

Parallel computing applications and financial modelling

Heather M. Liddell, D. Parkinson, G.S. Hodgson and P. Dzwig

Department of Computer Science, Queen Mary, University of London, Mile End Road, London E1 4NS, UK

E-mail: heather@dcs.qmul.ac.uk

Abstract. At Queen Mary, University of London, we have over twenty years of experience in Parallel Computing Applications, mostly on “massively parallel systems”, such as the Distributed Array Processors (DAPs).

The applications in which we were involved included design of numerical subroutine libraries, Finite Element software, graphics tools, the physics of organic materials, medical imaging, computer vision and more recently, Financial modelling. Two of the projects related to the latter are described in this paper, namely Portfolio Optimisation and Financial Risk Assessment.

Keywords: Parallel computing, massively parallel systems, distributed array processor, financial modelling, portfolio optimisation, risk assessment

1. Introduction

The UK has a long history of leading the world in innovative computing – this is particularly true in High Performance Computing and Communications (HPCC). Within this area there has been growing emphasis on the use of “massively parallel systems”, comprising many thousands of processors working together. Two systems capable of massive parallelism are the Distributed Array Processor (DAP) and the transputer; both were invented in the UK. The DAP activity at QMW formed a nucleus for parallel processing, from which evolved a number of research groups and specialist centres in the UK ; more importantly, a large community of users was stimulated to use parallel computers for a wide variety of problems in Science, Engineering and other disciplines. High Performance Computing (and hence massively parallel processing) is of strategic national importance, both as an industry in its own right and also because of the requirements of other industries. Senator Gore, the former US Vice president stated as long ago as 1989 that “the nation which most completely assimilates high-performance computing into its economy will emerge as the dominant intellectual, economic and technological force in the next (i.e. this!) century”.

The use of massively parallel systems required the development of new algorithms, software tools and application techniques; a reliable migration path must be provided from current systems – this is particularly important in the commercial world. These are the challenging areas we have been addressing in our research.

2. Parallel computing applications on the distributed array processor

The work we are going to describe is the result of over twenty years of research activity at Queen Mary. First the DAP Support Unit (DAPSU) was formed then the Centre for Parallel Computing, a unit established by Queen Mary in 1987. We were later joined by Prof. Yakup Paker and his group. Finally the London Parallel Application Centre (LPAC), a consortium comprising Queen Mary, Imperial College, University College London and City University was established. Our work has always been interdisciplinary – we worked with engineers, scientists and social scientists. We were also involved in many collaborative projects with other institutions in the UK, in Europe and on a wider international scene. An important aspect of our research concerns “technology transfer” to industry and commerce.

Our first “massively” parallel machine was the ICL 4096 processor DAP [1], funded by the UK Computer Board, which arrived in 1980. This marked the beginning of 12 continuous years of a national DAP service. The original ICL DAP was replaced by another 4096 processor machine, the AMT (Active Memory Technology) DAP 610, in 1988. A second generation ICL 1024 processor “mini-DAP” formed part of a grant from the SERC (Science and Engineering Research Council), the forerunner of EPSRC, the Engineering and Physical Sciences Research Council. One essential requirement for the growth of use and acceptance of parallel computing technology is the provision of efficient system software, portable software tools and applications software. Much of our application software took the form of specialist program libraries and tools. Our approach has been to attempt to provide libraries that have a similar user interface and documentation style to that available in the NAG library, but which are specifically designed for parallel machines. A major product was the DAP Subroutine library [2], marketed by AMT. The DAP Finite Element Library [3] was produced for engineering applications and demonstration programs were also developed. Various members of our Centre were involved in the design, implementation and validation of data-parallel languages – Fortran Plus, C++, AI/Logic languages, such as Lisp and Prolog, and the functional language HASKELL.

The focus of our activity has always been towards users of parallel systems. We have been actively engaged in many different application areas – engineering, physical sciences, computer vision, medical imaging and econometrics. One example was an Esprit Basic Research project, OLDS [4], in which we were working with colleagues in Physics and Chemistry at QMW and with European partners, exploring the potential use of organic materials in novel micro-electronics devices. Another was a collaborative medical imaging project, MIRIAD, in which we investigated methods for the display and segmentation of features in 3D images [5,6].

LPAC’s major objective was technology transfer, primarily via collaborative projects carried out with Industrial/Commercial partners and with applications experts in our own establishments, to produce useable applications software, portable to many different parallel computing systems, and to provide migration paths to these systems. We considered the European dimension to be particularly important. We continued to address the more traditional industrial and scientific computational problem areas, including industrial modelling,

use of neural networks and visualisation, but considerable emphasis was placed on the commercial and financial applications, partly because of our geographical proximity to the City of London. These were new fields for academic computing people, but we were fortunate in being able to collaborate with colleagues having expertise in these areas. The rest of this paper will concentrate on that aspect of our work.

3. Application of parallel and distributed computing to portfolio optimisation

The financial markets form one of the most important engines of industry. This project was particularly interested in understanding and improving the methodology and science behind the creation and maintenance of large portfolios of financial instruments (bonds, shares etc) through the deployment of large scale numerical computing techniques.

Portfolios are collections of financial instruments whose overall intention is to, for example, minimise the risk caused by exposure to market fluctuations. From our perspective they are a collection of instruments governed by non-linear dynamics that reflect market dynamics. The requirement for the financial engineer is to select a portfolio in an optimal way. In reality such solutions are, of course, sub-optimal.

The modelling of classical portfolios and understanding their dynamics has a long history, the landmark of which was Markowitz’s paper [7]. This involves creating a model in the form of a stochastic differential equation, whose stochastic terms are reflected in volatility. This volatility is conventionally taken as a constant. In prosaic terms an investor makes an estimate of what the growth of a particular component may be and this estimate becomes a parameter of the calculation. The solution to the optimisation problem can then be found by deploying conventional optimisation techniques. Naturally the growth is neither constant nor easily determined, although there may be good reason for believing that it will fall within certain bounds.

Much work has been done internationally on the modelling of individual instruments under varying conditions. In comparison little has been done on the optimisation of medium to large-scale portfolios with the use of high-density historical data sets to drive the inputs to the models. In all there may be several tens of such instruments in a portfolio (which is itself an instrument). The aim is then to model the behaviour of this problem that has perhaps fifty dimensions. In our

research we set out to apply the techniques of high performance computing, in which the investigators have a great deal of experience, to the problem of modelling such portfolios.

Much has also been written in recent years concerning large systems with deterministic constraints. This has proved to be a fruitful field for the application of large scale, and in particular, parallel computing. Conventional optimisation with non-stochastic inputs can normally be completed within reasonable run-times on PCs or similar, giving flexibility when building individual models to respond to client requests. With that approach typical PC systems permit handling of moderate sized portfolios in the region of, say, two hundred instruments. The introduction of the stochastic inputs and the associated Monte Carlo simulations makes the problem one in which issues such as parallel computation come to play a conspicuous role. Hence the problem demands the application of high performance computing [8].

In our research we investigated the numerical determination of portfolios with non-stochastic constraints combined with non-deterministic inputs, and also considered the stability of the resulting portfolios in a model which has been developed by the author and her co-workers. This required the application of parallel computing [9,10].

Normally it is assumed that the portfolio elements will behave as “Wiener processes” and have a predicted performance which is completely deterministic; that is, a particular share or bond will have a return of $x\%$. In our model we allowed not only for multivariate distributions but also for returns which have a range of values, whose size and characteristics are determined purely by empirical historical data.

3.1. Outline of the problem

The relative behaviour of securities is defined by their correlation information. Traditionally a single set of inputs for growth expectations of individual portfolio components is used in a single optimisation solution. We generalised this to allow stochastic inputs to represent the distribution of possible outcomes for the returns generated by individual portfolio components. These ‘scenarios’ are described in terms of specified rates of return and standard deviation and the statistical behaviour of these inputs can be selected from a set of standard statistical distributions.

The results of the many ensuing Monte-Carlo simulations contain much useful information that can be ex-

tracted by Principal Component Analysis. This information can be fed back into the model to improve the results. In this paper, some discussion on the stability of the results is included.

Rather than input a constant value for the expected rate of return of each input scenario, we allow a range of values with a generalised Wiener process being used to simulate each input scenario. Generating these Wiener processes forms the first stage of a four stage process for each market, and may be summarised as follows:

- Stage 1 For each simulation, generate a Wiener process for each market scenario.
- Stage 2 Solve the resulting optimisation problem at points on the “efficient frontier” [9] for each simulation portfolio) and apply Principal Component Analysis to the simulation portfolios.
- Stage 3 For those points on the efficient frontier, calculate an averaged portfolio over all the simulations (we call this the ‘mean’ portfolio).
- Stage 4 Calculate mean and volatility of rates of return for simulated scenarios and solve again to obtain an optimum market portfolio (we call this the ‘efficient’ portfolio). In this way a mean portfolio and an efficient portfolio is obtained for each market. We permit instruments to be grouped into a number of separate markets, for which historical correlation information between the securities of each such market is available. An overall solution can then be obtained
- Stage 5 A further optimisation stage (across markets), optimised by constraining the proportions of each market in the overall portfolio and by using correlation information generated by the Monte-Carlo simulations across the markets

Securities are considered as forming a single market if correlation information is available across that set of securities (see section 2.2). For each market, scenarios defining the anticipated performance of the market are specified by the user; each scenario directly influences the performance of one (and only one) sector of the market, although historical correlation information is required for the whole market (so sectors need not be independent).

During our research we investigated the numerical determination of portfolios with non-stochastic constraints combined with non-deterministic inputs, and also the stability of the resulting portfolios.

3.2. The Application of HPC

We have seen that the basic inputs to the model are securities which are grouped into markets – for each market there are a number of scenarios which define the anticipated performance of that particular market. Each market is conceived as consisting of a number of sectors; each sector may be influenced by different constraints. Hence the user can exert fine control over how each market is expected to perform. Individual markets are processed separately. An overall solution is obtained by applying the optimisation to the markets themselves.

Optimal portfolios calculated for each market are influenced by sets of constraints, which may be simple bounds or linear or non-linear constraints. The various types of constraint permit different optimum portfolios to be chosen on the efficient frontier; for example, a portfolio with minimum risk, or one with maximum return, or one with some intermediate choice of risk or return. Because we are performing many Monte Carlo simulations, we generate sets of optimal portfolios and sets of efficient frontiers. This technique has allowed us to predict portfolio behaviour that performs extremely well under most cases.

We have deployed a variant of PCA analysis to enable us to identify the behaviour of the portfolios under all the scenarios generated. In particular we have been able to use the technique to examine the components of the portfolios and their variation as the market behaviour changes. This has enabled us to identify various categories of portfolio component and to observe the role that they play in determining the behaviour of the portfolio as a whole.

The models under investigation can include differing types of securities including shares, bonds and options in (potentially geographically) diverse markets. These are defined by their correlation information. Traditionally a single set of inputs for growth expectations of individual portfolio components is used in a single optimisation solution. We generalise this to allow stochastic inputs to represent the distribution of possible outcomes for the returns generated by individual portfolio components.

3.3. Monte carlo simulations

A choice of statistical distribution for the simulated scenarios is provided. Each scenario in a simulation run is considered as a generalised Wiener process; the number of discrete time steps of each scenario is spec-

ified by the user. The stochastic values calculated for each index scenario are compared with the historic values calculated from the expected rates of return and standard deviations supplied for individual securities (in the relevant sector if related scenarios are provided for a market). The ratios so obtained are used to scale the historic rates of return and standard deviations of individual securities and optimum portfolios are then determined using the scaled values.

For example, consider a market where the given scenarios are the securities themselves. The (historic) correlation matrix of the market securities is also the correlation matrix for the scenarios. The covariance matrix is formed from the correlation matrix by scaling each row and column by the standard deviation of each security. We calculate a set of scaling factors for the first simulation as follows:

- (i) each Wiener process for a scenario/security price S over time t is of the form

$$\frac{\Delta S}{S} = \mu \Delta t + \sigma \varepsilon \sqrt{\Delta t},$$

where μ is the expected growth rate of the security per unit time, σ is the standard deviation of the security price and ε is a random number from a multivariate normal distribution (since the scenarios are related).

- (ii) Having generated all the Wiener processes for the simulation, we can calculate the perturbed values (growth rate and standard deviation) which result for each scenario.
- (iii) We then scale the historic correlation matrix by these standard deviations to produce the perturbed covariance matrix for the simulation and solve the appropriate optimisation problems.
- (iv) This process is then repeated for each simulation.

A general-purpose optimisation routine is used to determine an optimum portfolio, based upon the SQP (Sequential Quadratic Programming) method [12]. For each market, constraints may be supplied in three forms: simple bounds, linear constraints of the form $Ax \geq B_l$ and $Ax \leq B_u$, and smooth non-linear constraints (defined by a vector of constraints or optionally its Jacobian). (A is an $m * n$ matrix of coefficients, X the sum of the proportions of each security, B_{us} and B_l represent the upper and lower bounds of the inequalities, respectively).

Constraints will often need to be included to produce a balanced portfolio across the market. Without some constraints on the proportion of certain securities, the

portfolio can become dominated by a very few securities (especially for the case of maximum return). Constraints of this form can result in discontinuities of the efficient frontier. It is also usual to constrain the sum of the proportions of each security in the portfolio to one, i.e. $\sum_j X_j = 1$.

For each simulation run, five types of optimisation can be solved: the optimisations for minimum risk and for maximum return (which represent the end points of the efficient frontier); optimisations of intermediate risk and intermediate return; or other intermediate points on the efficient frontier defined using the parameterised objective function $\sum_i \sum_j X_i X_j C_{ij} - \lambda(\sum_j X_j E_j)$. This Monte Carlo simulation is repeated a specified number of times. Sufficient simulation runs should be performed to ensure the stochastic results have converged. Variance reduction techniques are used to improve convergence. A further set of optimisations is then solved using the observed means and standard deviations; analogous results are also calculated using the user-supplied values. Principal Component Analysis is then applied to the recorded portfolios. More details are given in [10].

3.4. Parallelisation of the system

There are two ways in which we can take advantage of the parallel aspects of this system.

First, the numerical algorithms used have substantial vector and matrix operations inherent in the linear algebra of the calculations. By utilising computationally efficient BLAS (Basic Linear Algebra Subroutines) in the numerical algorithms, we can ensure advantage is taken of the architecture of the hardware platform used (for example, matrix operations use optimum cache sizes and vector operations use pipelining efficiently). The effect of using such machine specific BLAS on the system is an approximate three-fold increase in computational speed on Intel PCs.

Second, the Monte-Carlo simulations are independent of one another and therefore can be executed in parallel without synchronisation. Currently available library software is not thread-safe for some of the numerical algorithms we require. A solution is to implement the system using PVM (Parallel Virtual Machine, see Geist et al. [11]), which divides the computation into separate system processes that communicate by message passing. Separate library routine calls cannot then interfere with one another since they each have their own address space. The danger of the PVM

approach is that the communication costs can become significant if message passing is too frequent.

Our application is ideally suited to the PVM approach. A master process controls the generation of the Wiener processes and the perturbed scenario inputs, it then spawns the appropriate number of slave processes. Each slave process is responsible for performing a number of simulations. The master process passes the perturbed scenario inputs and other historic data necessary for the optimisations to each slave process (only one message per slave process is needed for this). It then waits for the slaves to complete their computations and pass back the simulation portfolios (again only one message per slave is needed for this). The master process can then continue and analyse the results. Thus frequent message passing is avoided and synchronisation is kept to a minimum.

Whilst originally we believed that parallel computing would provide the only way forward for these calculations, recent development in PC-level architectures and performance level has meant that large-scale high power PCs can be coupled to offer an effective solution for some classes of problems.

Average computation times for each time-point of the set of tests described below (2000 simulations over 59 securities for one 'point' on the efficient frontier, but excluding the PCA analysis) were reduced from 13.2 minutes for the single process solution to 9.1 minutes on a twin processor Intel PC. This is a significant saving bearing in mind that the Wiener process initialisation time is not dependent on the number of points on the efficient frontier, hence multiple solutions on the efficient frontier will show more favourable savings in computational costs for solution.

3.5. Principal component analysis

We want to develop tools for understanding the composition and stability of portfolios under various assumptions about how errors in risk and return are measured.

Principal components are defined as linear combinations of the individual variables appearing in the portfolio. Thus for the securities X_k , the first principal component is of the form $Y_1 = \sum_k a_{1k} X_k$ with a_{1k} chosen to maximise the sample variance of Y_1 over all of the simulation portfolios subject to $a_1' a_1 = 1$. Further principal components Y_j are defined similarly, but are also required to be orthogonal to each of the previous principal components, thus $a_j' a_i = 0$ for $i < j$, [13].

Principal component analysis is essentially used to reduce the dimensionality of a given problem. The principal components are chosen so that the first component has the greatest impact on the analysis and each subsequent component has a decreasing impact. We are aiming to project the original data space onto a lower dimensional space, whilst not losing too much information. To do this we need to be able to measure the errors we are introducing by selecting the first $r (\leq n)$ principal components.

Applying principal component analysis to the covariance matrix (of the portfolio simulations) we see that the total variance of the original system is the sum of the diagonal elements $\omega_1^2 + \omega_2^2 + \dots + \omega_n^2$ or the trace of the matrix, which we denote as s_n^2 . The variance of the system that is 'explained' by the first r principal components is given by S_r^2 . A traditional measure of the 'goodness-of-fit' of the subspace projection defined by the first r principal components is the proportion S_r^2/S_n^2 . This measure turns out not to be well-suited to our application, hence we have designed three alternative measures based on Euclidean distance, the error in risk, and the error in return. These measures are defined either as means (ERR_1 , ERR_2 and ERR_3) or in least-squares form (ERR_4 , ERR_5 and ERR_6) over the simulations [10].

4. Financial risk assessment for portfolio management

This project aimed to build an understanding of the stability of portfolios using non-linear optimisation methods with stochastic inputs to develop the understanding of appropriate stability criteria in conjunction with a large international bank as "end users". Real market data (March 1997 – March 2001) was used to address the stability of solutions of these problems to variations in the holdings. This is important where the holder wishes to understand the sensitivity of the portfolio holdings to minor variations in the holdings. These may correspond to proportions of holdings or to variations in possible scenarios. It will also develop further the understanding of real issues in the solution of this class of non-linear optimisation problems on HPC systems in the context of real economic systems.

Portfolio managers would like to choose portfolios containing a mix of securities so that risk is controlled. The future values of securities are uncertain and these values will change relative to each other depending on future events. A UK portfolio manager might want to

choose a portfolio which will follow the future market independent of the choice of when, if ever, the UK joins the EURO. There are thus a number of scenarios which a manager might postulate, each giving a different value and variance of that value independent of the actual scenario for any given security. The task of the manager is to suggest a portfolio which independent of the actual scenario, will meet some given criterion, e.g. minimum risk, maximum return or some intermediate state.

The above process was tested with equity data based on the FTSE 100 Index on the London Stock Exchange. Even this relatively small data set has brought to light a number of detailed issues relating to the sensitivity which have to be investigated if the above process is to be applied more widely and to much larger markets. The use of Principal Component Analysis can highlight which predictions are important (and which are largely irrelevant) to the portfolio optimisation. We have thus created a numerical approach for the inclusion of predictions of expectations of return in portfolios, along with linear and non-linear constraints.

4.1. Alternative definitions of risk

The traditional definition of risk is $\sum_i \sum_j X_i X_j C_{ij}$, where X_j is the proportion of security j in the portfolio, and $C_{ij} = \rho_{ij} \sigma_i \sigma_j$ is the covariance between securities i and j . A portfolio manager may well compare performance with some benchmark portfolio (with components B_j), hence it is natural to modify the above definition, so risk is defined as: $\sum_i \sum_j C_{ij} (X_i - B_i)(X_j - B_j)$. Another use for this type of definition is where the benchmark represents an existing portfolio, hence the use of this form of definition gives a bias towards the existing portfolio.

The need to determine various measures of risk and return arises because of the non-deterministic inputs that our approach allows. The tool that we use to achieve this is again Principal Component Analysis.

5. The analysis of bond portfolios

In order to verify our approach we constructed a series of portfolios. The remaining sections of this paper briefly discuss the results obtained.

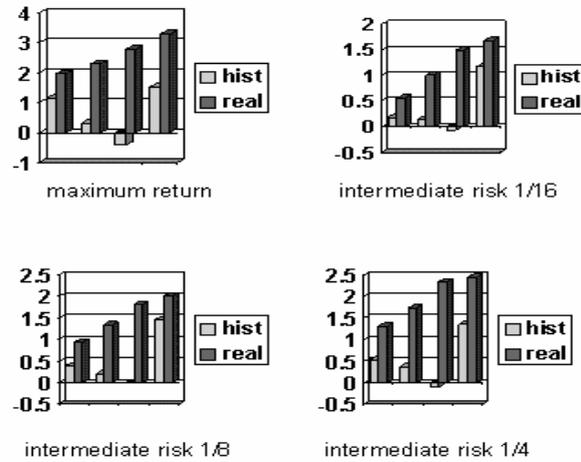


Fig. 1. Bond performance with risk relative to benchmark.

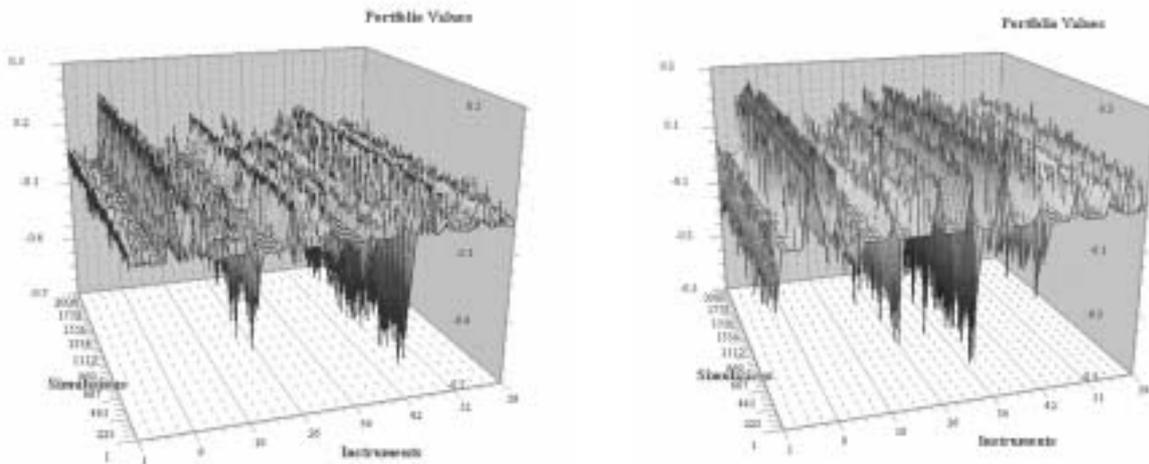


Fig. 2. Typical sets of simulation portfolios.

5.1. The portfolio data

Tests are based upon bond index data collected over a period of approximately five years. The data is close of day prices for bond indexes with five different sets of maturities (from short to long dated) across internationally diverse markets (e.g. USA, UK, Japan, Australasia and Europe), together with deposit rates and currency exchange rates. The period covered by the data runs from 1993 to March 1998. The data was initially pre-processed to remove spurious records from the time series and to interpolate any missing values. Apart from this little else was done in the way of preparation of the data and we have no reason to believe that it is other than a representative set of data.

Figure 1 shows the absolute performance of the selected portfolios of bond instruments, but with risk defined relative to a typical benchmark (with maturities weighted towards medium dated bonds).

The types of optimisations solved are maximum return and three intermediate risks (1/16, 1/8, 1/4 of maximum risk). Also an upper limit of 5% is imposed on the proportion of an individual instrument in the selected portfolio. The quarters used are 1 December 1996 to 1 December 1997, with predictions based on historical observation of each preceding quarter or on actual growth over the quarter. Both sets of results show positive performances, but the value of good predictions is clear.

Table 1

Risk of efficient portfolio	7.90E-4	7.43RE-04	6.87E-04
Risk of mean portfolio	7.04E-04	6.97E-04	6.78E-04
Return of efficient portfolio	0.43641	0.42894	0.42299
Return of mean portfolio	0.42797	0.42388	0.42108

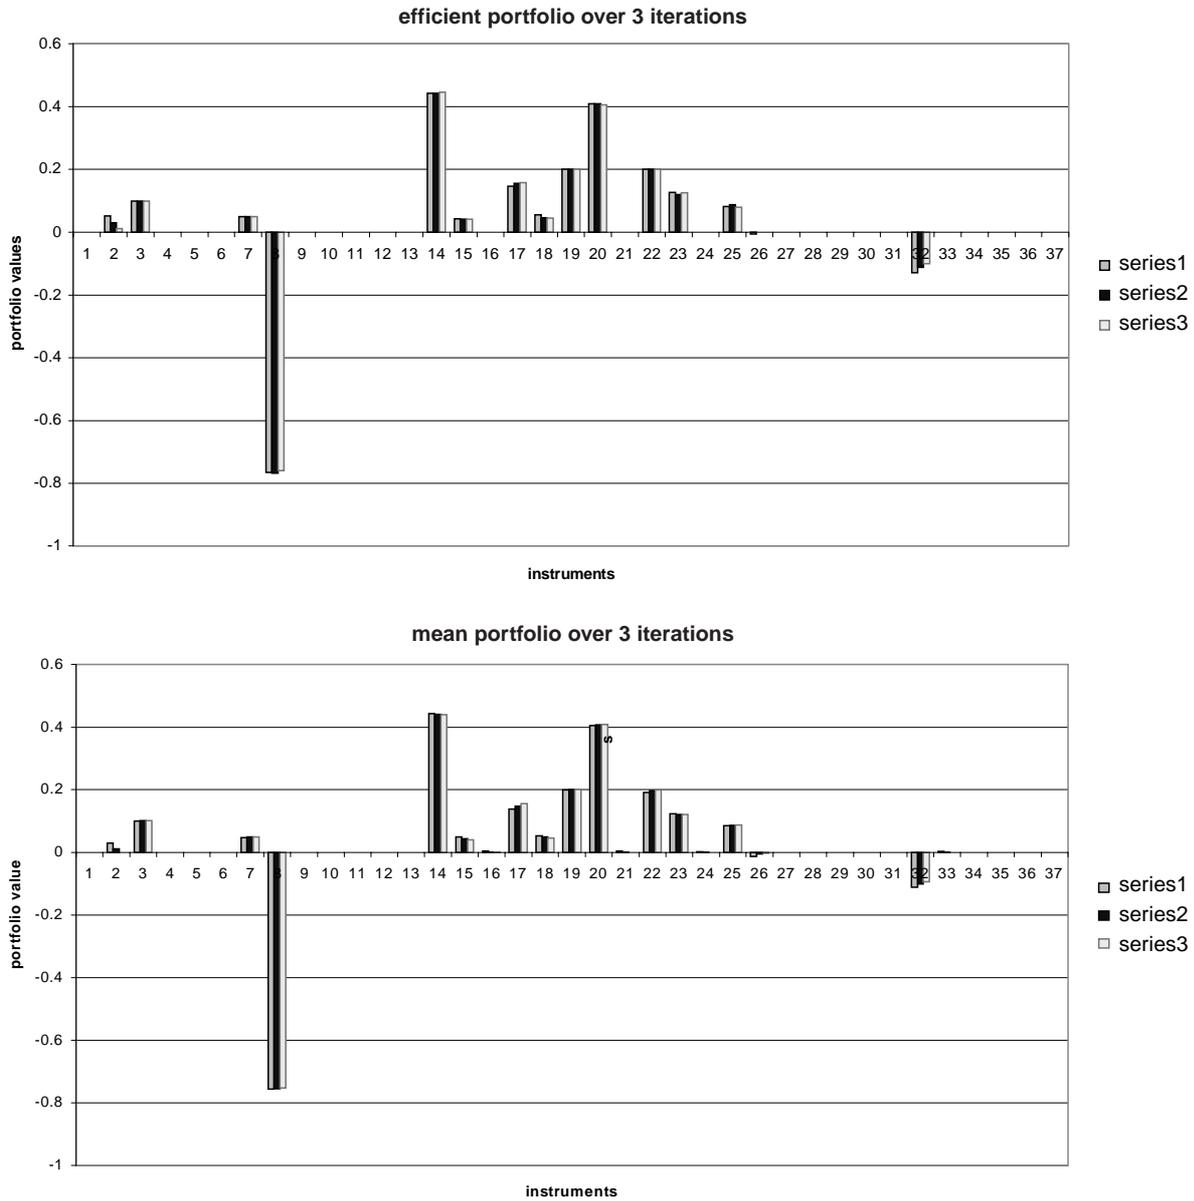


Fig. 3.

5.2. Sensitivity of portfolio to time-period of the historic data

The periods, from which data is taken, that are used in this process are very long. This has necessitated

the application of a model to assign significance to the data points used in the calculation of the portfolio. The longer the time series used, the less relevant to more recent portfolios are the data points obtained from the beginning of the series. In order to deal with this prob-

lem we have looked at various methods of assigning significance to data points. The one we have chosen is the scheme due to Joubert and Rogers [14], in which a geometric weighting is applied.

It is clear that the choice of weighting scheme must not be such as to substantially reduce the significance of recent major events or to completely obviate those which may have occurred during the period covered. We have tested this and other models and are satisfied that the model chosen and the detailed weighting scheme applied reflects this need.

Constraints have also been applied to our modelling system. A less risky portfolio can be achieved through diversity. We have therefore chosen to limit bond holdings to a minimum of 0% and a maximum of 20% of the portfolio. Deposit holdings are bounded by ± 1 (so “hedging” is allowed), but additional (linear) constraints require that (positive) deposit holdings must be matched by bond holdings and the sum of the deposit holdings is zero to limit currency speculation and to prevent gearing.

5.3. *Principal component analysis with predictions from historic data*

The dataset used for the principal component analysis is that specified above. There are 59 instruments selected, comprising 10 countries (Australia, Germany, Canada, Denmark, Japan, New Zealand, Sweden, Switzerland, UK, USA) with 10 deposit instruments and 49 bond index instruments. Points in time have been chosen at monthly intervals from March 97 to March 98. For each time-point, geometric weighting has been applied to the dataset to obtain the historic returns, standard deviations and covariance matrix for the 59 instruments. Significant variation in behaviour of the instruments can be observed over this period.

The stochastic inputs for these tests have been set to the historic returns and standard deviations (with geometric weighting) calculated at each of the time-points. Figure 2 shows two typical sets of simulation portfolios. Essentially there are three types of instruments present in these portfolios:- the ‘popular’ securities which are present in most portfolios; the ‘unpopular’ securities which are usually absent; and other ‘volatile’ securities which are present or absent because of the volatility built into the model by the standard deviations of the stochastic inputs. The first plot shows a limited number of popular and unpopular securities, with a large proportion of mid-height volatile securities; the second plot shows the portfolios concentrated into the more

popular securities. We will see these two plots represent the most volatile and least volatile respectively of the time-points we have considered; this will be confirmed by the principal component analysis.

Each simulation produces an efficient portfolio determined by the perturbed input values. These efficient portfolios have a limited number of instruments in each portfolio, but not necessarily the same volatile securities present in each such portfolio. On the other hand, the mean portfolio constructed by taking means of each instrument over all the simulations, has many more instruments present – the popular securities as well as all of the volatile securities. The definition of the mean portfolio ensures it is feasible (with respect to the bounds and linear constraints), but it is not necessarily efficient. Figure 3 shows three typical simulation portfolios; different portfolios clearly show different securities present and volatility in the proportions present. The second plot shows the corresponding mean portfolio, which has many more securities present than the corresponding efficient portfolio (the conventional Markowitz optimal portfolio generated by the means of the scenario inputs).

5.4. *Refinement of the mean portfolio*

Having performed the simulations, the mean rates of return and the mean of the standard deviations observed for the simulated scenarios (the Wiener processes) are calculated. A further optimisation problem is then solved using the observed means. The resulting portfolio is of course ‘efficient’. We also calculate the mean portfolio from the simulation portfolios; the mean portfolio is not, in general, efficient. The efficient portfolio will be mainly composed of the ‘popular’ securities, while the mean portfolio will also contain many of the ‘volatile’ securities.

Principal Component Analysis will have suggested that a number of the original instruments can be safely ignored. The ‘unpopular’ securities will have zero or very small holdings in the mean portfolio. Clearly, a first step in removing the unwanted securities from the mean portfolio is to remove the unpopular securities with values below some small threshold (uneconomic holdings).

The effect of removing such holdings and adding constraints as above appears to be to bring the iterated mean portfolio closer to the iterated efficient portfolio in risk/return space. Table 1 shows a typical example of this behaviour by repeating the simulation runs so that three sets of portfolios are available. Subtle variations

in the efficient portfolios occur; analogous subtle variations can also be observed in the impact in risk/return space, where the risks and returns of the efficient and mean portfolios can be seen to converge rapidly

6. Conclusions

We have created a numerical approach for the inclusion of predictions of expectation of return in portfolios, along with linear and non-linear constraints. We believe that we have been able to demonstrate that the method as described is capable of producing reliable portfolios, which are stable against perturbation indicating that the selection is at least close to optimal. The timing of our tests indicates that in order to obtain sufficient convergence it may be necessary to employ networked high-end PC technology, or other more powerful parallel architectures, for the solution of large problems. The technology has demonstrated its capability to include scenarios as inputs, dependent upon expected outcomes, and thus enhances the armoury of the portfolio manager.

Acknowledgements

We would like to acknowledge the invaluable input from Paul Conyers and Dominic Kini of ABN Amro and from Bernard Fromson formerly of ABN Amro. We also thank our former colleagues at the London Parallel Applications Centre, Queen Mary College and at Kings College. We would also like to thank NAg Ltd. for their assistance, in particular David Sayers. This work has been supported by the UK Engineering and Physical Sciences Research Council under grant number GR/M84282.

References

- [1] H.M. Liddell and D. Parkinson, The measurement of performance on a highly parallel processor, *Trans. on Computers C-32* **32** (1982).
- [2] H.M. Liddell and G.S.J. Bowgen, The DAP Subroutine Library, *Comput Physics Communicatuions* **26**(311) (1982).
- [3] H.M. Liddell and C.H. Lai, Finite elements using long vectors of the DAP, *Parallel Computing* **8** (1988), 351–361.
- [4] Organic Low Dimensional Structures (OLDS) ESPRIT BR 3200 30 months from 1/7/8 (collaborative with Universities of Edinburgh, Genoa, Max-Planck-Institute, Tubingen and Ernst-Leitz-Wetzlar GmbH)
- [5] Multidimensional and Multimodal medical imaging and display, collaborative IED project SERC GR/F34909, 1989–1991.
- [6] H.M. Liddell, D. Parkinson, P. Undrill et al., Integrated presentation of 3D imagery and anatomical atlases, using a parallel processing system, *SPIE Vol. 1653, Image Capture, Formatting and Display*, 1992.
- [7] H.M. Markowitz, Portfolio Selection, *Journal of Finance* **7**(1) (1952), 778.
- [8] S.A. Zenios, *Financial Optimization*, Cambridge University Press, 1993.
- [9] G.S. Hodgson, P. Dzwig, H.M. Liddell and D. Parkinson, Application of HPC to Medium-size Stochastic Systems with Non-linear Constraints in Finance, *Lecture Notes in Computer Science* **1401** (1998), 411–418.
- [10] G.S. Hodgson, H.M. Liddell, P. Dzwig and D. Parkinson, Medium-size Stochastic Systems and Principal Component Analysis in Portfolio Optimisation, paper presented to the European Working Group on Financial Modelling Meeting, New York, November 2000.
- [11] P. Gill, W. Murray and M. Wright, *Practical Optimisation*, Academic Press, 1981.
- [12] E. Geist et al., *A users guide and tutorial for networked parallel computing*, 1994.
- [13] W.J. Krzanowski, *Principles of Multivariate Analysis*, Oxford University Press, 1990.
- [14] A.W. Joubert and L.C.G. Rogers, Ticking over in real time in Quantitative and Computational Finance, UNICOM 1994, pp. 144.



Hindawi

Submit your manuscripts at
<http://www.hindawi.com>

