

Probabilistic models for access strategies to dynamic information elements

Hansen, Martin Bøgsted; Olsen, Rasmus L.; Schwefel, Hans-Peter

Publication date:
2006

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Hansen, M. B., Olsen, R. L., & Schwefel, H.-P. (2006). *Probabilistic models for access strategies to dynamic information elements*. Department of Mathematical Sciences, Aalborg University. Research Report Series No. R-2006-33

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

**Probabilistic models for access strategies to
dynamic information elements**

by

Martin B. Hansen, Rasmus L. Olsen and Hans-Peter Schwefel

R-2006-33


November 2006

DEPARTMENT OF MATHEMATICAL SCIENCES
AALBORG UNIVERSITY

Fredrik Bajers Vej 7 G ■ DK-9220 Aalborg Øst ■ Denmark

Phone: +45 96 35 80 80 ■ Telefax: +45 98 15 81 29

URL: www.math.auc.dk/research/reports/reports.htm



Probabilistic models for access strategies to dynamic information elements

Martin B. Hansen^a Rasmus L. Olsen^b Hans-Peter Schwefel^b

^a*Department of Mathematical Sciences, Aalborg University, Fredrik Bajers Vej 7G, DK-9220 Aalborg, Denmark*

^b*Center for TeleInFraStruktur, Aalborg University, Niels Jernes Vej 12, DK-9220 Aalborg, Denmark*

Abstract

In various network services (e.g., routing and instances of context-sensitive networking) remote access to dynamically changing information elements is a required functionality. 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 (3) triggered by changes of the information element. This paper develops probabilistic models for these scenarios, which allow to compute a number of performance metrics, with a special focus on the mismatch probability. In particular, we use matrix-analytic methods to obtain explicit expressions for the mismatch probability that avoid numerical integration. Furthermore, limit results for information elements spread over a large number of network nodes are provided, which allow to draw conclusions on scalability properties. The impact of different distribution types for the network delays as well as for the time between changes of the information element on the mismatch probability are obtained and discussed through the application of the model in a set of example scenarios. The results of the model application allow for design decisions on which strategy to implement for specific input parameters and specific requirements on the performance metrics.

Key words: Distributed systems; Remote access; Performance modelling

¹ *Email addresses:* mbh@math.aau.dk (M.B. Hansen), rlo@kom.aau.dk (R.L. Olsen), hps@kom.aau.dk (H.-P. Schwefel).

1 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 [14].
- 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 [16].
- Context-sensitive services require access to typically remotely obtained context information. Context information may thereby be used both during service execution [20] as well as for the service discovery process [7].
- 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 [21,5,3] .

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. In order to quantitatively support such design decisions, this paper presents the analysis of a set of important base cases for such remote information access. The abstracted scenario is thereby shown in Figure 1.




Fig. 1. Abstracted scenario for remote information access.

The node, called *requester*, on the left-hand side of Figure 1 has to execute a computation for which it needs N input variables, which are dynamically changing and whose values are known at corresponding N remote nodes. The latter are called *information providers* in this paper. The information provider can send messages to the requester (downstream) through an interconnecting network, which imposes some delay and possible message loss and re-ordering of the downstream messages. If needed, also the requester can send messages to the information providers, here referred to as 'upstream' communication.

The assumptions in this paper regarding the specific type of information are rather general, in particular, no assumptions on the semantics are made, see [13,18] for specific use-cases:

- The information element at the information provider changes value at discrete points in time. Thereby two cases are distinguished later: (1) The information element never changes back to a previous value, as e.g. occurring for monotonous changes such as time and (2) the information element takes a finite set of values and can also possibly change back to a previously taken value.
- Neither the requester nor the information provider can influence the timing of the changes of the information element. This is e.g. the case for environment information provided by sensor devices, and it needs to be distinguished from cases of distributed implementations of shared variables, which can benefit from commitment or concurrency control protocols [9].
- It is irrelevant for the analysis in the paper whether the information element is a single-valued real or integer variable or a very complex data-structure.

Two basic types of solutions for such remote access are well known, see e.g., [10,14,19]:

- (1) **Reactive**, 'on-demand' access: Whenever the requestor needs a certain piece of remote information, it sends a request message to the information provider, which responds by sending the value of the information element. This in principle implements a client-server architecture.
- (2) **Proactive** distribution of information: The information provider will proactively distribute updates of the value of the information element 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 Section 2.3.
 - (b) **Periodic** proactive updates: After certain time-intervals, the current value of the information element is distributed to potential request processes.

In [19] a methodology for analysis was developed and quantitative results were obtained in a number of special cases: Only the case of $N = 1$ information providers was considered analytically and mainly for Poisson assumptions on event process and network delay processes, see Section 2.2 for a summary. Furthermore, the analysis in [19] focused on scenarios in which the information element never changes back to previous values. This paper, however, provides generalizations in a number of ways. More general distributions for event as well as delay processes are treated. Moreover it is shown how an analysis can be obtained for non-monotonous event processes as well as multiple information sources.

The paper is organized as follows: In Sections 2 and 3, an overview is given of the various access strategies and previous results from [19] as well as some mathematical preliminaries and the system model including the event, update and request processes. The analysis is done on general distribution types with particular focus on processes with matrix-exponentially distributed marginal distributions in Section 4. In Sections 5 and 6, the multiple source case and the case of information elements that can change back to a previous value are treated, respectively. The paper is concluded with a number of quantitative results in Section 7 and a summary and outlook in Section 8.

2 Description of different access strategies and previous results

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 the so-called mismatch probability.

2.1 System abstraction

The abstracted model contains three parts:

- The N remote information elements are maintained by remote nodes (information providers) and they dynamically change their values at certain discrete points in (continuous) time, described by the point processes $E^{(n)} = \{E_i^{(n)}, i \in \mathbf{Z}\}$, where $E_i^{(n)}$ is an increasing sequence of event times for information element n , $n = 1, \dots, N$, numbered such that $E_0^{(n)}$ is the event just before 0. The process $E^{(n)}$ is called the **event process** for information element n . $E^{(n)}(t)$ denotes the value of the information element at time t .
- All N remote information elements are required by a certain entity (the requester/client) at certain moments in time, identified by the **request**

process, $R = \{R_k, k \in \mathbf{Z}\}$, which in turn is a point process denoted in the same way as the $E_i^{(n)}$'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 proactive approaches.

- Communication between the requesting entity and the information provider n is described by stochastically varying upstream delays, $\{U_k^{(n)}, k \in \mathbf{Z}\}$ and downstream delays, $\{D_k^{(n)}, k \in \mathbf{Z}\}$.

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

2.2 Base cases, performance metrics and previous results

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 for the remote access strategy.
- (2) Access delay: The time interval from the moment when the N information elements are needed at the requester until they are finally available for use. For the proactive access strategies, this delay is zero. Processing times are neglected (or assumed to be included in the communication delays) in this paper.
- (3) Mismatch probability: The probability that any of the N values of the information elements that are 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. [13,18] for use-cases and their interpretation.

Figures 2, 3, and 4 illustrate the three base cases, reactive access, proactive event-driven access, and proactive periodic, respectively, for the case of $N = 1$ information providers.




Fig. 2. Reactive access: In the example, the k 'th access, R_k , leads to a 'correct' value, while the $k + 1$ 'th access causes a mismatch event.




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




Fig. 4. Proactive periodic update using a deterministic period: R_k results in mismatch, while R_{k+1} leads to a correct result.

The proactive event-based strategy is further distinguished in [19] in a case with incremental updates as opposed to full update messages, with sizes $s_d^{(i)}$ and $s_d^{(f)}$, respectively. Incremental updates, e.g. sending only the difference to the last value, may lead to smaller message sizes, $s_d^{(i)} < s_d^{(f)}$, at the cost of increased mismatch probability for some cases, see Section 4.2.

A summary of results under special assumptions on the three base cases from [19] is given in the following table

	Reactive	Proact. event Full update	Proact. event Incremental	Proact. Periodic
mmPr				Numeric
Exponential delay	$\frac{\lambda}{\lambda+\nu}$	$\frac{\lambda}{\lambda+\nu}$	$1 - e^{-\lambda/\nu}$	solution of Markov chain
Deterministic delay	$1 - e^{-\lambda D}$	$1 - e^{-\lambda D}$	$1 - e^{-\lambda D}$	Not treated in [19]
Overhead	$\mu(s_u + s_d)$	$\lambda s_d^{(f)}$	$\lambda s_d^{(i)}$	$\tau s_d^{(f)}$
Access delay	$\mathbb{E}(U) + \mathbb{E}(D) > 0$	0	0	0

Thereby the limitations are: only $N = 1$ information providers are considered; the event process is a Poisson process with rate λ , the downstream delays are independent and identically exponentially distributed with rate ν (first row) or deterministic with constant value D (second row). The periodic approach uses

a Poisson process for the moments of sending updates. The request process is only relevant for overhead in the reactive strategy; it can be a general process with average inter-request time $1/\mu$. The upstream delay is only relevant for access delay in the reactive strategy; the only assumption is that its first moment exists.

2.3 Subcases and scope of this paper

As already identified in [19], the case of full update messages, containing the complete information, need to be distinguished from a strategy that uses incremental updates, e.g. transmitting only changes in the information element. Although such an incremental strategy could also be implemented for a reactive approach, the actual implementation would be rather complicated, since all information providers then need to keep track of the last requests for every requester. Therefore the incremental updates appear more useful for the proactive approaches. Another distinction is necessary in the proactive approach regarding the ordering of update messages at the requester, if the network allows for reordering (as we assume here). The requester could order the received updates according to receive time, or according to sender sequence numbers. The latter has the advantage that the latest sent message can be used, but at the cost of increased message sizes through sequence numbers. Hence such an approach reduces mmPr at the cost of overhead.

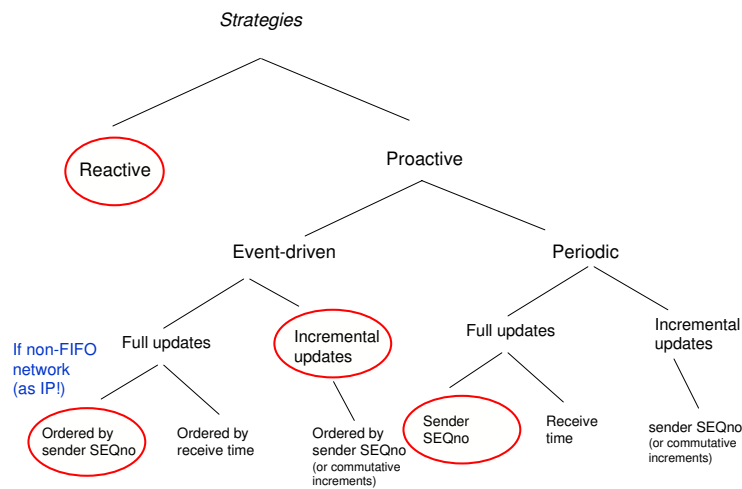


Fig. 5. Overview of the different cases. Those cases encircled are those we will focus on in this paper. The circled cases are the main focus of the paper.

Figure 5 summarizes the different cases in a tree-like structure. The incremental update cases are thereby only considered for messages ordered by sender sequence numbers, or for cases in which the actual order does not matter, e.g.

for additive increments. Reason for this limitation is that otherwise a single reordering event at any time in the past would lead to persistent mismatches in the future.

3 Notation and mathematical preliminaries

We first introduce a notation to refer to the different subcases of the remote access scenarios and subsequently introduce the relevant mathematical concepts.

3.1 Notation

For convenience we use the Kendall notation with / when we are classifying queueing systems and introduce a Kendall-alike notation with |, when we are classifying information retrieval systems.

More specifically, the remote access scenario is described by specifying the event process, the delay process, and the number of information providers as $E|D|N$. This part is the always present pre-fix, e.g., $M|M|1$ for the single information provider, Poisson event process, and exponentially distributed downstream delays. If nothing else is specified, then it is the event-driven, full update, ordered by sender sequence number case by default.

Otherwise, the prefix can be extended as $E|D|N|RS|US|order$ where 'RS' is the request strategy and can be either 'react(U)' with specification of the upstream delay process, or for the proactive strategies 'event', or 'periodic(P)', where P is specifying the period. 'US' is by default 'full' but can also be specified to be 'incr' (incremental). Ordering is by default by sender sequence numbers ('sSeq'), but can also be receive time ('rTime'). The event process by default never changes back to previous values ($E[mono]$) but can also be specified as $MAP[rec]$, which should stand for a recurrent Markovian process. Here recurrent is used in the sense that the process may return to at least some of the states.

Examples:

- $M|M|4|Periodic(M)$: Monotonous-type Poisson event process, exponentially distributed downstream delay, four information providers, proactive periodic strategy with Poisson period.
- $MAP[rec]|GI|1|React(M)$: Recurrent Markov process as event process, GI downstream delays, one source, reactive access strategy with exponential upstream delay.

- *ME|ME|4|event|incr|rTime* matrix-exponential renewal processes for event and downstream delay, event-driven strategy with incremental updates, ordered by receive time.

3.2 Mathematical preliminaries

Let $X = \{X_i, i \in \mathbf{Z}\}$ be the times of occurrences of some phenomenon, where X_i is an increasing sequence of event times numbered such that X_0 is the event just before 0. If we put $Y_i = X_i - X_{i-1}$, $i \in \mathbf{Z}$ and assume that the sequence of Y_i 's are independent and identically distributed (iid) then X is called a renewal process.

With an abuse of notation, a random variable with the distribution of the Y_i 's is denoted generically as X . We denote the cumulative distribution function (cdf) and the complementary cdf of X by F_X and $\bar{F}_X = 1 - F_X$, respectively. If the probability density function (pdf) of the distribution function exists, it is denoted by f_X .

If we further assume X is stationary (i.e. for every $r = 1, 2, \dots$ and all bounded (Borel) subsets A_1, \dots, A_r of the real line, the joint distribution of $\{|A_1 + t|, \dots, |A_r + t|\}$ does not depend on t ($-\infty < t < \infty$), [6, Definiton 3.2.I]). Then define the forward recurrence time $Y = X_1$ and the backward recurrence time $U = -X_0$ and their distribution functions as $B_X(t) = \mathbb{P}(Y \leq t)$ and $A_X(t) = \mathbb{P}(U \leq t)$. Whenever X is stationary the distribution of the backwards recurrence time is the same as the forward recurrence time, [1, p. 150]. Moreover,

$$b_X(t) = \frac{\bar{F}_X(t)}{\mathbb{E}(X)} \quad \text{and} \quad a_X(t) = \frac{\bar{F}_X(t)}{\mathbb{E}(X)}. \quad (1)$$

3.3 Matrix-exponential distributions

Consider now a vector-matrix pair $\langle \mathbf{p}, \mathbf{B} \rangle$ and a row-vector $\boldsymbol{\varepsilon}$ of ones. A distribution with cdf F which can be expressed as

$$F(t) = 1 - \mathbf{p} \exp(-t\mathbf{B})\boldsymbol{\varepsilon}'$$

is said to have a matrix-exponential representation with generator or representation $\langle \mathbf{p}, \mathbf{B} \rangle$, see [17,11].

Special cases of matrix-exponential distributions are Hyper-Exponential distributions and Erlangian distributions, see Appendix. A special version of the

former, namely truncated Power-Tail (TPT) distributions [8], are used in Section 7 to illustrate the mmPr behavior for scenarios with high variance in inter-event and downstream delay processes. The Erlangian distributions are used to illustrate the behavior when the variance of participating distributions is decreasing.

4 Analytic results for single information provider and monotonous-type event processes

Before addressing more general scenarios, we focus on the case of a single information provider ($N = 1$) for an information element that does not change back to a previous value.

4.1 Reactive on-demand access

In [19], the mismatch probability for the monotone event process and reactive strategy was derived. Assume E is a stationary event process and construct the right-continuous stochastic process $E(t) = i$, $t \in [E_i, E_{i+1})$. Formally, in this setting we attach to each arrival in the event process the state $E(E_i)$ and to each request the upstream and downstream delays (U_k, E_k) . This leads to a system consisting of the following two marked point processes $\{(E_i, E(E_i)), i \in \mathbf{Z}\}$ and $\{(R_k, U_k, E_k), k \in \mathbf{Z}\}$. Then by stationarity we have the following probability of mismatch upon reception of the message for any request at time R_k

$$\begin{aligned} \text{mmPr} &= \mathbb{P}(E(R_k + U_k + D_k) \neq E(R_k + U_k)) \\ &= \mathbb{P}(E(D_k) \neq E(0)) \\ &= 1 - \int \mathbb{P}(E(D_k) = E(0) | D_k = t) F_D(dt) \\ &= 1 - \int \bar{B}_E(t) F_D(dt). \end{aligned}$$

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

$$\text{mmPr} = 1 - \int \bar{B}_E(t) F_D(dt). \quad (2)$$

Although, the mmPr is independent of the request process, the process R will influence statistical properties of corresponding estimators of mmPr. Note also, the mmPr is not depending on the upstream delay process.

Equation (2) where $B_E(t)$ is the cdf of the backwards recurrence time of the event process, can be simplified for a large class of processes, with relevant examples listed below.

With the above Equation (2), any stationary event process is covered for a general downstream delay, i.e. the $G|G|1|react(G)$ case.

Notice that the case of correlated delays is also taken care of, since those do not influence the stationary mmPr. One should also notice that the case of dropped packets in principle can be treated with the set-up presented above. One possibility is to assume that a dropped packet has infinite delay time and then put an atom of the delay distribution at infinity.

4.1.1 Event process is Poisson

Under the assumption that the event process, E , is a Poisson process with rate λ we obtain:

$$\begin{aligned} \text{mmPr} &= 1 - \int e^{-\lambda t} F_D(dt) \\ &= 1 - L\{F_D\}(\lambda) \end{aligned} \quad (3)$$

where the last term is the Laplace-Stieltjes transform of the cdf of the downstream delay, evaluated at λ . Hence, Equation (3) covers the $M|G|1|react[G]$ case.

4.1.2 Event and delay processes are matrix-exponential

Assume E is a matrix-exponential (ME) event process with representation $\langle \mathbf{p}_E, \mathbf{B}_E \rangle$ and D is a ME process with representation $\langle \mathbf{p}_D, \mathbf{B}_D \rangle$. Then

$$\begin{aligned} \text{mmPr} &= 1 - \frac{1}{\mathbb{E}(E)} \int_0^\infty \int_t^\infty \bar{F}_E(s) ds F_D(dt) \\ &= 1 - \int_0^\infty \frac{\mathbf{p}_E \mathbf{V}_E}{\mathbf{p}_E \mathbf{V}_E \boldsymbol{\epsilon}'_E} \exp(-\mathbf{B}_E t) \boldsymbol{\epsilon}'_E \mathbf{p}_D \mathbf{B}_D \exp(-\mathbf{B}_D t) \boldsymbol{\epsilon}'_D dt. \end{aligned}$$

The integral can be solved using a Kronecker product representation (at the cost of expanding the matrix dimensions):

$$\text{mmPr}$$

$$\begin{aligned}
&= 1 - \frac{1}{\mathbf{p}_E \mathbf{V}_E \boldsymbol{\epsilon}'_E} \int_0^\infty [(\mathbf{p}_E \mathbf{V}_E) \otimes (\mathbf{p}_D \mathbf{B}_D)] \exp [-(\mathbf{B}_E \oplus \mathbf{B}_D)t] [\boldsymbol{\epsilon}'_E \otimes \boldsymbol{\epsilon}'_D] dt \\
&= 1 - \frac{(\mathbf{p}_E \mathbf{V}_E) \otimes (\mathbf{p}_D \mathbf{B}_D)}{\mathbf{p}_E \mathbf{V}_E \boldsymbol{\epsilon}'_E} [\mathbf{B}_E \oplus \mathbf{B}_D]^{-1} \boldsymbol{\epsilon}'_{\dim(B_E) \cdot \dim(B_D)}. \tag{4}
\end{aligned}$$

Hence, the equation above solves the $ME|ME|1|react[G]$ case. The analysis in Section 4.2 shows that the proactive event-driven approach with full updates (ordered by sender sequence numbers) actually results in exactly the same stationary mismatch probability. A slightly different, but equivalent matrix-exponential representation results from the derivation, see Equation (8).

Simplified formulas for Poisson case: In case one of the two processes, event or downstream delay, is a Poisson process, the ME representation of that renewal process reduces to scalars and Equation (4) simplifies accordingly:

Poisson event process: $M|ME|1|react$

$$\text{mmPr} = 1 - \mathbf{p}_D \mathbf{B}_D [\mathbf{B}_D + \lambda \mathbf{I}]^{-1} \boldsymbol{\epsilon}'_D.$$

Exponential downstream delay: $ME|M|1|react$

$$\text{mmPr} = 1 - \frac{\nu}{\mathbb{E}(E)} \mathbf{p}_E \mathbf{V}_E [\mathbf{B}_E + \nu \mathbf{I}]^{-1} \boldsymbol{\epsilon}'_E.$$

Delays with constant offset: The case of a shifted delay distribution is particularly interesting, since frequently there is a minimum network delay, d_0 , to which the variable delay, D_v , caused by congestion is added. This means the total delay D is given by

$$D = d_0 + D_v$$

whereby

$$f_D(t) = \begin{cases} 0 & \text{for } 0 \leq t < d_0 \\ f_{D_v}(t - d_0) & \text{for } t \geq d_0. \end{cases}$$

Hence, Equation (2) and its ME equivalent are modified correspondingly:

$$\text{mmPr} = 1 - \frac{1}{\mathbb{E}(E)} \int_0^\infty \left[\int_{t+d_0}^\infty \bar{F}_E(s) ds \right] f_{D_v}(t) dt,$$

which in the matrix-exponential case becomes

$$\text{mmPr} = 1 - \frac{(\mathbf{p}_E \mathbf{V}_E \exp[\mathbf{B}_E d_0]) \otimes (\mathbf{p}_D \mathbf{B}_D)}{\mathbf{p}_E \mathbf{V}_E \boldsymbol{\epsilon}'_E} [\mathbf{B}_E \oplus \mathbf{B}_D]^{-1} \boldsymbol{\epsilon}'_{\dim(B_E) \cdot \dim(B_D)}.$$

I.e. the only modification as opposed to Eq. (4) is the term $\exp(\mathbf{B}_E d_0)$ in the left term of the vector Kronecker product.

4.1.3 Events changes according to a Markov jump process

Assume E are the points of state-transitions of a Markov process with matrix \mathbf{Q} , which has steady state probability π_i and state-leaving rates $\lambda_i = -\mathbf{Q}_{i,i}$ for $i = 1, \dots, S$:

$$\text{mmPr} = 1 - \sum_{i=1}^S \pi_i \int_0^{\infty} e^{-\lambda_i t} F_D(dt). \quad (5)$$

This can be written integral-free when D has a matrix-exponential representation, i.e. for the $MJ|ME|1|react(G)$ case:

$$\text{mmPr} = 1 - \sum_{i=1}^S \pi_i \mathbf{p}_D \mathbf{B}_D [\mathbf{B}_D + \lambda_i \mathbf{I}]^{-1} \boldsymbol{\epsilon}'_D. \quad (6)$$

4.2 Proactive event based update

For the proactive event-based updates, we need to further distinguish between incremental and full updates. In both cases we assume that either the ordering of the messages at the requester does not matter (for the incremental case) or is performed via sender sequence numbers.

4.2.1 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.

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 and inter-change of integration

$$\begin{aligned}
\text{mmPr} &= 1 - \int_0^\infty \mathbb{P}(D \leq t | B = t) B_E(dt) \\
&= 1 - \int_0^\infty \int_0^t F_D(ds) B_E(dt) \\
&= 1 - \int_0^\infty \overline{B}_E(s) F_D(ds)
\end{aligned} \tag{7}$$

which is the same as Equation (2). This fact can also be explained intuitively as in the reactive case the mmPr is the probability that the forward recurrence time is less than the delay time. Moreover, in the proactive full updates case the mmPr is the probability that the backwards recurrence time is less than the delay time. As the forward and backward recurrence times are the same, the conclusion follows.

For the matrix-exponential case the mmPr then equals (4). It is also possible to use the following alternative representation and arrive at

$$\begin{aligned}
\text{mmPr} &= 1 - \int_0^\infty \mathbb{P}(D \leq t) B_E(dt) \\
&= 1 / \mathbb{E}(E) \int_0^\infty \overline{F}_D(t) \overline{F}_E(t) dt \\
&= \frac{1}{\mathbf{p}_E \mathbf{V}_E \boldsymbol{\varepsilon}'_E} \int_0^\infty \mathbf{p}_D \exp(-\mathbf{B}_D t) \boldsymbol{\varepsilon}'_D \mathbf{p}_E \exp(-\mathbf{B}_E t) \boldsymbol{\varepsilon}'_E dt,
\end{aligned}$$

which again can be simplified using product spaces:

$$\text{mmPr} = \frac{\mathbf{p}_D \otimes \mathbf{p}_E}{\mathbf{p}_E \mathbf{V}_E \boldsymbol{\varepsilon}'_E} [\mathbf{B}_D \oplus \mathbf{B}_E]^{-1} \boldsymbol{\varepsilon}'_{\dim(B_D) \cdot \dim(B_E)}. \tag{8}$$

This is the same result as Equation (4), so it is left to the reader to decide which one is more esthetic. See Section 4.1 for special cases.

As a general, important conclusion: The mmPr in the $G|G|1|react(G)$ case is exactly the same as in the corresponding $G|G|1|event|full|sSeq$ case and given by Equation (2).

4.2.2 Incremental updates

In this scenario, the requester only accesses the correct information, if all update messages from previous events have been successfully received. In this

case, a mismatch would occur, if **any** of the update messages is still in transit. In this system we consider the marked point process $\{(E_i, E(E_i), D_i), i \in \mathbf{Z}\}$ and the series of request times R .

Let the stochastic process Q_t denote the number of updates in transit at time t . Then the $G/G/\infty$ -queueing model can be used to describe the model, as the E_i 's are the arrivals times of customers, each arrival requires D_i amount of service, and Q_t is the number of busy servers. Henceforth, the mismatch probability at request R_k is the probability that Q_{R_k} is strictly greater than zero. By stationarity we have

$$\text{mmPr} = \mathbb{P}(Q_{R_k} > 0) = \mathbb{P}(Q_0 > 0).$$

This is equivalent to the probability that an $G/G/\infty$ 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

$$\text{mmPr} = \mathbb{P}(E/D/\infty \text{ queue is busy}). \quad (9)$$

By Little's formula the mismatch probability can, under suitable ergodicity assumptions, be calculated as

$$\text{mmPr} = \frac{\bar{b}}{\bar{b} + \bar{c}}, \quad (10)$$

where \bar{b} and \bar{c} are the mean busy period and mean busy cycle, respectively.

Busy period problems for infinite server queues have been studied quite extensively by various authors, as many problems can be recast as infinite server queues, see e.g. [15] for a recent application in network related areas. In [12] general relations are given for the length of the busy cycle, busy period and idle period. Renewal arguments are used to derive explicit formula for the Laplace transforms of a) $P_0^*(t)$, the probability that at time t the system is empty, b) the busy cycle, and c) busy period distribution functions. Although, \bar{b} and \bar{c} can be derived by the moment-generating properties of the Laplace transform, explicit formulas have only been obtained in the $GI/D/\infty$ and $M/GI/\infty$ cases, hence solving the $GI|D|1|event|incr$ and $M|GI|1|event|incr$ cases, see Equations (11) and (12) below, respectively.

Utilizing results for the queue-length probabilities of the $M/GI/\infty$ queue, the $M|GI|1|event|incr$ access strategy, with Poisson assumptions on E (with rate λ) and general independent (GI) assumptions for the downstream delay D (with mean \bar{D}), results in a mismatch probability of

$$\text{mmPr} = 1 - \exp(-\lambda\bar{D}). \quad (11)$$

4.3 Scenarios without message loss and reordering

In case of deterministic downstream delay, the mmPr actually is identical for the event-driven incremental strategy and the event-driven full update strategy, since no reordering and no message loss can occur. Also, the ordering of messages at the requester, rTim or sSeq, is equivalent in this case. Continuing the investigations on the queueing representation for the event-driven incremental case from the previous section, under the general independent (GI) assumption for E and a deterministic delay time of size D the mmPr is given by [12, Corollary 2]

$$\text{mmPr} = \frac{(1 - F_E(D))^{-1} \int_0^D x F_E(dx) + D}{\lambda^{-1}(1 - F_E(D))^{-1}} \quad (12)$$

where $\lambda = (\int_0^\infty x F_E(dx))^{-1}$ is the arrival rate of E . An equivalent representation of this result can be obtained also from Eq. 2.

Matrix-exponential event process: Assuming a matrix-exponential renewal event process with representation $\langle \mathbf{p}_E, \mathbf{B}_E \rangle$, and a deterministic downstream delay of constant size D , Equation (12) can be written integral-free as:

$$\begin{aligned} \text{mmPr} &= \lambda_E \int_0^D \mathbf{p}_E \mathbf{B}_E x \exp(-\mathbf{B}_E x) \boldsymbol{\varepsilon}' dx + D \lambda_E \mathbf{p}_E \exp(-\mathbf{B}_E D) \boldsymbol{\varepsilon}' \\ &= 1 - \lambda_E \mathbf{p}_E \left(D \mathbf{I} + \mathbf{B}_E^{-1} \right) \exp(-\mathbf{B}_E D) \boldsymbol{\varepsilon}' + D \lambda_E \mathbf{p}_E \exp(-\mathbf{B}_E D) \boldsymbol{\varepsilon}' \\ &= 1 - \lambda_E \mathbf{p}_E \mathbf{B}_E^{-1} \exp(-\mathbf{B}_E D) \boldsymbol{\varepsilon}'. \end{aligned} \quad (13)$$

This equation provides the integral-free expression to the $ME|D|1|event|incr$ scenario, which is identical to the $ME|D|1|event|full$ which in turn has the same mmPr as the $ME|D|1|reactive(G)$ strategy. Formulated more general: In scenarios in which no message loss and no packet ordering can result, all three strategies lead to the same mmPr.

4.4 Proactive, periodic update

In this scenario the information provider informs the request process periodically about the value of the information element. These periodic updates are subject to the downstream network delay. Formally, we consider

a system consisting of the marked point processes $\{(I_k, D_k), k \in \mathbf{Z}\}$ and $\{(E_i, E(E_i)), i \in \mathbf{Z}\}$. Define for any t the following stopping time (with respect to the natural filtration, [1, Chapter I.8]),

$$\tau_t = \max\{n \mid I_n + D_n \leq t\}.$$

Then by stationarity we have the following probability of mismatch upon request at time R_k :

$$\begin{aligned} \text{mmPr} &= \mathbb{P}(E(U_{\tau_{R_k}}) \neq E(R_k)) \\ &= \mathbb{P}(E(U_{\tau_0}) \neq E(0)) \\ &= \int_0^\infty \mathbb{P}(\text{no useful update received in } [0, t]) A_E(dt) \end{aligned}$$

Consider the $G|GI|1|periodic(M)|full|sSeq$ system. For this case a general formula is derivable: Consider an event process which is a stationary renewal process, whose backwards recurrence time has cdf A_E , iid downstream delays with cdf F_D , and the updates are assumed to be a stationary Poisson point process with intensity τ . Without loss of generality (by stationarity) we assume the request time is 0 and consider the point process of useful updates received before 0. Henceforth, they can be viewed as a thinned non-stationary Poisson point process where the thinning probability at time t is given by $F_D(t)$. The intensity function of the thinned Poisson point process is given by

$$\tau(t) = \tau F_D(t). \quad (14)$$

Hence, the general formula for the mismatch probability under Poisson assumption for the process of sending updates becomes

$$\text{mmPr} = \int_0^\infty \exp\left(-\int_0^t \tau F_D(s) ds\right) A_E(dt). \quad (15)$$

Hence, this is the solution for the mmPr of the $G|GI|1|periodic(M)$ system.

For matrix-exponential delay distributions, the inner integral in Equation (15) can be simplified to

$$\text{mmPr} = \int_0^\infty \exp\left[-\tau(t - \bar{D} + \mathbf{p}_D \mathbf{B}_D^{-1} \exp(-\mathbf{B}_D t) \boldsymbol{\epsilon}')\right] A_E(dt),$$

which under matrix-exponential assumption on the event process can be rewritten as

$$\text{mmPr} = \frac{e^{\tau \bar{D}}}{E} \int_0^\infty e^{-\tau t} \cdot \exp \left[-\tau \mathbf{p}_D \mathbf{B}_D^{-1} \exp(-\mathbf{B}_D t) \boldsymbol{\varepsilon}'_D \right] \cdot [\mathbf{p}_E \exp(-\mathbf{B}_E t) \boldsymbol{\varepsilon}'_E] dt.$$

This expression can be simplified further using Kronecker products, however without getting rid of the integral completely. Therefore, we use numeric integration later in Section 7. Note that in some cases, with a special structure on B_D , it may be possible to express the matrix-exponential as a closed-form expression involving exponential functions and thereby integral-free expressions can be derived.

4.4.1 Poisson event process and exponential downstream delays

Under further assumptions, namely for the $M|M|1|periodic(M)$ case, Equation (15) specializes to

$$\text{mmPr} = \lambda \int_0^\infty \exp \left(\frac{\tau}{\nu} (1 - e^{-\nu t}) \right) \exp(-(\tau + \lambda)t) dt, \quad (16)$$

whenever the event process is a Poisson process with intensity λ and the downstream delays are iid exponentially distributed with mean $1/\nu$.

The expression in (16) can alternatively be expressed as (substituting t with $e^{-\nu t}$)

$$\text{mmPr} = \frac{\lambda}{\nu} \exp \left(\frac{\tau}{\nu} \right) \frac{\Gamma \left(\frac{\tau + \lambda}{\nu} \right)}{\left(\frac{\tau}{\nu} \right)^{(\tau + \lambda)/\nu}} F_{\Gamma \left(\frac{\tau + \lambda}{\nu}, \frac{\tau}{\nu} \right)}(1), \quad (17)$$

where $F_{\Gamma(a,b)}$ is the cdf of a gamma distribution with parameters a and b .

4.4.2 Markov Model Approach

Another way of modelling the $M|M|1|periodic(M)$ case derived in (16) and (17) is to consider a stochastic process $\{X_t, t \in \mathbf{R}\}$ describing the result of a request at time t . The process is a Markov jump process with states (1) and $(0, i)$, $i = 0, 1, 2, \dots$ where state (1) denotes no mismatch and the state $(0, i)$ denotes a mismatch and number of updates, started after the last event which is in transit. The transition diagram is depicted in Figure 6. This




Fig. 6. Transition diagram for the $M|M|1|periodic(M)$ case.

model leads to the following generator matrix

$$Q = \begin{bmatrix} -\lambda & \lambda & 0 & 0 & 0 & \dots \\ 0 & -\tau & \tau & 0 & 0 & \dots \\ \nu & \lambda & -(\nu + \tau + \lambda) & \tau & 0 & \dots \\ 2\nu & \lambda & 0 & -(2\nu + \tau + \lambda) & \tau & \dots \\ 3\nu & \lambda & 0 & 0 & -(3\nu + \tau + \lambda) & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}.$$

From which the following steady-state probability vector π can be derived (solving $\pi Q = \pi$ and $\sum_{i=1}^{\infty} \pi_i = 1$).

$$\begin{aligned} \pi_1 &= \frac{\frac{\nu}{\lambda} \sum_{n=3}^{\infty} (n-2) \prod_{i=1}^{n-2} \frac{\tau}{i\nu + \tau + \lambda}}{1 + \frac{\nu}{\lambda} \sum_{n=3}^{\infty} (n-1) \prod_{i=1}^{n-2} \frac{\tau}{i\nu + \tau + \lambda}} = 1 - \text{mmPr} \\ \pi_2 &= \frac{1}{1 + \frac{\nu}{\lambda} \sum_{n=3}^{\infty} (n-1) \prod_{i=1}^{n-2} \frac{\tau}{i\nu + \tau + \lambda}} \\ \pi_n &= \frac{\prod_{i=1}^{n-2} \frac{\tau}{i\nu + \tau + \lambda}}{1 + \frac{\nu}{\lambda} \sum_{n=3}^{\infty} (n-1) \prod_{i=1}^{n-2} \frac{\tau}{i\nu + \tau + \lambda}}, \quad n \geq 3. \end{aligned} \tag{18}$$

4.5 Summary

The results in this section allow to compute analytically the mmPr for the scenario of $N = 1$ information providers, at which the information element never changes back to a previous value. The following cases can be analytically calculated:

- Reactive approach: Equation (2) is the general solution when the event

process is a general process. The upstream delay and the request process are both irrelevant. The downstream delay process is a general process. Equation (4) or equivalently Equation (8) is the integral-free representation for matrix-exponential processes.

- Proactive event-driven strategy with full updates (and sender sequence numbers for ordering): the mmPr is identical to the one in the reactive case, so all the statements from the item above hold.
- Proactive event-driven with incremental updates (and order irrelevant or can be re-created by sender sequence numbers): Equation (11) provides the mmPr for a Poisson event process and GI downstream delays.
- In scenarios without message loss and message reordering, all three strategies, reactive, proactive event-driven incremental, and proactive event-driven with full updates show exactly the same mmPr; the special case of deterministic downstream delay is treated as relevant example in (12). Also, the ordering of messages at the requester, rTim vs. cSeq, is irrelevant in this case.
- The proactive periodic strategy with full updates is covered by Equation (15) for the case of a GI event process and GI downstream delay, but Poisson assumptions on the period of sending updates. Integral-free representations exist in Equations (18) and (17) when all three participating processes are assumed to be Poisson.

The other performance metrics, access delay and networking overhead, are straightforward and were already presented in [19], see Section 2.2 for a summary. They hold for the general $G|G|1|x|y|z$ cases, hence are insensitive to the actual distribution types.

5 The case of multiple information providers

Now we consider the case of $N > 1$, i.e. the remote information elements are provided by N independent event sources $E^{(1)}, \dots, E^{(N)}$ and a correct computation at the requester is obtained only if the correct values are obtained from all N information providers. For all three basic access strategies, explicit results for the mmPr are obtained in Section 5.1. Furthermore, scenarios with very large N can be approximated using limit results that are derived in Section 5.2. Convergence to these limits is subject to appropriate scaling of system parameters. This scaling behavior allows to conclude on scalability of the strategies with respect to an increasing number of information providers.

5.1 Finite N

This time, we start with the proactive cases, since they are rather straightforward under appropriate independence assumptions.

5.1.1 Proactive cases

In all the proactive cases a mismatch is obtained if at least one of the information providers yields a mismatch. As the event sources are independent the probability of mismatch can be obtained by the following function of the individual mismatch probabilities,

$$\overline{\text{mmPr}}_N = 1 - \prod_{i=1}^N (1 - \text{mmPr}_i), \quad (19)$$

where each mmPr_i is obtained by the corresponding equations in Sections 4.2 and 4.4. We chose here and subsequently the notation $\overline{\text{mmPr}}_N$ for the total mmPr in this case of multiple information providers, to distinguish it for mmPr_N which reflects the mmPr for the part of the information element maintained by information provider N .

5.1.2 Reactive case

In the reactive scenario, the requestor sends out a request message to each information provider. The requests are sent out at the same time-instant (multicast), but may arrive at the information providers after different upstream delays, see Figure 7.




Fig. 7. Reactive access to multiple information providers: In the example, the k 'th access, R_k , leads to a 'correct' value, while the $k + 1$ 'th access causes a mismatch event.

Assume $t = 0$ is the time of sending out the request. Then the information will

be finally processed at the requester at time $\widetilde{M}_N = \max(U_1 + D_1, \dots, U_N + D_N)$, see Figure 7. The $\overline{\text{mmPr}}_N$ is:

$$\begin{aligned} \overline{\text{mmPr}}_N &= 1 - \mathbb{P}(E_1^{(1)}(\widetilde{M}_N) = E_1^{(1)}(U_1), \dots, E_1^{(N)}(U_N) = E_1^{(N)}(\widetilde{M}_N)) \\ &= 1 - \int_0^\infty \mathbb{P}(E_1^{(1)}(t - U_1) = E_1^{(1)}(0), \dots, E_1^{(N)}(t - U_N) = E_1^{(N)}(0) | \widetilde{M}_N = t) \\ &\quad F_{\widetilde{M}_N}(dt). \end{aligned}$$

As \widetilde{M}_N is dependent on the U_i 's it seems difficult to simplify this expression. However, if we assume that the same upstream delay is imposed on all request messages for all information providers, i.e. the request reaches all information providers at the same time instant, we can obtain rather explicit results and also limit theorems based on weak convergence and extreme value theory. See Section 5.2 for the latter. The assumption of identical upstream delay holds (approximately) e.g. for deterministic upstream delay, or if we assume a scenario of broadcast requests and assume that the broadcasting mechanism reaches all nodes at (approximately) the same time.

Reactive case with equal (deterministic) upstream delay

The requester will send a request that reaches all information providers at the same time instant. This scenario e.g. occurs, if we assume that the upstream delays are deterministic with value u . However, the information will only be processed when all answers reach the requester, e.g. at time \widetilde{M}_N , after the reading of the information at the providers. For the $G|G|N|react(D)$ -case, if M_N denotes $\max\{D_1, \dots, D_N\}$ and m_N denotes $\min\{E_1^{(1)}, \dots, E_1^{(N)}\}$ (maximum of the forward recurrence times for the N information elements), we get the following general formula

$$\begin{aligned} \overline{\text{mmPr}}_N &= 1 - \int_0^\infty \mathbb{P}(E_1^{(1)}(t) = E_1^{(1)}(0), \dots, E_1^{(N)}(t) = E_1^{(N)}(0) | M_N = t) F_{M_N}(dt) \\ &= 1 - \int_0^\infty \mathbb{P}(E_1^{(n)} \geq t_1, \dots, E_1^{(N)} \geq t) F_{M_N}(dt) \\ &= 1 - \int_0^\infty \overline{F}_{m_N}(t) F_{M_N}(dt). \end{aligned} \tag{20}$$

Notice that (20) resembles (2) as we can interpret the problem as a reactive

on-demand problem, where the event process has backward recurrence-time distribution \overline{F}_{m_N} and a delay process with cdf F_{M_N} .

If we utilize our standing assumption of independent delay and inter-event times between information providers, we get

$$F_{M_N}(t) = F_D(t)^N \quad \text{and} \quad \overline{F}_{m_N}(t) = B_E(t)^N,$$

which of course implies $f_{M_N}(t) = N F_D(t)^{N-1} f_D(t)$, whenever F_D has a pdf f_D . This in turn yields

$$\overline{\text{mmPr}}_N = 1 - N \int_0^\infty B_E(t)^N F_D(t)^{N-1} f_D(t) dt. \quad (21)$$

5.2 Limit results

In this subsection we analyze the behavior of $\overline{\text{mmPr}}_N$ for $N \rightarrow \infty$ under appropriate scaling on the event rates.

5.2.1 Proactive, periodic updates

If we consider Poisson event and information processes with rate λ and ν and iid exponentially distributed delays with mean $1/\tau$ and scale down the individual event processes with rate $\lambda_N = \lambda/N$ we obtain from Equations (17) and (19)

$$\lim_{N \rightarrow \infty} \overline{\text{mmPr}}_N = 1 - \exp \left(-\frac{\lambda}{\nu} \exp \left(\frac{\tau}{\nu} \right) \frac{\Gamma \left(\frac{\tau}{\nu} \right)}{\left(\frac{\tau}{\nu} \right)^{(\tau)/\nu}} F_{\Gamma \left(\frac{\tau}{\nu}, \frac{\tau}{\nu} \right)}(1) \right). \quad (22)$$

Hence, in this $M|M|N|periodic(M)$ case, we obtain the above limit for increasing N when scaling down the event rate linearly with N .

5.2.2 Proactive, event based updates

Now, for the proactive, event based scheme with incremental updates we get for a Poisson event process, $M|GI|N|event|incr$, and the scaling $\lambda_N = \lambda/N$

$$\lim_{N \rightarrow \infty} \overline{\text{mmPr}}_N = 1 - e^{-\lambda \overline{D}}. \quad (23)$$

Actually, this is not only a limit result, but it holds exactly for all N .

For the case of full updates and exponentially distributed downstream delays with mean $1/\nu$, $M|M|N|event|full$, we obtain for the same linear scaling $\lambda_N = \lambda/N$:

$$\lim_{N \rightarrow \infty} \overline{\text{mmPr}}_N = 1 - e^{-\lambda/\nu}.$$

Note that the scaling of the event-rate in the limit creates a scenario in which the update messages are not subject to reordering any more, since they are infinitely far apart. As a result, the proactive event-driven full and incremental strategies are in the limit equivalent except for cases with message loss.

5.2.3 Reactive approach

As opposed to the proactive approaches, the scaling properties of the reactive schemes are different. Assume for instance the delays to be iid with exponentially decaying tails, i.e. $\bar{F}_D(x) \sim e^{-\nu x}$, as $x \rightarrow \infty$. This case can conveniently be noted as $M|GI|N|React(D)$, where the GI has exponentially decaying tails. Then it can be proved that

$$\frac{M_N - b_N}{a_N} \quad \text{with} \quad a_N = 1/\nu, \quad b_N = \log(N)/\nu,$$

converges in distribution to the Gumbel distribution, [4], with cdf

$$F(x) = e^{-e^{-x}}.$$

Hence, if we scale down the individual event processes by

$$\lambda_N := \frac{\lambda}{N \log(N)}, \tag{24}$$

we obtain the following limit behavior of the $\overline{\text{mmPr}}_N$ (by the definition of weak convergence, see e.g. [2, (10)]),

$$\begin{aligned} \lim_{N \rightarrow \infty} \overline{\text{mmPr}}_N &= 1 - \lim_{N \rightarrow \infty} \mathbb{E}(\exp \left(\frac{-\lambda}{\log(N)} \left(\left(\frac{M_N - b_N}{a_N} \right) a_N + b_N \right) \right)) \\ &= 1 - e^{-\lambda/\nu}. \end{aligned} \tag{25}$$

Note, that this limit closely resembles the proactive cases, but under different scaling.

Now, alternatively assume the delays to be iid with polynomially tails, i.e. for some $\alpha > 0$, $\overline{F}_D(x) \sim x^{-\alpha}$ as $x \rightarrow \infty$. This case can conveniently be noted as $M|GI|N|React(D)$, where the GI has polynomially decaying tails. Then it can be proved that

$$\frac{M_N - b_N}{a_N} \quad \text{with} \quad a_N = N^{1/\alpha}, \quad b_N = 0,$$

converges in distribution to the Frechet distribution,[4], with cdf

$$F(x) = e^{-x^{-\alpha}}.$$

Hence, if we scale down the individual event processes by

$$\lambda_N := \frac{\lambda}{N^{1+1/\alpha}}, \tag{26}$$

and assume X has the Frechet distribution with cdf F , then we obtain the following limit behavior of the $\overline{\text{mmPr}}_N$ (by the definition of weak convergence, see e.g. [2, (10)])

$$\begin{aligned} \lim_{N \rightarrow \infty} \overline{\text{mmPr}}_N &= 1 - \mathbb{E}e^{-\lambda X} \\ &= 1 - \alpha \int_0^\infty e^{-(\lambda x + x^{-\alpha})} x^{-\alpha-1} dx. \end{aligned} \tag{27}$$

5.3 Summary

In this section, we consider the scenario in which the information element consists of a vector spread over $N > 1$ information providers. For an execution of a computation at the requester, all N parts of this vector are required, and a mismatch results if any of them does not correspond to the true current value. For the proactive cases, under assumption of mutual independence of the N downstream delay processes and the N event processes, the mmPr can simply be computed from a product expression, Equation (19). For the reactive case, under similar independence assumption and given that the upstream delays are identical for the request to reach all N nodes, Equation (21) allows to compute the mmPr, in most cases involving numerical integration.

Furthermore, limit results have been obtained for the three different cases, which allow to approximate via simple analytic expressions, under suitable scaling of the event rate, the $\overline{\text{mmPr}}_N$ for an increasing number of information

providers. The results in the reactive case are summarized in the following table:

	Exp. dec. delay	Pol. dec. delay
Scaling	$\lambda/(N \log(N))$	$\lambda/N^{1+1/\alpha}$
$\lim_{N \rightarrow \infty} \overline{\text{mmPr}}^N$	$1 - e^{-\lambda/\nu}$	$1 - \alpha \int_0^\infty e^{-(\lambda x + x^{-\alpha})} x^{-\alpha-1} dx$

For the proactive cases, the scaling of λ/N is the same for all sub-cases, and the results can be summarized as

	Proact. event Full	Proact. event Incremental	Proact. perioic (exp. period)
$\lim_{N \rightarrow \infty} \overline{\text{mmPr}}^N$			
Exp. delay	$1 - e^{-\lambda/\nu}$	$1 - e^{-\lambda/\nu}$	(22)
	Holds for <i>GI</i> delay		

As a summary remark on scalability for remote access to an increasing number of information providers: In the exponential $M|M|N|x$ case, the reactive approach requires to scale down the Poisson information change rate more strongly than the proactive scenarios, namely as $N \log N$. This stronger down-scaling of the event rate becomes even more pronounced for delay distributions with Pareto tails. This is a consequence of the extreme value statistics for the downstream delay that need to be applied in the reactive case.

6 Markov Event Processes

Now we remove the assumption that the information element cannot change back to a previous value. Instead, the information element is assumed to be described by the state of a Markov process with generator matrix \mathbf{Q} . State changes at instances of transitions of a Markov chain were already discussed earlier in Section 4.1.3, but then assuming that the information element is changing to a previously unobserved value.

For the analysis in this section, we focus on the reactive approach, with explicit integral-free results for the $MAP[rec]|ME|1|react(G)$ system. For the proactive approaches, we restrict ourselves to an outlook on the possible approaches.

As in Section 4, we again focus on the case of a single information provider, $N = 1$, but the general result of Equation (19) also applies in this case to the

proactive strategies. For the reactive case, the result obtained below can be extended analogously to Section 5.1.2, but without having compact integral-free representations.

6.1 Reactive approach

In the reactive approach the access leads to a mismatch, if after the downstream delay time, the Markov process \mathbf{Q} is in a different state as at the time when the update was sent out (assumed here to be $t = 0$). Due to stationarity, the probability of being in state i at time $t = 0$ is just the steady-state probability π_i . Hence, by conditioning on the downstream delay time, we obtain the following

$$\text{mmPr} = 1 - \int_{t=0}^{\infty} \sum_{i=1}^S \pi_i [\exp(\mathbf{Q}t)]_{i,i} f_D(t) dt \quad (28)$$

Hence, the above equation provides the solution to the $MAP(rec)|G|1|react(G)$ case. In the following, we treat two special cases for the down-stream delay, the exponential and subsequently the general matrix-exponential:

6.1.1 Exponential downstream delay

Now D is exponential distributed with rate ν and Equation (28) can be reduced to

$$\begin{aligned} \text{mmPr} &= 1 - \nu \sum_{i=1}^S \pi_i \left[\int_{t=0}^{\infty} \exp([\mathbf{Q} - \nu \mathbf{I}]t) dt \right]_{i,i} \\ &= 1 - \nu \sum_{i=1}^S \pi_i [(\mathbf{Q} - \nu \mathbf{I})^{-1}]_{i,i}. \end{aligned}$$

6.1.2 General matrix-exponential downstream delay

If the downstream delay has a marginal distribution that is matrix-exponential, i.e.

$$f_D(t) = \mathbf{p}_D \mathbf{B}_D \exp(-\mathbf{B}_D t) \boldsymbol{\epsilon}'_D,$$

integral-free expressions for Eq. (28) can be obtained as follows:

$$\begin{aligned}
\text{mmPr} &= 1 - \int_{t=0}^{\infty} \sum_{i=1}^S \pi_i [\exp(\mathbf{Q}t)]_{i,i} \mathbf{p}_D \mathbf{B}_D \exp(-\mathbf{B}_D t) \boldsymbol{\varepsilon}'_D dt \\
&= 1 - \sum_{i=1}^S \pi_i \int_{t=0}^{\infty} \mathbf{e}_i \exp(\mathbf{Q}t) \mathbf{e}_i' \mathbf{p}_D \mathbf{B}_D \exp(-\mathbf{B}_D t) \boldsymbol{\varepsilon}'_D dt \\
&= 1 - \sum_{i=1}^S \pi_i [\mathbf{e}_i \otimes (\mathbf{p}_D \mathbf{B}_D)] [(-\mathbf{Q}) \oplus \mathbf{B}_D]^{-1} [\mathbf{e}_i' \otimes \boldsymbol{\varepsilon}'_D]. \tag{29}
\end{aligned}$$

Hereby, \mathbf{e}_i is a row vector with all components zero excepts for the i -th component, which is equal to one.

6.2 Outlook on proactive strategies

The proactive cases can be further distinguished to event-driven and periodic. Starting with the periodic approach, $MAP[rec]|G|1|periodic(M)$, a general formula for the mmPr is obtainable if we condition on the time of the last event and split the probability of no useful update in the independent (by the Poisson assumption of the delay process) events 1) no update since last event and 2) for each possible change at last event π_i the last update does not carry this information. Formally, use the intensity function τ defined in Equation (14) and define the following stopping time (with respect to the natural filtration), which denote the last update before time t

$$\tau_t = \sup\{I_k | I_k \leq t, k \in \mathbf{Z}\},$$

where $\{I_k, k \in \mathbf{Z}\}$ is the update process defined in Section 4.4.

Furthermore, if we let $\widetilde{\mathbf{Q}}$ denote the time-reversed version of \mathbf{Q} (see [1, page. 58]), then

$$\begin{aligned}
\text{mmPr} &= \int_{-\infty}^0 \mathbb{P}(I \cap [t, 0] = \emptyset, E(\tau_t) \neq E(t)) A_E(dt) \\
&= \int_{-\infty}^0 \mathbb{P}(I \cap [t, 0] = \emptyset) \mathbb{P}(E(\tau_t) \neq E(t)) A_E(dt) \\
&= \int_{-\infty}^0 e^{-\int_t^0 \tau(s) ds} \sum_{i=1}^n (1 - \pi_i \mathbb{P}(E(\tau_t) = E(t) | E(t) = i)) A_E(dt) \\
&= \int_{-\infty}^0 e^{-\int_t^0 \tau_{FD}(-s) ds} \sum_{i=1}^n \left(1 - \pi_i \int_{-\infty}^t \left(\sum_{j \neq i} [e^{\widetilde{\mathbf{Q}}(t-s)}]_{ji} \tilde{q}_{ij} \right) F_{\tau_t}(ds) \right) A_E(dt).
\end{aligned}$$

Although, this formula apparently always require some sort of numerical integration, one can in special cases obtain considerably simplified expressions. A detailed study of this periodic case with numerical examples is left for future research.

For the proactive event-driven approaches with full or incremental updates, one approach is to analyze the process of receiving updates at the requester. This process is the output process of an $E/D/\infty$ queue. Unfortunately, the output process of this queue is not known for general E , so we have to limit ourselves to the Poisson case, i.e. the information element is a Markov process with generator \mathbf{Q} such that the state-leaving rate is the same for all states, i.e.

$$\text{diag}(\mathbf{Q}) = -\lambda \mathbf{e}'.$$

For this Poisson assumption, the output process of the $M/GI/\infty$ queue is also Poisson, hence the process of receiving updates is Poisson with rate λ and both its forward and backwards recurrence time are exponentially distributed with same rate λ . An extended approach as in Sections. 4.2.1 and 4.2.2 appears promising for such a scenario, but its study is left for future research.

6.3 Summary

In this section we consider ways to compute analytically the mmPr for the scenario of $N = 1$ information providers, at which the information element can change back to previous values, as described by a continuous time Markov process. In summary, we can conclude:

- Reactive approach: For $MAP(rec)|G|1|react(G)$, a general formula, Eq. (28) for the mmPr was derived and simplified to an integral-free expression, Eq. (29), in the case where the downstream delays are matrix-exponential distributed.
- Proactive strategies: For the periodic case a general formula was given for the $MAP[rec]|G|1|periodic(M)$ case, and it was indicated how one could proceed to get simplified expressions. This is left for future research. The event driven case was only touched upon as a more detailed analysis involves the study of output processes for general queueing systems.

7 Quantitative results and validation

This section uses the models from Sections 4, 5, and 6 to obtain and discuss numerical results for selected example scenarios.

7.1 Single information provider, monotonous type event process

First, we consider the case of a single information provider, but while varying the distributions of the inter-event process and downstream delay process. In particular, we want to investigate numerically the impact of different distribution types, namely Erlangian distributions with smaller coefficient of variation than an exponential, and truncated Power-Tail distributions with a tail-exponent of $\alpha = 1.4$ which leads to unboundedly growing variance for increasing number of phases. See the appendix for details about these distribution types.

7.1.1 Parametric study for the Poisson event process




Fig. 8. Comparison of mismatch probabilities in the different remote access strategies for a Poisson event process and varying downstream delay [19].

Exponential network delays: First, we focus on the case of a Poisson event process together with exponential or deterministic downstream delay of varying rate, i.e. the $M|M|1|x$ and $M|D|1|x$ cases: Figure 8 shows the results for the mmPr as computed by the analytic models for the different remote access strategies, for the assumption of a Poisson event process with rate $\lambda = 1$. In the proactive periodic case, the period is iid exponentially distributed with varying rate $\tau = 10^{-2}, \dots, 10$. As the analysis in Section 4.2.1 shows, the reactive and the proactive event-driven strategy with full updates lead to exactly the same mmPr (dashed curve). The proactive event-driven strategy with incremental updates shows a slightly higher mmPr (solid curve). According to the analysis in Section 4.2.2, this mmPr in the event-driven incremental case for the considered Poisson event case is insensitive of the delay distribution; consequently, the $M|D|1|event|incr$ case is also represented by

the same solid curve. Furthermore, according to Section 4.3, the event-driven strategy with full updates and hence the reactive strategy with deterministic downstream delay are also captured in this solid curve. Hence, only two curves are needed to represent the six cases for reactive strategies and proactive event-driven strategies.

An additional set of dashed-dotted curves reflects the mmPr of the periodic cases, $M|M|1|periodic(M)$, with different rate τ of the period. The following additional observations can be made from Figure 8:

- The reactive strategy in the case of deterministic downstream delays, $D \equiv 1/\nu$, (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. This observation is investigated further below via the use of Matrix-exponential distributions.
- For very short downstream delays (large ν) the mmPr of both the reactive and the proactive event-driven strategies decay asymptotically as $mmPr(\lambda, \nu) \sim \lambda/\nu$ 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. This is explainable with the arguments in Section 4.3, since no message reordering will occur for infinitely fast networks.
- In the limit case $\nu \rightarrow \infty$, the proactive periodic approach shows a limit of $\lim_{\nu \rightarrow \infty} mmPr(\lambda, \nu, \tau) = \lambda/(\lambda + \tau) > 0$. Consequently, for large ν eventually, the periodic approach will at some point always perform worse than the event-driven and reactive approaches.

Non-exponential network delays: Figure 9 shows the impact of different network delay distributions. For the reactive strategy, $M|ME|1|react(G)$, two curves are shown in the bottom of the figure, which are identical for the proactive event-driven strategy with full updates, $M|ME|1|event|full$. Since the down-stream delay distribution is irrelevant for the proactive event-driven incremental strategy, $M|GI|1|event|incr$, it results in a horizontal line, shown dashed-dotted in the figure.

The upper of the two curves for the reactive strategy (marked with circles) represents the case of an Erlangian- T delay distribution in the reactive/proactive-event-driven-full approach, i.e. with increasing T on the x-axis, the coefficient of variation is reduced as $1/T$ converging to a deterministic distribution. This decrease in variance actually results in an increased mmPr. The lower curve (marked with '+') shows the mmPr for a TPT- T distributed network delay. The TPT distribution uses a tail-exponent of $\alpha = 1.4 < 2$, hence for an in-




Fig. 9. Mismatch Probability for reactive and the two different proactive strategies for downstream delay process which are renewal processes with matrix-exponential representation: Shown for Erlangian and TPT distributions with increasing number of phases along the x -axis.

finite number of phases, it shows infinite variance. The mmPr decays with increasing T but appears to converge to a value slightly below 25%.

The upper two curves in Figure 9 show the periodic case, $M|ME|1|periodic(M)$, with a Poisson rate of $\tau = 2$ for the period. For this choice of τ , the mmPr values are always higher than for the other strategies. The qualitative behavior when increasing the number of phases of the Erlangian and TPT delay distribution is the same as for the reactive strategy, namely with increasing variance (TPT case), the mmPr drops but converges to a value slightly above 0.45; for decreasing variance (Erlangian case), the mmPr increases and converges towards a value close to 0.6, provided by the $M|D|1|periodic(M)$ case.

7.1.2 Matrix-exponential event processes

Similar qualitative behavior as in the previous section is observed when distribution of the inter-event times is varied, see Figure 10: The change from exponential towards a deterministic distribution (Erlangian with many phases) actually results in a significant increase of the mmPr for all strategies, while the use of TPT distributed inter-event times actually reduces the mmPr. The middle two curves (green colour) thereby represent the case of deterministic downstream delays, in which four of the strategies are actually equivalent, see Section 4.3.

When both processes, the event and the downstream delay, are represented by TPT respectively Erlangian distributions, the impact on the mmPr is




Fig. 10. Mismatch Probability for reactive and proactive (full) strategy for event process which are renewal processes with matrix-exponential representation: Shown for Erlangian and TPT distributions.




Fig. 11. Mismatch Probability for reactive and proactive (full) strategy for event and delay processes which are renewal processes with matrix-exponential representation: Shown for Erlangian and TPT distributions.

strongest, i.e. the worst case in the considered candidate set occurs, when both distributions are Erlangian, i.e. the $D|D|1|x$ case, see Figure 11.




Fig. 12. Mismatch Probability for reactive strategy for an event process which is an ON/OFF process with same average duration of the ON+OFF cycle.

7.2 Recurrent event process: ON/OFF process

We now look at the mmPr in cases that the information element can change back to a previous value. We use the example of a binary information element, e.g. the state of a device being either busy or idle, here also called ON and OFF. Therefore, the event process is a two-state continuous time Markov chain, where the average change rate is kept consistent with the parameter settings at the end of the previous settings, namely, $\overline{ON} + \overline{OFF} = 2$ so that the average inter-event time is still kept at $\bar{E} = 1$. However, ON-and OFF state leaving rates are varied so that they show different holding times, i.e. we vary the ratio

$$\kappa = \frac{\overline{OFF}}{\overline{ON}}$$

while keeping their sum constant.

Figure 12 shows the resulting mmPr for three different delay distributions: Erl-20 in the upper set of curves, exponential in the intermediate curves, and TPT-20 in the lower curves. For each delay distribution, three curves are given: the recurrent MAP process (solid) which shows a lower mmPr than a monotonous-type Markov Jump process (dashed), since there is some probability that an even number of changes has happened since sending the response, which then would lead to a match of the remote information element. The mmPr has a maximum when ON and OFF period show same average duration, at which point the information element changes form a homogeneous Poisson process, i.e. the $MJ|ME|1|react$ case is equivalent to a $M|ME|1|react$ case at $\kappa = 1$. The latter is shown as dotted horizontal line in Figure 12.




Fig. 13. Mismatch Probability for multiple information providers, all Processes are Poisson. The event-rate per IP is scaled down by a factor of N so that the total event rate remains constant.

When κ goes to zero or infinity, the mmPr approaches zero, since in these limit cases, the ON/OFF process is actually only dominated by one of the two states, namely OFF if $\kappa = 0$, and ON for $\kappa \rightarrow \infty$.

7.3 Multiple information providers

Finally, we provide numerical results for multiple information providers, $N > 1$ thereby also showing empirically the accuracy of the approximations by the limit theorems. We restrict ourselves to exponential distributions here, i.e. $M|M|N|x$ cases.

Figure 13 shows the resulting mmPr. Note that although the reactive case results in the same mmPr as the proactive event-driven case with full updates for $N = 1$, this is not the case any more for larger N . In fact, for the applied linear scaling in the figure, the mmPr of the reactive case grows to 1 when $N \rightarrow \infty$.

Also, it becomes obvious, that the difference between full updates and incremental updates for the event-driven strategies vanishes for larger N . Both cases have the same limit, see Section 5.2. Furthermore, the limit approximation is actually exact for the event-driven incremental strategy, while it is only an approximation for the case of full updates.

8 Summary and outlook

This paper has developed a methodology and explicit analytic solutions for the quantitative analysis of different strategies for remote access to dynamically changing information elements.

The analytic results lead to the following conclusions:

- For a single information provider and monotonous-type event processes, the mmPr of the reactive strategy and the proactive-event driven strategy with full updates are identical.
- The mmPr of the proactive event-driven strategy with incremental updates is smaller or equal to the full update case. In the case of a Poisson Event process the mmPr is independent of the downstream delay and described by the busy probability of an $M/G/\infty$ queue.
- For networks without loss and re-orderings (FIFO type networks), the reactive strategy and all proactive event-driven strategies lead to the same mmPr.
- For the proactive, periodic strategies, an explicit solution Equation (15) has been obtained for the scenario, when the instances of sending updates form a Poisson process. Integral-free solutions (17) or equivalently (18) result, when all participating processes are Poisson.
- The mmPr for the reactive case in case of recurrent Markov Event processes (which may change back to previous values) is obtained in general integral-free form, Equation (29). An outlook on how the proactive strategies may be approached is furthermore given in Section 6.2.
- The case of multiple information providers is treated for all strategies in Section 5 and limit results are obtained that allow to identify interesting differences in scaling behavior as well as allow to obtain simple approximative expressions.

The analysis has subsequently been applied to scenarios with general matrix-exponential distributed inter-event times and network delays. The numerical results show that for the given settings, the high-variance case (truncated Power-tail distributions for the events/delays) actually leads to smaller mismatch probability than the exponential case, except for the case of event-driven incremental updates, which is insensitive to the delay distribution for a Poisson event process. Analogously, the use of Erlangian distributions and in the limit deterministic distributions increases the mismatch probability for all schemes.

Furthermore, the analysis was applied to the case of a binary information element, which toggles between 2 values, e.g. an ON/OFF process. The mismatch probability is in this case smaller than for a monotonous-type event

process and it is largest, if the average of the ON duration and OFF duration are identical. Finally, also the case of multiple information elements has been analysed numerically, showing the different scaling behavior of the reactive strategy (factor $N \log N$ in case of exponential delays) and the proactive strategies (factor N), and the accuracy of approximations that were obtained from the limit theorems.

Other relevant scenarios, e.g. when ordering update messages according to receive time as opposed to using sequence numbers created at the information provider, will be considered in future work. Furthermore, the proactive cases for recurrent event processes have to be analysed further as outlined in Section 6.2. Finally, the application of the mmPr analysis to the actual use-cases of routing, context-sensitive networking, and replicant consistency for optimistic replication strategies will likely lead to further model refinements.

Acknowledgments

The authors would like to thank Søren Asmussen for hints on general limits of extreme values, Lester Lipsky for help on pinpointing Kronecker product representations of matrix-exponentials, and Thomas Mikosch for hints on Laplace transforms of extreme value distributions.

Parts of this research is performed in the framework of the EU-FP6 projects MAGNET (www.magnet-ist.org) and HIDENETS (www.hidenets.aau.dk). The authors would like to thank the consortia members for the fruitful discussions.

References

- [1] S. Asmussen. *Applied Probability and Queues*. Springer Verlag, New York, NY, 2nd edition, 2003.
- [2] P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, NY, 1968.
- [3] M. Bozinovski. *Fault-tolerant Platforms for IP-based Session Control Systems*. PhD thesis, Aalborg University, 2004.
- [4] E. Castillo, A. S. Hadi, N. Balakrishnan, and J. M. Sarabia. *Extreme Value and Related Models with Applications in Engineering and Science*. Wiley, New York, NY, 2005.
- [5] W. Chen, S. Toueg, and M. K. Aguilera. On the quality of service of failure detectors. *IEEE Transactions on Computers*, 51(5):561–580, May 2002.

- [6] D.J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes. Volume I: Elementary Theory and Methods*. Springer Verlag, New York, NY, 2nd edition, 2003.
- [7] M. G., R. L. Olsen, M. G. Genet, and R. Tafazolli. Service management platform for personal networks. *Proceedings of IST Summit, Dresden*, 2005.
- [8] M. Greiner, M. Jobmann, and L. Lipsky. The importance of power-tail distributions for telecommunication traffic models. *Operations Research*, pages 313–326, 1999.
- [9] A. Helal, A. Heddaya, and B. Bhargave. *Replication Techniques in Distributed Systems*. Kluwer Academic Publishers, Norwell, MA, 1996.
- [10] S. Lee, Y. Yu, S. Nelakuditi, Z.-L. Zhang, and C-N. Chuah. Proactive vs. reactive approaches to failure resilient routing. In *Proceedings of IEEE INFOCOM, Hong Kong*, 2004.
- [11] L. Lipsky. *Queueing Theory: A Linear Algebraic Approach*. McMillan and Company, New York, NY, 1992.
- [12] L. Liu and D.-H. Shi. Busy period in $GI^X/G/\infty$. *Journal of Applied Probability*, 33:815–829, 1996.
- [13] Y. Liu, R. L. Olsen, M.B. Hansen, and H.-P. Schwefel. Analytic computation of routing errors for source routing in dynamic topologies. In preparation.
- [14] Y. Liu and H.-P. Schwefel. An enhanced optimized link-state routing with delayed event driven link-state updating. In *Proceedings of the 16'th International Symposium on Personal, Indoor, and Mobile Radio Communication, Berlin*, 2005.
- [15] D. Miorandi and E. Altman. Connectivity in one-dimensional ad hoc networks: A queueing theoretical approach. *Wireless Networks*, 12:573–587, 2006.
- [16] H. Murakami, R.L. Olsen, H. P. Schwefel, and R. Prasad. Managing personal network specific addresses in naming schemes. In *Proceedings of WPMC, Aalborg, Denmark*, 2005.
- [17] M. F. Neuts. *Matrix-Geometric Solutions in Stochastic Models*. John Hopkins, University Press, Baltimore, 1981.
- [18] R. L. Olsen, Y. Liu, M.B. Hansen, and H.-P. Schwefel. Mismatch probabilities and their impact on context sensitive service-discovery. In preparation.
- [19] R. L. Olsen, H. P. Schwefel, and M. B. Hansen. Quantitative analysis of access strategies to remote information in network services. In *Proceedings of IEEE Globecom*, 2006.
- [20] S. Panagiotakis and A. Alonistioti. Intelligent service mediation for supporting advanced location and mobility aware service provisioning in reconfigurable mobile networks. *IEEE Wireless Communication*, 9(5):28–38, 2002.

- [21] T. Renier, E. Matthiesen, Hans Peter Schwefel, and Ramjee Prasad. Inconsistency evaluation in a replicated IP-based call control center. In *Proceedings of ISAS*, Helsinki, Finland, 2006.
- [22] H.-P. Schwefel and L. Lipsky. Impact of aggregated, self-similar ON/OFF traffic on delay in stationary queueing models (extended version). *Performance Evaluation*, 41:203–221, 2001.

A Candidate matrix-exponential distributions

A.1 Erlangian distribution

An Erlangian- T distribution is the convolution of T identical exponentials, i.e. the distribution of the sum of T iid exponential random variables. Its probability density function is:

$$f(x) = \mu \frac{(\mu x)^{T-1}}{(T-1)!} e^{-\mu x}. \quad (\text{A.1})$$

Note that the density at the origin $x = 0$ is $f(0) = 0$ as opposed to Hyper-exponential distributions, whose density has its maximum at the origin.

It follows for the first two moments of the Erlangian- T distribution:

$$\mathbb{E}(X) = \frac{T}{\mu} =: \bar{x}, \quad \mathbb{E}(X^2) = \frac{T^2 + T}{\mu^2}. \quad \Rightarrow \quad C^2 = \frac{1}{T} < 1.$$

An Erlangian- T distribution with high T can be used as approximation for a deterministic distribution. An Erlangian- T distribution can be represented as a matrix-exponential with matrix representation:

$$\mathbf{p} = [1, 0, \dots, 0], \quad \mathbf{B} = \mu \begin{bmatrix} 1 & -1 & & & \\ & 1 & -1 & & \\ & & \ddots & \ddots & \\ & & & 1 & -1 \\ & & & & 1 \end{bmatrix}.$$

A.2 Truncated power-tail distribution

In order to model distributions with large variance, frequently hyper-exponential distributions are used, whose pdf is a linear combination of different exponential densities. Choosing the weights and the rates of the exponential phases in a special way, namely both geometrically decaying, but with different factors,

$$R_{Y_T}(x) = \frac{1 - \theta}{1 - \theta^T} \sum_{i=0}^{T-1} \theta^i \exp \left[\frac{-\mu_T}{\gamma^i} x \right], \quad (\text{A.2})$$

the resulting complementary distribution functions show Power-law behavior, $R(x) \sim x^{-\alpha}$ for some orders of magnitude before they drop off exponentially, see [8]. The higher the number of phases, T , the later the drop-off occurs. The exponential drop-off is characterized in more detail in [22] by the so-called Power-Tail Range.

The variable θ can be chosen freely in the range $0 < \theta < 1$. For larger value of θ , more phases are necessary to obtain the same PT Range as for lower θ . In order to show Power-Law behavior with exponent α , and to have mean \bar{x} , the other constants in (A.2) have to be (see [8]):

$$\gamma = \left(\frac{1}{\theta} \right)^{1/\alpha},$$

$$\mu_T = \frac{1 - \theta}{1 - \theta^T} \frac{1 - (\theta\gamma)^T}{1 - \theta\gamma} \frac{1}{\bar{x}}.$$

The truncated powertail distribution admits the following matrix-exponential representation:

$$\mathbf{p}_T = \frac{1 - \theta}{1 - \theta^T} [\theta^0, \dots, \theta^{T-1}].$$

$$\mathbf{B}_T = \mu_T \begin{bmatrix} 1/\gamma^0 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & 1/\gamma^{T-1} \end{bmatrix}.$$