

HIGH PERFORMANCE SPREADSHEET SIMULATION ON A DESKTOP GRID

Juta Pichitlamken

Department of Industrial Engineering
Faculty of Engineering, Kasetsart University
Bangkok, 10900, THAILAND

Supasit Kajkamhaeng
Putchong Uthayopas

High Performance Computing and Networking Center
Faculty of Engineering, Kasetsart University
Bangkok, 10900, THAILAND

ABSTRACT

We present a proof-of-concept prototype for high performance spreadsheet simulation called S3. Our goal is to provide a user-friendly, yet computationally powerful simulation environment for end users. Our approach is to add power of parallel computing on Windows-based desktop grid into popular Excel models. We show that, by using standard Web Services and Service-Oriented Architecture (SOA), one can build a fast and efficient system on a desktop grid for simulation. The complexity of parallelism can be hidden from users through a well-defined computation template. This work also demonstrates that a massive computing power can be harvested by linking off-the-shelf office PCs into a desktop grid for simulation. The experimental results show that the prototype system is highly scalable. In the best case, the execution time can be reduced 13.6 times using 16 desktop PCs; the simulation time is dramatically reduced from 200 minutes to 14 minutes.

1 INTRODUCTION

Spreadsheet is a widely used computer application because of its versatility and ease of use in calculation or modeling a problem. Currently, the most popular spreadsheet program is Microsoft Excel on Windows since Windows shared more than 90% of the client operating system market (as of 2004) (Legard 2004). In fact, many textbooks use Excel as a calculation tool, e.g., Ragsdale (2004) and Stevenson and Ozgur (2007) which prove that Excel is sufficiently applicable to many types of analysis.

What-if analysis is used to assess how sensitive outputs are to changes in input values; e.g., how much total cost would increase if the project is delayed by x days. The analysis is inefficient if one parameter is changed at a time. Stochastic simulation is a widely-used technique for sensitivity analysis as it allows users to explore many more scenarios automatically. Applications of spreadsheet simulation include financial risk analysis (e.g., Paisittanand

and Olson 2006) and operational risk analysis (e.g., Shariff et al. 2006). Seila (2006) and Seila, Ceric, and Tadikamalla (2003) provide a comprehensive introduction to spreadsheet simulation.

When a problem size is large, computation power of a single computer might not be enough; users may have to wait for a few hours to see simulation results. Thus, accelerating computing speed for Excel in a transparent way is beneficial since it allows users to quickly “play” with the problem and gain insights. One approach to solve this problem is to enhance Excel using a parallel/distributed computing. The idea is to break a large simulation problem into many small sub-problems that can be executed concurrently on multiple computers. Excel can be customized through user-defined Visual Basic for Applications (VBA) macros or add-ins that allow users to include their own functions; therefore, it is possible to modify Excel to seamlessly use parallel computing to speed up its execution of a large simulation problem.

In this paper, we build a spreadsheet simulation system (S3) that is easy to use and computationally fast by utilizing power of parallel computing on a Windows-based desktop grid. When simulation runs are faster, an analyst can execute a simulation model under multiple sets of decision variables, or they may want to find the best set of decision variables for the system of interest (*optimization via simulation*). S3 allows users to execute a spreadsheet simulation model under multiple set of integer-valued decision variables by specifying their upper and lower bounds. For each combination of decision variables, S3 returns sample means and standard errors. Users can then explore these outputs by sorting them and examine which set of decision variables give the top $x\%$ performance.

We consider two key design problems:

1. To design spreadsheet *simulation infrastructure* which is inexpensive and easy to maintain.
2. To design an *Excel user's interface* in such a way that users can slightly modify their existing mod-

els or build their models without being aware of parallelism and infrastructure used. The parallel computation should be hidden from users as much as possible to make it user-friendly.

We address the first issue with a *desktop grid* which are built from commodity PCs, readily available in most organizations. Grid computing focuses on large-scale resource sharing. Grid architecture specifies protocols that define “the basic mechanisms by which users and resources negotiate, establish, manage, and exploit sharing relationships” (Foster, Kesselman, and Tuecke 2001).

Desktop grid computing generally consists of clients, workers, a manager, and servers. A client submits jobs that are executed by workers. A manager is responsible for job scheduling and resource management. Servers are used for data storage. A well-known example of desktop grids is SETI@home which is based on BOINC (Anderson 2004). Desktop grids are appealing for they allow intensive computation to be performed at low cost. However, the main challenge is stability and security. Given that our desktop grid consists of PCs within the same organization, the security issue may not be too prohibitive.

The second issue is more challenging. Typically, Excel evaluates formulas and display the results as values in the cells that contain the formulas (Ecklund 2007). A `.xls` file is an Excel *Workbook* which consists of *Worksheets*. By default, Excel has *automatic* calculation: it recalculates any cells that are dependent on other cells whose values have been changed. Abramson et al. (2001) calls it *sequential* calculation. This feature complicates evaluation of multiple cells that have built-in formulas in parallel.

As a result, the issue of hiding parallelism is resolved by separating random inputs and simulation outputs so that they are on separate *Worksheets*. We have a *Worksheet* template where a user specifies the number of inputs and their cell locations and similarly for outputs. When simulation runs are finished, users get outputs of each replication (if there is only one set of decision variables) or sample means and standard errors of each set of decision variables, i.e., a parametric sweep. Due to space limitation, we will only discuss the template and result display of the latter case.

This paper is organized as follows: We summarize related work in Section 2. We describe our system architecture in Section 3. We present our experimental results in Section 4, and we conclude with future research direction in Section 5.

2 RELATED WORK

In this section, we summarize works that address the issue of achieving parallelism, especially for spreadsheet.

Condor is a widely-known tool for maximizing utilization of computing resources (Litzkow, Livny, and Mutka

1988). The *Condor* scheduling system identifies idle machines and schedules background jobs on them. When those machines are used for non-*Condor* jobs, the *Condor* job is terminated and transferred to another machine, with the last system state before termination.

ActiveSheets is an application that allow parallel evaluation of spreadsheets. User interface is Microsoft Excel. *ActiveSheets* require custom functions for parallel calculation which are done at backend computers. When the computation is finished, the results are displayed on a spreadsheet. In Abramson et al. (2001), the backend platform is managed by EnFuzion (www.axceleon.com) on a high performance computing (HPC) system such as computer clusters or grids. On the other hand, Abramson et al. (2004) uses *NetSolve* (Agrawal et al. 2003), a grid middleware.

Nadiminti et al. (2004) introduce *ExcelGrid*, an open-source .NET plug-in that uses Excel as a front-end to a grid and perform user-defined calculations on it. Mustafee et al. (2006) show how *WinGrid* can enable *Witness*—a commercial simulation package—perform simulation replications in parallel on enterprise grid (linked resources within the same organization, as opposed to the public grid). In this work, users do not build a simulation model on *Witness* directly but specifying model parameters through Excel. The simulation results are also displayed on Excel.

ActiveSheets, *ExcelGrid* and *WinGrid* use Excel mainly for user interface; they do not perform computing-intensive calculation on Excel, and they are not designed specifically for spreadsheet simulation. On the other hand, our S3 still exploits Excel calculation, and it aims specifically at spreadsheet simulation applications. Hence, no complex programming is needed to benefit from our approach.

Microsoft also offers a tool for creating Excel Services for running a parametric sweep on a Excel 2007 Windows Compute Cluster Server (CCS) 2003 (Microsoft Corp. 2008). Excel Services is an architecture that allows Excel calculation on servers and enable applications to access Excel files. However, Excel Services and Windows CCS are not specifically designed for stochastic simulation. Thus, without further programming, users have no control over random number streams, and they have to implement their own algorithms for generating random variates.

Platform Symphony (www2.platform.com) is a commercial software designed to operate on enterprise grids. Platform introduces *Adapter* which enables Excel calculations to be run in parallel. Symphony’s *Adapter* is targeted for financial applications.

Widely used commercial spreadsheet simulation Excel add-ins also offer parallel versions: @Risk (www.palisade.com) has *RiskAccelerator*TM, and Crystal Ball (www.crystalball.com) provides *Crystal Ball Turbo*TM.

3 DESIGN AND ARCHITECTURE

We first describe a S3 spreadsheet simulation model. Then we provide details on system design and implementation.

3.1 Spreadsheet Simulation Models

By design, each simulation replication (i.e., runs) of the same model yield outputs that are independent and identically distributed (i.i.d.); therefore, we can achieve parallelism by assigning simulation replications that use different set of random inputs to each compute node, i.e., parallelism is achieved via *domain decomposition* (Quinn 2003). This technique relies on the execution of the same spreadsheet on multiple computers using different data set for each one. Once the calculation at all compute nodes are completed, simulation outputs are then aggregated and summarized, numerically via summary statistics, such as sample means and standard errors, or graphically, through Excel's charting tools.

It can be shown that random numbers of any parametric distributions can be transformed from uniform random numbers over the range $(0, 1)$ (see, for example, Banks et al. 2005 for proof). These standard uniform random variables are generated from mathematical algorithms, the so-called *random number generators* (RNGs). RNGs produce a very long sequence of pseudo-random numbers, and RNG seeds allow us to specify from which point in this sequence we get our numbers; we obtain the same sequence of numbers if the seeds are identical (see Henderson and Nelson 2006 for brief summaries on RNGs and random variate generation).

Excel has a built-in RNG, called `RAND()` which is not used in S3 for the following reasons: we do not know how to control sequence of numbers that `RAND()` produces. In addition, `RAND()` has some statistical deficiencies: Knusel (2005) and McCullough and Wilson (2005) discuss `RAND()` issues in Excel 2003. Therefore, we separate $(0, 1)$ uniform random numbers and Excel outputs from other Excel calculations.

We illustrate our approach via an example (more details in Section 4.1): We estimate value-at-risk (VaR) of stock portfolios by first simulating daily stock prices. In doing so, we need normally distributed random variables which in turn require uniform $(0, 1)$ random numbers. We fix the names of the following two Worksheets:

- `Model` contains calculation.
- `SimRun` is a template where a user specifies details on uniform $(0, 1)$ random numbers (number of rows and columns and their locations), outputs (number of outputs and their cell locations), and decision variables (number of decision variables, their location and their respective lower and upper bounds). See Figure 7. The location of some of

these cells are fixed, but some cell locations depend on locations of other cells. `SimRun` also holds uniform $(0, 1)$ random numbers that `Worksheet Model` needs. The RNG that our compute nodes use is Mersenne Twister ("mt19937") where we adapt to C# from C code in the GNU Scientific Library (Galassi et al. 2006). S3 knows the details about a simulation model through this worksheet.

In this VaR example, we have three simulation outputs that are estimated from multiple i.i.d. trials in `Worksheet Model`: VaR, average portfolio values on the next day and average profit/loss. S3 knows that they are outputs because we link them to `Worksheet SimRun`. Users can also execute simulation runs under multiple set of decision variables (DVs) by specifying the upper and lower bounds of each DV in `Worksheet SimRun` (see Figure 7). Currently, we only allow integer values with step size of 1. In this example, our DVs are the number of lots (one lot consists of 100 stocks) of each stock in our portfolio. DVs are specified on `Worksheet SimRun`, and they are used in `Worksheet Model` for calculation.

Once the calculation is completed, a user gets result on another `Worksheet` called `Output` (see Figure 8). For each combination of DVs in the specified range, sample means and standard errors are provided. (Standard error is a measure of how close a sample mean is to the unknown true mean. It is defined as a sample standard deviation divided by square root of the number of replications.)

3.2 Design and Implementation

We first explain the execution steps of S3 and discuss the details about system architecture.

The S3 system architecture is shown in Figure 9. The system is consisted of four main components:

1. *Users* upload Excel simulation models and download Excel output files when jobs are completed.
2. *Manager* is responsible for resource management (book keeping of status of workers), job management (job submission, job scheduling, and job allocation), and data management (managing data files).
3. *Workers* or compute nodes are PCs that execute Excel calculations. Currently, we have *dedicated workers* which are always available for *Manager* even if there are other jobs running on them.
4. *File Server* stores data files that are created during job execution. Users upload Excel files that contain their simulation models onto this *File Server* from which workers subsequently download. Once simulation is finished, users download Excel files that hold simulation results from the *File Server*.

The execution steps of S3 can be explained as follows (the number below correspond to ones in Figure 1).




Figure 1: System configuration.

1. A user creates a simulation model in Excel which has our add-in that connects to Manager through Web Services. Figure 2 shows the dialog box for job submission. Job description (such as the number of replications, number of workers and a seed number) is sent to Manager (1(a)) and the user's Excel file is sent to File Server (1(b)).
2. "Idle" Workers (not currently running Manager's jobs but maybe doing other jobs) periodically check with Manager to request jobs. If there are pending jobs, Manager sends them to Workers.
3. A Worker downloads an Excel file according to what Manager has assigned.
4. When a Worker completes his job, Worker uploads his job onto File Server.
5. Then Worker updates his status with Manager.
6. A user can check status of his submitted jobs through Manager (see Figure 3).
7. When user's job is completed, he downloads his output Excel file from File Server.

Figure 8 shows an example of output display.

4 EXPERIMENTAL RESULTS

We first describe our test problem and the experimental setups, followed by experimental results and discussion.

4.1 Value-at-Risk Test Problem

VaR is a risk metric that probabilistically describe market risks. Let L be investment profit or loss. Given some confidence level $\alpha \in (0, 1)$, the VaR of the portfolio at the confidence level α is given by the smallest number ℓ such that $\Pr\{L > \ell\} = 1 - \alpha$ (Jorion 2001). For a given




Figure 2: Dialog box when a user requests simulation runs.

combination of stocks, we simulate a value-at-risk over one-day period. We assume that stock prices follow a Brownian motion process. The calculation steps are (Hull 2003):

1. Compute the value of portfolio today using the current stock prices. Let S_{it} be the price of stock i on day t , and x_i is the number of stock i in the portfolio, $i = 1, 2, \dots, n$. The value of portfolio on day t is $P_t = \sum_i x_i S_{it}$.
2. Sample stock returns, $\mathbf{r} = [r_1, r_2, \dots, r_n]'$, from the multivariate normal distribution whose mean is zero and the covariance matrix is estimated from historical data.
3. Use r_i to determine the stock price on day $t + 1$: $S_{i(t+1)} = S_{it} e^{r_i}$. Revalue the portfolio value on day $t + 1$, $P_{t+1} = \sum_i x_i S_{i(t+1)}$.
4. Compute $\delta P = P_{t+1} - P_t$.
5. Repeat steps 2 to 4 m times to get samples of δP .

We estimate the $\alpha\%$ VaR as the α percentiles of m simulated values of δP . We use $m = 200$ and $n = 10$ stocks that are traded in the Stock Exchange of Thailand. The covariance matrix of stock returns are estimated from 50 trading days, from December 7, 2007 to February 20, 2008.

4.2 Experimental Setups

The experiment is done on a 19-PC system. The user's system is an AMD Athlon XP 1GHz system with 512MB RAM installing Windows XP. Both the manager and file server system are an AMD Athlon XP 2500+ system with 512MB RAM installing Windows Server 2003. The rest of computing nodes are an Intel Celeron 2.53 GHz system

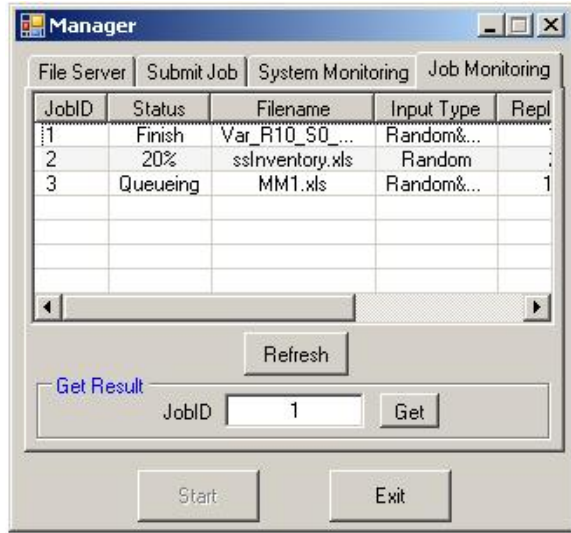


Figure 3: Dialog box when a user checks job status.

with 512MB RAM installing Windows XP. All machines are connected together using 100Mbps Fast Ethernet switch. In this work, all the softwares are developed with Visual Studio 2005 and Visual Studio Tool for Office (VSTO). The manager uses the FCFS (first come, first served) algorithm for job assignments.

We vary the number of combinations of decision variables (i.e., problem size) by changing the upper and lower bounds of each DVs, and the number of workers used. Each combination of decision variables gets 10 simulation replications. The performance measure is the computational time to complete the simulation runs. The test is done at 256, 512, 1024, 2048, 4096 combinations of decision variables. The number of workers used are 1, 2, 4, 8, and 16 workers.

4.3 Experimental Results

The runtime results are shown in Figure 4. We see that the run time decreases when number of processing nodes increases. This is due to the distribution of the processing tasks to multiple computing nodes simultaneously. However, benefits of increasing the number of workers diminish as the number of workers increases, e.g., the runtime decreases sharply when we include the second worker, but the benefit declines as we go from 8 to 16 workers.

We also consider a performance measure called *speedup* which is defined as a ratio between sequential runtime (runtime on one worker) and parallel runtime (runtime on multiple workers). Parallel run time consists of computation




Figure 4: Runtime on the desktop grid when the problem sizes and the number of workers vary.

time and communication overhead. Thus,

$$\text{Speedup} = \frac{T_{\text{Seq}}}{T_{\text{Parallel}}} = \frac{T_{\text{Seq}}}{T_{\text{comp}} + T_{\text{comm}}} \quad (1)$$

Speedup shows how many times faster the execution is when parallel computing is used; if we use n identical processors, the ideal speedup is n , i.e., when $T_{\text{comm}} = 0$, $T_{\text{comp}} = T_{\text{seq}}/n$ in (1). The actual speedup is lower due to communication overhead, which increases with the number of processors and the problem size. In our experiment, the communication overhead is mostly due to preparing communication channels between servers and workers, rather than in uploading or downloading files. Thus, for a given the number of processors, our communication overhead are relatively close across all problem sizes, and speedup for large problems is higher than for small problems. (that we consider).

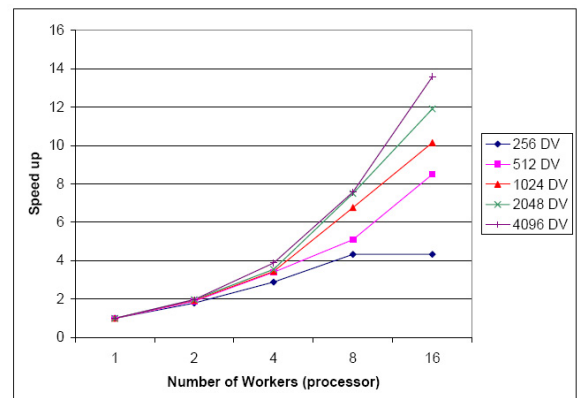


Figure 5: Speedup on the desktop grid when the problem sizes and the number of workers vary.

In Figure 5, we observe the following results:

- Speedup increases at a faster rate when the number of computing Service nodes are added. Maximum speedup gained in this experiment is 13.6; therefore, the application runs 13-14 times faster for the system of only 16 nodes.
- Speedup for large problems is higher than for small problems. The reason is that when more computing workload is available, the fraction of computing overhead (e.g., in load balancing and communication) to the computing workload is lower. Thus, our proposed system will work even better for large and complex problems. Note the seemingly unusual behavior of the results for 256 DVs, where the runtime for 8 and 16 processors are approximately 3 minutes. This is because the communication overhead is much larger than the parallel computational time, and this overhead is also close to the sequential run time. As a result, speedup is non-increasing (see (1)).

A fraction of runtime used as system overhead can be assessed through a ratio called *efficiency*. The efficiency of parallel computation is speedup divided by number of workers. Efficiency indicates the effectiveness that our computing systems are utilized to solve the problem. Due to communication overhead, efficiency is between 0 and 1 (100%), where being closer to one is desirable. Figure 6 shows the efficiency of our desktop grid. For a given problem size, the efficiency is high for a small number of workers since communication cost is low. As number of workers increases, the communication overhead increases as well. Thus, the efficiency is decreasing accordingly, and at a higher rate. In addition, running a larger simulation model is more efficient because the ratio of computation time to communication time is higher. We can also see that our implementation is very efficient since we can still maintain efficiency of more than 80% (0.8) for 16 nodes with the problem size of 4096 DVs.

5 FUTURE WORK

In this paper, we propose an architecture that allows spreadsheet simulation to use Windows-based desktop grids to accelerate the execution speed. Our approach is different in that we base almost all computations on Excel spreadsheets (except for random number generations). Thus, no complex programming is needed. We also show that the complexity of parallel computing can be mostly hidden from users through well-designed computation templates in Excel. With these templates, users can have a flexibility of modeling a simulation problem while enjoying massive computing power. From the experiments, we show that runtime can be reduced from 200 minutes to about 14 minutes. This speedup can




Figure 6: Efficiency on the desktop grid when the problem sizes and the number of workers vary.

make a huge difference in how a user analyzes problems; with desktop grids, they are able to consider many scenarios simultaneously or even to optimize a sizable model, thus gaining greater insights on the problem at hand.

Our work can be enhanced in many ways. More templates can be added for broader classes of problems. More transparency can be built so users are not aware of the manager existence by adding automatic job submission into Excel. For a longer term execution, some of the fault handling mechanism should be added to make it easy to use the system in a less reliable IT environment.

ACKNOWLEDGMENTS

This research is sponsored by the Thailand Research Fund Grant Number MRG4980121, and it is a part of the Thai National Grid Project supported by Ministry of Information and Communication Technology, the Royal Thai Government. The testing facility is supported by Department of Computer Engineering, Faculty of Engineering, Kasetsart University, Thailand. We would also like to thank the WSC referees for careful reading and helpful comments.

REFERENCES

- Abramson, D., J. Dongarra, E. Meek, P. Roe, and Z. Shi. 2004. Simplified grid computing through spreadsheets and NetSolve. In *Proceedings of the Seventh International Conference on High Performance Computing and Grid in Asia Pacific Region (HPCAsia'04)*. IEEE Computer Society.
- Abramson, D., P. Roe, L. Kotler, and D. Mather. 2001. ActiveSheets: Super-computing with spreadsheets. In *Proceedings of the High Performance Computing Symposium (HPC'01), Advanced Simulation Technologies Conference*, 110–115. San Diego, California: Society for Modeling and Simulation (SCS) Press.

- Agrawal, S., J. Dongarra, K. Seymour, and S. Vadhiyar. 2003. NetSolve: Past, present, and future; a look at a grid enabled server. In *Grid Computing: Making the Global Infrastructure a Reality*, ed. F. Berman, G. Fox, and T. Hey. John Wiley & Sons.
- Anderson, D. P. 2004. BOINC: A system for public-resource computing and storage. In *Proceedings of the Fifth IEEE/ACM International Workshop on Grid Computing*. Available via http://boinc.berkeley.edu/grid.paper_04.pdf [accessed March 11, 2008].
- Banks, J., J. S. Carson, B. L. Nelson, and D. M. Nicol. 2005. *Discrete-event system simulation*. 4th ed. Upper Saddle River, New Jersey: Prentice-Hall, Inc.
- Ecklund, P. 2007. Notes on excel calculations. Available via <http://faculty.fuqua.duke.edu/pecklund/ExcelReview/ExcelFormulasReview.pdf> [accessed March 11, 2008].
- Foster, I., C. Kesselman, and S. Tuecke. 2001. The anatomy of the Grid: Enabling scalable virtual organizations. *International Journal of High Performance Computing Applications* 15 (3): 200–244.
- Galassi, M., J. Davies, J. Theiler, B. Gough, G. Jungman, M. Booth, and F. Rossi. 2006. *GNU scientific library reference manual*. 2nd ed. Network Theory Ltd.
- Henderson, S. G., and B. L. Nelson. 2006. *Handbooks in operations research and management science, volume 13: Simulation*. North Holland: Springer-Verlag.
- Hull, J. C. 2003. *Options, futures & other derivatives*. 5th ed. Prentice Hall.
- Jorion, P. 2001. *Value at risk: The new benchmark for managing financial risk*. 2nd ed. McGraw-Hill Trade.
- Knusel, L. 2005. On the accuracy of statistical distributions in Microsoft Excel 2003. *Computational Statistics and Data Analysis* 48 (3): 445–449.
- Legard, D. 2004. IDC: Consolidation to windows won't happen. Available via <http://www.linuxworld.com.au/index.php/id;940707233;fp;2;fpid;1> [accessed March 11, 2008].
- Litzkow, M. J., M. Livny, and M. W. Mutka. 1988. Condor—a hunter of idle workstations. In *Proceedings of the 8th International Conference on Distributed Computing Systems*, 104–111.
- McCullough, B., and B. Wilson. 2005. On the accuracy of statistical procedures in Microsoft Excel 2003. *Computational Statistics and Data Analysis* 49 (4): 1244–1252.
- Microsoft Corp. 2008. Microsoft Excel running on Microsoft compute cluster. Available via <http://msdn.microsoft.com/en-us/library/bb463068.aspx> [accessed July 7, 2008].
- Mustafee, N., A. Alstad, B. Larsen, S. J. E. Taylor, and J. Ladbrook. 2006. Grid-enabling FIRST: Speeding up simulation applications using WinGrid. In *Proceedings of the 10th IEEE/ACM International Symposium on Distributed Simulation and Real-Time Applications (DS-RT 2006)*, 157–164.
- Nadiminti, K., Y.-F. Chiu, N. Teoh, A. Luter, S. Venugopal, and R. Buyya. 2004. ExcelGrid: A .NET plug-in for outsourcing Excel spreadsheet workload to enterprise and global grids. In *Proceedings of the 12th International Conference on Advanced Computing and Communication (ADCOM 2004)*. Available via <http://www.gridbus.org/papers/ExcelGrid.pdf> [accessed March 11, 2008].
- Paisittanand, S., and D. L. Olson. 2006. A simulation study of IT outsourcing in the credit card business. *European Journal of Operational Research* 175 (2): 1248–1261.
- Quinn, M. 2003. *Parallel programming in C with MPI and OpenMP*. McGraw-Hill.
- Ragsdale, C. T. 2004. *Spreadsheet modeling & decision analysis*. 4th ed. Mason, Ohio: South-Western (Thomson Learning).
- Seila, A. F. 2006. Spreadsheet simulation. In *Proceedings of the 2006 Winter Simulation Conference*, ed. L. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 11–18. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Seila, A. F., V. Ceric, and P. Tadikamalla. 2003. *Applied simulation modeling*. Thomson Learning.
- Shariff, A. M., R. Rusli, C. T. Leong, V. Radhakrishnan, and A. Buang. 2006. Inherent safety tool for explosion consequences study. *Journal of Loss Prevention in the Process Industries* 19 (5): 409–418.
- Stevenson, W. J., and C. Ozgur. 2007. *Introduction to management science with spreadsheets*. New York, NY: McGraw-Hill/Irwin.

AUTHOR BIOGRAPHIES

JUTA PICHITLAMKEN is an Assistant Professor in the Department of Industrial Engineering, Kasetsart University. Her research interests include ranking and selection procedures, simulation optimization, spreadsheet simulation, and stochastic processes. Her e-mail is juta.p@ku.ac.th.

SUPASIT KAJKAMHAENG is a Master student in High Performance Computing and Networking Center, Kasetsart University. His recent work has involved Cluster and Grid Computing. His email is kurata_sk@hotmail.com.

PUTCHONG UTHAYOPAS is an Assistant Professor in Department of Computer Engineering, Kasetsart University. His research interests include parallel/distributed computing, cluster and grid computing and parallel software tools. His email is pu@ku.ac.th.

	A	B	C	D	E	F	G	H	I
1	# of Random Input Column	10	# of Random Input Row	200	# of Output	3	# of Decision Variable	10	
2									
3	Parameter	Output	152469	94.46	-4499	Random Input	0.905667119	0.488449879	0.047825132
4							0.964480872	0.469999778	0.95441784
202							0.919439777	0.417144372	0.689897577
203						DV Input	1	1	1
204									
205	# of lot (100 stocks)	DV Input		ADVANC		BANPU		BBL	CPN
206		Start Range		1		1		1	1
207		End Range		4		1		4	17

Figure 7: Screenshot of Worksheet SimRun.

	A	B	J	K	L	M	N	O	P	Q		
1	Set	ADVANC	SCC	THAI	Average of Port value	Std. Err. of Port value	Average of Avg P/E	Std. Err. of Avg P/E	Average of value at risk	Std. Err. of value at risk		
2	1	1	1	1	152433	48.85342	57.98018	48.85342	-4235.35	118.1943		
3	2	1	1	2	155658.2	48.62206	58.20353	48.62206	-4273.16	118.5844		
4	3	1	2	1	173829.9	50.43734	54.87871	50.43734	-4340.3	118.7631		
5	4	1	2	2	177055.1	50.17006	55.10206	50.17006	-4367.48	127.5423		
36	1	1	1	1	1	1	176679.7	51.50393	54.70058	51.50393	-4391.98	127.3813

Figure 8: Screenshot of Worksheet Output.

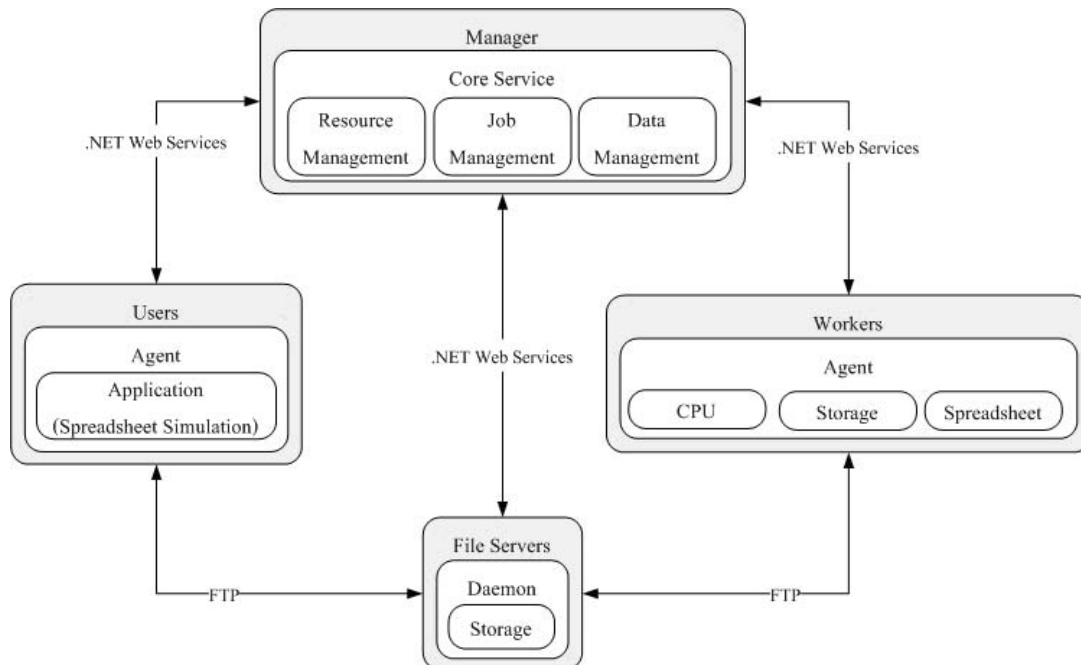


Figure 9: S3 architecture.