



Queueing Theory and Interdisciplinary Perspectives

Proceeding of the COLLOQUIUM, May 8th 2006
In celebration of the 40 years career of Lester Lipsky

Editors

Eike Jessen, Manfred Jobmann

Organisation

Chair for Network Architectures and Services
Department of Computer Science, Technische Universität München

Technische Universität München



Queueing Theory and Interdisciplinary Perspectives

Proceedings of the COLLOQUIUM, May 8th, 2006
In celebration of the 40 years career of Lester Lipsky

Eike Jessen, Manfred Jobmann (ed.)
Technische Universität München, TUM

Queueing Theory and Interdisciplinary Perspectives
Proceedings of the COLLOQUIUM, May 8th, 2006
In celebration of the 40 years career of Lester Lipsky

Editors:

Eike Jessen
Lehrstuhl Netzarchitekturen und Netzdienste (I8)
Technische Universität München
D-85748 Garching b. München, Germany
E-mail: jessen@in.tum.de
Internet: <http://www.net.in.tum.de/>

Manfred Jobmann
Lehrstuhl Netzarchitekturen und Netzdienste (I8)
Technische Universität München
D-85748 Garching b. München, Germany
E-mail: jobmann@net.in.tum.de
Internet: <http://www.net.in.tum.de/>

Cataloging-in-Publication Data

Queueing Theory and Interdisciplinary Perspectives
Proceedings of the COLLOQUIUM, May 8th, 2006
In celebration of the 40 years career of Lester Lipsky

Eike Jessen, Manfred Jobmann
ISBN: 3-937201-24-6
DOI: 10.2313/NET-2011-10-1
ISSN: 1868-2634 (print)
ISSN: 1868-2642 (electronic)
Lehrstuhl Netzarchitekturen und Netzdienste (I8) NET 2011-10-1
Series Editor: Georg Carle, Technische Universität München, Germany
© 2011, Technische Universität München, Germany

Preface

On May 8th, 2006, a group of scientists assembled at the University of Connecticut in Storrs to honor Lester Lipsky, who, having begun his scientific career as a physicist has become a world-level authority in computer performance modeling. Twelve of his scholars presented contributions related to the scientific work of Lester Lipsky; other scientists contributed results to Lester Lipsky at the occasion on this day which – though not part of the presentations of May 8th, 2006 – were devoted to Lester Lipsky and are included in this volume. Amir Faghri, Dean of the School of Engineering, delivered the greetings of the University. Lester Lipsky himself gave a biographical sketch by modeling 'Life as a Markov Chain'.

In addition to the wealth of scientific results provided as contributions to the colloquium, the event was a manifestation of the large family of friends and adherents of Lester Lipsky. He has played an exceptional role in attracting personalities into the world of his concepts and results, and into the warmth and cordiality of himself.

Note: The contributions in these proceedings are arranged in the sequence of the chronological order of the talks at the colloquium (marked as "Slides"). The layout of the slides may have been slightly changed by the editors in order to harmonize the layout of the proceedings.

Eike Jessen and Manfred Jobmann
Munich, August 2011

Welcome from the Dean

Welcome to the University of Connecticut and the School of Engineering. I am pleased that you have joined us here from all over the world to honor Professor Lipsky. Professor Lipsky is indeed an internationally distinguished scholar which is why we are gathered here to recognize him.

However, I should also note that he is a good mentor and teacher. We in the School of Engineering are very fortunate to have him as a mentor and colleague. He has significantly contributed to the mission of teaching, research and outreach of the Computer Science Department at the University of Connecticut. The quality of his archival publications and books are a testament to his outstanding contributions in the field.

I hope that you also have enough time to enjoy our beautiful campus. Best wishes for an enjoyable seminar to honor Professor Lipsky.

Regards,

Amir

Amir Faghri

United Technologies Endowed Chair Professor in Thermal-Fluids Engineering

Department of Mechanical Engineering

University of Connecticut

Contents

1	Jessen E.: <i>Quality of Service Revisited: A Practitioner's View (and the Power of Theory)</i>	5
2	Jessen E.: <i>Quality of Service Revisited: A Practitioner's View (and the Power of Theory) (Slides)</i>	9
3	Schwefel H.-P.: <i>Quantitative Analysis of Access Strategies to Remote Information in Network Services</i>	13
4	Schwefel H.-P.: <i>Quantitative Analysis of Access Strategies to Remote Information in Network Services (Slides)</i>	19
5	Jobmann M.: <i>Nonexponential Time Distributions in Biocatalytic Systems: Mass Service Replacing Mass Action</i>	27
6	Jobmann M.: <i>Effects of Molecular Timing on the Shape of Dose-Response Curves (Slides)</i>	37
7	Conneely M.: <i>Doubly and Triply Excited States of Atomic Systems</i>	47
8	Buzen J.: <i>Improving Simulation Accuracy through the Use of Synthetic Alignment Intervals</i>	51
9	Asmussen S.: <i>Distribution of Total Times for Tasks Subject to Failure and Restart (Slides)</i>	73
10	Nagy G.: <i>Ambiguities in clustering Subsequence Time Series (Slides)</i>	81
11	Antonios I.: <i>Performance-relevant Network Traffic Correlation</i>	93
12	Antonios I.: <i>Should Correlation Properties Govern Traffic Model Design? (Slides)</i>	101
13	Crovella M. E.: <i>Simulation with heavy-tailed workloads</i>	111
14	Ding Y.: <i>How does a physicist approach queueing theory? - A few examples (Slides)</i>	129
15	Seth S.: <i>Signal Probabilities in AND-OR Trees (Slides)</i>	137

QoS REVISITED: A PRACTITIONER'S VIEW (AND THE POWER OF THEORY)

Marco Hoffmann, Siemens AG

Eike Jessen, Technische Universität München

Manfred Jobmann, Technische Universität München

Quality of Service (QoS) has been a concern in communication technology from the beginning. It became a challenge, when the internet with its statistical packet multiplexing had to serve also for real time media and data transport, as e.g. video transmission and conferencing, telephone and other interactive collaboration techniques, and real time control. Whereas in classical communication technology availability of the channel and absence of noise and distortion were the main criteria for quality, in IP-networks low and constant delay with few packet losses are the targets, which widely can be realized by adequate bandwidth and buffer capacity.

Given the inherent characteristics of the IP transport mechanism, there are dynamic and static solutions to the problem of Quality of Service. During the nineties, two basic techniques have been developed, discussed, and analyzed widely, without a breakthrough in actual applications. First was the principle of bandwidth reservation, substantiated in Integrated Services (1990). For the particularly important scenario of media multicast, the resource reservation protocol RSVP was developed. These techniques rely on admission control, traffic shaping, and policing. They turned out, however, not to be feasible in large or heterogeneous networks. A simpler technique, more scalable to large and heterogeneous networks, mainly relied on prioritisation: Differentiated Services (1996). It mixes the priority scheme with some reservation mechanisms, not for single flows, but for classes of aggregated traffic. A particular choice of this technology has been implemented in some networks under the name "Premium Service".

Besides, there are some more or less static solutions for the QoS-problem. Simplest is overprovisioning: The network has to provide enough bandwidth to keep the utilisation low in any expectable condition of operation, at less than 20%, for instance. This is widely used in core networks where the bandwidth is cheap; it is hardly a solution for radio networks. For some important applications, and a classical technology, one implements equal size channels and assigns them on demand: This is the solution in the case of conventional telephone, and it has been transferred to GSM and UMTS. Finally, there is a growing form of network use for restricted communities, as e.g. in grid computing. Often, the load situation is foreseeable. Then building an exclusive network (maybe virtual) with guaranteed capacities for this restricted connectivity can be an economical solution. Such a particular form of a virtual private network may be based on devoted optical channels (Dense Wave Length Division Multiplexing, DWDM) or on layer two/layer three label switching.

It is worthwhile to look at the situation of the scientific networks, where our examples are taken from Europe, but the situation in USA is mainly the same. On the European level, there is the GEANT-2 network, with 10 Gbps trunks. The network is overprovisioned, but also

offers a premium service. Intra- and extra European end-to-end premium service is in experimental operation. On the national level, there are core networks with trunks between 1 and 10 Gbps, and regional structures for the last 0..30 miles. Few of them offer premium services. It is in access nets and the interconnections between core and access, that overload is frequent and QoS becomes an issue. Finally, within the scientific institutions there are campus or institute networks, with capacities of 0.1..10 Gbps; very few of them offer premium services.

The general situation may roughly be characterized by saying that the abundance and low price for bandwidth undermines motivation for the advanced quality assuring mechanisms. Nevertheless, DFN (Deutsches Forschungsnetz) wanted to study the feasibility of QoS in large networks; it gave a project to Technische Universität Darmstadt (Ralf Steinmetz) and to Technische Universität München in 2002, to build a small level-3 testbed, implement and study IP QoS-technologies there, and to design a methodology for the extrapolation of testbed results to larger networks, which was the work of Marco Hoffmann in Munich; primarily the larger network would be G-WiN, the Gigabit-Wissenschaftsnetz of DFN. To generalize the problem, we follow a technique frequently used in engineering. To get insight into a large and complex system, which is not yet real, we build a model, the testbed. It is small but may be configured and operated at our choice. We study the effect of QoS there under a load which is equivalent to the load foreseen in the G-WiN. Then we have to extrapolate our observations to the G-WiN. For this, we need a meta-level model, which is usable in the testbed and the G-WiN, a bridge between both. Adequate methods for this extrapolation are simulation fed by observations in the testbed, and in principle also analytic methods. In some scenarios, the model simply is the existing system, in which some changes may be made, the results of which are to be extrapolated to the system as a whole. These extrapolation problems are more frequent than the design of a system from the base.

In our particular project, simulation as an extrapolation instrument was roughly validated within the testbed. We used a simulator which, by construction, exactly realized the configuration and load in the testbed. Then we modified the testbed and checked that the simulator would deliver a traffic equal to our measurements in the modified testbed.

As it is hardly possible to simulate more than a part of a high-speed wide-area network, we also studied analytical methods for extrapolation and validated these by comparison to different simulators, with the following results: We found them applicable where the system behaviour is dominated by a small substructure, so in the case of a severe network bottleneck. In the simplest cases, G/GI/1-models were found to be useful. We evaluated their feasibility over a wide range of utilization and coefficient of variation. The applicability of these approximations is sharply reduced by lack of independence of the random variables and by long-term-dependencies. As to the parameter distributions, we found out that the power-tail distributions studied so intensely by Lester Lipsky were too hard for approximations; it is a merit of Lipsky's matrix-exponential methods that these cases can be calculated exactly.

Lester Lipsky made his first visit to our small group at Technische Universität München in 1994 and made us familiar with Power Tail Distributions, long-term dependencies, and

matrix-exponential modelling and evaluation technology. I think it is easy to see, how this world of innovative statistical concepts began to influence our more established methodology, in particular the naive belief that low utilization would cure all evils.

Far more important was the spirit of devotion to science that he brought to us, his contagious pleasure to approach hard problems, his persistency in work and his pride to present surprising results; he really could transform people into researchers. And all without even mentioning his deep concern for human, social, and political responsibility of all of us!

Reference:

Hoffmann, M.: Verfahren zur Übertragung von Dienstgüteaussagen, Doctoral Thesis, Technische Universität München, Fakultät für Informatik, 2005

Quality of Service Revisited: A Practitioner's View (and the Power of Theory)

Marco Hoffmann, Siemens AG

Eike Jessen, Technische Universität München

Manfred Jobmann, Technische Universität München

QoS in Communication

Classical: Availability, absence of noise and distortion and delay. Accentuated with

IP packet multiplexing: quasi random delay and loss, due to inadequate bandwidth and buffer capacity, and retransmission.

Basic dynamic solutions for the IP QoS problem:

- Bandwidth reservation: Integrated Services (IS, 1990), Resource Reservation Protocol RSVP. Admission Control, Traffic Shaping, Policing: not feasible in large or heterogeneous networks
- Prioritization: Differentiated Services (DS, 1996). Priority with some reservation for classes of aggregated traffic. „Premium Service“.

Basic static solutions for the (IP) QoS problem:

- Overprovisioning: Provide enough bandwidth to keep the utilization under e.g. 20%.
- Implement equal channels and assign one on demand: Conventional telephone, GSM, UMTS.
- Establish an exclusive network with guaranteed capacities: for a restricted connectivity. Can be economical solution for a grid, e.g. „VPN“, based on optical channels (DWDM) or layer 2/ layer 3 label switching.

Scientific Networks (in Europe)

European Level: GEANT-2 Network, 10 Gbps, overprovisioned, but Premium Service offered. Intra- and extra-European end-to-end experimental operation.

National Level: Core Networks 1..10 Gbps with last 0 .. 30 miles regional structures; few offer Premium Services.

Institutional Level: Campus Networks, 0,1 .. 10 Gbps; few offer Premium Services.

The Extrapolation Issue

To study the feasibility of QoS in a large network, DFN (Deutsches Forschungsnetz) gave a project to Technische Universität Darmstadt (Ralf Steinmetz) and to Technische Universität München in 2002, to build a small level 1..3 testbed, implement and study IP QoS there, and (Marco Hoffmann in Munich) design a methodology to extrapolate testbed results to larger networks (G-WiN Gigabit-Wissenschaftsnetz of DFN, e.g.)

Generalization:

Primary System: Model

- well known
- affinity to the secondary system
- may be modified (or not)
- behaviour is measurable

for us: testbed

Secondary System: Object of Interest

- partially known
- may not be modified for analysis (operational or non-existent)
- current behaviour is measurable (if existent), not future

for us: (part of) G-WiN

Extrapolation by Meta-Model →

- simulation or
- analytic methods

Often the primary system simply is to be transformed into the secondary: Frequent and important engineering task.

Extrapolation:

Simulation as an extrapolation instrument was validated within the testbed, realizing two different models on the basis of the testbed.

Approximate analytical methods for extrapolation were validated by comparison to simulation, with the result:

- applicable where the system behaviour is dominated by a small substructure (severe network bottleneck e.g.)
- in a wide range, good $G/GI/1$, approximations are available
- applicability of these approximations may be reduced sharply by excessive variance or by lack of independency of random variables and by long-term-dependencies.

Quantitative analysis of access strategies to remote information in network services

Rasmus L. Olsen¹, Hans-Peter Schwefel¹, Martin B. Hansen²

¹Center for TeleInFrastruktur, Aalborg University, Niels Jernes Vej 12, 9220 Aalborg, Denmark

² Dept. of Mathematical Sciences

Email:¹{hps|rlo}@kom.aau.dk, ²mbh@math.aau.dk

Abstract—Remote access to dynamically changing information elements is a required functionality for various network services, including routing and instances of context-sensitive networking. Three fundamentally different strategies for such access are investigated in this paper: (1) a reactive approach initiated by the requesting entity, and two versions of proactive approaches in which the entity that contains the information element actively propagates its changes to potential requesters, either (2) periodically or triggered by changes of the information element (3). This paper first develops a set of analytic models to compute different performance metrics for these approaches, with special focus on the so-called mismatch probability. The results of the analytic models allow for design decisions on which strategy to implement for specific input parameters (change rate of the information element, network delay characterization) and specific requirements on mismatch probability, traffic overhead, and access delay. Finally, the analysis is applied to the use-case of context-sensitive service discovery.

Keywords: Distributed systems, remote access, performance modelling, context-sensitive networking

I. INTRODUCTION

Timely, remote access to dynamically changing information elements is a common problem for a large range of functionalities in different layers of modern telecommunication networks:

- On the link-layer, efficient radio-resource management at base-stations requires information about channel state and buffer filling as measured in mobile devices.
- On the network layer, routing decisions require the knowledge about the existence and the characteristics of links between remote intermediate nodes. This is particularly relevant when topology changes are rather frequent such as in wireless multi-hop networks[1].
- Network Services, such as dynamic distributed data-bases as used in certain name-services in mobile networks, require knowledge about (remotely performed) updates of the name to address mapping [2].
- Context-sensitive services require access to typically remotely obtained context information. Context information may thereby be used both during service execution [3] as well as for the service discovery process [4].
- For highly dependable networks and services, resilience is obtained by replication of services, which requires state-updates at remote replicants in order to avoid inconsistency [5], [6], [7].

Common to all these use-cases of access to remote information is that basic design decisions on how to efficiently implement such access need to be taken. Efficiency is thereby typically measured by access delay, probability of using 'correct' information, and network traffic overhead created by the remote access strategy. Two basic types of solutions exist:

- 1) **Reactive**, 'on-demand' access: Whenever a certain piece of remote information is needed at the processing entity, it is actively obtained (request) from the remote entity that has access to this information. This in principle implements a client-server architecture.
- 2) **Proactive** distribution of information: The entity that has control of the information element will pro-actively distribute updates of its value to potential 'requesters'. Thereby, two underlying sub-strategies can be distinguished
 - a) **Event-Driven** proactive updates: Whenever the information element changes value, an update is triggered. For a further differentiation with respect to the semantics of these updates, see Sect. III-C.
 - b) **Periodic** proactive updates: After certain time-intervals, the current value of the information element is distributed to potential request processes.

In this paper, we provide the methodology and the results of the quantitative analysis of different performance metrics, including in particular the so-called mismatch probability for the different remote access strategies. Section III introduces the different analytic models, while the quantitative results of these analytic models and their validation via simulations are discussed in Sect. IV. Finally, the analysis is applied to a use-case scenario of context-sensitive service discovery in Personal Networks in Sect. V.

II. PROBLEM FORMALIZATION AND PERFORMANCE METRICS

This section provides an abstracted description of the access procedures to remote information using stochastic processes. This description allows to analytically obtain different performance metrics, in particular including the so-called mismatch probability.

The simplified model contains three parts:

- The information element is maintained by a remote node (information provider) and it dynamically changes

its value at certain points in (continuous) time. It is assumed here that these changes will always result in a value previously not observed, e.g. as it is the case for monotonic changes. Since the actual value of the information element is not relevant in this paper (only the fact whether it has changed), we will use a point process $\mathcal{E} = \{E_i, i \in \mathbf{Z}\}$, where E_i is an increasing sequence of event times numbered such that E_0 is the event just before 0. The process \mathcal{E} is called the **event process**. $E(t)$ denotes the value of the (monotonically increasing) information element at time t , see the Appendix.

- The remote information element is required by a certain entity (the requester/client) at certain moments in time, identified by the **request process**, $\mathcal{R} = \{R_k, k \in \mathbf{Z}\}$, which in turn is a point process denoted in the same way as the E_k 's. Depending on the selected update strategy, an event of the request process may trigger an actual request to the remote server (reactive approach), or it may lead to an instantaneous access to the local replication of the information element in the pro-active approaches.
- Communication between requesting entity and server is described by stochastically varying delays, the upstream delays, $\{U_k, k \in \mathbf{Z}\}$ between requester and server (only in case of pro-active approaches), and the downstream delays, $\{D_k, k \in \mathbf{Z}\}$. Messages are never lost, however, these delays are potentially unbounded. Messages are identified by sequence numbers, so that out-dated updates can be detected and discarded.

In this paper, we will limit our discussion to independent, identically distributed (iid) delay processes, corresponding in practice to cases in which the inter-message times are larger than the time-scales at which queues build up and drain in the network due to congestion. For fast core-networks, such an iid delay assumption is realistic. Furthermore, some of the performance metrics below, in particular the stationary mismatch probability in the reactive case, are insensitive to correlation properties of the delay processes.

Random variables with the upstream and downstream delay distributions are denoted generically as \mathcal{U} and \mathcal{D} , respectively. These delay distributions correspond to the end-to-end delays between information provider and requester, hence e.g. cases of wireless multi-hop communication can be included via appropriate choice of \mathcal{U} and \mathcal{D} . Also, message drops can be included via degenerated distributions (with probability mass at infinity).

In this paper, we consider the following three performance metrics, where focus is put on the mismatch probability:

- 1) Network overhead: The amount of data transmitted on the network in the remote access strategy.
- 2) Access delay: The time interval from the moment when a certain piece of information is needed at the requester until it is finally available for use. For the pro-active access strategies, this delay is zero. Processing times are neglected in this paper.

- 3) Mismatch probability: The probability that the value of the information element that is used at the requester does not match the current true value at the remote location. The consequence of such a mismatch depends on the specific application, see e.g. Sect. V.

III. ANALYTIC MODELS

A. Reactive, on-demand access

Figure 1 illustrates the message flows in the reactive approach. In this scheme a request is initiated by the client at




Fig. 1. Reactive access: In the example, the k -th access, R_k , leads to a 'correct' value, while the $k+1$ th access causes a mismatching value.

time R_k , which is received by the provider at time $R_k + U_k$. The provider creates a response message containing the value of the requested information element which is received at the requester at time $R_k + U_k + D_k$. In the shown example, the k -th request leads to a correct value, since no changes of the information element occurred, while the response is being transmitted. The $(k+1)$ -th request on the other hand leads to a mismatching value. For the assumption that the event process, \mathcal{E} , is a Poisson process with rate λ , the mismatch probability can be calculated as follows, see the appendix for the derivation:

$$\text{mmPr}_{\text{react}}(\lambda, \mathcal{D}) = 1 - \mathcal{L}\{f_{\mathcal{D}}\}(\lambda),$$

where the last term is the Laplace transform of the density of the down-stream delay, \mathcal{D} , evaluated at value λ . Note that the mmPr is independent of the request process, \mathcal{R} . However, \mathcal{R} will influence statistical properties of corresponding estimators of mmPR. Note also, that the mmPr is not depending on the upstream delay process. Two different cases for the downstream delay are interesting and considered later in this paper:

- Constant (deterministic) delay: $\mathcal{D} \equiv c$

$$\text{mmPr}_{\text{react}}^{(\text{det})}(\lambda, c) = 1 - \exp(-\lambda c). \quad (1)$$

- iid exponentially distributed delay with rate ν :

$$\text{mmPr}_{\text{react}}^{(\text{exp})}(\lambda, \nu) = 1 - \frac{\nu}{\lambda + \nu} = \frac{\lambda}{\lambda + \nu}. \quad (2)$$

Note that the mmPr assuming an exponential delay is in fact smaller than in the deterministic setting with $c = 1/\nu$ for all values of λ , and ν (since $e^{-x} \geq (1+x)^{-1}$ for all values of x). The network overhead, $V_{\text{react}}(T, s, \mu)$, in a time interval of duration T is depending on the message sizes $s \in \{s_u, s_d\}$ (upstream and downstream, respectively), and the rate μ of the (not necessarily Poisson) request process: $V_{\text{react}}(T, s, \mu) = \mu T(s_u + s_d)$. The average access delay is $E\{\mathcal{U}\} + E\{\mathcal{D}\}$.

B. Proactive - Periodic update

For the proactive case, no request messages are needed, but the remote node sends updates to the requester. First we discuss the 'periodic' version, i.e. the update is sent after some (potentially stochastically varying) time interval independent of event and request processes. See Figure 2 for an illustration. Assuming a Poisson process for the event process (with rate




Fig. 2. Proactive periodic update using a deterministic period: R_k results in mismatching value, while R_{k+1} leads to a correct value.

λ), iid exponentially distributed downstream delays with rate ν , and that the 'periodic' updates are determined by a third independent Poisson process with rate τ , the mismatch probability can be computed through the steady-state probabilities of the following 3-state continuous time Markov chain:

State 1:	Correct value at requester
transitions:	event \rightarrow S2, Update generated \rightarrow S1, Update arriving \rightarrow S1
State 2:	mismatch at requester, no correcting update in transit
transitions:	event \rightarrow S2, Update generated \rightarrow S3
State 3:	mismatch, correcting update in transit
transitions:	event \rightarrow S2, Update generated \rightarrow S3, Update arriving \rightarrow S1

Note that 'outdated' updates (which were sent out before the last event occurred) are irrelevant. If the delay-distribution is probabilistic, multiple updates in transit can occur, which would need to be counted in the state-space (transition marked with * above in the table). For simplicity, the table above and the generator matrix below do not implement this counting of updates in transit, although the numerical results in the subsequent sections include them. Under the Poisson assumptions on Event process, downstream delay, and update sending period, the state transitions can be described by the following generator matrix:

$$Q = \begin{bmatrix} -\lambda & \lambda & 0 \\ 0 & -\tau & \tau \\ \nu & \lambda & -\nu - \lambda \end{bmatrix}.$$

The mismatch probability is then the steady-state probability that the Markov process is in States S2 or S3, which has the following closed-form solution:

$$\text{mmPr}_{proact,periodic}^{(exp)}(\lambda, \nu, \tau) = \frac{\lambda [(\nu + \lambda)(2\tau + \nu + \lambda) + \tau^2]}{(\tau + \lambda)(\nu + \tau + \lambda)(\nu + \lambda)} \quad (3)$$

See [8] for the detailed analysis of the case with multiple updates in transit. The overhead can be computed as follows:

$$V_{proact,periodic}(T, s, \tau) = \tau T s_d.$$

The (average) access delay is 0.

C. Proactive - Event based update

In the proactive event driven update scheme, the provider sends an update to the client node, whenever an event has

happened, i.e. when the information element has changed value, see Figure 3. In order to investigate the mismatch




Fig. 3. Proactive event driven update: the request at time R_k results in a correct value, while R_{k+1} leads to a mismatch, since the updated value is in transfer when the user accesses the current value.

probability, two different cases with respect to the semantics of the update messages have to be distinguished:

Case I: Incremental Updates: In this scenario, the requester only accesses the correct information, if all update messages from previous events have been successfully received. These messages can be re-ordered by the network, but through the use of sequence numbers, the requester is able to put them back in the correct sequence. In this case, a mismatch would occur, if any of the update messages is still in transit. This is equivalent to the probability that an $\mathcal{E}/\mathcal{D}/\infty$ queue is in a busy period (a customer being served in the queue is equivalent to an update in transit). Hence, the mismatch probabilities can be computed as

$$\text{mmPr}_{proact,incr}^{(GI)}(\mathcal{E}, \mathcal{D}) = \Pr(\mathcal{E}/\mathcal{D}/\infty \text{ queue is busy}),$$

which reduces under Poisson assumptions for \mathcal{E} (with rate λ) and General Independent (GI) assumptions for the downstream delay \mathcal{D} (with mean \bar{D}) to

$$\text{mmPr}(\lambda, \bar{D}) = 1 - \exp(-\lambda \bar{D}).$$

The downstream delay \mathcal{D} can be GI, since the steady-state queue-length probabilities for the $M/GI/\infty$ queue are identical to the $M/M/\infty$ queue, see [9]. In case of constant (deterministic delay), the above mmPr is equivalent to the reactive case. Hence, for constant delay, re-active and proactive, incremental, event driven access strategies lead to the same mmPr.

Case II: Full updates: If a single update message contains all information so that previous updates are not needed at the requester, it is only important that the update message of the last event has reached the requester. Hence, the mmPr can be derived from a similar mapping to a queuing model, but here, only the last customer (event) is relevant. Hence, instead of an $\mathcal{E}/\mathcal{D}/\infty$ queue as in Case I, we are now in the setting of a finite $\mathcal{E}/\mathcal{D}/1/1$ queue with pre-emptive service and only a single customer in the system (a customer in service is pushed out and discarded by a newly arriving customer). Appendix B derives a general formula for the mmPr in that case, with the following two special cases:

- Constant (deterministic) delay: $\mathcal{D} \equiv c$

$$\text{mmPr}_{proact,full}^{(det)}(\lambda, c) = 1 - \exp(-\lambda c). \quad (4)$$

- iid exponentially distributed delay with rate ν :

$$\text{mmPr}_{proact,full}^{(exp)}(\lambda, \nu) = 1 - \frac{\nu}{\lambda + \nu} = \frac{\lambda}{\lambda + \nu}. \quad (5)$$

In both cases, the mmPr is identical to the corresponding reactive setting. In both cases, the overhead follows as:

$$V_{proact,event}(T, s, \lambda) = \lambda T s_d.$$

However, note that typically the message size for the incremental updates is (much) smaller than for the full updates. This difference depends on the complexity of the data structure of this information element, which is outside the scope of this paper. The (average) access delay is 0.

IV. ANALYTIC RESULTS AND VALIDATION VIA SIMULATION

A summary of selected analytic result is given in the following table: Figure 4 shows the results for the mmPr

	reactive	proact. event full update	proact. event incremental	proact. periodic
mmPr				
Exp. Delay	$\frac{\lambda}{\lambda+\nu}$	$\frac{\lambda}{\lambda+\nu}$	$1 - e^{-\lambda/\nu}$	$\approx \text{Eq. (3)}$
Det. Delay	$1 - e^{-\lambda c}$	$1 - e^{-\lambda c}$	$1 - e^{-\lambda c}$	see[8]
overhead access delay	$\mu(s_u + s_d)$ $E(\mathcal{U}) + E(\mathcal{D}) > 0$	$\lambda s_d^{(a)}$ 0	$\lambda s_d^{(i)}$ 0	$\tau s_d^{(a)}$ 0




Fig. 4. Comparison of mismatch probabilities in the different remote access strategies:

as computed by the analytic models for the different remote access strategies, for the assumption of a Poisson event process with rate $\lambda = 1$ and an iid exponentially distributed downstream delay with rate ν . In the proactive periodic case, the period is iid exponentially distributed with varying rate $\tau = 10^{-2}, \dots, 10$. The table and the figure allow to draw the following conclusions:

- For each of the cases, exponential delay and deterministic delay, the reactive approach leads to the same mmPr as the corresponding pro-active event driven approach with full updates. Hence, the latter is not shown in the figure.
- The reactive strategy in the case of deterministic downstream delays, $\mathcal{D} \equiv 1/\nu$, (solid line) leads to a higher mmPr than in the case of an exponentially distributed

delay with same mean (dashed line). In contrast to intuition from other analytic models, e.g. in queueing models in which deterministic delays typically lead to shorter waiting times, here the deterministic case is not the best case scenario!

- For scenarios of long exponentially distributed delays, $\nu \rightarrow 0$, the derivative $d/d\nu$ of mmPr at the value $\nu = 0$ is $-1/\lambda$ for the reactive approach, while the derivative is zero for the pro-active event-driven incremental approach. Hence, for small values of ν (corresponding to long delays), the reactive approach is always creating a smaller mmPr.
- For very short downstream delays (large ν) both the reactive and the pro-active event-driven strategies decay asymptotically as mmPr λ/ν for both deterministic and exponential delays, and also independently of incremental or full updates. Hence, asymptotically for $\nu \rightarrow \infty$, all proactive event-driven and reactive strategies behave equally.
- for large $\nu \rightarrow \infty$, the pro-active periodic approach shows a limit of $\lim_{\nu \rightarrow \infty} \text{mmPr}(\lambda, \nu, \tau) = \lambda/(\lambda + \tau) > 0$. Consequently, for large τ eventually, the periodic approach will at some point always perform worse than the event-driven and reactive approaches.

Validation by simulation




Fig. 5. Theoretical results vs. simulated results varying the down link rate. 95% confidence intervals are provided for most of the simulation estimates, however, those are obtained without consideration of correlation properties.

Figure 5 shows a comparison of the mismatch probability as estimated from simulation experiments for the various strategies for varying downstream delay rates, ν . The results are obtained from simulating 1000 requests and the comparison of simulation and analytic results validates the analytic formulas and provides a visual illustration of variance properties of the corresponding simulation estimator. In addition to the validation of the analytic results, the simulations in Figure 5 show also two cases that have not been treated analytically:

- The case of periodic updates with deterministic downstream delay and update time interval ($\tau = 1$ sec.)

(triangle): In contrast to the periodic update strategy with an exponential delay and update time interval, the deterministic delay and update time interval leads to a reduced mmPr.

- The case of multiple information providers: In this scenario, the information element is a tuple of which the different elements are provided at different entities. This setting will also be considered in the next section. If the different elements change according to independent Poisson processes, the tuple changes also according to a Poisson process with the sum of the individual rates. Detailed analytic modelling of this scenario is outside the scope of this paper, but the black star curve in Fig. 5 shows the simulation result for a 5-tuple with distributed providing entities and exponential downstream delays. See the next section for more discussion on this scenario.

V. APPLICATION TO CONTEXT-SENSITIVE SERVICE DISCOVERY/ROUTING

In this section, we discuss the impact of mismatching remote information for the example of context-sensitive service discovery in Personal Networks. Thereby, we also extend the performance metric from the pure mismatch probability to expected values of observed information deviation. The latter requires a semantic description of the information element, which here will be done based on a simplified setting for illustration purposes.

A. Context sensitive service discovery

A Personal Network (PN) [10] is a logical private network that interconnects the user's Personal Area Network (PAN) with remote nodes; the latter are typically grouped in so-called clusters. Figure 6 shows an example of a PN which consists of two interconnected clusters (one of which is the user's PAN). Since PNs may be large and geographically distributed,




Fig. 6. Context aware service discovery in Personal Network

and furthermore they could contain many devices and hence potentially many services, context-dependent service ranking may strongly increase the user-friendliness of the service-discovery process. An example of a hierarchical service discovery architecture is illustrated, [4], in in Figure 6. The user submits a service discovery request which is sent to the Service Management Node (SMN) in the PAN. The SMN obtains then the context information from information providers in the PAN or in other clusters using one of the three strategies discussed in earlier parts of this paper. After the collection of context information, the SMN can rank and filter the service(s), e.g. based on a calculated score. If the calculated score is below a

certain threshold, the service is considered not relevant and is not shown to the user. An example for such a score-function is

$$\text{score}(t) = \frac{\sum_{n=1}^N w^{(n)} f^{(n)}(E^{(n)}(t), E_{ref}^{(n)})}{\sum_{n=1}^N w^{(n)}}, \quad (6)$$

where $f^{(n)}$ is an a pre-defined function, which determines a score for the matching of the n -th context field, $E^{(n)}(t)$, in comparison to a reference value, $E_{ref}^{(n)}$. Here $w^{(n)}$ denotes a weighting factor for the different context elements. As Equation (6) shows, the score is based on up to n context values, which are processed at the SMN. There can be a mismatch between the context value used at the SMN for service ranking and the true value at the remote node, which in turn will lead to a wrongly calculated score value, leading to a possible wrong service (de)selection.

B. Impact of mismatch probability on score function

We use a simulation model that simulates up to n context providing nodes, in which the context information is assumed to be monotonically increasing. The same simple, linear scoring function $f^{(n)}(t) = 4E^{(n)}(t)$ is used for all context values with same weights $w^{(n)} = 1$. For each context access strategy and each parameter setting, 100.000 service discovery requests are simulated, and the average of the absolute error in score value, is calculated. All results are obtained using an exponentially distributed link delay. The proactive, event-driven scheme is based on full updates (Case II). The simulated results for the average error for different link delay rates are shown in Figure 7: Although the proactive event-driven strategy with full updates and the reactive strategy show the same mismatch probability, the average error for the reactive strategy is higher. This can be explained by the histograms of the error distribution in Fig. 8; the mismatch probability only indicates whether there is a deviation (height of bar at left hand). As for the mismatch probability, the periodic strategy leads to a higher average score error for small delays at the right end of Fig. 7, but it actually outperforms the reactive strategy for long delays.




Fig. 7. Average score error using the three update mechanisms.

Furthermore, the case of multiple (five) information providers increases the average score error for the reactive

strategy, while the impact on the other strategies seems less pronounced (not shown here). More detailed analysis of such scenarios with multiple information providers will be done in the future.




Fig. 8. Histogram of the absolute errors of the three updating strategies.

VI. SUMMARY AND OUTLOOK

This paper presents analytic models for the mismatch probability of the following different access strategies to dynamically changing, remote information elements: (1) Reactive access; (2) proactive, periodic access; (3) proactive event-driven access with the two sub-cases incremental and full update messages. The discussion of the analytic results focuses on information elements that never change back to a previous value, the change events form a Poisson process, and network delays are described by an exponentially distributed or constant random variables. The actual request process is irrelevant for the value of the mismatch probability but instead it has an impact on statistical estimation properties and on other performance metrics, such as generated network traffic. Simulation results are subsequently used to validate the analytic results and to provide quantitative results for the scenarios outside the scope of the analytic models treated in this paper. This includes multiple information elements provided by different entities in the network. Finally, the mismatch probability and its impact is discussed in a use-case of context sensitive service discovery in Personal Networks. The analytic models will be extended in future papers to include more general settings and in order to cover the scenario of multiple information sources. Furthermore, additional use-case scenarios such as link-state information in ad-hoc routing, binding tables in dynamic name services, and replication for resilience purposes will be investigated.

ACKNOWLEDGMENTS

The authors would like to thank Yaoda Liu (AAU), Homare Murakami (NICT, Japan), and Thibault Renier (AAU), for their helpful comments. Parts of this research is performed in the framework of the EU-FP6 projects MAGNET (www.magnet-ist.org) and HIDDENETS (www.hiddenets.aau.dk).

REFERENCES

- [1] Y. Liu and H.-P. Schwefel, "An enhanced optimized link state routing with delayed event driven link state updating," *16th IEEE Proceedings of Personal, Indoor and Mobile Radio Communications (PIMRC)*, September, 2005.
- [2] H. Murakami, R. Olsen, H. P. Schwefel, and R. Prasad, "Managing personal network specific addresses in naming schemes," *Proceedings of WPMC'05, Aalborg, Denmark*, 2005.
- [3] S. Panagiotakis and A. Alonistioti, "Intelligent service mediation for supporting advanced location and mobility aware service provisioning in reconfigurable mobile networks," *IEEE Wireless Communication*, vol. 9, no. 5, pp. 28–38, 2002.
- [4] M. Ghader, R. L. Olsen, M. G. Genet, and R. Tafazolli, "Service management platform for personal networks," *Proceedings of IST Summit 2005, Dresden*, 2005.
- [5] T. Renier, E. V. Matthiesen, H. P. Schwefel, and R. Prasad, "Inconsistency evaluation in a replicated IP-based call control center," *To be published in the proceedings of ISAS'06*, Helsinki, Finland, May 2006.
- [6] W. Chen, S. Toueg, and M. K. Aguilera, "On the quality of service of failure detectors," *IEEE Transactions on Computers*, vol. 51, no. 5, May 2002.
- [7] M. Božinovski, "Fault-tolerant platforms for IP-based session control systems," Ph.D. dissertation, Aalborg University, 2004.
- [8] M. B. Hansen, R. L. Olsen, and H.-P. Schwefel, "General scenarios of remote access strategies to information elements," *Submitted to Performance Evaluation*.
- [9] R. W. Wolff, *Stochastic Modeling and the Theory of Queues*. Upper Saddle River, NJ: Prentice Hall, 1989.
- [10] I. Niemegeers and S. H. de Groot, "From personal area networks to personal networks: A user oriented approach," *Kluwer Journal*, May 2002.

APPENDIX

A. Derivation of the mmPr for the reactive case

Assume \mathcal{E} is a stationary point process, then define the excess $Y = R_1$ and the age $U = -R_0$ and their distribution functions by $B(t) = P(R_1 \leq t)$ and $A(t) = P(-R_0 \leq t)$, [9]. The density functions are denoted by a and b . Furthermore, construct the stochastic process $E(t) = k$, $t \in [E_k, E_{k+1})$. Assume that event process \mathcal{E} is a Poisson process with intensity λ , then by stationarity we have the following probability of mismatch upon reception of the message for any request at time R_k

$$\begin{aligned} P(E(R_k + U_k + D_k) \neq E(R_k + D_k)) \\ &= P(E(D_k) \neq E(0)) \\ &= 1 - \int P(E(D_k) = E(0) | D_k = t) f_D(t) dt \\ &= 1 - \int e^{-\lambda t} f_D(t) dt = 1 - \mathcal{L}\{f_D\}(\lambda). \end{aligned}$$

As the mismatch probability does not depend on R_k we can define the mismatch probability in the reactive case to be

$$\text{mmPr}_{\text{react}} = 1 - \mathcal{L}\{f_D\}(\lambda).$$

B. Derivation of the mmPr for the pro-active, event-driven strategy with full updates

The probability of mismatch for the requesting time R_k is derived by conditioning on the situation that no event has happened in the interval $[t, R_k]$ and that the message is not delayed more than $R_k - t$ time units, consequently by stationarity

$$\begin{aligned} \text{mmPr}_{\text{proact,full}} &= 1 - \int_0^\infty P(\mathcal{D} \leq t | U = t) a(t) dt \\ &= 1 - \int_0^\infty P(\mathcal{D} \leq t) a(t) dt. \end{aligned}$$

Quantitative Analysis of access strategies to remote information in (wireless) network services

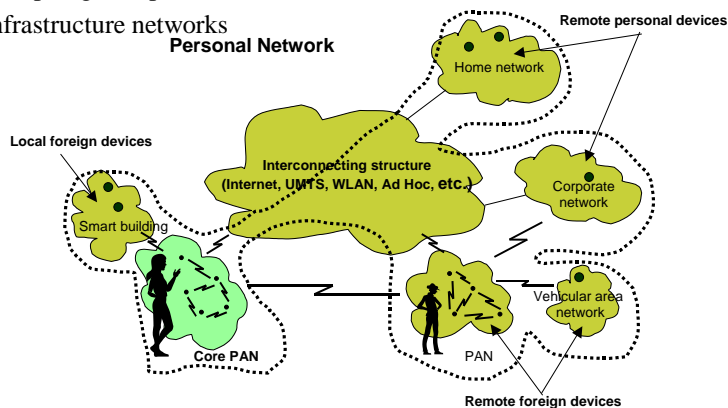
Rasmus Olsen, Martin B. Hansen, Hans-Peter Schwefel
{hps,rlo}@kom.aau.dk, mbh@math.aau.dk

Center for TeleInFrastruktur
Aalborg University
Denmark



Motivation: Context-sensitive Personal Networks


- Extension of PAN concept to Personal Networks
 - Logical networks, defined by appropriate security associations
 - Potential huge geographical/topological span
 - Consisting of ad-hoc and infrastructure networks
 - User centric definition



- Challenges
 - Network aspects: architecture, service discovery, naming, connectivity
 - User friendliness and efficiency via context-sensitive networking
 - Air interface/MAC: cross-layer optimization
 - Security: PAN/PN level authentication & authorization
 - Platforms: PAN devices, gateways, edge functionality, etc.



Context sensitive Service Discovery in MAGNET

- Context information (User context, network context, environment context)
 - potentially dynamically changing
 - obtained from remote entities, e.g. sensors for environment context
- Service Management Nodes (SMNs) need to evaluate the context information as
 - either reactively (on-demand) obtained
 - or pro-actively distributed by information provider
- Quality metrics that guide selection of access strategy
 - Network overhead (created data volume)
 - access delay
 - timeliness/correctness of information: mismatch probability



Other important examples of remote access


- Routing in (ad-hoc) networks
 - Example: source routing
 - Source needs to know links (or link-states) of remote nodes
 - Dynamically changing topologies
 - Pro-active (OLSR) vs reactive (AODV) as well as hybrid strategies (ZRP)
- Link-adaptation in cellular networks
 - Wireless link properties change over time (shadowing, multi-path propagation, interference)
 - Common Strategy: Base station adjusts downlink transmission parameters dynamically, e.g. coding scheme, transmission power, etc.
 - Based on current channel conditions
 - Which need to be measured in mobile station (in particular for FDD systems)



+ More examples: dynamic name services, data replication in fault-tolerant distributed systems, etc.


Different Access strategies

1. Reactive/on-demand




2. Pro-active (two cases: incremental vs. full updates)

2.1 Pro-active event-driven



2.2 Pro-active 'periodic'



Processes and Assumptions


Participating Processes

1. Event Process \mathcal{E} (when Poisson, with rate λ)
2. Request process \mathcal{R} (with rate μ)
3. Downstream Delay \mathcal{D} (when Poisson, with rate ν)
4. Upstream Delay \mathcal{U} (only in case of reactive approach)
5. Period of pro-active periodic approach \mathcal{P} (when Poisson, with rate τ)

Assumptions

- All processes are mutually independent
→ request process irrelevant for stationary mmPr
- $\mathcal{D}, \mathcal{U}, \mathcal{P}$ are independent, identically distributed (iid)
- Message re-ordering can occur in network, but requestor can put messages back in order (e.g. using sequence numbers)
- The information element never changes back to an earlier value, e.g. as for monotonous behavior


Analysis: Reactive Case



Assuming Poisson distributed Event process ($M|GI$ case)

$$\begin{aligned}
 &P(E(R_k + U_k + D_k) \neq E(R_k + D_k)) \\
 &= P(E(D_k) \neq E(0)) \\
 &= 1 - \int P(E(D_k) = E(0) | D_k = t) f_D(t) dt \\
 &= 1 - \int e^{-\lambda t} f_D(t) dt \\
 &= 1 - \mathcal{L}\{f_D\}(\lambda).
 \end{aligned}$$

Analysis: Proactive Event-Driven Case



Pro-active event-driven: Two cases

a) Incremental updates:

$$1 - \text{mmPr} = \Pr(\text{no message in transit at } R_k) = \Pr(E/D/\infty \text{ Queue is idle})$$

Assuming Poisson Event process: $\text{mmPr} = \Pr(M/GI/\infty \text{ queue is busy}) = 1 - \exp(-\lambda/\nu)$


b) Full updates

$$\text{mmPr} = \Pr(\text{last message still in transit}) = \Pr(E/D/1/1 \text{ Queue is busy})$$

equivalently:

$$\begin{aligned}
 \text{mmPr}_{\text{proact,full}} &= 1 - \int_0^\infty P(\mathcal{D} \leq t | U = t) a(t) dt \\
 &= 1 - \int_0^\infty P(\mathcal{D} \leq t) a(t) dt.
 \end{aligned}$$


Analysis: Proactive periodic case



Pro-active periodic:

- Assuming all processes are Poisson (M|.|M|M case)
- Representation as continuous-time Markov chain model
- Counting #msgs in transit
- Explicit solution:

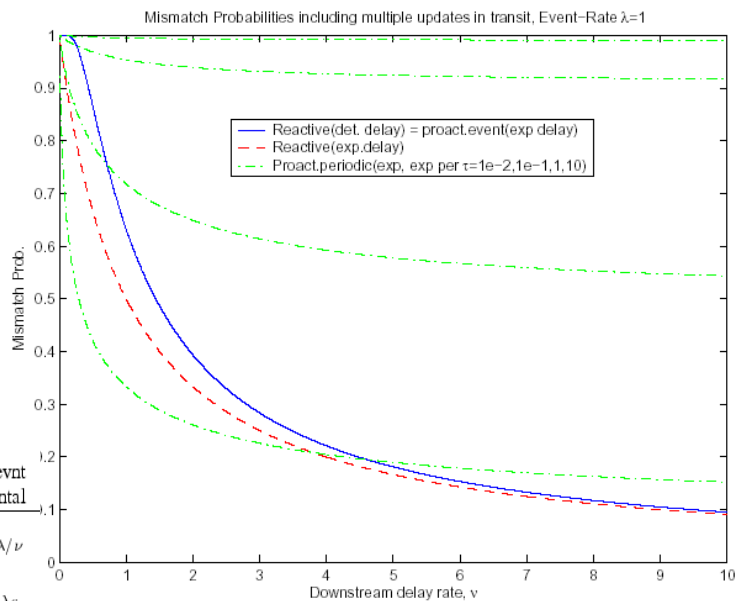
$$\pi_n = \frac{\prod_{i=1}^{n-2} \frac{\tau}{i\nu + \tau + \lambda}}{1 + \frac{\nu}{\lambda} \sum_{n=3}^{\infty} (n-1) \prod_{i=1}^{n-2} \frac{\tau}{i\nu + \tau + \lambda}}, \quad n \geq 3$$



- If #msgs in transit limited to max 1

$$\text{mmPr}_{\text{proact,periodic}}^{(\text{exp})}(\lambda, \nu, \tau) = \frac{\lambda [(\nu + \lambda)(2\tau + \nu + \lambda) + \tau^2]}{(\tau + \lambda)(\nu + \tau + \lambda)(\nu + \lambda)} \quad (3)$$

Quantitative results: Poisson Event Process




	reactive	proact. evt full update	proact. evt incremental
mmpr			
Exponential Delay	$\frac{\lambda}{\lambda + \nu}$	$\frac{\lambda}{\lambda + \nu}$	$1 - e^{-\lambda/\nu}$
Deterministic Delay	$1 - e^{-\lambda c}$	$1 - e^{-\lambda c}$	$1 - e^{-\lambda c}$
overhead access delay	$\mu(s_u + s_d)$ $E(U) + E(D) > 0$	$\lambda s_d^{(a)}$ 0	$\lambda s_d^{(i)}$ 0

Markov Event processes *(in progress)*

Scenario: information element may change back to a previous value later

- Modeling of event process as Markov model with generator Q
- Monotonous case is upper bound for mmPr



1. Reactive case

- Conditioning on downstream delay
- Instead of $\exp(-\lambda t)$ use probability that Event Markov process is in same state i at time 0 and t : $[\exp(-Qt)]_{ii}$


2.1 Pro-active event-driven case (full updates)

- Conditioning on downstream delay convoluted with backwards recurrence time of $\mathcal{E}+\mathcal{D}$
- Partial results based on using MAP/GI/ ∞ departure processes

2.2 Pro-active periodic case (exponential \mathcal{P}, \mathcal{D})

- Large Markov model (when multiple messages in progress are allowed)
- Numerical results hopefully soon...

Multiple Information Providers



Information Element may be vector (E_1, E_2, \dots, E_N)

- Each element provided by a different node
- Mismatch if ANY of the E_i are different then its true value at the provider at time of use

Re-active case:

- When $U = \text{const.}$ \rightarrow mismatch if any E_i changes in interval $\max(D_1, \dots, D_N)$
- Potential interesting approximations for large N (e.g. using Gumbel distribution)

Pro-active cases:

- Product solution due to independence assumptions

Summary and Outlook

- Scenario: Remote access to context (routing/channel state) information
- Quantitative Analysis of different strategies
 - Re-active
 - Pro-active event-drive (incremental/full) and periodic
- Set of results for rather general assumptions on $\mathcal{E}, \mathcal{D}, \mathcal{P}$

Long to-do list

- Markov Event Processes, multiple information providers (as explained)
- Estimation properties of mmPr
- More general cases: correlated delays, phase-type distributions
- Application to use-cases: service discovery, naming, routing, replication,...
- Analysis of hybrid strategies

... and much more...

1

Nonexponential Time Distributions in Biocatalytic Systems: Mass Service Replacing Mass Action

Peter W. Kühl¹ and Manfred Jobmann²

¹ Institute of Theoretical Biology, Schaulistr. 2, CH-4142 Münchenstein BL, Switzerland Peter-W.Kuehl@unibas.ch

² Department of Informatics, Technische Universität München, Boltzmannstr. 3, D-85748 Garching, Germany jobmann@in.tum.de

Summary. Formal kinetic methods to analyze biocatalytic systems are traditionally based on the law of mass action. This law involves the assumption that each molecular state has an exponentially distributed lifetime. We regard this assumption as unduly restrictive and advocate a more general, service theory based approach (termed mass service kinetics or briefly service kinetics). In service-theoretic terms biocatalysts are servers and their ligands are customers. The time intervals between arrivals of ligand molecules at special service loci (active or binding sites) as well as the service periods at these loci need not be exponentially distributed; rather, they may adopt any distribution (e.g., Erlangian, hyperexponential, variomorph). We exemplify the impact of nonexponential time distributions on a performance measure of wide interest: the steady-state throughput. For its computation we use matrix-analytic methods. Specifically, we show that nonexponential interarrival times convert hyperbolic mass action systems (whose characteristic is a hyperbolic velocity–concentration or dose–response curve) into nonhyperbolic mass service systems and that type and extent of their nonhyperbolicity are determined by type and parameters of the interarrival time distribution. Furthermore, we analyze the combined effect of a non-Poissonian arrival process and a waiting site near the catalyst’s active site on the throughput of the system. A major conclusion of our and other studies is that it is a questionable practice to use routinely and exclusively mass action kinetics for the interpretation and performance evaluation of biocatalytic systems.

Key words: Nonhyperbolic velocity–concentration curves, queueing theory, interarrival time, waiting space, steady-state throughput

1.1 Introduction

The kinetic behaviour of biocatalytic systems (i.e., systems containing biological macromolecules that catalyze chemical transformations or nutrient and

2 Peter W. Khl and Manfred Jobmann

metabolite transport or signal transduction) is traditionally described by mass action kinetics. The latter presupposes the validity of the law of mass action. Though usually formulated nonstochastically (by ordinary differential equations or algebraic expressions), the law of mass action clearly has stochastic roots and must be interpretable as the result of probabilistic phenomena at the molecular level by virtue of the fact that matter is not a continuum but consists of discrete units (atoms, molecules) with an intrinsically random behaviour. Already C. M. Goldberg (1836–1902) and P. Waage (1833–1900), the originators of the law of mass action, were aware of this and since then numerous authors, e.g., Rnyi [1], have investigated both the kinetic and equilibrium versions of the law of mass action in terms of probability and stochastic processes. It is beyond the scope of this article to compile and review the pertinent literature. We rather confine ourselves to mention two fundamental stochastic properties of the law of mass action: (i) ergodicity (= equality of time and ensemble averages) and (ii) Markovity (= lack of memory, complete forgetfulness). If a system is nonergodic and/or non-Markovian, it is by definition not a mass action system. Deviations from ergodicity, reported in experimental studies on conformational transitions in biological macromolecules (e.g., [2, 3]), are not a topic of the present paper. We rather focus on violation of the second above-mentioned property, the Markov property, by allowing for nonexponential interevent times at the molecular level. How to cope with biocatalytic systems containing non-Markovian elements and to what extent does the performance of such systems differ from that of conventional mass action systems? These are the main issues we address in the following.

1.2 The Service-Theoretic Approach

We propose to apply – instead of the traditional theory of mass action – the “theory of mass service” (also known under the names queueing theory and service theory) to biochemical kinetics for a number of conceptual and methodological reasons: (i) Biocatalysts are stochastic “servers” and their ligands are “customers” which are served at special “service loci” (active or binding sites). Various modes of operation (waiting or rejection, first come–first served, service interruptions, priority rules, etc.), being well known in man-made service systems, may also be encountered in biocatalytic systems. Furthermore, biocatalysts can build up chains and networks endowed with structural, dynamic and regulatory properties that are largely analogous to those of production lines or queueing networks of anthropic origin. (ii) For the quantitative analysis and performance evaluation of service systems and queueing networks, probability theorists have developed a rich repertoire of mathematical tools and methods. At least some of these are expected to be applicable and useful also in the nanoworld of enzymes, transporters and signal transducers. (iii) Nonexponential time distributions, being foreign to mass action systems, are quite common in mass service systems. It is primarily for

this reason why we advocate the application of the service-theoretic approach to biocatalytic systems and the build-up of a more general type of kinetics, termed – in contradistinction to mass action kinetics – *mass service kinetics* or briefly *service kinetics*. Biokinetics is thus put on a broader fundament, gains in flexibility and versatility and may lead to new interpretations of old phenomena. In the next section we use a very simple enzymic model in order to give an idea of the potentialities of the service-kinetic approach in a biochemical setting.

1.3 The Van Slyke–Cullen Scheme Modelled as a Service-Kinetic System

The simplest possible scheme of an enzymic reaction is



where E, S, ES and P stand for enzyme, substrate, enzyme-substrate complex and product, respectively. This scheme has been named after Van Slyke and Cullen [4] (VC) since these authors ignored in their mathematical analysis the (commonly included) reverse reaction $ES \rightarrow E + S$. We, too, ignore this reaction since its inclusion usually affects only the scale but not the shape of the $v(S)$ curve (v = reaction velocity; S = substrate concentration). We use the VC scheme as basic paradigm of a biocatalytic system.

The mass-action kinetic analysis of the VC scheme leads to only one type of $v(S)$ curves: the rectangular hyperbola. An entirely different situation is encountered when one subjects the VC scheme to a (mass-) service kinetic analysis: besides the familiar hyperbolic $v(S)$ curve a great variety of non-hyperbolic $v(S)$ curves can be obtained. This diversity is achieved by allowance for (i) nonexponential time distributions and/or (ii) a waiting room for substrate molecules which sit and wait there till clearance of the active site. In Figs. 1.1–1.3 some illustrative examples of nonhyperbolic $v(S)$ curves as consequences of (i) and/or (ii) are shown.

Let us have a closer look at the various situations giving rise to the $v(S)$ curves of Figs. 1.1–1.3.

Figure 1.1. When the interarrival times of the substrate molecules at the enzyme’s active centre are nonexponential, the resulting nonhyperbolic $v(S)$ curves are – compared with the classical hyperbola – either “lifted” or “depressed”, depending on whether the coefficient of variation of the interarrival time distribution is smaller or larger than 1 (see curves (a), (b) and (e), respectively). A special situation is represented by curve (d): here the stochastic arrival pattern varies with the arrival intensity (i.e., with the substrate concentration); specifically, its coefficient of variation gradually decreases with increasing substrate concentration in such a way that a sigmoidal $v(S)$ curve is generated. We regard arrival patterns with an intensity-dependent coefficient of variation (called by us “variomorphic” [5]) not as an exotic rarity but




Fig. 1.1. The effect of nonexponential interarrival time distributions on the throughput in a VC system without waiting facilities. $[S]$ and v stand for substrate concentration and reaction velocity, respectively. $[S]$ is given in units of the Michaelis constant K_M and v as fraction of the maximal throughput. The lower-case letters attached to the curves indicate various interarrival time distributions: (a) = deterministic, (b) = Erlangian of order 10, (c) = exponential, (d) = variomorphic with a coefficient of variation decreasing nonlinearly from 14.1 at $[S] = 0$ to about 0.7 at $[S] = 2.2$, and (e) = hyperexponential of order 2 with a coefficient of variation of 4. The computational techniques (based on matrix-analytic methods) used here as well as in Figs. 1.2 and 1.3 for generating the $v(S)$ curves are described in [5]. For readers familiar with the symbolic representation of service systems (explained in [5]) we add that curves (a), (b), (c), (d), and (e) represent the throughputs in the service systems $G/M/1/0$ and $G^{\text{vario}}/M/1/0$ with (a) $G = D$, (b) $G = E_{10}$, (c) $G = M$, (d) $G^{\text{vario}} = PH_4^{\text{vario}}$, and (e) $G = H_2$, respectively

as the rule in biocatalytic systems whose ligand arrival process is nonexponential. A more detailed description of nonexponential and variomorphic arrival processes and the resulting $v(S)$ or dose–response curves can be found in [5].

Figure 1.2. When near the enzyme’s active centre exist waiting spaces from which substrate molecules move to the active centre as soon as the latter is cleared, one usually obtains lifted nonhyperbolic $v(S)$ curves. Extent of lifting and nonhyperbolic “deformation” depends on (i) the number of available waiting spaces, (ii) the time needed for the substrate’s transit from the waiting space to the active centre, and – except the limiting case of an infinite number of waiting spaces – (iii) the probability distributions of interarrival and service




Fig. 1.2. The effect of waiting spaces on the throughput in a purely exponential VC system. The numbers attached to the curves indicate the number K of waiting spaces near the enzyme’s active centre. $[S]$ and v are normalized as in Fig. 1.1. For further details see the text. For readers familiar with the symbolic representation of service systems (explained in [5]) we add that the curves labelled 0 to 100 represent the throughputs in the service system $M/M/1/K$ with $K = 0, 1, 2, 3, 10,$ and $100,$ respectively. Note that $K = 100$ already approximates very well the case of $K = \infty$

times. In Fig. 1.2 we assumed that the mentioned transit time is negligibly small and both interarrival and service times are exponentially distributed. Under these conditions allowance for 1, 2, 3, 10, 100 or infinitely many waiting spaces yields the $v(S)$ curves shown. The larger the number of waiting spaces, the more approach the $v(S)$ curves the form of a ramp with a sharp transition at $[S] = K_M$, the Michaelis constant. The latter type of kinetics is phenomenologically equivalent to the so-called Blackman kinetics [6, 7] in biochemistry and physiology and the Holling type 1 “functional response” [8, 9] in ecology. The waiting room effect on the throughput in a purely exponential enzyme system was reported earlier by Trenkenshu [10, 11].

Figure 1.3. Here we illustrate the combined effect of two service-kinetic possibilities: a nonexponential arrival process and a waiting room with space for just one substrate molecule. Analogously to the purely exponential system of Fig. 1.2, also in this system allowance for one waiting space suffices to considerably increase the system’s throughput (see the $v(S)$ curves in pairs (b), (d) and (e) of Fig. 1.3). However this increase is not equal for systems with exponential and nonexponential interarrival time distributions – neither




Fig. 1.3. The effect of nonexponential interarrival time distributions and one waiting space on the throughput in the VC system. $[S]$ and v are normalized as in Fig. 1.1. In pairs (b), (d), and (e) the interarrival time distributions are identical with those used in Fig. 1.1 for generating the curves (b), (d), and (e), respectively. The upper (lower) curve in each pair of curves represents the throughput in a VC system with one (no) waiting space. When expressed symbolically [5], the upper and lower curves represent the throughputs in the service systems (b) $E_{10}/M/1/1$ and $E_{10}/M/1/0$, (d) $PH_4^{\text{vario}}/M/1/1$ and $PH_4^{\text{vario}}/M/1/0$, and (e) $H_2/M/1/1$ and $H_2/M/1/0$

absolutely nor as percentage value, regardless of being measured at selected substrate concentrations, e.g., at $[S] = K_M$, or when summed up over the whole substrate concentration range $(0, \infty)$.

The $v(S)$ curves shown in Figs. 1.1–1.3 all originate from a VC scheme with an exponential service time distribution. What do $v(S)$ curves look like when (case 1) solely the service time or (case 2) both the service and interarrival times are nonexponentially distributed? In answering this question we have to differentiate between VC systems with waiting facilities and those without them. *Answer in case 1.* In a VC system possessing no waiting room one always obtains one and the same curve shape: the classical hyperbola. In other words, this system is entirely insensitive to the stochastic character of the service process (provided the system is at steady state as we tacitly assume throughout this paper). However, this insensitivity is abolished when the enzyme molecule can offer a finite number of waiting spaces to substrate molecules arriving at an occupied active centre. In this case the VC system's throughput is the better (i.e., the $v(S)$ curves are the more lifted), the smaller

the coefficient of variation of the service time (data not shown). *Answer in case 2.* In a VC system devoid of waiting facilities the $v(S)$ curves often do not differ much when besides the interarrival time also the service time is nonexponential. However, special conditions – e.g., arrival and service processes being both Erlangian of high order – can give rise to quite unusual (undulatory or sawtooth-like) $v(S)$ curves which markedly differ from those in systems with an identical nonexponential arrival process but an exponential server; for illustrative examples see [5]. Finally, when comparing position and shape of $v(S)$ curves obtainable in VC systems equipped with a finite number of waiting spaces, one again finds differences between doubly (arrival *and* service) and singly (arrival *or* service) nonexponential systems. These differences become smaller with an increasing number of waiting spaces and totally vanish when this number goes to infinity. In the latter case the $v(S)$ curve adopts the ramp shape shown in Fig. 1.2 and it is irrelevant whether the interarrival and/or service time distributions are exponential or nonexponential (case of total insensitivity).

1.4 General Conclusions and Open Questions

Using the VC scheme as the basic paradigm of a biocatalytic system, we showed in the preceding section that a mass-service-theory inspired approach opens up prospects and possibilities that are closed to the traditional mass action approach. In particular we demonstrated that allowance for nonexponential time distributions and incorporation of waiting facilities for arriving substrate molecules yields a wide spectrum of $v(S)$ curves of which the classical hyperbola is merely a special case occurring under special, quite restrictive conditions. The mass service approach is of course also applicable to reaction schemes more complex than the VC scheme treated above. If in such schemes already conventional mass action kinetics leads to nonhyperbolic $v(S)$ or dose–response curves, the inclusion of mass-service specific elements may either reinforce or attenuate the system’s deviations from the hyperbolic response. However complicated a reaction scheme may be, it is clear that service-theoretic concepts and methods are apt to augment our understanding of biocatalytic systems. Time-honoured mass action kinetics should therefore be supplemented (if not replaced) by mass service kinetics. Other authors [12, 13, 14], using concepts and methodological approaches different from ours (e.g., a deterministic fractal approach [12] or Monte Carlo simulation algorithms [13, 14]), also emphasized that conventional mass action kinetics is often inadequate to describe biocatalytic systems, especially under *in vivo* conditions.

Finally, we would like to touch upon the question: What physical mechanisms can be envisaged to give rise to nonexponential time distributions at the molecular level? As far as the interarrival time distribution is concerned, we consider the following possibilities: (i) Substrate molecules are often products

8 Peter W. Khl and Manfred Jobmann

of a neighbouring enzyme in a metabolic chain. If these product molecules are released in a nonexponential manner and are – due to channelling – not (or not fully) thermally randomized, their arrival pattern at the next enzyme is also nonexponential. (ii) Coherent spiking of enzymic reactions in small volumes, as described by Mikhailov and Hess [15], may more or less derandomize the stream of arriving ligand molecules. (iii) The intracellular milieu which is characterized by extensive compartmentalization, macromolecular crowdedness, lacunarity and a non-Euclidian geometry may generate unusual, e.g. power-tailed, arrival patterns. (iv) The matrix process [16], i.e., the intramatrix migration of ligand molecules between the first (or last) contact point on the biocatalyst’s surface and the active centre, may not only change the intensity with which ligand streams arrive at the active centre but may also modify their stochastic pattern by various derandomizing mechanisms (e.g., overflow, regularizing filtering, alternative gated pathways). Direction and degree of nonexponentialization may often be dependent on the ligand concentration (e.g., exponential at low $[S]$ and increasingly hypo- or hyperexponential at rising $[S]$) and thus give rise to so-called variomorphic arrival patterns (see [5] and Sect. 1.3 above). Concerning the service time distribution, we mention two further feasible mechanisms of nonexponentialization: (v) The biocatalyst has to undergo a number of sequential conformational transitions or has to perform a number of sequential chemical or physical operations before the catalytic act (formation or fission of a covalent bond, emission of a signal, transportation of a nutrient across a membrane, etc.) can occur. Sequential multistage processes of this kind make the service time hypoexponential (Erlangian) though each individual stage is exponentially distributed. (vi) The biocatalyst’s conformational state required for the catalytic act may be reached by two or more parallel single- or multiple-step pathways which leads to a service time with hyperexponential or hyper-Erlang distribution.

The above-listed mechanisms of nonexponentialization in biocatalytic systems are largely hypothetical and direct experimental evidence supporting or invalidating them is scarce so far. However, it is expected that the rapidly advancing experimental techniques in single-molecule biochemistry (e.g., sophisticated optical techniques [17, 18] with high time resolution) will provide unambiguous data concerning nonexponential time distributions and enable open questions to be answered or reformulated.

References

1. Rnyi, A.: A discussion of chemical reactions using the theory of stochastic processes. In: Turn, P. (ed) Selected Papers of Alfrd Rnyi, Volume 1, 1948–1956, pp. 367–380. Akadmiai Kiad, Budapest (1976)
2. Shyamsunder, E.: Broken ergodicity in myoglobin. Ph. D. Thesis, University of Illinois at Urbana-Champaign (1986)
3. Wennmalm, S., Edman, L., Rigler, R.: Non-ergodic behaviour in conformational transitions of single DNA molecules. *Chem. Phys.*, **247**, 61–67 (1999)

4. Van Slyke, D.D., Cullen, G.E.: The mode of action of urease and of enzymes in general. *J. Biol. Chem.*, **19**, 141–180 (1914)
5. Köhl, P.W., Jobmann, M.: Receptor–agonist interactions in service-theoretic perspective. Effects of molecular timing on the shape of dose–response curves. *J. Recept. Sign. Transd.*, **26**, 1–34 (2006)
6. Blackman, F.F.: Optima and limiting factors. *Ann. Bot.*, **19**, 281–295 (1905)
7. Dabes, J.N., Finn, R.K., Wilke, C.R.: Equations for substrate-limited growth: the case for Blackman kinetics. *Biotechnol. Bioeng.*, **25**, 1159–1177 (1973)
8. Holling, C.S.: The components of predation as revealed by a small-mammal predation of the European pine sawfly. *Canad. Entomol.*, **91**, 293–320 (1959)
9. Taylor, R.J.: *Predation*. Chapman and Hall, New York and London (1984)
10. Trenkenschu, R.P.: Effect of metabolic bottleneck organization on kinetics of enzyme substrate conversion. *Mol. Biol.*, **22**, 1170–1177 (1988)
11. Trenkenschu, R.P.: Application of the queueing theory to biokinetics. In Shokin, Y.I. (ed) *Evolution Modeling and Kinetics: Collection of Scientific Works*, pp. 125–160. Nauka, Novosibirsk (1992) (in Russian)
12. Savageau, M.A.: Michaelis–Menten mechanism reconsidered: implications of fractal kinetics. *J. Theor. Biol.*, **176**, 115–124 (1995)
13. Berry, H.: Monte Carlo simulations of enzyme reactions in two dimensions: fractal kinetics and spatial segregation. *Biophys. J.*, **83**, 1891–1901 (2002)
14. Schnell, S., Turner, T.E.: Reaction kinetics in intracellular environments with macromolecular crowding: simulations and rate laws. *Progr. Biophys. Mol. Biol.*, **85**, 235–260 (2004)
15. Mikhailov, A., Hess, B.: Microscopic self-organization of enzymic reactions in small volumes. *J. Phys. Chem.*, **100**, 19059–19065 (1996)
16. Stetzkowski, F., Banerjee R., Marden, M.C., Beece, D.K., Bowne, S.F., Doster, W., Eisenstein, L., Frauenfelder, H., Reinisch, L., Shyamsunder, E., Jung, C.: Dynamics of dioxygen and carbon-monoxide binding to soybean leghemoglobin. *J. Biol. Chem.*, **260**, 8803–8809 (1985)
17. Min, W., English, B.P., Luo, G., Cherayil, B.J., Kou, S.C., Xie, X.S.: Fluctuating enzymes: lessons from single-molecule studies. *Acc. Chem. Res.*, **38**, 923–931 (2005)
18. Tinnefeld, P., Sauer, M.: Branching out of single-molecule fluorescence spectroscopy: challenges for chemistry and influence on biology. *Angew. Chem. Int. Ed.*, **44**, 2642–2671 (2005)

Effects of Molecular Timing on the Shape of Dose-Response Curves

Manfred Jobmann¹ Peter W. K uhl²

¹Department of Informatics
Technische Universit t M nchen
jobmann@in.tum.de

²Institute of Theoretical Biology
University of Basel

Colloquium in honor of Prof. Lester Lipsky

Chemical Processes in Biological Systems

Mass-kinetic Models


- Chemical description: $X_1 + X_2 \rightarrow X_3 + \dots$
- Mass-kinetic Model: Coupled System of differential equations:

$$\begin{aligned}\frac{d}{dt}[X_1] &= f_1([X_1], [X_2], \dots) \\ \frac{d}{dt}[X_2] &= f_2([X_1], [X_2], \dots)\end{aligned}$$

In equilibrium

$$0 = f_1([X_1], [X_2], \dots)$$

- Enzyme-Substrate and Receptor-Agonist reactions:




Solution of Mass-kinetic Models

Michaelis-Menten Equations

- ▶ Equilibrium "velocity", v , (better named) reaction rate, r :

$$v = \frac{v_{\max} \cdot [S]}{K_m + [S]} = \frac{v_{\max} \cdot \sigma}{1 + \sigma}, \quad \sigma = \frac{[S]}{K_m}$$

$$r = \frac{r_{\max} \cdot [A]}{A_{50} + [A]} = \frac{r_{\max} \cdot \alpha}{1 + \alpha}, \quad \alpha = \frac{[A]}{A_{50}}$$



Service-theoretic Models

Motivation:


- ▶ New view: Protein production in short bursts at random time intervals
- ▶ Mesoscopic view allows to keep track of few/every molecule.
- ▶ Then: \sqrt{n} law of statistical physics applies.
- ▶ Stochastic models instead of mass-kinetic models:

$A \xrightarrow{k} B$ is modeled as

$$(\#A, \#B) \xrightarrow{\#A \cdot k \cdot dt} (\#A - 1, \#B + 1)$$

Service-theoretic Model

Total Loss-System



Modeled: 1 binding site \implies Total Loss-System:
 $PH_k/M/1/0$:

$$\mathbf{Q} = \begin{pmatrix} \mathbf{B}_1 & \mathbf{B}_0 \\ \mathbf{A}_2 & \mathbf{A}'_1 \end{pmatrix}$$

Total Loss-System

Special case: $M/M/1/0$

$$\blacktriangleright \mathbf{Q} = \begin{pmatrix} -\lambda & \lambda \\ \mu & -\mu \end{pmatrix}$$

$$\blacktriangleright \pi_0 = \frac{\mu}{\lambda + \mu}, \quad \pi_1 = \frac{\lambda}{\lambda + \mu}$$

$$\blacktriangleright \text{Thruput: } \pi_1 \cdot \mu = \frac{\mu \cdot \rho}{1 + \rho} = \begin{cases} 0, & \rho = 0 \\ \frac{\mu}{2}, & \rho = 1 \\ \mu, & \rho \rightarrow \infty \end{cases}$$


Now associate $\rho \longrightarrow \sigma$, α and $\mu \longrightarrow v_{\max}$, r_{\max}

\implies Michaelis-Menten solution

Total Loss-System

Special cases: D/M/1/0, E_k /M/1/0:


D/M/1/0: Beněš [Beněš, 1959]



Total Loss-System


Special cases: H_2 /M/1/0

Models effects of *partial agonists* or *competitive inhibition*:



Total Loss-System

Case: PH₄/M/1/0



$$p = 0.99, \quad q_1 = 1.00, \quad q_2 = 1.00$$

$$\lambda_{11} = 10^9, 10^6, 10^3, 125.2, 37.34, 16.025, \dots$$


$$\lambda_{12} = 10^{1.3}$$

$$\lambda_{21} = 10^{-5}, 10^{-4}, 10^{-3}, 0.002, 0.00302, 0.0041, \dots$$

$$\lambda_{22} = 10^{1.3}$$


Total Loss-System

Case: PH₄/M/1/0




Total Loss-System

Special case: D/D/1/0




Total Loss-System

Case: D/E₁₀₀/1/0




Total Loss-System

Case: E₁₀₀/D/1/0



Total Loss-System





Case: E₅₀/E₅₀/1/0




Conclusions

- ▶ Classical mass-kinetic model \implies Michaelis-Menten equations
- ▶ Proposed here: service-theretic models
 - ▶ which include Michaelis-Menten's results but
 - ▶ allow for modeling additional phenomena of biocatalyst processes, e.g. effects of *partial agonists* and *competitive inhibition*.
- ▶ Further work:
 - ▶ multiple binding sites
 - ▶ temporary refractory states (by server with vacations?)

For Further Reading I


-  Lipsky, L. R. (1992).
Queueing Theory: a linear algebraic approach.
McMillan Publishing Company.
-  Lipsky, L. R. (2008).
Queueing Theory: a linear algebraic approach (2nd ed.).
Springer.
-  Beneš, V. E. (1959).
On Trunks with Negative Exponential Holding Times
Serving a Renewal Process.
The Bell System Technical Journal, 38:211–258.
-  Gaver, D. P., Jacobs, P. A., and Latouche, G. (1984).
Finite Birth-and-Death Models in Randomly Changing
Environments.
Advances in Applied Probability, 16:715–731.

For Further Reading II

-  Kuehl, P. W., Jobmann, M. (2006).
Receptor-agonist interactions in service-theoretic
perspective. Effects of molecular timing on the shape of
dose-response curves.
J. Recept. Sign. Transd.,

Doubly and Triply Excited States of Atomic Systems

The study of the spectra of multiply excited atoms and ions has been a very active area of research during the past decades. Most of these highly excited states are auto-ionizing states. Auto-ionization falls within the general class of phenomena called the Auger effect. In the Auger effect the system decays into a partition of its constituent parts. If the initial composite system is neutral, or positively charged, and its constituent decay particles are an electron and the residual ion, then the process is called auto-ionization. Auto-ionizing states are formed by scattering processes and by photoabsorption. For example we might have



where the auto-ionizing state He^{**} corresponds to a temporarily bound or resonant state of the compound system ($e^- + He^+$). Auto-ionizing states in atomic physics have been extensively studied both experimentally and theoretically from the early sixties. These states are conveniently described by using the projection operator formalism introduced by Feshbach¹. One of the first papers to use projection operators was by Lipsky and Russek² on the auto-ionizing states of Helium.

The Feshbach Formalism for two electron systems

Given a hydrogen like target and a second electron, the Schrodinger Equation is written in the usual way

$$(H - E)\Psi(\mathbf{r}_1, \mathbf{r}_2) = 0$$

where

$$H = -\frac{\hbar^2}{2m}\nabla_1^2 - \frac{\hbar^2}{2m}\nabla_2^2 - \frac{Ze^2}{r_1} - \frac{Ze^2}{r_2} + \frac{e^2}{r_{12}}$$

and Ψ is an eigenfunction of definite angular momentum, spin and parity.

The basic idea of the projection operator formalism is to introduce two operators P and Q which separate Ψ into scattering-like ($P\Psi$) and quadratically integrable ($Q\Psi$) parts:

$$\Psi = P\Psi + Q\Psi$$

and where

$$P + Q = 1, (\text{completeness}) \quad P^2 = P, \quad Q^2 = Q, (\text{idempotency}) \quad PQ = 0 (\text{orthogonality})$$

The following eigenvalue problem constitutes the heart of the projection operator formalism:

$$QHQ\phi_n = \epsilon_n \phi_n$$

Lipsky and Russek² showed that an upper bound to these eigenvalues ϵ_n could be obtained by calculating the matrix $H_{nm} = \langle u_n | H | u_m \rangle$ where

$$u_n = \frac{1}{\sqrt{2}} (R_{n_1\ell_1}(1)R_{n_2\ell_2}(2)L_{\ell_1\ell_2}^M(\hat{1}, \hat{2}) \pm R_{n_1\ell_1}(2)R_{n_2\ell_2}(1)L_{\ell_1\ell_2}^M(\hat{2}, \hat{1}))$$

and

$$L_{\ell_1 \ell_2}^M = \sum_m (\ell_1 \ell_2 L | m M - m) Y_{\ell_1}^m(\hat{1}) Y_{\ell_2}^{M-m}(\hat{2})$$

and where the $R_{n\ell}$ are the hydrogenic radial functions and where no $1sn\ell$ configuration is included. Of course only a finite number of basis functions $\{u_n\}$ could be included - hence the name truncated diagonalization method or TDM.

In the early seventies³ we redid these calculations again using a much larger basis set. We also computed the widths of these resonances from

$$\Gamma_n = 2\pi \left| \left\langle \psi(\epsilon_n) \left| \frac{1}{r_{12}} \right| \phi_n \right\rangle \right|^2$$

where

$$\psi(\epsilon) = C \left[R_{1s}(1) F_{\ell}(2) Y_{\ell}^m(\hat{2}) \pm R_{1s}(2) F_{\ell}(1) Y_{\ell}^m(\hat{1}) \right]$$

and where the $F_{\ell}(r)$ is the radial function of the autoionizing electron. This radial function was obtained from the well known static exchange approximation and involved the numerical solution of two coupled differential equations. From analyzing the wavefunctions $\{\phi_n\}$, energy levels ϵ_n , and the widths Γ_n we were able to arrange the levels into series converging on the $N=2$ threshold of the residual ion. We introduced our own notation for these series which is sometimes still used today although it is now superseded by a more sophisticated notation partially based on group theory. Our results^{4,5,6} on two electron systems H^- , He , Li^+ , Be^{2+} and B^{3+} became a standard reference for workers in this area. In actual fact, the energies ϵ_n which we calculated are not the ones that are actually observed. What is actually measured are

$$E_n = \epsilon_n + \Delta_n$$

where Δ_n is the shift in energy due to the coupling to the continuum and is given by

$$\Delta_n = \langle \phi_n | QHPGPHQ | \phi_n \rangle$$

where G is the Green's function $1/(H - E)$. This contribution to the energy is usually very small.

Three Electron Systems

The excited states of Lithium and Lithium-like ions are of fundamental interest because Lithium is the first element in the Periodic Table which has both an inner and an outer shell. The energy spectrum of lithium-like ions can be divided into three groups of levels. The first group, converging on the $(1s)^2 \ ^1S^e$ ground state of the corresponding two-electron residual ion consists of levels with a filled $1s$ shell together with one excited electron. They are the ground state and the singly excited states $(1s)^2 n\ell \ ^2L$. These states are well studied both experimentally and theoretically. The second group is composed of levels for which there is only one $1s$ electron. They are called the doubly excited states. The third spectral group, consists of states with configuration $n_1 \ell_1 n_2 \ell_2 n_3 \ell_3$ where n_1, n_2, n_3 are greater than, or equal to 2. These are described as the triply excited states.

Lester and Mohammed Ahmed⁷ in 1975 extended the TDM method to the 3-electron system and used the method to calculate the energies of some triply excited states of the Lithium isoelectronic sequence. The dimension of the matrices in

this case were large and computing power was limited at that time. In 1981, Lester and I tried again to do more extensive calculations, but due to computer limitations did not succeed. In 1991 much larger computers became available to us and it was possible for us to do a limited calculation of some doubly excited states^{8,9} of C^{3+} and also, together with Arnold Russek, we put the TDM method for 3-electron systems on a more secure mathematical foundation.

We have recently completed a comprehensive set of calculations^{10,11,12} on the triply excited $2\ell'2\ell''n\ell'''$ states of 3-electron systems using hydrogenic basis functions. This involved the diagonalization of 2000 dimensional matrices. We have reported and tabulated the energies, classification, effective quantum numbers, and configuration mixings of the triply excited ${}^{2,4}S^{e,o}$, ${}^{2,4}P^{e,o}$, ${}^{2,4}D^{e,o}$, ${}^{2,4}F^{e,o}$ states of lithium-like ions ($Z = 3, 10$). For each Z , there are a total of 72 Rydberg series converging on the doubly excited states of the residual 2-electron ion and we were able to find the lowest 10 to 12 levels for each series, giving a total of well over 5000 levels. The perturbation of a Rydberg series by an isolated state or by another Rydberg series was examined and discussed and our results compared with recent theoretical and experimental data.

Work in Progress

In all our calculation we have used hydrogenic basis functions. These functions do not form a complete set on the interval $[0, \infty]$. However, the associated Laguerre functions

$$R_{n\ell}(r) = \left(\frac{2}{\ell+1}\right)^{3/2} \sqrt{\frac{(n-\ell-1)!}{(n+\ell+1)!}} \left(\frac{2r}{\ell+1}\right)^\ell \exp\left(-\frac{r}{\ell+1}\right) L_{n-\ell-1}^{2\ell+2}\left(\frac{2r}{\ell+1}\right)$$

where L_β^α is the associated Laguerre polynomial, constitute a complete set of discrete orthonormal functions. Functions of this type are well known to give better results than hydrogenic functions for the lowest one or two levels of any symmetry, since they include (in a random way) pieces of the hydrogenic continuum functions. We have used these functions and the complex rotation method to calculate the energies and widths of the doubly excited states of 2-electron systems. Briefly, if T and V are the (real valued) matrices of the kinetic and potential energies, then the (complex symmetric) matrix of the rotated Hamiltonian H_θ is given by

$$H_\theta = e^{-2i\theta}T + e^{-i\theta}V$$

The eigenvalues of H_θ then include the complex resonance eigenvalues that are independent of the parameter θ in the limit that the basis is complete and we have for these eigenvalues

$$E = E_n - i\frac{\Gamma_n}{2}$$

This method is very attractive since it automatically gives us the width Γ and the shift due to the coupling to the continuum. We have found excellent agreement with Ho for the lowest ${}^1S^e$ doubly excited states of the Helium isoelectronic sequence using 980 basis functions.

References:

1. H.Feshbach, *Annals of Physics*, **5**,357(1958); **19**,287,1962
2. L. Lipsky and A. Russek. Auto-Ionizing States in Helium. *Physical Review*,

- 142**, pp.59-71,1965
3. L. Lipsky and M.J. Conneely. Calculation of Autoionization Widths of Two-Electron Systems. *Physical Review. A* **14**, pp.2193-2205, 1976.
 4. M.J. Conneely and L.Lipsky. Widths and Configuration Mixings of Two-Electron Systems Below the N=2 Threshold. *Journal of Physics B: At. Mol. Phys.* **11**,pp. 4135-4154,1978
 5. L. Lipsky, R. Anania and M.J. Conneely. Energy Levels and Classifications of Doubly Excited States in Two-Electron Systems(Z=1,2,3,4,5) Below the N=2 and N=3 Thresholds *Atomic Data and Nuclear Data Tables* **20**, pp.127-141, 1977.
 6. M.J. Conneely and L.Lipsky. Tables of Widths of Doubly Excited States in Two Electron Systems Below the N=2 Threshold. *British Library Supplementary Publication Scheme*, 1978
 7. Mohammed Ahmed and Lester Lipsky. Triply excited states of three-electron atoms systems. *Physical Review A* **12**,4, pp. 1176-1196,1975.
 8. R. Bruch, H. Wang, D. Schneider, L. Lipsky and M.J. Conneely. Doubly Excited States of 3-electron Systems(Z=2 to 10) below the $(1s2l)LS\pi$ Threshold, *Zeitschrift Fur Physik D*, **21**, pp.S227-229,1991.
 9. M.J. Conneely, L.Lipsky and A. Russek. Three-Electron Systems with Inner Shell Vacancies. *Physical Review A* **46**,7, pp. 4012-4025,1992
 10. M.J. Conneely and Lester Lipsky. Hollow States of Lithium *Physical Review A* **61**, 032506, 2000
 11. M.J. Conneely and Lester Lipsky. Energy Levels and Classifications of Triply Excited States of Li, Be⁺, B²⁺, F⁶⁺, and C³⁺. *Atomic Data and Nuclear Data Tables* **81**,No.1 pp115-190, 2002
 12. M.J. Conneely and Lester Lipsky. Energy Levels and Classifications of Triply Excited States of N⁴⁺, O⁵⁺, F⁶⁺, and Ne⁷⁺. *Atomic Data and Nuclear Data Tables* **83**,No.1,pp.115-190, 2004
 13. Y.K. Ho.Complex-coordinate calculations for doubly excited states of two-electron atoms.*Physical Review A*, **23**,pp. 3137-2149, 1981

IMPROVING SIMULATION ACCURACY THROUGH THE USE OF SYNTHETIC ALIGNMENT INTERVALS¹

Jeffrey P. Buzen
buzen@post.harvard.edu

Monte Carlo simulations are widely used in the analysis of stochastic models and their associated steady state distributions. This paper introduces a new procedure for improving the accuracy of such simulations. The approach is based on constructing synthetic alignment intervals that are appended to the output of the original simulations, creating extended simulations whose output conforms to certain mathematical relationships. Satisfying these relationships is shown to be sufficient to guarantee that the underlying steady state distributions have been computed accurately.

1. Introduction

Variability is present in the behavior of many real world systems. It is often convenient to regard this variability as the physical manifestation of some underlying stochastic process. This assumption leads to the creation of stochastic models of system behavior. These models can then be evaluated through analytic techniques, numerical methods, or Monte Carlo simulation.

In many cases, analysts are interested in determining the steady state distribution of the underlying stochastic process. When Monte Carlo simulation is being used for this purpose, the Ergodic Theorem ensures - with probability one - that the output of the simulation will provide an accurate characterization of the underlying steady state distribution if the simulation runs “long enough”, and if the random number generator is “good enough”. The first few sections of this paper present a new procedure for testing the output of a simulation to determine if an accurate characterization of the underlying steady state distribution has been obtained.

The tests are based on the equivalence of two different methods for deriving the equations that characterize the steady state distribution. The first employs the classical approach of setting the time derivative of the transient distribution equal to zero. The second method, which relies solely upon testable assumptions, is based on a direct analysis of the observable quantities that simulation programs actually generate. Both methods are shown to produce identical results. This equivalence provides the rationale for the new testing procedure. A simple example is presented to illustrate the issues involved.

The discussion then turns to cases where the output of the simulation does not satisfy the conditions sufficient to guarantee accuracy. It is shown that accuracy can be improved in such cases by appending specially constructed “alignment intervals” to the end of the original output trajectories.

¹ Preprint: CMG 2007 Conference Proceedings, San Diego, CA, Dec. 2-7, 2007 www.cmg.org

Since system behavior during these synthetic alignment intervals is driven by external calculations rather than calls on a random number generator, the essential stochastic nature of the simulation is clearly compromised. Nevertheless, the accuracy of the simulation can actually be improved. This apparent paradox is explained by introducing the distinction between distributional and trans-distributional properties. Synthetic alignment intervals are shown to preserve distributional properties (dependent solely upon steady state distributions) while at the same time altering other properties (such as busy period distributions) that are trans-distributional in nature.

2. Example – Simple Queuing Network

Begin by considering the simple queuing network shown in Figure 1. There are two queues, each served by a single server. A total of three customers circulate around the network, cycling between one server and the other.

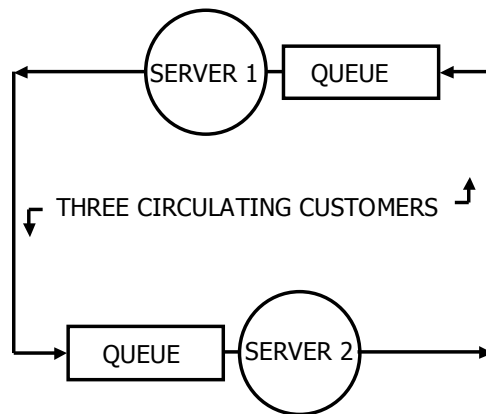


Figure 1: Simple Queuing Network

Assume that the service times at servers 1 and 2 are determined by sampling from exponentially distributed random variables with means of 2 seconds and 4 seconds respectively. The queuing network in Figure 1 can thus be regarded as the realization of a stochastic process. The structure of this process is identical to that of an M/M/1/3 queue.

Let $p(n)$ be the steady state probability that the number of customers at server 1 is equal to n (where $n = 0, 1, 2, 3$). Analytic expressions for $p(n)$ are well known to queuing theorists. However, suppose for purposes of this example that $p(n)$ is not known but is instead being evaluated through a Monte Carlo simulation.

Figure 2 depicts one possible trajectory that such a simulation might generate. The trajectory is exactly 30 seconds in duration. This is, of course, too short an interval to obtain reliable estimates of the underlying steady state distribution; however, it is entirely sufficient for purposes of this discussion.




Figure 2: Possible Simulation Output

For any simulation interval, the steady state probability $p(n)$ is estimated by $P(n)$, the fraction of the interval for which the number of customers at server 1 is equal to n . In the case of Figure 2, the values of $P(n)$ are computed as follows.

To compute $P(0)$, note that there are zero customers at server 1 for a total of $2 + 4 + 6 + 4 = 16$ seconds. Thus $P(0) = 16/30$.

Similarly, there is exactly one customer at server 1 for a total of $1 + 1 + 2 + 1 + 1 + 2 = 8$ seconds. Thus $P(1) = 8/30$.

Likewise, there are two customers at server 1 for a total of $2 + 1 + 1 = 4$ seconds, implying $P(2) = 4/30$.

Finally, there are three customers at server 1 for a total of 2 seconds, which implies $P(3) = 2/30$.

As already noted, expressions for $p(n)$, the steady state distribution of the underlying stochastic process, can also be derived analytically in this simple case. Using the standard notational conventions of queuing theory, assume that service times at server 1 are exponentially distributed with mean $1/\mu$ and that service times at server 2 are exponentially distributed with mean $1/\lambda$. Then the values of $p(n)$ can be expressed as well known functions of the ratio λ/μ .

$$p(0) = 1/[1 + \lambda/\mu + (\lambda/\mu)^2 + (\lambda/\mu)^3] \quad (1)$$

$$p(1) = (\lambda/\mu) p(0) \quad (2)$$

$$p(2) = (\lambda/\mu)^2 p(0) \quad (3)$$

$$p(3) = (\lambda/\mu)^3 p(0) \quad (4)$$

Note that equations (1) – (4) characterize the steady state distribution of the stochastic process associated with Figure 1. There is no guarantee that these equations will apply to the output of a Monte Carlo simulation, especially when the duration of the simulation interval is very short.

In this example, the mean service time at server 1 is 2 seconds, so $1/\mu = 2$. Also, the mean service time at server 2 is 4 seconds, so $1/\lambda = 4$. The ratio λ/μ is thus equal to $1/2$. Replacing λ/μ by $1/2$ in equations (1) – (4) yields the following steady state distribution:

$$\begin{aligned} p(0) &= 1/[1 + (1/2) + (1/2)^2 + (1/2)^3] \\ &= 8/15 \end{aligned}$$

$$\begin{aligned}
 &= 16/30 \\
 p(1) &= (1/2) \times 8/15 \\
 &= 4/15 \\
 &= 8/30 \\
 p(2) &= (1/2)^2 \times 8/15 \\
 &= 2/15 \\
 &= 4/30 \\
 p(3) &= (1/2)^3 \times 8/15 \\
 &= 1/15 \\
 &= 2/30
 \end{aligned}$$

As this computation demonstrates, the values of $P(n)$ that were generated by the Monte Carlo simulation (Figure 2) are identical to the values of $p(n)$ that were obtained from equations (1) – (4). This perfect alignment between observed and theoretical results is, of course, the goal of Monte Carlo simulation. However, in this case the 30 second simulation interval seems much too short to expect this goal to be satisfied.

Is the observed alignment between the simulator’s output and the analytic solution merely the consequence of a highly specialized and artfully constructed trajectory, or is it the result of fundamental principles that can be generalized to a broad class of possible trajectories? As this paper will demonstrate, fundamental principles are indeed involved, and these principles extend well beyond the simple example presented above.

3. Classical Analysis Approach

To proceed with the analysis, it is helpful to return to first principles and review the classical approach for deriving the steady state distribution of any continuous time Markov process.

The first step is to identify the states of the Markov process and the “permissible” transitions between these states. Essentially, a transition is “permissible” if the specification of the Markov process allows it to occur with a probability that is greater than zero.

In the example illustrated in Figure 1, the state of the process will simply be defined as an integer that represents the number of customers at server 1. Thus, state can be equal to 0, 1, 2 or 3.

Permissible transitions occur when a single customer completes service at one of the servers and proceeds to the queue at the other server. This results in a state change of plus or minus one. The four possible states and the six permissible transitions are illustrated in Figure 3.




Figure 3: State Transition Diagram for Figure 1

The state of this Markov process at any time t can be represented by a random variable with distribution $p(n,t)$. In other words, $p(n,t)$ is the probability that the Markov process is in state n at time t . In addition to depending on t , $p(n,t)$ also depends on the initial state of the process at time $t = 0$.

Intuitively, it is reasonable to believe that the dependence of $p(n,t)$ on the initial state will become negligible as more and more time passes. In the limit as t approaches infinity, $p(n,t)$ will then become completely independent of t . Stochastic processes that conform to this intuitive notion are said to become stationary or reach steady state.

From a mathematical perspective, when $p(n,t)$ becomes independent of t , the derivative of $p(n,t)$ with respect to t becomes equal to zero. As shown in standard texts on queuing theory, setting these derivatives equal to zero (for $n = 0, 1, 2$ and 3) yields the following equations:

$$\lambda p(0) = \mu p(1) \quad (5)$$

$$[\mu + \lambda]p(1) = \lambda p(0) + \mu p(2) \quad (6)$$

$$[\mu + \lambda]p(2) = \lambda p(1) + \mu p(3) \quad (7)$$

$$\mu p(3) = \lambda p(2) \quad (8)$$

Equations (5) – (8) are special cases of the classical Chapman-Kolmogorov equations and are commonly referred to as global balance equations.

To obtain a unique solution to these equations, it is necessary to account for one additional constraint: the sum of all possible state probabilities must be equal to one. Since there are only four possible states (0, 1, 2 and 3), this constraint implies

$$p(0) + p(1) + p(2) + p(3) = 1 \quad (9)$$

Solving equations (5) – (9) for $p(n)$ yields the well known queuing formulae in equations (1) – (4).

4. Alternative Analysis Approach

The analysis outlined in the preceding section is concerned with the mathematical properties of the stochastic process that provided the basis for the Monte Carlo simulation. The alternative approach presented here deals directly with observable properties of the state space trajectory that is actually generated by the simulation. The goal is to derive an expression for $P(n)$ without appealing to the notion of an underlying steady state stochastic process.

To begin, note that the state transition diagram shown in Figure 3 can be adapted to track the status of the associated simulation over time. Simply add a token to the diagram and assume that the token moves from circle to circle via one of the arrows whenever the number of customers at server 1 changes. The position of the token at any instant represents the current state of the system.

It should be immediately apparent that, during the course of any finite simulation interval, the number of transitions that the token makes out of a given state will be equal to the number of transitions it makes into that state. The only exceptions are the initial state, which has one extra transition out (just after the start of the simulation) and the final state, which has one extra transition in (just before the end of the simulation). If the initial and final states are identical, these two extra transitions will balance one another, implying that the number of transitions out is equal to the number of transitions in for all possible states. This condition is referred to as “flow balance”.

5. Implications of Flow Balance

To analyze the mathematical consequences of the flow balance assumption, let $C(n)$ denote the number of times during the execution of a simulation that a customer completes service at server 1 while the system is in state n . Similarly, let $A(n)$ denote the number of times that a customer completes service at server 2 while the system is in state n .

For state 0, transitions out occur only as a result of service completions at server 2 (causing a transition from state 0 to state 1). The number of times these transitions occur is equal to $A(0)$. Similarly, the only transitions into state 0 occur while the system is in state 1 and the customer being processed at server 1 completes its service. The number of times these transitions occur is equal to $C(1)$. The flow balance condition thus implies:

$$A(0) = C(1) \tag{10}$$

For state 1, there are two possible transitions in (a completion at server 2 while in state 0 or a completion at server 1 while in state 2). Similarly, there are two possible transitions out (a completion at server 1 while in state 1 or a completion at server 2 while in state 1). For state 1, flow balance implies:

$$A(1) + C(1) = A(0) + C(2) \tag{11}$$

Similar considerations regarding state 2 imply:

$$A(2) + C(2) = A(1) + C(3) \tag{12}$$

State 3 is similar to state 0:

$$C(3) = A(2) \tag{13}$$

The next step is to transform these four flow balance equations so they are expressed in terms of transition rates rather than raw counts. Suppose that $T(n)$ is the amount of time the system spends in state n during the simulation interval. Then $T = T(0) + T(1) + T(2) + T(3)$ must be equal to the total length of the simulation interval. Also, $T(n)/T$ must be equal to the proportion of time (during the simulation interval) that the system spends in state n . As already noted, this quantity is denoted by $P(n)$:

$$P(n) = \frac{T(n)}{T} \text{ for } n=0,1,2,3 \quad (14)$$

$$\text{where } T = T(0) + T(1) + T(2) + T(3) \quad (15)$$

Equation (14) can now be combined with equations (10) – (13) to show:

$$\frac{A(0)}{T(0)} \times P(0) = \frac{C(1)}{T(1)} \times P(1) \quad (16)$$

$$\left[\frac{A(1)}{T(1)} + \frac{C(1)}{T(1)} \right] \times P(1) = \frac{A(0)}{T(0)} \times P(0) + \frac{C(2)}{T(2)} \times P(2) \quad (17)$$

$$\left[\frac{A(2)}{T(2)} + \frac{C(2)}{T(2)} \right] \times P(2) = \frac{A(1)}{T(1)} \times P(1) + \frac{C(3)}{T(3)} \times P(3) \quad (18)$$

$$\frac{C(3)}{T(3)} \times P(3) = \frac{A(2)}{T(2)} \times P(2) \quad (19)$$

Note that the constraints on $P(0)$, $P(1)$, $P(2)$ and $P(3)$ that are expressed by equations (16) – (19) are valid for any trajectory generated by any simulation that conforms to the state transition diagram illustrated in Figure 3. The only requirement is that the trajectory must satisfy flow balance: that is, the initial and final states must be the same. No distributional assumptions of any type are required, and the concept of steady state (in the stochastic sense) does not arise.

6. Homogeneity Assumptions

The next step is to simplify equations (16) – (19) by introducing additional assumptions regarding the values of $A(n)/T(n)$ and $C(n)/T(n)$. Note that $C(n)$ is equal to the number of requests completed by server 1 while the system is in state n . In addition, $T(n)$ is the amount of time the system spends in state n . Thus, $C(n)/T(n)$ is the rate at which requests are completed by server 1 while the system is in state n .

Since no requests can be completed by server 1 while the system is in state 0, $C(0)$ must equal zero. Assuming $T(0) \neq 0$, this implies that $C(0)/T(0) = 0$.

Now consider $C(n)/T(n)$ for $n=1,2,3$. If the rate at which server 1 completes requests is independent of the number of customers at that server, it would seem reasonable to assume that all three values of $C(n)/T(n)$ are equal to each other:

$$\frac{C(1)}{T(1)} = \frac{C(2)}{T(2)} = \frac{C(3)}{T(3)} \quad (20)$$

Equation (20) is referred to as the assumption of homogeneous service rates at server 1. Note that this assumption concerns only the observable properties of the trajectory in Figure 2. It can be verified directly without assuming that the trajectory has been generated by a Monte Carlo simulation. However, as discussed in the Appendix, homogeneity is a reasonable assumption to make when analyzing the output of a Monte Carlo simulation of an M/M/1/3 queue.

Equation (20) can be rearranged to show:

$$\frac{C(n)}{T(n)} = \frac{C(1)+C(2)+C(3)}{T(1)+T(2)+T(3)} \text{ for } n=1,2,3 \quad (21)$$

Note that $C(1)+C(2)+C(3)$ is equal to the total number of requests completed by server 1. Also, $T(1)+T(2)+T(3)$ is the total amount of time that server 1 is actively processing requests. Thus, the expression on the right side of equation (21) is simply the request completion rate calculated over all time for which server 1 is active. Homogeneity is thus equivalent to assuming that the conditional completion rates on the left side of equation (21) are all equal to the unconditional completion rate on the right.

In the underlying M/M/1/3 stochastic process, the parameter μ represents the unconditional completion rate at server 1. Note that μ cannot be observed directly, but its estimate $\bar{\mu}$ can be obtained by running a Monte Carlo simulation and computing the expression on the right side of equation (21). In other words,

$$\bar{\mu} = \frac{C(1)+C(2)+C(3)}{T(1)+T(2)+T(3)} \quad (22)$$

Equations (21) and (22) then imply that the assumption of homogeneous service rates at server 1 is equivalent to assuming:

$$\frac{C(n)}{T(n)} = \bar{\mu} \text{ for } n=1,2,3 \quad (23)$$

Similar assumptions will now be introduced regarding the values of $A(n)/T(n)$ associated with server 2. Since n is the number of customers at server 1, and since there are a total of 3 customers in the network in Figure 1, there are $3-n$ customers at server 2 when the system is state n . Thus, $A(3)/T(3)=0$.

In this case, the assumption of homogeneous service rates at server 2 implies:

$$\frac{A(0)}{T(0)} = \frac{A(1)}{T(1)} = \frac{A(2)}{T(2)} \quad (24)$$

Equation (24) implies:

$$\frac{A(n)}{T(n)} = \frac{A(0) + A(1) + A(2)}{T(0) + T(1) + T(2)} \text{ for } n=0,1,2 \quad (25)$$

The right side of equation (25) is the unconditional completion rate at server 2, computed over all time for which server 2 is active. In terms of the underlying stochastic process, this measurable quantity is an estimate of the parameter λ in the underlying stochastic process. The symbol $\bar{\lambda}$ will be used to represent this estimate:

$$\bar{\lambda} = \frac{A(0) + A(1) + A(2)}{T(0) + T(1) + T(2)} \quad (26)$$

Equations (25) and (26) then imply that the assumption of homogeneous service rates at server 2 is equivalent to assuming:

$$\frac{A(n)}{T(n)} = \bar{\lambda} \text{ for } n=0,1,2 \quad (27)$$

Equations (16) - (19) can now be simplified by combining them with the homogeneity assumptions in equations (23) and (27). Simple substitution yields equations (28) – (31) below:

$$\bar{\lambda} P(0) = \bar{\mu} P(1) \quad (28)$$

$$[\bar{\mu} + \bar{\lambda}]P(1) = \bar{\mu} P(0) + \bar{\lambda} P(2) \quad (29)$$

$$[\bar{\mu} + \bar{\lambda}]P(2) = \bar{\mu} P(3) + \bar{\lambda} P(3) \quad (30)$$

$$\bar{\mu} P(3) = \bar{\lambda} P(2) \quad (31)$$

Note also that equations (14) and (15) imply that

$$P(0) + P(1) + P(2) + P(3) = 1 \quad (32)$$

which is the analog of equation (9).

Solving equations (5) – (9) for $p(n)$ will always yield equations (1) – (4). Similarly, it follows immediately that solving equations (28) – (32) for $P(n)$ will always yield equations (33) – (36).

$$P(0) = 1/[1 + \bar{\lambda}/\bar{\mu} + (\bar{\lambda}/\bar{\mu})^2 + (\bar{\lambda}/\bar{\mu})^3] \quad (33)$$

$$P(1) = (\bar{\lambda}/\bar{\mu})P(0) \quad (34)$$

$$P(2) = (\bar{\lambda}/\bar{\mu})^2 P(0) \quad (35)$$

$$P(3) = (\bar{\lambda}/\bar{\mu})^3 P(0) \quad (36)$$

Clearly, equations (33) – (36) have exactly the same form as equations (1) – (4) except that the abstract stochastic quantities $p(n)$, λ and μ have all been replaced by their observable counterparts $P(n)$, $\bar{\lambda}$ and $\bar{\mu}$ respectively.

To summarize the main point, equations (1) – (4) and (33) – (36) have just been derived under two very different sets of assumptions. First, it was shown that equations (1) – (4) are satisfied by the steady state distribution of the Markov process associated with Figure 1. Then, it was shown that equations (33) – (36) are satisfied by the output of any Monte Carlo simulation that is based on Figure 1, provided the output satisfies the directly observable conditions of flow balance and homogeneous completion rates at servers 1 and 2.

7. Application to Simulation Output

This conclusion leads to a simple test for simulation accuracy: if the output trajectory generated by a simulation of the system shown in Figure 1 satisfies the assumptions of flow balance and homogeneous service rates at servers 1 and 2, the values of $P(n)$ generated by the simulation are guaranteed to be accurate. In other words, the values of $P(n)$ generated by the simulation are guaranteed to equal the corresponding values of $p(n)$ obtained by evaluating equations (1) – (4) with μ set equal to $\bar{\mu}$ and λ set equal to $\bar{\lambda}$.

The flow balance assumption is clearly satisfied by the trajectory shown in Figure 2 (and replicated at the top of Figure 4) since the initial and final states are both equal to zero.

Next consider the homogeneity assumption for server 1. Note that the values of $T(0)$, $T(1)$, $T(2)$ and $T(3)$ are 16, 8, 4 and 2 respectively. Also, the values of $C(1)$, $C(2)$, and $C(3)$ are 4, 2 and 1 respectively.

Thus, the service completion rates for server 1 are:

$$C(1)/T(1) = 4/8 = 0.5$$

$$C(2)/T(2) = 2/4 = 0.5$$

$$C(3)/T(3) = 1/2 = 0.5$$

Since $C(n)/T(n)$ has the same value for $n=1, 2$ and 3 , homogeneity is satisfied at server 1.

To analyze server 2, note that the values of $A(0)$, $A(1)$ and $A(2)$ are equal to 4, 2 and 1 respectively. Thus, the service completion rates for server 2 are:

$$A(0)/T(0) = 4/16 = 0.25$$

$$A(1)/T(1) = 2/8 = 0.25$$

$$A(2)/T(2) = 1/4 = 0.25$$

Since $A(n)/T(n)$ has the same value for $n=0, 1$ and 2 , homogeneity is also satisfied at server 2.

Since the trajectory shown in Figure 2 (and the top of Figure 4) satisfies the required tests, the associated values of $P(n)$ are guaranteed to be accurate. This explains the surprising result discovered in Section 2.




Figure 4: Two Possible Trajectories

Of course, it is entirely possible for a Monte Carlo simulation to generate a trajectory that does not satisfy all the required tests. The lower trajectory in Figure 4 provides such an example. The two points of difference between the upper and lower trajectories are marked with vertical arrows

The lower trajectory clearly satisfies the assumption of flow balance; however, servers 1 and 2 fail to satisfy homogeneity as the following calculations show:

For server 1: $C(1)/T(1) = 4/7$

$$C(2)/T(2) = 2/5$$

$$C(3)/T(3) = 1/2$$

For server 2: $A(0)/T(0) = 4/16$

$$A(1)/T(1) = 2/7$$

$$A(2)/T(2) = 1/5$$

This failure to satisfy homogeneity should have an adverse impact on the accuracy of the simulation. To verify this, note that the values of $P(n)$ generated by the simulation are:

$$P(0) = 16/30$$

$$P(1) = 7/30$$

$$P(2) = 5/30$$

$$P(3) = 2/30$$

The “correct” values of $P(n)$ are obtained by determining the values of $\bar{\mu}$ and $\bar{\lambda}$ from the output trajectory and then applying equations (33) – (36).

$$\bar{\mu} = (4+2+1)/(7+5+2) = 0.5$$

$$\bar{\lambda} = (4+2+1)/(16+7+5) = 0.25$$

For the lower trajectory, $\bar{\lambda}/\bar{\mu}$ is again equal to $1/2$. Substituting this value into equations (33) – (36) yields the same “analytically correct” distribution seen earlier.

$$P(0) = 16/30$$

$$P(1) = 8/30$$

$$P(2) = 4/30$$

$$P(3) = 2/30$$

Note that the simulation has correctly determined the values of $P(0)$ and $P(3)$, but has generated incorrect values for $P(1)$ and $P(2)$.

8. Generalizations and Extensions

The discussion has concentrated thus far on the simple model shown in Figure 1 and its associated M/M/1/3 Markov process in Figure 3. It is a routine matter to apply exactly the same reasoning to any system being modeled by a continuous time Markov process. Simply draw the state transition diagram and identify the sets of conditional service rates that have to satisfy homogeneity conditions.

Once the homogeneity conditions are identified, simply run a simulation and examine the generated trajectory. If all homogeneity conditions are satisfied (which implicitly requires every state in the state transition diagram to be visited at least once), and if flow balance is also satisfied, it can be concluded with absolute certainty that the simulation has generated an accurate characterization of the steady state distribution of the underlying stochastic process.

Of course, the parameters of the underlying stochastic process must be set equal to values actually observed in the output trajectory generated by the simulation (rather than using the parameter values specified as input to the simulation program). Assuming that the simulation has run for a reasonably long interval of time, these two sets of values should be nearly identical.

Note that this entire process can be carried out in cases where the closed form analytic expression for the steady state distribution is unknown.

9. Synthetic Alignment Intervals

As this discussion has illustrated, analysts who use Monte Carlo simulations to evaluate steady state distributions should, under ideal circumstances, expect these simulations to generate trajectories that satisfy flow balance and various forms of homogeneity. In cases where these conditions are not satisfied, the accuracy of the simulation can generally be improved by extending the trajectory so that the conditions can be met.

There are two ways to extend the trajectory. The first is to allow the simulation program to run for an additional interval of time. As discussed in the Appendix, the required homogeneity assumptions will almost always be satisfied in the limit. Also, flow balance can be satisfied in principle by continuing to run the simulation until the initial state reappears once again. However, the time required to achieve these conditions may be unacceptably long.

As an alternative, consider the possibility of appending specially constructed “alignment intervals” to the end of the original simulations. The objective of these alignment intervals is to bring the extended trajectory into compliance with the required assumptions. Since synthetic alignment intervals are not generated by calling upon the random number generator that drives the simulation program, they constitute artificially tailored appendages to the original simulation. This implies that the extended trajectories are no longer “purely stochastic”, which raises legitimate concerns regarding the validity of this entire approach

To address these concerns, note first that the information represented within a steady state distribution is only a subset of the totality of information that can be obtained by analyzing the associated stationary stochastic process. For example, in the case of a single server queue, the steady state distribution contains detailed information about queue length, but contains no information at all about the lengths of busy periods. In fact, two stationary stochastic processes with different internal structures (and different busy period distributions) can have exactly the same steady state distribution [BUZE2006a].

For purposes of this discussion, properties of a stochastic process that depend only on its steady state distribution will be referred to as distributional properties. Conversely, properties of a stochastic process that are not uniquely determined by the steady state distribution will be referred to as trans-distributional properties.

This distinction is critically important because appending a non-random alignment interval to the output of a Monte Carlo simulation is almost certain to alter certain trans-distributional properties of the extended trajectory. Homogeneity and flow balance are not sufficient to ensure that all trans-distributional properties of the underlying stochastic process are preserved. Even though the extended trajectory is no longer “random”, and even though it no longer exhibits the totality of trans-distributional properties associated with the underlying stochastic process, the arguments presented in Sections 3 – 6 are still valid: they are sufficient to demonstrate that the extended trajectory provides an entirely accurate characterization of the underlying steady state distribution. This conclusion provides

the philosophical justification for appending a non-random alignment interval to the output of a Monte Carlo simulation.

In many cases, it is reasonable to assume that the original Monte Carlo simulation has generated a steady state distribution that is quite close to the exact solution. The approach outlined here can be used to inspect the output of the simulation and identify the most serious violations of homogeneity. Synthetic alignment intervals that correct only these most serious violations can then be appended to the original trajectory, resulting in an incremental improvement in accuracy.

Such a step-wise approach to improving simulation accuracy may ultimately prove to be the most effective procedure for applying these results in practice. It is important to note in this regard that simulation results can be surprisingly accurate even though the required homogeneity assumptions are not satisfied exactly [SURI1983].

A specific example of an alignment interval that correctly adjusts the lower trajectory in Figure 4 is presented in the next section. General algorithms for the construction of synthetic alignment intervals are currently under development.

10. Example – Alignment Interval

Figure 5 displays two trajectories, both having durations of 30 seconds. The upper trajectory in Figure 5 is identical to the lower trajectory in Figure 4.

Assume that the upper trajectory in Figure 5 has been generated by a Monte Carlo simulation of the stochastic process illustrated in Figure 1. As previously discussed, the service rates at servers 1 and 2 are not homogeneous for this trajectory. In particular,

$$\begin{aligned} \text{For server 1:} \quad & C(1)/T(1) = 4/7 \\ & C(2)/T(2) = 2/5 \\ & C(3)/T(3) = 1/2 \end{aligned}$$

$$\begin{aligned} \text{For server 2:} \quad & A(0)/T(0) = 4/16 \\ & A(1)/T(1) = 2/7 \\ & A(2)/T(2) = 1/5 \end{aligned}$$

As already noted, the values of $P(n)$ for this trajectory differ from the correct values obtained by setting $\bar{\lambda}/\bar{\mu} = 1/2$. in equations (33) – (36). In other words, the simulation results shown in the upper trajectory in Figure 5 do not accurately characterize the steady state distribution of the underlying stochastic process.

The lower trajectory in Figure 5 displays an alignment interval that can be used to correct this problem. If the lower trajectory is appended to the upper trajectory, the resulting 60 second trajectory will exhibit homogeneous completion rates for server 1 and server 2.




Figure 5: Simulated Trajectory & Alignment Interval

The following calculations, which are based on the complete 60 second trajectory, verify that completion rates for servers 1 and 2 are now homogeneous.

$$\begin{aligned} T(0) &= 2 + 4 + 7 + 3 + 2 + 5 + 5 + 4 \\ &= 32 \end{aligned}$$

$$\begin{aligned} T(1) &= 1 + 1 + 1 + 1 + 1 + 2 + 1 + 1 + 3 + 1 + 1 + 2 \\ &= 16 \end{aligned}$$

$$\begin{aligned} T(2) &= 2 + 1 + 2 + 1 + 1 + 1 \\ &= 8 \end{aligned}$$

$$\begin{aligned} T(3) &= 2 + 2 \\ &= 4 \end{aligned}$$

$$C(1) = 8$$

$$C(2) = 4$$

$$C(3) = 2$$

$$A(0) = 8$$

$$A(1) = 4$$

$$A(2) = 2$$

As shown below, server 1 satisfies homogeneity:

$$C(1)/T(1) = 8/16 = 0.5$$

$$C(2)/T(2) = 4/8 = 0.5$$

$$C(3)/T(3) = 2/4 = 0.5$$

As shown below, server 2 also satisfies homogeneity:

$$A(0)/T(0) = 8/32 = 0.25$$

$$A(1)/T(1) = 4/16 = 0.25$$

$$A(2)/T(2) = 2/8 = 0.25$$

Since flow balance is satisfied and completion rates at both server 1 and server 2 are homogeneous, the values of $P(n)$ must accurately characterize the steady state distribution of the underlying stochastic process.

To verify this claim, note that the parameters of the underlying stochastic process are estimated as follows for the 60 second trajectory:

$$\bar{\mu} = (8 + 4 + 2)/(16 + 8 + 4) = 0.5$$

$$\bar{\lambda} = (8 + 4 + 2)/(32 + 16 + 8) = 0.25$$

Once again, $\bar{\lambda} / \bar{\mu} = 1/2$. Substituting this value into equations (33) – (36) yields the same distribution seen earlier.

$$P(0) = 8/15$$

$$P(1) = 4/15$$

$$P(2) = 2/15$$

$$P(3) = 1/15$$

As shown below, the values of $P(n)$ obtained from the 60 second trajectory are identical to the “analytically correct” distribution that appears above.

$$P(0) = T(0)/T = 32/60 = 8/15$$

$$P(1) = T(1)/T = 16/60 = 4/15$$

$$P(2) = T(2)/T = 8/60 = 2/15$$

$$P(3) = T(3)/T = 4/60 = 1/15$$

In this example, the alignment interval and the original trajectory have exactly the same length. This is not a general requirement: shorter alignment intervals can in fact be constructed.

11. Conclusions

By testing the output of a Monte Carlo simulation to see if certain mathematical relationships are satisfied, it is possible to determine if the simulation has generated an accurate result (i.e., if the simulation has provided an accurate characterization of the steady state distribution of the underlying stochastic process).

If the output fails to pass the appropriate tests, the accuracy of the simulation can be improved by appending a synthetic (non-random) alignment interval to create an extended simulation interval that does in fact possess the desired characteristics. An example illustrating this procedure has been provided.

The development of algorithms for the construction of synthetic alignment intervals is – at present – an open research problem. If general algorithms can be developed, simulation times can be shortened and confidence in the accuracy of simulation results can be enhanced.

The approach presented here can be extended directly to any Monte Carlo simulation of a continuous time Markov process.

12. Bibliographic Notes

The material presented in this paper represents a new application of Operational Analysis. Introduced thirty years ago [BUZE1976a], operational analysis is concerned with the development of equations that characterize the observable behavior of systems as they operate over time. No assumptions are made regarding the existence of an underlying stochastic process. Instead, all assumptions are formulated in terms of relationships among quantities that can be observed and measured under normal operating conditions.

The concept of homogeneity as used in this paper – along with the derivations of $P(n)$ that are based on homogeneity and flow balance – closely parallel material that was originally presented in [BUZE1976b]. These derivations were subsequently extended to a broad class of queuing network models [DENN1978]. Suri’s analysis of the robustness of queuing network formulae was based upon operational analysis and the concept of homogeneity, but his work did not consider implications for Monte Carlo simulation [SURI1983].

The application of operational analysis to the output of Monte Carlo simulations is a recent development that has been characterized as Operational Analysis 2.0 [BUZE2006b]. Since Monte Carlo simulations can be regarded as explicit realizations of underlying stochastic processes, new issues that were not considered in the original formulation of operational analysis become relevant. These include the rationale for utilizing synthetic alignment intervals and the relationship between stochastic and operational assumptions (see Appendix).

A preliminary version of this paper was presented in June 2007 at the 21st European Conference on Modeling and Simulation (ECMS) held in Prague, CZ.

13. Acknowledgments

I am grateful to Rich Olcott for pointing out the need to introduce notational conventions that clearly distinguish between abstract stochastic quantities such as $p(n)$ and measurable values such as $P(n)$ that are associated with the output of a Monte Carlo simulation.

14. Bibliography

- [BUZE1976a] J.P. Buzen, “Fundamental Operational Laws of Computer System Performance,” *Acta Informatica*, 7, No. 2 (June 1976), 167-182.
- [BUZE1976b] J.P. Buzen, “Operational Analysis: The key to the new generation of performance prediction tools”. In *Proceedings of IEEE COMPCON 76* (Washington, DC, Sept. 1976) , 166-171
- [BUZE2006a] J.P. Buzen, “Counter intuitive aspects of statistical independence in steady state distributions”. In *Communication Networks and Computer Systems*, J.A. Barria (Ed.), Imperial College Press, London, 89 – 104.
- [BUZE2006b] J.P. Buzen, “New Perspectives on Benchmarking, Modeling and Monte Carlo Simulation: Operational Analysis 2.0” In *CMG 2006 Conference Proceedings* (Reno, NV, Dec. 4 – 8, 2006), Computer Measurement Group, Turnersville, NJ.
- [DENN1978] P.J Denning and J.P. Buzen, “The Operational Analysis of Queuing Network Models,” *ACM Computing Surveys*, 10, No. 3, (Sept. 1978), 225-261.
- [SURI1983] R. Suri, “Robustness of Queuing Network Formulas,” *Journal ACM*, 30, No. 3 (July 1983), 564-594.

Appendix

The assumptions of flow balance and homogeneous service times may seem somewhat arbitrary, but they are actually quite reasonable when dealing with the output of a Monte Carlo simulation that is based on an M/M/1/3 stochastic process. This Appendix outlines the basic arguments that support this claim.

In the case of the flow balance, the sole reason for assuming that the initial and final states of the trajectory are identical is justify the claim that the number of transitions into each state must be equal to the number of transitions out.

As previously noted in Section 4, failure to satisfy flow balance implies that there will be one extra transition out of the initial state and one extra transition into the final state. For example, suppose that state 1 is the initial state, and state 3 is the final state. As a result of this failure to satisfy flow balance, equations (10) – (13) become:

$$A(0) = C(1)$$

$$A(1) + C(1) = A(0) + C(2) + 1$$

$$A(2) + C(2) = A(1) + C(3)$$

$$C(3) = A(2) - 1$$

Equations (16) - (19) then become

$$\frac{A(0)}{T(0)} \cdot P(0) = \frac{C(1)}{T(1)} \cdot P(1)$$

$$\left[\frac{A(1)}{T(1)} + \frac{C(1)}{T(1)} \right] \cdot P(1) = \frac{A(0)}{T(0)} \cdot P(0) + \frac{C(2)}{T(2)} \cdot P(2) + \frac{1}{T}$$

$$\left[\frac{A(2)}{T(2)} + \frac{C(2)}{T(2)} \right] \cdot P(2) = \frac{A(1)}{T(1)} \cdot P(1) + \frac{C(3)}{T(3)} \cdot P(3)$$

$$\frac{C(3)}{T(3)} \cdot P(3) = \frac{A(2)}{T(2)} \cdot P(2) - \frac{1}{T}$$

As T increases, the value of $1/T$ clearly decreases. In the limit as the length of the simulation interval approaches infinity, the term $1/T$ becomes negligible and can be ignored. [Note that certain values of $A(n)/T(n)$ in the global balance equations must remain strictly greater than zero as T approaches infinity for this conclusion to be valid.]

Thus, the assumption of flow balance produces a solution for $P(n)$ that – except under special circumstances – is exactly correct in the limit as T approaches infinity and approximately correct in most cases where T is large.

Now consider the homogeneity assumptions for servers 1 and 2. These assumptions regarding the output of the simulation are a direct consequence of the assumption that service times are exponentially distributed in the underlying stochastic process.

To demonstrate the connection, assume that the underlying $M/M/1/3$ stochastic process is in steady state and consider its behavior during a sample path of length T . [A sample path is simply an idealized mathematical counterpart of the output of a perfect Monte Carlo simulation.]

As usual, let $p(0)$ be the steady state probability that there are no customers at server 1. Then

$$p(0) \cdot T = \text{the expected amount of time that server 1 is idle.}$$

$$[1-p(0)] \cdot T = \text{the expected amount of time that server 1 is busy.}$$

Since $1/\mu$ is the mean service time at server 1, the expected number of requests completed during the time server 1 is busy is

$$\frac{[1-p(0)] \cdot T}{1/\mu} = \mu \cdot [1-p(0)] \cdot T$$

To compute the expected rate at which requests are completed by server 1, this number must be divided by $[1-p(0)] \cdot T$, which is the expected amount of time that server 1 is busy. Thus the expected service rate is simply equal to μ .

This analysis is entirely straightforward and does not require service times at server 1 to be exponentially distributed. However, this analysis does depend upon one highly significant implicit assumption that needs to be examined before proceeding further.

As already noted, $[1-p(0)] \cdot T$ is the expected amount of time server 1 is busy during the sample path of length T . Note that this total busy time is divided into a number of distinct busy periods, each beginning when the server is idle and a new customer arrives (a transition from state 0 to state 1), and each concluding when a customer completes service and leaves no one behind (a transition from state 1 to state 0). Each customer served during a busy period contributes exactly one full service time to the total duration of the busy period. Thus the total duration of the busy period can be regarded as the sum of an integral number of samples from the probability distribution that characterizes the service time.

To determine if homogeneity is satisfied at server 1, it is necessary to evaluate the conditional service rates at server 1, computed over those times when the stochastic process is in state n (for $n=1,2,3$). This analysis is similar to the unconditional case that has just been presented: first determine the expected

time spent in state n ; then divide this time into the expected number of completions that occur while in state n .

It should be immediately apparent that the expected time spent in state n is equal to $p(n) \cdot T$. However, computing the expected number of completions while in state n requires a more subtle analysis. Unlike the previous case, the total time in state n cannot be regarded simply as the sum of a number of busy periods, where each busy period's length is the sum of an integral number of complete service requests. Instead, customers can be interrupted by new arrivals midway through their service, causing a change of state. Thus, the total service time for a single request can be split among two or more states. As a result, the expected number of requests completed while in state n cannot always be written as:

$$\frac{p(n) \cdot T}{1/\mu} = \mu \cdot p(n) \cdot T$$

However, if service times are exponentially distributed at server 1, the memoryless property of the exponential distribution implies that service completions while in state n can be characterized as a Poisson process with rate μ . This process operates whenever the system is in state n , even if this time is partitioned into sub-intervals with beginning and end points that may not correspond to the actual beginnings and ends of individual service requests.

As a result of these considerations, the assumption of exponential service times at server 1 implies that the expected number of completions while in state n is in fact equal to $\mu \cdot p(n) \cdot T$. The conditional completion rate at server 1 while the underlying stochastic process is in state n is therefore equal to:

$$\frac{\mu \cdot p(n) \cdot T}{p(n) \cdot T} = \mu$$

Since this expression is independent of n , the expected conditional completion rates at server 1 are independent of n for $n=1, 2$ and 3 . This is, of course, the stochastic counterpart of the assumption of homogeneous service rates.

A similar argument shows that the conditional completion rates at server 2 are independent of n for $n=0, 1$ and 2 if service times at server 2 are exponentially distributed.

These arguments demonstrate that homogeneity is not an arbitrary algebraic convenience, created merely to simplify the analysis. On the contrary, homogeneity assumptions can be expected to hold in any Monte Carlo simulation of a continuous time Markov process where service times are exponentially distributed and where the simulation is allowed to run for a sufficiently long period of time. In the limit as T approaches infinity, homogeneity assumptions can be expected to hold with probability one.

Distribution of Total Times for Tasks Subject to Failure and Restart

Søren Asmussen
Aarhus University, Denmark
<http://home.imf.au.dk/asmus/>

May 8, 2006

Joint work with Lester Lipsky, Robert
Sheahan and Pierre Fiorini¹

¹Full paper (with additional co-author Tomasz Rolski) can be downloaded from www.thiele.au.dk. Conditionally accepted for *Math. Oper. Res.*

T execution time of job $\sim F$

U failure time $\sim G$

D task duration $\sim H$

RESUME checkpoints

REPLACE start new job at failure

REPLACE start same job

$$\begin{array}{c}
 T \\
 \hline
 U_1 \\
 \hline
 U_2 \\
 \hline
 U_3 \\
 \hline
 U_N \\
 \hline
 \end{array}
 D = T + U_1 + \dots + U_{N-1}$$

Target: tail $\bar{H}(x) = P(D > x)$ of D

Theorem 1 T bounded $\Rightarrow \bar{H}(x) \approx e^{-\gamma x}$

Theorem 2 T unbd $\Rightarrow H$ heavy-tailed

LIGHT TAILS

1) exponential $\bar{F}(x) = e^{-\beta t}$

2) Gamma-like $\bar{F}(x) \approx ct^{\alpha-1}e^{-\beta t}$

3) Rayleigh $\bar{F}(x) = e^{-\beta t^2}$

4) LT Weibull $\bar{F}(x) = e^{-\beta t^\gamma}$,
 $\gamma \geq 1$

General: $\int_0^\infty e^{\epsilon t} f(t) dt < \infty$,
some $\epsilon > 0$

Failure rate $\frac{f(t)}{\bar{F}(t)} \rightarrow \infty$ or $c > 0$

HEAVY TAILS

General: $\int_0^\infty e^{\epsilon t} f(t) dt = \infty$,
all $\epsilon > 0$

Failure rate $\frac{f(t)}{F(t)} \rightarrow 0$

1) power tails $\bar{F}(x) = \frac{c}{t^\alpha}$

2) lognormal $\bar{F}(x) \approx ce^{-\alpha \log^2 t}$

3) HT Weibull $\bar{F}(x) = e^{-\beta t^\gamma}$,
 $\gamma < 1$

0) $\bar{F}(x) \approx \frac{c}{\log^\alpha t}$

2^δ) $\bar{F}(x) \approx e^{-\alpha \log^\delta t}$

$\delta = 1$: power tails

COMPARISON RESULTS

$$F = G \Rightarrow \bar{H}(x) \approx \frac{1}{x}$$

$$F = G \Rightarrow ED = \infty$$

$$\bar{F} \ll \bar{G} \Rightarrow ED < \infty$$

$$\bar{F} \gg \bar{G} \Rightarrow ED = \infty$$

$$\bar{F} \ll \bar{G}^2 \Rightarrow \text{Var } D < \infty$$

$$\bar{F} \gg \bar{G}^2 \Rightarrow \text{Var } D = \infty$$

$$\bar{F}_1 \ll \bar{F}_2 \Rightarrow H_{F_1, G} \ll H_{F_2, G}$$

$$\bar{G}_1 \ll \bar{G}_2 \Rightarrow H_{F, G_1} \gg H_{F, G_2}$$

TAIL OF H

4 examples of each of F, G :

LT Weibull

exponential

HT Weibull

power

$\bar{F}(t)$	e^{-t^2}	e^{-t}	$e^{-t^{1/2}}$	$\frac{1}{t^\alpha}$
e^{-u^2}	$\frac{1}{x}$	$e^{-\log^{1/2} x}$	$e^{-\log^{1/4} x}$	$\frac{1}{\log^{\alpha/2} x}$
e^{-u}	$e^{-\log^2 x}$	$\frac{1}{x}$	$e^{-\log^{1/2} x}$	$\frac{1}{\log^\alpha x}$
$e^{-u^{1/2}}$	$e^{-\log^4 x}$	$e^{-\log^2 x}$	$\frac{1}{x}$	$\frac{1}{\log^{2\alpha} x}$
$\frac{1}{t^\alpha}$	$e^{-x^{\frac{2}{2+\alpha}}}$	$e^{-x^{\frac{1}{1+\alpha}}}$	$e^{-x^{\frac{1/2}{1/2+\alpha}}}$	$\frac{1}{x}$

Precise meaning: logarithmic asymptotics

Constants omitted $e^{-c \log^{1/2} x}$

EXACT ASYMPTOTICS

$\gamma(t)$ solution of $\mathbf{1} = \int_0^t e^{\gamma(t)u} g(u) du$
 (geometric sums, renewal equation)

$$T \equiv t_0 \Rightarrow \overline{H}(x) \approx C(t_0) e^{-\gamma(t_0)x}$$

Bounded support:

$$f(t) \sim A(t_0 - t)^\alpha, \quad t \uparrow t_0$$

$$\overline{H}(x) \sim C_1(t_0) \frac{e^{-\gamma(t_0)x}}{x^\alpha}$$

$C_1(t_0)$ involves

$$\gamma(t_0), C(t_0), \Gamma(\alpha + 1), g(t_0)^{\alpha+1}$$

$$F \text{ Gamma-like: } \bar{F}(x) \sim At^\eta e^{-\delta t}$$
$$g(t) = \beta e^{-\beta t}$$


$$\bar{H}(x) \sim \frac{A\Gamma(\delta/\beta) \log^\eta x}{\beta^{\delta/\beta-1-\eta} x^{\delta/\beta}}$$

Ambiguities in clustering Subsequence Time Series

To Emeritus Professor Lester Lipsky,
our teacher, mentor, and friend

George Nagy

The problem

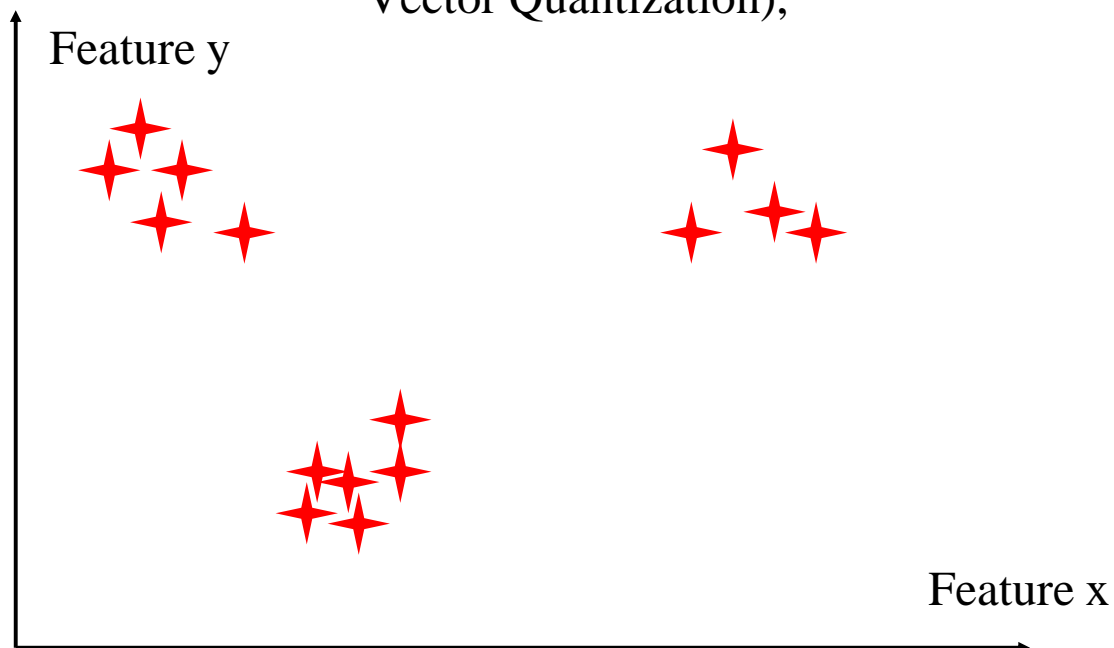


Conclusion

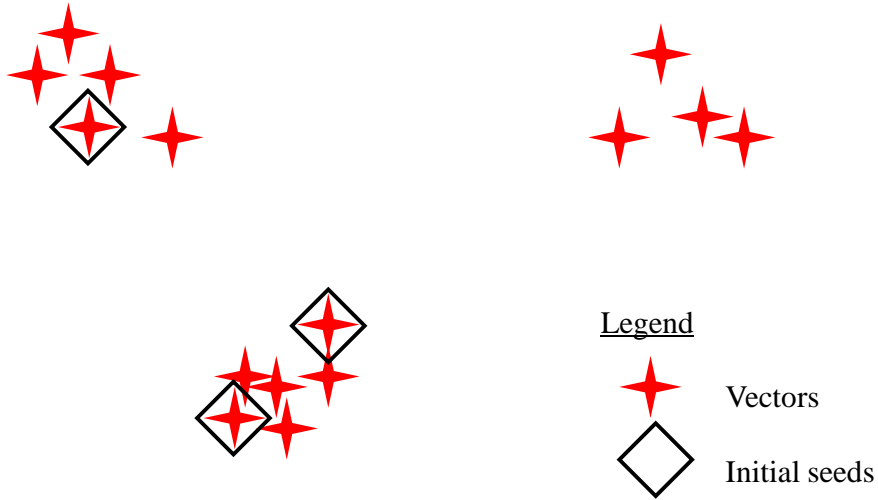
- Although many publications purport to show the contrary,

we cannot expect the *location* of the K-means centroids of shifted & overlapped subsequence vectors to discriminate such sequences.


K-means Clustering Algorithm
(a.k.a. McQueens Algorithm (1967) or
Vector Quantization),




Initial seeds selected randomly from vectors



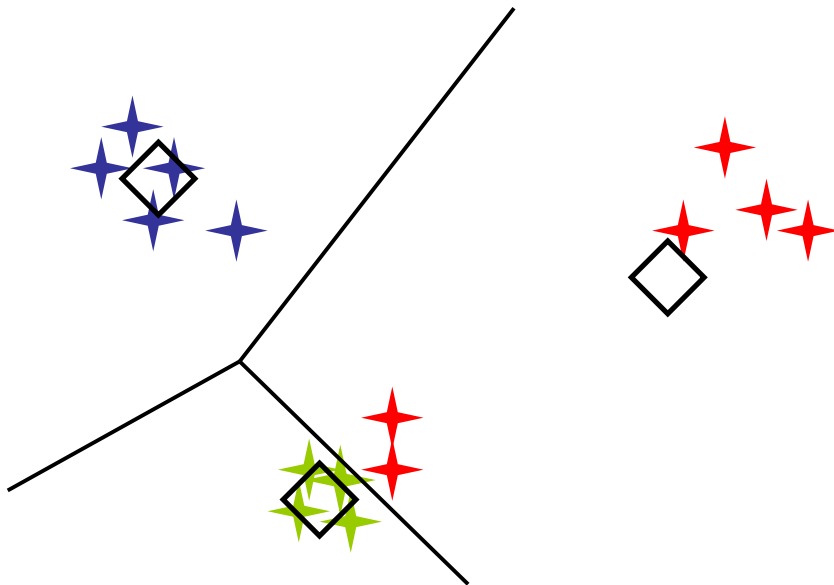
Initial Voronoi partition of space into proximity regions of seeds




Partition based on initial cluster centers (Step 1)



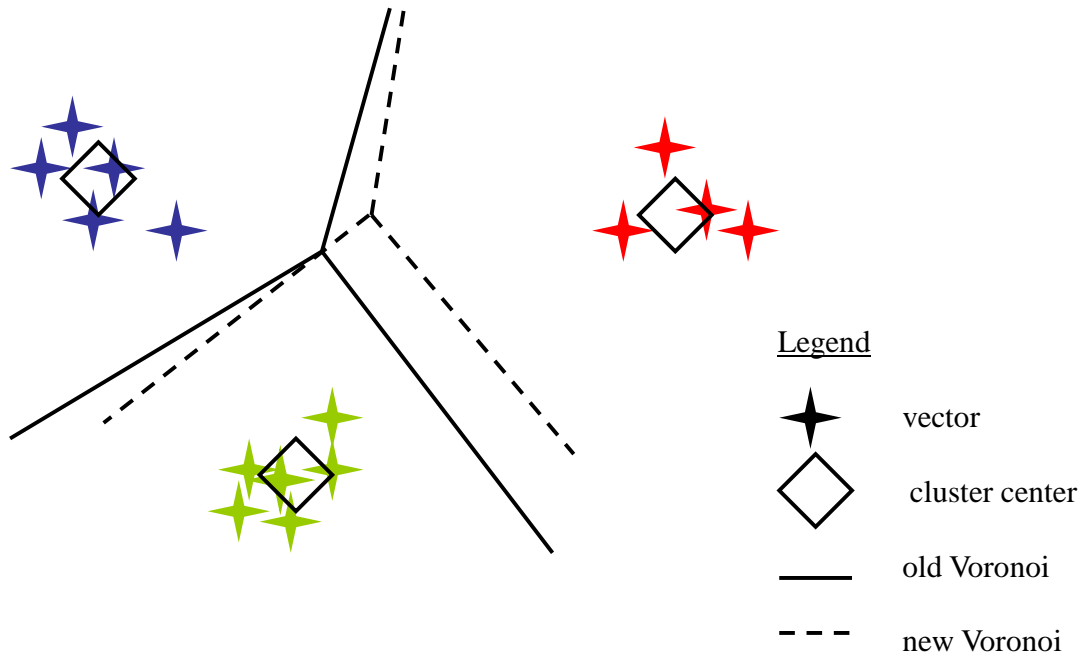
New *centroids* based on current partition (Step 2)



Next partition based on current centroids (Step 1)



New centroids based on current partition (Step 2)



Convergence!


- Further iterations will result in no change.
- A local minimum of the distances of the vectors to their cluster centers has been reached.

NB The K-means optimization criterion is often called *mean-squared error* (MSE).

Properties of K-means


- K-means eventually reaches a local minimum, because both steps decrease the MSE.
- Clusters may be dropped (“*falling between two stools*”).
- New clusters are never generated.
- The number of clusters is set by the number of initial seeds. If a cluster is dropped, new seeds may be chosen.
- More runs with fresh random seeds may decrease the MSE.
- Many variations of the algorithm proposed in the last 30 years.
- MATLAB has an efficient implementation, with options for choice of initial seeds, distance measures, max iterations, max replication.
- Equivalent cluster configuration with different cluster numbering can be remapped into one another.
- There may be multiple cluster configurations with the *same global minimum* that cannot be remapped

Example: A set of points in 2-D.
 How would we group them into 3 clusters?




An optimal cluster configuration with $K=3$.

The cluster centroids form a triangle.



An equivalent cluster configuration with $K=3$. The centroids cannot be remapped by renumbering them, but the triangles are *similar*.



Pairwise distances

- The K-means optimization criterion MSE, minimizing the intra-cluster radii, is equivalent to *maximizing* the average intra-cluster distances or similarities (*dot products*):

$$KSE = \sum_{k=1}^K \frac{1}{n^{(k)}} \sum_{i \in I(k)} \sum_{j \in I(k)} \mathbf{x}_i \cdot \mathbf{x}_j'$$

(Cosine and Euclidian distance are equivalent for same-length vectors)

Equivalent cluster configurations

- In the 2-D, $K=3$ examples, we saw an instance of a set of vectors that produce two *equivalent, non-remappable* cluster configurations.
- We will now show that this phenomenon inevitably occurs in clustering periodic time sequences.

STS Clustering

$S = 2, 3, 8, 4, 7, 4, 3, 9, 4, 7, 7, 3, 8, 5, 7$

Pattern vectors with $w=5$:

2, 3, 8, 4, 7
3, 8, 4, 7, 4
8, 4, 7, 4, 3
4, 7, 4, 3, 9
7, 4, 3, 9, 4
4, 3, 9, 4, 7
3, 9, 4, 7, 7
9, 4, 7, 7, 3
4, 7, 7, 3, 8
7, 7, 3, 8, 5
7, 3, 8, 5, 7

Repetitive subsequences

- $s = \{2, 3, 8, 4, 7\}$
- $S = \{2, 3, 8, 4, 7, 2, 3, 8, 4, 7, 2, 3, 8, 4, 7, \dots\}$

- Circulant matrix:

$$M = \begin{bmatrix} 2 & 3 & 8 & 4 & 7 \\ 3 & 8 & 4 & 7 & 2 \\ 8 & 4 & 7 & 2 & 3 \\ 4 & 7 & 2 & 3 & 8 \\ 7 & 2 & 3 & 8 & 4 \end{bmatrix}$$

All 5 vectors have the same length.

Maximum number of distinct distance values

- $S = \{2, 3, 8, 4, 7, 2, 3, 8, 4, 7, 2, 3, 8, 4, 7, \dots\}$
- $\leftarrow \quad w \quad \rightarrow$
- Max number of different distances between subsequences
 - = $w/2$ for w even,
 - = $(w-1)/2$ for w odd.
- (This is much less than the expected $w(w-1)/2=10$)
- Here ($w=5$) we have only dot product values of 104 and 113 between the 10 pairs of vectors.
- In feature space, the vectors form isosceles simplexes! These give rise to equivalent cluster configurations.

Aperiodic sequences (w=8)

- # 1 2 3 4 5 6 7 8 9 10 11 12 13 ...
 - S: {2 3 8 4 7 1 6 5 9 4 1 3 8 ...}
-
- (1) 2 3 8 4 7 1 6 5
 - (3) 8 4 7 1 6 5 9 4
 - (2) 3 8 4 7 1 6 5 9
 - (4) 4 7 1 6 5 9 4 1
- many almost identical pairwise distance values, giving rise to cluster configurations with almost the same MSE

Conclusion

- The K-means MSE criterion can be expressed in terms of pairwise distances.
- Multiple occurrences of the same distance value give rise to equivalent configurations of cluster centroids.
- Such equivalent cluster configurations have different memberships and different centroids, hence they are not mappable by renumbering.
- Therefore different random initializations are likely to result in different cluster configurations with the same MSE.
- Pairs of shifted subsequences of a periodic sequence share the same component values.
- The distance between such subsequences is restricted to only a few ($\sim w/2$) distinct values.
- So STS cluster centroids cannot be mapped 1:1 to periodic sequences.
- An approximate version of this phenomenon occurs in subsequence clustering of all “smooth” sequences.

A very simple example of equivalent cluster configurations



PERFORMANCE-RELEVANT NETWORK TRAFFIC CORRELATION

Hans-Peter Schwefel
Center for Teleinfrastruktur
Aalborg University
email: hps@kom.aau.dk

Imad Antonios
Dept. of Computer Science
Southern Connecticut State Univ.
email: antonios1@southernct.edu

Lester Lipsky
Dept. of Comp. Sci. & Eng.
University of Connecticut
email: lester@enr.uconn.edu

ABSTRACT

Correlation structure is an important metric to consider when modeling the performance of network traffic. Particularly, the presence of long-range dependence (LRD) in the input process may, under some circumstances, lead to poor queueing performance. In this paper, we first aim to characterize the conditions under which the presence of LRD is performance relevant. We define variations of an ON/OFF-type process that employ a truncated power-tail (TPT) distribution, and analyze their correlation structure in relation to queueing performance. Our analytic results show that the correlation structure in some cases is very sensitive to the presence of a background Poisson process, and that while other model variations exhibit LRD, it is only those with TPT-distributed ON times that queueing performance is poor. These results lead us to propose a procedure for extracting performance-relevant correlation properties, whose effectiveness we demonstrate via simulation experiments using synthetic and measured traffic.

Keywords: ON/OFF models, long-range dependence, queueing models, autocorrelation.

INTRODUCTION AND MOTIVATION

Innumerable studies of Internet traffic have shown that it exhibits *self-similarity* and *long-range dependent* (LRD) properties (see [Leland et al. 1994]), meaning that it is highly varying, and bursty over a wide range of time scales. This poses challenges in understanding the factors underlying it, and thus makes it difficult to develop models to predict network performance. Within the large body of research on the subject, a bulk of the work has concentrated on constructing statistical models that reproduce the correlation structure of the measured traffic [Melamed 1991, Norros 1995], and studying the effects of generated traffic on queueing perfor-

mance behavior via simulation. The main conclusion was that the presence of LRD in traffic carries important implications for network performance. Another body of research was concerned with uncovering the physical factors that underlie self-similarity and LRD [Crovella and Bestavros 1996, Veres and Boda 2001]. The factors identified include file-size distribution, file transmission time distribution, user think times, and TCP congestion control mechanism, which were used to construct analytic performance models with matching statistical properties of measured traffic. One of the important long-range dependent traffic models is the aggregation of ON/OFF sources with heavy-tailed ON periods [Crovella and Bestavros 1996, Willinger et al. 1995].

Since LRD properties have been recognized not to be sufficient predictors of queueing performance, it is important to characterize the conditions under which their presence has a significant bearing on network performance, which is the goal of this paper. The starting point of our inquiry is the definition of a rich analytic traffic model of ON/OFF-type sources with four variations whose correlation structure we examine. In particular, we consider the autocorrelation structure of packet inter-arrival times as well as the counting process associated with the model variations. To account for settings where non-LRD traffic is aggregated with LRD traffic, the model variations allow for a background Poisson process.

Summary of results: Our results of the coefficient of autocorrelation for the inter-packet times show that not all models exhibit LRD properties, whereas for the counting process LRD properties are present in all model variations with TPTs. Our calculations also indicate that the correlation structure is highly sensitive to the superposition of a background Poisson process only for inter-packet times, but not for the counting process. Our analytic results for the queueing behavior reveal that some models with LRD in their inter-packet times

or counting process do not show significant delays – it is only the models in which TPTs are used for the ON periods that poor queueing behavior, as measured by average queue lengths, can be observed for utilizations well below 1. Power tails in the OFF periods only become performance relevant when the utilization is close to saturation. The mere existence of LRD is hence not a performance-relevant aspect. This conclusion motivated the investigation of an alternative correlation metric that is based on a filtered counting process. An exploration of the proposed measure using both synthetically generated and measured traffic reveals its effectiveness in destroying the correlation structure for traces whose associated queueing performance is not poor, thus making it a more reliable predictor of performance than the presence of LRD properties.

A SIMPLE ANALYTIC ON/OFF MODEL

ON/OFF models have been widely used for modeling of bursty traffic for the last 20 years [Heffes 1980]. After the identification of long-range dependent phenomena in network traffic [Leland et al. 1994], such models have been extended to include heavy-tail distributions [Dumas and Simonian 2000] or their truncated counterparts [Schwefel and Lipsky 1999]. The latter reference introduces an aggregated ON/OFF model with general Matrix-Exponential (ME) ON periods (see Appendix or [Lipsky 1992] for a detailed treatment of ME distributions), represented by a complex, but structured Markov Modulated Poisson Process (MMPP). For the purpose of this paper, we will only present the structure of single-source ON/OFF models. These, however, can be generalized to describe the aggregation of traffic from multiple sources by expressing their structure as the product of the state spaces of single-source models as shown in [Schwefel 2000].

The model variations considered make use of truncated power-tail distributions introduced in [Greiner et al. 2000] (see Appendix for brief description); when these describe the ON periods only, the model is referred to as TPT-ON. If it is for the OFF periods only the model is TPT-OFF, and when both the ON and OFF periods are TPT distributed the model is TPT-ALL. As a case for comparison, a fully exponential ON/OFF model (EXP) is also used. It is important to note that for the range of TPT distributions considered, the results would be indistinguishable from using a full power-tail for finite sample experiments, whether these are based on measurement or simulation.

All model variations are special cases of the ME-ALL model as presented in the following. The ME-ALL model is an MMPP model which generates inter-packet times from a single ON/OFF source,

with peak-rate λ_p and long-term average rate $\lambda = \lambda_p * E\{ON\} / (E\{ON+OFF\})$ aggregated with some background Poisson process with rate λ_{bg} . The ME-ALL model uses Matrix-Exponential OFF periods, $\langle \mathbf{p}_{off}, \mathbf{B}_{off} \rangle$, and Matrix-Exponential ON periods with representation $\langle \mathbf{p}_{on}, \mathbf{B}_{on} \rangle$. The modulating Markov process then has the following generator matrix \mathcal{Q} , and the corresponding Poisson packet rates on the diagonal of \mathcal{L} :

$$\mathcal{Q} = \left[\begin{array}{c|c} -\mathbf{B}_{off} & \mathbf{B}_{off} \boldsymbol{\varepsilon}'_{off} \mathbf{p}_{on} \\ \hline \mathbf{B}_{on} \boldsymbol{\varepsilon}'_{on} \mathbf{p}_{off} & -\mathbf{B}_{on} \end{array} \right],$$

$$\mathcal{L} = \left[\begin{array}{c|c} 0 & \\ \hline & 0 \\ \hline & \lambda_p \mathbf{I} \end{array} \right] + \lambda_{bg} \mathcal{I}.$$

Note that the individual blocks may have different dimensions. $\boldsymbol{\varepsilon}'$ is a column vector with all elements equal to 1 of the corresponding dimension, and \mathcal{I} is the unit matrix of dimension $\dim(\mathbf{B}_{on}) + \dim(\mathbf{B}_{off})$. The ME-ALL process has a mean inter-packet time of $E\{X\} = 1/(\lambda + \lambda_{bg})$.

An ME-ALL/M/1 queue is a special case of an MMPP/M/1 queue and can be represented as a Quasi-Birth-Death Process with the standard matrix-geometric solution [Latouche and Ramaswami 1993]. LRD properties manifest themselves in the correlation structure of a stochastic process, which can be computed for the inter-packet times from an MMPP using standard methods, see [Meier-Hellstern and Fischer 1992].

In addition to the inter-packet time process (X_i), the counting process can also be used to describe traffic. In this paper, we use the interval-based version of the counting process, where $N_i(\Delta)$ is the random variable representing the number of packet arrivals in interval $[i\Delta, (i+1)\Delta]$. The covariance of the counting process $N_i(\Delta)$ of an MMPP for an interval size $\Delta > 0$ is obtained in [Neuts 1989].

The TPT-ALL model uses truncated power-tail distributions for both the ON period and the OFF period. The use of an exponential distribution $\langle p_{OFF} = 1, B_{off} = 1/Z \rangle$ for the OFF periods results in the TPT-ON model, and the use of an equivalent exponential distribution in the ON periods instead causes the model to reduce to the TPT-OFF model. When both ON and OFF periods are exponential, the model reduces to the classical 2-state ON/OFF model [Heffes 1980].

CORRELATION AND QUEUEING PERFORMANCE

Common approaches in traffic modeling typically attempt to fit the first two moments of the traffic trace; capturing the correct mean is essential, since it determines the utilization of any performance model, which in turn can have a critical impact on stationarity. Fitting the variance is fully justified for service time distributions, since the basic performance metric of average queue length in a simple queueing model such as M/G/1 is directly related to the variance of the service time. On the other hand, the variance of an arrival process is of little impact on performance behavior in realistic utilization ranges well below 1. Of primary interest are long-range dependent properties [Melamed 1991, Norros 1995], which have more recently been recognized as being critical in guiding the development of traffic models.

Correlation Structure and Long-Range Dependence

We consider the correlation structure of the four model variations both in terms of their inter-packet time and counting processes. While studies on traffic LRD have mostly focused on the latter, using the inter-packet time process is primarily motivated by its immunity to potential sampling problems caused by the additional counting process parameter of interval size.

We first examine the case for inter-packet times. Figure 1 shows the results of analytic computations of the autocorrelation coefficients when a background Poisson stream is merged with the output of each of the model variations. In all models, the arrival rate of the ON/OFF model is $\lambda = 1$, and the contribution of the background process is the additional rate $\lambda_{bg} = 0.1$. Except for EXP, in which short-range dependence (SRD) can be observed, all models exhibit LRD properties. Looking at the TPT-OFF model, we find that the magnitude of the autocorrelation coefficients is greater than that for TPT-ON. This is surprising since the TPT-OFF model without the presence of the background Poisson traffic is a renewal process: the samples are either created by a Poisson process of rate λ_p (during the ON period), or by a convolution of an exponential residual time of the ON period together with a TPT-distributed ON period and another exponentially distributed time until first packet in the ON period. The absence of background traffic also renders EXP a renewal process, while for TPT-ALL and TPT-ON LRD properties are still present. These results indicate that only when the ON period is TPT distributed that LRD properties appear for inter-packet times.

Additional experiments with very small values for λ_{bg} in the order of 10^{-5} and smaller show that even




Figure 1: **Correlation structure of the inter-packet times with 10% background Poisson traffic:** The TPT-OFF model shows positive correlation only in the presence of background Poisson traffic. The TPT-ALL and TPT-ON models show LRD in both scenarios. Curves without background traffic are not shown since two of them, EXP and TPT-ALL, are renewal.

such small background Poisson rates are sufficient to keep the LRD properties of the inter-packet times for the TPT-OFF model. Hence, any small overlay process will make these models appear to be LRD.

The results of the computation of the autocorrelation coefficient for the counting process are shown in Figure 2. The TPT-ALL, TPT-ON and TPT-OFF models all exhibit LRD properties, and the magnitude of the coefficients is in decreasing order for the respective model. The EXP model only exhibits short-range dependence. As opposed to the computations for the inter-packet times, the background Poisson stream has virtually no effect on the correlation structure.

The behavior of the different processes is summarized in the following table (LRD = long-range dependence, SR = short-range correlation, 0 = renewal process):

model / scenario	$\lambda_{bg} = 0$	$\lambda_{bg} \rightarrow 0$	$\lambda_{bg} > 0$
inter-packet times			
EXP	0	0	SR
TPT-OFF	0	LRD	LRD
TPT-ON	LRD	LRD	LRD
TPT-ALL	LRD	LRD	LRD
counting process			
EXP	SR	SR	SR
TPT-OFF	LRD	LRD	LRD
TPT-ON	LRD	LRD	LRD
TPT-ALL	LRD	LRD	LRD




Figure 2: **Correlation structure of the counting process for $\Delta = 1$:** All three models that use TPT distributions show LRD properties. The background Poisson traffic has negligible impact on the correlation structure: the curves for $\lambda_{bg} = 0.1$ are visually indistinguishable and hence not shown here.

Queueing Performance

Since all four models, TPT-ALL, TPT-ON, TPT-OFF, and EXP have an MMPP representation, analytic results for the queue length distribution and the average queue length when using these processes as arrival to an exponential server can be obtained via known matrix-algebraic methods [Latouche and Ramaswami 1993, Neuts 1981]. In this section, we analyze the different performance behavior as measured by average queue length for different server rates ν , corresponding to different utilizations $\rho = (\lambda + \lambda_{bg})/\nu$.

As Figure 3 shows, three different ranges of utilization values need to be distinguished:

1. For low values of $\rho < 0.1$ (in the given parameter setting), all four models show the same, low average queue length.
2. At $\rho = 0.1$ the normalized average queue length of the TPT-ALL and TPT-ON model increases rapidly. These so-called blow-up points were already identified for the TPT-ON model in [Schwefel and Lipsky 1999]. They occur when the packet-arrival rate during ON periods is higher than the service rate, namely when $\lambda_p + \lambda_{bg} > \nu$ or equivalently $\rho > (\lambda + \lambda_{bg})/(\lambda_p + \lambda_{bg})$. The TPT-OFF and the EXP model do not show these blow-up points.
3. For high utilization $\rho \rightarrow 1$, TPT-ON and EXP show a horizontal asymptote of the normalized mean queue length, indicating growth according to $(1 - \rho)^{-1}$. The TPT-OFF and TPT-ALL models however show blow-ups in their average queue lengths that grow faster than the M/M/1




Figure 3: **Queueing behavior of different ON-OFF/M/1 queues with 10% background traffic:** Shown is the average queue length normalized to the mean queue length $\rho/(1 - \rho)$ of an M/M/1 queue. The TPT-ON and TPT-ALL model show a blow-up around $\rho = 0.1$, while the TPT-OFF and EXP model show rather low queue lengths for typical network utilization values $\rho < 0.8$. The TPT distributions in the OFF periods only becomes relevant for $\rho \rightarrow 1$. The existence of background Poisson traffic does not have a relevant influence here, despite its strong influence on correlation structure of the inter-packet times.

queue when $\rho \rightarrow 1$. This growth can be explained by the analysis of GI/M/1 queues with GI=TPT as discussed in [Greiner et al. 2000].

Note that the existence of background Poisson traffic, despite its potentially large influence on correlation structure of the inter-packet times, has only limited impact on the queueing behavior. In the case where there is no background Poisson traffic, performance is nearly indistinguishable from that shown in Figure 3. Large amounts of background traffic may however shift the blow-up points of the TPT-ON and TPT-ALL model as quantitatively described in Item (2) above.

In summary, the analytic performance results in terms of average queue lengths clearly show that although the TPT-OFF model with some background Poisson traffic exhibits long-range dependence with a large tail-exponent (large absolute values of the autocorrelation curve) and extremely high variance, its queueing performance is rather well behaved. It is only for high utilization values $\rho \rightarrow 1$ that this model causes high queue lengths. The TPT-ON and TPT-ALL models on the other hand are nearly equivalent in terms of queueing behavior in that power tails in the ON periods are the performance dominating element and lead to these blow-up points when the utilization is well below 1.

A PERFORMANCE-RELEVANT CORRELATION METRIC

The previous sections show that autocorrelation of the inter-packet time process is in any case not a good indicator of potential poor queue performance, since the presence of LRD in this process is rather sensitive to the existence of a background Poisson process or any other overlapping process. Auto-correlation of the counting process in turn is less sensitive to such disturbance by a background process, but it still does not provide a performance-relevant metric, since either ON periods or OFF periods could cause such LRD if heavy-tail distributions are present, while only heavy-tailed ON periods may cause poor performance behavior for utilizations well below 100%. The reason for the poor performance, whether it is measured by average queue length/delay or by buffer overflow probability, is the existence of heavy-tailed over-saturation periods. Therefore, it is only the periods of high traffic that are performance critical, even though long-range dependence can also be created by subsequent intervals of heavy-tailed length with low traffic.

Measuring the empiric distribution of the duration of over-saturation periods directly may however not show heavy-tailed distributions due to potential sensitivity toward small-scale traffic variations. The advantage of the correlation properties is the robustness toward such small scale variations. As a consequence, this section proposes and evaluates an approach that is based on the autocorrelation properties of a filtered counting process, such that the impact of periods of low traffic is removed.

The Filtering Method

The approach works as follows: Any maximum-length sequence of subsequent values of the counting process with low traffic intensity $N_i(\Delta), \dots, N_{i+k}(\Delta) < N^*$ is replaced by a single value, for which we use the average of this sequence. In pseudocode this filtering of the measured counts $N_i(\Delta)$ into its counterpart $M_j(\Delta)$ can be expressed as:

Algorithm 0.1: FILTER($N[], M[], N^*$)

```

j ← 1; i ← 1;
while i ≤ length(N)
  if N[i] < N*
    then { M[j] = N[i]; j ← j + 1;
           i ← i + 1;
           k ← 1
           while N[i + k] < N*
             do k ← k + 1
           }
    else { M[j] = average(N[i], ...,
                        N[i + k - 1])
           j ← j + 1;
           i ← i + k
  }

```

Choice of threshold: An important parameter in the algorithm above is the threshold N^* , which determines whether the interval is identified as 'low traffic' or 'high traffic'. In the general case, when only the traffic measurements are available, this value is suggested to be set to $N^* = \lambda\Delta = E(N)$, the expected value of the number of arrivals in an interval of size Δ . In case that it is known that the traffic is being fed into a network with bottleneck capacity ν , the intuitively obvious choice of the threshold is $N^* = \nu\Delta$, which in stable scenarios, namely when $\lambda < \nu$, results in a larger threshold.

Instead of investigating LRD properties of the counts N_i , the LRD properties of the filtered (and somewhat shorter, except for pathological cases) sequence M_j should be investigated, in which LRD properties as caused by periods of low traffic have been removed.

The filtering approach above can alternatively be applied to the inter-arrival time process. In this case, periods of low traffic correspond to a sequence of relatively large inter-packet times, those that exceed $1/\lambda$, or $1/\nu$ if ν is known. When applied to inter-packet times, the algorithm uses the first value as the threshold and replaces sequences of inter-packet times exceeding it with a single value being the threshold.

Application to Synthetic Traces

Figure 4 demonstrates the effectiveness of this approach. The figure shows the analytically computed autocorrelation function of the counting process of both the TPT-ON and TPT-OFF models (dashed lines), together with the empiric correlation function for the original simulated counts (solid line) and the corresponding filtered counts (dotted). Figure 4 is plotted on a log-log scale to reveal the LRD properties, while the linear scale in Figure 5 accentuates the effect of the filtered count on the correlation structure. These results show that the correlation structure of the filtered counts M_j only exhibits LRD properties for models with TPT-distributed ON periods, exactly those with poor performance behavior.

Application to Traffic Measurements

In this section, we aim to further support the effectiveness of the filtering approach in identifying performance-relevant correlation by considering traffic traces and associating them to simulated traffic. We first computed the autocorrelation function for the interarrival process of the Bellcore trace (BC-pAug89) used in [Leland et al. 1994] and its filtered version. The correlation structure of the filtered process, as shown in Figure 6, strongly exhibits LRD, denoting its relevance for queueing performance. It even shows a much clearer straight line behavior in




Figure 4: Autocorrelation coefficient computations and simulations for the counting process of a 1-Burst model and its processed variant (log-log scale): LRD properties in the TPT-OFF model are destroyed by the pre-processing, but not in the TPT-ON model.

the log-log plot, indicating the presence of performance relevant correlation. This is confirmed by the results of a trace-driven queueing simulation for this measurement shown in the curve marked by squares in Figure 7. The figure shows the mean queueing delay as a function of utilization assuming exponential service times. The queueing delay caused by such a trace is comparable to that of an N -Burst simulation with 5 sources and TPT-distributed ON durations (marked by crosses). Note that the queueing delay is substantially larger (by a factor of 100 or more) than in a comparable simulation of Poisson traffic for utilizations above 50%. The sample size for the simulated data was chosen to be the same as the trace, $n = 1e6$.

Based on these results, it can be stipulated that the Bellcore data represents an example of a dataset with LRD caused by ON-OFF like behavior with power-tailed ON periods, which is consistent with the poor queueing behavior. In fact, without having done any detailed parameter fitting, the queueing behavior closely matched that of a TPT-ON model with a small number of sources.

A second set of packet data is also shown in Figure 7, namely a DEC trace of UDP packets (dec-pkt-4) available from [Danzig et al.]. The performance behavior of that trace as obtained in the trace-driven simulation (marked by circles) is very close to an $M/M/1$ queue, with some stronger deviations for ρ above 90%. Since the measurement originates from backbone traffic, an attempt to model it with a large number of TPT-ON sources, here $N = 500$ (marked by stars in Figure 7), yields remarkably close resemblance to the queueing delays of the UDP trace.

Since the method to identify the performance-




Figure 5: Same as Fig. 4 but on a linear scale to accentuate the effect of the filtering process on each of the models.

relevant correlation is designed for traffic generated from a small number of sources (N in the order of 1, ..., 20), it is well suited for analyzing traffic with a low degree of multiplexing (e.g. access networks). As we have shown in Figure 7 and previously in [Schwefel and Lipsky 1999], it is for such traffic that the performance impact of LRD is most detrimental.

SUMMARY AND CONCLUSIONS

In this paper, we examined the autocorrelation structure of a variation of analytical ON/OFF-type models that employ TPT distributions. Our computations indicate that the presence of a background Poisson process has a very strong impact on the coefficient of correlation of the inter-packet times for the TPT-OFF model. For the counting process of these models, the autocorrelation structure was found to exhibit LRD properties for the TPT-ON, TPT-ALL, and TPT-OFF models. The autocorrelation coefficients here are however rather insensitive to the background process. A queueing analysis of the traffic models revealed that only those with an ON period that is TPT distributed, namely TPT-ON and TPT-ALL, suffer from poor performance for a realistic utilization range where $\rho < 0.8$. We proposed and explored an alternative correlation metric that extracted performance-relevant aspects of the counting process. Through simulation experiments based on both measured and synthetic traffic traces, we demonstrated the effectiveness of this approach in destroying the correlation structure of models whose associated queueing performance are not poor, making it a more reliable metric than the mere recognition of LRD.




Figure 6: **Autocorrelation structure of actual and filtered Bellcore Ethernet packet trace (BC-pAug89) traffic:** The filtered variant clearly reveals performance-relevant correlation.

ACKNOWLEDGMENTS

The authors would like to acknowledge Robert Sheehan at the University of Connecticut for his helpful comments that stimulated the development of the filtering process.

APPENDIX

Matrix Exponential (ME) Distributions: Using the notation of [Lipsky 1992], the vector-matrix pair $\langle \mathbf{p}, \mathbf{B} \rangle$ represents a ME distribution with reliability function $R(x)$, and density function $f(x)$ in the following way:

$$\begin{aligned} R(x) &= \mathbf{p} \exp(-x\mathbf{B})\boldsymbol{\epsilon}', \\ f(x) &= -\frac{R(x)}{dx} = \mathbf{p}\mathbf{B} \exp(-x\mathbf{B})\boldsymbol{\epsilon}' \end{aligned}$$

where $\boldsymbol{\epsilon}'$ is a column-vector with all components being 1.

The moments of the distribution come out by integration:

$$\mathbb{E}\{X^k\} = k! \mathbf{p} \mathbf{B}^{-k} \boldsymbol{\epsilon}' \quad (1)$$

Truncated Power-Tail (TPT) Distributions: The T -phase ME representation of a TPT distribution with exponent α is given by [Greiner et al. 2000]: let

$$0 < \theta < 1, \quad \text{and} \quad \gamma := \left(\frac{1}{\theta}\right)^{1/\alpha} > 1.$$

Then, let $\langle \mathbf{p}_T, \mathbf{B}_T \rangle$ be the ME representation of a T -phase Hyperexponential distribution with

$$\mathbf{p}_T = \frac{1-\theta}{1-\theta^T} [1, \theta, \theta^2, \dots, \theta^{T-1}], \quad \text{and}$$




Figure 7: **Queueing performance of measured and synthetic traffic:** The queueing performance of the Bellcore trace (BC-pAug89) is comparable to that of an N -Burst model with $N = 5$, while the queueing behavior corresponding to the UDP packets from the dec-pkt-4 trace closely matches that of highly multiplexed synthetic traffic ($N = 500$) and the $M/M/1$ queue.

$$\mathbf{B}_T = \mu \text{diag} \left(1, \gamma^{-1}, \dots, \gamma^{-(T-1)} \right).$$

The parameter, μ , is a positive constant that can be chosen to set the mean of the distribution according to Eq. (2).

The expected value follows directly from Eq. (1)

$$\mathbb{E}(X_T) = \mathbf{p}_T \mathbf{B}_T^{-1} \boldsymbol{\epsilon}' = \frac{1}{\mu} \frac{1-\theta}{1-\theta^T} \frac{1-(\gamma\theta)^T}{1-\gamma\theta} \quad (2)$$

The so-called power-tail range of the TPT distribution is given in [Schwefel and Lipsky 1999]:

$$\text{Rng}(\mathbf{B}_T) = \frac{\gamma^{T-1}}{\mu} =: \bar{x}_T. \quad (3)$$

The following table shows the ratio $\text{Rng}(\mathbf{B}_T)/\mathbb{E}(X)$ for $\alpha = 1.4$ and $\theta = 0.5$ and for values of T , as used in the main part of the paper:

T	1	10	20
$\text{Rng}(\mathbf{B}_T) / \mathbb{E}(X)$	1	35.9	$4.46 \cdot 10^3$
...	30	40	50
...	$6.20 \cdot 10^5$	$8.74 \cdot 10^7$	$1.23 \cdot 10^{10}$

References

[Crovella and Bestavros 1996] Mark Crovella and Azer Bestavros: *Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes*. Proceedings of the ACM Sigmetrics, 1996.

- [Danzig et al.] P. Danzig, J. Mogul, V. Paxson, and M. Schwarz: *The Internet Traffic Archive*, <<http://ita.ee.lbl.gov/>>
- [Dumas and Simonian 2000] V. Dumas and A. Simonian: *Asymptotic bounds for the fluid queue fed by sub-exponential ON/OFF sources*. Advances in Applied Probability **32**, No 1. March 2000.
- [Greiner et al. 2000] Michael Greiner, Manfred Jobmann, and Lester Lipsky: *The Importance of Power-tail Distributions for Telecommunication Traffic Models*. Operations Research **47**, No. 2, pp 313-326, March 1999.
- [Heffes 1980] H. Heffes: *A class of data traffic processes - covariance function characterization and related queueing results*. Bell Syst. Tech. J. 59, pp. 897-930. 1980.
- [Latouche and Ramaswami 1993] Guy Latouche and V. Ramaswami: *A Logarithmic Reduction Algorithm for Quasi-Birth-Death Processes*. Journal of Applied Probability **30**, pp. 650-674, 1993.
- [Leland et al. 1994] Will Leland, Murad Taqqu, Walter Willinger, and Daniel V. Wilson: *On the Self-Similar Nature of Ethernet Traffic (Extended Version)*. Proc. of IEEE/ACM Trans. on Networking, **2**, 1. February 1994.
- [Lipsky 1992] Lester Lipsky: *QUEUEING THEORY: A Linear Algebraic Approach*. MacMillan Publishing Company, New York, 1992. 2nd edition to appear in Springer.
- [Meier-Hellstern and Fischer 1992] Kathy Meier-Hellstern and Wolfgang Fischer: *MMPP Cookbook*. Performance Evaluation **18**, p. 149-171. 1992.
- [Melamed 1991] B. Melamed: *TES: A class of methods for generating autocorrelated uniform variates*. ORSA Journal on Computing **3**, no. 4, pp. 317-329. 1991.
- [Neuts 1981] Marcel Neuts: *MATRIX-GEOMETRIC SOLUTIONS IN STOCHASTIC MODELS*. John Hopkins University Press, London, 1981.
- [Neuts 1989] Marcel Neuts: *STRUCTURED STOCHASTIC MATRICES OF M/G/1 TYPE AND THEIR APPLICATION*. Dekker, New York, 1989.
- [Norros 1995] I. Norros: *On the Use of Fractional Brownian Motion in the Theory of Connectionless Networks*. IEEE JSAC **13**, no. 6, pp. 953-962. 1995.
- [Schwefel and Lipsky 1999] Hans-Peter Schwefel and Lester Lipsky: *Impact of aggregated, self-similar ON/OFF traffic on delay in stationary queueing models (extended version)*. Performance Evaluation **43**, pp. 203-221, 2001.
- [Schwefel 2000] Hans-Peter Schwefel, *Performance Analysis of Intermediate Systems Serving Aggregated ON/OFF Traffic with Long-Range Dependent Properties*, PhD Dissertation, Technische Universität München, 2000.
- [Veres and Boda 2001] Andras Veres, and Miklós Boda, *The Chaotic Nature of TCP Congestion Control*. Proceedings of the ACM Sigcomm, 2001.
- [Willinger et al. 1995] Walter Willinger, Murad Taqqu, Robert Sherman, and Daniel Wilson: *Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level (Extended Version)*. Proceedings of the ACM Sigcomm, 1995.

AUTHOR BIOGRAPHIES

HANS-PETER SCHWEFEL is an Associate Professor and head of the IP Labs at Aalborg University, Denmark. His research focuses on IP based wireless networks with main interest in performance and dependability aspects. Before he joined Aalborg University, he was a project manager at Siemens Information and Communication Mobile, supervising research projects and responsible for the development of technical concepts for next generation mobile networks. He obtained his doctoral degree in the area of IP traffic and performance modeling from the Technical University in Munich, Germany. For his research activities, he also spent extended periods of time at the University of Florence, Italy, the University of Connecticut, USA, and at ATT Labs, USA.

IMAD ANTONIOS received his Ph.D. in Computer Science from the University of Connecticut under the guidance of Lester Lipsky. In 2003 he joined the Department of Computer Science at Southern Connecticut State University, USA, as Assistant Professor. His research interests are in the performance and dependability modeling of systems.


LESTER LIPSKY received his Ph.D. in Theoretical Atomic Physics from the University of Connecticut in 1965. In 1968 he joined the Department of Computer Science at the University of Nebraska, Lincoln, and became Full Professor in 1976. Presently, (as of 1987) he is Professor of Computer Science and Engineering at the University of Connecticut, becoming Emeritus in 2005. Professor Lipsky's major research interests are in the analytical modeling of the performance of computer systems and networks, and related mathematical and numerical problems. He has authored or co-authored over 150 research articles. His book, "QUEUEING THEORY: A Linear Algebraic Approach" was published by Macmillan, NY in 1992. The second edition will be published sometime in 2007 by Springer-Verlag.

Should Correlation Properties Govern Traffic Model Design?

Hans-Peter Schwefel
Imad Antonios
Lester Lipsky

Traffic is self-similar

- General Web traffic is *Self-Similar* [Leland '94]
- Symptoms: high variability over wide range of time scales, long-range dependence



Why is traffic self-similar?

- User Behavior
 - File-size distribution is heavy-tailed [Garg '92, Crovella '96, Cunha '95, Woodruff '96]
 - Transmission times and user think time distributions [Mogul '95]
- Protocol Behavior
 - TCP congestion control mechanisms, irrespective of file-size distribution, produces self-similarity [Veres '00]
- Internet Topology
 - Router in- and out-degrees are governed by power-laws [Pitkow '98]

Self-similarity and long-range dependence

- Long-range dependence (LRD) refers to the structure of the auto-correlation function lag k , $r(k)$
- LRD processes have slow-decaying $r(k)$ as opposed to SRD processes whose $r(k)$ decays exponentially



- New models were needed since the Poisson assumption is no longer applicable

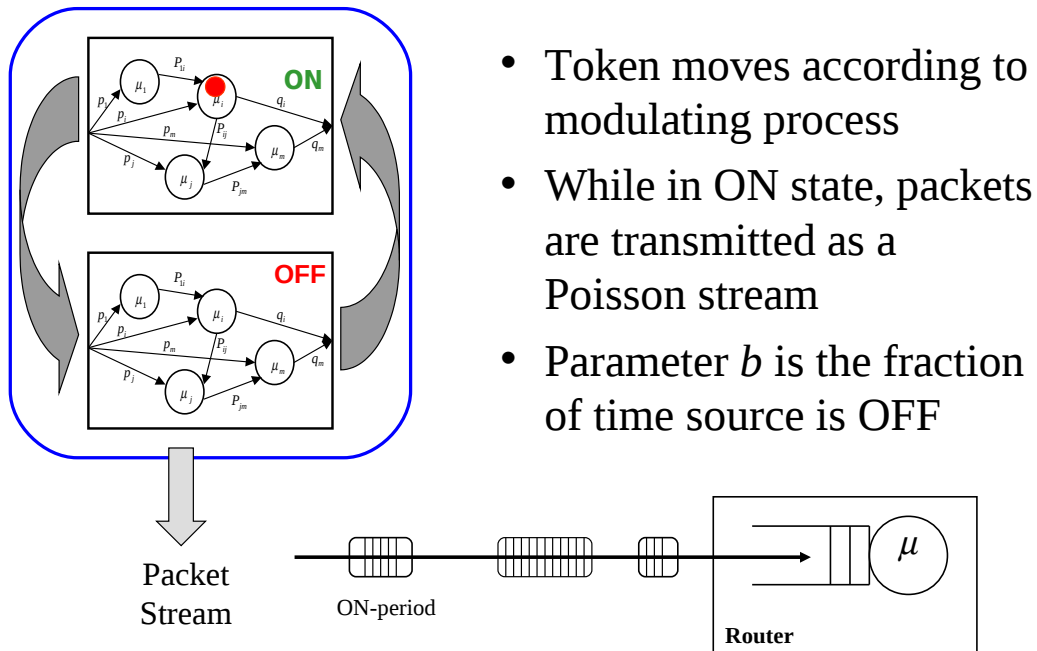
LRD Models

- Statistical approach
 - Based on fractional Gaussian noise (FGN) to generate synthetic traffic [Paxson '97, Ledesma '00]
- Fluid queueing model
 - Traffic modeled as a continuous flow [Dumas '00]
- Markovian models
 - In [Nain '01], semi-Markov process as input to queue, but results were very limited
 - [Schwefel and Lipsky '99-00] developed a model of N independent heavy-tailed sources (N -Burst model)

Is LRD really the culprit?

- *“It comes out that LRD is a good parameter for quantifying the level of QoS a network can provide in the transmission of the considered traffic: the higher the LRD, the worse the QoS [13] [15].”* -
Published in **ICCCN2005**
- Objectives of this work
 - Demonstrate that the presence of LRD is not a reliable predictor of performance
 - Suggest a different property to frame the problem

An On/Off traffic source: 1-Burst model




- Token moves according to modulating process
- While in ON state, packets are transmitted as a Poisson stream
- Parameter b is the fraction of time source is OFF

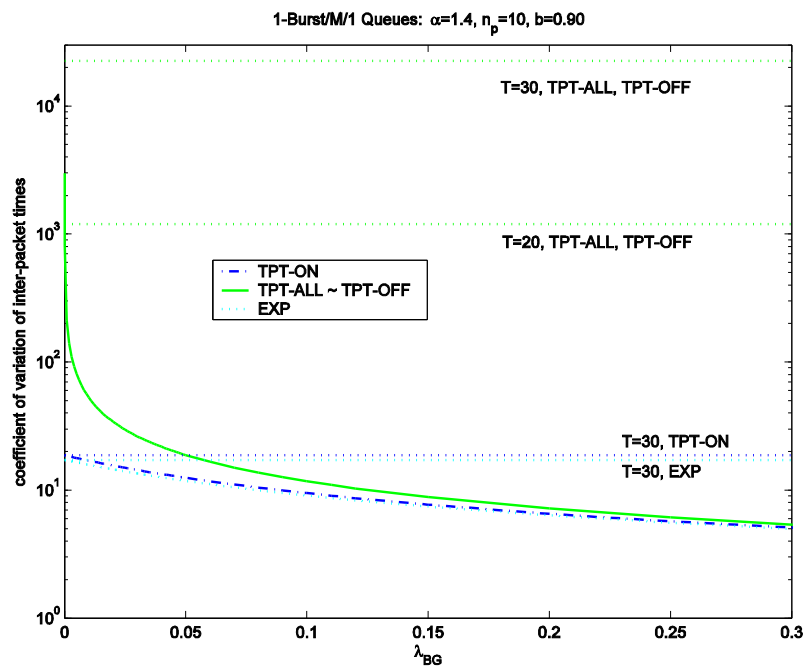
Model variations



- Both ON-time and OFF-time distribution exponential, or both Power tailed (TPT)
- Truncated power-tailed ON time and exponential OFF time, or vice versa
- Background Poisson process
- Models are MMPP's with their solutions formulated using a Linear Algebraic Queueing Theory approach

Coefficient of variation of counting process




Coefficient of variation of inter-packet times




$r(k)$ of inter-arrivals – no background traffic $r(k)$ of inter-arrivals – with background traffic

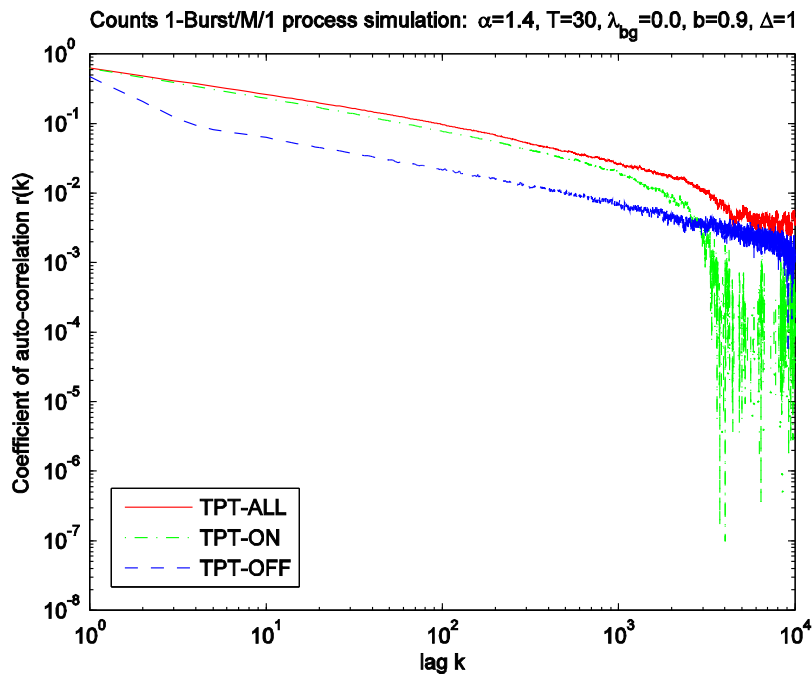
$r(k)$ of counting process – insensitive to background




Queueing delays with or without background traffic



Simulation results – no background traffic



More simulations



Summary of autocorrelation structure

model / scenario	$\lambda_{bg} = 0$	$\lambda_{bg} \rightarrow 0$	$\lambda_{bg} > 0$
inter-packet times			
EXP	0	0	SR
TPT-OFF	0	LRD	LRD
TPT-ON	LRD	LRD	LRD
TPT-ALL	LRD ²	LRD	LRD
counting process			
EXP	SR	SR	SR
TPT-OFF	LRD	LRD	LRD
TPT-ON	LRD	LRD	LRD
TPT-ALL	LRD	LRD	LRD

Summary and Conclusions

- Investigated the relationship between correlation structure of different input processes and router delay
- Presented models that produce LRD in inter-packet times and/or counting process that do not exhibit poor performance
- Suggested that ON-time distribution is a more reliable way to frame router performance

Simulations With Heavy-Tailed Workloads

MARK E. CROVELLA and LESTER LIPSKY*

Abstract

Recent evidence suggests that some characteristics of computer and telecommunications systems may be well described using *heavy tailed* distributions — distributions whose tail declines like a power law, which means that the probability of extremely large observations is non-negligible. For example, such distributions have been found to describe the lengths of bursts in network traffic and the sizes of files in some systems. As a result, system designers are increasingly interested in employing heavy-tailed distributions in simulation workloads. Unfortunately, these distributions have properties considerably different from the kinds of distributions more commonly used in simulations; these properties make simulation stability hard to achieve. In this paper we explore the difficulty of achieving stability in such simulations, using tools from the theory of stable distributions. We show that such simulations exhibit two characteristics related to stability: slow convergence to steady state, and high variability at steady state. As a result, we argue that such simulations must be treated as effectively

*Mark Crovella is with the Department of Computer Science, Boston University and Lester Lipsky is with the Department of Computer Science and Engineering, University of Connecticut. This research was supported in part by NSF grant CCR-9501822 and by a grant from Hewlett-Packard Company.


always in a transient condition.

1 Introduction


Recently the phenomenon of network traffic *self-similarity* has received significant attention in the networking community [10]. Asymptotic self-similarity refers to the condition in which a timeseries's autocorrelation function declines like a power-law, leading to positive correlations among widely separated observations. Thus the fact that network traffic often shows self-similarity means that it shows noticeable bursts at a wide range of time scales—typically at least four or five orders of magnitude. A related observation is that file sizes in some systems have been shown to be well described using distributions that are *heavy tailed*—distributions whose tails follow a power-law—meaning that file sizes also often span many orders of magnitude [3].

Heavy tailed distributions behave quite differently from the distributions more commonly used to describe characteristics of computing systems, such as the Normal distribution and the exponential distribution, which have tails that decline exponentially (or faster). In contrast, because their tails decline relatively slowly, the probability of very large observations occurring when sampling random variables that follow heavy tailed distributions is non-negligible. In fact, the distributions we discuss in this chapter have *infinite variance*, reflecting the extremely high variability that they capture.

As a result, designers of computing and telecommunication systems are increasingly interested in employing heavy-tailed distributions to generate workloads for use in simulation. However, simulations employing such workloads may show unusual characteristics; in particular, they may be much less stable than simulations with less variable inputs. In this chapter we discuss the kind of instability that may be expected in simulations with heavy-tailed inputs and show that they may exhibit two features: first, they will be very slow to con-



(a)



(b)

Figure 1: Sample Data from Heavy Tailed Distribution with $\alpha = 1.2$

verge to steady state; and second, they will show highly variable performance at steady state. To explain and quantify these observations we rely on the theory of *stable* distributions [4, 15].

These problems are not unique to simulation of telecommunications systems, arising also in risk and insurance modeling [2]. Solutions to certain aspects of these problems have been proposed drawing on rare event simulation and variance reduction techniques [8, 14].

In general however many of the problems associated with the simulations using heavy-tailed workloads seem quite difficult to solve. This chapter does not primarily suggest solutions but rather draws attention to these problems, both to yield insight for researchers using simulation and to suggest areas in which more research is needed. As a result we conclude with a summary of the issues that should be addressed when using simulations with heavy-tailed workloads.

2 Heavy Tailed Distributions

2.1 Background

Let X be a random variable with cdf $F(x) = P[X \leq x]$ and complementary cdf (ccdf) $\bar{F}(x) = 1 - F(x) = P[X > x]$. We say here that a distribution $F(x)$ is *heavy tailed* if

$$\bar{F}(x) \sim cx^{-\alpha} \quad 0 < \alpha < 2 \quad (1)$$

for some positive constant c , where $a(x) \sim b(x)$ means $\lim_{x \rightarrow \infty} a(x)/b(x) = 1$. (We note that more general definitions of heavy tails are common; see for example [6].) If $F(x)$ is heavy tailed then X shows very high variability. In particular, X has infinite variance, and if $\alpha \leq 1$, X has infinite mean. Section 2.2 will explore the implications of infinite moments in practice; here we note simply that if $\{X_i, i = 1, 2, \dots\}$ is a sequence of observations of X then the sample

variance of $\{X_i\}$ as a function of i will tend to grow without limit, as will the sample mean if $\alpha \leq 1$.

The simplest heavy tailed distribution is the *Pareto* distribution which is power-law over its entire range. The Pareto distribution has pmf

$$p(x) = \alpha k^\alpha x^{-\alpha-1} \quad 0 < k \leq x$$

and cdf

$$F(x) = P[X \leq x] = 1 - (k/x)^\alpha \quad (2)$$

in which the positive constant k represents the smallest possible value of the random variable.

In practice, random variables that follow heavy tailed distributions are characterized as exhibiting many small observations mixed in with a few large observations. In such datasets, most of the observations are small, but most of the contribution to the sample mean or variance comes from the few large observations.

This effect can be seen in Figure 1, which shows 10,000 synthetically generated observations drawn from a Pareto distribution with $\alpha = 1.2$ and mean $\mu = 6$. In Figure 1(a) the scale allows all observations to be shown; in Figure 1(b) the y axis is expanded to show the region from 0 to 200. These figures show the characteristic, visually striking behavior of heavy tailed random variables. From the left plot it is clear that a few large observations are present, some on the order of hundreds to one thousand; while from the right plot it is clear that most observations are quite small, typically on the order of tens or less.

An example of the effect of this variability on sample statistics is shown in Figure 2. This figure shows the running sample mean of the data points from Figure 1, as well as a level line showing the mean of the underlying distribution (6). Note that the sample mean starts out well below the distributional mean, and that even after 10,000 observations it is not close in relative terms to the




Figure 2: Running Mean of Data from Figure 1

distributional mean.

2.2 Heavy Tails in Computing Systems

A number of recent studies have shown evidence indicating that aspects of computing and telecommunication systems can show heavy tailed distributions. Measurements of computer network traffic have shown that autocorrelations are often related to heavy tails; this is the phenomenon of *self similarity* [5, 10]. Measurements of file sizes in the Web [1, 3] and in I/O patterns [13] have shown evidence that file sizes can show heavy tailed distributions. And finally, the CPU time demands of Unix processes have also been shown to follow heavy tailed distributions [9, 7].

The presence of heavy tailed distributions in measured data can be assessed in a number of ways. The simplest is to plot the cdf on log-log axes, and visually inspect the resulting curve for linearity over a wide range (several orders of magnitude). This is based on Equation 1, which can be recast as:

$$\lim_{x \rightarrow \infty} \frac{d \log \bar{F}(x)}{d \log x} = -\alpha$$




Figure 3: Log-log Complementary Distribution of Sizes of Files Transferred Through the Web

so that for large x , the cdf of a heavy tailed distribution should appear to be a straight line on log-log axes with slope $-\alpha$.

An example empirical dataset is shown in Figure 3, which is taken from [3]. This figure is the cdf of file sizes transferred through the network due to the Web, plotted on log-log axes. The figure shows that the file size distribution appears to show power law behavior over approximately three orders of magnitude. The slope of the line fit to the upper tail is approximately -1.2 , yielding $\hat{\alpha} \approx 1.2$.

3 Stability In Systems With Heavy Tailed Workloads

As heavy tailed distributions are increasingly used to characterize workload characteristics of computing systems, researchers interested in simulating such systems are beginning to use heavy tailed inputs to simulations. For example,

[12] describes methods for generating self-similar time series for use in simulating network traffic and [11] uses heavy-tailed file sizes as inputs to a network simulation. However, an important question arises: how stable are such simulations? This can be broken down into two questions:

1. How long until such simulations reach steady state?, and
2. How variable is system performance at steady state?

In this section we will show that if simulation outputs are dependent on all the moments of the distribution F then the answers to the above questions can be surprising. Essentially, we show that such simulations can take a very long time to reach steady state; and that such simulations can be much more variable at steady state than is typical for traditional systems.

Note that some simulation statistics may not be directly affected by all the moments of the distribution F , and our conclusions do not necessarily apply to those cases. For example, the mean number of customers in an $M/G/\infty$ queueing system may not show unusual behavior even if the service time distribution F is heavy tailed because that statistic only depends on the mean of F .

Since not all simulation statistics will be affected by heavy tailed workloads, we choose a simple statistic to show the generality of our observations: the sample mean of the heavy tailed inputs. Since our results apply to the sample mean of the input, we expect that any system property that behaves like the sample mean should show similar behavior. For example, assume we want to achieve steady state in a particular simulation. This implies that the measured system utilization $\lambda\bar{x}$ (where λ^{-1} is the average interarrival time and \bar{x} is the sample mean of service times over some period) should be close to the desired system utilization ρ . For this to be the case, \bar{x} must be close to its desired mean μ .

To analyze the behavior of the sample mean, we are concerned with the convergence properties of sums of random variables. The normal starting point for such discussions would be the Central Limit Theorem (CLT). Unfortunately,

the CLT applies only to sums of random variables with finite variance, and so does not apply in this case. In the place of the CLT we instead have limit theorems for heavy tailed random variables first formulated by Lévy [4, 15].

To introduce these results we need to define the notation $A \xrightarrow{d} B$ which means that the random variable A converges in distribution to B (roughly, has distribution B for large n). Then the usual CLT can be stated as: for X_i i.i.d. and drawn from some distribution F with mean μ and variance $\sigma^2 < \infty$, define

$$A_n = \frac{1}{n} \sum_{i=1}^n X_i$$

and

$$Z_n = n^{1/2}(A_n - \mu); \quad (3)$$

then

$$Z_n \xrightarrow{d} \mathcal{N}(0, \sigma^2) \quad (4)$$

where $\mathcal{N}(0, \sigma^2)$ is a Normal distribution.


However, when X_i are i.i.d. and drawn from some distribution F that is heavy tailed with tail index $1 < \alpha < 2$, then if we define

$$Z_n = n^{1-1/\alpha}(A_n - \mu) \quad (5)$$

we find that

$$Z_n \xrightarrow{d} \mathcal{S}_\alpha \quad (6)$$


where \mathcal{S}_α is an α -Stable distribution. The α -Stable distribution has four parameters: α , a location parameter (analogous to the mean), a scale parameter (analogous to the standard deviation), and a skewness parameter. Based on the value of the last parameter, the distribution can be either skewed or symmetric. A plot of the symmetric α -Stable distribution with $\alpha = 1.2$ and location zero is shown in Figure 4. From the figure it can be seen that this distribution has a bell-shaped body much like the Normal distribution but that it has much heavier tails. In fact the α -Stable distribution has power-law tails that follow

Figure 4: Pmf of an α -Stable Distribution


the same α as that of the distribution F from which the original observations were drawn.

From Equations 5 and 6 we can make two observations about the behavior of sums of heavy tailed random variables. First, Equation 5 states that such sums may converge much more slowly than is typical in the finite variance case. Second, Equation 6 states that even after convergence, the sample mean will show high variability—it follows a heavy tailed distribution.

These effects can be seen graphically in Figure 5. This figure shows histograms of A_n for varying values n . On the left we show the case in which the X_i s were drawn from an Exponential distribution; on the right we show the case in which the X_i s were drawn from a strictly positive heavy tailed distribution with $\alpha = 1.4$; in both cases the mean of the underlying distribution was 1. The plot on the left shows that the most likely value of the sample mean is equal to the true mean, even when summing only a small number of samples. In addition, it shows that as one sums larger numbers of samples, the sample mean converges quickly to the true mean. However, neither of these observations are



(a)



(b)

Figure 5: Histogram of A_n as n varies for Exponential (a) and Heavy Tailed (b) Random Variables.

true for the case of the heavy tailed distribution on the right. When summing small numbers of samples, the most likely value of the sample mean is far from the true mean, and the distribution progresses to its final shape rather slowly.

Thus we have seen that the convergence properties of sums of heavy tailed random variables are quite different from those of finite variance random variables. We relate this to steady state in simulation as follows: presumably for a simulation to reach steady state, it must at a minimum have seen enough of the input workload to observe its mean. Of course it may be necessary for much more of the input to be consumed before the simulation reaches steady state, so this condition is a relatively weak one. Still, we show in the next two subsections that this condition has surprising implications for simulations.

3.1 Slow Convergence to Steady State

Equation 6 states that for large n , Z_n converges in distribution. Thus another way of formulating Equation 5 is:

$$|A_n - \mu| \sim n^{1/\alpha-1}.$$

In this form it is more clear how slowly A_n converges to μ . If α is close to 1, then the rate of convergence, measured as the difference between A_n and μ , is very slow—until, for $\alpha = 1$, the average does not converge at all, reflecting the fact that the mean is infinite.

Suppose one would like to use A_n to form a estimate of the mean μ that is accurate to k digits. Alternatively, one might state that a simulation has reached steady state when the observed mean of the input A_n agrees with μ to k digits. Then we would like

$$|A_n - \mu|/\mu \leq 10^{-k}.$$

Now, as a rough approximation:

$$|A_n - \mu| = c_1 n^{1/\alpha-1}$$

for some positive constant c_1 . Then we find that:

$$n \geq c_2 10^{\frac{k}{1-1/\alpha}}.$$

We can say that given this many samples, k digit accuracy is “highly likely.”

For example, assume we would like 2-digit accuracy in A_n , and suppose $c_2 \approx 1$. Then the number of samples n necessary to achieve this accuracy is shown in Table 1. This table shows that as $\alpha \rightarrow 1$, the number of samples necessary to obtain convergence in the sample mean explodes. Thus, it is *not feasible* in any reasonable amount of time to observe steady state in such a simulation as we have defined it. Over any reasonable time scale, such a simulation is *always in transient state*.

Table 1: Number of Samples Necessary to Achieve 2 Digit Accuracy in Mean as a Function of α

α	n
2.0	10,000
1.7	72,000
1.5	1,000,000
1.2	10^{12}
1.1	10^{22}

3.2 High Variability at Steady State

Equation 6 shows that even at steady state, the sample mean will be distributed according to a heavy tailed distribution, and hence will show high variability. Thus, the likelihood of an erroneous measurement of μ is still non-negligible. Equivalently, the simulation still behaves erratically.

To see this more clearly, let us define a *swamping* observation as one whose presence causes the estimate of μ to be at least twice as large as it should be.




Figure 6: Probability of a Swamping Observation in 10^5 Inputs as a Function of α

That is, if we happen to encounter a swamping observation in our simulation, the observed mean of the input will have a relative error of at least 100%.

In a simulation consisting of n inputs, a swamping observation must have value at least $n\mu$. Let us assume that the inputs are drawn from a Pareto distribution. Such a distribution has $\mu = k\alpha/(\alpha - 1)$. Then the probability $p_{n\mu}$ of observing a value of $n\mu$ or greater is

$$p_{n\mu} = P[X > n\mu] = \left(\frac{k}{nk\alpha/(\alpha - 1)} \right)^\alpha = \left(\frac{\alpha - 1}{n\alpha} \right)^\alpha$$

and the probability p of observing such a value at least once in n trials is

$$p = 1 - (1 - p_{n\mu})^n.$$

Figure 6 shows a plot of p as a function of α for $n = 10^5$. (The figure is not significantly different for other values of n , *e.g.*, $10^6, 10^7$.) It shows that even in a relatively long simulation, the probability of a swamping observation is not negligible; when α is below about 1.3, such an observation could occur more often than once in a hundred simulations. The probability declines very

rapidly for $\alpha < 1.1$ not because the variability of the simulation is declining, but because of the way we have defined the swamping observation: in terms of the distributional mean. When $\alpha = 1$, the mean is infinite, and so it becomes impossible to observe a value greater than the mean.

Taken together, Table 1 and Figure 6 also provide some insight into the value of α above which it may be possible to obtain convergent, consistent simulations. The table shows that simulation convergence becomes impractical when α is somewhere in the region between 1.7 and 1.5; and the Figure shows that simulations become erratic at steady state in approximately the same region. As a result, we can conclude that the difficulties inherent in simulations with heavy tailed inputs are likely to be particularly great when α is less than about 1.7; and that when α is greater than or equal to about 1.7 it may be feasible (given sufficient computing effort) to obtain consistent steady state in simulation.

4 Conclusions

We have shown that a difficult problem arises when simulating systems with heavy tailed workloads. In such systems, steady-state behavior can be elusive, because average-case behavior depends on the presence of many small observations as well as a few large observations.

This problem has two implications. First, since a number of large but rare observations must occur before average case behavior is evident, convergence of a simulation to steady state may be slow. It may not be possible in any reasonable time to achieve steady state. Second, since many small observations must occur to balance the presence of large observations, large observations can have a dominating effect on performance results even at steady state. Simulations may still behave erratically even at steady state.

References

- [1] Martin F. Arlitt and Carey L. Williamson. Web server workload characterization: The search for invariants. *IEEE/ACM Transactions on Networking*, 5(5), 1997.
- [2] S. Asmussen and K. Binswanger. Simulation of ruin probabilities for subexponential claims. Preprint, 1995.
- [3] Mark E. Crovella and Azer Bestavros. Self-similarity in World Wide Web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, December 1997.
- [4] William Feller. *An Introduction to Probability Theory and Its Applications*, volume II. John Wiley and Sons, second edition, 1971.
- [5] Mark Garrett and Walter Willinger. Analysis, modeling, and generation of self-similar VBR video traffic. In *Proceedings of ACM SIGCOMM '94*, September 1994.
- [6] Charles M. Goldie and Claudia Kluppelberg. Subexponential distributions. In Robert J. Adler, Raisa E. Feldman, and Murad S. Taqqu, editors, *A Practical Guide To Heavy Tails*, pages 435–460. Chapman & Hall, New York, 1998.
- [7] M. Harchol-Balter and A. Downey. Exploiting process lifetime distributions for dynamic load balancing. *ACM Transactions on Computer Systems*, 15(3), 1997.
- [8] P. Heidelberger. Fast simulation of rare events in queueing and reliability models. *Lecture Notes in Computer Science*, 729:165–202, 1993.
- [9] W. E. Leland and T. J. Ott. Load-balancing heuristics and process behavior. In *Proceedings of Performance and ACM Sigmetrics*, pages 54–69, 1986.

-
- [10] W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2:1–15, 1994.
 - [11] Kihong Park, Gi Tae Kim, and Mark E. Crovella. On the relationship between file sizes, transport protocols, and self-similar network traffic. In *Proceedings of the Fourth International Conference on Network Protocols (ICNP'96)*, pages 171–180, October 1996.
 - [12] Vern Paxson. Fast approximation of self-similar network traffic. Technical Report LBL-36750, Lawrence Berkeley National Laboratory, April 30 1995.
 - [13] David L. Peterson. Data center I/O patterns and power laws. In *CMG Proceedings*, December 1996.
 - [14] R. Y. Rubinstein and B. Melamed. *Efficient Simulation and Monte Carlo Methods*. John Wiley and Sons, Inc., 1997.
 - [15] Gennady Samorodnitsky and Murad S. Taqqu. *Stable Non-Gaussian Random Processes*. Stochastic Modeling. Chapman and Hall, New York, 1994.

How does a physicist approach queueing theory? - A few examples

Yiping Ding
BMC Software

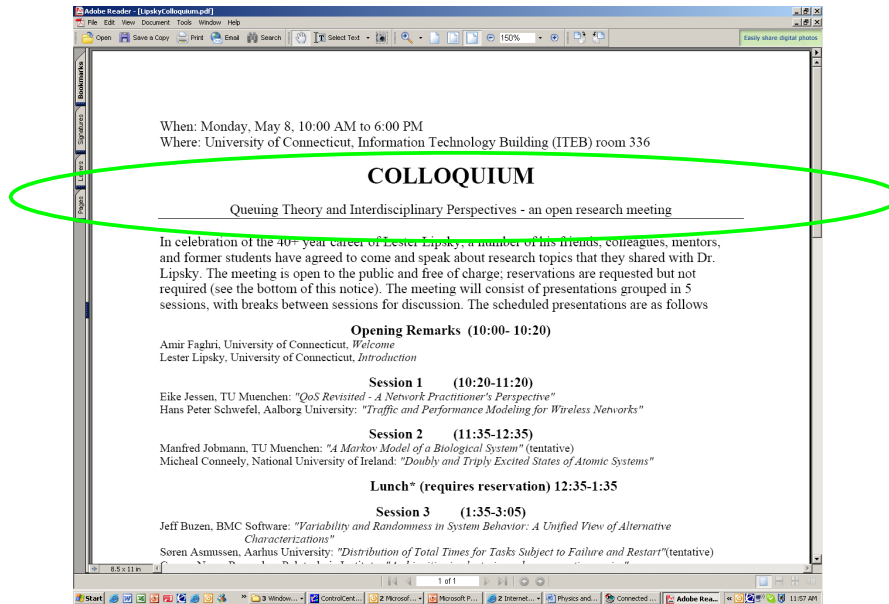


Professor Lester Lipsky is retiring ...

- However, he has students and will still be active:
- The server will not be idle if there are jobs in the queue.



Queueing Theory and Interdisciplinary Perspectives: Physics and Queueing Theory



Dr. Lipsky was and still is a physicist

- How does a physicist like Dr. Lipsky get into queueing theory?
- ◆ Waiting is a natural phenomenon
- ◆ People studying physics often have to carry out resource intensive computations
- ◆ For some physicists, the waiting becomes so painful that they decide to study why



A Physicist can provide a unique interdisciplinary perspective to queueing theory

- ◆ Use the mathematical tools they are familiar with
 - ◆ Use the thought experiments they are comfortable with
 - ◆ Use the experiment results to validate the theory
- Let's see some examples ...




A Linear Algebraic Approach to Queueing Theory

Yahoo! SEARCH Results for "A Linear Algebraic Approach" (Results 1 - 10 of about 640,000)

1. Find in a Library. Queueing theory. **a linear algebraic approach** ... Title: Queueing theory ... **a linear algebraic approach** - Author: Lester R Lipsky ... English • Document Type: Book • Subjects: Linear, Queueing theory, Algebra • Library Count: world5libraries.org/wcpcp/ow/c600a1448b70d4d419a6bd4d094526.html - [More from this site](#) - [Save](#)
2. **A modified linear algebraic approach to electron scattering using cubic splines** ... A modified **linear algebraic approach** to electron scattering using cubic splines A modified **linear algebraic approach** to the solution of the Schrödinger equation for ... stacks.iop.org/0022-3708/19/4/27 - [More from this site](#) - [Save](#)
3. **A Linear Algebraic Approach to Metering Schemes** ... **A Linear Algebraic Approach to Metering Schemes** **A Linear Algebraic Approach to Metering Schemes** A modeling agency is able to measure the interaction between servers and clients during a certain number of time ... cseeser.ist.psu.edu/459707.html - [More from this site](#) - [Save](#)
4. **A linear algebraic approach to electron-molecule collision** ... **A linear algebraic approach** to electron-molecule collision **A linear algebraic approach** to the solution of the Schrödinger equation for low-energy electron ... stacks.iop.org/0022-3708/14/101 - [More from this site](#) - [Save](#)
5. Queueing Theory **A Linear Algebraic Approach** - Book at Yahoo! Shopping ... Yahoo! Shopping is the best place to comparison shop for Queueing Theory. **A Linear Algebraic Approach** - Book. Compare products, compare prices, read reviews and merchant ratings: shopping.yahoo.com/Queueing Theory. **A Linear Algebraic Approach** 30... - [More from this site](#) - [Save](#)
6. **Algebraic Approach for Quasi-Linear Constrained Differential Equations** ... **Algebraic Approach for Quasi-Linear Constrained Differential Equations** **Algebraic Approach for Quasi-Linear Constrained Differential Equations** To study the existence of solutions to constrained quasi-linear differential equations with polynomial ... cseeser.ist.psu.edu/38002.html - [More from this site](#) - [Save](#)
7. **Linear functionals on idempotent spaces: an algebraic approach** ... **Linear functionals on idempotent spaces: an algebraic approach** **Linear functionals on idempotent spaces: an algebraic approach** In this paper, we present an **algebraic approach** to idempotent functional analysis, which is an abstract version of ... cseeser.ist.psu.edu/38002.html - [More from this site](#) - [Save](#)
8. **An Algebraic Approach to Linear-Optical Schemes for Deterministic Quantum Computing** ... **An Algebraic Approach to Linear-Optical Schemes for Deterministic Quantum Computing** **An Algebraic Approach to Linear-Optical Schemes for Deterministic Quantum Computing** Comment: new version with minor modifications **Linear-Optical Passive** ... cseeser.ist.psu.edu/38002.html - [More from this site](#) - [Save](#)
9. **An algebraic approach to time scale analysis of singularly perturbed linear systems** ... **An algebraic approach to time scale analysis of singularly perturbed linear systems** **An algebraic approach to time scale analysis of singularly perturbed linear systems** Bibliography, p. 41. Supported by the Air Force Office of Scientific Research ... hdl.handle.net/1721.1.20957 - [More from this site](#) - [Save](#)

“Journal of Physics B:
Atomic and Molecular Physics,”





DEDICATION

Physicists value
different ideas, different views,
and different approaches

A Path to Discovery

"Theories of the known which are described by different ideas, may be equivalent in all their predictions and are hence scientifically indistinguishable. However, they are not psychologically identical when trying to move from that base into the unknown. For different views suggest different kinds of modifications which might be made. Therefore, a good scientist today might find it useful to have a wide range of viewpoints and mathematical expressions of the same theory available to him. This may be asking too much of one man. Then new students should as a class have this. If every individual student follows the same current fashion in expressing and thinking about the generally understood areas, then the variety of hypotheses being generated to understand the still open problems is limited. Perhaps rightly so, ... BUT if it is in another direction, who will find it?"

- Richard P. Feynman

"So spoke an honest man, the outstanding intuitionist of our age and a prime example of what may lie in store for anyone who dares to follow the beat of a different drum."

- Julian Schwinger

From a special issue on Richard Feynman (who died on 15 February 1988) in: PHYSICS TODAY, February 1989. Feynman's quote (slightly paraphrased here) was taken from his Nobel Lecture in June 1965.

[Note: Feynman and Schwinger shared the Nobel prize with S. Tomonaga in 1965 for their work on Quantum Electrodynamics in the late forties. Working independently, and using radically different methods, they ended up with mathematically equivalent theories. Schwinger and Tomonaga were the "mainstreamers", but everyone calculates using Feynman's method to this day.]




Figure 2.3.1: Time-dependent state transition diagram for a closed M/M/1/N loop, describing the mean time $\tau_u(n)$ for a queue to grow by one customer.

it will occur written near it. For instance, the system can go from $n \rightarrow n+1$ in one step, with probability $\lambda/(\mu+\lambda)$, with a mean time delay of $1/(\mu+\lambda)$. A wavy arrow corresponds to the sum of all possible ways the system can get from the tail to the head for the first time, irrespective of the number of transitions taken. Thus the arrow labeled " $\tau_u(n)$ " includes not only the direct transition ($n \rightarrow n+1$), but also ($n \rightarrow n-1 \rightarrow n \rightarrow n+1$), and ($n \rightarrow n-1 \rightarrow n-2 \rightarrow n-1 \rightarrow n \rightarrow n+1$), and the infinite number of other sequences that eventually lead to $n+1$.

Our ability to represent an infinite number of sequences by a single symbol is the key to setting up a soluble set of recursive relations. If the system starts with n at S_1 , an event will occur in mean time $1/(\mu+\lambda)$. That event can be one of two things. Either the queue will go directly to $n+1$, or it will drop to $n-1$, in which case it will take time $\tau_u(n-1)$ to get back to n , and a further $\tau_u(n)$ to finally get to $n+1$. Mathematically we can write

$$\tau_u(n) = \frac{\lambda}{\mu+\lambda} \cdot \frac{1}{\mu+\lambda} + \frac{\mu}{\mu+\lambda} \left[\frac{1}{\mu+\lambda} + \tau_u(n-1) + \tau_u(n) \right],$$

where $\tau_u(0) = 1/\lambda$. For convenience, drop the subscript u when no confusion is likely to arise. The two terms without a τ in them combine to yield the following:

$$\tau(n) = \frac{1}{\mu+\lambda} + \frac{\mu}{\mu+\lambda} [\tau(n-1) + \tau(n)]. \tag{2.3.1a}$$

We interpret this as follows. It takes a mean time of $1/(\mu+\lambda)$ for something to happen. If the event was an arrival, we are done. The probability that it was not an arrival is $\mu/(\mu+\lambda)$, in which case the queue will have dropped back to $n-1$ and take a mean time of $[\tau(n-1) + \tau(n)]$ to first get back to n and then to $n+1$.

Note that $\tau(n)$ appears on both sides of the equation, indicating that the system got back to where it started, and that is what we mean by a regenerative process. We will derive equations for more complicated processes in just this way, so the reader should expect to return to this section for reference.

We next solve for $\tau(n)$ and get the following recursive equation.

$$\tau(n) = \frac{1}{\lambda} [1 + \mu\tau(n-1)], \quad n > 0, \tag{2.3.1b}$$

and $\tau(0) = 1/\lambda$. By direct substitution into (2.3.1b) it follows that $\tau(1) = (1+1/\rho)/\lambda$ and $\tau(2) = (1+1/\rho+1/\rho^2)/\lambda$. One can guess that the general expression for $\tau(n)$ is

$$\tau(n) = \frac{1}{\lambda} \sum_{j=0}^n \frac{1}{\rho^j}, \tag{2.3.2a}$$

which can be proven by induction to be the solution of (2.3.1b)[†]. Equation (2.3.2a) is the well-known (certainly by now) geometric series, for which a closed-form expression exists.

$$\tau(n) = \frac{1}{\lambda} \sum_{j=0}^n \frac{1}{\rho^j} = \frac{1/\mu}{1-\rho} \left(\frac{1}{\rho^{n+1}} - 1 \right) \quad \text{for } \rho \neq 1 \tag{2.3.2b}$$

and

$$\tau(n) = \frac{n+1}{\mu} \quad \text{for } \rho = 1. \tag{2.3.2c}$$


We are now ready to find the time it takes for a queue to grow to a given length.

Definition 2.3.2

$\tau_u(0 \rightarrow n)$:= mean first-passage time for the queue at S_1 to grow from 0 to n customers. The queue could drop to 0 many times before finally reaching the goal. \diamond

[†] We interject a word or two about *guessing*. If science was merely a sequence of deductions, we all would already be replaced by computers. Research is a creative process. The imaginative scientist, mathematician, or engineer *plays* with the tools of the trade and regularly makes guesses at what is correct. (These guesses are often credited to *intuition*.) Most guesses that prove wrong never come to public light. You, the reader, only see the successes and thus may think that there is some secret process going on to which you will never be privy. Nonsense. The creative person who plays long enough with the relevant material will ultimately make many correct guesses. Remember, proof by induction does not require that we defend the source of the guess. It must only prove that the guess is correct (if it is).

How did a physicist handle a similar situation ?



Stephen Hawking's phenomenal, multi-copy bestseller, *A Brief History of Time*, introduced the ideas of this brilliant theoretical physicist to readers all over the world. Now, in a major publication event, Hawking returns with a lavishly illustrated sequel that unravels the mysteries of the breakthroughs that have occurred in the years since the release of his acclaimed first book.

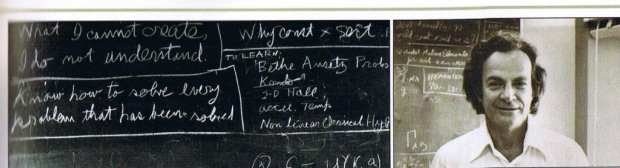


THE UNIVERSE IN A NUTSHELL

ONE OF THE MOST INFLUENTIAL THINGS of our time, Stephen Hawking is an intellectual icon, known not only for the adventurousness of his ideas but for the clarity and wit with which he expresses them. In this new book Hawking returns to the cutting edge of theoretical physics. The truth is often stranger than fiction, to explain the universe in men's terms the principles that control our universe. Like many in the community of theoretical physicists, Professor Hawking is seeking to uncover the grail of science—the elusive Theory of Everything that lies at the heart of the cosmos. In an accessible and often playful style, he guides the reader through his search to uncover the secrets of the universe, from supergravity to supersymmetry, from quantum field theory to M-theory, from holography to dual string theory. This book takes us to the wild frontiers of science, where superstring theory and p-branes may hold the key to the puzzle. And he lets us behold some of the most exciting intellectual



at a jagged edge where space and time were crumpled together. If the universe was infinite, it would be very difficult to define boundary conditions. ... a colleague named Jim Hartle and I realized there was a possibility. Maybe the universe has no boundary in space and time. This seems to be in direct contradiction with the Penrose and I proved, which showed that the universe had a beginning, a boundary in time. However, as I explained in chapter 2, there is another kind of time, called imaginary time. If you look at the ordinary real time that we feel going by, the universe in real time determines its history in space, and vice versa, but the two kinds of history can be



The blackboard at Caltech at the time of Feynman's death in 1988. Richard Feynman.

FEYNMAN STORIES

Born in Brooklyn, New York, in 1918, Richard Feynman completed his Ph.D. under John Wheeler at Princeton University in 1942. Shortly afterward, he was drawn into the Manhattan Project. There he was known for both his exuberant personality and practical jokes—at the Los Alamos labs, he enjoyed cracking the top-secret safes—and for being an exceptional physicist: he became a key contributor to atomic bomb theory. Feynman's perpetual curiosity about the world was the very root of his being. It was not only the engine for his scientific success, it led him to numerous astonishing exploits, such as deciphering Mayan hieroglyphics.


In the years following World War II, Feynman found a powerful new way of thinking about quantum mechanics, for which he was awarded the Nobel Prize in 1965. He challenged the basic classical assumption that each particle has one particular history. Instead, he suggested that particles travel from one location to

another along every possible path through spacetime. With each trajectory Feynman associated two numbers, one for the size—the amplitude—of a wave and one for its phase—whether it is at a crest or a trough. The probability of a particle going from A to B is found by adding up the waves associated with every possible path that passes through A and B.

Nevertheless, in the everyday world it seems to us that objects follow a single path between their origin and their final destination. This agrees with Feynman's multiple history (or sum-over-histories) idea, because for large objects his rule for assigning numbers to each path ensures that all paths but one cancel out when their contributions are combined. Only one of the infinity of paths matters as far as the motion of macroscopic objects is concerned, and this trajectory is precisely the one emerging from Newton's classical laws of motion.

very different. In particular, the universe need have no beginning or end in imaginary time. Imaginary time behaves just like another direction in space. Thus, the histories of the universe in imaginary time can be thought of as curved surfaces, like a ball, a plane, or a saddle shape, but with four dimensions instead of two (see Fig. 3.9, page 84).

If the histories of the universe went off to infinity like a saddle or a plane, one would have the problem of specifying what the boundary conditions were at infinity. But one can avoid having to specify boundary conditions at all if the histories of the universe in imaginary time are closed surfaces, like the surface of the Earth. The surface of the Earth doesn't have any boundaries or edges. There are no reliable reports of people falling off.



In Feynman's path integral a particle takes every possible path.

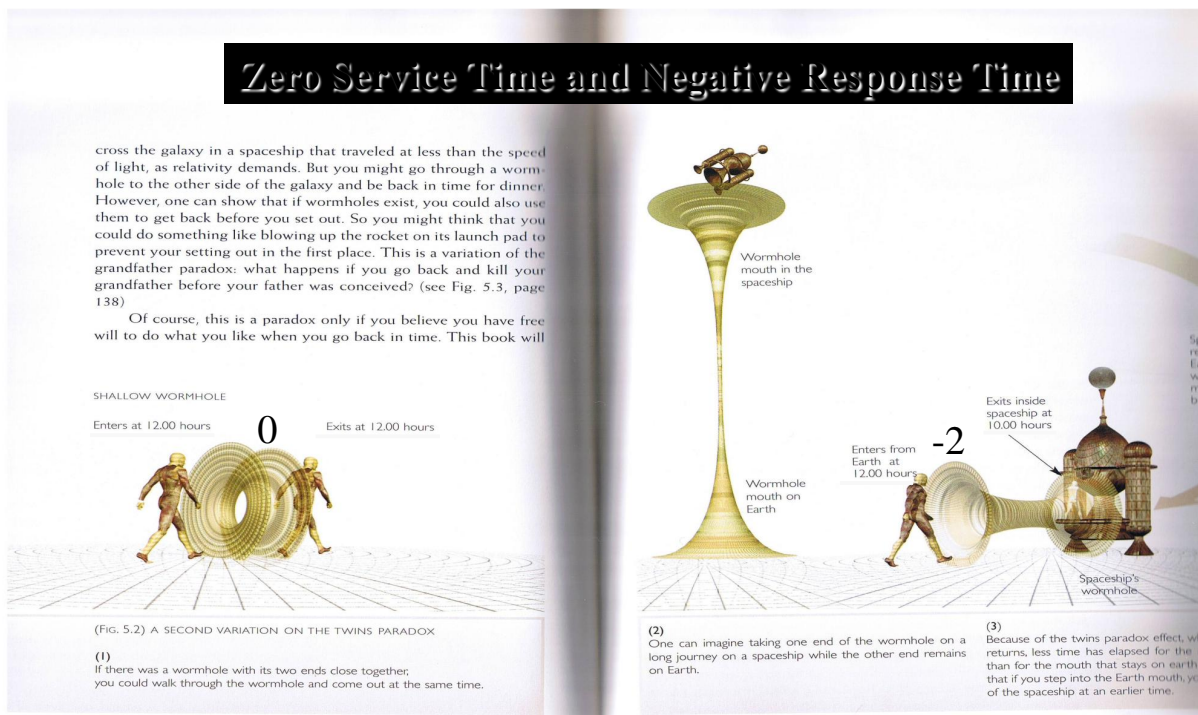
Some fundamental laws are very similar

- Examples


- The quantized energy of the photons is
- $E = hf$,
- where h is the Planck's constant and f is the radiation frequency of the photons. This is very much like a utilization formula in queueing theory,
- $U = s x$,
- if we view the Planck's constant as a "service time," s , and the radiation frequency as the throughput, x :
- How busy you are in a given interval, U , is pretty much the same concept as how much energy that you have!



Queueing theory puts wormholes in perspective !



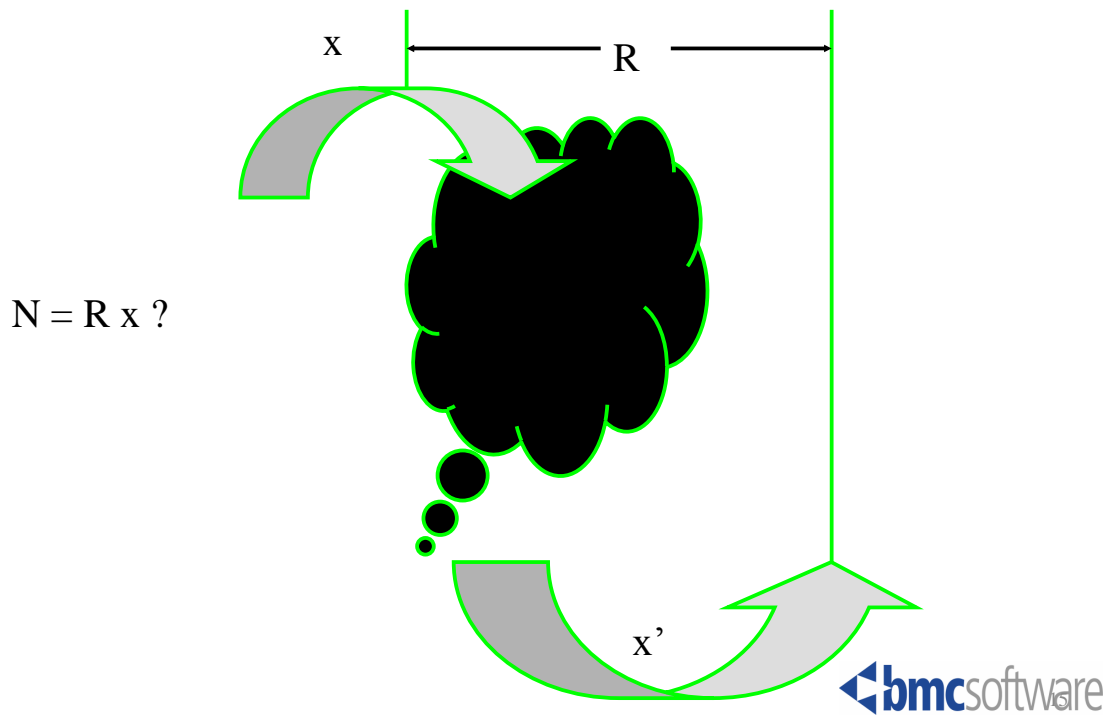
A customer arrives before it departs ...



Physical Laws and Queueing Theory

- ◆ Will computer performance and capacity still be a challenge to us in the future?
 - It will be in our lifetime
- ◆ What are the performance implications of a quantum computer?
 - 1kg has $E = mc^2 = 8.9874 \times 10^{16}$ joules energy
 - A 1kg computer can perform 5.4258×10^{50} ops/sec
- ◆ Where will bottlenecks reside in future computers?
 - The speed of light is limited: 2.9979×10^8 meters / sec.
- ◆ Will performance / queueing theory related formulas and algorithms hold in the future?
 - Some will and some will not

Little's Law and black holes



Dr. Lipsky:
We wish you the best!

Signal Probabilities in AND-OR Trees

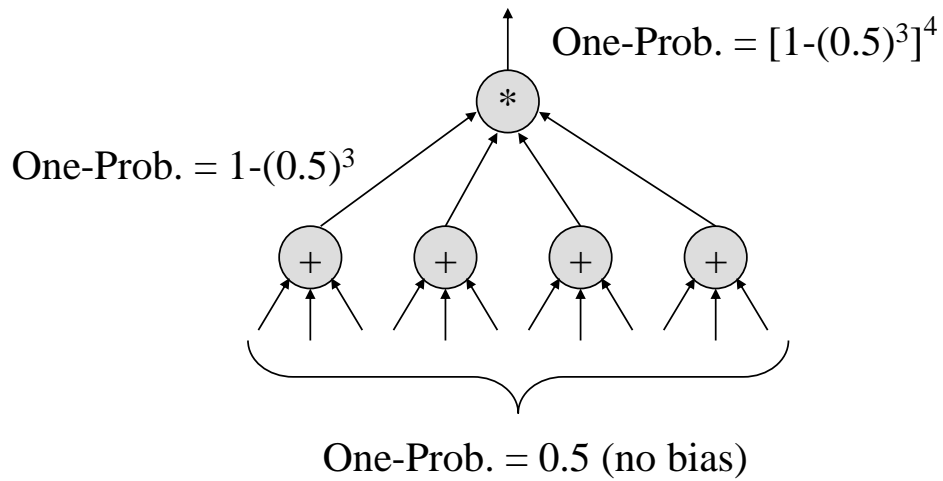
by
Sharad Seth

In Celebration of
Lester Lipsky
Great Mentor and Friend

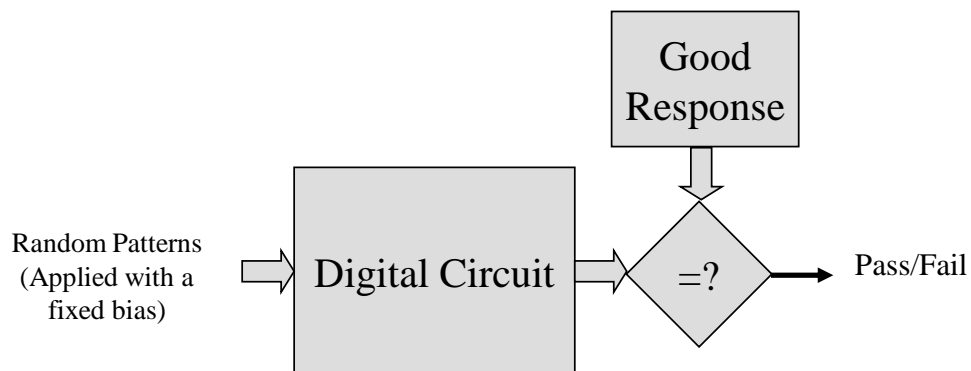
Outline

- Background
- The Problem
- Salient Results
- Conclusion

Signal Probabilities in a 2-Level Tree



Random-Pattern Testing



Prior Work

- Considered trees of NAND gates:
 - Fixed gate fan-in: n
 - Number of levels: l
 - Input One-probability: x (aka the *bias probability*)
- Showed, for large l , the output signal probabilities alternated between close to zero and one.
- For trees of 2-input NANDs ($n=2$), x_{opt} , the bias probability at which the detectability of primary input faults is maximum, is very sensitive to the value of x , with a sharp peak at $x_{opt} = 0.617$.


Prior Work (Contd.)

- Practical application was shown by doing Monte Carlo experiments on several Illiac IV processing boards, with logic approximated by NAND trees of fan-in between 2 and 3.
- Experiments verified:
 - random-pattern detectability is very sensitive to the value of x .
 - Optimum results occur close to $x = 0.617$.

The suggestion of using a bias probability derived from the average fan-in is purely intuitive; could lead to erroneous results if the signal probabilities (and detectabilities) in a tree were very sensitive to the changes in gate fan-ins.


Our paper provided a theoretical answer to the question of sensitivity

Transfer Function of 2-Level AND-OR Trees



- S-shaped; exact shape depends on the fan-ins m , and n .
- Properties:
 - $g(x)$ is monotonic
 - Three *fixed points*, where $g(x) = x$: 0 , 1 , and c .
 - First n (m) derivatives 0 at $x=0$ ($x=1$)
 - $g(x) > x$ above c and $g(x) < x$ below c

Multilevel Trees



A four-level AND-OR tree viewed as a two-layer tree

The transfer function of an l -layer tree denoted as $g_l(x)$

$$g_l(x) = g_{l-1}(g(x))$$

where, $g(x)$ is the transfer function of a 2-level tree

- Asymptotic Behavior: Not hard to show that:

$$\lim_{l \rightarrow \infty} g_l(x) = 0 \quad \text{for } 0 \leq x < c \quad \text{and}$$

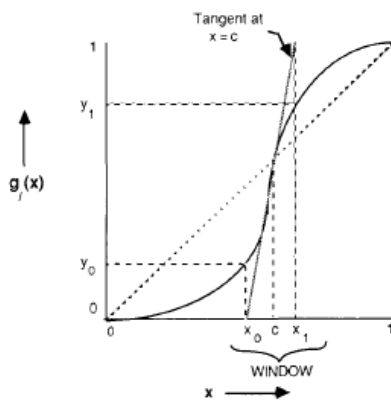
$$\lim_{l \rightarrow \infty} g_l(x) = 1 \quad \text{for } c < x \leq 1$$

- Therefore, $g_l(x)$ is a unit step function in the limit, with c as its *firing point*.

Firing Point vs. Gate Fan-ins

		n						
		2	3	4	5	6	7	8
m	2	0.618	0.848	0.921	0.951	0.967	0.977	0.982
	3	0.389	0.682	0.805	0.887	0.902	0.925	0.940
	4	0.282	0.580	0.724	0.803	0.850	0.881	0.905
	5	0.220	0.511	0.666	0.755	0.810	0.847	0.873
	6	0.181	0.461	0.622	0.717	0.778	0.820	0.849
	7	0.153	0.423	0.587	0.687	0.752	0.800	0.829
	8	0.133	0.392	0.558	0.661	0.730	0.777	0.812

Window of Uncertainty



Window of uncertainty is the inverse of the slope at c

Window Width vs. Tree Depth
(Binary Trees, $m=n=1$)

No. of Layers	Window Width	x_0	x_1	y_0	y_1
1	0.655	0.214	0.868	0.089	0.939
2	0.428	0.353	0.782	0.107	0.922
3	0.280	0.445	0.725	0.110	0.914
4	0.184	0.505	0.688	0.110	0.910
5	0.120	0.544	0.664	0.108	0.908
6	0.079	0.569	0.648	0.107	0.907
7	0.051	0.586	0.638	0.105	0.906
8	0.034	0.597	0.631	0.105	0.905
9	0.022	0.604	0.626	0.104	0.905
10	0.014	0.609	0.624	0.104	0.905

Window width decreases exponentially: $w_i = (w_1)^i$

Conclusion

- AND-OR trees generalize the fixed fan-in NAND trees discussed in the literature.
- Input vectors composed of bits independently chosen to be 1 with probability x .
- Basis for analysis: The probability transfer function of two-level AND-OR trees; finds the output probability as a function of the input bias x .
- The transfer function of every iterated tree approaches the step function, where the step is at $c=g(c)$.
- Conjecture: Behavior of any large AND-OR tree would resemble a step function occurring at a position that is highly dependent on the specifics of the structure closer to the leaves.
- A fun project, started by a casual conversation about collaboration.

Reference

Lester Lipsky and Sharad C. Seth,
“Signal Probabilities in AND-OR Trees”
IEEE Transactions on Computers, 38(11),
November 1989, pp. 1558-1563.

ISBN 3-937201-24-6
DOI 10.2313/NET-2011-10-1

ISSN 1868-2634 (print)
ISSN 1868-2642 (electronic)