

APPROXIMATE AND SIMULATION BASED ANALYSIS FOR DISTRIBUTED OBJECT SOFTWARE PERFORMANCE MODELS

Panajotis Katsaros and Constantine Lazos
Department of Informatics
Aristotle University of Thessaloniki
54006 Thessaloniki
Greece
E-mail: {katsaros,clazos}@csd.auth.gr

KEYWORDS

synchronous object invocation, nested object invocations, simultaneous resource possession, approximate MVA, flow-equivalent approximation, regenerative method

ABSTRACT

In complex software systems, the effectiveness of model based performance predictions is limited by the availability of appropriate solution techniques. These techniques should allow to take into account the software components interaction effects. In distributed object systems, the main problem is the simultaneous resource possession caused by the synchronous, often nested object invocations, which block the callers, until they get the replies. This paper provides a review of the analysis techniques, which address that fact, while preserving the abstract system view, offered by a queuing network representation. Two of these techniques, were proposed for solving a general class of models, with one or more layers of software servers and a third technique was designed specifically for distributed object software performance models. The advent of an extended flow-equivalent approximation, which is also described, opens new prospects for the development of efficient solution algorithms. Finally, simulation based estimation is discussed, in respect with the applicability of the well-founded and accurate, single-run regenerative method.

1. INTRODUCTION

Classical queuing network solution techniques, assume that a job can only use a single resource at a time and simultaneous resource possession is not possible. Moreover, most of them may be only used for analyzing typical “flat” system representations and do not directly support the solution of excessively complex queuing networks, like the ones obtained in realistic performance models of distributed software systems.

Early contributions to the problem of analyzing simultaneous resource possession, were based either on

- iterative estimation techniques with unknown calculation time and not guaranteed convergence for the performance measures of interest or
- low-level Markov based approximations that are neither practical nor scalable enough for software engineering purposes.

Recent advances, like the advent of an extended flow-equivalent approximation (Kurasugi and Kino 1999), constitute a new theoretical basis for the development of efficient solution algorithms.

In this paper, we review those techniques that we believe, they may have an impact in the practice of distributed object systems performance modeling and we comment on the new prospects opened by the last developments.

Accurate simulation based analysis of such complex models is an absolute necessity, since it is basically the only alternative for checking the validity of the approximations to be used. Thus, we also review theoretical results on the applicability of the single-run regenerative approach, in hierarchical queuing network models.

2. ITERATIVE PERFORMANCE EVALUATION ALGORITHMS

The Stochastic Rendezvous Networks

The Stochastic Rendezvous Networks (SRVNs) were first introduced in (Woodside 1989) and constitute a practical and scalable technique for the analysis of systems, with software components interactions. The rendezvous mechanism can be used to model different interaction types present in software, including all the types of object invocations (Woodside et al. 1995) taking place in distributed object systems. The solution for a SRVN can be found with an iterative algorithm that computes a series of intermediate solutions, using an MVA approximation and continues until the estimated throughputs converge.

The primitive structural entity of a SRVN model is the *task*. Each task may represent a single software or hardware component that communicates with other tasks, by messages of a request-wait-reply style. The task, which sends the message (*client task*), requests the rendezvous and blocks until it gets the reply. The receiving task (*server*) accepts the request and executes two (or more) service phases. The client task is blocked only for the time period of the first phase. Thus, a synchronous object invocation may be modeled by a rendezvous request with zero service requirements, for its second phase. On the contrary, asynchronous (one-way) object invocations are modeled as rendezvous requests with zero service requirements, for their first phase.

A SRVN differs from a typical queuing network, in that a server may also act as a client requesting service from one

or more lower layer tasks, in any phase of a rendezvous service execution. Each server may provide more than one service, modeled as separate task *entries*, with their own parameters. However, each task owns a single message queue with a specific service discipline. Entries correspond to distinguished object method invocations, since messages are directly addressed to the entries and not to the tasks.

There are three types of tasks, namely:

- the *pure clients*, labeled as 1, 2, . . . , R, which may only initiate requests,
- the *active servers*, labeled from R+1 to K, which may accept requests and initiate new ones and
- the *pure servers*, labeled from K+1 to N, which may only provide service to accepted requests.

Processors are pure server pseudo-tasks that provide hardware resource service to one or more active server tasks. The entries of the tasks, assigned to them, request service one slice at a time and each request and service are modeled in the SRVN as a rendezvous between the requesting task and the processor pseudo-task.

The SRVN solution algorithm, proposed in (Woodside et al. 1995), is summarized in the following steps:

1. Carry out the software model transformation to create tasks that represent processors, for co-allocated tasks. This involves:

- a. define a new pseudo-task entry ε for each entry e of all the tasks co-allocated at the same processor and an arc from e to ε for its processing requests

- b. label the arc with the mean number of processing slices, in phase p , calculated as

$$y_{ep} = 1 + \sum_{d \in \mathcal{E}} y_{edp}$$

where y_{edp} are the parameters denoting the mean number of entry e requests to d , during phase p

- c. if s_{ep} is the input denoting the mean service time provided by entry e , during phase p , set

$$s_{\varepsilon 1} = \frac{s_{ep}}{1 + \sum_{d \in \mathcal{E}} y_{edp}}$$

and $s_{ep}=0$, since the processor pseudo-task is a pure server and each entry ε of it owns only phase 1

2. Construct the so-called *task request graph*, with a node for each task and a directed arc from i to j , if the entry e in task i sends requests to any entry d in task j . This graph is used to determine the order of computation in the next steps. If there is a cycle in the task request graph, it is not possible to determine such an ordering.

3. Let us denote:

λ_e , the unknown throughput of entry e in messages/sec

λ_{0e} , the parameter specifying the arrival rate, if there is an external stream of requests to entry e

$Y_{de} = \sum_p y_{dep}$, the mean number of entry d requests to entry e , during all phases p

Then, $\lambda_e = \lambda_{0e} + \sum_{d \in \mathcal{E}} Y_{de} \cdot \lambda_d$, for all e from R+1 to N

and this constitutes a set of $N - R$ equations, with N unknowns. This may be also written as

$$\lambda_e = \sum_{r=1}^R \alpha_{er} \cdot \lambda_r + \sum_{d=R+1}^N \beta_{ed} \cdot \lambda_{0d} \quad (1)$$

for all e from R+1 to N, where the coefficients α and β are determined by applying Gaussian elimination to the original set of equations. Set $w_{kl} = 0$, with w_{kl} denoting the mean waiting time seen by entry k requests, when sending messages to entry l , for all k, l with $Y_{kl} > 0$.

4. Calculate mean service times (including the queuing delays) and throughputs, in the order found at step 2: For each task i , the mean service time is given as

$$x_{ep} = s_{ep} + \sum_{d \in \mathcal{E}} y_{edp} \cdot (t_{ed} + \delta_{ed})$$

with t_{ed} the mean delay for entry e , when sending a message to entry d and δ_{ed} is the specified mean round-trip communication delay from entry e to d .

For FIFO tasks, $t_{ed} = w_{ed} + x_{bl}$, since entry's e task remains blocked only for the duration of the first phase of the rendezvous request. Throughputs are given as

$$\lambda_r = \frac{1}{\sum_p x_{rp}}$$

for all entries r from 1 to R and from (1), for all entries from R+1 to N.

5. Estimate the mean waiting times for all entry requests directed to other entries, based on the MVA expression provided in (Woodside et al. 1995) and the alternative approximations proposed, in order to determine the task queue properties, at the instant of the rendezvous request arrival.
6. If throughputs are all sufficiently close to the previous iteration, then stop. Otherwise go to step 4.

The results obtained include throughputs, mean waiting times, entry service times and entry and task utilizations. The described solution technique is closely related to the Bard-Schweitzer approximate MVA (Schweitzer 1979) and the heuristics provided in (Woodside et al. 1995) support either FIFO or Preemptive Priority scheduling and stochastic or deterministic phase types, with possibly varied coefficients of variation, in the second case.

The SRVN based solution variants cannot be applied to models with multi-threading tasks.

The Method of Layers

In the method of layers (Rolia and Sevcik 1995), the system performance model is viewed as a sequence of layers, where processes request service only from one level lower, in the hierarchy. Processes with statistically identical behavior form a *group* or a class of processes. To identify the level of each group, in the hierarchy, a topological sort of the process groups may be first developed. Thus, each group is placed at exactly one level and the resulted graph is assumed not to contain any

cycles. If requests for service span more than one level, virtual flow-equivalent groups modeling the requests and their service rates have to be introduced, in all intermediate layers (this is not required for the SRVNs based solution techniques).

Each pair of successive levels, in the hierarchy, defines a submodel. The response time of a process that is considered as a client in one submodel defines its service time, when it is considered as a server in its alternate submodel. The performance estimates for the submodels are found using a modified version of the Linearizer algorithm (Chandy and Neuse 1982). The solution of all of the submodels represents an iteration of the algorithm. The algorithm is applied iteratively, until the changes in estimated mean response times, between successive iterations, is below a specified tolerance.

When solving the hierarchy of models only software contention is taken into account. After the algorithm terminates, a second model, in which software contention is ignored, is created. Each process is included in a single hardware resource contention queuing network and the amount of time that a process does not spend competing for devices is used as a think time in the model.

With new estimates of device responsiveness, another solution for the software contention model is found. The method of layers alternates between software and hardware contention models, until the estimates for mean response times of non-serving processes differ by less than some tolerance. This approach is similar to the method of complementary delays (Heidelberger and Trivedi 1983) that proceeds by using closed queuing networks, augmented by a virtual delay service center.

The method of layers supports the use of multiple entry FIFO rendezvous servers, as in the case of the SRVNs based solution techniques. Additionally, the use of the so called multi-servers permit the representation of multiple thread servers (as opposed to the SRVNs based solution techniques).

The combination of the use of the Linearizer, which often provides superior accuracy, when compared to the Bard-Schweitzer algorithm and the simultaneous solution of all servers in a layer yields, in some cases, more accurate results, compared to the ones obtained by the SRVN based solution variants. In the sequel, the basic method of layers estimation is outlined:

1. Initialize the response time estimates for process groups, with the mean service times, assuming no hardware or software contention. The initial response times are computed level by level in a bottom-up manner.
2. WHILE group response times have not reached a fixed point DO
 - WHILE group response times have not reached a fixed point DO
 - FOR each software submodel selected in a top-down manner DO
 - solve the submodel using Linearizer with the following residence time expressions (Rolia and Sevcik 1995): FIFO, rendezvous, multiple entry, multi-server and delay

```

END WHILE
solve the hardware resource contention submodels
using Linearizer with the following residence time
expressions (Rolia and Sevcik 1995): PS, FIFO,
LIFO, priority preemptive resume, and delay
Update the group response time estimates
END WHILE

```

The Method of Decomposition

The method of decomposition, introduced in (Kahkipuro 2000), was designed specifically for distributed object software performance models. It includes approximations, which overcome the limitation of not being applicable (as in the case of the methods already described) in models with cyclic graph dependencies. Cycles may be introduced, when two objects are allowed to make synchronous invocations to each other, as for example in the case of the callback interaction pattern. This type of interaction is sometimes used as a means to avoid extensive blocking times that may be caused by the use of a single synchronous object invocation. Recursive invocations to the same object are allowed, since this is also possible in real world applications.

In the sequel, we provide a description of the core solution technique. The overall approach lies on the well-known method of surrogate delays (Jacobson and Lazowska 1982), in order to decompose the model into multiple queuing networks, so that the primary and the secondary (software or hardware) resources of a blocking access are always in different networks. As a result, we obtain a set of multi-chain product form queuing networks.

For each case of simultaneous resource possession, the primary network is created by removing all secondary resources and all accesses to them and then by including a surrogate delay resource. The following transformations are carried out:

- for each access to the primary resource, its service demand is increased by the sum of the response times obtained for all blocking invocations to the secondary resources
- for each chain (class), the service demand of the surrogate delay server is defined to be the sum of the response times for all non blocked accesses of the secondary resources used by the chain

For each primary network, a set of secondary networks is created. These networks contain the secondary resource itself and an auxiliary delay server, for modeling the time the jobs spend elsewhere in the system. Each chain in the original queuing network generates at most two secondary chains:

- a closed one, if there are blocking accesses, with its population defined as the maximum number of accesses that can reach the secondary resource in parallel (multithreading server) and
- a secondary chain that corresponds to a chain with non blocking accesses to the resource, with the same parameters (equal arrival rates in case of open chains and equal populations in case of closed ones)

The service demands for all non blocking and blocking secondary resource accesses remain the same, unless there are nested blocking accesses to other resources. In this case, the service demand is increased by the sum of the response times obtained for the nested accesses. For closed secondary chains the service demand imposed to the auxiliary delay server is calculated through the Little's rule, by using the throughput of the calling resource, in the corresponding primary network. For open secondary chains there is no need to specify a service demand for the auxiliary delay server.

The resulted networks are to be solved by an approximate Bard-Schweitzer MVA algorithm, modified appropriately (by the load concealment transformation, as described in Agrawal 1985) for the simultaneous solution of the open and the closed chains. Input parameters for some networks require the existence of a solution for some other networks and it is even possible to obtain cyclic dependencies. Thus, the solution is carried out iteratively and in each repetition the input parameters of the networks are adjusted to the outcome of the previous repetition, until the throughputs of the secondary resources are sufficiently close to the throughputs of the corresponding primary ones.

To conclude, the method of decomposition does not support priority scheduling, as opposed to the SRVNs based solution, proposed in (Woodside et al. 1995) and the method of layers.

3. AN EXTENDED FLOW-EQUIVALENT APPROXIMATION

The standard decomposition or flow-equivalent method (Courtois 1977 and Chandy et al. 1975) is a well known approximation, based on the replacement of a subset of the model's queuing centers, with a single, flow-equivalent center, characterized by a calculated state-dependent service rate.

Flow-equivalent based performance estimation techniques have been used in various simultaneous resource possession models (Sauer 1981). However, these methods could not be applied to models, where jobs that possess different types of passive resources (i.e. different sets of blocked object servers, as in the case of synchronous nested object invocations) compete with other jobs for the same active resources.

Recently, in (Kurasugi and Kino 1999), the authors introduced an extended flow-equivalent approximation, for overcoming this particular restriction. The proposed approximation is described in the frame of a general two-layer queuing network, where the upper layer represents software resource contention and the lower layer, the caused hardware resource contention (Figure 1).

The existence of a single upper layer does not provide adequate support for the representation of nested object invocations, where multiple object servers may be blocked at the same time. However, the advent of such an extended flow-equivalent approximation opens new prospects for the development of efficient non-iterative estimation techniques and for this reason, we proceed to the description of it.

Let us consider a queuing model, where the upper layer consists of R FIFO software servers labeled as $1, 2, \dots, R$ and the lower layer consists of N queues representing hardware resources that are labeled as $1, 2, \dots, N$. Each software server owns $k_i (\geq 1)$ threads, $i \in 1, \dots, R$ and one queue.

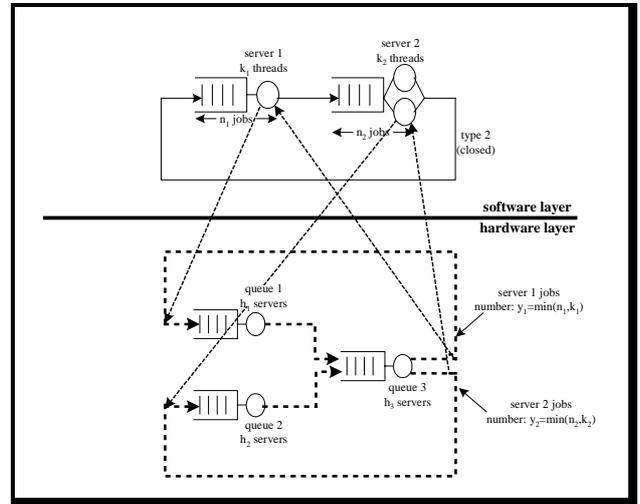


Figure 1: A two-layer queuing network with simultaneous resource possession

A job (object invocation) that arrives at server i will immediately be routed to the lower layer upon its arrival, if there is an available thread at the server. The time period in which a job passes to and returns from the lower layer is referred to as the server's service time. Each queue $j \in 1, \dots, N$ in the lower layer owns $h_j (\geq 1)$ servers and one queue. Each job routed to the lower layer from software server i will travel through the queues at that layer, according to a specified Markovian routing chain, associated with the software server i . Subsequent routing of a job (object invocation) through the various software servers, in the upper layer, can be conducted according to a number of possible (open or closed) routing chains, that represent the desired software functionality. In this article, our description, for the upper layer, is restricted to the simple case of a single closed routing chain, with v_i denoting the relative frequency for a job to visit the software server i .

Let $x_{j(i)}$ be the total number of jobs of software server i , placed at the lower layer queue, j . The number of jobs of software server i , in the lower layer is given as,

$$y_i = \sum_{j=1}^N x_{j(i)}$$

and the aggregate queue occupancy vector is defined as

$$\mathbf{y} = (y_1, y_2, \dots, y_R)$$

If n_i is the total number of jobs (object invocations) at the software server i , then $y_i = \min(k_i, n_i)$.

For any job arriving at the queue j , the required service time will depend on the particular software server the job comes from and is subject to an exponential distribution with parameter $\mu_{j(i)}$. The service discipline at each lower layer queue can either be FIFO, PS (Processor Sharing) or LIFOPR (LIFO Preemptive Resume). If the discipline is

FIFO, the parameters $\mu_{j(i)}$ are the same for any i , i.e. $\mu_{j(i)} = \mu_{j(h)}$, for any $h \neq i$.

The extended flow-equivalent method, proposed in (Kurasugi and Kino 1999), is summarized in the following steps:

1. Calculate the numbers of jobs $\mathbf{y} = (\min(n_1, k_1), \dots, \min(n_R, k_R))$ in the lower layer, for all possible software server vectors, $\mathbf{n} = (n_1, n_2, \dots, n_R)$.
2. Each \mathbf{y} -job based model is created from the original two-layer model, by disregarding the behavior of jobs in the upper layer and setting a short circuit at each software server, so that each job arriving at the server from the lower layer returns immediately to it. Thus, each short-circuited server can be effectively replaced by a state-dependent virtual service center, whose service rate is calculated in the following manner:
 - a. Solve the shaped \mathbf{y} -job based models, for all \mathbf{y} , by using the convolution algorithm, in order to obtain the normalization constants $G(\mathbf{y})$.
 - b. Approximate the service rate $\mu_i^*(\mathbf{n})$ at the software server i as

$$\mu_i^*(\mathbf{n}) = v_i \cdot \frac{G(\mathbf{y} - \mathbf{e}_i)}{G(\mathbf{y})}$$

where $\mathbf{e}_i = (0, \dots, 0, 1, 0, \dots, 0)$, with the unit placed in the i -th position.

In (Kurasugi and Kino 1999), the authors also propose a Markov based approximation and an alternative closed form approximation, for the upper layer, which we think are inappropriate for the performance evaluation of nested object invocations. Thus, we omit their description.

The extended flow-equivalent approximation of (Kurasugi and Kino 1999) may be used as is, in the hybrid simulation setting, in the way described in the next section. On the other side, we continue working to take advance of it, in order to provide a novel estimation technique that will successfully address the needs for efficient performance prediction, under the existence of synchronous, nested object invocations, in the software layer. The software servers' workload derivation procedure, to be used, will be based on the one introduced in (Vetland et al. 1993).

4. SINGLE-RUN SIMULATION BASED PERFORMANCE EVALUATION

Recently, (Nikolaidou and Anagnostopoulos 2003) introduced a network/application oriented modeling approach, for distributed system simulation. In (Katsaros and Lazos 2002) we focused on the development of (hybrid) simulation queuing models, for object based distributed software. Alternatively, (Brosso et al. 2002) introduces the use of XML Petri Nets (PNML), as a means for structured Petri-net model development. Also, recent advances (Anagnostopoulos and Nikolaidou 2003) in Faster-than-Real-Time simulation opened new prospects in model to system validation. In this section, we are restricted to the single-run output analysis of complex hierarchical model configurations, like the ones described.

We propose the use of the so-called regenerative method for sequentially controlled estimation of steady-state performance measures, for the following reasons:

- it is based on a sound theoretical foundation,
- it is not bound to the problem of the "system initialization bias",
- it is characterized by appealing asymptotic properties,
- its accuracy has been tested both in "flat" (Katsaros and Lazos 2003a), as well as, in hierarchical queuing network models (Blum et al. 1985) and
- it is also applicable in the hybrid simulation setting, with state-dependent service rates (Schedler 1993).

Although this method is not in widespread use, due to an inherent difficulty in identifying regeneration system states, there is significant theoretical progress towards this direction, in the last few years. The most notable advances rely on a generalized semi-Markov (GSMP) system representation and conclude to general results that may be utilized, for the regeneration state selection, at the more abstract level of a queuing network model (Schedler 1993). In (Katsaros and Lazos 2003a), we have experimentally evaluated the effectiveness of alternative selection strategies, in the sequentially controlled setting, in terms of both the obtained accuracy and the observed efficiency behavior.

5. CONCLUSIONS AND FURTHER RESEARCH WORK

This work is focused on the available approximate solution methods that we believe they may have an impact in the practice of distributed object systems performance modeling. They are all based on iterative estimations that scale well in large software performance models, but are also characterized by unknown calculation time and not guaranteed convergence, for the measures of interest.

The SRVNs based solution techniques lack support for cyclic object invocation dependencies, recursion and multithreading servers, since they were initially designed for the performance analysis of general software systems. The method of layers allows the use of multithreading servers, but requires that processes request service only from one level lower, in the hierarchy. Both approaches offer support for synchronous, as well as asynchronous object method invocations, by the use of the two-phase rendezvous type interaction mechanism and multiple entries software servers.

Priority scheduling is mainly required for the modeling of real-time distributed object systems (Katsaros and Lazos 2003b) and is also included as an open possibility for the two aforementioned methods. On the other side, it is not supported by the method of decomposition, which, as opposed to the two other alternatives, allows cyclic object invocation dependencies (callbacks) and recursion.

Although there are still great challenges for improvements in the existing core estimation techniques and the specialized approximation heuristics used by them, the advent of the extended flow-equivalent approximation,

described in section 3, opens new prospects for the development of new, efficient solution approaches.

Significant progress has been also carried out in the field of model parameter capture (Chatzigeorgiou 2003 and Vetland et al. 1993).

Simulation based analysis for hierarchical model structures like the ones used in the described analysis methods is an absolute necessity, since it is basically the only alternative for checking the validity of the used approximations. The proposed single-run output analysis method is based on a sound theoretical foundation, it is not bound to the problem of the “system initialization bias” and it is also characterized by appealing asymptotic properties. Furthermore, its accuracy has been tested in “flat” (Katsaros and Lazos 2003a), as well as, in hierarchical queuing network models (Blum et al. 1985) and we have also theoretically checked its applicability in the hybrid simulation setting, with state-dependent service rates.

REFERENCES

- Agrawal, S. C. 1985. *Metamodeling: A study of approximations in queuing models*. MIT Press, Cambridge, MA, USA.
- Anagnostopoulos, D. and M. Nikolaidou. 2003. “Timing Issues and Experiment Scheduling in Faster-than-Real-Time Simulation”. (To appear in) *Transactions of the Society for Modeling and Simulation*. SCS.
- Brosso, I.; G. Bressan and W. V. Ruggiero. 2002. “Simulation e-business applications using PNML”. In *Proceedings of the 2nd WSEAS International Conference on Simulation, Modeling and Optimization (ICOSMO)* (Skiathos, Greece, Sep. 25-28). WSEAS.
- Blum, A.; L. Donatiello; P. Heidelberger; S. Lavenberg and E. A. MacNair. 1985. “Experiments with decomposition of extended queuing network models”. In *Modeling Techniques and Tools for Performance Analysis*, D. Potier (Ed.). Elsevier Science, 623-640.
- Chandy, K. M. and D. Neuse. 1982. “Linearizer: A heuristic algorithm for queuing network models of computing systems”. *Communications of the ACM* 25, No. 2, 126-134.
- Chandy, K. M.; U. Herzog and L. Woo. 1975. “Approximate analysis of general queuing networks”. *IBM Journal of Research and Development* 19, 43-49.
- Chatzigeorgiou, A. 2003. “Performance and power evaluation of C++ object-oriented programming in embedded processors”. *Information and Software Technology* 45, 195-201.
- Courtois, P. J. 1977. *Decomposability: Queuing and computer system applications*, Academic Press.
- Heidelberger, P. and K. S. Trivedi. 1983. “Analytical queuing models for programs with internal concurrency”. *IEEE Transactions on Computers* 32, No. 1, 73-82.
- Jacobson, P. A. and E. D. Lazowska. 1982. “Analyzing queuing networks with simultaneous resource possession”. *Communications of the ACM* 25, No. 2, 142-151.
- Kahkipuro, P. 2000. *Performance modeling framework for CORBA based distributed systems*. PhD thesis, Department of Computer Science, University of Helsinki, Finland.
- Katsaros, P. and C. Lazos. 2002. “Structured performance modeling and analysis for object based distributed software systems”. In *Proceedings of the 15th ISCA International Conference on Parallel and Distributed Computing Systems* (Louisville, Kentucky, USA, Sep. 19-21). ISCA, 96-102.
- Katsaros, P. and C. Lazos. 2003. “Return state selection for improved effectiveness in sequentially controlled regenerative simulation”, (submitted)
- Katsaros, P. and C. Lazos. 2003. “A simulation test-bed for the design of dependable e-services”, (To appear in the) *3rd WSEAS International Conference on Simulation, Modeling and Optimization* (Rethymno, Greece, Oct. 12-15). WSEAS
- Kurasugi, T. and I. Kino. 1999. “Approximation methods for two-layer queuing models”. *Performance Evaluation* 36-37, 55-70.
- Nikolaidou, M. and D. Anagnostopoulos. 2003. “A distributed system simulation modeling approach”. *Elsevier Simulation Modelling – Practice and Theory* 11, No. 3-4, 251-267.
- Rolia, J. A. and K. C. Sevcik. 1995. “The Method of Layers”, *IEEE Transactions on Software Engineering* 21, No. 8, 689-700.
- Sauer, C. H. 1981. “Approximate solution of queuing networks with simultaneous resource possession”. *IBM Journal of Research and Development* 25, No. 6, 894-903.
- Schweitzer, P. 1979. “Approximate analysis of multiclass closed networks of queues”. In *Proceedings of the International Conference on Stochastic Control and Optimization* (Amsterdam, The Netherlands).
- Shedler, G. H. 1993. *Regenerative Stochastic Simulation*. Boston, Academic Press.
- Vetland, V.; P. Hughes and A. Solvberg. 1993. “Improved parameter capture for simulation based on composite work models of software”. In *Proceedings of the 1993 Summer Computer Simulation Conference* (Boston, Massachusetts, USA), 110-115.
- Woodside, C. M. 1989. “Throughput calculation for basic Stochastic Rendezvous Networks”, *Performance Evaluation* 9.
- Woodside, C. M.; J. E. Neilson; D. C. Petriu and S. Majumdar. 1995. “The Stochastic Rendezvous Network Model for performance of synchronous client-server-like distributed software”. *IEEE Transactions of Computers* 44, No. 1, 20-34.

PANAJOTIS KATSAROS holds a BSc degree in Mathematics from the Aristotle University of Thessaloniki - Greece, an MSc degree in Software Engineering from the University of Aston, in Birmingham – UK and a PhD in Computer Science, from the Aristotle University of Thessaloniki. He is a lecturer of Computer Science at the Aristotle University of Thessaloniki – Greece. His research interests are focused on queuing network-based performance analysis of alternate load distribution and fault tolerance strategies, in distributed object systems. He has also published research articles on simulation methodology issues and one book on Programming Language Compilers. e-mail: katsaros@csd.auth.gr url: <http://delab.csd.auth.gr/~katsaros/index.html>

CONSTANTINE LAZOS holds a BSc degree in Mathematics from the University of Athens - Greece, an MSc degree in Computer Science from the University of Birmingham - UK and a PhD in Computer Science from the University of Southampton - UK. He was a lecturer for two academic years, at the University of Birmingham and he is a Professor of Computer Science at the Aristotle University of Thessaloniki, since 1980. His published work and research interests are focused on computer systems performance analysis and simulation techniques and he is the author of numerous educational books on Databases, Compilers and Programming Languages (C++, FORTRAN etc). e-mail: clazos@csd.auth.gr