Economic Aspects of a Utility Computing Service

Jean Paul Degabriele
Trusted Systems Lab
Hewlett-Packard Laboratories,
Filton Road, Stoke Gifford,
Bristol BS34 8QZ, UK
+44 (0) 117 312 8658
jeanpaul.degabriele@hp.com

David Pym
Trusted Systems Lab
Hewlett-Packard Laboratories,
Filton Road, Stoke Gifford,
Bristol BS34 8QZ, UK
+44 (0) 117 312 8012
david.pym@hp.com

ABSTRACT

We present a case study of business and systems modelling for a Utility Computing service. Our analysis is concerned mainly with service pricing, resource flexing, and costs related to preventive security measures. We explain the conceptual and mathematical basis of our modelling approach, describing the details of a particular (discrete event, executable) model of a Utility Computing service, and show how the information obtained from such a model can aid business and design decisions.

Categories and Subject Descriptors
C.4 [**Performance of Systems**] – *design studies, performance attributes, reliability, availability, and serviceability.*

General Terms
Economics, Security, Performance, Management, Design, Experimentation.

Keywords
Utility Computing, Economics, Modelling, Pricing, Security, Resource Flexing, DEMOS 2000.

1. INTRODUCTION

Utility Computing is expected by many technologists to be one of the next major sources of income in the IT services market. While a considerable number of providers already hit the market, others are still studying the full potential of utility computing services. Most of the research is focused on the possible business models and their pricing, architecture of the infrastructure, and ways of securing utility computing services. Because of the many possible alternatives in designing and offering a utility computing service, we propose a modelling methodology to explore these possibilities and their interaction. We make use of Demos 2000 [6] as our modelling and simulation platform. Our analysis is mainly focused on system scrubbing, flexing, and utility computing service pricing.

GridNets 2007 October 17-19, 2007, Lyon, France.
 ISBN 978-963-9799-07-3.
DOI 10.4108/gridnets.2007.2223

2. UTILITY COMPUTING

Utility Computing is more of a different approach to computing resource than a new computing technology. The main idea is to offer computing resource as a utility on a pay-per-use basis, similarly to electricity and gas. Thus one no longer needs to invest in infrastructure, run it, maintain it, and secure it, in order to have computer resource at one's disposal. A number of services are possible that fall under the Utility Computing paradigm. *Data-oriented* services offer bulk storage and bandwidth — ideal for backup purposes for instance. *Computation-oriented* services are, however, more common. These offer computational power normally on a per CPU-hour basis, well-suited for computer graphics rendering, analysis of DNA and protein sequences, and other computationally intensive applications. At a higher level stand *Application-oriented* services. Here the service provider offers some proprietary software together with the necessary computer hardware on which to run it. Common examples of the applications offered are Customer Relationship Management Software (CRM), Database Management Systems, and e-Accounting software.

Generally the infrastructure is located in a *Data Centre* where it can be managed and maintained easily by the service provider, and the client can access the resource remotely. In Computation-oriented services (on which we focus mainly in this text) the client can have different degrees of remote access. At one extreme is the *Farm Renting* model, in which the client is allocated a network of machines (alternatively referred to as a Farm) over which he has complete control. At the other extreme is the *Job Submission* model, in which the client is presented with a web interface with which he submits his application together with a control script. The client can then retrieve the results of the computation from the web interface. The Farm Renting model allows the client to debug his applications before and during the job execution, whereas in the Job Submission model the application needs to be free of any bugs. On the other hand, the Job Submission model presents less exposure and hence better security. Moreover, it allows for better use of the infrastructure, because farms need not be allocated in fixed sizes. A *Resource Flexing* scheme can be used to allocate idle resources for job computation during periods of low load. In general this should increase the amount of available resources for when the next job request arrives, thereby maximizing the overall throughput of the infrastructure.

3. SECURING THE UTILITY COMPUTING INFRASTRUCTURE

The Utility Computing market is already a rather competitive market with a number of providers that offer a range of different services. Security is not yet, however, a major issue. Many providers claim that their services are secure but none of them specifies in detail the security measures that they employ. It is almost certain that when utility computing becomes more ubiquitous, a number of security incidents will occur in which clients will experience significant financial losses and service providers will suffer damage to their reputations. So, at a time when the utility computing market will be even more competitive, trust might be the discriminating factor between service providers. A number of security measures are possible in order to secure a utility computing infrastructure. Amongst others, there are: scanning of uploads against malware, encryption and authentication services between machines (through IPsec, TLS, or SSH for instance), farm separation (through routers/firewalls and/or vlans), IDS/IPS/AIS, and system scrubbing. In our analysis, we have focused mainly on system scrubbing, and categorized it on four different levels.

- *System Reuse*: This offers the least amount of protection. Here the machine is presented to the client with a used system where the client is allocated a new user account on the system with the previous users' accounts disabled. Vulnerabilities in the operating system could allow previous users to infect the system with malware, or allow the current user to access previous users' data. In the former case, the risk can be mitigated by scanning the uploaded data. Apart from confidentiality breaches, malware can also degrade the machines' performance. It is worth mentioning that if previous users delete (without overwriting) their data, then it becomes accessible to subsequent users. While if they retain it, it can't be accessed directly by another user.
- *System Refreshing*: Here the client is offered a machine with a freshly installed system. This in principle should eliminate the risk that the system is infected with malware and ensures that the system operates at its maximum performance. However the client may still recover data stored by the previous clients.
- *Clearing*: This refers to the removal of data so that it may not be reconstructed using normal system capabilities (i.e., not through physical access to the media). For this purpose, a single overwrite of all the memory space is normally enough. This process is more intensive in comparison to a system refresh where only the *File Allocation Table* is erased before installing the operating system.
- *Sanitization*: This is intended to protect against data recovery even in the event that the attacker has physical access to the media. This is normally accomplished through multiple overwrites or degaussing.

Clearing and Sanitization do not refer exclusively to secondary storage. Main memory and other memory buffers (such as the network card's memory buffer) may contain sensitive information as well. In [9], the United States Department of Defence defines a Sanitization Matrix which lists procedures suitable for clearing and sanitizing several storage media.

4. PRICING A UTILITY COMPUTING SERVICE

As in any other utility service, the business model employed and the service pricing play a pivotal role in its success. The business model has to suit the customers' needs in terms of accessibility, scalability, and usability amongst others. Most prominent are the *Subscription* and the *Metered Usage* business models. Hybrids of these models are also common – *Metered Subscription* is commonly employed by Internet Service Providers and Mobile Network Providers. A more detailed discussion of business models for utility computing services can be found in [14].

In [10], Low and Byde propose a pricing strategy based on auction. It is claimed that this scheme should keep a balance between supply and demand. There are, however, some subtle differences between traditional auction-based markets (such as a vegetable market) and the utility computing market. In utility computing, the consumption of goods may stretch over a considerable time span. Thus the present demand would affect the future supply and not just the present one. An auction based market is comparable to a feedback control system. A delay in the feedback loop reduces the system's stability resulting in an oscillatory behaviour which deviates from the equilibrium point. Another difference is that buyers may arrive sparsely in which case a buyer is unlikely to have any competitors or else the bidding stage has to be prolonged thereby degrading the service accessibility. Additionally, an auction scheme fails to take into account the risk incurred in accepting a relatively small job request, and then be unable to fulfil a more sizeable (and more profitable) job request at a later stage. In [11], Paleologo points out the inadequacy of a traditional *Cost-Plus* pricing methodology for utility services. Paleologo suggests a *Price-at-Risk* pricing methodology which takes into account the uncertainty in the pricing decision.

Our contribution is to demonstrate how computer simulation, based on executable models [6], can aid the pricing decision stage. We have built a model of a utility computing service. The model can be used for instance to calculate parameters such as the *Capacity* and the *Multiplexing Gains* used in the Price-at-Risk methodology. It can be used to quantify the gain obtained from a flexing strategy, or the costs of including one of the system scrubbing schemes discussed earlier. In our model, we associate a *Service Level* with each job request. This determines the amount of resources to be allocated for the job. In general the faster a job is computed the better for the client. Thus clients will demand services that require large amounts of resource for relatively short periods. Such a demand profile tends to reduce the effective capacity of the infrastructure. In view of this we recognise some appropriate properties for a just pricing scheme.

- It should encourage a constant load and discourage an erratic load pattern.

- The cost to compute a job should be proportionate to its size.
- The cost to compute a job should increase as the quality of service improves.

We have implemented two basic pricing schemes that follow these guidelines.

4.1 Scheme #1

A price is paid per CPU-hour which depends on the Service Level. The Service Level in turn determines how many machines (on average) will be allocated for the job. A possible flaw of this scheme is that a client can split his job into smaller ones, submit them simultaneously as separate jobs, and process them at a lower quality of service. Hence he effectively gets the same Service Level at a lower price. For instance assume that:

Table 1. Example Pricing Table for Scheme #1

Service Level	Machines Allocated	Cost/CPU-hour
1	10	£1.00
2	50	£1.20

Client A submits a 1000 CPU-hour job at Service Level 2. His job takes 20 hours and is charged 1200 pounds. Client B splits his 1000 CPU-hour job into five 200 CPU-hour jobs and submits them simultaneously at Service Level 1. This takes 20 hours and he is charged only a 1000 pounds.

4.2 Scheme #2

A price is paid per CPU-hour which depends on the Service Level. The Service Level determines the amount of time taken to compute the job. So the number of allocated machines depends also on the job size. Of course, not every job request can be supplied with the highest Service Level. Bigger jobs get better value for money. This might seem unfair but is reasonable if we consider that the client should always get a better service than he would get if he were to invest in infrastructure of his own.

5. OUR MODELLING TECHNOLOGY

Our modelling philosophy is that of classical applied mathematics, using tools from algebra, logic, computation theory, queuing theory, and probability theory [2,6,12,13,15]. Essential to our approach is the need to construct models at levels of abstraction that are appropriate for answering the questions of interest.

In general, we consider the concepts discussed below to be the essential components of a system model [2,6,12,13,15].

- Environment: Systems exist in complex, dynamic, possibly highly structured environments. Mathematically, the structural aspects of the environment are described in the same way as the system of interest (see below) and the dynamic aspects — here the main issue is generate events that incident upon the system — are captured stochastically using appropriate probability distributions and queues.
- Locations: Systems are, typically, highly distributed across spatial, temporal, or more abstract spaces. Mathematically, we capture the notion of location as a collection of connected 'places' that supports a notion of substituting one sub-location for another.
- Resources: The actions performed by systems are dependent upon the availability of accessible resources; moreover, as the system performs its actions, resources are modified (e.g., consumed, generated). Resources are, typically, associated with locations. Mathematically, we capture resources as collections of entities that may be combined and compared.
- Processes: A system's purpose is to execute processes in order to perform task, provide services, and so on. Processes access and manipulate resources. Mathematically, we capture the notion of process using the concept of a 'process algebra' (see, for example, [12,13,7]) which provides a collection of combinators, such as basic notion of prefix, notions of hiding, choice and concurrent composition, and recursion, which allow the definition of complex processes from a collection of basic actions.

These conceptual structures are partially captured by Demos2000 (henceforth 'Demos2k'), a semantically justified modelling language developed by Birtwistle, Christodolou, Taylor, and Tofts [6]. Specifically, Demos2k captures our conceptual set-up to the following extent:

- Environment: Demos2k provides a complete account of queues and a comprehensive collection of probability distributions;
- Locations: Demos2k does not provide a notion of location. Some aspects of location can be represented implicitly (though see Section 9);
- Resources: Demos2k has a basic but highly effective treatment of resource, lacking the richness of our mathematical notion but entirely adequate for our present purposes.
- Processes: Demos2k provides a rich and effective notion of process. The details are beyond out present scope.

The mathematical semantics of Demos2k has been given in [3,4].

6. THE MODEL

The model, depicted in Figure 1, was implemented using the Demos2k modelling language [6]. The Demos2k code for the model is given in full in [8], and is available at http://www.hpl.hp.com/techreports/2007/HPL-2007-101.pdf.

Our model consists of four main entities: Clients, Farms, Scrubbing Processes, and the Allocator. In terms of our conceptual organization, all of these are processes, which access various kinds resources (e.g., actual machines). The incidence of job requests is generated stochastically, using a negative exponential probability distribution. We do not make any use of the concept of location here (though see Section 9).

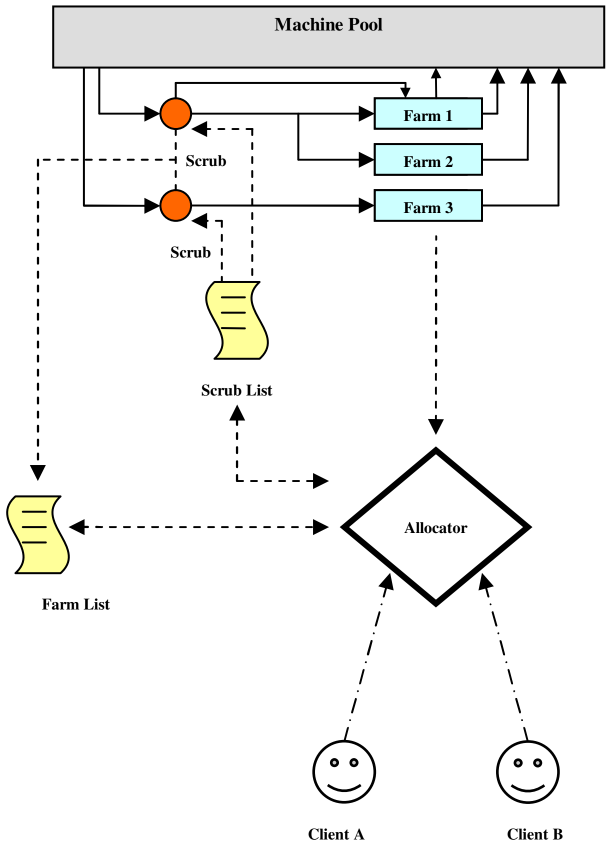

Figure 1. Basic Operation of The Model

The client's sole purpose is to generate job requests. Each client makes requests independently, where the time interval between each request is governed by an exponential distribution. A job request consists of two parameters: a job size W in CPU-hours and a Service Level *QoS*. The *QoS* is interpreted in accordance with the pricing scheme employed, but in general the higher the *QoS* the faster the job computation.

Farms are created by the Allocator entity in response to successful job requests from clients. Each farm corresponds to a single job, after which the farm is released on job completion. The farm process is implemented as a loop which continuously waits for instructions (in the form of syncs) mainly from the Allocator. Five instructions are defined.

1. *Lease Expired* signals that the time allocated for the job has come to an end. In response to this, the farm releases its machines and halts.

2. *Work Completed* signals the event that the job has been completed, in which case the farm releases its machines but continues to exist until a Lease Expired message is received. A state parameter is included to indicate the freshness of the message. The signal is fresh if the state parameter matches the current state of the farm.

3. *Resource Instruction* instructs the farm to take or release a number of machines. A Resource Instruction increments the current state of a farm and reschedules a new Work Completed instruction.

4. *Flexing Query* interrogates the farm for any flexing machines that it currently holds and can return back to the Allocator in order to fulfill job requests.

5. *Work Query* interrogates the farm for the amount of work in CPU hours needed to complete the job. This information is used by the Allocator to determine to which farm the idle machines should be allocated for flexing.

Scrubbing Processes are created by the Allocator in order to scrub a set of machines, before they are allocated to a farm. Machines are grabbed from the resource pool and held by a Scrubbing Process for a period of time defined by the model parameter *scrubTime*. Each of the four scrubbing schemes described earlier can be modeled by varying this parameter. On completion the Scrubbing Process would consult the *Scrub List* to determine to which farm(s) it should forward the machines. Finally the Scrubbing Process would release the machines, send the corresponding *Resource Instruction* signals to the intended farms, and terminate.

The *Allocator* is the central entity in the model, and controls most of the remaining entities. Its main responsibilities are: processing job requests, managing and flexing resources, and maintain records (such as the *Farm List* and the *Scrub List*). The flexing strategy adopted in the model uses a single-farm flexing algorithm. In the event that a set of machines are idle, the Allocator consults the Farm List and interrogates every farm with a Work Query. The idle machines are all allocated to the farm with the longest life expectancy. The life expectancy L of a farm F having a workload W_F and M_F machines is given by:

$$L = (W_F - M_F \times T_S) / (M_F + M_I)$$

where M_I = Idle Machines and T_s = Scrub Time. The scheme is zero-risk as each farm is always allocated enough machines to complete the job in the allotted time span.

The basic operation of the model is depicted in Figure 1, showing the flow of information between entities and how machines progress through their cycle. Clients submit job requests (W, *QoS*) to the Allocator. The Allocator determines the number of machines required to fulfill the request and checks how many machines are available. It starts by checking how many machines are idle. If these are not enough it consults the Scrub List for any machines destined for flexing. If the machines destined for flexing together with the idle machines do not add up to the required amount, the Allocator consults the Farm List for farms which have flexing machines and sends them a Flexing Query. The farms reply with the amount of machines that they can release (and still complete their job in time). If the total number of machines is enough to fulfill the request a farm is allocated, otherwise the job request is denied. A Scrubbing Process is created to clean the idle machines and the machines retrieved from other farms, while the machines retrieved from other Scrubbing Processes are redirected to the new farm by amending the Scrub List. The Allocator loops indefinitely waiting for job requests and other messages (such as job completion notifications). Each time it is inquired it checks for any idle machines and goes through the flexing subroutine.

7. RESULTS

The model has various potential applications. Some of these are:

- To determine the probability that an amount of machines will be enough to cater for a certain demand distribution;
- To calculate the gain in capacity that can be attained by a particular flexing strategy;
- To determine how distinct Service Levels should be priced;
- To quantify the impact of implementing a scrubbing scheme in a flexed architecture;
- To estimate parameters necessary for the Price-at-Risk methodology;
- To provide insight for SLA design.

Table 2. Sample Model Parameters

QoS	Job Duration	Probability	Price
1	6 months	0.05	£0.10
2	4 months	0.15	£0.14
3	2 months	0.15	£0.20
4	1 month	0.20	£0.30
5	2 weeks	0.25	£0.40
6	1 week	0.10	£0.60
7	3 days	0.06	£0.90
8	1 day	0.04	£1.50

The model (with the Demos2k code as given in [8]) was run a number of times, each with a different amount of machines. The Pricing scheme used was that described in scheme #2 where the price of each Service Level is listed in Table 2. The probability associated with each Service Level is the probability that a job request demands that Service Level. Figure 2 depicts a plot of the revenue attained by each run as the number of machines is increased. The experiment was repeated for a Scrub Time value of 1 hour instead of 5 hours, and once again with no flexing. Figure 3 represents the same three experiments but the abscissa represents the percentage of jobs denied.

It can be seen from Figure 2 that the red (flexed, Ts=1hr) and blue (flexed, Ts=5hrs) plots attain the maximum possible revenue while the green (not flexed, Ts=5hrs) plot approaches this maximum more slowly. A small increase in revenue due to a reduced Scrub Time is also evident. Instead of contrasting the revenue attained by each configuration, we can make a comparison in terms of investment. In particular we can compare the amount of infrastructure required to fulfill the same amount of requests and hence attain the same revenue. For instance from Figure 3 we can see that for a flexed configuration with 5000 machines the job denial ratio is 7.9 %. On the other hand the non-flexed configuration requires 6000 machines to attain an almost equivalent job denial ratio of 6.7 %. Hence in this scenario an increase in investment of 20% is required to attain the same revenue as a flexed architecture.

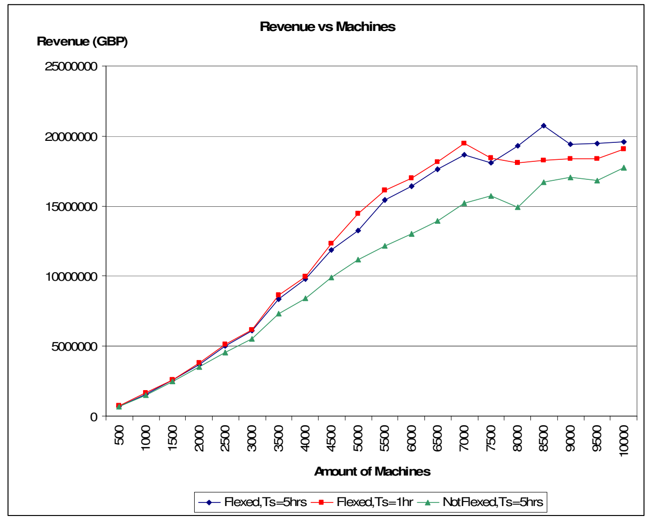

Figure 2. Revenue vs Number of Machines.

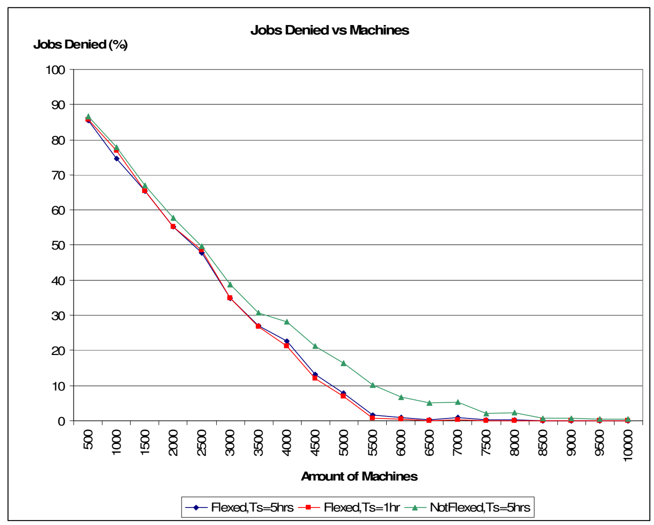

Figure 3. Percentage of Jobs Denied vs Number of Machines.

As mentioned earlier, high QoS values diminish the effective capacity of the infrastructure. This, of course, should be reflected in the service pricing. Our model can be used to quantify the effective capacities attained by each QoS, and hence determine a fair price for each Service Level. The model was adjusted such that each QoS is priced at 10p/CPUhr. Then it was run eight times, where in each run all job requests were processed at one particular Service Level. Figure 4 shows the revenue attained for each Service Level. As a starting point for a fair pricing scheme, we could adjust the price of each Service Level such that all bars attain the same revenue. Then according to these results, Service Level 8 for instance should cost three times as much as Service Level 7.

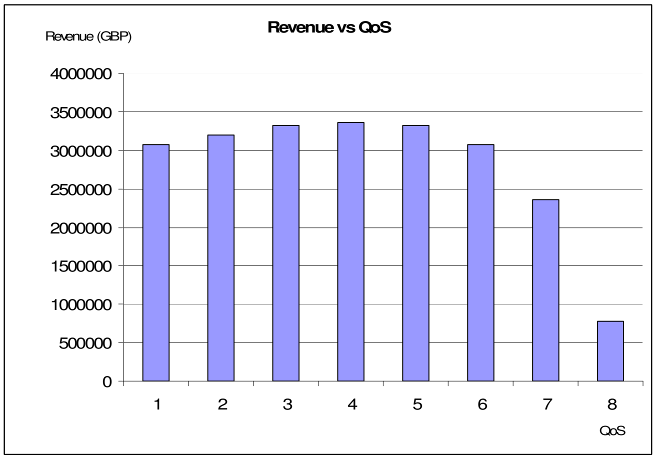

Figure 4. Comparing Revenues for Distinct Service Levels.

A rather unexpected effect, portrayed in Figure 4, is that in going from Service Level 4 down to Level 1, the revenue is seen to decrease. In our model, the client pays for the service on job completion. The simulation runtime here is two years, and the job processing times vary from 1 day to 6 months, depending on the Service Level. Thus, at lower Service Levels there is a greater amount of work which is not yet paid for at the end of the run. This situation might amount to a serious cash-flow problem. On top of this, one can add inflation and maintenance costs, which also increase as computation intervals get longer. So, if clients are to pay on job completion, it may not be viable to offer Service Levels between 1 and 4. Alternatively, one could circumvent this problem by changing to a business model in which the client pays either in advance or periodically until the job is completed.

8. RELATED WORK

A good summary of the various pricing and business models that have been proposed for utility computing can be found in [5]. Oceano, described in [1], is a utility computing infrastructure prototype developed at IBM. Including features such as scrubbing and flexing, Oceano proves to be quite in line with our utility computing model, although it is intended for web services rather than computation services. Yu, Buyya, and Tham [16] suggest a cost-based scheduling scheme in order to improve the internal workflow of a job. Potentially there might be scope to augment this scheme with a flexing strategy.

9. CONCLUSIONS AND FUTURE WORK

We demonstrated some of the potential that discrete event modelling holds for economic studies in relation to Information Technology and Information Security. In this paper we examined the case of Utility Computing, and showed how our simple model can help in business decisions, as well as exploring a bigger fraction of the possibility space.

In our analysis we aimed to keep the model as simple as possible. Two issues that are overlooked in this model are *Farm Fragmentation* and the *Atomic Transaction* nature of distributed jobs. The flexing algorithm moves machines from farm to farm without any notion of machine location, resulting in Farm Fragmentation. Mainly this is due to the difficulty in associating location with resource elements in DEMOS 2000. Thus our model does not represent the network traffic overhead which may result from a particular flexing algorithm. Secondly, distributed jobs are normally composed of smaller atomic transactions. However in the model it is assumed that job computation can be split into arbitrarily small transactions. By stepping the Allocator entity in discrete time and some other modifications it should be easy to include this in the model. Another limitation is that the implemented flexing strategy assumes jobs are arbitrarily distributable. Unfortunately this applies only to a limited set of problems, such as an exhaustive search of a cryptographic key. Therefore the 20% increase in infrastructure utilization mentioned in section 6 is essentially an upper bound of the flexing gain that the implemented flexing algorithm can attain. Thus if the cost of implementing such a flexing scheme is more than the cost of increasing the computing infrastructure by 20%, then such a scheme is obviously not viable.

Possible directions for future work could be to amend the model to portray these factors. The flexing strategy could be upgraded to a multiple-farm flexing scheme where idle machines are distributed among multiple farms rather than one farm. Risky flexing, where some jobs may not be completed on time, can also be investigated. The model could be amended to include maintenance and upgrading costs for a more thorough business continuity study. Jobs could be portrayed by a more general model, such as the one presented in [16]. Finally, it should be noted that the model can be easily adapted to portray other utility computing services, such as Web Services.

Turning to the modeling framework itself, one aspect of our current work is to build a version of Demos2k that supports an explicit notion of location which captures at least some of the properties of our conceptual framework. Such a language would allow us to extend the model described herein to capture more aspects of the distributed architecture supporting a Utility Computing service, such as Farm Fragmentation for instance.

10. ACKNOWLEDGMENTS

We are grateful to Chris Dalton, Brian Monahan, Chris Tofts, and Mike Yearworth for their advice on various aspects of this paper.

11. REFERENCES

[1] Appleby, K., Fakhouri, S., Fong, L., Goldszmidt, G., Kalantar, M., Krishnakumar, S., Pazel, D., Pershing, J., and Rochwerger, B. Oceano – SLA Based Management of A Computing Utility. In *7th IFIP/IEEE Intl. Symp. on Integrated Network Management*, May. 2001, pp. 855-868.

[2] Birtwistle, G., *Demos - Discrete event modelling on Simula*. Macmillan, 1979.

[3] Birtwistle, G., Tofts, C., The Operational Semantics of Demos: part I, *Transactions of the Simulation Society* 10(4):299 – 333, 1993.

[4] Birtwistle, G., Tofts, C., The Operational Semantics of Demos: part II, *Transactions of the Simulation Society*, 11(4):303 – 337, 1995.

[5] Buyya, R., Abramson, D., Giddy, J., Stockinger, H. Economic models for resource management and scheduling in Grid computing. *Concurrency and Computation: Practice and Experience* 2002, 14(13-15):1507-1542

[6] Christodolou, A., Taylor, R., and Tofts, C., 2000. Demos 2000. http://www.demos2k.org

[7] Collinson, M., Pym, D., Tofts, C. Errata for Formal Aspects of Computing (2006) 18:495–517 and their Consequences. To appear, *Formal Aspects of Computing*, 2007.

[8] Degabriele, J. P., Pym, D. Economic Aspects of a Utility Computing Service. *Technical Report HPL-2007-101, HP Laboratories, Bristol*, available from: http://www.hpl.hp.com/techreports/2007/HPL-2007-101.pdf

[9] DoD 5220.22-M *National Industrial Security Program Operating Manual (NISPOM)*, available from: http://www.usaid.gov/policy/ads/500/d522022m.pdf

[10] Low, C., and Byde, A. Market-Based Approaches to Utility Computing. *Technical Report HPL-2006-23, HP Laboratories, Bristol*, available from: http://www.hpl.hp.com/techreports/2006/HPL-2006-23.pdf

[11] Paleologo, G. A. Price-at-Risk: A methodology for pricing utility computing services. *IBM Systems Journal*, Vol. 43, No 1. 2004, available from: http://www.research.ibm.com/journal/sj/431/paleologo.html

[12] Pym, D. and Tofts, C. A calculus and logic of resources and processes. *Formal Aspects of Computing.*, 18(4): 495-517, 2006.

[13] Pym, D. and Tofts, C. Systems Modelling via Resources and Processes: Philosophy, Calculus, Semantics, and Logic. *Electronic Notes in Theoretical Computer Science*., 172, 545-587, 2007, available from: http://www.sciencedirect.com/science/journal/15710661

[14] Rappa, M. A. The utility business model and the future of computing services. *IBM Systems Journal*, Vol. 43, No 1. 2004, available from: http://www.research.ibm.com/journal/sj/431/rappa.html

[15] Yearworth, M., Monahan, B. Pym, D. Predictive Modelling for Security Operations Economics. *Proc.I3P Workshop on the Economics of Securing the Information Infrastructure*. Washington DC, 23-24 October, 2006. Available from: http://www.hpl.hp.com/techreports/2006/HPL-2007-125.pdf

[16] Yu, J., Buyya, R., and Tham, C. K. Cost-based Scheduling of Scientific Workflow Applications on Utility Grids. *Proc. 1st International Conference on e-Science and Grid Computing* (e-Science 2005), December 2005, 140-147.