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-168

Grid Computing Infrastructures and their Value for Risk Management

by

Wolfgang Hackenbroch, Matthias Henneberger

Januar 2006

to be presented at: Fortieth Annual Hawaii International Conference on System Sciences (HICSS-40), Waikoloa, Hawaii, (USA), January 2006

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

Grid Computing Infrastructures and their Value for Risk Management

Wolfgang Hackenbroch, Matthias Henneberger
University of Augsburg,
Competence Center IT & Financial Services.
86135 Augsburg, Germany,
www.wi-if.de
{wolfgang.hackenbroch,matthias.henneberger}@wiwi.uni-augsburg.de

Abstract

Risk/return management has evolved as one of the key success factors for enterprises especially in the financial services industry. It is highly demanding in terms of business requirements and technical resources, making it an almost ideal application for grid computing concepts. In this paper we analyze the value proposition of grid computing for risk/return management. We focus on a specific problem—the estimation of covariance matrices—and propose a model to quantify the benefits and cost of the corresponding calculations. Our model not only makes a contribution to understand the business value of grid computing in the domain of risk/return management, it also constitutes a building block for the development of economic resource allocation mechanisms.

1. Introduction

Whereas in the beginning, grid computing concepts were restricted to large-scale scientific applications like in high-energy physics, astronomy or biology, it has over the past years evolved to an increasingly relevant technology for the commercial sector as well. At the same time, it has become difficult to some extent to precisely differentiate grid computing from the related concepts of e.g. distributed computing, cluster computing or utility computing.

The meaning we intend to convey by our use of the term grid computing is best captured by a definition in [8]: "A Grid is a type of parallel and distributed system that enables the sharing, selection, and aggregation of geographically distributed, autonomous resources dynamically at runtime depending on their availability, capability, performance, cost, and users' quality-of-service requirements". An overview of the status quo and current applications of grid computing provide e.g. [5], [14], and [1].

The availability of grid enabled business applications seems to be a critical success factor for the wide adoption and further development of this powerful technology. One of the most promising application domains of grid computing concepts is the financial services industry with its information-driven business models, time-critical mathematical calculations and accordingly high needs for computing power. In fact, this sector is often mentioned among the key industries for grid computing applications, see for example [16] or [19].

In this context resource-intensive risk/return management applications seem to be especially suitable. With the potentially huge amount of computing capacity a grid infrastructure offers (embracing resources of the whole enterprise or even of external resource providers) such applications can possibly be accelerated dramatically. Yet even on a grid infrastructure resources are not unlimited. Thus for grid-based risk/return management it is necessary to find an economic optimum between the benefits of fast calculations on the one hand and the cost for allocated computing capacity on the other. We will therefore develop an economic model quantifying the functional relationship between computing capacity provided by a grid infrastructure and the corresponding economic value. For our analysis we focus on a specific problem in risk/return management—the calculation of covariances, which is a very complex and time-consuming assignment. With the paper at hand we are striving to provide the missing link between the capabilities of grid computing and its economically reasonable application in risk/return management.

The remainder of this text is organized as follows: In section 2 we give a short overview of risk/return management and the value proposition of grid computing. Section 3 examines how the economic value gained from fast covariance calculations can be quantified. We point out results and provide an interpretation in section 4. Section 5 concludes our considerations.

2. Grid-based Risk/Return Management

Risk/return management is concerned with the evaluation of risk and return associated with investment decisions. Whereas literature most often focuses solely on risk management, we will rather speak of risk/return management emphasizing an integrated view because risk management can only unfold its potential in combination with the management of the corresponding return. In this section we will give a brief review of risk/return management and its application on a grid infrastructure.

2.1. Principles of Risk/Return Management

Enterprises are investing capital into risky investment objects in order to generate cash inflows and subsequently increase the return of the invested capital. Generally higher return is systematically associated with higher risk. This connection is theoretically explained by economic models like the "Capital Asset Pricing Model", that was developed by Lintner, Sharpe, and Mossin and empirically verified later on. Following the argumentation of [21, pp. 194] it is therefore crucial for the survival (i.e. not going bankrupt) and success (i.e. achieving the most out of the invested capital) of an enterprise to be able to allocate the available capital to the right combination of investment objects taking into account their specific contributions to the overall risk and return. This is especially important for the financial services industry where risk taking is an essential part of the business model. Yet its application domain is not confined to the traditional asset management context. In fact almost all business transactions are associated with uncertainty and thus contribute to an enterprise's overall risk exposure. Thus the management of risk and return is a fundamental business function in every enterprise.

One major goal of risk/return management as pointed out above is the prevention of bankruptcy by restricting potential losses resulting from risky investment objects.[1] The increasing importance of this goal is emphasized by a growing number of rules and regulations that require enterprises to hold an adequate part of their available capital to back their risky investments (see e.g. [18, pp. 8]). This share of the available capital then makes less or no contribution to the overall earnings. In the financial services industry rules and regulations are especially tight (e.g. Basel II, Solvency II or the German rules MaH and MaK) and institutions are required to keep significant capital reserves to retain their solvency even in extreme market situations. By management decisions these restrictions are broken down along the organizational hierarchies into guidelines representing limits for the maximum risk a business unit is willing (or able) to take. In order to abide by the given risk limits the knowledge of the current overall risk position is an essential prerequisite. Consider for example a trading unit in an investment bank that needs exact and timely information about its risk exposure when deciding on security or option trades. It is one fundamental challenge of risk/return management to provide accurate measures of risk to assist these decisions.

Suitable and widely accepted measures of risk are variance, standard deviation (as the positive square root of variance) or synonymously volatility (as discussed in [20, pp. 154]), and especially for defining regulatory risk limits the Value-at-risk (VaR). Finance offers a set of instruments to calculate these measures, ranging from Monte Carlo simulation to regression analysis of historical data. Some of them seem to be promising with regard to grid computing concepts as well, see e.g. [19]. Yet in this paper we focus on the "variance-covariance approach" and especially on the estimation of covariance matrices containing the pairwise covariances between investment objects. They are used to determine the variance of a portfolio of investment objects taking into account correlation effects that exist between them. Although we concentrate on portfolio variance as a measure of risk it is worth mentioning that the variance-covariance approach is also deployed to determine—under certain assumptions—the VaR of a portfolio, making it a fairly fundamental concept.

2.2. Value Proposition of Grid Computing

Mainly due to computing capacity restrictions[2] risk/return management calculations are conducted at most on a daily or overnight basis. Practitioners like [19] even report that they are usually performed weekly or even over longer time intervals. 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, using a grid, the necessary underlying calculations like covariance estimations can possibly be dramatically accelerated without the need to invest into additional cost-intensive computing infrastructure.

Accordingly, the relevant, yet basic, value proposition of grid computing is almost canonical and already numerously stated: 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

[1] We do not restrict our analysis to securities but also associate with the term investment object e.g. customers or projects.

[2] Obviously there exist other restrictions as well like the availability and quality of input data.

distributed in more or less autonomous units, provoking a high percentage utilization. In the following we elaborate in more detail the specific value proposition of grid computing in the context of risk/return management:

1. The input data needed for estimation and forecasting is usually geographically and/or logically distributed. This matches the fundamental structure of a grid system and thus can be exploited for distributed processing because the corresponding parts of the covariance matrix are estimated where the data is available. No centralization of data (causing communication and management complexity) is necessary and the advantages of a grid infrastructure are fully exploited. For instance, portfolio information of a globally acting enterprise might be scattered over several trading units in different locations.

2. There is a trend towards higher frequency of input data first observed and published by [11]. Today, for example stock market data is widely available 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. Grid computing can contribute its share to storing and processing this huge amount of data providing comprehensive and up-to-date information.

3. Because of the permanent movement and development especially of financial markets the demand for risk/return management 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 risk/return management 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" .

4. There exists a variety of 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 grid infrastructure in the background. With the available (and varying) computing capacity they can be performed continuously including "slow" business hours like e.g. overnight and are still incessantly adding value for the enterprise. Next to covariance estimation eligible candidates range from the calculation of β-factors up to Value-at-Risk calculations and back-testing. The key point here is the automatic allocation of resources whenever they are available, increasing the flexibility and manageability of existing infrastructure.

5. Last but not least grid computing offers new possibilities for intra- and inter-organizational collaboration regarding the integration, coordination and usage of resources. In this context especially the possibility to contract additional computing capacity from an external provider on a pay-per-use basis seems to be promising. As already stated above the demand for computing capacity is changing over time. Following the grid paradigm enterprises can buy resources on-demand paying only for computing capacity that is actually used.

Concerning this value proposition one question inevitably arises: What benefits and cost are associated with grid-based risk/return management? Two different perspectives can be distinguished. In the long run an enterprise decides on how much capacity it wants to acquire, e.g. by investing into grid infrastructure or by contracting with external service providers. In the short run it has to choose (ideally in real-time) how much of its available resources—usable for various tasks like calculating customers' portfolios or storing operational data—it wants to allocate (and pay for). In combination with the corresponding cost (as proposed in subsection 3.2) our valuation model enables a mechanism to ensure an economically efficient planning and allocation of required computing capacity provided by grid resources.

Although the quantification we suggest is to some extent also applicable outside the domain of grid computing, it complements existing economic grid resource allocation mechanisms. It does not deal with physical resource allocation but with the valuation of computing capacity on a higher abstraction level, as proposed e.g. by [10].

Furthermore, grid-based covariance estimation can be regarded as a "service" which is employed by several other services or financial applications. The set of services, together with means for provisioning and pricing, may constitute a "service market" as laid out in more detail by [13]. At the same time services demand for different kinds of resources (like CPU time, data storage capacity, software licenses etc.), which are provided on a corresponding grid "resource market". The abstraction essentially implies that the service consumer has no concrete knowledge of physical resources necessary to solve the problem at hand. Yet, for resource allocation, a valuation of the service from the perspective of the service consumer is mandatory.

The research questions addressed in this paper fit seamlessly into this setting: We consider a "covariance matrix estimation service" that provides its user transparently with up-to-date covariance data for the relevant investment universe. By using a distributed, grid-based computation service, covariance matrices can be estimated in parallel, which increases the update frequency and actuality of the data. Thereby we try to narrow the gap between mainly economical questions (concerning the service market) and technical questions (relevant for the resource market).

3. A Valuation Model for Fast Covariance Calculations

For our valuation model we consider an enterprise frequently calculating its risk position by computing the covariance matrix of all investment objects it is engaged in. 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 of the covariance matrix cannot currently be accomplished in real-time the covariances at hand are always significantly outdated. 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", to ensure that the overall risk remains below the given risk limit at any time.[3] They are doing so by using the capital allocation between risky and risk-free investment objects for balancing their overall risk position.[4] 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. Hence the faster it can calculate covariances the smaller the safety margin can be. This text contains a shortened version of our model. The complete exposition can be found in the working paper [7, pp. 7] of the same authors.

3.1. The Risk-at-Risk Approach

The discrete time horizon at hand consists out of $m > 1$ equidistant periods. We will write e.g. r_t to indicate the value of a parameter at the end of period $t = 0, 1, \ldots, m$. We are considering an enterprise equipped with a total capital of $K,\ K > 0$, which has access to and is engaged in a set of n risky investment objects and an additional risk-free investment alternative. (Dis-)investments are performed in the standard Markowitz setting, i.e. for example the enterprise is generally risk-averse and striving for efficient combinations of investment objects, which are perfectly divisible and traded on a no-frictions market. We assume that the available capital K is always completely allocated to the risky portfolio and/or to the risk-free alternative and will denote the risky portion of K by x (with $x \in \mathcal{R},\ x \geq 0$).

Future returns of the portfolio are modelled as independent stochastic variables. Their probability distribution for each period can be characterized by mean and standard deviation. This obviously implies that the investment objects can be "marked to market", i.e. that there is a (current and historical) price attached to them.[5] Following the idea of random walks, historical portfolio returns can be used for estimating mean and standard deviation (or volatility) of future portfolio returns. The risky portion of the enterprise's capital yields the expected return μ, the risk-free investment pays the time-invariant risk-free interest rate i, which is equal to the borrowing rate. We assume that always $\mu > i > 0$. Thus the total benefit B per period can be written as

$$B(x) = K(i + x(\mu - i)). \tag{1}$$

The overall risk position of the enterprise (we perceive the enterprise as the weighted "sum" of its investment objects) is expressed by the portfolio risk, measured by the variance of the portfolio returns, σ^2. We can calculate the 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.[6] The covariance of an investment object to itself is its variance, i.e. $\mathrm{Cov}_{ii} = \sigma_i^2$. Because of the symmetry of the matrix the total number of different (co)variances is according to $\frac{1}{2}n(n+1)$.

For the investment portfolio of the enterprise the covariance matrix is not previously known. Covariances can be empirically estimated, or forecasted, by analyzing historical data. Although we do not want to focus on the different methods of volatility and correlation forecasting (see e.g. [17] or [2] for details) it becomes clear that the corresponding computations are not trivial. Depending on the estimation method used, the calculation of these matrices is a very resource and time intensive problem.[7]

Furthermore 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 we have a complete covariance matrix every kT periods ($k \in \mathcal{N} := \{1, 2, 3, \ldots\}$), see fig-

[3] In the regulatory context this is often called "haircut", like in [4, pp. 29]

[4] 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.

[5] It is common practice to use some variation of a random walk model for the price movement on security or commodity markets. This approach goes ultimately back to the early and path-breaking contribution of Louis Bachelier who proposed his findings for the stock market—comparing price movements with a "drunkard's walk"—in his thesis paper [3].

[6] The aggregated risk can be calculated according to this formula irrespective of the return distribution. Nevertheless using the variance as risk measure has implications on the presumed return distribution.

[7] In fact the data and computing intensity needed for covariance calculation is often considered a major disadvantage of the variance-covariance approach.

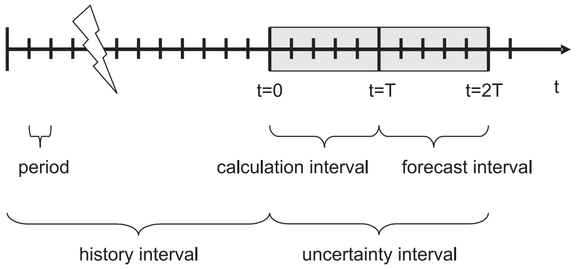

Figure 1. Period model and relevant time intervals

ure 1.[8] Whenever a covariance matrix is ready the input data used for its estimation is exactly T periods old. We can immediately determine the portfolio risk by appropriately summing up the covariances in the matrix. The latter can then be used for a risk adjustment decision as well as for portfolio optimization until the next matrix is finished. The moment before this happens the input data is already $2T$ periods outdated. Therefore the uncertainty interval that has to be taken into account spans $2T$ periods: in the worst case the risk has been going up over $2T$ periods before the enterprise realizes that it is exceeding the maximum risk it is willing (or able) to take.

The portfolio risk is fluctuating during this time in a non-predictable way because of changing (market) values of the underlying investment objects and varying covariances of their periodical returns respectively. We therefore model the portfolio risk at time t, denoted as σ_t, as well as a stochastic variable, governed by a random walk. We assume in particular that the σ_t are approximately normally distributed with (strictly positive) mean μ_σ and standard deviation σ_σ (both tagged with an index σ indicating the fact that the parameters apply to the portfolio *risk*), i.e. $\sigma_t \sim \nu(\mu_\sigma, \sigma_\sigma)$.[9] This relates to a phenomenon known from the empirical behavior of stock market prices called heteroscedasticity, see e.g. [20, pp. 152]. It is characteristic for the spirit of ARCH and GARCH models where the covariance matrix is time-varying as described e.g. in [6] or [12].

As discussed above the enterprise is required to abide by a given risk limit. We follow an approach comparable to the VaR to account for this fact. On a formal level, the VaR is the α quantile of a financial return, r_t, and is defined as the value q_α, for which the probability $P(r_t \leq q_\alpha) = \alpha$. The VaR for a portfolio associated with a holding period and a confidence level α can therefore be characterized as the expected loss that is after the holding period only exceeded with the (small) probability of $(1-\alpha)$. In analogy to the VaR we assume that the overall risk of the enterprise should exceed the risk limit $\overline{\sigma}$, $\overline{\sigma} > 0$, only with the (small) probability of $(1-\alpha)$.

We speak of the Risk-at-risk over a holding period and with the confidence level α, $0 \ll \alpha < 1$ in this context. When the overall risk position of the enterprise is calculated in period T, the portion x of risky capital for the next T periods is adjusted considering the Risk-at-risk over the uncertainty interval of $2T$. That means the overall risk of the enterprise, $x\sigma_t$, determines the value of x: Quantifying the risk limitation as $P(x\sigma_t \leq \overline{\sigma}) \geq \alpha$ for $t = 2T$ it follows—because of the $\nu(\mu_\sigma, \sigma_\sigma)$ distribution of σ_t and with q_α as the (onesided) α-quantile of the standardized normal distribution—that

$$\frac{\overline{\sigma} - x\mu_\sigma}{x\sigma_\sigma\sqrt{2T}} = q_\alpha. \tag{2}$$

Solving this equation for x and inserting it into equation (1) gives us the benefits subject to T. We will now direct our attention towards the corresponding cost associated with covariance estimation when they are performed on a grid infrastructure. Both elements are then used to deduce an optimal, risk-adjusted allocation of grid capacity.

3.2. Optimization

Since in this paper our focus lies on the economic aspects of grid computing technology, we will in the following apply a straightforward cost function for the (opportunity) cost side of grid-based covariance calculations:[10] The enterprise decides on a computing capacity of z, $z > 0$, units of computing capacity[11] per period used for covariance calculations. The amount of computing capacity needed for the calculation of one single covariance is denoted by w, $w > 0$. The workload for one covariance basically depends on two factors: The estimation method used and the length of the history interval considered. Both have significant influence on the complexity and the resource consumption of the corresponding calculations. Abstracting from a specific procedure we assume a given estimation method with a (moving) history interval of fixed length, and therefore w to be constant. Additionally we restrict ourselves to the consideration of the time needed for computation, neglecting e.g. latency or transmission times, and correspondingly to the cost for computation which occurs in the form of an (internal or external) factor price p per unit of grid

[8] It is important to distinguish the *calculation interval* T from the length of the *history interval* used for estimation.

[9] This can be seen as a consequence of the central limit theorem. In practice the rigorous distribution assumption would be relaxed by using the true distribution of the σ_t delivered by the grid calculated sequence of standard deviations.

[10] Refer to the working paper [15] for thoughts on grid computing technology in risk/return management as well as for an adequate basis for developing a more elaborate cost model.

[11] Computing capacity could be measured, for instance, in the often-used unit of "floating point operations" per period. Nevertheless, in this text we abstract from a concrete measure.

computing capacity. Thus we can state for the total cost of computation that

$$C(z) = p \cdot z. \tag{3}$$

There is a functional relationship between T and z describing the computing capacity per period necessary to complete the covariance matrix within the interval T. With $\frac{1}{2}n(n+1)$ (co)variances we have

$$T(z) = \frac{n(n+1)w}{2z}. \tag{4}$$

The enterprise is striving to maximize its profit by adjusting the grid capacity utilized. Thus z forms the decision variable of our objective function. Putting it all together enables the formulation of the objective function, $Z(z)$, as[12]

$$Z(z) := Ki + \frac{K\overline{\sigma}(\mu - i)\sqrt{z}}{q_\alpha \sigma_\sigma \sqrt{n(n+1)w}} - pz \quad \rightarrow \text{max!} \tag{5}$$

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

$$z^* = \frac{K^2(\mu - i)^2\overline{\sigma}^2}{4n(n+1)p^2 w q_\alpha^2 \sigma_\sigma^2}. \tag{6}$$

Since $Z''(z) < 0\ \forall z$ the so defined z^* is a global maximum of the objective function and thus determines the optimal grid allocation and—using equation (4)—the corresponding interval T^* the enterprise should comply with for grid-based covariance estimation. Exceeding T^* will result in less than optimal capital allocation in the risky yet profitable investment objects while falling short of T^* will entail larger than optimal cost for managing the risk.

4. Results and Interpretation

We will now take a closer look at our model and the implications for grid based risk/return management. Most of the results determined by equation (6) have readily intuitive explanations. For example the more capital the enterprise has to its disposal the more (in absolute terms) it will invest into risky investment objects.[13] Higher risk exposure in turn increases the importance of risk/return management 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 grid resources to risk/return management 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 larger share of risky capital, x. In order to manage the consequently more voluminous portfolio our model proposes that additional grid capacity is necessary. A most interesting and to some extent counter-intuitive result is produced in combination with the parameter σ_σ. One would possibly expect that, with increasing volatility of the portfolio risk (expressed by a larger σ_σ), the optimal grid allocation goes up as well. That this is not true shows equation (6). Basically for a given risk limit $\overline{\sigma}$ a higher volatility of the portfolio risk can be leveraged by faster covariance calculations (so that the enterprise can still get close to the risk limit without hazardously exceeding it during the uncertainty interval). The cost on the other side are depending on T in a reciprocal fashion. As a consequence, for higher volatility of the portfolio risk the cost are increasing more quickly than the benefits leading ultimately to an increasing T^* and decreasing z^*, respectively.[14]

We modelled the cost side by a price p per grid computing capacity. It can be interpreted differently depending on the way resources are provisioned: Firms could either build their own grid infrastructure embracing resources of the whole enterprise or computing capacity could be provided by external service providers. When resources are provisioned externally the price is given by the provider and the enterprise needs to decide on the amount of resources it wants to allocate and pay for. In an internal scenario prices need to be interpreted as opportunity cost or transfer prices.[15] Our model then is valuable to give an indication for the economically efficient allocation of the shared resources depending on the (variable) demand. We regard this as characteristic for the application of grid technology: depending on current requirements a service user (here: the estimation of covariances) can consume exactly the amount of resources needed (provided "on demand" by a service provider like the enterprise itself and/or external supplementation). In the contrary for a dedicated system the resource allocation, once determined, can not be adjusted. Instead it would be optimal to use all of the available resources at a time as opportunity cost are zero.

We will conclude our considerations by a short comparison of grid computing and "traditional" server based computing with respect to our model. We stated that grid com-

[12] Equations (1), (2), and (4) deliver the benefits subject to z. $Z(z)$ is then the result of an analytically necessary and numerically justifiable approximation of the difference between $B(z)$ and $C(z)$. Refer to the working paper [7, p. 14] for a complete deduction. In this paper we also provide the rationale for the approximation.

[13] This is always true under the assumptions made in this context: risky investment objects have a higher return than the risk-free alternative and the enterprise is trying to maximize its overall return. Note that the *share* of capital allocated to risky investments, x, is independent of the enterprise's total capital, K, because the risk limit, which is maintained using x, is a given, exogenous parameter.

[14] Such a situation could be observed e.g. during a "regime switch" between two volatility clusters where a time interval with relatively low volatility switches into an interval with higher volatility or vice versa.

[15] In fact economic allocation mechanisms determining prices based on supply and demand are widely discussed in the context of grid computing, see e.g. [9].

puting delivers high-performance computing capacities at low cost. We can now render this aspect more precisely with regard to the calculation of covariance matrices. The corresponding computations can be distributed on several nodes, as all pairwise covariances can be calculated independently from each other. Thus 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 server. Moreover higher utilization levels can be expected due to on-demand allocation of resources as described above. In our model lower cost are reflected by a decreased price p leading to an increasing optimal capacity z^* and a shorter calculation time T. Therefore grid computing—apart from other aforementioned benefits—reaches out for risk/return management and qualifies as a suitable infrastructure to perform the corresponding calculations. In the case of covariance matrix estimation we expect calculations to be performed more frequently allowing enterprises to better exploit risk limits as the current risk position is determined in near- or even real-time.

The discussed results are focussing on the economic view on grid resource allocation. Regarding concrete application scenarios of our model, additional aspects should be considered. Our model could not only be employed to ensure an economically efficient allocation of computing capacity within an enterprise, it could also be utilized to ensure adequate pricing of an external grid-based covariance estimation service.[16] Today financial software or data suppliers like Reuters, Bloomberg or RiskMetrics already offer online services e.g. for the calculation of portfolio risk. From here it is only a small step towards corresponding grid-based services. Questions concerning the technical implementation, security aspects or the specific characteristics of grid networks in this context provide room for future research.

5. Conclusion

In this paper we restricted our analysis to one well-defined problem: the grid-based estimation of covariance matrices. Although covariances are widely used in financial applications, we thereby covered only a small subset of risk/return management methods and algorithms. Other approaches and applications for grid computing (like e.g. Monte-Carlo simulations which have a high parallelization potential as well) still have to be evaluated regarding benefits and cost. We demonstrated how the value derived from risk/return management calculations can be measured. We developed a model considering an enterprise that has to decide on the amount of capital it wants to allocate to cover potential losses resulting from a risky investment portfolio. Several assumptions and an approximation (for the objective function) were necessary to get to an analytical solution. Moreover we targeted our model to the calculation of covariances from scratch. In reality one would probably calculate covariances in a continuous fashion, i.e. reuse the results of the previous calculation or compute only the covariances that have changed during the last period thus reducing the necessary grid capacity and therefore ultimately the cost of calculation. Moreover it is a simplification to assume that it is possible to attach a well defined factor price to computing capacity. In a realistic scenario our abstract view on resources needs to be broken down to actual physical resources.

Despite these limitations, the basic principles introduced in this paper can be adapted to other scenarios in more sophisticated and complex surroundings, including e.g. load balancing or security issues. We are convinced that an economic analysis is relevant for many grid computing applications. It complements emerging concepts of market-oriented resource allocation, e.g. in the context of grid-based service and resource markets. Market mechanisms in connection with individual utility functions that value benefits and cost of resource consumption guarantee an economically efficient allocation of resources within an enterprise. Thus paving the way to grid-enable risk/return management and providing rules for the flexible and economically efficient allocation of grid resources will in our opinion lay the foundation for a truly real-time and service-oriented enterprise supporting all kinds of business critical decisions.

[16] Of course in the case of an externally provided service enterprise specific parameters of the model like the overall capital or the individual risk limit need to be estimated or the model need to be adapted.

References

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

[2] 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.

[3] L. Bachelier. Théorie de la spéculation. *Annales Scientifiques de l'École Normale Supérieure*, 17:21–86, 1900.

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

[5] 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.

[6] T. Bollerslev, R. F. Engle, and J. M. Wooldridge. A capital asset pricing model with time-varying covariances. *Journal of Political Economy*, 96:116–131, 1988.

[7] 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.

[8] R. Buyya. Grid Computing Info Centre: Frequently Asked Questions. http://gridcomputing.com/gridfaq.html, 2004.

[9] 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.

[10] A. Dan, C. Dumitrescu, and M. Ripeanu. Connecting client objectives with resource capabilities: An essential component for grid service management infrastructures. In *IC-SOC'04*, New York, 11 2004. ACM.

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

[12] R. F. Engle and K. F. Kroner. Multivariate simultaneous generalized ARCH. *Econometric Theory*, 11:122–150, 1995.

[13] 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.

[14] I. Foster and C. Kesselman, editors. *The Grid 2: Blueprint for a new Computing Infrastructure*, San Francisco, 2003. Morgan Kaufmann Publishers.

[15] G. Fridgen and W. Hackenbroch. On the complexity of grid-based risk management calculations. http://www.wi-if.de, 2005.

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

[17] 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.

[18] P. Jackson, D. J. Maude, and W. Perraudin. Back capital and Value at Risk. *Bank of England working papers*, 79, 1998.

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

[20] K. Spremann. *Portfoliomanagement*. Oldenbourg Wissenschaftsverlag GmbH, München, 2 edition, 2003.

[21] T. C. Wilson. Calculating risk capital. In C. Alexander, editor, *The Handbook of Risk Management and Analysis*, pages 193–232. John Wiley & Sons Ltd., West Sussex, 1996.