

# A Basic Calculus for Modeling Service Level Agreement<sup>★</sup>

Rocco De Nicola<sup>1</sup>, Gianluigi Ferrari<sup>2</sup>, Ugo Montanari<sup>2</sup>,  
Rosario Pugliese<sup>1</sup>, and Emilio Tuosto<sup>2</sup>

<sup>1</sup> Dipartimento di Sistemi e Informatica,  
Università di Firenze, Via C. Lombroso 6/17, 50134 Firenze – Italy  
email: {denicola,pugliese}@dsi.unifi.it

<sup>2</sup> Dipartimento di Informatica,  
Largo Pontecorvo 1, 56127 Pisa – Italy  
email: {giangi,ugo,etuosto}@di.unipi.it

**Abstract.** The definition of suitable abstractions and models for identifying, understanding and managing Quality of Service (QoS) constraints is a challenging issue of the Service Oriented Computing paradigm. In this paper we introduce a process calculus where QoS attributes are first class objects. We identify a minimal set of primitives that allow capturing in an abstract way the ability to control and coordinate services in presence of QoS constraints.

## 1 Introduction

*Service Oriented Computing* (SOC) [14] has been proposed as an evolutionary paradigm to build wide area distributed systems and applications. In this paradigm, services are the basic building blocks of applications. Services are heterogeneous software components which encapsulate resources and deliver functionalities. Services can be dynamically composed to provide new services, and their interaction is governed in accordance with programmable coordination policies. Examples of SOC architectures are provided by WEB services and GRID services.

The SOC paradigm has to face several challenges like service composition and adaptation, negotiation and agreement, monitoring and security. A key issue of the paradigm is that services must be delivered in accordance with *contracts* that specify both client requirements and service properties. These contracts are usually called Service Level Agreements (SLA). SLA contracts put special emphasis on Quality of Service (QoS) described as a set of non functional properties concerning issues like response time, availability, security, and so on.

The actual metric used for evaluating QoS parameters is heavily dependent on the chosen level of abstraction. For instance, when designing network infrastructures, performance (with some probabilistic guarantees) is the main QoS metric. When describing multimedia applications, visual and audible qualities would be the crucial parameters. Instead, for final users, the perceived QoS is not just a matter of performance

---

<sup>★</sup> Work partially supported by EU-FET Project AGILE, EU-FET Project MIKADO, EU-FET Project PROFUNDIS, and MIUR project SP4 Architetture Software ad Alta Qualità di Servizio per Global Computing su Cooperative Wide Area Networks.

but also involves availability, security, usability of the required services. Moreover, the user would like to have a certain control on QoS parameters in order to customize the invoked services, while network providers would like to have a strict control over services. The resolution of this tension will be inherently dynamic depending on the run-time context.

In our view, it is of fundamental importance to develop formal machineries to describe, compose and relate the variety of QoS parameters. Indeed, the formal treatment of QoS parameters would contribute to the goal of devising robust programming mechanisms and the corresponding reasoning techniques that naturally support the SOC paradigm. In this paper we face this issue by introducing a process calculus where QoS parameters are used to control behaviours, i.e. QoS parameters are first class objects.

The goal of the present paper is to identify a minimal set of constructs that provide an abstract model to control and coordinate services in presence of QoS constraints. This differentiates our proposal from other approaches. In particular, process calculi have been designed to model QoS in terms of performance issues (e.g., the probabilistic  $\pi$ -calculus [15]). Other process calculi have addressed the issues of failures and failure detection [13]. Process calculi equipped with powerful type systems have also been put forward to describe the behavioral aspects of contracts [11, 10, 9].

Some preliminary results towards the direction of this paper can be found in [3, 4]. Cardelli and Davies [3] introduced a calculus which incorporates a notion of communication rate (bandwidth) together with some programming constructs. In [17, 8, 4] a (hyper)graph model to control explicitly QoS attributes has been introduced. The graphical semantics allows us to describe interactions in accordance with the agreed QoS level as optimal paths in the model thus creating a bridge between formal models and the protocols used in practice. Here, we elaborate on [4] with the aim of bridging further the gap between formal theories and the pragmatics of software development.

Fundamental to our approach is the notion of QoS values; a QoS value is a *tuple* of values and each component of the tuple indicates a QoS dimension. The values of the fields can be of different kind, for instance, the value along the latency dimension could be a numerical value but the security values could have the form of sets of capabilities indicating the permissions to perform some operations on given resources, e.g. read or write a file. Compositionality of QoS values is therefore a key element of our approach: the composition of QoS values will be a QoS value as well. Indeed, one might want to build a QoS value based on latency, availability and access rights of a service.

To guarantee compositionality of QoS parameters, we shall require QoS values to be elements of suitable algebraic structures, called *constraint semirings* (c-semirings, for short), consisting of a domain and two operations, the *additive* (+) and the *multiplicative* ( $\cdot$ ) operations, satisfying some properties. The basic idea is that the former is used to select among values and the latter to combine values. C-semirings were originally proposed to describe and program constraints problems [2]. Several semirings have been proposed to model QoS issues. For instance, general algorithms for computing shortest paths in a weighted directed graph are based on the structure of semirings [12]. The modelling of trustness in an ad hoc networks exploits the semiring structure [16]. C-semiring based methods have a unique advantage when problems with multiple

QoS dimensions must be tackled. In fact, it turns out that cartesian products, exponentials and power constructions of c-semirings are c-semirings as well.

Our process calculus,  $\mathcal{KoS}$ , builds on  $\mathcal{K}$  (*Kernel Language for Agent Interaction and Mobility*) [5].  $\mathcal{K}$  is an experimental kernel programming language specifically designed to model and program wide area network applications with located services and mobility.  $\mathcal{K}$  naturally supports a *peer-to-peer* programming model where interacting peers (nodes in  $\mathcal{K}$  terminology) cooperate to provide common sets of services.  $\mathcal{KoS}$  primitives handle QoS values as first class entities. For instance, an overlay network is specified by creating nodes ( $node_\kappa(t)$ ) and new links ( $s \xrightarrow{\kappa} t$ ) and indexing them with the QoS value  $\kappa$  of the operation. Thus, for instance the expression  $s \xrightarrow{\kappa} t$  states that  $s$  and  $t$  are connected by a link whose QoS parameters are given by  $\kappa$ .

The operational semantics of  $\mathcal{KoS}$  ensures that the QoS values are respected during system evolution. Suppose for example that node  $s$  would interact by an operation whose QoS value is  $\kappa'$  with node  $t$  along the link  $s \xrightarrow{\kappa} t$ . This interaction will be allowed provided that the SLA contract of the link is satisfied, namely,  $\kappa' \leq \kappa$ .

We shall illustrate the expressiveness of the calculus through several examples. This can appear as an exercise in coding a series of linguistic primitives into our calculus notation, but it yields much more because the encodings offer a practical illustration of how to give a precise semantic interpretation of QoS management. Indeed, the main contribution of this paper is the careful investigation of a minimal conceptual model which provides the basis to design programming constructs for SOC. We focus on the precise semantic definition of the calculus because it is a fundamental step to design programming primitives together with methods supporting the correct usages of the primitives and the formal verification of the resulting applications.

The rest of the paper is organized as follows. In the next section we illustrate a motivating example and, in Section 3, we introduce syntax and semantic of  $\mathcal{KoS}$ . In Section 4, we deal with expressivity issues and in the subsequent one we present a more complex scenario and show how it can be tackled by following our approach.

## 2 A motivating example

Before introducing the formal definition of  $\mathcal{KoS}$ , we prefer to show its usefulness by modelling a realistic, but simplified, example. Our purpose here is to give a flavour of the underlying programming paradigm. We consider a scenario where  $n$  servers provide services to  $m$  clients and we focus on balancing the load of the servers. Clients and servers are located on different nodes; a generic client node has address  $c_i$  while a generic server node has address  $s_j$ . Clients issue requests to servers by spawning process  $R$  from their node to a server node. For simplicity, we abstract from the actual structures of QoS values, and we assume that clients and servers “knows” each other and cannot be created dynamically. Adding dynamicity is straightforward.

A generic client node  $M_i$ , for  $i \in \{1, \dots, m\}$ , is described by the following term:

$$M_i \stackrel{\text{def}}{=} c_i :: \langle s_1, \kappa_1 \rangle \mid \dots \mid \langle s_n, \kappa_n \rangle \mid !C_\delta.$$

Intuitively,  $M_i$  represents a network component with address  $c_i$ , containing tuples of the form  $\langle s_j, \kappa_j \rangle$ , for  $j \in \{1, \dots, n\}$ , and running process  $!C_\delta$ . Each tuple  $\langle s_j, \kappa_j \rangle$  represents

the load  $\kappa_j$  of the server  $s_j$  that the client perceives, thus the whole set of tuples represents a sort of *directory service* containing the SLA contract with the available servers. Operator  $!$  is the replication operator:  $!C_\delta$  represents an unbounded number of concurrent copies of process  $C_\delta$ . Finally, process  $C_\delta$  specifies the behaviour of the client and is defined as follows:

$$C_\delta \stackrel{\text{def}}{=} (?u, ?v). \varepsilon_v[R]@u. \text{con}_{v,\delta}\langle u \rangle. \langle u, v \cdot \delta \rangle.$$

Initially, the client selects a server by non-deterministically inputting a tuple by means of the operation  $(?u, ?v)$ . Once the input is executed, variables  $u$  and  $v$  are instantiated with the server name and its load, respectively. Afterward, the client tries to spawn process  $R$  to the selected server  $u$ . Execution of  $\varepsilon_v[R]@u$  takes place only if a “suitable” link toward  $u$  exists. What here is meant for “suitable” is that the load  $v$  of the client must not exceed the value on the link. Then, since remote spawning consumes the links traversed during the migration, the client attempts to re-establish a connection with  $u$  by executing  $\text{con}_{v,\delta}\langle u \rangle$ . Notice that the operation  $\text{con}_{v,\delta}\langle u \rangle$  is used by the client to ask for a link with a QoS value increased of a quantity  $\delta$ . Once the connection has been established, the client updates its SLA view of the servers load by inserting tuple  $\langle u, v \cdot \delta \rangle$  into its local directory service.

A generic server  $N_j$ , for  $j \in \{1, \dots, n\}$ , is described as follows:

$$N_j \stackrel{\text{def}}{=} s_j :: \langle h \rangle \mid \langle c_1, \kappa'_1 \rangle \mid \dots \mid \langle c_m, \kappa'_m \rangle \mid \\ !(S \ c_1 \ s_j) \mid \dots \mid !(S \ c_m \ s_j).$$

Similarly to clients,  $N_j$  encapsulates a directory service containing SLA data about the clients. This directory service is formed by tuples of the form  $\langle c_i, \kappa'_i \rangle$ , for  $i \in \{1, \dots, m\}$ , each recording the QoS value  $\kappa'_i$  assigned to the link towards node  $c_i$ , and by the current load of the server, represented by a tuple containing a natural number  $\langle h \rangle$ . For any client  $c_i$  there is a *load manager*  $S \ c_i \ s_j$  which decides whether a link with  $c_i$  can be re-established or not. Process  $S \ c \ s$  is written as follows:

$$S \ c \ s \stackrel{\text{def}}{=} (?l). \langle l \rangle. \text{If}_s \ l < \max \\ \text{then } (c, ?v). \text{acc}_{f(v,l)}\langle c \rangle. \langle c, f(v, l) \rangle.$$

The load manager repeatedly acquires the tuple  $\langle h \rangle$  (current load) and compares it with the maximum admissible load (max). Then, the process decides whether to accept requests for new connections coming from the client: the link is created only when  $h$  is less than max. The QoS value of the new link is computed by a function  $f$  and depends on both the old QoS value and the current load.

Finally, we assume that process  $R$  representing clients service requests is a sequential process of the form

$$R \stackrel{\text{def}}{=} (?x). \langle x + 1 \rangle \dots \text{actual request} \dots (?y). \langle y - 1 \rangle,$$

Namely,  $R$  has a prologue and an epilogue which respectively increments and decrements the counter that measures the server load.

### 3 The calculus

This section introduces  $\mathcal{KoS}$  a calculus that provides a set of basic primitives for modelling and managing QoS values. A  $\mathcal{KoS}$  term represents a net made of *nodes* which model places where computations take place or where services can be allocated/accessed. We assume as given a set of nodes  $\mathcal{S}$  (ranged over by  $s, t, \dots$ ) that are connected by *links* representing the middleware infrastructure, i.e., the interactions between two nodes can take place only if they are connected by a sequence of links. Links are weighted by “measures” that represent the QoS value of the connections.

#### 3.1 QoS Values as constraint semirings

We assume existence of a set of *QoS values*  $C$ , ranged over by  $\kappa$ , that forms a *constraint semiring* [2] (*c-semiring*).

**Definition 1 (C-semiring).** *An algebraic structure  $\langle A, +, \cdot, \mathbf{0}, \mathbf{1} \rangle$  is a c-semiring if  $A$  is a set ( $\mathbf{0}, \mathbf{1} \in A$ ), and  $+$  and  $\cdot$  are binary operations on  $A$  that satisfy the following properties:*

- $+$  (additive operation) is commutative, associative, idempotent,  $\mathbf{0}$  is its unit element and  $\mathbf{1}$  is its absorbing element;
- $\cdot$  (multiplicative operation) is commutative, associative, distributes over  $+$ ,  $\mathbf{1}$  is its unit element, and  $\mathbf{0}$  is its absorbing element.

Operation  $+$  induces a partial order on  $A$  defined as  $a \leq_A b \iff a + b = b$ . The minimal element is thus  $\mathbf{0}$  and the maximal  $\mathbf{1}$ .  $a \leq_A b$  means that  $a$  is more constrained than  $b$ .

An example of c-semiring is  $\langle \omega, \min, +, +\infty, 0 \rangle$ , where  $\omega$  is the set of natural numbers, the minimum between natural numbers is the additive operation and the sum over natural numbers is the multiplicative operation. Notice that in this case the partial order induced by the additive operations is the inverse of the ordinary total order on natural numbers. Another example of c-semiring is  $\langle \wp(\{A\}), \cup, \cap, \emptyset, A \rangle$ , where  $\wp(A)$  is the powerset of a set  $A$ , and  $\cup$  and  $\cap$  are the usual set union and intersection operations.

$\mathcal{KoS}$  does not take a definite standing on which of the many c-semiring structures to use. The appropriate c-semiring to work with should be chosen, from time to time, depending on the kind of QoS dimensions one intends to model. Below, we introduce some c-semiring structures together with the QoS dimension they handle:

- $\langle \{true, false\}, \vee, \wedge, false, true \rangle$  (boolean): Network and service availability.
- $\langle \text{Real}^+, \min, +, +\infty, 0 \rangle$  (optimization): Price, propagation delay.
- $\langle \text{Real}^+, \max, \min, 0, +\infty \rangle$  (max/min): Bandwidth.
- $\langle [0, 1], \max, \cdot, 0, 1 \rangle$  (probabilistic): Performance and rates.
- $\langle [0, 1], \max, \min, 0, 1 \rangle$  (fuzzy): Performance and rates.
- $\langle 2^N, \cup, \cap, \emptyset, N \rangle$  (set-based, where  $N$  is a set): Capabilities and access rights.

C-semiring based methods have a unique advantage when problems with multiple QoS criteria must be tackled. In fact, it turns out that cartesian products, exponentials and power constructions of c-semirings are c-semirings as well.

$N, M ::=$	<b>N</b>	
$\mathbf{0}$		<i>Empty net</i>
$  s :: P$		<i>Located Process</i>
$  s \xrightarrow{\kappa} t$		<i>Link</i>
$  (v\ s)N$		<i>Node restriction</i>
$  N \parallel M$		<i>Net composition</i>
$P, Q ::=$	<b>P</b>	
$\mathbf{0}$		<i>Null process</i>
$  \gamma.P$		<i>Action prefixing</i>
$  (v\ s)P$		<i>Restriction</i>
$  P \mid Q$		<i>Parallel process</i>
$  !P$		<i>Iteration</i>
$\gamma ::=$	<b>P</b>	
$node_{\kappa}\langle t \rangle$		<i>Node creation</i>
$  con_{\kappa}\langle t \rangle$		<i>Connection request</i>
$  acc_{\kappa}\langle t \rangle$		<i>Connection acceptance</i>
$  (T)$		<i>Input</i>
$  \langle v_1, \dots, v_n \rangle$		<i>Output</i>
$  \varepsilon_{\kappa}[P]@t$		<i>Remote process spawning</i>
$T ::= \varepsilon \mid v \mid ?x \mid \neg v \mid T, T$	<b>I</b>	

**Table 1.**  $\mathcal{KoS}$  Syntax

### 3.2 Syntax

The syntax of  $\mathcal{KoS}$  is presented in Table 1. Other than the existence of  $C$ , existence of a set of *names*  $\mathcal{N}$  (ranged over by  $r, s$  and  $t$ ) is assumed. First-class *values*, ranged over by  $u$  and  $v$ , can be either QoS values or names.

The syntax for nets permits the (possibly empty) parallel composition of *located processes* and *links*. A located process  $s :: P$  consists of a name  $s$ , called the *address* of  $P$ , and the process  $P$  running at  $s$ . A link  $s \xrightarrow{\kappa} t$  states that  $s$  and  $t$  are connected by a link whose QoS value is  $\kappa$ . The net  $(v\ s)N$  is a net that declares  $s$  as restricted in  $N$ , which is the scope of the restriction.

The syntax for processes is standard. The symbol  $\mathbf{0}$  overloads the symbol for empty nets; however, the contexts will clarify whether it refers to processes or nets. Prefixes  $\gamma$  encompass actions for

- creating a node ( $node_{\kappa}\langle t \rangle$ ) or a connection to/from another node ( $con_{\kappa}\langle t \rangle, acc_{\kappa}\langle t \rangle$ ),
- exchanging tuples of values ( $(T)$  and  $\langle v_1, \dots, v_n \rangle$ ),
- remotely spawning a process ( $\varepsilon_{\kappa}[P]@t$ ).

Links are oriented, indeed  $s \xrightarrow{\kappa} t$  allows a process to be spawned from  $s$  to  $t$  but not the viceversa. The creation of new links is obtained by synchronising actions  $con_{\kappa}\langle t \rangle$  and  $acc_{\kappa'}\langle s \rangle$  performed at  $s$  and  $t$ , respectively.

Communication involves exchange of tuples (i.e. finite sequences) of values that are retrieved via pattern matching. Input prefixes use *templates*  $T$ , namely finite sequences

of values or placeholders (written as  $?x$ ). Execution of an output prefix causes generation of a tuple of values  $v_1, \dots, v_n$ . Both the empty template and the empty tuple are denoted by  $\varepsilon$ . Hereafter, we let  $\mathbf{t}$  range over tuples of values and, given a template  $T$  and a tuple  $\mathbf{t}$ , we let  $T_i$  and  $\mathbf{t}_i$  denote the  $i$ -th element of  $T$  and  $\mathbf{t}$ , respectively.

The placeholder  $?x$  binds the occurrences of  $x$  in the rest of the template, namely, in  $?x.T$ , the scope of  $?x$  is  $T$ . The set  $\text{bn}(T)$  collects the names bound in  $T$  while  $\text{fn}(T)$  denotes the names having free occurrences in  $T$ ; their definitions are standard. We consider as equivalent those templates that differ only for renaming of bound names. The template  $\neg v$  tests for inequality, namely, it requires the matching tuple to contain a value different from  $v$  (see Definition 7). The only binders of the calculus are the placeholder  $?x$  and the node restriction  $\nu s$ . Note that node names might be QoS values (e.g., for specifying access rights), hence, we write  $\text{fn}(\kappa)$  to denote the names appearing in  $\kappa$ . Moreover, we require that QoS values do not bind node names, therefore,  $\text{bn}(\kappa)$  is empty, for any QoS value  $\kappa$ . We formally define *free* and *bound* names of nets and processes as follows. In the following we write  $\text{fn}(\_, \_)$  (resp.  $\text{bn}(\_, \_)$ ) as an abbreviation for  $\text{fn}(\_) \cup \text{fn}(\_)$  ( $\text{bn}(\_) \cup \text{bn}(\_)$ , respectively).

**Definition 2 (Free and bound names).** *The free names of prefix actions are defined as expected:  $\text{fn}(\gamma) = \text{fn}(\kappa) \cup \{s\}$ , if  $\gamma \in \{\text{node}_\kappa\langle s \rangle, \text{con}_\kappa\langle s \rangle, \text{acc}_\kappa\langle s \rangle\}$ ,  $\text{fn}((T)) = \text{fn}(T)$ ,  $\text{fn}(\langle v_1, \dots, v_n \rangle) = \text{fn}(v_1) \cup \dots \cup \text{fn}(v_n)$  and  $\text{fn}(\varepsilon_\kappa[P]@s) = \text{fn}(\kappa, P) \cup \{s\}$ . Bound names of  $\gamma$  are defined similarly, e.g.,  $\text{bn}((T)) = \text{bn}(T)$  and  $\text{bn}(\varepsilon_\kappa[P]@s) = \text{bn}(P)$  (while in the remaining cases is the empty set).*

*The sets  $\text{fn}(\_)$  and  $\text{bn}(\_)$  of free and bound names of processes and nets are defined accordingly. The only non-standard case is that for links where we let  $\text{fn}(r \xrightarrow{\kappa} s) = \text{fn}(\kappa) \cup \{s, r\}$  and  $\text{bn}(r \xrightarrow{\kappa} s) = \emptyset$ .*

As usual, processes or nets obtained by  $\alpha$ -converting bound names are considered equivalent. Moreover, we assume the following structural congruence laws.

**Definition 3 (Structural congruence).** *The relation  $\equiv_P \subseteq P \times P$  is the least equivalence relation on processes (containing  $\alpha$ -conversion and) satisfying the following axioms:*

- $(P, \mid, \mathbf{0})$  is a commutative monoid;
- $!P \equiv_P P \mid !P$ .

*The relation  $\equiv_N \subseteq N \times N$  is the least equivalence relation on nets (containing  $\alpha$ -conversion and) satisfying the following axioms:*

- $(N, \parallel, \mathbf{0})$  is a commutative monoid;
- if  $P \equiv_P Q$  then  $s :: P \equiv s :: Q$ ;
- $s :: P \mid Q \equiv s :: P \parallel s :: Q$ ;
- $s :: (\nu t)P \equiv (\nu t)(s :: P)$ , if  $t \neq s$ ;
- $(\nu s)(N \parallel M) \equiv N \parallel (\nu s)M$ , if  $s \notin \text{fn}(N)$ ;
- $(\nu s)(\nu t)N \equiv (\nu t)(\nu s)N$ .

The last axiom of Definition 3 states that the order of the restrictions is irrelevant, hence we can write  $(\nu s_1, \dots, s_n)N$  instead of  $(\nu s_1) \dots (\nu s_n)N$ .

### 3.3 Semantics

We define the operational semantics of  $\mathcal{KoS}$  by means of a labelled transition system that describes the evolution of nets. In the semantic clauses, it is useful to define a function that, given a net  $N$ , yields the names that are used as node addresses in the net.

**Definition 4 (Addresses).** Let  $\text{addr}$  be the function given by:

$$\text{addr}(N) = \begin{cases} \emptyset, & N = \mathbf{0} \vee N = s \xrightarrow{\kappa} t \\ \{s\}, & M = s :: P \\ \text{addr}(M) \setminus \{s\}, & N = (\nu s)M \\ \text{addr}(N_1) \cup \text{addr}(N_2), & N = N_1 \parallel N_2. \end{cases}$$

Notice that  $\text{addr}(N) \subseteq \text{fn}(N)$ , but not necessarily  $\text{addr}(N) = \text{fn}(N)$ , for instance if  $N = s :: \langle t \rangle. \mathbf{0}$  then  $\text{fn}(N) = \{s, t\}$  while  $\text{addr}(N) = \{s\}$ . Basically,  $\text{addr}(N)$  collects those free names of  $N$  that effectively occur in  $N$  as address of some node.

**Definition 5 (Localized Actions).** Let  $\gamma$  be a prefix, then the localized prefix  $\gamma@s$  is defined as follows:

$$\gamma@s = \begin{cases} s \varepsilon_\kappa \langle P \rangle @t & \text{if } \gamma = \varepsilon_\kappa \langle P \rangle @t \\ s \gamma & \text{otherwise} \end{cases}$$

The syntax of localized actions  $\alpha$  is given below:

$$\alpha ::= \gamma@s \mid s \text{ link } t \mid \tau$$

We let  $\text{fn}(\gamma@s) = \text{fn}(\gamma) \cup \{s\}$  and  $\text{bn}(\gamma@s) = \text{bn}(\gamma)$ .

**Definition 6 (Nets semantics).** The operational semantics of nets is given by the relation  $\rightarrow \subseteq N \times (\alpha \times C) \times N$ . Relation  $\rightarrow$  is defined by the rules in Table 2 and the following standard rules:

$$\begin{array}{ll} \text{( ) } \frac{N \xrightarrow{\tau} M}{(\nu s)N \xrightarrow{\tau} (\nu s)M} & \text{( ) } \frac{N \equiv N' \xrightarrow{\alpha} M' \equiv M}{N \xrightarrow{\alpha} M} \\ \text{( ) } \frac{N \xrightarrow{\alpha} N'}{N \parallel M \xrightarrow{\alpha} N' \parallel M} \text{ if } \begin{cases} \text{bn}(\alpha) \cap \text{fn}(M) = \emptyset \wedge \\ (\text{addr}(N') \setminus \text{addr}(N)) \cap \text{addr}(M) = \emptyset \end{cases} & \end{array}$$

Intuitively,  $N \xrightarrow[\kappa]{\alpha} M$  states that the net  $N$  can perform the transition  $\alpha$  to  $M$  by exposing the QoS value  $\kappa$ . Clearly, all local transitions (communications, node or link creations) have unitary QoS value, while the only non-trivial QoS values appear on the transitions that spawn processes or show the presence of links. Let us give more detailed comments on the rules in Table 2.

Rule ( ) states that a link within a net disappears once it has been used. These transitions are used in the premises of rules ( ) and ( ) for establishing a path between two nodes such that a remote evaluation can take place.

Rule ( ) accounts for action prefixing; node creation, however, deserves a specific treatment that is defined in rule ( ). The side condition of ( ) also states that



---

( )	$s \xrightarrow{\kappa} t \xrightarrow{s \text{ link } t}_{\kappa} \mathbf{0}$
( )	$s :: \gamma.P \xrightarrow[\mathbf{1}]{\gamma@s} s :: P, \gamma \notin \{node_{\kappa}(t), con_{\kappa}(s), acc_{\kappa}(s)\}$
( )	$s :: node_{\kappa}(t).P \xrightarrow[\mathbf{1}]{node(t)} s :: P \parallel s \xrightarrow{\kappa} t \parallel t :: \mathbf{0}, s \neq t$
( )	$\frac{N \xrightarrow[\mathbf{1}]{s \text{ con}_{\kappa}(t)} N' \quad M \xrightarrow[\mathbf{1}]{t \text{ acc}_{\kappa'}(s)} M'}{N \parallel M \xrightarrow[\mathbf{1}]{\tau} N' \parallel M' \parallel s \xrightarrow{\kappa} t} \kappa \leq \kappa'$
( )	$s :: \varepsilon_{\kappa}[Q]@s.P \xrightarrow[\mathbf{1}]{\tau} s :: P \parallel s :: Q$
( )	$\frac{N \xrightarrow[\kappa']{r \varepsilon_{\kappa}^s(P)@t} N' \quad M \xrightarrow[\kappa'']{r \text{ link } r'} M' \quad \kappa' \cdot \kappa'' \leq \kappa}{N \parallel M \xrightarrow[\kappa' \cdot \kappa'']{r' \varepsilon_{\kappa}^s(P)@t} N' \parallel M'} , t \neq r'$
( )	$\frac{N \xrightarrow[\kappa']{r \varepsilon_{\kappa}^s(P)@t} N' \quad M \xrightarrow[\kappa'']{r \text{ link } t} M' \quad \kappa' \cdot \kappa'' \leq \kappa}{N \parallel M \xrightarrow[\kappa' \cdot \kappa'']{\tau} N' \parallel M' \parallel t :: P}$
( )	$\frac{N \xrightarrow[\mathbf{1}]{s(T)} N' \quad M \xrightarrow[\mathbf{1}]{s \text{ t}} M' \quad \bowtie(T, \mathbf{t}) = \sigma}{N \parallel M \xrightarrow[\mathbf{1}]{\tau} N' \sigma \parallel M'}$

---

**Table 2.** Network semantics

no link from  $s$  to itself can be created. Indeed, we assume that transitions that involve only the local node have unitary QoS value and are always enabled.

Rule ( ) allows a process allocated at  $s$  to use a name  $t$  as the address of a new node and to create a new link from  $s$  to  $t$  exposing the QoS value  $\kappa$ . The side condition of ( ) prevents that new nodes (and links) are created by using addresses of existing nodes.

Rule ( ) adds a new link between two existing addresses  $s$  and  $t$ ; the link is created only if the processes at  $s$  and  $t$  satisfy the SLA contract. More precisely, the accepting node  $t$  is willing to connect only to those nodes that declare a QoS value lower than  $\kappa'$ . If this condition holds, a new link is added to the net, such link has the QoS value exposed by  $s$ . One can think of  $s$  as asking for the connection with *at least* some characteristics

expressed by  $\kappa$  and  $t$  establishes the connection only when it can enforce the requirement of  $s$ , namely  $\kappa \leq \kappa'$ .

Rule ( ) states that the local spawning of a process is always enabled while rules ( ) and ( ) control process migration and require more detailed explanations. A remote spawning action  $\varepsilon_\kappa[P]@t$  consists of the migrating process  $P$ , the arrival node  $t$  and a QoS value  $\kappa$  expressing that  $P$  must be routed on a path exposing a QoS value<sup>1</sup> at most  $\kappa$ . Differently from the local spawning of processes, remote spawning is not always possible, it is indeed mandatory that the net contains a path of links from the starting node  $s$  to the arrival node  $t$ . Moreover, the SLA contract of the path between  $s$  and  $t$  must not exceed the value  $\kappa$  that the spawner has declared. Notice that this semantically describes the SLA agreement on the mobility of processes. This is formally achieved by rules ( ) and ( ). More specifically, rule ( ) states that, if the migrating process can go through an intermediate node  $r$  and a link from  $r$  to a node  $r' \neq t$  exists, the QoS value  $\kappa'$  of the partial path from  $s$  to  $r$  composed with the value  $\kappa''$  of the link from  $r$  to  $r'$  must be lower than  $\kappa$ . If this is the case, a transition can be inferred stating that  $P$ , spawned from  $s$ , can go through  $r'$  exposing the QoS value  $\kappa' \cdot \kappa''$ . Rule ( ) is similar to ( ) but describes the last hop of  $P$ , namely when the target node  $t$  is reached. In this case,  $P$  is spawned at  $t$ , provided that the QoS value of the whole path that has been found is lower than  $\kappa$ .

Rule ( ) establishes that a synchronization takes place provided that sender and receiver are allocated at the same node and that the template and the tuple match according to the definition below. Hereafter, we use  $\sigma$  to denote a substitution, i.e. a map from names to names and QoS values, and  $\sigma[\sigma']$  to denote the composition of substitutions, i.e. the substitution  $\sigma''$  defined as follows:  $\sigma''(x) = \sigma'(x)$  if  $x \in \text{dom}(\sigma')$ ,  $\sigma''(x) = \sigma(x)$  if  $x \in \text{dom}(\sigma) - \text{dom}(\sigma')$ .

**Definition 7 (Pattern matching).** A template  $T$  and a tuple  $\mathbf{t}$  match when the following function is defined

$$\bowtie(T, \mathbf{t}) = \begin{cases} \varepsilon & \text{if } (T = \varepsilon \wedge \mathbf{t} = \varepsilon) \vee (T = v \wedge \mathbf{t} = v) \\ \varepsilon & \text{if } T = \neg v \wedge \mathbf{t} = v' \wedge v \neq v' \\ \{v/x\} & \text{if } T = ?x \wedge \mathbf{t} = v \\ \sigma[\sigma'] & \text{if } T = F, T' \wedge \mathbf{t} = v, \mathbf{t}' \wedge \bowtie(F, v) = \sigma \wedge \bowtie(T', \mathbf{t}') = \sigma' \end{cases}$$

where the application of a substitution to a template,  $T\sigma$ , is defined as follows:

$$T\sigma = \begin{cases} \varepsilon & \text{if } T = \varepsilon \\ v, T'\sigma & \text{if } T = x, T' \wedge \sigma(x) = v \\ x, T'\sigma & \text{if } T = x, T' \wedge x \notin \text{dom}(\sigma) \\ ?x, T'\sigma\{x/x\} & \text{if } T = ?x, T'. \end{cases}$$

Under the conditions of ( ), the substitution  $\bowtie(T, \mathbf{t})$  is applied to the receiver. Note that  $\bowtie$  may not be defined, for instance  $\bowtie(\neg s, s)$  does not yield any substitution and, therefore, the match in such a case does not hold.

<sup>1</sup> The QoS value of a path  $s_0 \xrightarrow{\kappa_1} s_1 \dots s_{n-1} \xrightarrow{\kappa_n} s_n$  is defined as  $\kappa_1 \cdot \dots \cdot \kappa_n$ .

## 4 Examples

In this section we present some specification examples. To make the presentation more readable let us introduce some notational conventions. First, we avoid writing trailing  $0$  processes, second, we write  $\varepsilon[P]@r$  instead of  $\varepsilon_1[P]@r$  and similarly for  $node_1\langle t \rangle$ ,  $con_1\langle t \rangle$  and  $acc_1\langle t \rangle$ .

*Boolean expressions* Booleans are encoded as processes that allocate a pair of names to a node:

$$\begin{aligned} True\ r &\stackrel{\text{def}}{=} (\nu t)\varepsilon[\langle t, t \rangle]@r \\ False\ r &\stackrel{\text{def}}{=} (\nu f, f')\varepsilon[\langle f, f' \rangle]@r. \end{aligned}$$

The truth and the falsity are tested by checking that the names in a pair are equal or different, respectively. The following process tests for the equality of two names:

$$Test\ x\ y\ r \stackrel{\text{def}}{=} (\nu t)(node\langle t \rangle.\varepsilon[Eval\ y\ r \mid \langle x \rangle]@t),$$

where  $Eval\ y\ r \stackrel{\text{def}}{=} (y).True\ r \mid (\neg y).False\ r$ . Process *Test* spawns the tuple  $\langle x \rangle$  and the *Eval* process onto a newly generated node so that the first or the second component of *Eval* have exclusive access to  $\langle x \rangle$ . Notice that only one of the components can consume the tuple, indeed, either  $x = y$  (and only the pattern  $(y)$  matches  $\langle x \rangle$ ) or  $x \neq y$  (and only the pattern  $(\neg y)$  matches  $\langle x \rangle$ ). Finally, *True* or *False* allocates on node  $r$  the truth value corresponding to evaluation of  $x = y$ . Assuming the encoding of booleans, we can represent standard control structures such as **if-then-else** and **while**.

The encoding of boolean values is indeed an example of a standard programming metaphor for finding and handling services. Assume that we want to describe a *look-up* mechanism for discovering distributed services. For instance, the *web services* technology allows deploying new services by gluing together those that have been published. Web service composition, however, requires a look-up phase where the available service must be discovered. In the boolean example, processes *True* and *False* are the services that have been published and composed together to provide the *Test* service. Notice that the look-up phase does not require the knowledge of the service name but only that of the “schema” of the service. For instance, whenever a new “true” service is published it suffice to generate a new name and use it for building the “schema” for the true service (i.e., a pair of two equal names).

*Public, private, permanent and stable links* Links in  $\mathcal{KoS}$  are public entities: when available they can be exploited by all processes. Consider the following  $\mathcal{KoS}$  net:

$$N \stackrel{\text{def}}{=} s :: \varepsilon_3[P]@t \parallel s \overset{1}{\frown} r \parallel r :: con_2\langle t \rangle.\varepsilon_2[Q]@t \parallel t :: acc_2\langle r \rangle,$$

where QoS values are the c-semiring of natural numbers. Net  $N$  has three nodes  $s$ ,  $r$  and  $t$  and, initially, only  $s$  and  $r$  are connected by a link with QoS value 1. Node  $s$  is trying to spawn  $P$  on  $t$  which is not possible because there is no path from  $s$  to  $t$ . Node  $r$  is willing to spawn a process  $Q$  on  $t$ , as well; however,  $r$  is aware that a link must be

first created. Node  $t$  simply accepts requests for establishing a link from  $r$ . Initially, it is only possible to synchronize  $con_2\langle t \rangle$  and  $acc_2\langle r \rangle$  which, by applying rule ( ) leads to

$$N' \stackrel{\text{def}}{=} s :: \varepsilon_3[P]@t \parallel s \stackrel{1}{\curvearrowright} r \parallel r :: \varepsilon_2[Q]@t \parallel r \stackrel{2}{\curvearrowright} t \parallel t :: \mathbf{0}.$$

Now, applying rules ( ), ( ) and ( ) we derive

$$N' \xrightarrow[2]{\tau} s :: \varepsilon_3[P]@t \parallel s \stackrel{1}{\curvearrowright} r \parallel r :: \mathbf{0} \parallel t :: Q.$$

Notice that the link between  $r$  and  $t$  is consumed by the migration of  $Q$  hence  $P$  cannot reach  $t$ . However,  $N'$  can also evolve differently, in fact, both the two spawning actions are enabled, because the creation of the link between  $r$  and  $t$  has also provided a path from  $s$  to  $t$  exposing the QoS value 3. Hence, by rules ( ), ( ), ( ) and ( ) we can also derive

$$N' \xrightarrow[3]{\tau} s :: \mathbf{0} \parallel r :: \varepsilon_2[Q]@t \parallel t :: P.$$

Noteworthy, the migration of  $P$  prevents  $Q$  to be spawned because the link created by  $r$  has been used by  $P$ .

In general, this kind of interference should be avoided and this can be done in  $\mathcal{KoS}$  by expressing *private links* which can be specified by exploiting the properties of c-semirings. The intuition is that the use of a link is allowed only whether the traversing process has the appropriate “rights”. If we represent access rights as sets of names, then a process must “know” all the names needed for traversing the link. For instance, consider the following net:

$$s :: \varepsilon_{\{r,s\}}[P]@t \parallel s \stackrel{\{r\}}{\curvearrowright} s',$$

process  $P$  can traverse the link  $s \stackrel{\{r\}}{\curvearrowright} s'$  because it “knows”  $r$ , that is the only name required to traverse the link. Noteworthy,  $P$  could not traverse  $s \stackrel{\{r,u\}}{\curvearrowright} s'$  because it does not expose name  $u$ .

We consider the c-semiring  $\mathcal{R} = \langle \wp_{\text{fin}}(\mathcal{S}) \cup \{\mathcal{S}\}, glb, \cup, \mathcal{S}, \emptyset \rangle$  to represent access rights (recall that  $\mathcal{S}$  is the set of sites). It is straightforward to prove that  $\mathcal{R}$  is a c-semiring; moreover, the order induced by the additive operation of  $\mathcal{R}$  is the inverse of the set inclusion (i.e.,  $X \leq Y \iff Y \subseteq X$ ).

Therefore, a private link between the nodes  $s$  and  $t$  can be specified as

$$(\nu p)(s :: P \parallel s \stackrel{\{p\}}{\curvearrowright} t \parallel t :: Q),$$

indeed, in order to pass through link  $s \stackrel{\{p\}}{\curvearrowright} t$ , a process must exhibit the “password”  $p$ . The knowledge of  $p$  is handled by enlarging the scope of the restriction and communicating it.

We conclude by illustrating how one could implement *permanent* links, i.e. links that are always available, by exploiting replication:

$$s :: !con_k\langle t \rangle \parallel t :: !acc_{k'}\langle s \rangle$$

A slight variation are *stable* links, which are links existing until a given condition is satisfied.

$$Stable_s G t \stackrel{\text{def}}{=} !con_k\langle t \rangle \mid \varepsilon[While G do acc_{k'}\langle s \rangle od \mathbf{0}]@t$$

*Cryptography* By exploiting private links,  $\mathcal{KoS}$  can encode standard encryption/decryption mechanisms usually adopted for expressing security protocols in process calculi (see e.g. [1]). Consider the following net:

$$(\nu k, s_k)(i :: P \parallel i \xrightarrow{\{k\}} s_k \parallel s_k :: M \parallel s_k \xrightarrow{\{k\}} r \parallel r :: Q), \quad (1)$$

and assume that the only links from/to  $s_k$  are those appearing in (1). Net (1) aims at representing the initiator  $i$  and the responder  $r$  of a protocol that share a key  $k$ . According to (1) a key is modelled by means of a pair made of a name and a node which roughly speaking contains those messages that are encrypted with  $k$ .

The intuition is that encrypting corresponds to allocating a message on  $s_k$  while decrypting corresponds to the possibility of “jumping” on  $s_k$  and reading a message or, in other words, to the knowledge of  $k$  for traversing links  $i \xrightarrow{\{k\}} s_k$  or  $s_k \xrightarrow{\{k\}} r$ .

## 5 Composing Overlay Networks

We consider a scenario where a service is replicated over the nodes of an overlay network and can be invoked through a *unique* handler  $H$  that manages the requests of the clients. This kind of architectures is adopted from Internet Service Providers (ISP) that offer dial-up connection to end-users (EU). In this case a telecommunication company (TC) handles the phone overlay networks. The EU connects to the “nearest” ISP server by dialing a single (country-wide) number. The TC takes care of dispatching the call to the closest ISP server on the overlay network. There are (at least) two possible way of connecting the EU and the ISP server. Either the TC establishes a direct connection between the EU and the ISP, or the TC act as a gateway between the phone overline network and the ISP overlay network. Both solutions can be easily expressed in  $\mathcal{KoS}$  in the logical architecture of the system: the handler  $H$  manages the requests (e.g., controls the access rights of the client), looks for a suitable server, and forwards the request, while trying to balance the load of any replica of the server. Hence, the request of a client  $C$  might not be forwarded to the “best” server from the client’s point of view. In this case,  $H$  provides another server to  $C$ , however, the client may or may not commit to use it.

The simplest way to model this composed overlay network is to assume that the link between  $C$  and  $H$  have QoS values expressing the access rights of  $C$ . When a server  $s$  meeting both the request of  $C$  and the load constraints is found,  $H$  replies to  $C$  and tells  $s$  to accept a (private) link from  $C$ . Hereafter, we assume that  $h$  is the node address of  $H$ . We detail the client first:

$$\begin{aligned} C \kappa pr c &\stackrel{\text{def}}{=} (\nu r)(\varepsilon_k[\langle \text{“connect”}, c, r \rangle]@h. \\ &\quad (r, ?s, ?pr', ?p). \\ &\quad \text{If}_c pr' < pr \\ &\quad \text{then } con_{\{r,p\}}\langle s \rangle.\mathcal{E}_{\{p,r\}}[R]@s \\ &\quad \text{else } con_{\{r,p\}}\langle s \rangle.\mathcal{E}_{\{p,r\}}[\langle r, \text{“to-much”} \rangle]@s). \end{aligned}$$

Process  $C$  requests  $H$  to find a server and waits for the response. The request contains  $c$ , the node address of  $C$ , and a private name  $r$ . Name  $r$  can be thought of as the unique

marker of the request so that only  $C$  will acquire data corresponding to request  $r$ . Possibly,  $H$  returns a response (marked with  $r$ ) containing the server address  $s$ , the price  $pr'$ , and the password  $p$  for the private link. Finally,  $C$  establishes a private link with  $s$  and, depending on the price  $pr'$  required by the server, either raises its request  $R$  ('then' branch) or notifies  $s$  that the service is too expensive ('else' branch).

The definition of  $H$  requires the following auxiliary processes.

$$\begin{aligned} Rd(T) &\stackrel{\text{def}}{=} (T).\langle t_T \rangle \\ Lt_s r i &\stackrel{\text{def}}{=} Rd(r, ?j). \text{If}_s j \leq i \text{ then True else False} \end{aligned}$$

Process  $Rd(T)$  looks for a tuple matching  $T$  and immediately re-generates the consumed tuple; this is denoted by  $t_T$  which is obtained from  $T$  by removing all the '?' occurring in its placeholders. Then, process  $Lt_s r i$ , interpreting  $\langle r, v \rangle$  as a "cell" having address  $r$  and containing value  $v$ , reads the value in  $r$  and establishes if it is less than/equal to  $i$ .

$$\begin{aligned} H &\stackrel{\text{def}}{=} !(\text{"connect", } ?x, ?r).\langle r, 1 \rangle \\ &\quad \text{While}_h Lt_h r \text{ nserv} \\ &\quad \text{do} \\ &\quad \quad (r, ?i).Rd(pref(x, i), ?l, ?pr.) \\ &\quad \quad \text{If}_h l \geq \max \\ &\quad \quad \quad \text{then } \langle r, i + 1 \rangle \\ &\quad \quad \quad \text{else} \\ &\quad \quad \quad (v p)(\varepsilon[\langle \text{"newlink", } x, r, p \rangle] @ pref(x, i). \\ &\quad \quad \quad \varepsilon[\langle r, pref(x, i), pr, p \rangle] @ x. \\ &\quad \quad \quad (x, ?v).acc_{f(v, l)} \langle x \rangle. \langle x, f(v, l) \rangle. \\ &\quad \quad \quad (pref(x, i), ?l, ?pr).\langle pref(x, i), l + 1, pr \rangle) \\ &\quad \text{od } \varepsilon[\langle r, \text{"no-server-available"} \rangle] @ x. \end{aligned}$$

Process  $H$  is continuously listening for a connection request. Once such a request is issued from a client at  $x$ ,  $H$  starts scanning the server list (nserv is the number of servers). For each server  $s$ , node  $h$  contains a tuple  $\langle s, l, pr \rangle$  where  $l$  is an estimation of the load of  $s$  and  $pr$  is the price for using  $s$ . Also, for each client  $x$ ,  $h$  maintains a tuple  $\langle x, \kappa \rangle$  that reports the connection between  $H$  and  $x$  (as done in Section 2). Moreover,  $H$  uses a function  $pref$  that, given the client address  $x$  and the index  $i$ , yields the  $i$ -th server "preferred" by the client. At the  $i$ -th iteration of the while loop,  $H$  reads the information of the  $i$ -th server preferred by  $x$  and, if the load of such a server is too high, the cycle is repeated provided that more servers are left ('then' branch); otherwise, a password  $p$  for a private link is generated and communicated to both  $x$  and the selected server. The server will accept a private link creation from  $x$  so that the client owning the password  $p$  can perform a request at  $s$ . Finally,  $H$  re-establish a link with  $x$  according to the new load of the servers by exploiting function  $f$  and reflecting this changes in the tuple corresponding to  $x$  (i.e.,  $\langle x, f(v, l) \rangle$ ), as in Section 2. Indeed, the mechanism of load balancing is the one defined in Section 2, the only difference being that now  $H$  is the unique handler that manages the connections with the clients.

Given  $H$  and  $C$ , the servers must simply wait for a connection request (issued from  $H$ ) and establish the private connection with the client:

$$S \ s \stackrel{\text{def}}{=} !((\text{"newlink"}, ?x, ?r, ?p).acc_{\{r,p\}}\langle x \rangle \dots \text{wait \& execute} \dots \\ \varepsilon[(s, ?l, ?pr).\langle s, l-1, pr \rangle]@h)).$$

Once the request has been served,  $S$  simply updates the load of  $s$ .

A net where  $C$ ,  $H$  and  $S$  work can be defined as follows.

$$\parallel_{i=1,\dots,m} x_i :: C \ \kappa_i \ pr_i \ x_i \mid !acc\langle h \rangle \mid con_{\kappa'_i}\langle h \rangle \\ \parallel_{i=1,\dots,m} h :: !con\langle x_i \rangle \parallel h :: H \parallel \\ \parallel (\nu s_1, \dots, s_n)(\parallel_{j=1,\dots,n} h :: !con\langle s_j \rangle \mid \langle s_j, l_j, pr_j \rangle \parallel s_j :: S)$$

where  $\parallel_{i=1,\dots,m} N_i$  shortens  $N_1 \parallel \dots \parallel N_n$ .

The other solution touched upon at the beginning of this section can be achieved by exploiting the possibility offered by  $\mathcal{KoS}$  of “connecting” links to form paths between nodes. More precisely, instead of connecting directly the client’s node  $x$  and (the node of) the selected server  $s$ , we can connect  $h$  and  $s$  so that the client’s request at  $s$  is routed through  $h$ .

## 6 Conclusion

We have formally defined  $\mathcal{KoS}$  a process calculus that provides basic primitives to describe QoS requirements of distributed applications. We demonstrated the applicability of the approach by specifying some expressive case studies.

Our research program is to provide a solid foundation to drive the design of languages and middleware having application-oriented QoS mechanisms. The work reported here is a preliminary step in this direction. In terms of calculus design, the current definition of  $\mathcal{KoS}$  assumes that links are the basic construct to manage QoS interactions and cooperation. This is a reasonable assumption for several cases. For instance, in this paper we handled the QoS composition between different overlay networks by suitable links. However, one could interpret QoS composition of overlay networks in a more general sense than adding suitable links. An interesting challenge for future research is to extend  $\mathcal{KoS}$  with more general mechanisms for composing overlay networks than simple parallel composition via links.

There are a number of ways in which our setting can be extended. For instance, it would be interesting to develop type systems which would allow determining QoS properties of processes. We plan to extend types for access control of [7, 6] to deal with QoS attributes. In particular, it would be interesting to exploit such types to capture the notion of contract. Another direction for future research is developing observational semantics for  $\mathcal{KoS}$  based on the idea of observing QoS values. These abstract theories could permit reasoning on  $\mathcal{KoS}$  nets and comparing them on the basis of the perceived QoS values.

## References

1. M. Abadi and A. Gordon. A Calculus for Cryptographic Protocols: The Spi Calculus. *Information and Computation*, 148(1):1–70, January 1999.
2. S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based constraint satisfaction and optimization. *Journal of the ACM*, 44(2):201–236, March 1997.
3. L. Cardelli and D. Rowan. Service combinators for web computing. *Software Engineering*, 25(3):309–316, 1999.
4. R. De Nicola, G. Ferrari, U. Montanari, R. Pugliese, and E. Tuosto. A Formal Basis for Reasoning on Programmable QoS. In N. Dershowitz, editor, *International Symposium on Verification – Theory and Practice – Honoring Zohar Manna’s 64th Birthday*, volume 2772 of *Lecture Notes in Computer Science*, pages 436 – 479. Springer-Verlag, 2003.
5. R. De Nicola, G. Ferrari, and R. Pugliese. KLAIM: A kernel language for agents interaction and mobility. *IEEE/ACM Transactions on Networking*, 24(5):315–330, 1998.
6. R. De Nicola, G. Ferrari, and R. Pugliese. Programming access control: The KLAIM experience. In *International Conference in Concurrency Theory*, *Lecture Notes in Computer Science*. Springer-Verlag, 2000.
7. R. De Nicola, G. Ferrari, R. Pugliese, and B. Venneri. Types for access control. *Theoretical Computer Science*, 240(1):215–254, June 2000.
8. G. Ferrari, U. Montanari, and E. Tuosto. Graph-based Models of Internetworking Systems. In T. Aichernig, Bernhard K. Maibaum, editor, *Formal Methods at the Crossroads: from Panaces to Foundational Support*, volume 2757 of *Lecture Notes in Computer Science*, pages 242 – 266. Springer-Verlag, 2003.
9. A. Igarashi and N. Kobayashi. A generic type system for the Pi-calculus. *Theoretical Computer Science*, 311(1–3):121–163, Jan. 2004.
10. N. Kobayashi. Type Systems for Concurrent Processes: From Deadlock-Freedom to Livelock-Freedom, Time-Boundedness. In J. van Leeuwen, O. Watanabe, M. Hagiya, P. D. Mosses, and T. Ito, editors, *Theoretical Computer Science: Exploring New Frontiers of Theoretical Informatics, Proceedings of the International IFIP Conference TCS 2000 (Sendai, Japan)*, volume 1872 of *Lecture Notes in Computer Science*, pages 365–389. IFIP, Springer-Verlag, Aug. 2000.
11. G. Meredith and S. Bjorg. Service-Oriented Computing: Contracts and Types. *Communications of the ACM*, 46(10):41 – 47, October 2003.
12. M. Mohri. Semiring frameworks and algorithms for shortest-distance problems. *Journal of Automata Languages and Combinatorics*, 7(3):321–350, 2002.
13. U. Nestmann and R. Fuzzati. Unreliable failure detectors with operational semantics. In *Proc.ASIAN 2003*, *Lecture Notes in Computer Science*. Springer-Verlag, 2003.
14. M. Papazoglou and D. Georgakopoulos. Special issue on service oriented computing. *Communications of the ACM*, 46(10), 2003.
15. C. Priami. Stochastic  $\pi$ -calculus. *The Computer Journal*, 38(6):578–589, 1995.
16. G. Theodorakopoulos and J. Baras. Trust Evaluation in AdHoc Networks. In *WiSe '04: Proceedings of the 2004 ACM workshop on Wireless security*, pages 1–10. ACM Press, 2004.
17. E. Tuosto. *Non-Functional Aspects of Wide Area Network Programming*. PhD thesis, Dipartimento di Informatica, Università di Pisa, May 2003. TD-8/03.