

Towards Self-Configuration and Management of e-Service Provisioning in Dynamic Value Constellations

Jaap Gordijn
Department of Computer
Science
Free University Amsterdam
The Netherlands
gordijn@cs.vu.nl

Hans Weigand
Department of Information
Management
Tilburg University
The Netherlands
h.weigand@uvt.nl

Manfred Reichert,
Roel Wieringa
Department of Electrical
Engineering, Mathematics and
Computer Science
University of Twente
The Netherlands
{m.u.reichert,roelw}@cs.utwente.nl

ABSTRACT

Networked value constellations are collections of enterprises that jointly satisfy complex consumer needs. Increasingly, such needs are satisfied by e-services, i.e. commercial services that can be ordered and provisioned via the Internet. Current research in dynamic web-service composition has yielded run-time platforms to dynamically compose complex web services, but there is still a considerable gap between web services and commercial e-services. To compose e-services, an estimation of commercial profitability must be made, which is absent from web service composition. In this paper, we extend our earlier approach to e-service composition with a dynamic part, that ensures that a commercial e-service can be dynamically composed from other commercial e-services, and can be mapped on a web service composition process.composition of lower-level web services. We propose a skeleton-oriented approach, that first composes a network of enterprises, jointly satisfying need, based on commercial considerations. Second, given a set of such candidate value constellations, the business processes providing the services can be dynamically configured. We illustrate this skeleton-driven composition of networked value constellations by using a case study of clearing and repartitioning of Intellectual Property Rights (IPR).

General Terms

Value web, value constellation, coordination process, e-service, web service, dynamic composition

Keywords

Business networking, inter-organizational collaboration, conceptual modeling

1. INTRODUCTION

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC'08 March 16-20, 2008, Fortaleza, Ceará, Brazil

Copyright 2008 ACM 978-1-59593-753-7/08/0003 ...\$5.00.

Enterprises increasingly participate in networked value constellations [15]. By doing so, these enterprises can jointly satisfy a more complex consumer need than they could ever do on their own. Well known examples include Amazon.com, Cisco and Dell; these enterprises have strong partner relationships with other enterprises to provide their products. Many of these constellations actually provide *services*: activities and deeds of a mostly intangible nature [9]. A key problem in Service Sciences [14] is how to actually compose and manage such service-provisioning value constellations. In this paper, we concentrate on the following three main issues.

Firstly, we focus on *e-services* specifically. These are commercial services that can be ordered and provisioned online. As an aside, we note that this is different from web services, which are usually web API's that can be accessed over a network and are sometimes defined as software systems that support interoperability over a network, usually using XML technology. Web services are an implementation technique for e-services.

Secondly, we specifically deal with *dynamic* value constellations. In practice, enterprises enter and leave the constellation over time, services evolve, and parties may play different roles in different periods. Service provisioning constellations therefore have to deal in an adequate way with these dynamics, e.g. by reconfiguring themselves on-the-fly.

Thirdly, we consider *self-configuring* networked value constellations. In classical work on coordination [10], a distinction is made between hierarchical and multi-lateral organized constellations. In a hierarchical organization, top-level enterprises determine the selection of enterprises forming a constellation (e.g. Cisco, Dell). In contrast, in multi-lateral configurations, enterprises have no control over each other, and the configuration emerges in a self-organizing way (e.g. Skype). Both configurations can be empirically observed in practice, and occur also in combination. We want to focus on multi-lateral configurations. This means that our constellations must be self-configuring without central coordinator.

The contribution of this paper is a *skeleton-driven* approach to the dynamic (re)composition of networked value constellations for the provisioning e-services. This is inspired by the skeleton planning approach used by Omelayenko [11] for dynamic composition of semantic web services. Our approach differs from Omelayenko's because we apply it to *value-based e-service composition*. To do so,

we propose two different, but related, skeletons, that each can be instantiated, and together represent a generated networked value constellation for service provisioning. These skeletons are: (1) the *business value* skeleton (section 4), (2) the *business process* skeleton (section 5).

2. CASE STUDY: IPR CLEARANCE AND REPARTITIONING

We explore our skeleton-approach by using a real case study on Intellectual Property Rights (IPR). In The Netherlands, we have a long-standing relationship with IPR societies, who are operating in a large value web consisting of other societies, IPR owners, and IPR users.

IPR consists of many elements, and in this paper we focus on one right only, namely the *right to make content public*, henceforth called *publication right*. Commercial entities (radio & television stations, bars, discotheques, etc.) have to pay to IPR owners (artists, producers, text & song writers) a fee for using intellectual property (e.g. a music track). IPR societies act as a “man in the middle”; they collect this money from the IPR users (called *right clearance*), and they pay money to the IPR owners (called *repartitioning*). As there exist many IPR owners and -users, and there is one IPR society per type of [IPR owner, right] combination, multiple societies are involved in clearing and repartitioning the money for one specific track used. For instance, in The Netherlands, to clear and repartition neighboring rights, there is the SENA for artists and producers, and BUMA for song- and text writers. Even more societies will be involved in case of music tracks with IPR owners of different nationalities, since each country has its own IPR societies.

The clearance and repartitioning activities are commercial *e-services* provided by IPR societies, because they are done automatically. Payment for the provisioning of these *e-services* is also automated. Moreover, the constellation becomes increasingly *dynamic*. Currently, IPR handling is in many countries a monopolistic activity as IPR societies are appointed and controlled by a local government, but due to EU-deregulation, private entities will also be allowed to develop IPR activities. IPR users and owners may then select their preferred IPR-society for clearance and repartitioning. This will result in a situation that for *each* music track, a *different* constellation has to be created, since each different IPR owner of the track may select its own IPR society, as well as the IPR user may do. Due to the large number of music tracks played daily by different rights users, this should be automatically done *on-the-fly*. Moreover, IPR owners and users may change their preference over time. This results in dynamic constellations, in terms of societies and right-stakeholders leaving and entering a constellation.

In theory, building a IPR clearance and repartitioning network can be done by a centralized authority. However, such a solution is not acceptable to the existing IPR societies, as a centralized authority is considered as too powerful. Consequently, a *self-organizing & managing* solution is considered as more promising by the industry. Here, a decentralized mechanism forms and manages the IPR clearing and repartitioning constellation, thereby spreading power sufficiently.

3. SKELETONS AND INSTANTIATIONS

We assume that there is a given set of IPR users, societies, and owners. We call this set a *value web*. We do

allow value webs to grow and shrink, but we do not consider here the mechanism by which this happens. We also assume that the members of a value web have compatible run-time platforms that can dynamically search, compose, and execute web services. Several run-time platforms for dynamic service composition are now available, that allow dynamic discovery and matching of requested and offered web services, including negotiation of service level agreements (SLAs), monitoring of quality-of-service and dynamic reconfiguration in case of SLA violation [1, 2, 5, 6]. Different actors in the value web may have the same platform, or they may have different platforms that satisfy the relevant standards.

A value web exists to satisfy certain consumer needs, and we consider consumers of this need to be part of the value web too. Whenever a consumer need occurs, a particular subset of the value web will provide services to each other, and by doing so create a complex service that satisfies the need of the consumer. We call a subset of the value web that satisfies a particular consumer need a *value constellation*. The purpose of a value constellation is to satisfy a particular consumer need occurring at a particule time. Value constellations are formed dynamically, so that different needs of the same consumer may lead to different value constellations, the same need occurring in different consumers may lead to different value constellations, and even different occurrences of the same need in the same consumer may lead to different constellations at different times. For example, IPR societies have a need to audit IPR users, but they may use different auditors for this at different times. The determination which value constellation must be formed for a given need must be done *on-the-fly* in an automated manner. Our goal is to develop techniques to do this in a value-driven way, using currently available for dynamic service composition allow dynamic binding of services, service level negotiation, and quality of service (QoS) monitoring.

Our solution consists of the following elements.

- Value perspective: In our approach each value constellation is described by a *value model*, that shows which actors are participating in this constellation, and who is providing services to whom and against which reciprocal service. We use the e^3 value notation [7] to describe value models. Furthermore, for each kind of consumer need we will define a *value model skeleton*, or value skeleton for short, that will be instantiated to a value model for each consumer need occurrence. We will extend the e^3 value notation to be able to represent these value skeletons.
- Process perspective: For each value skeleton we define a *coordination process skeleton*, or coordination skeleton for short, that represents the same kinds of actors as the value configuration skeleton, and tells us which kinds of activities need to be performed by these kinds of actors, and how these must be composed to deliver a *e-service* to the consumer. A process skeleton can be instantiated to a process model that describes for a particular consumer need what coordination process must be performed by which individual actors and how this delivers the required *e-service* to a particular consumer.

Instantiation of value and process skeletons is context sensitive. In particular, the consumer need occurrence

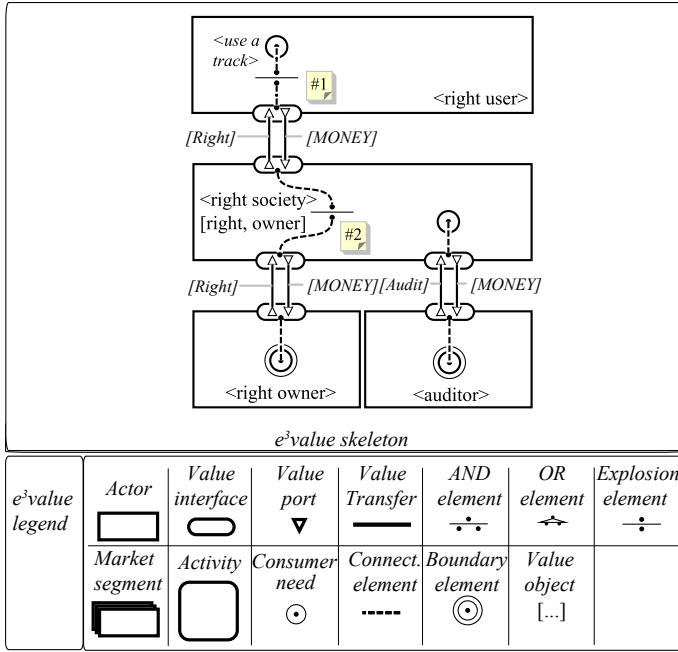


Figure 1: A skeleton e^3 value model for the IPR case. vervang “right” door “IPR”

should provide sufficient information to allow us to instantiate the value model and process model skeletons, respectively.

In the next two subsections we elaborate these two perspectives.

4. BUSINESS VALUE PERSPECTIVE

4.1 Value Skeleton

The value skeleton shows, given a consumer need, how that need is satisfied by providing valuable e-services by multiple enterprises. The focus is on *what* is exchanged in terms of *economic value*, and not yet how this is accomplished (this is considered by the business process perspective, see Sec. 5).

To specify a business value skeleton, we use the e^3 value technique [7]. Elements in the template that need to be instantiated are indicated by angled brackets ‘<’ and ‘>’. This is an extension to the e^3 value notation.

For the case at hand, we have the following skeleton actors. First, the <IPR user> is a financially independent actor, making money by playing a music tracks. Examples include a radio station (usually earning money by advertisements), a bar (music creates a certain atmosphere, thereby attracting visitors), and a restaurant. An <IPR user> has a need to <use a track>. Which track specifically is required, is part of the instantiation process for the skeleton, and not of the skeleton itself. To satisfy this need, the <IPR user> usually has to obtain a *set* of rights (to use a track), and in return pay for this. In Fig. 4.1, we have used an e^3 value explosion element (annotated with ‘#1’) to denote that *one* need for a track results in a set of (different) rights. So, a series of rights can be transferred between the <IPR user> and the <IPR society>(s).




Figure 2: An instantiated e^3 value model for the IPR case.

Second, an <IPR society> is an independent organization, who for a specific IPR owner clears a specific right. Because an IPR society clears only a subset of all existing rights for a subset of all existing right owners, an instantiation of this template typically requires more than one society. We will demonstrate this in Sec. 4.2 further. Furthermore, for a particular track, an <IPR society> usually clears rights for more than one IPR owner. This is for example the case if an orchestra played a track and some musicians in the orchestra have the same nationality. Consequently, another explosion element (annotated ‘#2’) has been added to model that an issued usage right for a track to an <IPR user> requires that multiple rights (from multiple <IPR owners>) need to be obtained.

Third, an <IPR owner> is an independent entity having one or more rights on the used track. Examples include artists, text and song writers, and record companies.

The value skeleton also shows an <auditor> that satisfies a need of an <IPR society>. This is a different need and it will be instantiated at different times, depending on the goals of the IPR society, say once a month or once a year. But in order to satisfy this need, an <IPR user> has to keep logs of the music tracks played. These logs must be reported to the <IPR society>, so that the clearing and repartitioning process can be carried out. As these logs are the source for track clearing and therefore result in financial transactions, it is attractive for the <IPR user> to commit a fraud (e.g. to log not all the used tracks, but just a subset). To accommodate this risk, the <IPR society> hires an auditing service from an <auditor>, who occasionally checks the logs of the <IPR user> for correctness. The operational details of this must be specified in the process skeleton for auditing.

4.2 Value Model

To obtain a working value constellation that clears and repartitions a particular music track at a particular time, we have to instantiate the skeleton as introduced in Sec. 4.1. The following description mentions some elements of an *instantiation process* that takes a consumer need and a value skeleton as input and produces the instantiated value model

of figure 4.2 as output. The instantiation process will itself be a coordination process implemented by means of web services. The process will be decentralized: It can be started by any partner in the value web and it can involve any partner in the value web. We describe in the following the instantiation process for our example value skeleton.

This instantiation process is triggered by the IPR user. In Fig. 4.2, the parameter $\langle \text{IPR user} \rangle$ has been bound to ‘TopPop radio’, a radio station, who wants to clear the track ‘BiCycle’ by ‘The Cyclists’.

For a given track (such as ‘Bicycle’ by ‘The Cyclists’), the instantiation process must obtain the list of IPR societies required to clear this track. We assumed in Sec. 3 that participating actors have registered themselves in the value web beforehand, and, at that time have provided sufficient information to guide the instantiation process. For this specific case, the registration process includes - per society - the address of a web-service that can be used as a directory service to look up whether the society is capable of clearing (a subset of) the rights for a specific track. The directory service itself can be distributed and replicated. A distributed and replicated solution is required to control power of individual enterprises. If the directory would be fully centralized, a single enterprise (owning the database) would be able to control the instantiation process of a value web into a value constellation. IPR clearance societies have made it clear to us that this would not be an acceptable solution for them.

The international ISRC code may be used to identify music tracks. The web-services of the various societies are used to instantiate the skeleton value model. As a result, in our example, the ‘SENA’ and the ‘BUMA’ are found as the two IPR societies that are required to clear the music track ‘Bicycle’ by ‘TheCyclists’.

Note that *two* IPR societies are selected, whereas the skeleton only shows *one* society. This is the result of the explosion element ($\#1$) as shown in the skeleton model; this element says that generally, multiple rights (from multiple IPR owners) have to be cleared. It corresponds to the directory lookup in the instantiation process that we described above. In Fig. 4.2, the explosion element in Fig. 4.1 is ‘expanded’ by using an e^3 value interface (denoted as a bounded rectangle with arrowheads inside).

A value interface models that an actor wants to exchange something of value (here: money) for something else that carries economic value (here: a series of rights on a music track). As such, the value interface models the notion of ‘economic reciprocity’: stating that in order to obtain something of value, something else of value should be offered in return. For the example, the value interface of ‘TopPop radio’ bundles four value object transfers: A right with ‘SENA’ (a IPR society) for money, and a right with ‘BUMA’ (another society), also for money.

The IPR societies themselves have to look-up into their own local directories the IPR owners for the track to be cleared. Again, this is part of the value model instantiation process and it precedes the execution of the coordination process discussed later. In the example, ‘SENA’ finds ‘John Doe’, ‘Roger Cha’ and ‘CDgram’, and ‘BUMA’ finds ‘Chris Cross’ and ‘M. Oldfield’ and they return these answers to the instantiation process, which then includes these in the value model being generated. The directories of the IPR societies can be centralized, meaning that they are operated by ‘SENA’ and ‘BUMA’. This reflects the hierarchical nature

of these societies; they operate on behalf of right-owners, who registered on beforehand at ‘SENA’ and ‘BUMA’.

The value perspective is useful for the following reasons. First, the instantiation process of a value skeleton into a value model, describing a particular value constellation satisfying a need, selects enterprises and the services they offer based on *economical* motivations. From a consumer perspective, the instantiation process should therefore select a bundle of services that reasonably satisfies the stated consumer need. This means that services and needs should be described such that an informed value-based service selection can be made. Current approaches, e.g. on web-service composition, do not consider services as commercially valuable entities, and therefore can not select services on an economic-value criterion. Rather, such web-services are selected on grounds of interoperability concerns, or technical quality concerns. From a supplier perspective, the instantiation process should select those suppliers that reasonably can make a profit by participating in the constellation. Additionally, the instantiation process must balance the interests of both the consumer and suppliers, such that the provided bundle of services can be provisioned in a sustainable way.

Second, by taking a business value perspective, reasoning about formation of a constellation can be done more efficiently. If such reasoning would start at the coordination process level, many possibilities would have to be considered that from a commercial perspective would not be of interest (e.g. because one participant can not make a profit by participating in the constellation at hand). Such alternatives are discarded beforehand, by pre-selecting only commercially viable constellations.

5. DISCUSSION AND FURTHER RESEARCH

5.1 Coordination process instantiation

A value model is not a process model because a value model states the objects, services, rights, experiences and other things of economic value are exchanged among economic actors, but does not tell us what they should do to realize these exchanges [8]. Fig. 3 shows a process skeleton that is corresponds to the value skeleton by means of an activity diagram. It says that when a track is scheduled, certain payments have to be performed; the explosion element in the value skeleton has been translated into a multi-instance activity in the process skeleton.

The instantiation process to create the process instance takes the generated instance (value model) from the value skeleton as a starting point. First, the value model gives the enterprises involved in the constellation. who, from an economic value perspective, have agreed to participate in the constellation at hand. Second, the value model lists the selected commercial services, which need to be provisioned by the coordination process.

The instantiation process that creates the process instance from the process skeleton is analogous to the one for value models. Figure 4 shows an instantiated process that corresponds to the value model of figure 3.

5.2 Profitability estimations

This small example illustrates several issues to be resolved to make this approach work. Our e^3 value models come with techniques to make estimations of the number of value trans-

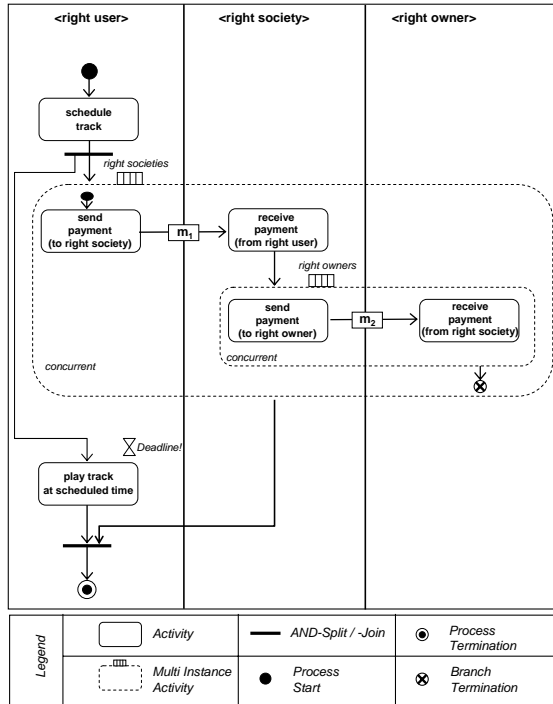


Figure 3: Process skeleton for the IPR case.

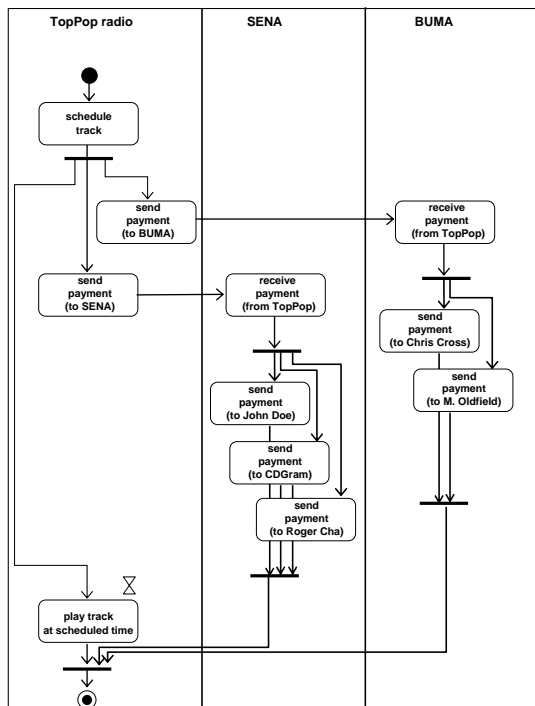


Figure 4: Instantiated process model for IPR case.

actions over a time period, and tools to make net present value computations of the profitability for a business based on these estimations. This allows a business to decide whether or not to join a value web on business-economic grounds. However, these techniques are geared to the static case. In order to decide whether or not to join a value web in which value constellations are formed dynamically, a business needs to estimate how often it will be part of a value constellation within that web over a period of time. This requires simulation techniques that are currently not yet part of e^3 value .

5.3 Process skeleton design

We currently have methods to design a process model from a value model. Weigand et al. [16] use a language-action approach to suggest patterns for inter-business communication that support the economic value exchanges represented in a value model (rights and money, in our case). This is typically done by means of a request/accept communicative transaction. Pijpers and Gordijn [12] propose a step-wise transformation approach that first considers ownership rights transfers, then possession right transfers and finally the time ordering of these transfers. These approaches work for the design of coordination processes from value models, but we here need to design a process skeleton from a value skeleton. Guidelines for this design process are closely intertwined with the directives for instantiating a process model (or value model) from a process skeleton (value skeleton) and we need to look into this in more detail.

5.4 Consistency of views

Second, we do have criteria for checking the consistency of a coordination process with a value model in monetary terms [3, 4]. Basically, these criteria say that the values estimated to be exchanged in a time period by the value model, are indeed exchanged in the coordination process. But these criteria work at the value model level and we need to lift them to the skeleton level.

5.5 Quality and risk

Third, different ways of realizing a value model in a process model differ in efficiency (do we perform some actions in a batched way?) and risk (do we perform audits once a week or once a year?). These issues now come up due to the possibility of automatic self-reconfiguration. The general form of the problem is that there are objectives to realize and choice of means to realize them. Since we want to work in a decentralized manner, the best we can get is local optimization and mutual adjustment on the basis of negotiations. Risk estimations and trust assumptions play a central role here: If we trust that violations by IPR users hardly occur, then we can audit once a year; if we do not trust IPR users, we may want to audit once a day but then the costs for IPR societies will be correspondingly larger. Current work on quality of service negotiations in dynamic web service composition provides techniques to compare and negotiate QoS levels of different services [17, 13], but it does not take the risk assumptions of the service delivery process into account. This needs to be dealt with to make dynamic service composition an economic reality.

Acknowledgements

The authors thank Bert Hazelaar of SENA for offering the IPR case study and the fruitful discussions about the IPR case study.

6. REFERENCES

- [1] L. Baresi, E. D. Nitto, C. Ghezzi, and S. Guinea. A framework for the deployment of adaptable web service compositions. *Service Oriented Computing and Applications*, 1(1):75–91, April 2007.
- [2] B. Benatallah, M. Dumas, and Q. Sheng. Facilitating the rapid development and scalable orchestration of composite web services. *Distributed and Parallel Databases*, 17:5–37, 2005.
- [3] L. Bodenstaff, P. Ceravolo, C. Fugazza, and A. Wombacher. Toward semantic-aware representation of digital business processes. In *accepted for the Second Knowledge Management in Organization Conference (KMO)*, 2007.
- [4] L. Bodenstaff, A. Wombacher, M. U. Reichert, and R. J. Wieringa. Monitoring collaboration from a value perspective. In E. Chang and F. K. Hussain, editors, *2007 Inaugural IEEE International Conference on Digital Ecosystems and Technologies, Cairns, Australia*, volume 1 of *IEEE-DEST 2007*, pages 134–140, Los Alamitos, February 2007. IEEE Computer Society Press.
- [5] M. Colombo, E. D. Nitto, and M. mauri. SCENE: a service composition execution environment supporting dynamic changes disciplined through rules. In A. Dan and W. Lamersdorf, editors, *International Conference of Service-Oriented Computing (ICSOC 2006)*, pages 191–202. Springer, 2006. LNCS 4294.
- [6] A. Dan, D. Davis, R. Kearney, A. keller, R. King, D. Kuebler, H. Ludwig, M. Polan, M. Spreitzer, and A. Youssef. Web services on demand: WSLA-driven automated management. *IBM Systems Journal*, 43(1):136–158, 2004.
- [7] J. Gordijn and H. Akkermans. Value based requirements engineering: Exploring innovative e-commerce idea. *Requirements Engineering Journal*, 8(2):114–134, 2003.
- [8] J. Gordijn, H. Akkermans, and H. Van Vliet. Business modelling is not process modelling. In *Conceptual Modeling for E-Business and the Web, ECOMO 2000*, volume 1921 of *LNCS*, pages 40–51. Springer, 2000.
- [9] C. Grönroos. *Service Management and Marketing: A Customer Relationship Management Approach, 2nd edition*. John Wiley & Sons, Chichester, UK, 2000.
- [10] T. W. Malone, J. Yates, and R. I. Benjamin. Electronic markets and electronic hierarchies. *Communications of the ACM*, 30(6):484–497, 1987.
- [11] B. Omelayenko. *Web Service Configuration on the Semantic Web: Exploring How Semantics Meets Pragmatics*. PhD thesis, Vrije Universiteit, Amsterdam, NL, 2005.
- [12] V. Pijpers and J. Gordijn. Bridging business value models and business process models in aviation value webs via possession rights. In *Proceedings of the 20th Annual Hawaii International Conference on System Sciences*, page cdrom. Computer Society Press, 2007.
- [13] S. Shen, G. Yu, T. Nie, R. Li, and X. Yang. Modeling QoS for semantic equivalent web services. In *Advances in Web-Age Information Management (WAIM04)*, pages 478–488. Springer, 2004. LNCS 3129.
- [14] J. Spohrer and D. R. (eds). Special issue on services science. *Communications of the ACM*, 49(7):30–87, 2006.
- [15] D. Tapscott, D. Ticoll, and A. Lowy. *Digital Capital - Harnessing the Power of Business Webs*. Nicholas Brealy Publishing, London, UK, 2000.
- [16] H. Weigand, P. Johannesson, B. Andersson, M. Bergholtz, A. Edisuriya, and T. Ilayperuma. Value object analysis and the transformation from value model to process model. In G. Doumeingts, J. Mueller, G. Morel, and B. Vallespir, editors, *Enterprise Interoperability: new challenges and approaches*, pages 55–65. Springer-Verlag, London, 2007.
- [17] L. Z. Yutu Liu, Anne Ngu. QoS computation and policing in dynamic web service selection. In *Proc WWW2004, ACM, pp.66-73*, 2004.