

Towards an Ontological Foundation for Services Science

Roberta Ferrario and Nicola Guarino
ISTC-CNR, Laboratory for Applied Ontology
Via alla Cascata 56C, 38100 Trento
{ferrario, guarino}@loa-cnr.it

Abstract. Most of the efforts conducted on services nowadays are focusing on aspects related to data and control flow, often disregarding the main goal of the future *Internet of services*, namely to allow the smooth interaction of people and computers with services in the actual world. Our main claim is that it is crucial, to achieve such goal, to build a *global* service framework able to account for complex processes involving people and computers, which however have always *people* at their ends. That's why in this paper we mostly emphasize the role of *social* and *business-oriented* services, whose consideration is needed to evaluate the global quality of e-services in relation to their ultimate social benefits, taking the overall impact on the organizational structure into account. Along these lines, the contribution of this proposal is a first concrete step towards a unified, rigorous and principled ontology centred on the notion of *service availability*, which results in useful distinctions between *service*, *service content*, *service delivery* and *service process*. Services are modelled by means of a layered set of interrelated events, with their own participants as well as temporal and spatial locations.

Keywords. Ontology, Services Science, Social Service, Service Content, Service Process, Service Description, e-Government.

1 Introduction and state of the art

Despite the ubiquity of the notion of service and the recent proposals for a unified *Services Science* [1], multiple inconsistencies between definitions of service from different disciplines (and even within the same discipline) still exist ([2], [3]). In particular, despite the general goal of the future *Internet of services* is –arguably– to allow people and computers to smoothly interact with services in the real life, both traditional Web services approaches, as well as the more recent Semantic Web Services (SWS) proposals, seem to focus mainly on the aspects related to *data and control flow*, considering services as *black boxes* whose main characteristic is to interoperate in a well-specified way (see, for instance, [4], [5] [6], [7]). This is certainly very useful, but, according to a recent paper by Petrie and Bussler [8], apparently it seems to work well only within *service parks*, where run-time interoperability is technically feasible because services are very constrained. As the authors put it, “some interoperability among service parks might emerge, but it could take a long time”.

Overall, the current state of the art is well described in a recent note by Katia Sycara [9], who observed that, on one hand, “current Web services proposals don't enable the semantic representation of business relations, contracts, or business rules in a machine-understandable way”, while, on the other hand, “current business-process languages [...] are at a low abstraction level and don't provide formal business semantics”. In conclusion, “a need exists to model informal business requirements in ways that make it feasible to translate them into precise business-service specifications, including operational interfaces and rules for procedures, timing, integrity, and quality. Such modelling must be driven from the top down, directly from business requirements [...]. The modelling would provide a functionality that's entirely understandable from a business perspective; it would depend on business context, goals, and operational standards, but shouldn't depend on the technology used to implement them. The models would provide business value directly relating to business purposes and could be understood and used without knowledge of underlying IT artefacts”.

This is exactly the perspective we are adopting in this paper, which certainly calls for a broad, interdisciplinary effort such as that envisioned by services science [1]. For sure, a proper, general ontological foundation for the notion of service is a fundamental requirement for such endeavour. This is the long term goal of our work.

The present paper aims at establishing a first step in this long-term process, presenting the foundations of a new ontology of services that intends to establish a common, unifying framework for representing services according to different views, based on a vision that considers services as complex systems of commitments and activities, involving real people, organizations, and actual circumstances. In other words, we believe it is crucial, for the future internet of services, to take into account the whole *service system* [2] that interacts with Web

services, through complex chains involving people and computers, which however have always *people* at their ends. That's why in this paper we mostly emphasize the role of *social* and *business-oriented* services, although the approach we describe should be ultimately general enough to account for any kind of service. In doing so, we adopt a *global view* of services, which, in a sense, goes against the strict separation between the external and the internal view advocated by SWS standards such as WSMO [10, 11]. The main reason is that the *terminology* needed to properly expose, retrieve and interact with a service, and especially that needed to understand and negotiate Service Level Agreements (SLAs), unavoidably requires a common understanding of the general structure of the service process, and the related activities involving the value exchange process between the producer¹ and the customer. Of course, in some cases services producers may have very good reasons for not exposing their internal workflow details, but the point is that, in general, SLAs *may* refer to some details concerning the *way* the service is implemented and the nature of such details is not specified in advance. So, since the boundaries between the external and the internal service description cannot be defined in advance for all kinds of service, so that a global approach seems to be the only viable alternative for a foundational ontology of services.

A further reason for a global approach lies in the fact that, in many cases, it is important to account for the way a service-based architecture impacts the organizational architecture (indeed, service process re-engineering typically impacts organizational re-engineering). In this case, it is crucial to model in the proper way the links between services, people and organizations. This is especially relevant for application areas such as eGovernment, which is currently our major concern², where it is crucial to evaluate the global quality of e-services in relation to their ultimate social benefits, taking the overall impact on the organizational structure into account.

Modelling services according to this global view is not an easy task, however, mainly because the notion of service is a subtle and ambiguous one, so that many researchers simply have given up adopting a clear definition, relying on a variety of intuitive notions mainly coming from practical considerations, which lack unfortunately a coherent framework. In other words, we are still facing the general question: *what is a service?* Is there a single notion behind this term? And if there are multiple aspects, how are they related? How is the internal view of services as *business processes* related to the external view of Web services as exposed (aggregates of) functionalities?

In this paper we shall introduce a novel, general approach to service modelling founded on the basic principles of ontological analysis³, centred on the notion of *service commitment* as a temporal state resulting from an agent's promise to perform certain actions in the interest of potential beneficiaries in correspondence of certain triggering events. In this view, services are modelled by means of a layered set of interrelated temporal activities, each one with its own participants and spatiotemporal location. This approach shares many similarities, in its main inspirations, with Alter's work on service systems [2], as well as O'Sullivan's work on non-functional requirements for services [12], and Baida, Gordijn and Akkerman's work on the service value chain [13]. While these approaches are however relatively informal, we believe that this proposal is a first concrete step towards a unified, rigorous and principled ontology.

2 The proposed approach

2.1 The Basic Idea

If we start from the simple question "what is a service?", it is immediately very evident that there is a huge confusion, not only in the layman's common sense, but also in the way the term is used in the literature. Sometimes the term "service" is used to indicate an *action* (actually performed by somebody), or a generic *type of action* (including in this category data manipulation procedures such as those typically described as Web services) or perhaps the *capability* to perform some action; other times it refers to the *result* of such action, which is typically a *change* affecting an object or a person, or just the (subjective) *value*, or utility of such change; moreover, in certain settings (like Public Administrations) the term denotes an *organization* acting (or in charge of acting) in a certain way in the interest of somebody.

¹In literature it is often used the term "provider", where we have chosen to use "producer". The reason for adopting such term instead of the more common "provider" is given by the fact that the latter term is used ambiguously, sometimes to indicate what we have called "producer" and sometimes what we have named "trustee".

²See www.lego-lab.org

³Similar efforts are being carried out under the scope of the European project SUPER (<http://www.ip-super.org/>), which has, among its aims, that of adding an ontological foundation to IT artefacts for business process management, so that these may become more transparent to business people.

In our opinion, all these notions are somehow connected, and contribute to better specify the notion of service, but none of them can be properly identified with what we believe people are commonly referring to when asking for a service. More or less “official” definitions occurring in the ICT literature do not help much.

To start with, the definitions in the W3C glossaries, referred to Web services, present evident ambiguities:

“[...] a software system designed to support interoperable machine-to-machine interaction over a network”.
Web Services Glossary

but also

“[...] an abstract resource that represents a capability of performing tasks that form a coherent functionality from the point of view of providers entities and requesters entities. To be used, a service must be realized by a concrete provider agent.”
Web Services Glossary

It is easy to see that, even in the case of restricted technical domains such as that of Web services, services are sometimes seen as processes, other times they rather resemble agents. WSMO, on the other hand, introduces a radical distinction between Web services and services, stating that the former are “computational entities able (by invocation) to achieve a user’s goal”, while the latter are “the actual value provided by this invocation” [11]. From a more general perspective, the “Research manifesto for services science” [1] cites Ted Hill’s definition [14]:

“a service is a change in the condition of a person, or a good belonging to some economic entity, brought about as the result of the activity of some other economic entity, with the approval of the first person or economic entity”

On the same line, from the business-oriented perspective, we may mention the comprehensive list of “service concepts” involved in the value chain process discussed in [2], and the service ontology proposed in [13], which however don’t attempt a precise definition of service.

Finally, O’Sullivan and colleagues make a suggestion that fits very much our intuition (although the former author avoids to build on it in his recent PhD thesis [12]):

“a service instance is essentially a promise by one party (the provider) to perform a function on behalf of another party at some time and place, and through some channel”.

[15]

To see how these various definitions are related, let us start with some simple questions, focusing on very general public services such as fire-and-rescue services, snow removal services, children care services, etc. What do we *pay for*, when we fund such services with our taxes? What does it mean that, for instance, in a municipality there are such services, at a certain time? Is anybody extinguishing a fire or removing some snow *at that time*? No, certainly not necessarily. We can legitimately say that *here and now* both a fire-and-rescue service and a snow removal service are *present*, even though there are no lit fires, nor is it snowing. It suffices to say that there is someone (firemen, snow removal operators) who is *prepared* to perform precise actions in case something happens (fire, snow). So our core notion of service is based on the following statement:

A service is present at a time t and location \perp iff, at time t , an agent is explicitly committed to guarantee the execution of some type of action at location \perp , on the occurrence of a certain triggering event, in the interest of another agent and upon prior agreement, in a certain way.

So, in a sense, at the core of any service there is a *commitment* situation in which someone (the service *trustee*) guarantees the execution of some kind of *action(s)* (by means of a *service producer*, which may coincide with the trustee or be *delegated* by him) in the interest of somebody who agrees (the *service customer*), at a certain cost and in a certain way. From the ontological point of view, this situation is a static temporal entity, i.e. a *static event* in the sense of the DOLCE ontology [16]⁴, which involves the participation of a single agent, the *trustee*. This commitment state typically starts at the time of the commitment act, and its duration is determined by the commitment’s act itself⁵, which typically specifies some constraints concerning the way the commitment

⁴ Although the term “event” has often a dynamic connotation, we use such term in the more general sense of *entity which occurs in time* (also called *perdurant* in the DOLCE ontology). In this understanding, states and processes are considered as special event kinds.

⁵ We assume that the commitment act (the *speech act*) is instantaneous, and occurs at a time which does not necessarily coincide with the beginning of the availability state.

will be fulfilled.

As we shall argue in the rest of the paper, service commitment needs to be distinguished from *service content*, which concerns the kind of action(s) the trustee commits to, and *service process*, which is a set of business processes implementing the service commitment (see Figure 1). In turn, we distinguish service commitment from *service availability*, which involves a service process running at a certain time and location: this allows us to account for malfunctioning periods or working pauses, where the commitment still holds but the service is not available. Following [17], [18], [19] and [20], the commitment act can be seen as a *speech act* that most of the times is codified in a *document*, i.e. in an institutional object that can assume many different forms: a contract, an official declaration or deliberation, a service level agreement⁶, etc.

In institutional settings, the *trustee*, the agent who commits, is typically a Public Administration. On the other hand, the *service producer* may not necessarily coincide with the trustee, and can be either a PA or another kind of (private) organization, delegated by the trustee; in some exceptional cases even an individual agent. The same holds for the *service customer*, who can in turn be a PA, an organization, or an individual agent, the latter being much more common than in the previous case.

The last element present in the definition is the *triggering event*; two kinds of triggering events can be singled out. The first one, more trivial, is a simple request made directly by the customer (like a parent in need who requires children care); in this case the *service invocation* coincides with the triggering event. The second one is the occurrence of a particular event kind, like the lighting of a fire in a wood, or a difficult situation observed by a social assistant, that triggers the action⁷. Of course, since the occurrence of the triggering event is not known in advance, the action time is in general much shorter than the availability time, so a service may be available at a certain time even if none of its foreseen actions do actually occur.

It is worth stressing an original feature of our definition, namely the inclusion of the triggering event. Traditionally, approaches on services are goal oriented; take for instance this definition from [21]: “A service delivers a process to achieve a certain goal by using resources”. Note however that actually the goals may in some cases be just implicit, or even different if you compare the producer’s perspective with the customer’s perspective. In such cases, specifying the service also in terms of the triggering event and the action to be performed in correspondence of such event seems to be less ambiguous. The service’s goal doesn’t disappear in our approach, it is present in what has been called the service content, but the triggering event allows to justify the passage from service availability to service invocation. Moreover, note that a triggered action may not necessarily succeed. What the trustee guarantees may in some cases be only the action’s performance, not its result. This changes also the mechanisms for the evaluation of service quality, which must distinguish between actions/processes and resulting states.

It is interesting to compare our definition with the second W3C definition reported above. We can observe many similarities, the most obvious being the presence of producers and requesters, and the distinction between two levels: an *abstract* level, where functional capabilities find their place, and a *concrete* one, where the functionalities are *realized*. However, our notion of commitment is different from an abstract capability in two ways. First, an agent may be capable of doing something without being committed to do so (for instance, a Web service may be potentially operational but not activated). Second, our definition involves the producer agent (more exactly, the trustee) already in the notion of service, instead of confining it to the service realization only. This means that, in our approach, different agents will always guarantee different services (possibly with the same *content* – see below). This choice seems more intuitive to us: when asking “how many telephone services are there in this country?” the answer can be “Two, but they deliver the same content”. Moreover, W3C’s definition is terminologically ambiguous, since it adopts the same term for the abstract and the concrete level. For us, services are always concrete: so, for instance, a C procedure, even if endowed with a standard Web interfacing mechanism, is not a service, but rather a *service process description*. When a computer runs this procedure, then it produces a service that complies with such description.

In a sense, our definition binds together the abstract and the concrete levels, which are comprehended in an articulated unitary framework. It is important to underline that this binding is very different from collapsing the two levels: they are kept separated, but they are both taken into explicit account, while in most cases the concrete level, even when mentioned, is then neglected.

⁶ In the actual practice, the term “service level agreement” may be typically used to refer to the negotiation that the producer conducts with the user; here we are using the locution in a coarser sense, which includes also the agreements between trustee and producer and trustee and user and, possibly, between trustee and the community to which services are provided.

⁷ To be more precise, it is the *observation* of such event that triggers the action. It is worth noting that, for this reason, many services include among their supporting activities an explicit monitoring activity, which can be executed by the producer itself or delegated to another agent.

2.2 Services and goods

To better understand the nature of our proposal – that services are temporal entities (events) based on *commitments* – let us briefly discuss the difference between services and goods. According to the World Trade Organization, services are a sort of intangible goods, so that a service might be defined as anything you can buy, but “you can’t drop on your foot” [Shrybman 2001]. Yet, Ted Hill insists on the fact that services are not a special kind of goods, because goods and services belong to quite different ontological categories: goods are both *transactable* and *transferable*, while services are transactable, but not transferable. In Hill’s own words, “a surgical operation is not some kind of immaterial drug”: when you buy the drug you become an *owner* of it, in the sense that you can decide about its behaviour (i.e., assuming it in your body), while when you pay for the surgical operation you are not actually becoming the owner of it. In support to this argument, we argue that the ontological reason why services are not transferable is exactly because they are events: you cannot *own* an event, since if owning implies being in control of temporal behaviour, then, strictly speaking (at the token level), the temporal behaviour of an event is already determined, and changing it would result in a different event. So events are not transferable simply because they are not “ownable”. Since services are events, they are not transferable as well.

So, in conclusion, it seems legitimate to assume that goods are *objects* (endurants, in DOLCE’s terms), while services are *events* (perdurants). One may observe however that our economy is full of examples of transactions involving services, where the service seems to pass from hand to hand: certainly somebody may buy AMAZON, for instance: our point is that in this case the transaction involves the AMAZON *company*, not AMAZON’s *service*: there is a change of ownership concerning the service producer, but not the service itself, which remains the same.

2.3 The basic ontological structure of services

Let us continue our analysis with another question: what’s happening when a service is *produced*? As we have seen, a service may be concretely available even if it is not actually delivered, or maybe will be never delivered: we keep paying the firemen even if no fires occur. So, in our approach, a service has to be distinguished from its actual *delivery* to a particular customer. Indeed, typically the same service guarantees multiple deliveries. By the way, to avoid confusions, we propose an important terminological distinction: strictly speaking, it is not *the service* which is delivered, but its *content*, i.e., the actions intended to be performed in the interest of the customer. So a service is the concrete commitment (guaranteed by a *trustee*) to produce a certain content, consisting in actions of a certain kind executed in a certain way. Altogether, the various actions that ultimately lead to service production (performed by the service *producer* on behalf of the trustee) constitute the *service process*. We shall say that a service process *implements* a service. The concrete delivery to a particular customer presupposes however a *service acquisition* activity engaged by the latter, which typically negotiates a *service offer* resulting from *service bundling and presentation* activities on the producer’s side. Finally, to complete the picture, we have to take into account the activities related to the value exchange chain, which include the service exploitation from the customer’s (*customer’s exploitation*), the sacrifice needed to access to the service (*customer’s sacrifice*), as well as the corresponding activities from the producer’s side.

So, as illustrated in Figure 1, a service is conceived as a complex event, with five main proper parts: service commitment, service presentation, service acquisition, service process, and service value exchange. In the following, we shall discuss these notions in more detail, with the aim to establish the basis for an ontology of services able to account both for *service descriptions* from an external point of view (typical of Web services and Service Oriented Architectures) and for *service processes* from an internal, business modelling point of view.

First of all, let us remark that all the blocks described in Figure 1 are *events* (*perdurants*, in DOLCE’s terminology). This means that they can be characterized, roughly, by their *temporal location* and by their *participants*, linked to the event by means of what are usually called *thematic relations*: *agent*, *patient*, *theme*, *instrument*... Specifying a service (or a service kind) amounts to constraining these events by imposing suitable restrictions on their temporal locations and thematic relations. So, for instance, non-functional requirements such as those discussed in [12] are represented as attributes of specific service components, each involving a particular aspect (participant/thematic relation) of a particular service event. The resulting analysis, which we cannot discuss in detail for reasons of space, looks very similar to Alter’s *service responsibility tables* [2], where the rows correspond to service components (events), and the columns to specific aspects to be considered (thematic relations). We give therefore a clean ontological foundation to a business-oriented proposal. Moreover, specifying the agents involved in each event allows for a fine-grained account of the *organizational impact* of a certain service. Note that, although the relationships between these various events (for instance, whether or not they involve the same agents) may vary according to the nature of the service specified, there exists a systematic ordering relationship between them, so that a service has a *layered structure*. This ordering relationship is not so much a temporal precedence (indeed most of these events are temporally overlapping), but

rather an (existential) *ontological dependence* relationship: in order for an event at a certain layer to occur, some event at the higher level has to occur. Ultimately, all the events belonging to the service process presuppose some acquisition event, which in turn presupposes the service commitment.

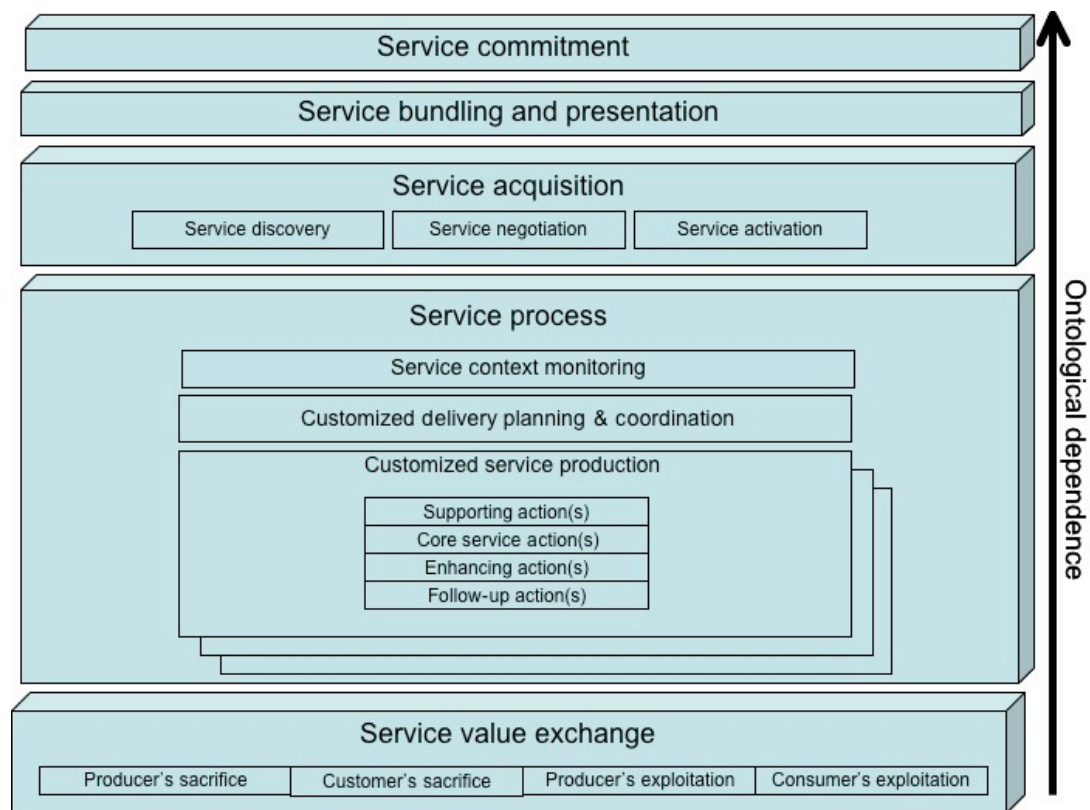


Figure 1: The layered structure of service activities

At this point it is important to notice the central role played by actions and events in this account for the description of services; this is in contrast with the major trend, which is to describe services in terms of pre-conditions and post-conditions, like in WSMO, where processes are represented as transitions between states [11]. Under a different perspective, the two approaches could also work in conjunction.

There are several reasons why in our opinion it is important to explicitly represent events; first of all, though for the front-office in most cases it is enough to know which is the starting state and which the desired one, for the back-office it is important that the whole process be transparent (to know who does what and when), especially when a failure is at stake. But malfunctioning is not the only reason: without events one does not have sufficient expressiveness to distinguish two different commitments with the same content but different deadlines. Even if one sees these deadlines as non functional properties, it is hard to use them, say, for expressing a SLA without a clear semantics.

Again, in [22] a list of service quality's determinants is given, in which at least a couple of these determinants are strongly space-time dependent: responsiveness, connected with timeliness of service and access, defined by (among others) three items: "waiting time to receive a service [...]; convenient hours of operation; convenient location of service facility."

By using only pre- and post-conditions other subtle but important differences are lost. Take for example these two scenarios: 1. an unemployed woman who becomes pregnant and 2. a pregnant woman who becomes unemployed. In our account, the two scenarios can be distinguished by the fact that in 1 the pregnancy is the triggering event, while in 2 it is the unemployment. This difference may result in the activation of different services (for instance, a financial aid in 1 and a help in searching a new job in 2, or a legal enquiry on the employer if there's a suspicion of unfairness). In a pre- post-conditions framework both scenarios have the same pre-conditions and thus should activate the same services. Notably, the literature in economics has since long recognized that comparing the outcomes of services is not enough in order to evaluate their quality. See for instance [22] (similar distinctions appear in [23], [24], and [25]):

Quality evaluations are not made solely on the outcome of a service; they also involve evaluations of the *process* of service delivery.

Let us now consider the various events constituting the service process internal structure. In Fig. 1, the containment relationship between the various blocks represents the parthood relationship. The core constituent of a service process is a set of basic activities (each called *customized service production*⁸), centered around the delivery of service content to a *single customer*. In addition to the *core service action(s)* depending on the service nature, a customized service production may include *enhancing actions* intended to increase the service value or differentiate it from those of competitors [14] as well as *supporting actions* needed to enable the core service consumption and *follow-up* actions intended to monitor the core action's results. In addition to customized delivery activities, the service process includes various back-office activities concerning *customized delivery planning and coordination*, plus an activity we have labelled as *service context monitoring* – which seems to be neglected by most current approaches – which involves the various actions necessary to detect the event that triggers service production, which can be an external situation or a customer's request: without an explicit modelling of such activity, there would be no way to account for delays or improper management of triggering events.

As a presupposition to service production, typically some *service acquisition* activities are required from the side of the customer⁹. These include *service discovery*, which is the event where the service trustee (or producer) and the service customer first meet together; *service negotiation*, which involves an agreement between the two parties, and *service invocation*, which refers to the event where the customer agrees to the service (not necessarily implying immediate production).

On the other hand, the service production results in a complex chain of transfers of value, which are represented in Figure 1 as the event Service value exchange. With a simplification, this is decomposed in *Producer's cost*, *Customer's cost*, *Producer's revenue* and *Customer's revenue*.

While in the case of the producer, most of the times both for cost and for revenue the value has to be intended in terms of money, in the customer's case things are more complicated. For instance, especially for services in the social domain, the customer's cost can be seen as an action whose results go somehow against some of the recipient's desires, but which the customer is still willing to perform, like *service sacrifice* as specified in [11]. Also the customer's revenue sometimes is not expressible in monetary terms, but only as some wellness state. Moreover, even though there's always an ultimate recipient of a service, we could also have indirect recipients, like the community that pays with its taxes for the service and benefits in terms of enhancement of its social conditions.

In conclusion, we can say that a service is characterized in a *prescriptive way (commitment level)*, while a service process in a *descriptive way (implementation level)*. The commitment level is where the "rules of the game" are established: what types of action compose the service, what types of agents are entitled to execute those actions, what types of agents may qualify as recipients, what types of events can become triggering situations. It is also the level where legal responsibility is at stake. In fact, from the point of view of the service offering, it is not important who in particular executes certain actions, but rather that a certain kind of action is executed in a certain way, by an agent who displays certain features and has some competences. The agent who is responsible that the required conditions are met is usually an organization, such as for instance a public administration. Such responsibility is usually distributed and assigned according to some structural constraints, i.e. by devising a structure of roles and sub-organizations. The ontological analysis of organizations is thus a topic tightly connected to the ontological analysis of services.

When we come to the actual service process, the various *kinds* mentioned at the commitment level need to be instantiated in concrete *tokens*. Individual agents are those who realize the core actions of service production, whose recipients are, ultimately, concrete agents (citizens); also the triggering situation is the occurrence of a precise (instance of) event. The service production level is thus the *descriptive* level, the one the data that are recorded and transferred belong to.

Finally, let us mention the issue of spatio-temporal location of services. In very general terms, one could say that in most cases when a somebody makes available a service, this availability spans over a spatio-temporal region which includes the spatio-temporal region in which the core service actions will (possibly) be executed; in rare cases, the two can coincide. For some special services, the analysis can be further complicated by the fact that the service may be delivered in a place and at a time and received in another place at another time. We won't enter into these details at present, but the issue needs to be investigated.

⁸ In the context of public services, a single event of customized service production is often called an *intervention*.

⁹ Even in the case of free, public services, it is difficult to imagine a case where the **customer** is not required to actually discover the service, or make a minimal sacrifice to exploit it.

3 A revised version of Alter's responsibility tables

In a recent article, Steven Alter [2] has presented a conceptual instrument that he calls "service responsibility tables" (SRT); these are aimed at facilitating a better understanding of services primarily based on the responsibilities assigned to each role; moreover, Alter suggests to add as many columns as necessary in order to address different aspects of analysis. For instance, in Table 1 below, we have in the rows the various activities constituting a service, in the first two columns the responsibilities that the main actors have been assigned to in each constituting activities (provider in the first and customer in the second), while in the third column Alter records the problematic situations that may emerge in the execution of some activities composing a service. For what concerns the inclusion of additional columns, in another table Alter collects other possible topics, among which: participant or interpersonal issues, goals and requirements, preconditions, triggers, post-conditions, benefits provided to the customer, duration (cycle time), provider cost, customer cost, etc.

In practice, Alter isolates two orthogonal components of services: the constituting actions and, for each of these actions, the responsibilities of the involved stakeholders; he thus describes *how* such stakeholders participate to the various events. These *modes* of participation individuate the *role* the participants play in the various events constituting the service.

Even though the topics suggested by Alter are heterogeneous and sometimes confusing, we are interested in his idea of representing the events composing a service and the role participants have in these events.

In order to represent all this, we take inspiration from a notion introduced in linguistics to account for the internal structure of events: so-called *thematic relations*, expressing the nature of the relationship between an event and its participants. Adding thematic relations to those linking an event to its own qualities (such as temporal and spatial location) we have a full set of attributes at our disposal, among which the following ones appear to be as especially relevant for our purposes:

- Agent (the active role, the one who acts in the event)
- Theme/Patient (the one who undergoes the event; the patient changes its state, the theme does not)
- Goal (what the event is directed towards – typically a desired state of affairs)
- Recipient/Beneficiary (the one who receives the effects of the event)
- Instrument (something that is used in the performance of the event)
- Location (where the event takes place)
- Time/duration (when the event takes place, or how long it lasts)

As a result, in the table we have the main composing events in the rows (service commitment, service acquisition, service process, service value exchange...) and the thematic relations in the columns.

In order to give an idea of the approach, we take an example and we represent it using the tables. The example is directly taken from [14] and it is about a guy who goes to the mechanic's garage to have his car repaired. The aim of Table 2 is that of representing in an explicit way the fundamental constraints that need to be specified in an actual service description. This can bring many advantages both in the comprehension of the service's features and in the many different evaluations of service quality that can be made under various viewpoints.

The table describes the events in which a generic car repair service is articulated. The values we put in the various cells allow us to express the relevant constraints that distinguish this service from others.

We start with service commitment. During the commitment event, that chronologically comes first and is the one that all the other events depend on, the mechanic commits with a Public Administration (for instance the Chamber of Commerce) with a subscription act and his commitment consists in guaranteeing that he will execute a certain type of job (illustrated in the job description, on which he commits) according to the local rules. This commitment is valid in the whole Province (for instance) and starting from that very moment on.

After the commitment, we have the service acquisition, which in turn is composed by three different events: discovery, negotiation and activation. During discovery the customer looks for a mechanic (that is then the theme of his search) with the goal of having his car repaired. Note that not all the cells in this line are filled, meaning that, for instance, the instrument used for the discovery activity is not specified. Should we describe a service based (exclusively) on a certain mediator for the discovery process, the name of such mediator would be specified in the "Instrument" cell.

After the service is discovered, the negotiation between customer and mechanic starts; the goal is (probably) an agreement and the negotiation is on the customization of the service (in other words, how the service type in the job description is tailored to the customer's needs). At that point the mechanic activates the service, i.e., the related scheduling and organization activities. The last two events usually take place in the garage and the whole service acquisition event is performed after the commitment has been taken and before the occurrence of the actual repair.

Table 1: Alter's Service Responsibility Tables

Provider Activity or Responsibility	Customer Activity or Responsibility	Problems or Issues
Loan officer identifies businesses that might need a commercial loan.		Loan officers are not finding enough leads
Loan officer contacts potential loan applicant.	Potential loan applicant agrees to discuss the possibility of receiving a loan.	
Loan officer discusses loan applicant's financing needs and possible terms of the proposed loan.	Potential loan applicant discusses financial needs.	Loan officer is not able to be specific about loan terms, which are determined during the approval step, which occurs later.
Loan officer helps loan applicant compile a loan application.	Loan applicant compiles loan application.	Loan applicant and loan officer sometimes exaggerate the applicant's financial strength and prospects.
Loan officer and senior credit officer meet to verify that the loan application has no glaring flaws.		20% of loan applications have glaring flaws.
Credit analyst prepares a loan write-up summarizing the client's financial history, providing projections of sources of funds for loan payments, etc.		10% rate of significant errors, partly because credit analysts use an error prone combination of several spreadsheets and a word-processing program. Much rework due to experience of credit analysts.
Loan officer presents the loan write-up to a senior credit officer or loan committee.		Meetings not scheduled in a timely manner. Questions about exaggerated statements by some loan officers.
Senior credit officer or loan committee makes approval decision.		Excessive level of non-performing loans. Rationale for approval or refusal not recorded for future analysis.
Loan officer informs loan applicant of the decision.	Loan applicant accepts or declines an approved loan.	25% of refused applicants complain reason is unclear. 30% of applicants complain the process takes too long.
Loan administration clerk produces loan documents for an approved loan that the client accepts.		

Table 2: The Garage example

		Agent	Theme/Patient	Goal	Recipient/Beneficiary	Instrument	Location	Time/duration....
Service Commitment		Mechanic	Job description		PA (Chamber of Commerce)	Subscription act	Province/ Region?	Starting from a fixed date before the opening of the garage and until the duration of the license
Service Acquisition	Discovery	Customer	Mechanic	Car repaired				After opening and before actual repair
	Negotiation	Customer, Mechanic	Service customization	(Agreement)			Garage	
	Activation	Mechanic	Internal execution plan				Garage	
Service Process		Mechanic	Car	Car repaired	Customer	Mechanic's tools	Garage	Period in which the repair actually occurs
Service Value Exchange	Producer's sacrifice	Mechanic	Working hours	Being payed			Garage, bank...	A certain time (usually) after that the car has been repaired
	Customer's sacrifice	Customer	Money, car's unavailability, time needed to pick-up car...	Car repaired				
	Producer's exploitation	Mechanic	Money					
	Customer's exploitation	Customer	Car repaired/car availability					

The actual service process (as can be noted from Figure 1) is a very complex one, consisting of a lot of interconnected activities; here, for simplicity reasons, we choose to represent only the service's core actions.

In the service process event, the mechanic, with his tools and in his garage, performs some actions on the car aimed at having it repaired; this in the interest of the customer.

Finally, there is an articulated service value exchange event, which is constituted by a bunch of activities corresponding to what counts as a "sacrifice" or an "exploitation" from the producer's and consumer's points of view. This is a complex topic, that deserves a more thorough examination, because both the components of cost and those of revenue can be many and different evaluations can be conducted with different purposes.

Simplifying a lot, here we can say that the mechanic counts as a sacrifice his working hours with the goal of being paid, while the customer counts as a sacrifice the money he pays, the time to go to the garage, the time the car is unavailable and so on with the goal of having the car repaired; the mechanic earns money, while the customer's revenue consists in having his car available again.

There are some remarks that can be made; first of all, from the knowledge representation point of view, one thing that can be easily observed is that some values must be the same across multiple cells; for instance, the mechanic plays a role of agent both in the service commitment and process, while he plays the role of patient in service acquisition. This might be a problem as most languages ordinarily used to talk about services (like those based on description logics) are not expressive enough to account for co-reference between variables.

Another remark – a methodological one – is that these tables can be further refined, for example by decomposing the service process event in its internal layers.

Even though the example is quite elementary, it is already possible to see how much additional information the table can convey.

4 Concluding Remarks and Future Issues

In this paper we have proposed a novel framework aimed at constituting a common ontological foundation for services science. Let us briefly discuss what the main contributions of this approach are, and what future research directions we are considering.

1. *Revisitation of the difference between internal and external service views.* We have seen that the black box model of services based on external behaviour is too limited, and that a higher expressivity is necessary both to describe services in terms of their internal structure and to properly characterize SLAs and non-functional attributes.
2. *Improvement of the classic definition of services coming from economics.* We have seen that Hill's definition based on change is not general enough, since, for instance, it does not allow to consider services which do not necessarily produce a change, such as fire control.
3. *Focus on core actions instead of pre- and post conditions.* We have seen how pre- and post- conditions cannot by themselves capture important aspects of services, related to the way the service process is performed.
4. *Activity-based service representation.* We have seen how to describe a service in terms of a layered structure of related activities (events, in the most general sense of this term). The separation of the various activities described in Fig. 1 allows us to properly account for non-functional properties, which instead of generically belonging to the service as a whole are attributes that characterize specific activities. In this way, it is possible to determine what aspect of a given service implementation is responsible for a certain service property. In particular, spatio-temporal attributes can be easily taken into account.
5. *Comprehensive business-oriented approach.* We have introduced a clear distinction between service commitment, service process, and service content, taking also into account important issues affecting service quality and evaluation, such as bundling and presentation activities, acquisition activities, and actions related to the service value chain.
6. *Common framework to describe service according to different views,* in terms of more or less general constraints among the various service activities, providing an ontological foundation to the technique of *responsibility tables* introduced by Alter.

Given the preliminary nature of the present paper, many are the directions in which the analysis can be extended and enriched.

For sure, in order to be effective, this exploratory work needs to result in a formal model, that will constitute an ontology of services that, as a component of a modular social ontology, should be in the end connected with an ontology of organizations.

Acknowledgements

This work is carried out under the scope of the activities of the LEGO lab (www.lego-lab.it), a joint e-government initiative located in Trento, as well as the TOCA.IT project, founded by the Italian Ministry of Research (Tecnologie Orientate alla Conoscenza per Aggregazioni di Imprese su Internet) and the project “CSS – Cartella Socio-Sanitaria” founded by the Autonomous Province of Trento. The initial ideas at the basis of this project have emerged from a fruitful collaboration with “Servizio Politiche Sociali e Abitative” of the Autonomous Province of Trento concerning the revision of a catalog of social services, to be shared among different Public Administrations. The first author is funded by a PostDoc grant from the Autonomous Province of Trento.

References

1. Chesbrough, H. and J. Spohrer, *A Research Manifesto for Services Science*. Communications of the ACM, 2006. **49**(7): p. 35-40.
2. Alter, S., *Service system fundamentals: Work system, value chain, and life cycle*. IBM Systems Journal, 2008. **47**(1): p. 71-85.
3. Baida, Z., *Software-aided Service Bundling - Intelligent Methods & Tools for Graphical Service Modeling*. 2006, Vrije Universiteit Amsterdam.
4. Janssen, M. and R. Wagenaar. *From Legacy to Modularity: a Roadmap Towards Modular Architectures Using Web Services Technology*. in *Electronic Government*. 2003: Springer.
5. Papazoglou, M.P. and D. Georgakopoulos, *Service-Oriented Computing*. Communications of the ACM, 2003. **46**(10): p. 25-28.
6. Traverso, P. and M. Pistore. *Automated Composition of Semantic Web Services into Executable Processes*. in *International Semantic Web Conference (ISWC'04)*. 2004. Hiroshima, Japan.
7. Vetere, G. and M. Lenzerini, *Models for semantic interoperability in service-oriented architectures*. IBM Systems Journal, 2005. **44**(4): p. 887-903.
8. Petrie, C. and C. Bussler, *The Myth of Open Web Services: The Rise of the Service Parks*. IEEE Internet Computing, 2008. **12**(3): p. 94-96.
9. Sycara, K., *Unthethering Semantic Web Services*, in *Semantic Web Services, Part 2*, D. Martin and J. Domingue, Editors. 2007, IEEE Intelligent Systems. p. 11-13.
10. Fensel, D. and C. Bussler, *The Web Service Modeling Framework WSMF*. Electronic Commerce Research and Applications, 2002. **1**: p. 113-137.
11. Roman, D., et al., *Web Service Modeling Ontology*. Applied Ontology, 2005. **1**(1): p. 77-106.
12. O'Sullivan, J., *Towards a Precise Understanding of Service Properties*, in *Faculty of Information Technology*. 2006, Queensland University of Technology. p. 232.
13. Baida, Z., J. Gordijn, and H. Akkermans, *Service Ontology*. 2001, Free University Amsterdam.
14. Hill, T.P., *On Goods and Services*. Review of Income and Wealth, 1977. **23**(4): p. 315-338.
15. Dumas, M., et al. *Towards a semantic framework for service description*. in *Data Semantics 9: Semantic Issues in E-Commerce 239*. 2003. Hong Kong: Kluwer.
16. Masolo, C., et al., *The WonderWeb Library of Foundational Ontologies and the DOLCE ontology. WonderWeb Deliverable D18, Final Report (vr. 1.0. 31-12-2003)*. 2003.
17. Castelfranchi, C., *Grounding We-Intention in Individual Social Attitudes: On Social Commitment Again*, in *Realism in Action - Essays in the Philosophy of Social Sciences*, M. Sintonen and K. Miller, Editors. 2003: Dordrecht.
18. Jennings, N.R., *Commitment and conventions: The foundation of coordination in multi-agent systems*. The Knowledge Engineering Review, 1993. **8**(3): p. 223-250.
19. Verdicchio, M. and M. Colombetti. *A logical model of social commitment for agent communication*. in *AAMAS 2003*. 2003: Elsevier.
20. Singh, M.P., *An ontology for commitments in multiagent systems, toward an unification of normative concepts*. Artificial Intelligence and Law, 1999. **7**: p. 97-113.
21. Cauvet, C. and G. Guzelian. *Business Process Modeling: A Service-Oriented Approach*. in *HICSS '08, 41st Annual Hawaii International Conference on System Sciences*. 2008: IEEE Computer Society.
22. Parasuraman, A., V.A. Zeithaml, and L.L. Berry, *A Conceptual Model of Service Quality and Its Implications for Future Research*. Journal of Marketing, 1985. **49**(4): p. 41-50.

23. Sasser, W.E.J., R.P. Olsen, and D.d. Wyckoff, *Management of Service Operations: Text and Cases*. 1978, Boston, MA: Allyn & Bacon.
24. Gronroos, C., *A Service-Oriented Approach to Marketing of Services*. *European Journal of Marketing*, 1978. **12**(8): p. 588-601.
25. Lehtinen, U. and J.R. Lehtinen, *Service Quality: A Study of Quality Dimensions*. 1982, Service Management Institute: Helsinki.