Financial Computations on Clusters Using Web Services

Shirish Chinchalkar, Thomas F. Coleman, and Peter Mansfield

Cornell Theory Center, Cornell University,
55 Broad Street, Third Floor, New York, NY 10004, USA
{shirish, coleman, peterm}@tc.cornell.edu

Abstract. The pricing of a portfolio of financial instruments is a common and important computational problem in financial engineering. In addition to pricing, a portfolio or risk manager may be interested in determining an effective hedging strategy, computing the value at risk, or valuing the portfolio under several different scenarios. Because of the size of many practical portfolios and the complexity of modern financial instruments the computing time to solve these problems can be several hours. We demonstrate a powerful and practical method for solving these problems on clusters using web services.

1 Introduction

The problems of financial engineering, and more generally computational finance, represent an important class of computationally intensive problems arising in industry. Many of the problems are portfolio problems. Examples include: determine the fair value of a portfolio (of financial instruments), compute an effective hedging strategy, calculate the value-at-risk, and determine an optimal rebalance of the portfolio. Because of the size of many practical portfolios, and the complexity of modern financial instruments, the computing time to solve these problems can be several hours.

Financial engineering becomes even more challenging as future ‘scenarios’ are considered. For example, hedge fund managers must peer into the future. How will the value of my portfolio of convertibles change going forward if interest rates climb but the underlying declines, and volatility increases? If the risk of default of a corporate bond issuer rises sharply over the next few years, how will my portfolio valuation be impacted? Can I visualize some of these dependencies and relationships evolving over the next few years? Within a range of parameter fluctuations, what is the worst case scenario?

Clearly such “what if” questions can help a fund manager decide today on portfolio adjustments and hedging possibilities. However, peering into the future can be very expensive. Even “modest” futuristic questions can result in many hours of computing time on powerful workstations. The obvious alternative to waiting hours (possibly only to discover that a parameter has been mis-specified), is to move the entire portfolio system to a costly supercomputer. This

is a cumbersome, inefficient, and “user unfriendly” approach. However, there is good news: most of these practical problems represent loosely-coupled computations and can be solved efficiently on a cluster of processors in a master-worker framework.

We have been investigating the design of effective parallel approaches to the problems of financial engineering, and computational finance, on clusters of servers using web services. Our particular approach is to represent the portfolio in Excel with the back-end computing needs satisfied by a cluster of industry standard processors running in web services mode. The user environment we have used is Microsoft’s .NET.

2 Introduction to Web Services

A web service is a piece of functionality, such as a method or a function call, exposed through a web interface ([1]). Any client on the internet can use this functionality by sending a text message encoded in XML to a server, which hosts this functionality. The server sends the response back to the client through another XML message. For example, a web service could compute the price of an option given the strike, the stock price, volatility, and interest rate. Any application over the internet could invoke this web service whenever it needs the price of such an option. There are several advantages in using web services to perform computations:

1. XML and HTTP are industry standards. So, we can write a web service in Java on Linux and invoke it from a Windows application written in C# and vice a versa.
2. Using Microsoft’s .NET technology, we can invoke web services from office applications such as Microsoft Excel. This feature is especially useful in the financial industry, since a lot of end-user data is stored in Excel spreadsheets.
3. No special-purpose hardware is required for running web services. Even different types of computers in different locations can be used together as a web services cluster.
4. Since the web service resides only on the web server(s), the client software does not need to be updated every time the web service is modified. (However, if the interface changes, the client will need to be updated).
5. The web service code never leaves the server, so proprietary code can be protected.
6. Web services can be accessed from anywhere. No special purpose interface is necessary. Even a hand-held device over a wireless network and the internet can access web services.

In the context of large-scale financial computations, there are some limitations to the utilization of web services as well:

1. There is no built-in mechanism to communicate with other web services. This limits the use to loosely coupled applications.

2. The results of any computation performed during a single web service call can only be sent to the client at the end of that web service call. Thus, there is no mechanism for sending partial results while the computation is going on, without using another technology such as MPI.
3. Since messages are sent using a text format over the internet, this is not a viable computational technique for “short” computations involving a lot of data to be communicated.

Writing a web service is no different from writing a function or a method that performs the same computation. Other than a few declarative statements, there is no difference between a web service and an ordinary function. There is no reference to message passing, converting data into XML, and so on. These details are hidden from the programmer.

Similarly, invoking the web service from a client program is no different from making a function call within the same process. The relatively minor difference involves specifying in the main program the location of the web service.

3 Cluster Computing Using Web Services

A typical portfolio manager could have a large portfolio of complex instruments. These instruments may have to be priced every day. Often, several future scenarios of the stock market or interest rates may have to be simulated and the instruments may have to be priced in each scenario. Clearly, a lot of computing power is necessary. If the instruments can be priced independently of one another, we can make use of web services to perform this computation.

The entire computation can be partitioned into several tasks. Each task can consist of the pricing of a single instrument. We can have a separate web service to price each instrument. The client then simply needs to invoke the appropriate web service for each instrument. We can use other models of computation as well. For instance, in case of Monte Carlo simulation, we could split the simulations among the processors.

Figure 1 shows the overall organization of our architecture. The front-end is a typical laptop or a desktop running Excel. Data related to the portfolio is available in an Excel spreadsheet. This front-end is connected over internet or a LAN to a cluster of nodes, each of which runs a web server. When a large computation is to be performed, it is broken into smaller tasks by the Excel front-end. Each task is then shipped to an individual node which works on it independent of the other nodes. The nodes send results back to Excel, which is used to view results.

What are the advantages in using web services for cluster computing instead of MPI? There are several. First, MPI is suitable for tightly coupled problems which require significant communication between participating processes. For the financial problem described above, no inter-process communication is needed because the problem is loosely coupled. So, web services prove adequate for this task. Second, web services offer a high degree of fault tolerance. Processors can be added or taken down at any time; if a task is aborted, or if a processor is too

Fig. 1. Overview of computing architecture

slow, that task can be redirected to another processor. Third, in a commercial setting, the computational service (such as pricing of a financial instrument) is provided by one organization, whereas the user could be another organization or an individual. In such a setting, the user need not (and indeed should not) know the internal details of the cluster set-up, nor should he have to install any special-purpose software to use the cluster. This is possible with a web services interface.

4 Load Balancing

Given a set of tasks, we can distribute them across a .NET web services cluster in two different ways. We could send all the tasks at once to the main cluster node which uses Network Load Balancing (NLB) to distribute the tasks. However, the NLB monitors network traffic and considers those nodes that are actively communicating as busy and those that are not as idle. This is reasonable in transaction processing applications where each task can be processed quickly and the number of tasks is very large. For the problems we are interested in, we have a relatively small number of tasks, each of which takes seconds, minutes, or hours of computational time. For such problems, a node which is not sending messages might be busy doing computation and might be wrongly classified as idle by NLB. In such cases, the following approach is more suitable: We first assign tasks, one to each processor. When any task finishes, the next task is sent to the node which finished that task. This algorithm works well in practice provided there is only one application running on the cluster. If multiple applications need to run simultaneously, a centralized manager is necessary.

The load balancing mechanism described above can be implemented as a generic class, provided all data for all tasks is known before any task is executed. Fortunately, most finance applications that involve pricing portfolios of

instruments fall in this category. By using a load balancing template, we can remove from the user application, most of the low-level “plumbing” related to multi-threaded operation. This makes applications significantly easier to program and promotes code reuse.

All code related to invoking the web service asynchronously on a multi-node cluster, determining free nodes, using locks for multi-threaded operation, sending inputs, receiving results, and generating timing and speedup information is handled by the load balancing class. If the user wishes to process results as they are returned, he will need to write an application-specific callback. Again, this callback does not involve any lower-level message passing related code.

We have incorporated several useful features in the load balancing class. First, we can use multiple clusters simultaneously. These clusters can be in different locations, use different operating systems, and have different web servers. Second, if one of the nodes is down, tasks sent to that node get aborted. Once this is detected, that node is marked as down and those tasks are re-routed to one of the other nodes. Nodes marked down do not participate in further computation. Third, if, for any reason, one of the nodes is too slow, there is a provision for automatically timing out a task and re-routing it to another node. Such features are commonly required in all parallel programs and in the web services setting, they can be hidden from the application programmer very easily.

5 An Example

Cluster computing using web services as outlined above can be used to price portfolios comprising different types of instruments such as risky bonds, convertible bonds, and exotic options. We give an example which involves pricing a portfolio of callable bonds with Excel as a front-end and using multiple clusters.

A typical corporate bond has a face value, a fixed coupon, and a maturity date. Such a bond pays a fixed amount of interest semi-annually until maturity. At maturity, the face value or principal is returned[3]. A callable bond has an additional feature - the bond may be ‘called back’ by the issuing company by offering the bond holder or the investor an amount equal to the face value of the bond. This buy-back can be made on any of the coupon payment dates. Whether it is optimal for the issuing company to call in the bond or not depends on the prevailing interest rates and predictions of future interest rates. For example, if interest rates drop, it may be in the best interests of the issuing company to buy back the bond. If interest rates are high, the issuing company is unlikely to call in the bond. This presents two problems - first, future interest rates must be simulated, and second, the decision to buy the bond or not should be made at each coupon date, depending on the prevailing interest rate and the prediction of future interest rates.

For this work, we have used the Vasicek model for simulating interest rates. In this model, changes in interest rates are given by the formula

$$dr = a(\bar{r} - r)dt + \sigma dW \quad (1)$$

where dr is the change in the interest rate in a small time interval, dt , a is the mean reversion rate, $\bar{r}$ is the mean reversion level, and σ is the volatility. dW is a small increment of the Brownian motion, W (see [4] for more details). Given an initial interest rate, r_0 , we can easily simulate future interest rates using the above equation. For valuation of callable bonds and the calculation of greeks (see below), we need several tens of thousands of simulations.

Optimal exercise along each interest rate path is determined using the Least Squares Monte Carlo algorithm, which involves the solution of a linear regression problem at each coupon date and discounting of cashflows along each interest rate path. Details of this algorithm can be found in Longstaff and Schwartz[2].

We illustrate a few additional computations for a single bond. They can be extended to a portfolio quite easily. Along with the price of the bond, we also want the bond's 'greeks'; for example bond delta and bond gamma. Delta is the first derivative of the bond price with respect to the initial interest rate ($\partial B/\partial r$) and gamma is the second derivative of the bond price with respect to the initial interest rate ($\partial^2 B/\partial r^2$), where B is the price of the bond. In this work, we have computed them using finite differences as follows

$$\Delta = \left. \frac{\partial B}{\partial r} \right|_{r=r_0} \approx \frac{B(r_0 + dr) - B(r_0 - dr)}{2dr} \quad (2)$$

$$\Gamma = \left. \frac{\partial^2 B}{\partial r^2} \right|_{r=r_0} \approx \frac{B(r_0 + dr) - 2B(r_0) + B(r_0 - dr)}{dr^2} \quad (3)$$

The above calculations require the pricing of the bond at two additional interest rates, $r_0 + dr$ and $r_0 - dr$. For all three pricing runs, we use the same set of random numbers to generate the interest rate paths (see [4]).

Once the greeks are computed, we can approximate the variation of the bond price by the following quadratic

$$B(r) \approx B(r_0) + \Delta(r - r_0) + \frac{1}{2}\Gamma(r - r_0)^2 \quad (4)$$

A risk manager would be interested in knowing how much loss this bond is likely to make, say, 1 month from now. This can be characterized by two metrics: Value at Risk (VaR) and Conditional Value at Risk (CVaR). These can be computed from the above approximation by another Monte Carlo simulation. For an introduction to VaR and CVaR see [4].

The portfolio price, V , is simply a linear function of the individual bond prices

$$V = \sum_1^n w_i B_i \quad (5)$$

where the portfolio consists of n bonds, with w_i number of bonds of type i . The greeks can be computed analogously, and VaR and CVaR can be determined easily once the greeks are known.

Figure 2 shows the Excel front-end developed for this example. This interface can be used to view bond computing activity, cluster utilization and efficiency, a plot of portfolio price versus interest rate, and portfolio price, Value at Risk (VaR), Conditional Value at Risk (CVaR), and portfolio delta and gamma.

Fig. 2. Callable bond pricing on multiple clusters using Excel

In our example, the web service computes the bond price and bond greeks, whereas the Excel front-end computes the portfolio price, greeks, VaR, and CVaR. Our experiments with portfolios of as few as 200 instruments show that on 20 processors in 3 clusters at 3 different locations, we get speedups of more than 15. On a 64 processor Dell 2450 cluster consisting of 32 dual Pentium III 900 MHz processors, we have obtained speedups in excess of 60 relative to computation on a single processor for portfolios consisting of 2000 instruments, reducing 9 hours of computation on a single processor to about 9 minutes on the cluster.

6 Conclusion

Parallel computing, used to speed up a compute-intensive computation, has been under development, and in use by researchers and specialists, for over a dozen

years. Because a parallel computing environment is typically an isolated and impoverished one (not to mention very costly!), general industry has been slow to adopt parallel computing technology. Recent web services developments suggest that this situation is now improving, especially for certain application classes, such as portfolio modeling and quantitative analysis in finance. The work we have described here illustrates that a powerful analytic tool can be designed using web services technology to meet some of the computational challenges in computational finance and financial engineering.

Acknowledgements

This research was conducted using resources of the Cornell Theory Center, which is supported by Cornell University, New York State, and members of the Corporate Partnership Program. Thanks to Dave Lifka and his systems group, Yuying Li, and Cristina Patron for their help.

References

1. A. Banerjee, et. al. *C# web services - building web services with .NET remoting and ASP.NET*. Wrox Press Ltd., Birmingham, UK, 2001.
2. F. Longstaff and E. Schwartz. Valuing American Options by Simulation: A Simple Least Squares Approach. *The Review of Financial Studies*, 14:113-147, 2001.
3. Z. Bodie, A. Kane, and A.J. Markus, *Investments*. McGraw Hill. 2001.
4. P. Wilmott. *Paul Wilmott on Quantitative Finance, Volume 2*. John Wiley and Sons, New York, 2000.