University of Augsburg
Prof. Dr. Hans Ulrich Buhl
Research Center
Finance & Information Management
Department of Information Systems
Engineering & Financial Management

Universität
Augsburg
University

Discussion Paper WI-180

Service-Oriented Computing for Risk/Return Management

by

Wolfgang Hackenbroch, Matthias Henneberger

September 2006

in: Obwerweis, A., Weinhardt, C., Gimpel, H., Koschmider, A., Pankratius, V., Schnizler, B., ed., Wirtschaftsinformatik 2007 – eOrganisation: Service-, Prozess-, Market-Engineering, Karlsruhe, 2007, Universitätsverlag, Karlsruhe, March 2007, p.459-512

Universität Augsburg, D-86135 Augsburg
Visitors: Universitätsstr. 12, 86159 Augsburg
Phone: +49 821 598-4801 (Fax: -4899)
www.fim-online.eu

Service-Oriented Computing for Risk/Return Management

Wolfgang Hackenbroch, Matthias Henneberger

Competence Center IT & Financial Services, University of Augsburg
86135 Augsburg, Germany, www.wi-if.de
(wolfgang.hackenbroch,matthias.henneberger)@wiwi.uni-augsburg.de

Abstract. Financial applications in the field of risk/return management demand for powerful and at the same time flexible information technology resources. These requirements seem to be highly suitable for service-oriented computing concepts. In this paper we identify promising risk/return management services and analyze the specific value proposition of service-oriented computing in this context. Taking the calculation of covariance matrices as an example we propose a model to quantify the economic value of such a service and thereby make a contribution towards understanding the business value of service-oriented computing. Moreover, our quantification model can be applied for the pricing of services and for the allocation of resources to services in a market-based environment.

1 Introduction

Risk/return management (RRM) is an important business function, especially for enterprises in the financial services industry because in an increasingly volatile and competitive market environment financial positions are affected by a high number of interacting economic factors that need to be permanently evaluated. To gain competitive advantage banks and other financial services firms apply sophisticated and complex models for the quantification of risk and the optimization of risk and return of portfolios containing all kinds of investment objects. This task is time-critical and therefore demands for powerful computational resources. Service-oriented computing based on grid technologies promises to overcome current restrictions by offering high-end computational capabilities at moderate cost and supporting new ways of intra- and inter-enterprise collaboration.

Although the financial services sector is often mentioned as a key industry for service-oriented computing based on grid technologies, see e.g. Friedman [2003], Middlemiss [2004] or Ricadela [2002], there is a lack of publications concerning applications in this area and important questions are still to be answered. In this paper we are striving to develop an understanding for the applicability of the general concept of service-oriented computing in the specific area of RRM by answering the following research questions: i) Which business functions and concrete methodologies are promising for the application of service-oriented computing in the area of RRM? What is the specific value proposition of service-oriented computing in this domain? ii) How can benefits and costs of service-oriented computing be quantified? For our analysis we will identify and discuss grid services in the domain of RRM from a management perspective. To answer the second research question we will concentrate on one concrete service—the estimation of covariance matrices—and propose

an economic model for the quantification of benefits and costs. Based on this example we will discuss how economic models can be applied in the context of service and resource markets.
The remainder of this text is structured as follows: In section 2 we describe our notion of RRM and provide fundamental principles. Section 3 is dedicated to the concept of service orientation and derives the value proposition of service-oriented computing in the domain of RRM. In section 4 we provide a quantification of benefits and costs for the covariance estimation service.

2 Risk/Return Management

In this section we will first provide basic principles of RRM. We will then describe specific aspects that prove to be complex and resource-intensive in practice and therefore motivate the application of grid services in this context.

2.1 Definition and Objectives

Almost all value generating business transactions are associated with some kind of uncertainty and thus contribute to an enterprise's overall risk exposure. Enterprises currently face rising pressure from competition on a global scale, delivered on integrated markets (especially integrated financial markets) where more and more complex products are traded on. Hence, in corporate practice an integrated view on risk and return is mandatory. Considering this integrated view, we propose the following definition: RRM is concerned with identification, quantification and control of risk together with the corresponding return associated with all kinds of corporate investment decisions. It is necessary to point out that, accordingly, the scope of RRM is by far larger than risk management alone. It covers all business functions that are concerned with investment decisions and therefore need to evaluate the risk and return of the available alternatives. Therefore, for example, also portfolio management can be seen as a part of an integrated RRM.
It is an important objective of RRM to satisfy the manifold information demands inside an enterprise. Current as well as future risks and chances have to be listed in risk reports in order to provide documentation and to improve transparency inside an enterprise. Moreover, ex ante decision support for the planning of investment or disinvestment decisions has to be provided, answering the question whether (and to what degree) an intended investment or disinvestment will lead to value added for the enterprise. To this end, accurate and up-to-date risk and return measures have to be quantified on the transactional level as well as on various aggregation levels along the organizational hierarchies like departments or business units. Especially for financial services institutions with their typically high exposure to market and credit risks this can be considered as crucial.
Prominent side conditions of RRM are the growing number of laws, rules and regulations aimed at the prevention of illiquidity and bankruptcy by restricting potential losses resulting from risky investment objects. Besides the German law concerning

"Control and Transparency in Corporations" (KonTraG) or the Anglo-American Corporate Governance Codex there exists a multitude of supervisory regulations. In the financial services industry, regulations are especially tight (e.g. as in the Basel II or Solvency II accord or in the Sarbanes-Oxley Act) and financial institutions are required to abide by certain well-defined risk limits , see e.g. Schierenbeck [2003, p. 505]. They keep significant capital reserves in order to retain their solvency even in extreme market situations. These capital reserves then restrict the overall earnings in terms of return from risky investments. Therefore it is an essential task for an integrated RRM to determine a suitable capital allocation against the background of risks being taken.

2.2 Complexity of Risk/Return Management Methods

RRM is a time-critical and resource-intensive[1] undertaking. It is time-critical, because regulatory constraints as well as internal information needs have to be fulfilled in certain, well-defined time frames, ensuring timeliness and relevance of the information delivered. It is resource-intensive, because there is a general trade-off between timeliness and precision of the underlying methods: Usually it is only possible—ceteris paribus—to improve in one dimension at the expense of the other. Excellence in both dimensions can only be achieved employing powerful resources. In the following we will consider fundamental and resource-intensive procedures in the area of RRM.

It is well known that the *estimation of distribution parameters* for risky investment objects is a fundamental problem in RRM, which is in fact often cited as "one of the important problems in finance" [Elton and Gruber, 1972, p. 409]. Many models of portfolio theory or capital market theory require for their practical application the estimation of input parameters. Usually, these parameters are then "fitted" by complex procedures to empirical values drawing on a multitude of historical data (the "history"). For example, Markowitz' "Portfolio Selection" model relies on expectations and covariances of the considered investment objects' returns in order to select an investor's optimal portfolio. Since the "true" return distribution parameters are unknown, this constitutes in fact an estimation problem. In this context, higher estimation precision can be achieved only with higher effort, be it by considering longer estimation time intervals or by implementing more complex computation procedures. We refer to e.g. Black and Litterman [1992], Hull and White [1998] or Kupiec [1995] and the literature referenced therein. Alexander [1996, pp. 233] provides an overview of techniques used in covariance estimation while some more recent approaches are presented e.g. in Taylor [2005].

A more specific task is the *estimation of covariance matrices*, containing all pairwise covariances between the investment objects' returns. In fact, covariances are used for risk quantification and aggregation in almost all areas of RRM, thus constituting a fundamental building block for many financial applications. Depending on the

[1] Here in terms of information technology (IT) resources, see subsection 3.

(possibly high) number of investment objects, these matrices typically contain a large count of covariances which all have to be individually estimated. Namely, we have for n investment objects $n(n+1)/2$ covariances, i.e. the number of matrix elements is increasing with approximately $n^2/2$. Assuming a history with length $N = 250$ working days per year for each investment object, we have for e.g. $n = 1.000$ investment objects already approximately $2Nn^2/2 = Nn^2 = 250.000.000$ historical values to take into account. It becomes apparent that, depending on the estimation method used, especially the estimation of covariance matrices is a very resource and time intensive problem.

A wide-spread method of RRM to deal with uncertain future values is *historical or stochastic simulation*. The latter is often also called “Monte-Carlo simulation”. In this context both forms have in common, that for a possibly large number of investment objects first the empirical return matrix has to be determined. For m so called “market factors” with N empirical (historical) prices each this is a $m \times N$ matrix. In the case of historical simulation one assumes that future values are independent and identically distributed conforming to the distribution of a representative historical sample. Stochastic simulation closely resembles this approach with the difference that future values are determined stochastically on the basis of an assumed distribution. In order to reduce calculation effort, most often the normal distribution is applied [Völker, 2001, pp. 76]. When a simulation approach is chosen, one usually relies on the “law of large numbers”, i.e. a very large number of simulation runs has to be performed in order to obtain reliable and statistically significant results. For instance, for $m = 100$ market factors with a history of 1 year each, the empirical return matrix with daily returns contains $mN = 25.000$ elements (assuming $N = 250$ working days per year). Performing e.g. 100.000 simulation runs this already leads to 2.500.000.000 values whose processing requires not only a lot of storage space but also significant computation capacity. We will point out in the following section how these requirements are addressed by service-oriented computing.

3 Service-Oriented Computing in Risk/Return Management

In academic research as well as in corporate practice a new paradigm for the utilization of information technology resources is en vogue. We use the term “service-oriented computing” in this text to address several (in fact converging) aspects: Elements from different domains like service-oriented architectures (abstract concept of services), grid technologies (grid middleware), distributed computing (load balancing, scheduling), or even from agent technologies (agent behavior, decentralized control) contribute to enable and facilitate the service-oriented computing principle. Services as well-defined software components are the pivot element of service-oriented architectures, whereas virtualization of infrastructure and resource sharing are primarily in the focus of approaches related to grid technologies or distributed computing. In our view the combination of both aspects sets the frame for service-oriented computing.

3.1 From Service-Oriented Architectures to Grid Services

According to Dostal et al. [2005, p. 7] we define service-oriented architectures (SOA) as the "abstract concept of a software architecture with focus on providing, searching and consuming so-called "services" over a network". Before we confine our notion of service-oriented computing, it is necessary to characterize two central terms in this context: resource and service. Following Neumann et al. [2006, pp. 206] we regard a *resource* as the representation of a logical or physical entity (like computing or data capacity, software licenses, hardware or network infrastructure). A *service* on the other hand provides a specific functionality and therefore aggregates the use of different (and in the context of service-oriented computing most often distributed) resources. Accordingly, a service in this sense is a software component designed to enable or support (part of) an enterprise's business process. Services are mainly employed by other services to perform a superior task. It makes sense to distinguish between two types of services: One speaks of "complex" or "composite" services on a high aggregation level in comparison to "basic" or "atomic" services on a low aggregation level as e.g. in Eymann et al. [2006, pp. 44] or Dodani [2006, p. 12]. Accordingly, it is characteristic for SOA that services are "loosely coupled" with each other, implying that they can be reused in a manifold way and that business processes can be dynamically configured with suitable, "best-of-breed" services. We will follow this perspective of "coarse granularity" and will provide examples for both service types offered in a SOA that can be used to "orchestrate" business processes, here in the context of RRM.

Regarding the related concept of grid computing, interestingly enough, the available definitions are mostly of descriptive nature and provide little more than certain essential characteristics. They are in most cases based on the seminal papers of Foster and Kesselman [1998] and Foster et al. [2001]. Various proponents have thereupon agreed that "a computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive and inexpensive access to high-end computational capabilities". It is an important characteristic, that grid computing reaches beyond administrative domains and thereby defines the so-called "virtual organization" as proposed by Foster et al. [2001]. For grid technologies, an evolution towards service-oriented architectures can be observed [Longworth, 2004]. We therefore speak of "grid services" in this text, denoting software components that realize resource intensive services based on grid technology. This term also conveys the close relationship between grid services and web services. The latter can be seen as a concrete implementation of SOA, based upon a set of various well-defined, mostly XML-based, standards, e.g. for message-based interaction (SOAP) and service description (WSDL). Grid services are based on specific web service standards as well, e.g. the specifications (OGSA) and (WSRF). They extend web services insofar as they imply the dynamic, yet for the user transparent, allocation of (physical) resources to services by some kind of a grid middleware. Our value proposition for service-oriented computing relies heavily on the virtualization of the underlying infrastructure.

3.2 Services for Risk/Return Management

In this section we will provide some examples for resource intensive services for RRM and classify them into complex and basic services. From a business perspective, at least two types of services can be distinguished: services for risk quantification and services (using risk/return information) for RRM, e.g. in the form of portfolio management.

A *risk quantification service* delivers one or several risk measures on different aggregation levels ranging from single investment objects up to the whole enterprise. Which measure should be applied depends on the type of risk (e.g. calculating credit risk differs significantly from calculating market risk) and on the objectives pursued. In corporate practice the probably most widespread risk measure is Value-at-Risk (VaR)[2], mainly because of its ease-of-use and simplicity. Furthermore, regulatory constraints require a frequent calculation of VaR on the enterprise level. Besides VaR, there exist a variety of other risk measures like variance or down-side risk measures, e.g. semi-variance. In a service-oriented environment a collection of services is conceivable, each calculating a specific risk measure.

Especially for trading securities on financial markets decision support systems are essential prerequisites that not only deliver precise up-to-date estimations of market parameters but also provide sophisticated optimization algorithms. In a service-oriented environment this could be achieved by a *portfolio optimization service*, after the necessary input parameters, e.g. variances and covariances, are determined. This essentially implies the determination of the efficient frontier, which is a NP-hard problem. Nevertheless, a numerical solution can be found for smaller problems using quadratic programming or applying heuristic approaches. The required precision of these heuristics, can only be achieved with significant computational effort. Especially on the trading floor it is furthermore mandatory, that information is available in real-time so that traders can quickly analyze different scenarios and react accordingly.

So far we distinguished services from a business perspective. Yet, as already stated in section 2, most business functions in RRM rely on fundamental methods like parameter estimation or simulation that could as well be realized as separated services and seized by complex services. For example, VaR can either be deduced analytically from the variances and covariances or can be determined using historical or Monte-Carlo simulation. Concrete examples for *basic services for parameter estimation and simulation* may be a historical data service, delivering e.g. an empirical return matrix

[2] According to Duffie and Pan [1997, p. 3] VaR is the “loss in market value over a time horizon H that is exceeded with probability $1 - \alpha$”.

consisting of the historical returns of the relevant investment objects, a regression analysis service, a stochastic or historical simulation service, a covariance matrix estimation service or a time series generation service (e.g. on the basis of ARCH or GARCH models). In section 4 we will describe one selected basic service—the covariance matrix estimation service—in more detail, also evaluating its application for the overall risk quantification of the enterprise.

3.3 Value Proposition of Service-Oriented Computing

The relevant, yet basic, value proposition of service-oriented computing concepts is almost canonical and to date already numerously stated, see e.g. Abbas [2004], Berman et al. [2003]or Foster and Kesselman [1998]. It delivers on demand computing power at transparent and relatively low cost (in comparison to dedicated, server-based computing) in combination with increased flexibility, scalability and a robust behavior against failure. This is achieved by using existing and/or standardized resources which are geographically and/or logically distributed in more or less autonomous units provoking high percentage utilization. In this context, Neumann et al. [2006, p. 206] distinguish between the possibility to perform computational intensive jobs in a reduced time frame, even enabling the solution of problem classes which so far could not be examined, and the possibility to share (not fully utilized) resources across different organizations. Yet, service-oriented computing is not primarily concerned with the utilization of free CPU cycles (which has been scrutinized for about 10 years by the disciplines of cluster computing and distributed computing). Instead, it provides new capabilities for intra- and inter-enterprise collaboration, enabling for example [Grauer, 2006, p. 71] world-wide communication and collaboration, access to high-performance computers for simulation or the verification of hypotheses, utilization of remote data sources and advanced visualization of research analyses. In the following we will elaborate on the basis of subsection 2 in more detail the specific value proposition of service-oriented computing in the context of RRM:

1. Because of computational constraints, RRM (especially the quantification of the overall risk position and adjustment of the capital allocation) to date is performed in relatively large and fixed time intervals. For instance, some enterprises are accumulating risk reports containing all relevant risk (and return) information only once every quarter or once a month, as reported by practitioners like e.g. Middlemiss [2004]. In other cases regulatory requirements demand that new information is available the next morning (after overnight batch runs) regarding the enterprise's risk position of the day before, as e.g. in Basel Committee on Banking Supervision [2004]. Whereas this might be sufficient in special cases, it is considered a disadvantage in general because investment decisions then are based upon outdated information. We figure that with service-oriented computing the necessary underlying calculations can possibly be dramatically accelerated without the need to invest into additional cost-intensive computing infrastructure. In this context we have the fol-

lowing scenario in mind: Because of integrated and possibly highly volatile markets, an integrated and up-to-date RRM is mandatory. It implies "sensitive behavior" regarding expected and unexpected market reactions as well as real-time determination of the current overall risk position. In order to achieve this goal, in this scenario RRM takes place (not in fixed time intervals but) dynamically and "event-driven", in much shorter time intervals. In fact we envision the "real-time enterprise" which is able to react appropriately almost immediately to all relevant and/or unforeseen market movements because it has timely and accurate risk/return information at its disposal. As soon as one calculation run is completed, immediately the next calculation starts. Depending on the current capacity of IT involved and the complexity of computation procedures the RRM time intervals are varying over time.

2. The input data needed for parameter estimation and forecasting in RRM is typically geographically and/or logically distributed. For instance, information about the portfolio of a globally investing enterprise might be scattered over several trading units in different locations. A common infrastructure based on grid services standards not only allows to share the data across organizational boundaries. The distribution of data also matches the fundamental structure of service-oriented computing and thus can be exploited for distributed processing where the data is available. In this case no centralization of data (causing communication and management complexity) is necessary and the advantages of service-oriented computing are fully exploited.

3. There is a trend towards higher frequency of input data first observed and published by Engle [2000]. Today, for example stock market data is widely available in a granularity down to the transactional level. The analysis of this "Ultra-High-Frequency" market data is a promising new area with implications for risk management not yet fully discovered. Service-oriented computing can contribute its share to storing and processing this huge amount of data providing up-to-date information and comprehensive analyses.

4. Because of the permanent movement and development of (financial) markets the demand for RRM calculations is itself far from constant. Using a dedicated infrastructure the enterprise is therefore committed to provide at any time a computing capacity aligned to the maximum demand during peak times. Additionally there is always the trade-off between RRM and other operations which have to be performed by the resources at hand. This challenge is met when unused resources can be seized at any time for additional speed and/or accuracy and in turn are available for daily operations in "quiet times".

5. There exists a variety of RRM applications that do not by all means require high-performance computing power at any given time. In contrast to the typical batch-processing mode they can possibly be executed cost-effectively on a service-oriented infrastructure in the background. With the available (and varying) computing capacity they can be performed continuously, including "slow business hours" like e.g. overnight time, and are still incessantly adding value for the enterprise. The

key point here is the automatic allocation of resources whenever they are available, increasing flexibility and manageability of existing corporate infrastructure.
6. Grid services offer new possibilities for intra- and inter-organizational collaboration regarding the integration, coordination and usage of resources and services. Concerning the inter-organizational provision of services and resources we need to clearly distinguish between high-level services that realize specific business functions (e.g. the quantification of risk) hiding the underlying grid-based calculations from the end-user and low-level services providing only the necessary computational capabilities. As Grauer et al. [2006, pp. 4] point out, especially for small and medium enterprises low level services are generally not satisfactory as these institutions most often lack the expertise necessary to implement complete RRM algorithms. Nevertheless we argue that in the financial services industry there is room for both approaches: For large financial services corporations that face fluctuating demand for computing capacity as discussed above, it may be interesting to contract additional computing capacity (on demand) from an external provider. Standardized solutions are not appropriate in this case as these institutions precisely try to get a competitive advantage on financial markets by applying proprietary methods and algorithms. Providers like SUN already offer the possibility to use their grid network to run resource intensive calculations on a pay-per-use basis. On the other hand small and medium enterprises demand for standardized high-level services in the area of portfolio management and risk quantification. Financial software or data suppliers like Reuters, Bloomberg and RiskMetrics already offer internet-based services e.g. for the calculation of portfolio risk, see for instance RiskMetrics Group [2006]. From here it is only a small step towards corresponding grid services.
Summarizing, we can state that service-oriented computing provides a highly suitable basis for the implementation of corporate RRM services. Exploiting the characteristic properties of grid services (like virtualized hardware at geographically distinct locations) an enterprise can benefit from more efficient risk/return calculations and more effective management procedures. At least for inter-enterprise collaboration there seems to be high potential for coarse-granular business services.

4 The Covariance Estimation Service

As pointed out above, covariances are an essential prerequisite for all kinds of financial risk calculations. In the following section we will consider a covariance estimation service that provides its user transparently with up-to-date covariance data for the relevant investment universe. To this end, we will propose an economic model for the quantification of benefits and costs that can be attributed to this service.

4.1 Economic Value of a Covariance Estimation Service

We consider an enterprise that frequently recalculates its risk position and therefore employs a service responsible for the estimation of covariance matrices. To this end we assume that the only purpose of covariance estimation is risk quantification.

Thus the service's economic value can be directly derived from the benefits of risk quantification alone. We perceive the enterprise as the weighted "sum" of its investment objects, i.e. its overall risk position is expressed as the portfolio risk, measured by the variance σ^2 of portfolio returns. We can calculate portfolio risk, resulting from n investment objects (numbered from 1 to n), using the covariance matrix, as $\sigma^2 = \sum_{i=1}^{n} \sum_{j=1}^{n} \mathrm{Cov}_{ij}$ with Cov_{ij} denoting the weighted covariance between investment objects i and j.
Since the enterprise is acting in an uncertain and dynamic environment its risk position is changing willingly (by making investment decisions) or unwillingly (by "movement" of the underlying markets). Because the estimation covariance matrices cannot currently be accomplished in real-time, the covariances at hand are always significantly outdated and therefore the variance calculated from the covariance matrix does not reflect the current risk position. We are in the following recurring to the fact that enterprises are adjusting their risk position to a value somewhere below a certain threshold thus constituting a "safety margin" (in the regulatory context often also called "haircut", see e.g. Basel Committee on Banking Supervision [2004, pp. 29]). They are doing so by using the capital allocation between risky and risk-free investment objects for balancing their overall risk position.[3] Our basic modelling approach is in the following: whenever covariances are available the safety margin can be adjusted immediately in a way that the resulting (and over time changing) overall risk position of the enterprise with high probability does not exceed the given risk limit at any time.
The estimation of one covariance matrix is assumed to take exactly T periods. The estimation of a new covariance matrix begins immediately after finishing the previous matrix, thus the service can provide a complete covariance matrix every kT periods ($k \in \mathcal{N} := \{1, 2, 3, \ldots\}$). According to [Buhl et al., 2005, pp. 7], the benefits B of risk quantification subject to T can then be calculated as[4]

$$B(T) = \left(i + \frac{\overline{\sigma}(\mu - i)}{q_\alpha \sigma_\sigma \sqrt{2T}} \right) \cdot K. \tag{1}$$

The parameters of the model are defined as follows. $K > 0$ denotes the enterprise's total capital, that is always completely allocated to a portfolio containing risky investment objects and/or a risk-free alternative. The portion of the enterprise's capital allocated to the risky portfolio yields the expected return μ, whereas the risk-free investment pays the time-invariant risk-free interest rate i. As we assume that always $\mu > i > 0$, the enterprise would fully allocate its capital to risky investment objects. Yet the enterprise is required to abide by the given risk limit $\overline{\sigma} > 0$

[3] At this point it is important to understand that the model presented in this text is not addressing the evaluation of the efficient set of investment objects or portfolio optimization (both problems also require the calculation of covariances), but the aggregation and management of the risk position of an enterprise.

[4] This text contains a shortened version of our model. The complete exposition including the assumptions of the model can be found in the working paper [Buhl et al., 2005, pp. 7] of the same authors.

and therefore needs to invest a certain share of the available capital into the risk-free alternative, depending on the intended safety margin. The safety margin itself is not modelled explicitly as a separate parameter, but results from the volatility of portfolio risk (denoted as σ_σ) and additionally from the frequency of the estimation, as higher frequency implies a more accurate determination of the current risk position. Consequently, the faster covariance estimations are accomplished (reflected by a smaller T), the smaller the safety margin can be, the less capital needs to be invested into the risk-free alternative and, ultimately, the higher are the overall benefits of the enterprise.

4.2 Value Proposition for the Covariance Estimation Service

We can now analyze our model with respect to the value proposition developed in section 3 and confine our findings for the special case of covariance matrix estimations. The first value proposition we identified was the possibility to generate additional value by accelerating RRM calculations because high-end computing capabilities are available for relatively moderate cost. It is difficult to quantify costs of our grid service as they depend on various situation-specific parameters and may comprise costs of physical resources as well as costs for the implementation, management and maintenance of the service itself, including e.g. user support by business experts. Yet in the following we are only considering computational resources as a cost driver, thus leaving the (end-user) perspective of a coarse granular service. More precisely, we assume costs to be proportional subject to the computing capacity needed for the estimation, reflected by a factor price p. Additionally, we restrict ourselves to the consideration of the time needed for computation, neglecting e.g. latency or transmission times. It is reasonable to assume that the so defined costs are moderate compared to server-based computing: In the case of covariance estimations the corresponding computations can be distributed on several resources, as all pairwise covariances can be calculated independently from each other. Efficiency losses are low and cost advantages actually take effect, as a number of low cost standardized components can provide the same capacity for covariance calculations than an expensive "traditional" server. Moreover, higher utilization levels can be expected due to on-demand allocation of resources. We will therefore analyze the effect of a decrease in costs on the overall benefits. There is a functional relationship between T and the computing capacity $z > 0$ necessary to complete the covariance matrix within the time frame T. With $n(n+1)/2$ covariances and w denoting the workload per covariance we have

$$T(z) = \frac{n(n+1)w}{2z}. \tag{2}$$

A larger z results in a smaller calculation time frame T which in turn leads to increasing benefits. Together with the cost side, expressed by pz, we can formulate the objective function Z (with decision variable z) as the difference of benefits and costs on the basis of equations (1) and (2) as[5]

$$Z(z) := Ki + \frac{K\overline{\sigma}(\mu - i)\sqrt{z}}{q_\alpha \sigma_\sigma \sqrt{n(n+1)w}} - pz \quad \rightarrow \max! \tag{3}$$

Applying the standard optimization procedure (i.e. solving $Z'(z) = 0$ for z) delivers as a distinct solution

$$z^* = \frac{K^2(\mu - i)^2 \overline{\sigma}^2}{4n(n+1)p^2 w q_\alpha^2 \sigma_\sigma^2}. \tag{4}$$

Since $Z''(z) < 0 \; \forall z$, the so defined z^* is a global maximum of the objective function. It also determines—using equation (2)—the corresponding time frame T^* our service should comply with, considering benefits as well as costs. Lower costs in equation (4) are reflected by a smaller price p leading to an increasing optimal capacity z^* and a shorter calculation time T^*. Thus it is economically reasonable to allocate more capacity, resulting in increasing overall earnings. More precisely, as the computing capacity depends quadratically on the price p, even a small decrease of p has high impact on capacity and calculation time. For covariance matrix estimation we therefore expect calculations to be performed more frequently, allowing enterprises to better exploit risk limits when the current risk position is determined in near- or even real-time.

We also observed in section 3 that service-oriented computing is advantageous concerning the processing of high data volumes. This constitutes an important aspect in the case of covariance matrices as well: First, the number of input values for the calculation is substantial as pointed out in section 2. Moreover, these values are usually distributed within the enterprise, e.g. geographically distributed according to different financial markets where, third, necessary market data is often provided by different information providers, e.g. Reuters or Bloomberg. A service-oriented infrastructure based on common standards facilitates the integration of data across organizational boundaries. Covariances could be calculated independently, e.g. on portfolio level, and then be aggregated as needed.

With our model we can also analyze the effect of changing market conditions on the demand for computing capacity, an aspect already mentioned by point 4 and 5 in section 3. As can be observed by equation (4), the demand for computing capacity is influenced by market- or enterprise-specific parameters. For example, the more capital the enterprise has to its disposal the more (in absolute terms) it will invest into risky investment objects. Higher risk exposure in turn increases the importance of RRM which is correctly reflected by a larger value for z^*. The same argumentation holds when the enterprise faces a higher risk limit $\overline{\sigma}$. In this case it should allocate more resources to RRM applications, which is consistently leading to an increasing z^* in our model. Eventually, when the risk premium $(\mu - i)$ rises, investing into risky objects becomes more attractive and profitable, resulting in a

[5] The objective function is the result of an analytically necessary and numerically justifiable approximation, see the working paper [Buhl et al., 2005, p. 14] for a complete deduction.

larger share of risky capital. In order to manage the consequently more voluminous portfolio our model proposes that additional computing capacity is necessary. In contrast to server-based concepts where resource allocation is fixed (at least for a certain time interval), with service-oriented computing resource allocation can be adjusted dynamically, realizing an optimal allocation at any point in time.

In practical settings the probably most important advantage of service-oriented infrastructures is the possibility to provide services across organizational boundaries. In section 3 we argued, that especially standardized services in the area of risk quantification seem to be well suited. We therefore envision a covariance estimation service that can be accessed via a service-oriented environment. As observed by Grauer et al. [2006, pp. 8] for the manufacturing industry, such a service-oriented environment needs to fulfill several other requirements, e.g. regarding security, accounting and billing or end-user access.

4.3 Model Application

One important characteristic of service-oriented computing is the possibility to share services and resources within and across organizational boundaries. Especially for resource sharing, grid computing concepts propose the application of market mechanisms to ensure an efficient allocation of available resources, see e.g. Wolski et al. [2001] or Buyya et al. [2002]. Yet a direct mapping of complex services to resources is not reasonable, as service consumers in general do not have concrete knowledge of the resources necessary to solve the problem at hand. To this end, Eymann et al. [2006, pp. 44] propose a two-tiered market. The set of services, together with means for provisioning and pricing, constitute the "service market". At the same time services demand for different kinds of resources, which are provided on the corresponding "resource market". Basic services are responsible for purchasing resources on the resource market.

This is an important application of economic models like the one we described in this section because for the pricing of services as well as for resource allocation it is necessary to quantify the economic value associated with service and resource consumption. In fact our model perfectly fits in this scenario: Covariance estimation can be regarded as a basic service whose benefits are derived when it is deployed for risk quantification, which may in turn be accomplished by a complex service. On the other hand the covariance estimation service needs to seize physical resources on the resource market to perform its calculations. Here, the model can be employed for resource allocation issues on the resource market. Regarding the pricing of a covariance estimation service it delivers the economic value of covariance estimations depending on the calculation frequency. This can be considered as a quality attribute of the service being the subject of service level contracts. For instance, when the service is provided externally, a concrete calculation time may be contracted. Considering the benefits as an upper bound (an enterprise would not pay more than its benefits), the price of the service is set depending on this calculation time.

5 Conclusion

In this paper we addressed two research questions concerning service-oriented computing and its applications in the field of RRM. First, we argued why complex and resource-intensive RRM calculations seem to be highly suitable to be performed on service-oriented infrastructures: Their specific properties and structure from our perspective almost ideally match with service-oriented infrastructure. We identified a set of potential services off the beaten track, like e.g. the estimation of covariances, to illustrate this close relationship. Second, we provided a model addressing the economic aspects of service-oriented computing.

We moreover argued, that our model can be applied in a market-oriented scenario where benefits and costs have to be evaluated. In fact, this in our opinion constitutes a prerequisite for the further development of adequate market-oriented approaches. For instance, auction mechanisms appear to be well-suited to ensure an economically efficient allocation of services and resources on the respective markets [Eymann et al., 2006, pp. 44]. We are currently working on the transformation of the model results for representing the valuation attributed to covariance estimation by market participants using a bidding language. Thereby we enable the regulation of access to services and resources depending on the individual priority, measured by the reservation price, ascribed by the service consumer, see Neumann et al. [2006, p. 207].

Apart from such market-oriented approaches there are open issues that need to be considered. With the exemplary list of services established in this text we intend to foster the development of service-oriented computing in the domain of RRM. Future work is planned to provide the proof-of-concept for basic services (parameter estimation or simulation) as well as for more complex RRM services. Here, problems concerning the technical implementation or security issues still need to be addressed. Additionally, economic effects and new business models resulting from the adoption of service-oriented computing in the area of RRM need to be analyzed. The availability of business applications, e.g. for RRM, is a critical success factor for the wide adoption and further development of service-oriented computing. From our perspective, there is plenty of room for further research in this new area, intertwining some of the most interesting ideas at the intersection of information systems and business domains.

Bibliography

A. Abbas. *Grid Computing: A Practical Guide to Technology and Applications*. Charles River Media, Hingham, 2004.

C. Alexander. Volatility and correlation forecasting. In C. Alexander, editor, *The Handbook of Risk Management and Analysis*, pages 233–260. John Wiley & Sons Ltd., West Sussex, 1996.

Basel Committee on Banking Supervision. International convergence of capital measurement and capital standards. http://www.bis.org/publ/bcbs107.pdf, 2004.

F. Berman, G. C. Fox, and A. J. G. Hey, editors. *Grid Computing; making the global infrastructure a reality*, West Sussex, 2003. John Wiley & Sons Ltd.

F. Black and R. Litterman. Global portfolio optimization. *Financial Analysts Journal*, 48(5):28–43, 10 1992.

H. U. Buhl, G. Fridgen, W. Hackenbroch, and M. Henneberger. An economic model for the allocation of grid resources to risk/return management. http://www.wi-if.de, 2005.

R. Buyya, D. Abramson, J. Giddy, and H. Stockinger. Economic models for resource management and scheduling in grid computing. *The Journal of Concurrency and Computation: Practice and Experience*, 14(Special Issue on Grid Computing Environments):1507–1542, 2002.

M. Dodani. SOA in an enterprise architecture. http://www-5.ibm.com/ch/events/soaas/, 06 2006.

W. Dostal, M. Jeckle, I. Melzer, and B. Zengler. *Service-orientierte Architekturen mit Web Services*. Elsevier Spektrum Akademischer Verlag, 2005.

D. Duffie and J. Pan. An overview of value at risk. *Journal of Derivatives*, 4:7–49, 1997.

E. J. Elton and M. J. Gruber. Earnings estimates and the accuracy of expectational data. *Management Science*, 18(8):B409–B424, 1972.

R. F. Engle. The econometrics of ultra high frequency data. *Econometrica*, 68(1): 1–22, 2000.

T. Eymann, D. Neumann, M. Reinicke, B. Schnizler, W. Streitberger, and D. Veit. On the design of a two-tiered grid market structure. In F. Lehner, H. Nösekabel, and P. Kleinschmidt, editors, *Multikonferenz Wirtschaftsinformatik 2006—Creating and Improving Business with Information Technology*, pages 41–52, Passau, 02 2006.

I. Foster and C. Kesselman, editors. *The Grid: Blueprint for a New Computing Infrastructure*, San Francisco, 1998. Morgan Kaufmann Publishers.

I. Foster, C. Kesselman, and S. Tuecke. The anatomy of the grid—enabling scalable virtual organizations. *International Journal of High Performance Computing Applications*, 15(3):220–222, 2001.

M. Friedman. Grid computing: accelerating the search for revenue and profit for financial markets. *Building an Edge*, pages 143–145, 2003.

M. Grauer. IT aus der Steckdose und Wirtschaftsinformatik. *Wirtschaftsinformatik*, 48(1):71–72, 01 2006.

M. Grauer, J. Reichwald, D. Christian, and T. Barth. The potential of service oriented computing for small and medium enterprises—analysis of value chains. In F. Lehner, H. Nösekabel, and P. Kleinschmidt, editors, *Multikonferenz Wirtschaftsinformatik 2006 – Creating and Improving Business with Information Technology*, Passau, 02 2006.

J. Hull and A. White. Incorporating volatility updating into the historical simulation method for value-at-risk. *Journal of Risk*, 1(1):5–19, 1998.

P. H. Kupiec. Techniques for verifying the accuracy of risk management models. *Journal of Derivatives*, 3:73–84, 1995.

D. Longworth. Grid lock-in on route to SOA. http://www.looselycoupled.com-/stories/2004/grid-lock-infr0831.html, 08 2004.

J. Middlemiss. Gearing up for grid. http://www.wallstreetandtech.com, 02 2004.

D. Neumann, C. Holtmann, and C. Orwat. Grid-economics. *Wirtschaftsinformatik*, 48(3):206–209, 2006.

A. Ricadela. Living on the grid. *InformationWeek*, 893:30–35, 2002.

RiskMetrics Group. Managed services. http://www.riskmetrics.com, 2006.

H. Schierenbeck. *Ertragsorientiertes Bankmanagement – Risiko-Controlling und integrierte Rendite-/Risikosteuerung*, volume 2. Gabler, Wiesbaden, 8 edition, 2003.

J. W. Taylor. Generating volatility forecasts from value at risk estimates. *Management Science*, 51(5):712–726, 2005.

J. Völker. *Value-at-Risk-Modelle in Banken*. Berlin Verlag, Berlin, 2001.

R. Wolski, J. S. Plank, J. Brevik, and T. Bryan. Analyzing market-based resource allocation strategies for the computational grid. *International Journal of High Performance Computing Applications*, 15(3):258–281, 2001.