, although not deterministic (e.g., its first step may either be the invocation of service b or of service c), is confluent, and it is easy to check that its behaviour reflects the one described on the right of Fig. 16.

A first optimization. When participant B receives the value sent by A, it may immediately send it to participant C, by calling it (and thus opening a subsession). One simply has to perform a "local" transformation on B. The resulting diagram is on the left of Fig. 17, and it is implemented in SSCC as process SC′, where we denote by E the new instance of B.

$$\begin{aligned} \mathsf{SC}' &\triangleq (\nu\, b, c)\,(\,\mathsf{A} \mid \mathsf{E} \mid \mathsf{C}\,) \\ \mathsf{E} &\triangleq b{\Rightarrow}\,(x)(\nu\, b_1)(c{\Leftarrow}\, x.(z)v.(y)b_1 \Uparrow y.R \mid b_1 \Downarrow (y)y.Q) \end{aligned}$$

Still, participant E (which replaces B from the previous version) needs to pass the value sent by C in their session, to its session with A, and a communication based on an auxiliary service is used.

[4] Since to the best of our knowledge no extension of Sequence Diagrams with session information exists, we introduce in SD a notation for it.

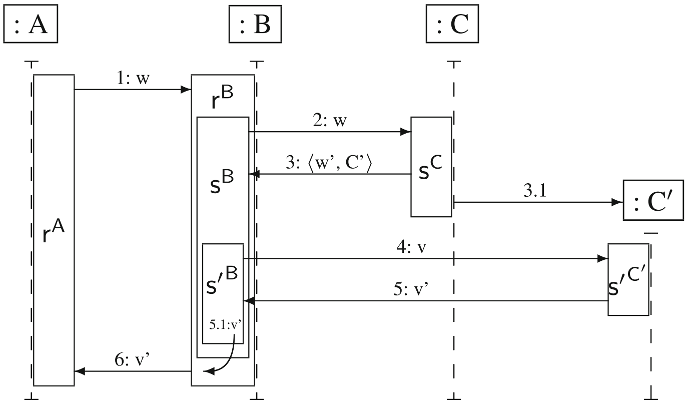

Fig. 18. SD for communication pattern: using subsessions and continuations

A second optimization. The transfer of a value from a subsession to its parent session is, in the previous implementation, not straightforward, since it requires the use of local services. In fact, it is more convenient to use a trans-session construct like **stream**. Now participant F (initially B), passes the value received from C, from the subsession to the main session, using another communication construct—a stream—instead of using a local service. This is implemented in process SC″.

$$\mathsf{SC}'' \triangleq (\nu\, b,\, c)\,(\,\mathsf{A} \mid \mathsf{F} \mid \mathsf{C}\,)$$

where the new code for B (now F) is as follows.

$$\mathsf{F} \triangleq b{\Rightarrow}\,(x)\,(\textbf{stream}\ c{\Leftarrow}\, x.(z)v.(y)\textbf{feed}\ y.R\ \textbf{as}\ f\ \textbf{in}\ f(y).y.Q)$$

The corresponding diagram is on the right of Fig. 17. The two diagrams in Fig. 17 are quite similar: the one on the left uses message passing (denoted by the straight arrow from the inner to the outer session), whereas the one on the right uses stream communication (described by the curved arrow).

Implementing the diagram. The previous diagram, and the corresponding SSCC process, model the pattern at hand in a service-oriented, session-based style. However, current Web Service technologies do not provide support for a complex mechanism such as sessions, considering instead request (one-way communication) and request-response (two one-way communications in opposite directions) only—see, e.g., the definition of WSDL [CCMW01]. It is easy to see that these are particular cases of sessions, where protocols are composed respectively by one output or by one output followed by one input. The new communication pattern is described in Fig. 18, and reads as follows.

Participant B receives in session r^{B} the value w, and then forwards it through its session s^{B} with participant C to C itself. Inside the same session C sends to B value w′ together with the name of a freshly generated service C′ to continue the conversation on. Now B invokes C′ creating a new subsession s′ of session s and, inside s′, B sends v and receives as answer v′. Participant B now forwards the value v′ back to A, passing it from session s′ to session r.

This pattern can be implemented as:

$$\mathsf{SC}''' \triangleq (\nu\, b,\, c)\,(\,\mathsf{A} \mid \mathsf{G} \mid \mathsf{D}\,)$$

where the new codes for B (now G) and C (now D) are below. To write these new codes we need to consider polyadic inputs and outputs, denoted respectively by $(x_1, \ldots, x_n)$ and $\langle v_1, \ldots, v_n\rangle$. They can be easily accommodated in the theory.

$$\mathsf{G} \triangleq b{\Rightarrow}\,(x)(\textbf{stream}\ c{\Leftarrow}\, x.(z,\, c')c'{\Leftarrow}\, v.(y)\textbf{feed}\ y.R\ \textbf{as}\ f\ \textbf{in}\ f(y).y.Q)$$
$$\mathsf{D} \triangleq c{\Rightarrow}\,(x)(\nu\, c')\langle w',\, c'\rangle.c'{\Rightarrow}\,(y)v'.S$$

Naturally, one asks whether the transformations of SC into SC′, SC″, and, finally, SC‴ are correct, not changing the observable behaviour of processes. Next we show how to give a positive answer to this question.

5.2. Soundness of the transformations

We now prove that the transformations presented are actually correct with respect to full weak bisimilarity. Interestingly, they are also transparent for process A, i.e., A needs not to be changed when the transformation is applied. To prove this we show that the three equations below hold.

$$(\nu\, \mathsf{c})(\mathsf{B} \mid \mathsf{C}) \;\approx_{\mathrm{f}}\; (\nu\, \mathsf{c})(\mathsf{E} \mid \mathsf{C}) \tag{9}$$

$$(\nu\, \mathsf{c})(\mathsf{E} \mid \mathsf{C}) \;\approx_{\mathrm{f}}\; (\nu\, \mathsf{c})(\mathsf{F} \mid \mathsf{C}) \tag{10}$$

$$(\nu\, \mathsf{c})(\mathsf{F} \mid \mathsf{C}) \;\approx_{\mathrm{f}}\; (\nu\, \mathsf{c})(\mathsf{G} \mid \mathsf{D}) \tag{11}$$

The correctness of the whole transformations, i.e., $\mathsf{SC} \approx_{\mathrm{f}} \mathsf{SC}' \approx_{\mathrm{f}} \mathsf{SC}'' \approx_{\mathrm{f}} \mathsf{SC}'''$ follows by closing under contexts the equations above.

Note that the transformation of auxiliary communications (passing value w from r^{B} to s^{B} and value v from s^{B} to r^{B}) into normal communications are not correct with respect to full strong bisimilarity. In fact, auxiliary communications require a few more steps, and leave behind them empty sessions and streams, which have to be garbage collected. The correctness of garbage collection is based on Axioms 1 and 2 (cfr. Sect. 4.3).

Proof (*Soundness of Equation* 9) The proof can be easily obtained by exhibiting a bisimulation including the two processes. For simplicity we will not detail it, but just highlight a few important points. For easier reading we recall B, E and C below.

$$\mathsf{B} \triangleq (\nu\, b_1, b_2)\,(\, b{\Rightarrow}\,(x)b_1 \Uparrow x.b_2 \Downarrow (y)y.Q \mid c{\Leftarrow}\; b_1 \Downarrow (x)x.(z)v.(y)b_2 \Uparrow y.R\,)$$
$$\mathsf{E} \triangleq b{\Rightarrow}\,(x)(\nu\, b_1)(c{\Leftarrow}\, x.(z)v.(y)b_1 \Uparrow y.R \mid b_1 \Downarrow (y)y.Q)$$
$$\mathsf{C} \triangleq c{\Rightarrow}\,(x)w'.(y)v'.S$$

The two processes can mimic each other even if they are nondeterministic, since the nondeterminism comes from τ steps, whose order is not important, given that processes are confluent. We thus just consider a possible order of execution of τ steps. Both the processes can execute the following sequence of actions: a service invocation at b giving rise to session r, an input on x taking value w, a service invocation at c giving rise to session s. Then B can execute the auxiliary communication along b_1. Then both B and C can execute their protocol with C. By executing the two missing auxiliary communications $(\nu\, \mathsf{c})(\mathsf{B} \mid \mathsf{C})$ and $(\nu\, \mathsf{c})(\mathsf{E} \mid \mathsf{C})$ reduce respectively to:

$$(\nu\, \mathsf{s})(r \rhd \mathsf{Q}[^{w}\!/\!_{x}][^{v'}\!/\!_{y}] \;\mid\; s \lhd \mathsf{R}[^{w}\!/\!_{x}][^{w'}\!/\!_{z}][^{v'}\!/\!_{y}] \;\mid\; s \rhd \mathsf{S}[^{w}\!/\!_{x}][^{v}\!/\!_{y}])$$
$$(\nu\, \mathsf{s})(r \rhd (s \lhd \mathsf{R}[^{w}\!/\!_{x}][^{w'}\!/\!_{z}][^{v'}\!/\!_{y}]) \;\mid\; \mathsf{Q}[^{w}\!/\!_{x}][^{v'}\!/\!_{y}] \;\mid\; s \rhd \mathsf{S}[^{w}\!/\!_{x}][^{v}\!/\!_{y}])$$

where we used structural congruence (which is included in full bisimilarity, according to Theorem 6 in Appendix D), garbage collection Axioms 1 and 2, and closure under contexts to remove the garbage produced by auxiliary communications.

The processes above can be proved equivalent using structural congruence, session independence (Axiom 3) and closure under contexts. □

Proof (*Soundness of Equation* 10) To prove the correctness of Eq. 10 it is enough to prove $E \approx_{\mathrm{f}} F$, then the thesis follows from closure under contexts. For easier reading we recall E and F below.

$$\mathsf{E} \triangleq b{\Rightarrow}\,(x)(\nu\, b_1)(c{\Leftarrow}\, x.(z)v.(y)b_1 \Uparrow y.R \mid b_1 \Downarrow (y)y.Q)$$
$$\mathsf{F} \triangleq b{\Rightarrow}\,(x)(\textbf{stream}\ c{\Leftarrow}\, x.(z)v.(y)\textbf{feed}\ y.R\ \textbf{as}\ f\ \textbf{in}\ f(y).y.Q)$$

Actually, in general we can prove

$$(\nu\, \mathsf{a})(\mathcal{C}'[\![a \Uparrow v.P]\!] \mid \mathcal{C}''[\![a \Downarrow (y).Q]\!]) \approx_{\mathrm{f}} \textbf{stream}\ \mathcal{C}'[\![\textbf{feed}\ v.P]\!]\ \textbf{as}\ f\ \textbf{in}\ \mathcal{C}''[\![f(y).Q]\!] \tag{12}$$

provided that neither a nor f occur elsewhere and P and $\mathcal{C}'$ contain no feeds. The thesis will follow by instantiation and closure under contexts.

The proof shows that the three pairs below, together with a few other pairs differing from these because of τ transitions (corresponding to intermediate steps), form a full bisimulation.

$$
\begin{array}{l}
((\nu\, \mathsf{a})(\mathcal{C}'\llbracket a \Uparrow v.P\rrbracket \mid \mathcal{C}''\llbracket a \Downarrow (y).Q\rrbracket), \textbf{stream}\ \mathcal{C}'\llbracket \textbf{feed}\ v.P\rrbracket\ \textbf{as}\ f\ \ \textbf{in}\ \mathcal{C}''\llbracket f(y).Q\rrbracket) \\
(\mathcal{C}'\llbracket P\rrbracket \mid \mathcal{C}''\llbracket \textbf{stream 0 as}\ f' = \langle v\rangle\ \textbf{in}\ f'(y).Q\rrbracket, \\
\qquad\qquad\qquad\qquad\qquad \textbf{stream}\ \mathcal{C}'\llbracket P\rrbracket\ \textbf{as}\ f = \langle v\rangle\ \textbf{in}\ \mathcal{C}''\llbracket f(y).Q\rrbracket) \\
(\mathcal{C}'\llbracket P\rrbracket \mid \mathcal{C}''\llbracket Q\rrbracket, \textbf{stream}\ \mathcal{C}'\llbracket P\rrbracket\ \textbf{as}\ f\ \ \textbf{in}\ \mathcal{C}''\llbracket Q\rrbracket)
\end{array}
$$

In each process considered in the relation, a, f, and f' do not occur elsewhere, and P and $\mathcal{C}'$ do not contain feeds. The only difficult part is when on the right the feed and the read from stream are executed and, correspondingly, the two auxiliary communications on the left. Actually, both transitions on the right amount to τ actions.

$$\textbf{stream}\ \mathcal{C}'\llbracket \textbf{feed}\ v.P\rrbracket\ \textbf{as}\ f\ \ \textbf{in}\ \mathcal{C}''\llbracket f(y).Q\rrbracket \xrightarrow{\tau} \textbf{stream}\ \mathcal{C}'\llbracket P\rrbracket\ \textbf{as}\ f = \langle v\rangle\ \textbf{in}\ \mathcal{C}''\llbracket f(y).Q\rrbracket$$

$$\textbf{stream}\ \mathcal{C}'\llbracket P\rrbracket\ \textbf{as}\ f = \langle v\rangle\ \textbf{in}\ \mathcal{C}''\llbracket f(y).Q\rrbracket \xrightarrow{\tau} \textbf{stream}\ \mathcal{C}'\llbracket P\rrbracket\ \textbf{as}\ f\ \ \textbf{in}\ \mathcal{C}''\llbracket Q[{}^{v}\!/\!_{y}]\rrbracket$$

The first transition is matched as follows.

$$
\begin{array}{l}
(\nu\, \mathsf{a})\mathcal{C}'\llbracket a \Uparrow v.P\rrbracket \mid \mathcal{C}''\llbracket a \Downarrow (y).Q\rrbracket = \\
(\nu\, \mathsf{a})\mathcal{C}'\llbracket \textbf{stream}\ a \Leftarrow v.\textbf{feed unit as}\ f''\ \ \textbf{in}\ f''(x).P\rrbracket \mid \\
\qquad\qquad\qquad \mathcal{C}''\llbracket \textbf{stream}\ a \Rightarrow (z)\textbf{feed}\ z\ \textbf{as}\ f'\ \ \textbf{in}\ f'(y).Q\rrbracket \xrightarrow{\tau}{}^{*} \\
(\nu\, \mathsf{a}, \mathsf{r})\mathcal{C}'\llbracket \textbf{stream}\ r \lhd \textbf{0 as}\ f''\ \ \textbf{in}\ P\rrbracket \mid \mathcal{C}''\llbracket \textbf{stream}\ r \rhd \textbf{0 as}\ f' = \langle v\rangle\ \textbf{in}\ f'(y).Q\rrbracket \sim_{\mathrm{f}} \\
\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \mathcal{C}'\llbracket P\rrbracket \mid \mathcal{C}''\llbracket \textbf{stream 0 as}\ f' = \langle v\rangle\ \textbf{in}\ f'(y).Q\rrbracket
\end{array}
$$

In the first step we used the definitions of auxiliary communications. The sequence of τ actions includes service invocation at a creating session r, communication of v along the session, the two feeds and the read from stream f''. In the last step we used Axioms 1 and 2 for garbage collection. For the challenge from the left, note that only τ actions are involved, thus they can be matched by the right term staying idle, and the order in which they are executed is not relevant.

The second transition from the right is matched by:

$$
\begin{array}{l}
\mathcal{C}'\llbracket P\rrbracket \mid \mathcal{C}''\llbracket \textbf{stream 0 as}\ f' = \langle v\rangle\ \textbf{in}\ f'(y).Q\rrbracket \xrightarrow{\tau} \\
\qquad\qquad\qquad\qquad \mathcal{C}'\llbracket P\rrbracket \mid \mathcal{C}''\llbracket \textbf{stream 0 as}\ f'\ \ \textbf{in}\ Q[{}^{v}\!/\!_{y}]\rrbracket \sim_{\mathrm{f}} \mathcal{C}'\llbracket P\rrbracket \mid \mathcal{C}''\llbracket Q[{}^{v}\!/\!_{y}]\rrbracket
\end{array}
$$

where we used Axiom 2 again. This concludes the proof. □

Equation 11 only holds when processes are sequential. Therefore, we adapt the type system for that purpose and present then the envisaged result.

5.3. Breaking sequential sessions

The equations proved in the previous section hold for arbitrary processes, and allow us to prove the correctness of the first two optimizations. However to prove the correctness of the implementation step they are not enough. In fact, it is not easy to break a session allowing the conversation to continue in a freshly generated new session, since, in general, communication patterns inside sessions can be quite complex, e.g., since sessions may include many ongoing concurrent communications. However, a small class of sequential sessions captures the most interesting within-session behaviours. Since such a class can be identified by a type system, we start by presenting the type system for *session sequentiality* (similar to the type system for single threadness for mobile safe ambients in [LS03]), and then we present properties of well-typed processes that allow us to prove the correctness of the last transformation.

The new type system is a simplification of the one previously presented, which guarantees also protocol compatibility. Moreover, we consider here just finite types, thus session protocols should be finite. Notice that this constraint does not forbid infinite behaviours, but just infinite sessions. In particular, if a process is typable according to the general type system and all the involved types are non recursive, then the process is typable according to this type system.

$$\frac{P\colon U}{v.P\colon !.U} \qquad \frac{P\colon U}{(x)P\colon ?.U} \qquad \frac{P\colon U}{a \Rightarrow P\colon \textsf{end}} \qquad \frac{P\colon U}{a \Leftarrow P\colon \textsf{end}} \qquad (\text{T-SEND, T-RECEIVE, T-DEF, T-CALL})$$

$$\frac{P\colon U}{r \rhd P\colon \textsf{end}} \qquad \frac{P\colon U}{r \lhd P\colon \textsf{end}} \qquad \frac{P\colon U}{\textsf{feed}\, v.P\colon U} \qquad \frac{P\colon U}{f(x).P\colon U} \qquad (\text{T-SESS-S, T-SESS-C, T-FEED, T-READ})$$

$$\frac{P\colon U \quad Q\colon \textsf{end}}{P|Q\colon U} \qquad \frac{P\colon \textsf{end} \quad Q\colon U}{P|Q\colon U} \qquad (\text{T-PAR-L, T-PAR-R})$$

$$\frac{P\colon U \quad Q\colon \textsf{end}}{\textsf{stream}\, P\, \textsf{as}\, f = \vec{v}\, \textsf{in}\, Q\colon U} \qquad \frac{P\colon \textsf{end} \quad Q\colon U}{\textsf{stream}\, P\, \textsf{as}\, f = \vec{v}\, \textsf{in}\, Q\colon U} \qquad (\text{T-STREAM-L, T-STREAM-R})$$

$$\frac{P\colon \textsf{end}}{\textsf{rec}\, X.P\colon \textsf{end}} \qquad \frac{P\colon U}{(\nu\, n)P\colon U} \qquad X\colon \textsf{end} \qquad \mathbf{0}\colon \textsf{end} \qquad (\text{T-REC, T-RES, T-VAR, T-NIL})$$

Fig. 19. Type system for session sequentiality

We consider typed processes of the form $P\colon U$ where U is the protocol type. We consider as types $?.U$, $!.U$, and **end**, denoting respectively a protocol that performs an input and then continues as prescribed by U, a protocol that performs an output and then continues as prescribed by U, and the terminated protocol. It is clear that in this setting a request is a session with protocol $!.$**end** (and complementary protocol $?.$**end**), while a request-response has protocol $!.?.$**end** (and complementary protocol $?.!.$**end**).

Definition 13 (**Type system for session sequentiality**). The type system is inductively defined by the rules in Fig. 19.

Under the typability assumption, SSCC sessions are sequential in a very strong sense: we can statically define a correspondence between inputs and outputs such that each input is always matched by the corresponding output. We show how to break sessions, i.e., how to make the conversation continue on a freshly created new session. The general law is presented under Theorem 3. The two pieces of the broken session have protocols that are simpler than the original one, thus by repeatedly applying the transformation we can reduce any protocol to a composition of request and request-response patterns. We formalize the procedure described so far.

Definition 14 Let P be a process. An input/output prefix inside P is *at top-level in* P if it is neither inside a service definition/invocation nor inside a session. Given a process P we can assign sequential indices to top-level input/output prefixes in P according to the position of their occurrence in the term, starting from 1. Thus the *i-th top-level prefix in* P is the top-level prefix in P that occurs in i-th position.

For instance, let P be $a.(x).\textsf{stream}\,(y).\textsf{feed}\,y\,\textsf{as}\,f\;\textsf{in}\,f(y).y\Leftarrow a.(z).\textsf{feed}\,z$. Then P annotated with indices on its top-level prefixes is:

$a : 1.(x) : 2.\textsf{stream}\,(y) : 3.\textsf{feed}\,y\,\textsf{as}\,f\;\textsf{in}\,f(y).y\Leftarrow a.(z).\textsf{feed}\,z$

Definition 15 Given a process P with a subterm Q, we say that Q is enabled in P if $P = \mathcal{C}[\![Q]\!]$ for some active context $\mathcal{C}[\![\bullet]\!]$. The same definition holds also for prefixes.

Intuitively a subterm is enabled when it can execute.

Lemma 3 *Let P be a process typable with the type system for session sequentiality. If P has type* **end***, then it has no top-level enabled prefixes, otherwise it has exactly one top-level enabled prefix, and this has index* 1.

Proof By induction on the typing proof of P. The thesis follows trivially for rules T-SEND, T-RECEIVE, T-DEF, T-CALL, T-SESS-S, T-SESS-C, T-FEED, T-READ, and T-NIL. For rule T-RES, it follows by inductive hypothesis. For rules T-PAR-L, T-PAR-R, T-STREAM-L, and T-STREAM-R, it follows from the observation that one of the two sides has no enabled prefix, thus the inductive hypothesis can be applied on the other side. □

Definition 16 Given a transition $P \xrightarrow{\alpha} Q$ and a prefix in P we say that the prefix is consumed if, in the derivation of the transition, rule L-SEND or L-RECEIVE is applied to the prefix.

Notice that the consumed prefix does not occur in Q.

Lemma 4 *Let $P_0 = (\nu\, a)\mathcal{D}\llbracket a\!\Rightarrow P,\, a\!\Leftarrow Q\rrbracket$ be a typed process such that a does not occur in $\mathcal{D}\llbracket\bullet, \bullet\rrbracket$. Suppose $P_0 \xrightarrow{\alpha_0} \ldots \xrightarrow{\alpha_{n-1}} P_n$. Then either:*

- *for all $i \in \{1, \ldots, n\}$ we have $P_i = (\nu\, a)\mathcal{D}'_i\llbracket a\!\Rightarrow P,\, a\!\Leftarrow Q\rrbracket$ for some $\mathcal{D}'_i\llbracket\bullet, \bullet\rrbracket$ or*
- *there exists $j \in \{1, \ldots, n\}$ such that for each $i < j$ we have $P_i = (\nu\, a)\mathcal{D}'_i\llbracket a\!\Rightarrow P,\, a\!\Leftarrow Q\rrbracket$ for some $\mathcal{D}'_i\llbracket\bullet, \bullet\rrbracket$, and for each $k \geq j$ we have $P_k = (\nu\, r)\mathcal{D}'_k\llbracket r \rhd P'_k,\, r \lhd Q'_k\rrbracket$ for some $\mathcal{D}'_k\llbracket\bullet, \bullet\rrbracket$, processes P'_k and Q'_k and session name r. Furthermore, in a transition $P_k \xrightarrow{\alpha_k} P_{k+1}$ the m-th prefix (according to Definition 14) in P'_k is consumed if and only if the m-th prefix in Q'_k is consumed, and the two are synchronized.*

Proof The proof is by induction on n. The base case ($n = 0$) is trivial. Let us consider the inductive case. When $P_i = (\nu\, a)\mathcal{D}'_i\llbracket a\!\Rightarrow P,\, a\!\Leftarrow Q\rrbracket$ the only possible transitions (since a does not occur in $\mathcal{D}'_i\llbracket\bullet, \bullet\rrbracket$) are transitions involving only the context, which leave the process in the same form, or the interaction between the service invocation on a and the service definition of a in the holes. If all the transitions are of the first category then we are in the first case of the lemma. Otherwise, let $j - 1$ be the first transition of the second category. This leads to a process of the form $P_j = (\nu\, r)\mathcal{D}'_j\llbracket r \rhd P'_j,\, r \lhd Q'_j\rrbracket$. Let us consider now processes of this form. In order to prove the condition on prefixes we show that the following property holds: at each step either (i) all the prefixes preserve their indices, or (ii) prefixes with index 1 are consumed and the indices of all other prefixes decrease by 1. For transitions not involving prefixes, the thesis follows trivially. Suppose now that, e.g., in P'_k an output prefix is consumed (the case where the prefix is an input prefix or the consumed prefix is in Q'_k is symmetric). Thus $P'_k \xrightarrow{\uparrow v} P'_{k+1}$ and $r \rhd P'_k \xrightarrow{r \rhd \uparrow v} r \rhd P'_{k+1}$. Since r is private, this label should interact with a label of the form $r \lhd \downarrow v$. Since sessions are only runtime syntax (cfr. Lemma 5 in Appendix A) such a label can be generated only by $r \lhd Q'_k$, and thanks to Lemma 3 this must be the prefix with index 1. Thus, in both P'_{k+1} and Q'_{k+1} prefix indices are equal to the indices in P'_k and Q'_k minus 1. This proves the thesis. □

We now have all the tools required to prove the correctness of the session breaking technique.

Theorem 3 *Let $(\nu\, a)\mathcal{D}\llbracket a\!\Leftarrow \mathcal{C}\llbracket (x).P\rrbracket,\, a\!\Rightarrow \mathcal{C}'\llbracket v.Q\rrbracket\rrbracket$ be a typed process such that a does not occur in $\mathcal{D}\llbracket\bullet, \bullet\rrbracket$. Suppose that there exists i such that $(x).P$ and $v.Q$ are the i-th top-level prefixes in $\mathcal{C}\llbracket (x).P\rrbracket$ and in $\mathcal{C}'\llbracket v.Q\rrbracket$, respectively. Let $y \notin \mathrm{fn}(P)$, $b \notin \mathrm{fn}(Q)$. Then:*

$$(\nu\, a)\mathcal{D}\llbracket a\!\Leftarrow \mathcal{C}\llbracket (x).P\rrbracket,\, a\!\Rightarrow \mathcal{C}'\llbracket v.Q\rrbracket\rrbracket \approx_{\mathrm{f}} (\nu\, a)\mathcal{D}\llbracket a\!\Leftarrow \mathcal{C}\llbracket (x, y).y\!\Leftarrow P\rrbracket,\, a\!\Rightarrow \mathcal{C}'\llbracket (\nu\, b)\langle v, b\rangle.b\!\Rightarrow Q\rrbracket\rrbracket$$

Proof We show that the following relation is a full weak bisimulation,

$$\begin{aligned}
&\{((\nu\, a)\mathcal{D}\llbracket a\!\Leftarrow \mathcal{C}\llbracket (x).P\rrbracket,\, a\!\Rightarrow \mathcal{C}'\llbracket v.Q\rrbracket\rrbracket, \quad (\nu\, a)\mathcal{D}\llbracket a\!\Leftarrow \mathcal{C}\llbracket (x, y).y\!\Leftarrow P\rrbracket,\, a\!\Rightarrow \mathcal{C}'\llbracket (\nu\, b)\langle v, b\rangle.b\!\Rightarrow Q\rrbracket\rrbracket)\\
&((\nu\, r)\mathcal{D}\llbracket r \lhd \mathcal{C}\llbracket (x).P\rrbracket,\, r \rhd \mathcal{C}'\llbracket v.Q\rrbracket\rrbracket, \quad (\nu\, r)\mathcal{D}\llbracket r \lhd \mathcal{C}\llbracket (x, y).y\!\Leftarrow P\rrbracket,\, r \rhd \mathcal{C}'\llbracket (\nu\, b)\langle v, b\rangle.b\!\Rightarrow Q\rrbracket\rrbracket)\\
&((\nu\, r)\mathcal{D}\llbracket r \lhd \mathcal{C}\llbracket P\rrbracket,\, r \rhd \mathcal{C}'\llbracket Q\rrbracket\rrbracket, (\nu\, r, b)\mathcal{D}\llbracket r \lhd \mathcal{C}\llbracket b\!\Leftarrow P\rrbracket,\, r \rhd \mathcal{C}'\llbracket b\!\Rightarrow Q\rrbracket\rrbracket)\\
&((\nu\, r)\mathcal{D}\llbracket r \lhd \mathcal{C}\llbracket P\rrbracket,\, r \rhd \mathcal{C}'\llbracket Q\rrbracket\rrbracket, (\nu\, r, r')\mathcal{D}\llbracket r \lhd \mathcal{C}\llbracket r' \lhd P\rrbracket,\, r \rhd \mathcal{C}'\llbracket r' \rhd Q\rrbracket\rrbracket)\}
\end{aligned}$$

where all the names and processes are universally quantified, $y \notin \mathrm{fn}(P)$, $b \notin \mathrm{fn}(Q)$, and $(x).$ and $v.$ have the same index.

Processes in the first pair can move only to processes of the same shape or to processes in the second pair because of Lemma 4. Similarly, processes in the second pair can move to processes of the same form or to processes in the third pair since prefixes $(x).$ and $v.$ have the same index, and thus are consumed together, again thanks to Lemma 4. In the third pair the only transition that can change the structure of the processes is the invocation of service b on the right, but since this is a τ step, the left part can answer by staying idle. Processes in the last pair only evolve to processes of the same shape. This concludes the proof. □

We are now in a position to prove the soundness of the last transformation.

Proof (*Soundness of Equation 11*) We have to prove that $(\nu\, \mathsf{c})(\mathsf{F} \mid \mathsf{C}) \approx_{\mathrm{f}} (\nu\, \mathsf{c})(\mathsf{G} \mid \mathsf{D})$. For easier reading we recall F, C, G and D below.

$$\begin{aligned}
\mathsf{F} &\triangleq b\!\Rightarrow (x)(\textbf{stream } c\!\Leftarrow x.(z)v.(y)\textbf{feed } y.R \textbf{ as } f \textbf{ in } f(y).y.Q)\\
\mathsf{C} &\triangleq c\!\Rightarrow (x)w'.(y)v'.S\\
\mathsf{G} &\triangleq b\!\Rightarrow (x)(\textbf{stream } c\!\Leftarrow x.(z, c')c'\!\Leftarrow v.(y)\textbf{feed } y.R \textbf{ as } f \textbf{ in } f(y).y.Q)\\
\mathsf{D} &\triangleq c\!\Rightarrow (x)(\nu\, c')\langle w', c'\rangle.c'\!\Rightarrow (y)v'.S
\end{aligned}$$

It is easy to verify that $(\nu\, \mathsf{c})(\mathsf{F} \mid \mathsf{C})$ can be typed according to the type system for session sequentiality. By considering:

$$\mathcal{D}[\![\bullet_1, \bullet_2]\!] = b\Rightarrow (x)(\textsf{stream } \bullet_1 \textsf{ as } f \textsf{ in } f(y).y.Q) \mid \bullet_2$$

$$\begin{array}{ll} \mathcal{C}[\![\bullet]\!] = x.\bullet & \quad P = v.(y)\textsf{feed } y.R \\ \mathcal{C}'[\![\bullet]\!] = (x)\bullet & \quad Q = (y)v'.S \end{array}$$

we can apply Theorem 3 to get the thesis, since prefixes (z) (in process F) and w' (in process C) have both index 2. □

Notice that the technique above can be extended to break protocols of services with more than one definition/invocation: simply break all of them at the same point.

Let us consider the following example. Let

$$P \triangleq a\Leftarrow v.(x)\textsf{stream } b\Leftarrow x.(y).\textsf{feed } y \textsf{ as } f \textsf{ in } f(z).z.(w)$$

together with servers for a and b

$$A \triangleq a\Rightarrow (x_1).v_1.(x_2).v_2 \qquad\qquad B \triangleq b\Rightarrow (x_3).v_3$$

It is easy to check that $P: !.?.!.?.\textsf{end}$, $A: ?.!.?.!.\textsf{end}$, and $B: ?.!.\textsf{end}$.

While B can be easily programmed as a solicit-response (the complementary operation to request-response), P and A require a transformation. Notice that input (x). in P and output v_1. in A have both index 2. We can apply the transformation, thus obtaining processes Q and C equivalent to P and A, respectively, while B is unchanged.

$$Q \triangleq a\Leftarrow v.(x,\, s)s\Leftarrow \textsf{stream } b\Leftarrow x.(y).\textsf{feed } y \textsf{ as } f \textsf{ in } f(z).z.(w)$$
$$C \triangleq a\Rightarrow (x_1).(\nu\, c)\langle v_1,\, c\rangle.c\Rightarrow (x_2).v_2$$

All the services in the new system have type $!.?.\textsf{end}$ or $?.!.\textsf{end}$, thus they can be implemented as request-responses or solicit-responses. The correctness of the transformation, when done on a closed system, is proved by applying Theorem 3.

6. Concluding remarks

SSCC is a typed language aiming at flexibly describing *services*, *conversations*, and *orchestration*, with a restricted set of constructors. The expressiveness of the language is witnessed by the simple implementation of all workflow patterns in [vdAtHKB03] (except for the ones that require process termination) and by the examples in Sects. 2.1 and 3.2. We have shown instead in Sect. 5 how to exploit formal techniques to define correct program transformations relating different styles of programming used in the field of service-oriented systems, namely object-oriented, session-based and request/request-response based. This allows to exploit the different techniques available in each field, and still get a system implemented using the desired technology. In addition to that, we have illustrated the expressiveness of SSCC also in a field like object-oriented programming, for which it was not conceived. We have further demonstrated the benefit of working with sequential sessions, where more powerful transformations are available.

As future directions of work, we plan to further study the behavioural theory of SSCC, and to consider other analysis techniques. For the former, we want to investigate the relationships between contextual equivalence and bisimilarity, to look for up-to techniques for bisimilarity, and to try to extend the axioms in Sect. 4.3 to a complete axiomatization. Concerning the latter, we are interested in more refined techniques for proving service availability (e.g., linearity of service invocation and definition) and in proofs of deadlock freedom for large classes of protocols.

Also, we intend to investigate the possibility of extending the session breaking transformation to larger classes of systems. We are aware of the fact that parallel communications make the agreement between the client and the server on where to change session more difficult. A promising approach to avoid this problem is to perform a preliminary transformation turning arbitrary sessions into sequential ones.

Another thread for future development concerns the definition of a *compensation* mechanism to recover from failures, and the study of its behavioural theory.

Acknowledgements

This work was partially supported by the Fundaçã para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) through the CITI and the LASIGE strategic projects PEst–OE/EEI/UI0527/ 2011 and PEst–OE/EEI/UI0408/2011 and by the projects PTDC/EIA-CCO/117513/2010 and PTDC/EIA-CCO/122547/ 2010.

A. Subject reduction

Lemma 5 *For each session name r and each process P, in P there are at most two sessions with name r. If there are exactly two such sessions, then they are one client session $r \lhd P'$ and one server session $r \rhd P''$, they are both in the scope of a binder for r and they are not nested. The three occurrences of r (client, server, and binder) are the only occurrences of r in P.*

Proof By induction on the length of the computation creating P. The thesis is true for computations of length 0 (sessions do not appear in the static syntax). Sessions are created only by service definitions and service invocations. When a session is created its name is bound, thus different from other session names. Because of that, different service invocations/definitions cannot create sessions with the same name. We have the creation of two sessions with the same name r only when a service invocation on r interacts with the corresponding service definition. In this case the two created sessions are a client session and a server session, they are not nested, and they are both in the scope of a restriction on r. Since r is bound, no other occurrences of r are created in the rest of the computation. □

Lemma 6 **(Substitution lemma).** *If $\Gamma, x\colon T' \vdash P\colon (U, T)$ and $\Gamma \vdash v\colon T'$ then $\Gamma \vdash P[{}^{v}/_{x}]\colon (U, T)$.*

Proof By induction on the typing proof. All the cases are simple. □

Lemma 7 *If $\Gamma \vdash P\colon (\textsf{end}, T)$ then P has no transitions of the form $P \xrightarrow{\mu} P'$ with $\mu \in \{\uparrow v, \downarrow v, (v) \uparrow v\}$.*

Proof The only way to obtain such transitions is to have processes of the form $\mathcal{C}[\![v.P]\!]$ or $\mathcal{C}[\![(x)P]\!]$ where $\mathcal{C}[\![-]\!]$ is composed only by streams, parallel compositions, and restrictions. Let us consider four cases according to the top-level operator in $\mathcal{C}[\![]\!]$.

In the base case we have to use rule T-SENDor T-RECEIVE. These rules do not allow $(\textsf{end}, T)$ as resulting type.

In the case of stream we have to use rule T-STREAM-R or T-STREAM-L. We consider just the first case, the second being symmetric. The stream has type $(\textsf{end}, T)$ only if the second argument has the same type. Since also the first argument has type $(\textsf{end}, T')$ we know by induction that neither of the arguments can do the communication transitions, thus P cannot do them too.

In the case of parallel composition we have to use rule T-PAR-R or T-PAR-L. We consider the first case, the second one being symmetric. Process P has type $(\textsf{end}, T)$ only if the second argument has the same type. Since also the first argument has type $(\textsf{end}, T)$ we know by induction that neither of the arguments can do the communication transitions, thus P cannot do them too.

In the case of restriction we have to use rule T-RES. Process P has type $(\textsf{end}, T)$ only if the restricted process has the same type. We know by induction that the argument cannot do the communication transitions, thus P cannot do them too. □

Lemma 8 **(Weakening).** *If $\Gamma \vdash P\colon (U, T)$ and $n \notin \mathrm{fn}(P)$ then $\Gamma, n\colon T' \vdash P\colon (U, T)$, for all T'.*

Proof Simple, by induction on the derivation of the typing judgment. □

Lemma 9 **(Strengthening).** *If $\Gamma, n : T' \vdash P\colon (U, T)$ and $n \notin \mathrm{fn}(P)$ then $\Gamma \vdash P\colon (U, T)$.*

Proof Simple, by induction on the derivation of the typing judgment. □

Lemma 10 **(Subject congruence).** *If $\Gamma \vdash P\colon (U, T)$ and $P \equiv Q$ then $\Gamma \vdash Q\colon (U, T)$.*

Proof It is enough to show that structural congruent terms can be given the same type using the same assumptions. It is enough to show this for the LHS and the RHS for each structural congruence rule, then the thesis follows by induction (the congruence axioms are simple). All the cases but the one for recursion are easy. We show just this case. Suppose that $\Gamma \vdash \mathsf{rec}\ X.P : (U, T)$. Then, by hypothesis $\Gamma, X : (U, T) \vdash P : (U, T)$. By structural induction on P we can prove that if $\Gamma, X : (U, T) \vdash P : (U, T)$ then $\Gamma \vdash P[{}^{\mathsf{rec}\ X.P}/_{X}] : (U, T)$. This holds for the case of $P = X$ and is preserved by all the contexts (notice in fact that the assumptions about different occurrences of the same variable are compatible). The proof is similar in the opposite direction. □

Let $\Gamma[{}^{[U']}/_{r}]$ denote the substitution on Γ of $[U']$ for $\Gamma(r)$.

Theorem 4 *Let P be a process such that $\Gamma \vdash P : (U, T)$. Then:*

- *if $P \xrightarrow{\uparrow v} P'$ then $U = !T'.U'$, $\Gamma \vdash v : T'$ and $\Gamma \vdash P' : (U', T)$;*
- *if $P \xrightarrow{(v)\uparrow v} P'$ then $U = !T'.U'$ and $\Gamma, v : T' \vdash P' : (U', T)$;*
- *if $P \xrightarrow{\downarrow v} P'$ then $U = ?T'.U'$ and $\Gamma, v : T' \vdash P' : (U', T)$;*
- *if $P \xrightarrow{a \Leftarrow (r)} P'$ then $\Gamma \vdash a : [U']$ and $\Gamma, r : [U'] \vdash P' : (U, T)$;*
- *if $P \xrightarrow{a \Rightarrow (r)} P'$ then $\Gamma \vdash a : [U']$ and $\Gamma, r : [U'] \vdash P' : (U, T)$;*
- *if $P \xrightarrow{\Uparrow v} P'$ then $\Gamma \vdash v : T$ and $\Gamma \vdash P' : (U, T)$;*
- *if $P \xrightarrow{(v)\Uparrow v} P'$ then $\Gamma, v : T \vdash P' : (U, T)$;*
- *if $P \xrightarrow{f \Downarrow v} P'$ then $\Gamma \vdash f : \langle T \rangle$ and $\Gamma, v : T \vdash P' : (U, T)$;*
- *if $P \xrightarrow{r \rhd \uparrow v} P'$ then $\Gamma \vdash r : [!T'.U']$, $\Gamma \vdash v : T'$ and $\Gamma[{}^{[U']}/_{r}] \vdash P' : (U, T)$;*
- *if $P \xrightarrow{(v) r \rhd \uparrow v} P'$ then $\Gamma \vdash r : [!T'.U']$ and $\Gamma[{}^{[U']}/_{r}], v : T' \vdash P' : (U, T)$;*
- *if $P \xrightarrow{r \rhd \downarrow v} P'$ then $\Gamma \vdash r : [?T'.U']$ and $\Gamma[{}^{[U']}/_{r}], v : T' \vdash P' : (U, T)$;*
- *if $P \xrightarrow{r \lhd \uparrow v} P'$ then $\Gamma \vdash r : [?T'.U']$, $\Gamma \vdash v : T'$ and $\Gamma[{}^{[U']}/_{r}] \vdash P' : (U, T)$;*
- *if $P \xrightarrow{(v) r \lhd \uparrow v} P'$ then $\Gamma \vdash r : [?T'.U']$ and $\Gamma[{}^{[U']}/_{r}], v : T' \vdash P' : (U, T)$;*
- *if $P \xrightarrow{r \lhd \downarrow v} P'$ then $\Gamma \vdash r : [!T'.U']$ and $\Gamma[{}^{[U']}/_{r}], v : T' \vdash P' : (U, T)$;*
- *if $P \xrightarrow{r\tau} P'$ then $\Gamma \vdash r : [!T'.U']$ or $\Gamma \vdash r : [?T'.U']$ and $\Gamma[{}^{[U']}/_{r}] \vdash P' : (U, T)$;*
- *if $P \xrightarrow{\tau} P'$ then $\Gamma \vdash P' : (U, T)$.*

Proof The proof is by induction on the derivation of the transition. We perform a case analysis on the last rule used in the derivation.

L-SEND: P has the form $v.P'$. This can be typed only using rule T-SEND and this requires $U = !T'.U'$, $\Gamma \vdash P' : (U', T)$ and $\Gamma \vdash v : T'$. This is exactly as desired.

L-RECEIVE: P has the form $(x)P''$ and $P' = P''[{}^{v}/_{x}]$. P can be typed only using rule T-RECEIVE and this requires $U = ?T'.U'$ and $\Gamma, x : T' \vdash P' : (U', T)$. Thanks to Lemma 6 we also have $\Gamma, v : T' \vdash P'[{}^{v}/_{x}] : (U', T)$.

L-CALL: P has the form $a \Leftarrow P''$ and $P' = r \lhd P''$. P can be typed only using rule T-CALL and this requires $U = \mathbf{end}$, $\Gamma \vdash P'' : (U', T)$ and $\Gamma \vdash a : [\overline{U'}]$. Using rule T-SESS-C (and thanks to Lemma 8) one can derive $\Gamma, r : [\overline{U'}] \vdash r \lhd P'' : (\mathbf{end}, T)$.

L-DEF: P has the form $a \Rightarrow P''$ and $P' = r \rhd P''$. P can be typed only using rule T-DEF and this requires $U = \mathbf{end}$, $\Gamma \vdash P'' : (U', T)$ and $\Gamma \vdash a : [U']$. Using rule T-SESS-S (and thanks to Lemma 8) one can derive $\Gamma, r : [U'] \vdash r \rhd P'' : (\mathbf{end}, T)$.

L-FEED: P has the form $\mathbf{feed}\ v.P'$. This can be typed only using rule T-FEED and this requires $\Gamma \vdash P' : (U, T)$ and $\Gamma \vdash v : T$. This is exactly as required.

L-READ: P has the form $f(x).P''$ and $P' = P''[{}^{v}/_{x}]$. P can be typed only using rule T-READ and this requires $\Gamma, x : T' \vdash P'' : (U, T)$ and $\Gamma \vdash f : \langle T' \rangle$. From Lemma 6 we have $\Gamma, v : T' \vdash P''[{}^{v}/_{x}] : (U, T)$ as required.

L-STREAM-PASS-P: P has the form $\mathsf{stream}\ P''\ \mathsf{as}\ f = \vec{v}\ \mathsf{in}\ Q$ with $P'' \xrightarrow{\mu} P'''$ and we have $P' = \mathsf{stream}\ P'''\ \mathsf{as}\ f = \vec{v}\ \mathsf{in}\ Q$. There are two cases according to the last rule used to type P. We consider rule T-STREAM-R first and rule T-STREAM-L later. Thanks to Lemma 7, $\mu \notin \{\uparrow v, \downarrow v, (v) \uparrow v\}$. Also, $\mu \notin \{\Uparrow v, (v) \Uparrow v\}$. By hypothesis all the assumptions on f, $\vec{v}$ and Q are satisfied. By induction hypothesis in all the cases but $r \bowtie\uparrow v$, $(v)r \bowtie\uparrow v$, $r \bowtie\downarrow v$ and $r\tau$ we have that $\Gamma' \vdash P''' : (\mathsf{end},\ T)$ for some extension Γ' of Γ. Thanks to Lemma 8, Γ' can be used to derive $\Gamma' \vdash P' : (U,\ T)$ as required. Notice also that the assumptions on Γ' are satisfied by induction hypothesis since the label is unchanged. For the other cases the problem is that the assumption about r is changed. However, thanks to Lemma 5 there are two cases. If there is just one occurrence of r, thus the assumption is never used outside P''', Lemma 9 can be used to drop the old assumption and Lemma 8 to add the new one, and the thesis follows. If there are three occurrences two should be in opposite session constructs and the third in a restriction binding them. The only label of these that can cross the restriction is $r\tau$, thus no occurrence of r can be in Q, since otherwise we can not obtain this label. Thus r is not used in Q and we can derive $\Gamma' \vdash Q : (U,\ T)$ as required, using again lemmas 9 and 8. Thus we can also derive $\Gamma' \vdash P' : (U,\ T)$ and the thesis follows.
Let us consider the second case. Notice that $\mu \notin \{\Uparrow v, (v) \Uparrow v\}$. Now both U and μ are preserved from the premise, thus in most of the cases the thesis follows immediately from the induction premise (when a new assumption is needed in Γ, such as in extrusions, Lemma 8 can be used, and the compatibility of the new assumption is guaranteed by the side condition on bound names of the typing rule). The only tricky cases concern labels $r \bowtie\uparrow v$, $(v)r \bowtie\uparrow v$, $r \bowtie\downarrow v$ and $r\tau$, but the same reasoning above applies. The thesis follows.

L-STREAM-PASS-Q: P has the form $\mathsf{stream}\ P''\ \mathsf{as}\ f = \vec{v}\ \mathsf{in}\ Q$ with $Q \xrightarrow{\mu} Q'$ and we have $P' = \mathsf{stream}\ P''\ \mathsf{as}\ f = \vec{v}\ \mathsf{in}\ Q'$. By hypothesis all the assumptions on P, f and $\vec{v}$ are satisfied. Also, $\Gamma, f : \langle T' \rangle \vdash Q : (U,\ T)$. By induction hypothesis $\Gamma', f : \langle T' \rangle \vdash Q' : (U',\ T)$ where Γ' and U' are defined by the statement of the theorem. Notice that Γ' verifies all the assumptions of rule T-STREAM-L (resp. T-STREAM-R) since it is either an extension of Γ (and in this case Lemma 8 can be used), or it changes the assumption about some session r, and in this case the same reasoning done for rule L-STREAM-PASS-P can be used. Thus one can use rule T-STREAM-L (resp. T-STREAM-R) to derive $\Gamma' \vdash P' : (U',\ T)$ as required.

L-STREAM-FEED: P has the form $\mathsf{stream}\ P''\ \mathsf{as}\ f = \vec{w}\ \mathsf{in}\ Q$ with $P'' \xrightarrow{\Uparrow v} P'''$ and we have $P' = \mathsf{stream}\ P'''\ \mathsf{as}\ f = v :: \vec{w}\ \mathsf{in}\ Q$. There are two cases corresponding to rules T-STREAM-R and T-STREAM-L. We consider the first one, the second being similar. By hypothesis $\Gamma \vdash P'' : (\mathsf{end},\ T')$, $\Gamma, f : \langle T' \rangle \vdash Q : (U,\ T)$ and $w' \in \mathrm{Set}(\vec{w}) \Rightarrow \Gamma \vdash w' :\ T'$. By induction hypothesis $\Gamma \vdash v :\ T'$ and $\Gamma \vdash P''' : (\mathsf{end},\ T')$. Thus using rule T-STREAM-R we can prove $\Gamma \vdash P' : (U,\ T)$ (notice, in particular, that the assumption about $v :: \vec{w}$ can be proved from the assumptions about v and $\vec{w}$).

L-STREAM-CONS: P has the form $\mathsf{stream}\ P''\ \mathsf{as}\ f = \vec{w} :: v\ \mathsf{in}\ Q$ with $Q \xrightarrow{f \Downarrow v} Q'$ and we have $P' = \mathsf{stream}\ P''\ \mathsf{as}\ f = \vec{w}\ \mathsf{in}\ Q'$. There are two cases corresponding to rules T-STREAM-R and T-STREAM-L. We consider the first one, the second being symmetric. By hypothesis $\Gamma \vdash P'' : (\mathsf{end},\ T')$, $\Gamma, f : \langle T' \rangle \vdash Q : (U,\ T)$ and $w' \in \mathrm{Set}(\vec{w} :: v) \Rightarrow \Gamma \vdash w' :\ T'$. By induction hypothesis $\Gamma, f : \langle T' \rangle, v :\ T' \vdash Q' : (U,\ T)$. Since $\Gamma, v :\ T'$ is an extension of Γ we can use it (thanks to Lemma 8) in all the premises of rule T-STREAM-R and finally derive $\Gamma, v :\ T' \vdash P' : (U,\ T)$.

L-PAR: the reasoning is as for rule L-STREAM-PASS-P, but there is no stream here.

L-SESS-VAL: we consider just the cases for $\triangleleft$, the other being simpler. P has the form $r \triangleleft P''$. By hypothesis $U = \mathsf{end}$, $\Gamma \vdash P'' : (U',\ T)$ and $\Gamma \vdash r : [\overline{U'}]$. Let us consider the case $P'' \xrightarrow{\uparrow v} P'''$ before. This implies $P' = r \triangleleft P'''$. By induction hypothesis $U' = !T'.U''$, $\Gamma \vdash v :\ T'$ and $\Gamma \vdash P''' : (U'',\ T)$. Using rule T-SESS-C we can prove $\Gamma[^{[\overline{U''}]}/_{r}] \vdash r \triangleleft P''' : (\mathsf{end},\ T)$ as required since this is the only place where the assumption about r is used inside the term thanks to Lemma 5, thus it can be changed using lemmas 9 and 8. Let us now consider the case $P'' \xrightarrow{\downarrow v} P'''$. Again $P' = r \triangleleft P'''$. By induction hypothesis $U' = ?T'.U''$, $\Gamma, v :\ T' \vdash P''' : (U'',\ T)$. Using rule T-SESS-C we can prove $\Gamma[^{[\overline{U''}]}/_{r}], v :\ T' \vdash r \triangleleft P'''$ as required, since this is the only place where the assumption about r is used inside the term thanks to Lemma 5, thus it can be changed using lemmas 9 and 8.

L-SESS-PASS: we consider just the cases for $\triangleleft$, the others being simpler. P has the form $r \triangleleft P''$ with $P'' \xrightarrow{\mu} P'''$ and $P' = r \triangleleft P'''$. By hypothesis $U = \mathsf{end}$, $\Gamma \vdash P'' : (U',\ T)$ and $\Gamma \vdash r : [\overline{U'}]$. Notice that $\mu \neq\updownarrow v$. Thus for all the cases but session communication labels we have $\Gamma' \vdash P''' : (U',\ T)$ for some extension Γ' of Γ. In the case of session communication labels instead the assumption about r' is changed from Γ to Γ'. Notice that thanks to Lemma 5 $r \neq r'$, thus in both the cases we can use rule T-SESS-C to derive $\Gamma' \vdash r \triangleleft P''' : (\mathsf{end},\ T)$ as required, since the label of the new transition is equal to the label of the premise, thus the assumptions on Γ' coincide.

L-SESS-COM-STREAM: P has the form **stream** P'' **as** $f = \vec{w}$ **in** Q with $P'' \xrightarrow{r \rhd \uparrow v} P'''$, $Q \xrightarrow{r \lhd \downarrow v} Q'$ and $P' =$ **stream** P''' **as** $f = \vec{w}$ **in** Q' (the other cases are similar). There are two cases corresponding to rules T-STREAM-R and T-STREAM-L. We consider just the first one, the second being similar. By hypothesis $\Gamma \vdash P\colon (\textbf{end}, T')$ and $\Gamma, f\colon \langle T' \rangle \vdash Q\colon (U, T)$. By induction hypothesis on the first transition $\Gamma \vdash r\colon [!T''.U']$, $\Gamma \vdash v\colon T''$ and $\Gamma[{}^{[U']}/_{r}] \vdash P'''\colon (U, T)$. From the second transition we have a redundant hypothesis on r and $\Gamma[{}^{[U']}/_{r}], f\colon \langle T' \rangle, v\colon T' \vdash Q'\colon (U, T)$. Notice that $\Gamma[{}^{[U']}/_{r}], f\colon \langle T' \rangle, v\colon T' = \Gamma[{}^{[U']}/_{r}], f\colon \langle T' \rangle$ since $\Gamma[{}^{[U']}/_{r}], f\colon \langle T' \rangle \vdash v\colon T'$. Thus we can apply rule T-STREAM-R to derive $\Gamma[{}^{[U']}/_{r}] \vdash P'\colon (U, T)$ as required.

L-SERV-COM-STREAM: P has the form **stream** P'' **as** $f = \vec{w}$ **in** Q with $P'' \xrightarrow{a \Rightarrow (r)} P'''$, $Q \xrightarrow{a \Leftarrow (r)} Q'$ and $P' = (\nu\, r)$**stream** P''' **as** $f = \vec{w}$ **in** Q' (the symmetric case is similar). There are two cases corresponding to rules T-STREAM-R and T-STREAM-L. We consider just the first one, the second being similar. By hypothesis $\Gamma \vdash P''\colon (\textbf{end}, T')$ and $\Gamma, f\colon \langle T' \rangle \vdash Q\colon (U, T)$. By induction hypothesis (on both the transitions) $\Gamma \vdash a\colon [U']$ and $\Gamma, r\colon [U'] \vdash P''\colon (\textbf{end}, T')$ and $\Gamma, f\colon \langle T' \rangle, r\colon [U'] \vdash Q'\colon (U, T)$. Using rule T-STREAM-R we can derive the judgment $\Gamma, r\colon [U'] \vdash$ **stream** P''' **as** $f = \vec{w}$ **in** $Q'\colon (U, T)$. Then we can use rule T-RES to derive $\Gamma \vdash P'\colon (U, T)$ as desired.

L-SESS-COM-PAR: the reasoning is as for rule L-SESS-COM-STREAM, but there is no stream here.

L-SERV-COM-PAR: the reasoning is as for rule L-SERV-COM-STREAM, but there is no stream here.

L-RES: P has the form $(\nu\, n)P''$ with $P'' \xrightarrow{\mu} P'''$ and $P' = (\nu\, n)P'''$. By hypothesis $\Gamma, n\colon T' \vdash P''\colon (U, T)$ for some T'. By induction hypothesis $\Gamma', n\colon T' \vdash P'''\colon (U', T)$ where Γ' and U' are as defined by the statement of the theorem. Thus we can apply rule T-RES to derive $\Gamma' \vdash (\nu\, n)P'''\colon (U', T)$ since the label is unchanged thus Γ' and U' are as before.

L-EXTR: P has the form $(\nu\, a)P''$ with $P'' \xrightarrow{\mu} P'$. By hypothesis $\Gamma, a\colon T' \vdash P''\colon (U, T)$. Thanks to the induction hypothesis $\Gamma', a\colon T' \vdash P'\colon (U', T)$ where Γ' and U' are as described in the statement of the theorem. This is exactly as required, given the different requirements between each action and the corresponding extruding action.

L-SESS-RES: P has the form $(\nu\, r)P''$ with $P'' \xrightarrow{r\tau} P'''$ and $P' = (\nu\, r)P'''$. By hypothesis $\Gamma, r\colon [U'] \vdash P''\colon (U, T)$ (the type of r should be a protocol since r is a session). By the induction hypothesis $\Gamma, r\colon [U''] \vdash P'''\colon (U, T)$. Then we can use rule T-RES to derive $\Gamma \vdash P'''\colon (U, T)$ as required.

L-STRUCT: By Lemma 10. □

Proof of Theorem 1 (Subject Reduction) The thesis follows from Theorem 4 and the characterization of reductions as transitions with labels τ given in Theorem 2. □

Proof of Theorem 2 (Type Safety) The proofs of all the cases are by contradiction. We suppose that such a subterm exists and we show that it is not typable. We consider the two different cases:

Protocol: Let us consider the first case. Here $v.P$ and $u.Q$ have types of the form $([!T.U], T'')$ and $([!T'.U'], T''')$ respectively. One can prove by structural induction on the context that the protocol part of the type is preserved (only the session construct can change it, but the side condition forbids sessions around the hole). Thus the two session constructs require $r\colon [!T.U]$ and $r\colon [?T'.\overline{U'}]$ (supposing that the first one is a server session, the symmetric otherwise). Since $\mathcal{D}[\![,]\!]$ does not bind r the assumptions are preserved, and at top level they should agree since the same Γ is used to type the two sides of parallel composition or stream. This is not the case and we have the required contradiction. The other cases are similar, with just **end** protocol for **0** and $([?T.U], T'')$ for input.

Sequentiality: In all the cases the two terms inserted into the double context have non **end** protocol. The property is preserved by the context (since there are no sessions around the hole). At top level we have two non **end** protocols, but the rules for parallel composition and stream can not be applied because of this. Since no other rules can type a parallel composition or a stream we have the desired contradiction. □

B. Workflow patterns in SSCC

This section completes Sect. 3.1, by modelling in SSCC the remaining workflow patterns from [vdAtHKB03].

WP2: Parallel split

“A point in the workflow process where a single thread of control splits into multiple threads of control which can be executed in parallel, thus allowing activities to be executed simultaneously or in any order. Example: after registering an insurance claim, two parallel subprocesses are triggered: one for checking the policy of the customer and one for assessing the actual damage.”

Parallel composition is built-in. The same in SCC and in Orc.

WP5: Simple merge

“A point in the workflow process where two or more alternative branches come together without synchronization. It is an assumption of this pattern that none of the alternative branches is ever executed in parallel. Example: after the payment is received or the credit is granted, the car is delivered to the customer.”

```
merge :: Bool → (ε → T) →...→ Bool → (ε → T) → Unit
merge ⤖ (b1)(a1)...(bn)(an)
         (if b1 then call a1 | ... | if bn then call an) >1> unit
```

This is more in line with van der Aalst specification than the corresponding model in Orc, since SSCC is able to model the fact that only some of the activities are activated. Notice that, given the assumptions, replacing $>^{\mathbf{1}}>$ with $>>$ will not change the behaviour.

WP6: Multi-choice

“A point in the workflow process where, based on a decision or workflow control data, a number of branches are chosen. Example: after executing the activity *evaluate_damage*, the activity *contact_fire_department* or the activity *contact_insurance_company* is executed. At least one of these activities is executed. However, it is also possible that both need to be executed.”

Notice that the code resulting from the application of this WP is not an activity, since it may produce more than one result. The more natural implementation of this WP in SSCC is:

```
multiChoice ⤖ (b1)(a1)...(bn)(an)
               (if b1 then call a1 >1 x > x | ... |
                if bn then call an >1 x > x)
```

A similar implementation is possible in Orc and in SCC.

However, this service is not typable, since all the results are sent in parallel inside the same session. To solve this problem we can sequentialize them by writing:

```
multiChoice :: [?Bool.?(ε → T). ... .?Bool.?(ε → T).!T. ... !T.end]
multiChoice ⤖ (b1)(a1)...(bn)(an)
                stream
               (if b1 then call a1 | ... | if bn then call an)
                as f in
                f(y1).y1 ....f(yn).yn
```

Note that the service may not produce the number of results specified by its type, since the stream may not supply enough elements.

WP7: synchronizing merge

“A point in the workflow process where multiple paths converge into one single thread. If more than one path is taken, synchronisation of the active threads needs to take place. If only one path is taken, the alternative branches should reconverge without synchronization. It is an assumption of this pattern that a branch that has already been activated, cannot be activated again while the merge is still waiting for other branches to complete. Example: extending the example of WP6 (Multi-choice), after either or both of the activities *contact_fire_department* and *contact_insurance_company* have been completed (depending on whether they were executed at all), the activity *submit report* needs to be performed (exactly once).”

```
syncMerge::  (ε → Bool) → (ε → Unit) →...→
             (ε → Bool) → (ε → Unit) → Unit
syncMerge *⇒ (b_1)(a_1)...(b_n)(a_n)
             call sync(ifSignal_b_1_a_1 ,..., ifSignal_b_n_a_n) >^n > unit
ifSignal_b_i_a_i :: ε → Unit
ifSignal_b_i_a_i ⇒ IfSignal(b_i , call a_i >^1 x > x)
```

where

```
IfSignal(b,P) = if b then P else unit
```

Essentially, we reuse WP3 (Synchronization) on services ifSignal_b_i_a_i. Service ifSignal_b_i_a_i invokes a_i and gives back its result if b_i is true, it immediately returns **unit** otherwise. When all the results from the invoked services have been collected, a **unit** value is returned. It can be used to trigger the final activity.

Similar to Orc and SCC.

WP8: Multi-merge

“A point in a workflow process where two or more branches reconverge without synchronization. If more than one branch gets activated, possibly concurrently, the activity following the merge is started *for every activation of every incoming branch*. Example: two activities *audit_application* and *process_application* running in parallel, which should both be followed by an activity *close_case*.”

```
merge  :: [?Bool.?(ε → T).  ...  .?Bool.?(ε → T).?(ε → T1).!T1.....!T1.end]
merge *⇒ (b_1)(a_1)...(b_n)(a_n)(c)
         stream
         (if b_1 then call a_1 | ... | if b_n then call a_n)
         as f in
            stream
            f(y_1).call c(y_1)....f(y_n).call c(y_n)
            as g in
            g(z_1).z_1.....g(z_n).z_n
```

This is not an activity, since it provides multiple replies. Similar to the Orc implementation. In SCC one can use the technique of WP1 (Sequence). Notice also that now the behaviour of the synchronization is the expected one (one instance is launched for each value).

WP11: Implicit termination

“A given subprocess should be terminated when there is nothing else to be done. In other words, there are no active activities in the workflow and no other activity can be made active (and at the same time the workflow is not in deadlock).”

This pattern is not meaningful in process calculi. In fact, in calculi the standard behaviour is that processes naturally terminate when they have finished their activity, not when a final state is reached by one of their components. This is the case for both SCC and SSCC, while this is not the case in workflow managers.

WP12: Multiple instances without synchronization

"Within the context of a single case (i.e., workflow instance) multiple instances of an activity can be created, i.e., there is a facility to spawn new threads of control. Each of these threads of control is independent of other threads. Moreover, there is no need to synchronise these threads. Example: a customer ordering a book from an electronic bookstore such as Amazon may order multiple books at the same time. Many of the activities (e.g., billing, updating customer records) occur at the level of the order. However, within the order, multiple instances need to be created to handle the activities related to one individual book (e.g., update stock levels, shipment). If the activities at the book level do not need to be synchronized, this pattern can be used."

Multiple instances of the same service can be executed concurrently without any particular problem. An unbounded number of instances can be created by persistent services. The same in SCC.

WP13: Multiple instances with a priory design time knowledge

"For one process instance an activity is enabled multiple times. The number of instances of a given activity for a given process instance is known at design time. Once all instances are completed, some other activity needs to be started. Example: the requisition of hazardous material requires three different authorizations."

```
sync_n :: (ε → T) → Unit
sync_n *⇒ (a) call sync (a,...,a) >^1 x > x
```

Since the number of instances (calls to) of service a is known to be n, it is enough to pass n arguments to service sync (cfr. WP3: Synchronization).

A similar approach can be used in Orc. A possible implementation in SCC is:

```
sync_n ⇒ (a) sync {a. ... .a.(x) return x} ⇐ a
```

where one instance of a is passed as invocation parameter, and the orher n−1 instances are passed inside the session protocol.

WP14: Multiple instances with a priory run-time knowledge

"For one case, an activity is enabled multiple times. The number of instances of a given activity for a given case varies and may depend on characteristics of the case or availability of resources, but is known at some stage during run-time, before the instances of that activity have to be created. Once all instances are completed some other activity needs to be started. Example: when booking a trip, the activity *book_flight* is executed multiple times if the trip involves multiple flights. Once all bookings are made, the invoice is to be sent to the client."

We treat this case as a particular case of WP15. See below for the discussion.

WP15: Multiple instances without a priory run-time knowledge

"For one case an activity is enabled multiple times. The number of instances of a given activity for a given case is not known during design time, nor is it known at any stage during run-time, before the instances of that activity have to be created. Once all instances are completed, some other activity needs to be started. The difference with WP14 is that even while some of the instances are being executed or already completed, new ones can be created. Example: for the processing of an insurance claim, zero or more eyewitness reports should be handled. The number of eyewitness reports may vary. Even when processing eyewitness reports for a given insurance claim, new eyewitnesses may surface and the number of instances may change."

Invoke service a as long as service c replies true. Instances are executed in parallel: the first instance is launched in parallel with parloop_c_a. Termination of an instance is checked together with the termination of the parloop launched together.

```
parloop_c_a :: ε → Unit
parloop_c_a *⇒ call c >^1 b >
                 IfSignal(b, call sync(a,parloop_c_a)) >^1 x > x
```

For simplicity we have chosen a loop service specific for c and a. To write a generic loop service that accepts two parameters (c and a) we have to customize sync to invoke services with parameters. We leave the exercise to the reader. Similar implementations can be done in Orc and in SCC.

As far as WP14 is concerned, the main choice is how to represent the run-time knowledge about the required number of instances to be executed, i.e., how to represent state. Possibilities include taking advantage of the number of values in a stream, of the number of available instances of a service, or of the number of values in a session protocol.

WP16: Deferred choice

"A point in the workflow process where one of several branches is chosen. In contrast to the XOR-split, the choice is not made explicitly (e.g., based on data or on a decision), but several alternatives are offered to the environment. However, in contrast to the AND-split, only one of the alternatives is executed. This means that once the environment activates one of the branches, the other alternative branches are withdrawn. It is important to note that the choice is delayed until the processing in one of the alternative branches is actually started, i.e., the moment of choice is as late as possible. Example: after receiving products there are two ways to transport them to the department. The selection is based on the availability of the corresponding resources. Therefore, the choice is deferred until a resource is available."

Requires a means to kill unwanted computations, such as the **where** operator in Orc.

C. On the labelled transition system of SSCC

C.1. The Harmony Lemma

In the proof below we use $\overline{a}$ to denote any label in $\{\uparrow a, r \bowtie\!\uparrow a, \Uparrow a\}$.

Theorem 5 (Harmony Lemma). *Let P and Q be processes with $P \equiv Q$. If $P \xrightarrow{\alpha} P'$, then $Q \xrightarrow{\alpha} Q'$ with $P' \equiv Q'$, and vice-versa.*

Proof By induction on the proof that $P \equiv Q$.

- *Equivalence relation*
  - Reflexivity. Immediate, taking Q' to be P'.
  - Symmetry. Immediate consequence of the induction hypothesis, since the thesis of the theorem is symmetric.
  - Transitivity. Assume $P \equiv Q$ because $P \equiv R$ and $R \equiv Q$, and suppose that $P \xrightarrow{\alpha} P'$. By induction hypothesis, $R \xrightarrow{\alpha} R'$ with $P' \equiv R'$; hence, again by induction hypothesis, $Q \xrightarrow{\alpha} Q'$ with $R' \equiv Q'$. $P' \equiv Q'$ follows by transitivity of $\equiv$.
- *Congruence properties*
  - Parallel composition. Suppose $P \equiv Q$. For each of the possible transitions of $P \mid R$, it is straightforward to verify that $Q \mid R$ can simulate them, possibly using the induction hypothesis; similarly, $R \mid Q$ can simulate $R \mid P$. Notice that the side conditions in the transition rules always hold since structurally congruent processes have the same free names.
  - Composition with stream. Analogous.
  - Name restriction. Suppose $P \equiv Q$ and let a be a name. For each of the possible transitions of $(\nu a)P$, it is easy to check that $(\nu a)Q$ can simulate them, possibly using the induction hypothesis.
  - Session input/output. Straightforward, observing (for input) that structural congruence is closed under substitution.
  - Stream input/output. Analogous.
  - Service definition/invocation. Straightforward.

- *Monoid structure*
  - Unit. Let Q be $P \mid 0$. If $P \xrightarrow{\alpha} P'$, then by rule L-PAR also $P \mid 0 \xrightarrow{\alpha} P' \mid 0$, since 0 has no free names. Also, $P' \mid 0$ is congruent to P'. Reciprocally, if $P \mid 0 \xrightarrow{\alpha} P$’, then the only rule that can have been applied is L-PAR (since $0 \not\rightarrow$), whence P' is $P'' \mid 0$ with $P \xrightarrow{\alpha} P''$.
  - Commutativity. Assume P is $R \mid S$ and Q is $S \mid R$. Take any proof of $R \mid S \xrightarrow{\alpha} T$ and replace occurrences of L-PAR by L-PAR’, of L-SESS-COM-PAR by L-SESS-COM-PAR’, of L-SERV-COM-PAR by L-SERV-COM-PAR’ and vice-versa; it is straightforward to verify that this yields a proof that $S \mid R \xrightarrow{\alpha} T'$ with $T \equiv T'$. The converse is analogous.
  - Associativity. Let P be $R \mid (S \mid T)$ and Q be $(R \mid S) \mid T$. Suppose that $P \xrightarrow{\alpha} P'$; there are six rules that can be used to infer this transition. For simplicity, in the proofs below we omit side conditions related to bound names, since it is simple to verify that they always follow from the assumptions.
    - · L-PAR: then $R \xrightarrow{\alpha} R'$ and P' is $R' \mid (S \mid T)$. The proof below shows that $(R \mid S) \mid T \xrightarrow{\alpha} (R' \mid S) \mid T$, which establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{\alpha} R'}{R \mid S \xrightarrow{\alpha} R' \mid S}\ \text{L-PAR}}{(R \mid S) \mid T \xrightarrow{\alpha} (R' \mid S) \mid T}\ \text{L-PAR}$$

    - · L-PAR’: then $S \mid T \xrightarrow{\alpha} U$; there are six sub-cases, according to the rule used to derive this transition.
      - ∗ The rule applied is L-PAR, so $S \xrightarrow{\alpha} S'$ and U is $S' \mid T$; then the following proof establishes the thesis.

$$\dfrac{\dfrac{S \xrightarrow{\alpha} S'}{R \mid S \xrightarrow{\alpha} R \mid S'}\ \text{L-PAR’}}{(R \mid S) \mid T \xrightarrow{\alpha} (R \mid S') \mid T}\ \text{L-PAR}$$

      - ∗ The rule applied is L-PAR’, so $T \xrightarrow{\alpha} T'$ and U is $S \mid T'$; then the following proof establishes the thesis.

$$\dfrac{T \xrightarrow{\alpha} T'}{(R \mid S) \mid T \xrightarrow{\alpha} (R \mid S) \mid T'}\ \text{L-PAR’}$$

      - ∗ The rule applied is L-SESS-COM-PAR, so $S \xrightarrow{r \bowtie \Updownarrow v} S'$, $T \xrightarrow{r \overline{\bowtie} \overline{\Updownarrow} v} T'$, U is $S' \mid T'$ and α is $r\tau$ for some fresh r; then the following proof establishes the thesis.

$$\dfrac{\dfrac{S \xrightarrow{r \bowtie \Updownarrow v} S'}{R \mid S \xrightarrow{r \bowtie \Updownarrow v} R \mid S'}\ \text{L-PAR’} \qquad T \xrightarrow{r \overline{\bowtie} \overline{\Updownarrow} v} T'}{(R \mid S) \mid T \xrightarrow{r\tau} (R \mid S') \mid T'}\ \text{L-SESS-COM-PAR}$$

      - ∗ The rule applied is L-SERV-COM-PAR, thus we have $S \xrightarrow{a \Leftrightarrow (r)} S'$, $T \xrightarrow{a \overline{\Leftrightarrow} (r)} T'$, U is $(\nu r)(S' \mid T')$ and α is τ; then the following proof establishes the thesis, since $(\nu r)((R \mid S') \mid T') \equiv R \mid ((\nu r)(S' \mid T'))$ as r is not a free name of R.

$$\dfrac{\dfrac{S \xrightarrow{a \Leftrightarrow (r)} S'}{R \mid S \xrightarrow{a \Leftrightarrow (r)} R \mid S'}\ \text{L-PAR’} \qquad T \xrightarrow{a \overline{\Leftrightarrow} (r)} T'}{(R \mid S) \mid T \xrightarrow{\tau} (\nu r)((R \mid S') \mid T')}\ \text{L-SERV-COM-PAR}$$

      - ∗ The rule applied is L-PAR-CLOSE, so $S \xrightarrow{r \bowtie (a) \overline{a}} S'$, $T \xrightarrow{r \overline{\bowtie} \downarrow a} T'$, U is $(\nu a)(S' \mid T')$ and α is $r\tau$. This case is analogous to that of L-SESS-COM-PAR, the extra name restrictions in the resulting processes posing no additional problem.

  - ∗ The rule applied is L-PAR-CLOSE’, so $S \xrightarrow{r\bowtie\downarrow a} S'$, $T \xrightarrow{r\overline{\bowtie}(a)\overline{a}} T'$, U is $(\nu a)(S' \mid T')$ and α is $r\tau$. This case is analogous to the previous one.

- · L-SESS-COM-PAR: then $R \xrightarrow{r\bowtie\updownarrow v} R'$ and $S \mid T \xrightarrow{r\overline{\bowtie}\overline{\updownarrow}v} U$; there are two similar sub-cases, according to whether the last transition is proved via L-PAR or via L-PAR’. Without loss of generality, assume that the former is the case; then the following proof establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{r\bowtie\updownarrow v} R' \quad S \xrightarrow{r\overline{\bowtie}\overline{\updownarrow}v} S'}{R \mid S \xrightarrow{r\tau} R' \mid S'}\ \text{L-SESS-COM-PAR}}{(R \mid S) \mid T \xrightarrow{r\tau} (R' \mid S') \mid T}\ \text{L-PAR'}$$

- · L-SERV-COM-PAR: then $R \xrightarrow{a\Leftrightarrow(r)} R'$ and $S \mid T \xrightarrow{a\overline{\Leftrightarrow}(r)} U$; again there are two similar sub-cases, according to whether the last transition is proved via L-PAR or via L-PAR’. Without loss of generality, assume that the former is the case; then the following proof establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{a\Leftrightarrow(r)} R' \quad S \xrightarrow{a\overline{\Leftrightarrow}(r)} S'}{R \mid S \xrightarrow{\tau} (\nu r)(R' \mid S')}\ \text{L-SERV-COM-PAR}}{(R \mid S) \mid T \xrightarrow{\tau} (\nu r)(R' \mid S') \mid T}\ \text{L-PAR'}$$

  Since r is not a free name of T, the latter process is structurally congruent to process $(\nu r)(R' \mid (S' \mid T))$.

- · L-PAR-CLOSE: then $R \xrightarrow{r\bowtie(a)\overline{a}} R'$ and $S \mid T \xrightarrow{r\overline{\bowtie}\downarrow a} U$. This case is analogous to that of L-SESS-COM-PAR, the extra name restrictions in the resulting processes posing no additional problem.

- · L-PAR-CLOSE’: then $R \xrightarrow{r\bowtie\downarrow a} R'$ and $S \mid T \xrightarrow{r\overline{\bowtie}(a)\overline{a}} U$. This case is again analogous to the previous one.

The case when $Q \xrightarrow{\alpha} Q'$ is dealt with by a similar case analysis.

- *Name restriction*

  – Parallel composition. Suppose P is $((\nu n)R) \mid S$ and Q is $(\nu n)(R \mid S)$. Assume first that $P \xrightarrow{\alpha} P'$; there are three different cases, according to which transition rule was used.

  - · L-PAR: then $(\nu n)R \xrightarrow{\alpha} R'$. There are three possible sub-cases.

    - ∗ Suppose $(\nu n)R \xrightarrow{\alpha} R'$ follows by L-RES. Then n is not a name in α, R' is $(\nu n)R''$ and $R \xrightarrow{\alpha} R''$. Since n is also not a name in S, the following derivation establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{\alpha} R''}{R \mid S \xrightarrow{\alpha} R'' \mid S}\ \text{L-PAR}}{(\nu n)(R \mid S) \xrightarrow{\alpha} (\nu n)(R'' \mid S)}\ \text{L-RES}$$

    - ∗ Suppose $(\nu n)R \xrightarrow{\alpha} R'$ follows by L-SESS-RES. Then α is τ, R' is $(\nu n)R''$ and $R \xrightarrow{n\tau} R''$. Again, since n is also not a name in S, the following derivation establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{n\tau} R''}{R \mid S \xrightarrow{n\tau} R'' \mid S}\ \text{L-PAR}}{(\nu n)(R \mid S) \xrightarrow{\tau} (\nu n)(R'' \mid S)}\ \text{L-SESS-RES}$$

- Suppose $(\nu n)R \xrightarrow{\alpha} R'$ follows by L-EXTR. Then α is $(n)\overline{n}$ and $R \xrightarrow{\overline{n}} R'$. Again, since n is also not a name in S, the following derivation establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{\overline{n}} R'}{R \mid S \xrightarrow{\overline{n}} R' \mid S}\ \text{L-PAR}}{(\nu n)(R \mid S) \xrightarrow{(n)\overline{n}} (R' \mid S)}\ \text{L-EXTR}$$

In either case, it is easy to verify that $(\nu n)(R \mid S)$ evolves to a process structurally congruent to the evolution of $((\nu n)R) \mid S$.

- L-PAR': then $S \xrightarrow{\alpha} S'$. Since $(\nu n)R \mid S$ is well-formed, n does not occur in S; therefore n cannot occur in α. Then the following derivation establishes the thesis.

$$\dfrac{\dfrac{S \xrightarrow{\alpha} S'}{R \mid S \xrightarrow{\alpha} R \mid S'}\ \text{L-PAR'}}{(\nu n)(R \mid S) \xrightarrow{\alpha} (\nu n)(R \mid S')}\ \text{L-RES}$$

- L-SESS-COM-PAR: then α is $r\tau$, $(\nu n)R \xrightarrow{r\bowtie\Updownarrow v} R'$ and $S \xrightarrow{r\overline{\bowtie}\overline{\Updownarrow}v} S'$. Then necessarily R' is $(\nu n)R''$ and the former transition is inferred via L-RES. The following derivation establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{r\bowtie\Updownarrow v} R'' \quad S \xrightarrow{r\overline{\bowtie}\overline{\Updownarrow}v} S'}{R \mid S \xrightarrow{r\tau} R'' \mid S'}\ \text{L-SESS-COM-PAR}}{(\nu n)(R \mid S) \xrightarrow{r\tau} (\nu n)(R'' \mid S')}\ \text{L-RES}$$

The cases when the rule applied is L-SERV-COM-PAR, L-PAR-CLOSE or L-PAR-CLOSE' are similar, except that further applications of S-SWAP may be necessary to verify that both processes evolve to structurally congruent processes.

Assume now that $Q \xrightarrow{\alpha} Q'$. Since the top-level constructor in Q is name restriction, there are three possible cases.

- Assume the last rule applied is L-RES. Then $R \mid S \xrightarrow{\alpha} T$, with Q' being $(\nu n)T$ and n a name not occurring in α. There are six sub-cases, corresponding to the six different rules that may be used to infer the transition of $R \mid S$.
  - L-PAR: then T is $R' \mid S$ with $R \xrightarrow{\alpha} R'$; then the following proof establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{\alpha} R'}{(\nu n)R \xrightarrow{\alpha} (\nu n)R'}\ \text{L-RES}}{((\nu n)R) \mid S \xrightarrow{\alpha} ((\nu n)R') \mid S}\ \text{L-PAR}$$

  - L-PAR': then T is $R \mid S'$ with $S \xrightarrow{\alpha} S'$; the following proof establishes the thesis.

$$\dfrac{S \xrightarrow{\alpha} S'}{((\nu n)R) \mid S \xrightarrow{\alpha} ((\nu n)R) \mid S'}\ \text{L-PAR'}$$

  - L-SESS-COM-PAR: then α is $r\tau$, $R \xrightarrow{r\bowtie\Updownarrow v} R'$, $S \xrightarrow{r\overline{\bowtie}\overline{\Updownarrow}v} S'$ and T is $R' \mid S'$. Care must be taken to distinguish whether n is v.
    If n is not v, then the following derivation establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{r\bowtie\Updownarrow v} R'}{(\nu n)R \xrightarrow{r\bowtie\Updownarrow v} (\nu n)R'}\ \text{L-RES} \qquad S \xrightarrow{r\overline{\bowtie}\overline{\Updownarrow}v} S'}{((\nu n)R) \mid S \xrightarrow{r\tau} ((\nu n)R') \mid S'}\ \text{L-SESS-COM-PAR}$$

If n is v, then by well-formedness the process performing the output must be R (otherwise S would contain a binder n, which violates the assumption that all bound names in $R \mid S$ are distinct); the following proof establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{r\bowtie\Updownarrow v} R'}{(\nu n)R \xrightarrow{r\bowtie(n)\Updownarrow v} R'}\ \textsc{L-extr} \qquad S \xrightarrow{r\overline{\bowtie}\overline{\Updownarrow} v} S'}{((\nu n)R) \mid S \xrightarrow{r\tau} (\nu n)(R' \mid S')}\ \textsc{L-sess-close}$$

- L-SERV-COM-PAR: then α is τ, $R \xrightarrow{a\Leftrightarrow(r)} R'$, $S \xrightarrow{a\overline{\Leftrightarrow}(r)} S'$ and T is $(\nu r)(R' \mid S')$. Notice that from the hypothesis it follows that n is distinct from r. Consider the following derivation.

$$\dfrac{\dfrac{R \xrightarrow{a\Leftrightarrow(r)} R'}{(\nu n)R \xrightarrow{a\Leftrightarrow(r)} (\nu n)R'}\ \textsc{L-res} \qquad S \xrightarrow{a\overline{\Leftrightarrow}(r)} S'}{((\nu n)R) \mid S \xrightarrow{\tau} (\nu r)(((\nu n)R') \mid S')}\ \textsc{L-serv-com-par}$$

Finally, using rules S-EXTR-PAR and S-SWAP, it follows that $(\nu r)(((\nu n)R') \mid S') \equiv (\nu r)(\nu n)(R' \mid S') \equiv Q'$.

- L-PAR-CLOSE: then α is $r\tau$, $R \xrightarrow{r\bowtie(a)\overline{a}} R'$, $S \xrightarrow{r\overline{\bowtie}\downarrow a} S'$ and T is $(\nu a)(R' \mid S')$. Again, by well-formedness, a is distinct from n. The following derivation establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{r\bowtie(a)\overline{a}} R'}{(\nu n)R \xrightarrow{r\bowtie(a)\overline{a}} (\nu n)R'}\ \textsc{L-res} \qquad S \xrightarrow{r\overline{\bowtie}\downarrow a} S'}{((\nu n)R) \mid S \xrightarrow{r\tau} (\nu a)(((\nu n)R') \mid S')}\ \textsc{L-sess-close}$$

- L-PAR-CLOSE’: then α is $r\tau$, $R \xrightarrow{r\bowtie\downarrow a} R'$, $S \xrightarrow{r\overline{\bowtie}(a)\overline{a}} S'$ and T is $(\nu a)(R' \mid S')$. Again, by well-formedness, a is distinct from n. The following derivation establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{r\bowtie\downarrow a} R'}{(\nu n)R \xrightarrow{r\bowtie\downarrow a} (\nu n)R'}\ \textsc{L-res} \qquad S \xrightarrow{r\overline{\bowtie}(a)\overline{a}} S'}{((\nu n)R) \mid S \xrightarrow{r\tau} (\nu a)(((\nu n)R') \mid S')}\ \textsc{L-sess-close'}$$

- Assume the last rule applied is L-SESS-RES. This case is very similar to the previous one, but simpler: since session names may not be communicated, there are less possible cases and no need arises to use close rules.
- Assume the last rule applied is L-EXTR. Then α is $(n)\,\mu$, where μ is an output (session or stream). By well-formedness, n does not occur in S, whence it follows that necessarily $R \xrightarrow{\mu} R'$ and Q' is $R' \mid S$. Then the following proof shows that $P \xrightarrow{\alpha} Q'$.

$$\dfrac{\dfrac{R \xrightarrow{\mu} R'}{(\nu n)R \xrightarrow{(n)\mu} R'}\ \textsc{L-extr}}{((\nu n)R) \mid S \xrightarrow{(n)\mu} R' \mid S}\ \textsc{L-par}$$

- Composition with stream. There are two congruence rules for this case; both of them require a case analysis that is completely similar to that in the previous case (since composition with a stream is very similar to parallel composition). The extra case arising from L-FEED-CLOSE is similar to the other close rules.
- Session. Assume P is $r \bowtie ((\nu a)R)$ and Q is $(\nu a)(r \bowtie R)$. Suppose first that $P \xrightarrow{\alpha} P'$; there are two different cases.
  - L-SESS-VAL: then α is $r \bowtie \mu$, where μ is an input/output action. There are two possible sub-cases, according to how the transition of $(\nu a)R$ is inferred (since L-SESS-RES does not apply).

      - ∗ L-RES: then P' is $r \bowtie ((\nu a)R')$ with $R \xrightarrow{\mu} R'$, and the following derivation establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{\mu} R'}{r \bowtie R \xrightarrow{r \bowtie \mu} r \bowtie R'}\ \textsc{L-sess-val}}{(\nu a)(r \bowtie R) \xrightarrow{r \bowtie \mu} (\nu a)(r \bowtie R')}\ \textsc{L-res}$$

      - ∗ L-EXTR: then P' is $r \bowtie R'$, μ is $\overline{a}$, and the following derivation establishes the thesis.

$$\dfrac{\dfrac{R \xrightarrow{\overline{a}} R'}{r \bowtie R \xrightarrow{r \bowtie \overline{a}} r \bowtie R'}\ \textsc{L-sess-val}}{(\nu a)(r \bowtie R) \xrightarrow{r \bowtie (a)\overline{a}} r \bowtie R'}\ \textsc{L-extr}$$

    - · L-SESS-PASS: this case is very similar with only two differences. In the case of L-EXTR, μ is now $\Uparrow a$, and the rest follows as before. There is also the extra case of L-SESS-RES, which is straightforward.

  Assume now that $Q \xrightarrow{\alpha} Q'$. The proof is very similar, so we will only sketch it; there are three cases.

    - · L-RES: then Q' is $(\nu a)S$ with $r \bowtie R \xrightarrow{\alpha} S$. There are two cases for the latter transition; in either of them, S must be of the form $r \bowtie R'$ and the thesis follows by swapping the application of the two rules.
    - · L-SESS-RES: similar, but now there is only one sub-case, corresponding to L-SESS-PASS.
    - · L-EXTR: then $r \bowtie R \xrightarrow{\mu} Q'$ and either α is $(a)\,\mu$ or α is $s \bowtie (a)\overline{a}$ and μ is $s \bowtie a$ for some session name s. Again there are two cases for the latter transition, and a straightforward swapping of the two rules yields the proof that $P \xrightarrow{\alpha} Q'$.

  - – Commutativity. Straightforward, since two different names are involved and well-formedness of the processes guarantees that all side conditions in the relevant rules will hold.
  - – Zero. Straightforward, since $(\nu a)0 \not\rightarrow$ and $0 \not\rightarrow$.

- *Recursion* This case is completely straightforward: if **rec** $X.R \xrightarrow{\alpha} P'$, then the only rule that can have been used to infer that transition is L-REC, whence it immediately follows that $R\left[{}^{\mathbf{rec}\,X.R}/_{X}\right] \xrightarrow{\alpha} P'$. Reciprocally, if the latter condition holds, then by L-REC also **rec** $X.R \xrightarrow{\alpha} P'$. □

C.2. Derivability of the new transition rules in the original LTS

Lemma 11 *Rule* L-PAR' *is admissible in the original LTS for* SSCC.

Proof The following derivation shows that any instance of L-PAR' can be derived in the original LTS.

$$\dfrac{\dfrac{Q \xrightarrow{\mu} Q' \qquad \mathsf{bn}(\mu \cap \mathsf{fn}(P)) = \emptyset}{Q \mid P \xrightarrow{\mu} Q' \mid P}\ \textsc{L-par} \qquad Q \mid P \equiv P \mid Q \qquad Q' \mid P \equiv P \mid Q'}{P \mid Q \xrightarrow{\mu} P \mid Q'}\ \textsc{L-struct}$$

□

Lemma 12 *Rule* L-REC *is admissible in the original LTS for* SSCC.

Proof The following derivation shows that any instance of L-REC can be derived in the original LTS.

$$\dfrac{P\left[{}^{\mathsf{rec}\ X.P}/_{X}\right] \xrightarrow{\mu} P' \qquad P\left[{}^{\mathsf{rec}\ X.P}/_{X}\right] \equiv \mathsf{rec}\ X.P \qquad P' \equiv P'}{\mathsf{rec}\ X.P \xrightarrow{\mu} P'}\ \textsc{L-struct}$$

□

The following lemma shows a property of transitions derived in the new LTS.

Lemma 13 *Let P and P' be processes and a be a name.*

1. *If $P \xrightarrow{(a)\uparrow a} P'$ then $P \equiv (\nu a)R$ for some R such that $R \xrightarrow{\uparrow a} P'$.*
2. *If $P \xrightarrow{r\bowtie(a)\uparrow a} P'$ then $P \equiv (\nu a)R$ for some R such that $R \xrightarrow{r\bowtie\uparrow a} P'$.*
3. *If $P \xrightarrow{(a)\Uparrow a} P'$ then $P \equiv (\nu a)R$ for some R such that $R \xrightarrow{\Uparrow a} P'$.*

Proof All three parts of the lemma are proved by induction on the proof of the transition.

For (i), the base case is when rule L-EXTR is applied. Then the thesis follows immediately from the premise of the rule and reflexivity of $\equiv$. The induction cases are when one out of L-PAR, L-PAR', L-STREAM-PASS-P, L-STREAM-PASS-Q, L-RES or L-REC is applied.

The first four cases are analogous. Suppose rule L-PAR was applied; then P is $P_1 \mid P_2$, P' is $P_1' \mid P_2$, $P_1 \xrightarrow{(a)\uparrow a} P_1'$ and a is not a free name of P_2. By induction hypothesis, $P_1 \equiv (\nu a)R$ with $R \xrightarrow{\uparrow a} P_1'$; also $P_1 \mid P_2 \equiv (\nu a)R \mid P_2$. By rule L-PAR, $R \mid P_2 \xrightarrow{\uparrow a} P'$. The case of L-REC is also straightforward: since $\mathsf{rec}\ X.P \equiv P\left[{}^{\mathsf{rec}\ X.P}/_{X}\right]$, the induction hypothesis immediately establishes the result. Finally, for L-RES, simply apply the induction hypothesis and use S-SWAP to conclude the thesis.

The proof of (ii) is completely similar except for the base case. Here, the rule being applied may also be L-SESS-VAL, in which case P is $r \bowtie Q$ and P' is $r \bowtie Q'$. By (i), also $Q \equiv (\nu a)R$ with $R \xrightarrow{\uparrow a} Q'$, whence $r \bowtie R \xrightarrow{r\bowtie\uparrow a} P'$. Since $r \bowtie (\nu a)R \equiv (\nu a)r \bowtie R$, the thesis follows.

The last case is analogous to the first. □

Lemma 14 *Rules* L-PAR-CLOSE *and* L-PAR-CLOSE' *are admissible in the original LTS for* SSCC.

Proof Suppose $P \mid Q \xrightarrow{r\tau} (\nu a)P' \mid Q'$ by L-PAR-CLOSE. By part (ii) of Lemma 13, $P \equiv (\nu a)R$ for some R such that $R \xrightarrow{r\bowtie\uparrow a} P'$. The following derivation shows that this instance of L-PAR-CLOSE can be derived in the original LTS.

$$\dfrac{\dfrac{\dfrac{R \xrightarrow{r\bowtie\uparrow a} P' \qquad Q \xrightarrow{r\bowtie\downarrow a} Q'}{R \mid Q \xrightarrow{r\tau} P' \mid Q'}\ \text{L-PAR}}{(\nu a)R \mid Q \xrightarrow{r\tau} (\nu a)P' \mid Q'}\ \text{L-RES} \qquad (\nu a)R \equiv P \quad Q \equiv Q}{P \mid Q \xrightarrow{r\tau} (\nu a)P' \mid Q'}\ \text{L-STRUCT}$$

For rule L-PAR-CLOSE', apply the previous construction with L-PAR' instead of L-PAR and invoke Lemma 11. □

Lemma 15 *Rules* L-SESS-CLOSE *and* L-SESS-CLOSE' *are admissible in the original LTS for* SSCC.

Proof Analogous to the previous one. □

Lemma 16 *Rule* L-FEED-CLOSE *is admissible in the original LTS for* SSCC.

Proof Suppose $\textbf{stream}\ P\ \textbf{as}\ f = \vec{w}\ \textbf{in}\ Q \xrightarrow{\tau} (\nu a)\textbf{stream}\ P'\ \textbf{as}\ f = a :: \vec{w}\ \textbf{in}\ Q$ by L-FEED-CLOSE. By part (iii) of Lemma 13, $P \equiv (\nu a)R$ for some R such that $R \xrightarrow{\Uparrow a} P'$. The following derivation shows that this instance of L-PAR-CLOSE can be derived in the original LTS. Note, in fact, that the application of L-STRUCT is sound, since from $(\nu a)R \equiv P$ it follows that $(\nu a)\textbf{stream}\ R\ \textbf{as}\ f = \vec{w}\ \textbf{in}\ Q \equiv \textbf{stream}\ P\ \textbf{as}\ f = \vec{w}\ \textbf{in}\ Q$ (for the left-hand-side), and since $\equiv$ is reflexive (for the right-hand-side).

$$\dfrac{\dfrac{\dfrac{R \xrightarrow{\Uparrow a} P'}{\textbf{stream}\ R\ \textbf{as}\ f = \vec{w}\ \textbf{in}\ Q \xrightarrow{\tau} \textbf{stream}\ P'\ \textbf{as}\ f = a :: \vec{w}\ \textbf{in}\ Q}\ \text{L-STREAM-FEED}}{(\nu a)\textbf{stream}\ R\ \textbf{as}\ f = \vec{w}\ \textbf{in}\ Q \xrightarrow{\tau} (\nu a)\textbf{stream}\ P'\ \textbf{as}\ f = a :: \vec{w}\ \textbf{in}\ Q}\ \text{L-RES}}{\textbf{stream}\ P\ \textbf{as}\ f = \vec{w}\ \textbf{in}\ Q \xrightarrow{\tau} (\nu a)\textbf{stream}\ P'\ \textbf{as}\ f = a :: \vec{w}\ \textbf{in}\ Q}\ \text{L-STRUCT}$$

□

D. On the bisimilarities of SSCC

D.1. Strong bisimilarity

We study here strong bisimilarity, hereafter referred to simply as "bisimilarity", as defined in Definition 10. Remember that bisimilarity can be obtained as the union of all bisimulations or as a fixed-point of a suitable monotonic operator; also it is well defined, as the next result shows.

Theorem 6 *Structurally congruent processes are bisimilar.*

Proof It suffices to show that $\equiv$ is a bisimulation, which is an immediate consequence of the Harmony Lemma. □

We now show that bisimilarity is a non-input congruence, just as in π-calculus. The strategy of the proof is the same as in [SW01], based on the notion and properties of a relation *progressing* to another relation.

Definition 17 A relation $\mathcal{R}$ on processes *strongly progresses* to another relation $\mathcal{S}$, denoted $\mathcal{R} \rightsquigarrow \mathcal{S}$, if, whenever $P\mathcal{R}Q$, $P \xrightarrow{\alpha} P'$ implies $Q \xrightarrow{\alpha} Q'$ for some Q' with $P'\mathcal{S}Q'$, and vice-versa.

Definition 18 A function $\mathcal{F}$ on processes is *strongly safe* if $\mathcal{R} \subseteq S$ and $\mathcal{R} \rightsquigarrow \mathcal{S}$ imply $\mathcal{F}(\mathcal{R}) \subseteq \mathcal{F}(\mathcal{S})$ and $\mathcal{F}(\mathcal{R}) \rightsquigarrow \mathcal{F}(\mathcal{S})$.

Lemma 17 *If $\mathcal{F}$ is strongly safe and $\sim \subseteq \mathcal{F}(\sim)$, then $F(\sim) = \sim$.*

Proof See [SW01]. □

Given a function $\mathcal{F}$, define $\mathcal{F}^*$ such that $\mathcal{F}^*(\mathcal{R})$ is the transitive closure of $\mathcal{F}(\mathcal{R})$.

Lemma 18 *If $\mathcal{F}$ is such that $\mathcal{R} \subseteq S$ and $\mathcal{R} \rightsquigarrow \mathcal{S}$ imply that $\mathcal{F}(\mathcal{R}) \subseteq \mathcal{F}^*(\mathcal{S})$ and $\mathcal{F}(\mathcal{R}) \rightsquigarrow \mathcal{F}^*(\mathcal{S})$, then $\mathcal{F}^*$ is strongly safe.*

Proof See [SW01]. □

The proof relies on defining functions $\mathcal{F}_{\mathrm{ni1}}$ and $\mathcal{F}_{\mathrm{ni}}$ like those for π-calculus; however, the definition of the former has to be slightly adapted.

Definition 19
- An n-ary *multi-hole context* C is a process where some occurrences of 0 have been replaced by holes $[\cdot]_i$; each hole may occur zero or more times. Given n processes $P_1, \ldots, P_n$, $C[P_1, \ldots, P_n]$ is the process obtained by uniformly replacing all occurrences of all holes in C by the corresponding process.
- A (multi-hole) context is said to be *non-input* if no hole occurs under an input prefix (x) or $f(x)$.
- Functions $\mathcal{F}_{\mathrm{ni1}}$ and $\mathcal{F}_{\mathrm{ni}}$ are defined as follows.

$$\begin{aligned}\mathcal{F}_{\mathrm{ni1}}(\mathcal{R}) &= \{\langle C[P], C[Q]\rangle \,[]\, P\mathcal{R}Q \text{ and } C \text{ is a non-input context}\}\\ \mathcal{F}_{\mathrm{ni}}(\mathcal{R}) &= \{\langle C[P_1, \ldots, P_n], C[Q_1, \ldots, Q_n]\rangle \,[]\, P_i\mathcal{R}Q_i \text{ and } C \text{ is an } n\text{-ary non-input context}\}\end{aligned}$$

Lemma 19 $\mathcal{F}_{\mathrm{ni}} = \mathcal{F}^*_{\mathrm{ni1}}$.

Proof As for π-calculus. □

Lemma 20 *Function $\mathcal{F}_{ni}$ is strongly safe.*

Proof Applying Lemma 18, one must show that, whenever $\mathcal{R} \subseteq \mathcal{S}$ and $\mathcal{R} \rightsquigarrow \mathcal{S}$, both $\mathcal{F}_{\mathrm{ni1}}(\mathcal{R}) \subseteq \mathcal{F}_{\mathrm{ni1}}(\mathcal{S})$ and $\mathcal{F}_{\mathrm{ni1}}(\mathcal{R}) \rightsquigarrow \mathcal{F}_{\mathrm{ni}}(\mathcal{S})$. The first of these is trivial by definition of $\mathcal{F}_{\mathrm{ni1}}$.

Assume that $P\mathcal{R}Q$. One must show that, for every context C, if $C[P] \xrightarrow{\alpha} P'$, then $C[Q] \xrightarrow{\alpha} Q'$ for some P' and Q' such that there exist an n-ary context C' and processes $P_1\mathcal{S}Q_1, \ldots, P_n\mathcal{S}Q_n$ for which P' is $C'[P_1, \ldots, P_n]$ and Q' is $C'[Q_1, \ldots, Q_n]$.

The proof is by induction on the derivation tree for $C[P] \xrightarrow{\alpha} P'$. In all steps, there are two cases to consider, according to whether C is $[\cdot]$ or not; the former case is always trivial, since the hypothesis $\mathcal{R} \rightsquigarrow \mathcal{S}$ establishes the thesis. Therefore, we always assume below that C is not $[\cdot]$. The proof looks at the last rule being applied.

- L-SEND: then C is $v.C_0$ and α is $\uparrow v$ for some v. Furthermore, $v.C_0[Q] \xrightarrow{\uparrow v} C_0[Q]$; since C_0 is also a multi-hole context and $\mathcal{R} \subseteq \mathcal{S}$, it follows that $\langle C_0[P], C_0[Q]\rangle \in \mathcal{F}_{\mathrm{ni}}(\mathcal{S})$, hence the thesis holds.

- L-RECEIVE: then C is $(x)C_0$, and since C is a non-input context (by definition of $\mathcal{F}_{\mathrm{nil}}$), it follows that C_0 does not contain holes; hence in this case $C[P]$ and $C[Q]$ coincide, and the result is trivial.
- L-FEED: then C is $\textbf{feed}\ v.C_0$ and α is $\Uparrow v$ for some v. Furthermore, $\textbf{feed}\ v.C_0[Q] \xrightarrow{\Uparrow v} C_0[Q]$; since C_0 is also a multi-hole context and $\mathcal{R} \subseteq \mathcal{S}$, it follows that $\langle C_0[P], C_0[Q]\rangle \in \mathcal{F}_{\mathrm{ni}}(\mathcal{S})$, hence the thesis holds.
- L-READ: then C is $f(x).C_0$, and since C is a non-input context (by definition of $\mathcal{F}_{\mathrm{nil}}$), it follows that C_0 does not contain holes; hence in this case $C[P]$ and $C[Q]$ coincide, and the result is trivial.
- L-CALL: then C is $a \Leftarrow C_0$ and α is $a \Leftarrow (r)$ for some r not occurring free in $C_0[P]$. Furthermore, $a \Leftarrow C_0[Q] \xrightarrow{a \Leftarrow (r)} r \lhd C_0[Q]$, since by definition of bisimulation r does not occur free in $C_0[Q]$. Taking C' to be the context $r \lhd C_0$ establishes the thesis.
- L-INV: analogous.
- L-PAR: there are two cases to consider.
  - If C is $C_0 \mid R$, then $C_0[P] \xrightarrow{\alpha} P'$ and α and R share no bound names. By induction hypothesis there exists a process Q' such that $C_0[Q] \xrightarrow{\alpha} Q'$, and P', Q' are respectively $C_0'[P_1, \ldots, P_n]$ and $C_0'[Q_1, \ldots, Q_n]$ for some n-ary multi-hole context C_0' and processes $P_1\mathcal{S}Q_1, \ldots, P_n\mathcal{S}Q_n$. Thus $C_0[Q] \mid R \xrightarrow{\alpha} Q' \mid R$, hence taking C' to be $C_0' \mid R$ establishes the thesis.
  - If C is $R \mid C_0$, then $R \xrightarrow{\alpha} R'$ and $C_0[P]$ and R share no bound names. By the hypothesis of L-PAR, $C_0[Q]$ and R also share no bound names, hence $R \mid C_0[Q] \xrightarrow{\alpha} R \mid Q'$, and taking C' to be $R \mid C_0'$ establishes the thesis.
- L-PAR’: analogous (the two cases are reversed).
- L-STREAM-PASS-P and L-STREAM-PASS-Q: analogous to L-PAR and L-PAR’, respectively.
- L-STREAM-FEED: there are two cases to consider.
  - If C is $\textbf{stream}\ C_0\ \textbf{as}\ f = \vec{w}\ \textbf{in}\ R$, then $C_0[P] \xrightarrow{\Uparrow v} P'$; by induction hypothesis, there exists a process Q' such that $C_0[Q] \xrightarrow{\Uparrow v} Q'$, and P' and Q' are respectively $C_0'[P_1, \ldots, P_n]$ and $C_0'[Q_1, \ldots, Q_n]$ for some n-ary multi-hole context C_0' and processes $P_1\mathcal{S}Q_1, \ldots, P_n\mathcal{S}Q_n$. Thus $\textbf{stream}\ C_0[Q]\ \textbf{as}\ f = \vec{w}\ \textbf{in}\ R \xrightarrow{\tau} \textbf{stream}\ Q'\ \textbf{as}\ f = v :: \vec{w}\ \textbf{in}\ R$, hence taking C' to be $\textbf{stream}\ C_0'\ \textbf{as}\ f = v :: \vec{w}\ \textbf{in}\ R$ establishes the thesis.
  - If C is $\textbf{stream}\ R\ \textbf{as}\ f = \vec{w}\ \textbf{in}\ C_0$, then $R \xrightarrow{\Uparrow v} R'$, hence $\textbf{stream}\ R\ \textbf{as}\ f = \vec{w}\ \textbf{in}\ C_0[Q] \xrightarrow{\tau} \textbf{stream}\ R'\ \textbf{as}\ f\ \ \textbf{in} = v :: \vec{w}\, C_0[Q]$, hence taking C' to be $\textbf{stream}\ R'\ \textbf{as}\ f = v :: \vec{w}\ \textbf{in}\ C_0$ establishes the thesis.
- L-STREAM-CONS: analogous (the two cases are reversed).
- L-SESS-VAL: then C is $r \bowtie C_0$, α is $r \bowtie \mu$ for some μ, and $C_0[P] \xrightarrow{\mu} P_0$. By induction hypothesis, $C_0[Q] \xrightarrow{\mu} Q_0$ for some Q_0 such that there exist a multi-hole context C_0' and processes $P_1\mathcal{S}Q_1, \ldots, P_n\mathcal{S}Q_n$ for which P_0 is $C_0'[P_1, \ldots, P_n]$ and Q_0 is $C_0'[Q_1, \ldots, Q_n]$. Taking C' to be $r \bowtie C_0'$ establishes the thesis, since then $r \bowtie C_0[Q] \xrightarrow{r \bowtie \mu} C'[Q_1, \ldots, Q_n]$.
- L-SESS-PASS: then C is $r \bowtie C_0$, α is neither an input nor an output, and $C_0[P] \xrightarrow{\alpha} P_0$. The proof then follows as above except that the action does not change when the session is added to C_0'.
- L-SESS-COM-PAR: C is either $C_0 \mid R$ or $R \mid C_0$; the two cases are analogous, so assume the first holds. Then $C_0[P] \xrightarrow{r \bowtie \Updownarrow v} P'$, $R \xrightarrow{r \overline{\bowtie} \Updownarrow v} R'$, α is $r\tau$ for some r and $C_0[P] \mid R \xrightarrow{r\tau} P' \mid R'$. By induction hypothesis there exists a process Q' such that $C_0[Q] \xrightarrow{r \bowtie \Updownarrow v} Q'$, and P', Q' are respectively $C_0'[P_1, \ldots, P_n]$ and $C_0'[Q_1, \ldots, Q_n]$ for some n-ary multi-hole context C_0' and processes $P_1\mathcal{S}Q_1, \ldots, P_n\mathcal{S}Q_n$. Then $C_0[Q] \mid R \xrightarrow{r\tau} Q' \mid R'$, hence taking C' to be $C_0' \mid R'$ establishes the thesis.
- the cases concerning rules L-SERV-COM-PAR, L-SESS-COM-STREAM, L-SERV-COM-STREAM, L-PAR-CLOSE, L-PAR- CLOSE’, L-FEED-CLOSE, L-SESS-COM-CLOSE and L-SERV-COM-CLOSE are all very similar to the previous one.

- L-RES: then C is $(\nu a)C_0$, $(\nu a)C_0[P] \overset{\alpha}{\rightarrow} (\nu a)P'$, a is not a name in α and $C_0[P] \overset{\alpha}{\rightarrow} P'$. By induction hypothesis there exists a process Q' such that $C_0[Q] \overset{\alpha}{\rightarrow} Q'$, and P', Q' are respectively $C_0'[P_1, \ldots, P_n]$ and $C_0'[Q_1, \ldots, Q_n]$ for some n-ary multi-hole context C_0' and processes $P_1 \mathcal{S} Q_1, \ldots, P_n \mathcal{S} Q_n$. Thus $(\nu a)C_0[Q] \overset{\alpha}{\rightarrow} (\nu a)Q'$, hence taking C' to be $(\nu a)C_0'$ establishes the thesis.
- L-SESS-RES: analogous, only now α is τ and the induction hypothesis is applied to a transition via $a\tau$.
- L-EXTR: analogous, only now α is an action extruding a and the induction hypothesis is applied to a transition without the extrusion; furthermore, the context C' is simply C_0'.
- L-REC: C is $\mathsf{rec}\, X.C_0$ and $C_0[P]\left[{}^{\mathsf{rec}\, X.C_0[P]}/_{X}\right] \overset{\alpha}{\rightarrow} P'$. Since P is well-formed, it contains no free occurrences of X; hence there exists a context C_1 such that $C_1[P]$ is $C_0[P]\left[{}^{\mathsf{rec}\, X.C_0[P]}/_{X}\right]$ and $C_1[Q]$ is $C_0[Q]\left[{}^{\mathsf{rec}\, X.C_0[Q]}/_{X}\right]$. Hence the induction hypothesis applies, and there exist a process Q' and a context C' such that $C_1[Q] \overset{\alpha}{\rightarrow} Q'$, P' is $C'[P]$ and Q' is $C'[Q]$. Therefore also $\mathsf{rec}\, X.C_0[Q] \overset{\alpha}{\rightarrow} Q'$, and C' is the required context. □

Theorem 7 *Bisimilarity is a non-input congruence.*

Proof Straightforward consequence of Lemmas 17 and 20. □

At this point we can explain in more detail why the original LTS for SSCC had to be changed. Consider any derivation containing an application of L-STRUCT.

$$\frac{P \overset{\mu}{\rightarrow} P' \qquad P \equiv Q \qquad P' \equiv Q'}{Q \overset{\mu}{\rightarrow} Q'}\ \text{L-STRUCT}$$

In general, the induction hypothesis will not be applicable to the subtree that shows $P \overset{\mu}{\rightarrow} P'$, since there is no obvious relationship between P and Q; furthermore, the thesis of the induction hypothesis does not help in establishing the final result, since again there is no obvious relationship between Q' and P'.

Observe also that this is not a problem of this particular proof technique. Whether the induction were on the derivation tree (as above), on contexts (as the proof for π-calculus, see [SW01]) or on processes (arguably an alternative) the same problem would arise, since the issue arises from the fact that the theorem assumes hypotheses on the actual process performing the transition. This justifies the attempt to eliminate L-STRUCT from the LTS altogether.

D.2. Weak bisimilarity

We turn now to weak bisimilarity, defined according to Definition 12. Weak bisimilarity treats internal actions as irrelevant. Again, weak bisimilarity can be obtained as the union of all weak bisimulations or as a fixed-point of a suitable monotonic operator.

Theorem 8 *Let $\simeq$ be the largest relation such that, whenever $P \simeq Q$, for every process P' and action α, if $P \overset{\alpha}{\rightarrow} P'$, then $Q \overset{\alpha}{\Rightarrow} Q'$ for some Q' with $P' \simeq Q'$ and vice-versa. Then $P \approx Q$ iff $P \simeq Q$.*

Proof The direct implication is straightforward, since $P \overset{\alpha}{\rightarrow} P'$ implies that $P \overset{\alpha}{\Rightarrow} P'$. For the converse, assume that $P \overset{\alpha}{\Rightarrow} P'$. If P' is P and α is τ, then result is trivial; otherwise, there exist processes $P_1, \ldots, P_n$ and $P_1', \ldots, P_m'$ such that P is P_1, $P_i \overset{\tau}{\rightarrow} P_{i+1}$ for $i < n$, $P_n \overset{\alpha}{\rightarrow} P_1'$, $P_j' \overset{\tau}{\rightarrow} P_{j+1}'$ for $j < m$ and P_m' is P'. By hypothesis, there exist processes $Q_1, \ldots, Q_n$ and $Q_1', \ldots, Q_m'$ (not necessarily distinct) such that $Q_i \overset{\tau}{\Rightarrow} Q_{i+1}$ for $i < n$, $Q_n \overset{\alpha}{\Rightarrow} Q_1'$ and $Q_j' \overset{\tau}{\Rightarrow} Q_{j+1}'$ for $j < m$; furthermore, $P_i \simeq Q_i$ and $P_j' \simeq Q_j'$ for all $i \leq n$ and $j \leq m$. In particular, $Q \overset{\alpha}{\Rightarrow} Q_m'$ and $P' \simeq Q_m'$, so $\simeq$ is a weak bisimulation. □

The reason for introducing $\simeq$ is that this relation is simpler to work with when proving properties by induction. In turn, the definition of $\approx$ is more symmetric and its relationship with $\sim$ is immediate.

We now show that bisimilarity is a non-input congruence, again like in π-calculus. The strategy of the proof is once more the same as in [SW01].

Definition 20 A relation $\mathcal{R}$ on processes *progresses* to another relation $\mathcal{S}$, denoted $\mathcal{R} \rightsquigarrow \mathcal{S}$, if, whenever $P\mathcal{R}Q$, $P \overset{\alpha}{\rightarrow} P'$ implies $Q \overset{\alpha}{\Rightarrow} Q'$ for some Q' with $P'\mathcal{S}Q'$, and vice-versa.

Definition 21 A function $\mathcal{F}$ on processes is *safe* if $\mathcal{R} \subseteq S$ and $\mathcal{R} \rightsquigarrow \mathcal{S}$ imply $\mathcal{F}(\mathcal{R}) \subseteq \mathcal{F}(\mathcal{S})$ and $\mathcal{F}(\mathcal{R}) \rightsquigarrow \mathcal{F}(\mathcal{S})$.

Lemma 21 *If $\mathcal{F}$ is safe and $\approx \subseteq \mathcal{F}(\approx)$, then $F(\approx) = \approx$.*

Proof See [SW01]. □

As is the case with π-calculus, proving that $\mathcal{F}_{\mathrm{ni}}$ is safe must be done directly, since chaining is not secure.

Lemma 22 *Function $\mathcal{F}_{ni}$ is safe.*

Proof Let $\mathcal{R} \subseteq \mathcal{S}$ and $\mathcal{R} \rightsquigarrow \mathcal{S}$. It must be shown that $\mathcal{F}_{\mathrm{ni1}}(\mathcal{R}) \subseteq \mathcal{F}_{\mathrm{ni1}}(\mathcal{S})$ and $\mathcal{F}_{\mathrm{ni}}(\mathcal{R}) \rightsquigarrow \mathcal{F}_{\mathrm{ni}}(\mathcal{S})$. As before, the first of these is trivial by definition of $\mathcal{F}_{\mathrm{ni1}}$.

Assume that $P_i \mathcal{R} Q_i$ for $i = 1, \dots, n$. It must be shown that, for every multi-context C, if there is a transition $C[P_1, \dots, P_n] \xrightarrow{\alpha} P'$, then $C[Q_1, \dots, Q_n] \overset{\alpha}{\Rightarrow} Q'$ for some P' and Q' such that there exist another multi-context C' and processes $P_1 \mathcal{S} Q_1, \dots, P_m \mathcal{S} Q_m$ for which P' is $C'[P_1, \dots, P_m]$ and Q' is $C'[Q_1, \dots, Q_m]$.

Once more we proceed by induction on the derivation tree for $C[P_1, \dots, P_n] \xrightarrow{\alpha} P'$. Most cases are very similar to the proof of Lemma 20; however, since the induction hypothesis now gives a weak transition some care must be taken.

As before, the case when C is $[\cdot]$ is straightforward; furthermore, the cases of rules L-SEND, L-RECEIVE, L-FEED, L-READ, L-CALL, L-INV, L-SESS-VAL, L-SESS-PASS, L-SESS-COM-PAR, L-SERV-COM-PAR, L-SESS-COM-STREAM, L-SERV-COM-STREAM, L-PAR-CLOSE, L-PAR-CLOSE', L-FEED-CLOSE, L-SESS-COM-CLOSE, L-SERV-COM-CLOSE, L-RES, L-SESS-RES, L-EXTR and L-REC are dealt with as in the proof of Lemma 20, with an extra step at the end (to take care of the possible extra τ steps) similar to the cases detailed above.

The only remaining cases are those when C is either a parallel composition or a stream composition, since now both subprocesses may be contexts.

- L-PAR: then C is $C_1 \mid C_2$, $C_1[P_1, \dots, P_n] \xrightarrow{\alpha} P'$, and α and $C_2[P_1, \dots, P_n]$ share no bound names. By induction hypothesis there exists a process Q' such that $C_1[Q_1, \dots, Q_n] \overset{\alpha}{\Rightarrow} Q'$, and P', Q' are respectively $C_1'[P_1', \dots, P_m']$ and $C_1'[Q_1', \dots, Q_m']$ for some multi-hole context C_1' and processes $P_1' \mathcal{S} Q_1', \dots, P_m' \mathcal{S} Q_m'$. By applying L-PAR to all steps of the sequence of transitions $C_1[Q_1, \dots, Q_n] \xrightarrow{\tau} \dots \xrightarrow{\alpha} \dots \xrightarrow{\tau} Q'$, we conclude that
$$C_1[Q_1, \dots, Q_n] \mid C_2[Q_1, \dots, Q_n] \overset{\alpha}{\Rightarrow} C_1'[Q_1', \dots, Q_m'] \mid C_2[Q_1, \dots, Q_n],$$
hence taking C' to be $C_1' \mid C_2$ establishes the thesis.[5]
- L-PAR': analogous (the roles of C_1 and C_2 are reversed).
- L-STREAM-PASS-P and L-STREAM-PASS-Q: as before, these are analogous to L-PAR and L-PAR', respectively.
- L-STREAM-FEED: then C is $\textsf{stream}\ C_1\ \textsf{as}\ f = \vec{w}\ \textsf{in}\ C_2$ and $C_1[P_1, \dots, P_n] \xrightarrow{\Uparrow v} P'$. By induction hypothesis, there exists a process Q' such that $C_1[Q_1, \dots, Q_n] \overset{\Uparrow v}{\Longrightarrow} Q'$, and P' and Q' are respectively $C_1'[P_1', \dots, P_n']$ and $C_1'[Q_1', \dots, Q_n']$ for some n-ary multi-hole context C_1' and processes $P_1' \mathcal{S} Q_1', \dots, P_n' \mathcal{S} Q_n'$. In other words, $C_1[Q_1, \dots, Q_n] \overset{\tau}{\Rightarrow} Q^* \xrightarrow{\Uparrow v} Q^{**} \overset{\tau}{\Rightarrow} Q'$; using L-STREAM-PASS-P for the τ transitions, we conclude that
$$\begin{aligned} &\textsf{stream}\ C_1[Q_1, \dots, Q_n]\ \textsf{as}\ f = \vec{w}\ \textsf{in}\ C_2[Q_1, \dots, Q_n] \\ &\qquad \overset{\tau}{\Rightarrow} \textsf{stream}\ Q^*\ \textsf{as}\ f = \vec{w}\ \textsf{in}\ C_2[Q_1, \dots, Q_n] \\ &\qquad \xrightarrow{\tau} \textsf{stream}\ Q^{**}\ \textsf{as}\ f = v :: \vec{w}\ \textsf{in}\ C_2[Q_1, \dots, Q_n] \\ &\qquad \overset{\tau}{\Rightarrow} \textsf{stream}\ C_1'[Q_1', \dots, Q_n']\ \textsf{as}\ f = v :: \vec{w}\ \textsf{in}\ C_2[Q_1, \dots, Q_n], \end{aligned}$$
hence taking C' to be $\textsf{stream}\ C_1'\ \textsf{as}\ f = v :: \vec{w}\ \textsf{in}\ C_2$ establishes the thesis.
- L-STREAM-CONS: analogous (the roles of C_1 and C_2 are reversed). □

Theorem 9 *Weak bisimilarity is a non-input congruence.*

Proof Straightforward consequence of Lemmas 21 and 22. □

[5] We assume that a n-hole context does not have to contain occurrences of all its holes, so in particular C_1 and C_2 are n-hole contexts in which some holes may not occur.

References

[AAA$^+$07] Alves A, Arkin A, Askary S, Barreto C, Bloch B, Curbera F, Ford M, Goland Y, Guízar A, Kartha N, Liu CK, Khalaf R, König D, Marin M, Mehta V, Thatte S, van der Rijn D, Yendluri P, Yiu A (2007) Business Process Execution Language for Web Services. Version 2.0

[ACKM03] Alonso G, Casati F, Kuno H, Machiraju V (2003) Web services—concepts, architectures and applications. Springer, Berlin

[Amb04] Ambler SW (2004) The Object Primer: agile model-driven development with UML 2.0. Cambridge University Press, Cambridge

[Bar84] Barendregt HP (1984) The lambda calculus: its syntax and semantics. North Holland, Amsterdam

[BBC$^+$06] Boreale M, Bruni R, Caires L, De Nicola R, Lanese I, Loreti M, Martins F, Montanari U, Ravara A, Sangiorgi D, Vasconcelos V, Zavattaro G (2006) SCC: a service centered calculus. In: Bravetti M, Núñez M, Zavattaro G (eds) Proc. of WS-FM 2006. LNCS, vol 4184. Springer, Berlin, pp 38–57

[BBNL08] Boreale M, Bruni R, De Nicola R, Loreti M (2008) Sessions and pipelines for structured service programming. In: Barthe G, de Boer FS (eds) Proc. of FMOODS'08. LNCS, vol 5051. Springer, Berlin, pp 19–38

[BBNL14] Boreale M, Bruni R, De Nicola R, Loreti M (2014) Caspis: a calculus of sessions, pipelines and services. Math Struct Comput Sci (to appear)

[BCC$^+$04] Bellwood T, Capell S, Clément L, Colgrave J, Dovey MJ, Feygin D, Hately A, Kochman R, Macias P, Novotny M, Paolucci M, von Riegen C, Rogers T, Sycara K, Wenzel P, Wu Z (2004) UDDI Version 3.0

[BFF$^+$06] Banci M, Fantechi A, Ficarra M, Giannini S, Santanni F (2006) Automotive case study: a UML description of scenarios. Internal report from the Sensoria EU IST project

[BGG$^+$06] Busi N, Gorrieri R, Guidi C, Lucchi R, Zavattaro G (2006) SOCK: a calculus for service oriented computing. In: Dan A, Lamersdorf W (eds) Proc. of ICSOC'06. LNCS, vol 4294. Springer, Berlin, pp 327–338

[BHTY10] Bocchi L, Honda K, Tuosto E, Yoshida N (2010) A theory of design-by-contract for distributed multiparty interactions. In: Gastin P, Laroussinie F (eds) Proc. of CONCUR 2010. LNCS, vol 6269. Springer, Berlin, pp 162–176

[BLMT08] Bruni R, Lanese I, Melgratti HC, Tuosto E (2008) Multiparty sessions in SOC. In: Lea D, Zavattaro G (eds) Proc. of COORDINATION'08. LNCS, vol 5052. Springer, Berlin, pp 67–82

[BM08] Bruni R, Mezzina LG (2008) Types and deadlock freedom in a calculus of services, sessions and pipelines. In: Meseguer J, Rosu G (eds) Proc. of AMAST'08. LNCS, vol 5140. Springer, Berlin, pp 100–115

[CCMW01] Christensen E, Curbera F, Meredith G, Weerawarana S (2001) WSDL: Web Services Definition Language. World Wide Web Consortium

[CFLM$^+$07] Cruz-Filipe L, Lanese I, Martins F, Ravara A, Vasconcelos VT (2007) Bisimulations in SSCC. DI/FCUL TR 07–37, Department of Informatics, Faculty of Sciences, University of Lisbon

[CFLM$^+$08] Cruz-Filipe L, Lanese I, Martins F, Ravara A, Vasconcelos VT (2008) Behavioural theory at work: program transformations in a service-centred calculus. In: Barthe G, de Boer F (eds) Proc. of FMOODS'08. LNCS, vol 5051. Springer, Berlin, pp 59–77

[CHY07] Carbone M, Honda K, Yoshida N (2007) Structured communication-centred programming for web services. In: De Nicola R (ed) Proc. of ESOP'07. LNCS, vol 4421. Springer, Berlin, pp 2–17

[CM13] Carbone M, Montesi F (2013) Deadlock-freedom-by-design: multiparty asynchronous global programming. In: Giacobazzi R, Cousot R (eds) Proc. of POPL 2013. ACM Press, New York, pp 263–274

[CNP$^+$11] Caires L, De Nicola R, Pugliese R, Vasconcelos VT, Zavattaro G (2011) Core calculi for service-oriented computing. In: Wirsing M, Hölzl MM (eds) Results of the SENSORIA Project. LNCS, vol 6582. Springer, Berlin, pp 153–188

[CPM06] Cook WR, Patwardhan S, Misra J (2006) Workflow patterns in Orc. In: Ciancarini P, Wiklicky H (eds) Proc. of COORDINATION'06. LNCS, vol 4038. Springer, Berlin, pp 82–96

[CV10] Caires L, Vieira HT (2010) Conversation types. Theoret Comput Sci 411(51–52):4399–4440

[GH05] Gay SJ, Hole MJ (2005) Subtyping for session types in the pi calculus. Acta Inform 42(2–3):191–225

[GHM$^+$07] Gudgin M, Hadley M, Mendelsohn N, Moreau J-J, Nielsen HF, Karmarkar A, Lafon Y (2007) Simple Object Access Protocol (SOAP) 1.2. World Wide Web Consortium

[GtBB$^+$06] Gnesi S, ter Beek M, Baumeister H, Hoelzl M, Moiso C, Koch N, Zobel A, Alessandrini M (2006) D8.0: Case studies scenario description. Deliverable from the Sensoria EU IST project

[HVK98] Honda K, Vasconcelos VT, Kubo M (1998) Language primitives and type disciplines for structured communication-based programming. In: Hankin C (ed) Proc. of ESOP'98. LNCS, vol 1381. Springer, Berlin, pp 22–138

[HYC08] Honda K, Yoshida N, Carbone M (2008) Multiparty asynchronous session types. In: Necula GC, Wadler P (eds) Proc. of POPL'08. ACM Press, New York, pp 273–284

[Jol] Jolie website. http://www.jolie-lang.org/

[KCM06] Kitchin D, Cook WR, Misra J (2006) A language for task orchestration and its semantic properties. In: Baier C, Hermanns H (eds) Proc. of CONCUR'06. LNCS, vol 4137. Springer, Berlin, pp 477–491

[LGMZ08] Lanese I, Guidi C, Montesi F, Zavattaro G (2008) Bridging the gap between interaction- and process-oriented choreographies. In: Cerone A, Gruner S (eds) Proc. of SEFM'08. IEEE Computer Society, Washington, DC, pp 323–332

[LPT07] Lapadula A, Pugliese R, Tiezzi F (2007) A calculus for orchestration of web services. In: De Nicola R (ed) Proc. of ESOP'07. LNCS, vol 4421. Springer, Berlin, pp 33–47

[LRV11] Lanese I, Ravara A, Vieira HT (2011) Behavioral theory for session-oriented calculi. In: Wirsing M, Hölzl MM (eds) Results of the SENSORIA Project. LNCS, vol 6582. Springer, Berlin, pp 189–213

[LS03] Levi F, Sangiorgi D (2003) Mobile safe ambients. ACM Trans. Program. Lang. Syst. 25(1):1–69

[LVMR07a] Lanese I, Vasconcelos VT, Martins F, Ravara A (2007) Disciplining orchestration and conversation in service-oriented computing. In: Proc. of SEFM 2007. IEEE Computer Society Press, Washington, DC, pp 305–314

[LVMR07b] Lanese I, Vasconcelos VT, Martins F, Ravara A (2007) Disciplining orchestration and conversation in service-oriented computing. DI/FCUL TR 07–3, Department of Informatics, Faculty of Sciences, University of Lisbon

[MC07] Misra J, Cook WR (2007) Computation orchestration: a basis for wide-area computing. J Softw Syst Model 6(1):83–110

[MC11] Montesi F, Carbone M (2011) Programming services with correlation sets. In: Kappel G, Maamar Z, Motahari-Nezhad HR (eds) Proc. of ICSOC 2011. LNCS, vol 7084. Springer, Berlin, pp 125–141

[MGZ07] Montesi F, Guidi C, Zavattaro G (2007) Composing services with JOLIE. In: Proc. of ECOWS'07. IEEE Computer Society Press, Washington, DC, pp 13–22

[Pie02] Pierce BC (2002) Types and Programming Languages. MIT Press, Cambridge

[pro] Sensoria project. Software engineering for service-oriented overlay computers. http://www.sensoria-ist.eu/

[PT95] Pierce BC, Turner DN (1995) Concurrent objects in a process calculus. In: Ito T, Yonezawa A (eds) Proc. of TPPP'94. LNCS, vol 907. Springer, Berlin, pp 187–215

[SW01] Sangiorgi D, Walker D (2001) The π-calculus: a theory of mobile processes. Cambridge University Press, Cambridge

[THK94] Takeuchi K, Honda K, Kubo M (1994) An interaction-based language and its typing system. In: Halatsis C, Maritsas DG, Philokyprou G, Theodoridis S (eds) Proc. of PARLE'94. LNCS, vol 817. Springer, Berlin, pp 398–413

[VCS08] Vieira HT, Caires L, Seco JC (2008) The conversation calculus: a model of service-oriented computation. In: Drossopoulou S (ed) Proc. of ESOP'08. LNCS, vol 4960. Springer, Berlin, pp 269–283

[vdAtHKB03] van der Aalst W, ter Hofstede A, Kiepuszewski B, Barros AP (2003) Workflow patterns. Distrib Parallel Databases 14(1):5–51

[VGR06] Vasconcelos VT, Gay S, Ravara A (2006) Typechecking a multithreaded functional language with session types. Theoret Comput Sci 368(1–2):64–87

[YV06] Yoshida N, Vasconcelos VT (2006) Language primitives and type discipline for structured communication-based programming revisited: two systems for higher-order session communication. In: Proc. of SecReT'06. ENTCS, vol 171, issue 4, pp 73–93

Received 23 February 2012
Revised 18 March 2013
Accepted 18 March 2013 by J. Parrow
Published online 6 August 2013