

Probabilistic models for access strategies to dynamic information elements

Bøgsted, Martin; Olsen, Rasmus Løvenstein; Schwefel, Hans-Peter

Published in:
Performance Evaluation

DOI (link to publication from Publisher):
[10.1016/j.peva.2009.08.015](https://doi.org/10.1016/j.peva.2009.08.015)

Publication date:
2010

Document Version
Early version, also known as pre-print

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Bøgsted, M., Olsen, R. L., & Schwefel, H.-P. (2010). Probabilistic models for access strategies to dynamic information elements. *Performance Evaluation*, 67(1), 43-60. <https://doi.org/10.1016/j.peva.2009.08.015>

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

Martin B. Hansen ^{a,b} Rasmus L. Olsen ^b Hans-Peter Schwefel ^{b,c}

^a*Department of Haematology and Centre for Cardiovascular Research, Aalborg Hospital, Aarhus University Hospital, Denmark*

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

^c*Forschungszentrum Telekommunikation Wien (ftw), Donau-City-Strasse 1, 1220 Vienna, Austria*

Abstract

In various network services 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 on mismatch probability of different distribution types for the network delays as well as for the time between changes of the information element 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:* martin.boegsted@rn.dk (M.B. Hansen), rlo@es.aau.dk (R.L. Olsen), hps@es.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, [1].
- 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 [15].
- 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 [8].
- For highly dependable networks and services, resilience is obtained by replication of services, which requires state-updates at remote replicas in order to avoid inconsistency [21,6,4] .

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 change happens due to certain external events. The nodes sensing these events are called *information providers* in this paper. The information provider can send messages to the requester (downstream) through an interconnecting network, which imposes a delay and possible message loss and re-ordering of the downstream messages. If needed, the requester can also send messages to the information providers, here referred to as 'upstream' communication.

Two basic types of solutions for such remote access are well known, see e.g., [12,17]:

- (1) **Reactive**, 'on-demand' access: Whenever the requester 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.
 - (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 consider the following three performance metrics in relation to the access strategies:

- (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** (mmPr): 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.

Depending on the requirements of the specific use-case, the goal may be to minimize one of these metrics while constraining the others, see [25] for an example of optimization in a such scenario. Other scenarios which makes use

of mismatch probability has been shown in [22,23].

The assumptions in this paper regarding the specific type of information are rather general, in particular, no assumptions on the semantics are made:

- 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 [11].
- It is irrelevant for the analysis in the paper whether the information element is a single-valued integer variable or a complex data-structure.
- The information provider will mark messages with a monotonously increasing sequence number, which is used by the requester to reorder received messages accordingly.

The rest of the paper is organized as follows: In Section 2, we define our notation and describe some mathematical preliminaries. The first strategy we describe are the proactive event-driven and periodic strategies which are treated in Section 3 and Section 4, respectively. This is followed by the reactive access strategy which is treated in Section 5. Finally, the paper provides and discusses a set of numerical results in Section 6 and concludes with a summary and outlook in Section 7.

2 Mathematical preliminaries

Let $\tau = \{T_i, i \in \mathbf{Z}\}$ be the times of occurrences of some phenomenon, where T_i is an increasing sequence of event times numbered such that T_0 is the event just before 0. If we put $X_i = T_i - T_{i-1}$, $i \in \mathbf{Z}$ and assume that the sequence of X_i 's are independent and identically distributed (i.i.d.) then τ is called a renewal process, see [7].

With an abuse of notation, a random variable with the same distribution as the X_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 .

We call τ stationary, if 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 ($-\infty < t < \infty$), [7, Definiton 3.2.I]. We define the forward recurrence time $V = T_1$ and the backward recurrence time $U = -T_0$ and their distribution functions as $B_X(t) = \mathbb{P}(V \leq t)$ and $A_X(t) = \mathbb{P}(U \leq t)$. Whenever τ is stationary the distribution of the backwards recurrence time is the same as the forward recurrence time, [2, p. 150]. Moreover,

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

2.1 Matrix-exponential distributions

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

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

is said to have a matrix-exponential representation with generator or representation $\langle \mathbf{p}, \mathbf{B} \rangle$, see [16,13]. Notice that we use the notation of [13] in which \mathbf{B} has a positive diagonal, and is negative or zero outside the diagonal. Its moments are expressed as

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

Special cases of matrix-exponential distributions are Hyper-Exponential distributions and Erlangian distributions. A special version of the former, namely truncated Power-Tail (TPT) distributions [10], are used in Section 6 to illustrate the mmPr behavior for scenarios with high variance in inter-event and downstream delay processes, see Appendix A. The Erlangian distributions are used to illustrate the behavior when the variance of participating distributions is decreasing in Section 6.

In case of delay offsets, which frequently happens to some extend, a shifted delay distribution is considered in Appendix B.

2.2 Kronecker product and sums

The integral representations of mmPr for Phase-type distributed event and delay processes in later sections can in some cases be written in integral-free

form utilizing Kronecker product representations. We follow here the notation for *Kronecker-Products* and *Kronecker-Sums*, with $\mathbf{A} \in \mathbf{R}^{m \times m}$ and $\mathbf{B} \in \mathbf{R}^{n \times n}$, as used in [9]:

$$\mathbf{A} \otimes \mathbf{B} := \begin{bmatrix} a_{11}\mathbf{B} & \dots & a_{1m}\mathbf{B} \\ \vdots & & \vdots \\ a_{m1}\mathbf{B} & \dots & a_{mm}\mathbf{B} \end{bmatrix},$$

$$\mathbf{A} \oplus \mathbf{B} := \mathbf{A} \otimes \mathbf{I}_{\dim(B)} + \mathbf{I}_{\dim(A)} \otimes \mathbf{B}.$$

Using these Kronecker representations, integrals of products of Matrix-exponential distributions can be simplified as follows,

$$\int \mathbf{p}_1 \exp(-\mathbf{B}_1 t) \boldsymbol{\varepsilon}'_1 \mathbf{p}_2 \exp(-\mathbf{B}_2 t) \boldsymbol{\varepsilon}'_2 dt = \mathbf{p}_1 \otimes \mathbf{p}_2 \int \exp(-(\mathbf{B}_1 \oplus \mathbf{B}_2)t) dt \boldsymbol{\varepsilon}'_{\dim(B_1) \cdot \dim(B_2)},$$

which is later used to achieve integral-free representations of the mismatch probability.

3 Analytic results for the proactive, event-driven access strategy

Now we introduce the first of the three access strategies. We start by presenting the abstracted model, and define the involved stochastic processes, whereafter we describe the performance metrics involved. In this study, we first investigate the case of $N = 1$ information provider, followed by the case of $N > 1$.

3.1 Abstraction model of event based update

The proactive, event-driven access strategy assumes that the requester has subscribed to the information providers, which will notify the requester when their part of the information element changes value. Figure 2 shows the abstracted model for this access strategy.

At discrete points in (continuous) time the requester needs the value of the information element consisting of the N parts, $E^{(n)}(t)$, $n = 1, \dots, N$. The time of requests are identified by

the **request process**, $R = \{R_k, k \in \mathbf{Z}\}$.




Fig. 2. Proactive event-driven update: the request at time R_k results in a correct value, while R_{k+1} leads to a mismatch for $E^{(1)}$, since the update is in transit at request time. This is not the case for $E^{(N)}$ though. For R_{k+2} the opposite is the case.

At the information provider, the needed information element changes value at certain points in time, determined by

the **event process**, $E^{(n)} = \{E_i^{(n)}, i \in \mathbf{Z}\}$, $n = 1, \dots, N$,

where $E_i^{(n)}$ is an increasing sequence of event times, numbered such that $E_0^{(n)}$ is the event just before 0. For the N event processes.

The time between sending a notification from the information provider until receiving it at the requester is described by

the **downstream delay**, $D^{(n)} = \{D_i^{(n)}, i \in \mathbf{Z}\}$, $n = 1, \dots, N$.

These delays correspond to the end-to-end delays between information provider n and the requester. Message drops can be included via degenerated distributions (with probability mass at infinity).

Unless otherwise specified, we assume joint independence between the event and delay processes, as well as joint stationarity.

Sub cases: With respect to the content of the update messages, we need to further distinguish between [17]

- Incremental updates: Only the difference of the value of the information element since the previously sent update is transmitted.
- Full updates: The complete information element is provided in each update.

In both cases we assume that the ordering of the messages at the requester does not matter (e.g. if for the incremental case, the updates are commutative) or that re-ordering of messages is performed via sender sequence numbers.

Notation: We introduce a Kendall-alike notation using $|$ for describing the access strategy, and the respective process specification, based on specifying first 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. Otherwise, the prefix can be extended as $E|D|N|RS|US$, where 'RS' is the request strategy and US indicates the content of the update messages. For the proactive, event-driven case, $RS='event'$. The update strategy 'US' is by default 'full', but can also be specified to be 'incr' (incremental). If 'RS' and 'US' are not specified, then by default, an event-driven strategy with full update is assumed.

For instance, $ME|ME|4|event|incr$ specifies a remote access to four information providers, at which the event processes and the corresponding downstream delays are matrix-exponential renewal processes; the system uses an event-driven access strategy with incremental updates.

The value of the information element by default never changes back to previous assumed values. In order to allow for other cases (which we here call 'recurrent event processes'), later we also allow the value of the information element to be described by the state of an ergodic Markov Process; this type of event process is denoted by 'MP'. If we only consider the times when the Markov process leaves a state, but the information element does not take any previous values (e.g. because its value increases monotonically) we call it a Markov Jump process (MJ).

3.2 Access delay and network overhead

The network overhead of the proactive, event-driven strategy is determined by the event process and the amount of information providers sending updates denoted by s_d

$$V_{[G|G|N|event|\{incr,full\}]} = \sum_{i=1}^N \frac{s_d^{(i)}}{\mathbb{E}(E^{(i)})}. \quad (2)$$

Depending on the data structure of the information element, typically the incremental update will lead to less traffic than the full update, since in most cases $s_d^{(i)}(Incr) \leq s_d^{(i)}(Full)$. Since the requester will get the locally stored update at time of request, the access delay is zero.

3.3 Mismatch probability for a single information provider ($N=1$)

Now we turn our attention to the mismatch probability, which needs to be treated differently whether full or incremental updates are used.

3.3.1 Mismatch probability using 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 no mismatch for the request at 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, with $B_E(t)$ being the cdf of the backward recurrence time, and \bar{B}_E its complement (see Section 2):

$$\begin{aligned}
\text{mmPr}[G|1|event|full] &= 1 - \int_0^\infty \mathbb{P}(D_i \leq t | B = t) B_E(dt) \\
&= 1 - \int_0^\infty \int_0^t F_D(ds) B_E(dt) \\
&= 1 - \int_0^\infty \bar{B}_E(s) F_D(ds) \\
&= 1/\mathbb{E}(E) \int_0^\infty \bar{F}_D(t) \bar{F}_E(t) dt.
\end{aligned} \tag{3}$$

For Matrix-Exponential Event and Delay processes, Equation (3) results in:

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

For the expression using a shifted time delay, see Appendix B.

Under the assumption that the event process, E , is a Poisson process with rate λ we can obtain from Equation (3):

$$\begin{aligned}
\text{mmPr}[M|GI|1|event|full] &= 1 - \int_0^\infty e^{-\lambda t} F_D(dt) \\
&= 1 - \mathcal{L}\{F_D\}(\lambda)
\end{aligned} \tag{5}$$

where the last term is the Laplace-Stieltjes transform of the cdf of the downstream delay, evaluated at λ . For the cases of a deterministic delay of value

$1/\nu$, and for exponential delay with rate ν , the mmPr can be expressed as [17]:

$$\begin{aligned}\text{mmPr}_{[M|D|1|event|full]} &= 1 - e^{-\lambda/\nu} \\ \text{mmPr}_{[M|M|1|event|full]} &= \frac{\lambda}{\lambda + \nu}.\end{aligned}$$

Note that the ratio of the previous two expressions for $\nu \rightarrow \infty$ converges to 1, i.e. the case of exponential and deterministic downstream delay show the same limit behavior for small delays.

3.3.2 Mismatch probability using incremental updates

When update messages only contain incremental information, the requester accesses the correct information, only 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.

The transmission network itself can be viewed as a $G/G/\infty$ -queueing model, for which the arrival process is the event process (as at these moments, an update message is sent out by the information provider). The service time of the queueing system is the downstream delay, and as delays imposed by the network to the downstream messages are independent of the number of messages in transit, an infinite server queue has to be used. Therefore, the transmission network can be modeled by an $E/D/\infty$ queue.

Let the stochastic process Q_t denote the number of updates in transit at time t . Thus, a customer being served models an update message in transit. Henceforth, the mismatch probability at request R_k is the probability that Q_{R_k} is strictly greater than zero. By stationarity we have that

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

Utilizing the well-known results for the queue-length probabilities of the $M/GI/\infty$ queue, [2], [18], 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 (with mean \bar{D}), results in a mismatch probability of

$$\text{mmPr}_{[M|GI|1|event|incr]} = 1 - \exp(-\lambda\bar{D}). \quad (6)$$

3.3.3 Scenarios without message loss and reordering

In case of deterministic downstream delay, the mmPr is actually identical for the event-driven incremental strategy and the event-driven full update strategy, since no reordering and no message loss can occur. Thus the mmPr can be derived from Equation (3), which becomes

$$\text{mmPr}_{[ME|D|1|event|\{incr,full\}]} = \frac{\mathbf{p}_E \mathbf{V}_E (\mathbf{I} - \exp(-\mathbf{B}_E \bar{D})) \boldsymbol{\varepsilon}'}{\mathbf{p}_E \mathbf{V}_E \boldsymbol{\varepsilon}'_E}. \quad (7)$$

This equation provides the integral-free expression to the $ME|D|1|event|incr$ scenario, which is identical to the $ME|D|1|event|full$. Later in Section 5, we will see that the reactive strategy also has the same mmPr, hence as conclusion: In scenarios in which no message loss and no packet reordering can result (i.e. the transmission network shows first-in-first-out behavior), these three strategies lead to the same mmPr.

3.4 Mismatch probability for multiple information providers ($N > 1$)

In the case where more than one information element is required, a mismatch is obtained if at least one of the information providers yields a mismatch. As all $2N + 1$ participating stochastic processes are considered pairwise mutually independent, the probability of mismatch can be obtained by the following function of the individual mismatch probabilities

$$\text{mmPr}_{[GI|GI|N|event|.]} = 1 - \prod_{i=1}^N (1 - \text{mmPr}_{[GI|GI|1|event|.]}^{(i)}). \quad (8)$$

Mismatch probability limits for multiple information providers

Analyzing the behavior of the mismatch probability for large N is useful for two purposes: (1) it provides insight on scaling properties of the systems, and (2) such limit results can be used for computationally efficient approximations of the mmPr for scenarios with large N . These are in particular useful for on-line calculations of the mmPr by participating nodes, see [22,23] for examples. The computational gain by these approximations turns out to be in fact larger for other access strategies, see Section 5.3.

Consider the case that the event process is Poisson and that for increasing N , the Poisson rate of the individual event processes is scaled down as $\lambda_N = \lambda/N$. Then, this scaling of the event-rate in the limit creates a scenario, in which the update messages sent by the same information provider are not subject

to reordering any more, since the time interval between updates becomes infinite as N goes to infinity. As a result, the proactive event-driven full and incremental strategies are in the limit equivalent, if there is no message loss. Therefore we obtain using Equation (8) and (6):

$$\begin{aligned}
\lim_{N \rightarrow \infty} \text{mmPr}[M|GI|N|event|incr] &= \text{mmPr}[M|G|N|event|full] \\
&= 1 - \prod_{i=1}^N (1 - (1 - \exp(-\lambda_N \bar{D}))) \\
&= 1 - e^{-\lambda \bar{D}}.
\end{aligned}$$

4 Proactive, periodic update

Now we turn our attention to the case when updates are sent periodically. In the following we present first the abstracted model of the access strategy, followed by the performance analysis of the case with $N = 1$ information providers, and subsequently the case of $N > 1$ for where we are evaluating the limit cases which allows simpler computations for potential online evaluation of multiple information provider scenarios.

4.1 Abstraction model of periodic update

The requester has made a subscription to the information provider, which will subsequently send updates to the requester in periodic intervals as illustrated in Figure 3.




Fig. 3. Proactive, periodic update: R_k results in mismatch from $E^{(1)}$, while R_{k+1} leads to a correct result for the case $N = 1$, but leads to a mismatch, if $E^{(N)}$ is also needed.

In this scenario the information provider reads out the information element and sends update messages to the request process at the moments described by the stochastic

update process $I^{(n)} = \{I_j^{(n)}, j \in \mathbf{Z}\}, n = 1, \dots, N.$

These N update processes are assumed jointly independent to the event, delay, and request processes. We consider the case when updates contain all information (full updates), and are ordered by the requestor using sequence numbers.

Notation: For this strategy we use in our Kendall-alike notation $E|D|N|RS|US$ the identification 'periodic(I)' for the 'RS' field. Thereby, I is specifying the update time interval process, e.g. $M|M|4|Periodic(M)$ stands for $N = 4$ Poisson event processes, exponentially distributed downstream delays, and proactive periodic access strategies with Poisson periods.

4.2 Access delay and network overhead

Network overhead is for this case entirely determined by the time interval between updates being used, i.e.

$$V_{[G|G|N|Periodic(G)]} = \sum_{i=1}^N \frac{s_d^{(i)}}{\mathbb{E}(I^{(i)})}. \quad (9)$$

The access delay is zero as the information is obtained from a local cache at the requester.

4.3 Mismatch probability for single information provider ($N=1$)

Formally, we consider a system consisting of the marked point processes $\{(I_j, D_j), j \in \mathbf{Z}\}$ and $\{(E_i, E(E_i)), i \in \mathbf{Z}\}$, with $E(E_i)$ denoting the value of the information element at time E_i . Define for any t the number, p_t , to be the index of the last correctly received update before time t

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

Consider an event process which is a stationary renewal process, whose backwards recurrence time has cdf A_E . Then, $I_{p_{R_k}}$ is the last correctly received information before time R_k . Without loss of generality (by stationarity) we assume the request time is 0 and consider the point process of useful updates correctly received before 0. Then we have the following probability of mismatch upon request at time R_k :

$$\text{mmPr}_{[G|G|1|Periodic(G)]} = \mathbb{P}(E(I_{p_{R_k}}) \neq E(R_k))$$

$$\begin{aligned}
&= \mathbb{P}(E(I_{p_0}) \neq E(0)) \\
&= \int_0^\infty \mathbb{P}(\text{no correct update received in } [0, t]) A_E(dt).
\end{aligned}$$

We shall focus on a system which uses a Poisson distributed random update time interval to reflect that updates being transmitted may not happen exactly at deterministic time intervals, but will be exposed to different aspects such as clock drift, operating system scheduling delays, etc., or following our notation the $G|GI|1|periodic(M)|full$ system. For this model we assume i.i.d. downstream delays with cdf F_D , and the updates are assumed to be a stationary Poisson point process with intensity τ . A point originating at time t provides a correct update if the delay time is no longer than t . This happens with probability $1 - F_D(t)$. Henceforth, the useful updates 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

$$\tau(t) = \tau F_D(t). \quad (10)$$

Hence, by the formula for void probabilities for in-homogenous point processes

$$\mathbb{P}(\text{no useful update received in } [0, t]) = \exp \left(- \int_0^t \tau(s) ds \right). \quad (11)$$

Remark that the application of void probabilities is a well-known method in applied probability, see e.g. [7] and [19] for a similar application in resequencing.

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

$$\text{mmPr}_{[G|GI|1|Periodic(M)]} = \int_0^\infty \exp \left(- \int_0^t \tau F_D(s) ds \right) A_E(dt). \quad (12)$$

For matrix-exponential delay distributions, the inner integral in Equation (12) can be simplified to

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

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

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

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 6 for evaluation. 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.

Poisson event process and exponential downstream delays: Under further assumptions on the event and delay process, Equation (12) specializes to

$$\text{mmPr}[M|M|1|Periodic(M)] = \lambda \int_0^\infty \exp \left(\frac{\tau}{\nu} (1 - e^{-\nu t}) \right) \exp(-(\tau + \lambda)t) dt, \quad (13)$$

with the event process being a Poisson process with intensity λ and the downstream delays are i.i.d. exponentially distributed with mean $1/\nu$.

The expression in (13) can alternatively be expressed as (by change of variable t with $e^{-\nu t}$, and using the notation $\phi = \lambda/\nu$; $\psi = \tau/\nu$)

$$\text{mmPr}[M|M|1|Periodic(M)] = \phi e^\psi \frac{\Gamma(\psi + \phi)}{\psi^{\psi + \phi}} F_{\Gamma(\psi + \phi, \psi)}(1), \quad (14)$$

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

Dominating the function $(1 - e^{-\nu t})$ in (13) by νt we get by Lebesgue's Dominated Convergence Theorem the following conclusion from (13). Then for very short downstream delays (large ν 's) the mismatch probability approach a limit of

$$\lim_{\nu \rightarrow \infty} \text{mmPr}[M|M|1|Periodic(M)] = \frac{\lambda}{\lambda + \tau}.$$

4.4 Mismatch probability for multiple information providers ($N > 1$)

Keeping the assumption of events at each information provider being independent of each other, delays and request process, the mismatch probability for accessing multiple information from the requester, is calculated in the same way as for the proactive event-driven strategy, given in Equation (8), but leading to different limits.

Limits for multiple information providers

If we consider a Poisson event process with rate λ and i.i.d. exponentially distributed delays with rate ν , and scale down the individual event processes with rate $\lambda_N = \lambda/N$ we obtain from Equations (8) and (14) (with $\phi = \lambda/\nu$; $\psi = \tau/\nu$)

$$\lim_{N \rightarrow \infty} \text{mmPr}[M|M|N|Periodic(M)] = 1 - \prod_{i=1}^N \left(1 - \frac{\lambda}{N\nu}\right) \frac{\Gamma(\psi + \lambda/(N\nu))}{\psi^{\psi + \lambda/(N\nu)}} \cdot F_{\Gamma(\psi + \lambda/(N\nu), \psi)}(1) \quad (15)$$

$$= 1 - \exp\left(-\phi e^{\psi} \frac{\Gamma(\psi)}{\psi^{\psi}} F_{\Gamma(\psi, \psi)}(1)\right). \quad (16)$$

Hence, in this case, we obtain the above limit for increasing N when scaling down the rate of the individual event process linearly with N .

5 Analytic results for reactive on-demand access

Finally we now present the reactive access strategy, in which the communication is initiated by the requester. First we introduce the abstraction models, whereafter we present the performance models for the cases, $N = 1$, $N > 1$ and the case when information elements may return to previous values.

5.1 Abstracted model of reactive access

For the reactive strategy, it is the request process that sends a request message to the information provider. Upon receiving the request message, the information provider reads out the value of the information element and sends it in a response message. Figure 4 shows the abstracted model for the reactive case.




Fig. 4. 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.

In this scenario, we consider only the case when the full information element is sent in each response. The communication time between the requesting entity and the information provider is described by the stochastically varying upstream and downstream delays. Compared to a reactive strategy, the additional stochastic processes,

the **upstream delay**, $U^{(n)} = \{U_k^{(n)}, k \in \mathbf{Z}\}$, $n = 1, \dots, N$,

are introduced, under the assumption of joint stationarity and joint independence.

Notation: We use 'react(U)' to specify this access strategy in the Kendall-like notation, where U specifies the type of the upstream delay process. For instance, $MJ|ME|1|React(G)$ denotes a reactive access strategy, in which the single event process is a Markov Jump process, the downstream delay distribution is matrix-exponential, and general upstream delay.

5.2 Access delay and network overhead

The network overhead is determined by the request process and the sizes of request and response messages

$$V_{[G|G|N|React(G)]} = \frac{1}{\mathbb{E}(R)} \sum_{i=1}^N (s_u^{(i)} + s_d^{(i)})$$

where, s_u is the size of the upstream request message, and s_d the size of the downstream response message.

The mean access delay is determined by the upstream and downstream delays, i.e.

$$A_{[G|G|N|React(G)]} = \mathbb{E}(\max_{i=1, \dots, N} \{U^{(i)} + D^{(i)}\}).$$

5.3 Mismatch probability using single or multiple information provider ($N \geq 1$)

For the reactive access strategy a mismatch happens if an event occurs while the response is in transit to the requester. This is exactly the same situation as for the proactive, event-driven scenario using full updates (see Section 3.3.1). Hence, the general mismatch probability for the reactive strategy has already been given in Equation (3), which thus also covers all $G|G|1|react(G)$ cases. Note that the upstream delay process does not influence the mmPr due to the stationarity and independence assumptions. In the case that more than one information is needed, i.e. $N > 1$, the requesting entity sends a request to each of the Information providers. The requests are sent out at the same time-instant (multi-cast), but may arrive at the information providers after different upstream delays. As the requester is in need of all information, it has to wait until all responses have arrived to carry out its computation. Assume $t = 0$ is the time of sending out the request. Then the information will be finally processed at the requester at time $\tilde{M}_N = \max(U_1^{(1)} + D_1^{(1)}, \dots, U_1^{(N)} + D_1^{(N)})$, see Figure 4. The mismatch probability is given by:

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

As \tilde{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, which will be addressed in the coming.

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

$$\begin{aligned} & \text{mmPr}_{[G|G|N|React(D)]} \\ &= 1 - \int_0^\infty \mathbb{P}(E_1^{(1)}(t) = E_1^{(1)}(0), \dots, E_1^{(N)}(t) = E_1^{(N)}(0) | M_N = t) F_{M_N}(dt) \\ &= 1 - \int_0^\infty \mathbb{P}(E_1^{(1)} \geq t, \dots, E_1^{(N)} \geq t) F_{M_N}(dt) \end{aligned}$$

$$= 1 - \int_0^\infty \bar{F}_{m_N}(t) F_{M_N}(dt). \quad (17)$$

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

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

which by differentiation of F_{M_N} 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

$$\text{mmPr}[GI|GI|N|React(D)] = 1 - N \int_0^\infty B_E(t)^N F_D(t)^{N-1} f_D(t) dt. \quad (18)$$

Limit behavior for multiple information providers

Now we focus on the limits of the mmPr expressions for the reactive scenario. Here, as it will be shown, under the right assumptions, the computational effort in calculating the mmPr will be reduced, which is in particular interesting for online computation of the mmPr, see e.g. [22,23] for use cases.

Delays are i.i.d. with exponential tails: As opposed to the proactive approaches, the scaling properties of the reactive schemes are different. Assume for instance the delays to be i.i.d. with exponentially decaying tails, i.e. $\bar{F}_D(x) \sim e^{-\nu x}$, as $x \rightarrow \infty$. If we recall that a sequence of random variables X_N is said to converge weakly to a random variable X if $\lim_{N \rightarrow \infty} \mathbb{E}f(X_N) = \mathbb{E}f(X)$, for any bounded and continuous function f , [3, (10)]. Then it has been proven in [5] that

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

converges weakly to a Gumbel distributed random variable with cdf

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

Hence, if we scale down the individual event processes by

$$\lambda_N := \frac{\lambda}{N \log(N)}, \quad (19)$$

we obtain the following limit behavior of the mismatch probability (using Slutsky's Theorem, i.e. if X_N converges weakly to X , $a_N \rightarrow 0$, and $b_N \rightarrow b$, then $X_N a_N + b_N$ converges weakly to b , see e.g. [3] for more details)

$$\begin{aligned} & \lim_{N \rightarrow \infty} \text{mmPr}[M|GI|N|React(D)] \\ &= 1 - \int_0^\infty e^{-NM_n \lambda_N} F_{M_N}(dt) \end{aligned} \quad (20)$$

$$\begin{aligned} &= 1 - \lim_{N \rightarrow \infty} \mathbb{E} \left(\exp \left(\frac{-\lambda}{\log(N)} \left(\left(\frac{M_N - b_N}{a_N} \right) a_N + b_N \right) \right) \right) \\ &= 1 - e^{-\lambda/\nu}. \end{aligned} \quad (21)$$

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

Delays are i.i.d. with polynomial tails: Now, alternatively assume the delays to be i.i.d. with polynomially tails, i.e. for some $\alpha > 0$, $\bar{F}_D(x) \sim x^{-\alpha}$ 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 has been proved in [5] 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, with cdf

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

Hence, if we scale down the individual event processes by

$$\lambda_N := \frac{\lambda}{N^{1+1/\alpha}}, \quad (22)$$

and assume X has the Frechet distribution with cdf F , then we obtain the following limit behavior of the mismatch probability (by use of a weak convergence result similar to the one leading to Equation (21))

$$\begin{aligned} \lim_{N \rightarrow \infty} \text{mmPr}[M|GI|N|React(D)] &= 1 - \mathbb{E}(e^{-\lambda X}) \\ &= 1 - \alpha \int_0^\infty e^{-(\lambda x + x^{-\alpha})} x^{-\alpha-1} dx. \end{aligned} \quad (23)$$

5.4 Recurrent Event Processes with one information provider ($N=1$)

Now we remove the assumption that information element cannot change back to a previous value. Instead, one single information element is assumed to be described by the state of a Markov process with generator matrix \mathbf{Q} .

In the reactive approach, the access leads to a mismatch, if after the downstream delay time, the Markov process, represented by a state transition rate matrix, \mathbf{Q} , and state probability vector, $\boldsymbol{\pi}$, 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

$$\text{mmPr}[MP(ree)|GI|1|React(G)] = 1 - \int_{t=0}^{\infty} \sum_{i=1}^S \pi_i [\exp(\mathbf{Q}t)]_{i,i} f_D(t) dt \quad (24)$$

General matrix-exponential downstream delay: If the downstream delay is a matrix-exponential renewal process, integral-free expressions for Eq. (24) can be obtained as follows:

$$\begin{aligned} \text{mmPr}[MP|ME|1|React(G)] &= 1 - \int_{t=0}^{\infty} \sum_{i=1}^S \pi_i [\exp(\mathbf{Q}t)]_{i,i} \mathbf{p}_D \mathbf{B}_D \exp(-\mathbf{B}_D t) \boldsymbol{\epsilon}'_D dt \\ &= 1 - \sum_{i=1}^S \pi_i \int_{t=0}^{\infty} \mathbf{e}_i \exp(\mathbf{Q}t) \mathbf{e}_i' \mathbf{p}_D \mathbf{B}_D \exp(-\mathbf{B}_D t) \boldsymbol{\epsilon}'_D dt \\ &= 1 - \sum_{i=1}^S \pi_i [\mathbf{e}_i \otimes (\mathbf{p}_D \mathbf{B}_D)] [(-\mathbf{Q}) \oplus \mathbf{B}_D]^{-1} [\mathbf{e}_i' \otimes \boldsymbol{\epsilon}'_D]. \end{aligned} \quad (25)$$

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

Exponential downstream delay: Further simplifications result if the downstream delay is exponentially distributed with rate ν . Equation (24) can be reduced to

$$\text{mmPr}[MP|M|1|React(G)] = 1 - \nu \sum_{i=1}^S \pi_i \left[\int_{t=0}^{\infty} \exp([\mathbf{Q} - \nu \mathbf{I}]t) dt \right]_{i,i}$$

$$= 1 - \nu \sum_{i=1}^S \pi_i \left[(\mathbf{Q} - \nu \mathbf{I})^{-1} \right]_{i,i}.$$

6 Numerical results

This section uses the models from Sections 3, 4, and 5 to obtain and discuss numerical results for selected example scenarios.

6.1 Single information provider, non-recurrent event process

The analytic models derived allow to compute the mmPr for the scenario of $N = 1$ information providers, at which the information element never changes back to a previous value.

We consider in the following numerically the case of a single information provider, with different distributions of the inter-event process and downstream delay process.

6.1.1 Exponential network delays and Poisson event process

The first scenario 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 5 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 i.i.d. exponentially distributed with varying rate $\tau = 10^{-2}, 0.1, 1, 10$. As the analysis in Section 3.3.1 shows, the reactive and the proactive event-driven strategy with full updates lead to exactly the same mmPr (dotted curve). The proactive event-driven strategy with incremental updates shows a slightly higher mmPr (solid curve). According to the analysis in Section 3.3.2, the 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 3.3.3, 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




Fig. 5. Comparison of mismatch probabilities in the different remote access strategies for a Poisson event process and varying downstream delay [17]. Notice the topmost curve for the periodic case with $\tau=0.01$.

cases, $M|M|1|periodic(M)$, with different rate τ of the period. The following additional observations can be made from Figure 5:

- 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 in Section 6.1.2 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 to zero, see Section 3.3. This is explainable with the arguments that no message reordering will occur for infinitely fast networks.
- For the case of the periodic update strategy, it was shown in Section 5.3 that for small downstream delays (large ν) the mmPr converges to the value $\lambda/(\lambda + \tau) > 0$. Consequently, for large ν eventually, the periodic approach will at some point always perform worse than the event-driven and reactive approaches.

Figure 5 shows only the mmPr. One approach to make a fair comparison between the respective strategies and their resulting mmPr in Figure 5, is to compare scenarios in which the network overhead and access delay are identical. However, for the reactive strategy the access delay is always larger than zero, hence is always larger than for the proactive strategies. A caching strategy may be applied at the cost of an increased mmPr in order to decrease

the mean access delay, see [25]. For the network overhead, each strategy should produce same amount of traffic, which happens if $\mu(s_d + s_u) = \lambda s_d = \tau s_d$, which for the scenario shown in Figure 5 happens i.e. for $\tau = 1$ and $\mu = 0.5$ (when assuming equal sizes of the update messages).

6.1.2 Matrix-Exponential network delays

To investigate numerically the impact of different distribution types, we use in the following the Erlangian distributions to represent delay distributions with smaller variance than exponential. At the other end for large variance, we use Truncated Power-Tail (TPT) distributions, see Appendix A, with a tail-exponent of $\alpha = 1.4$ which leads to unboundedly growing variance for increasing number of phases.

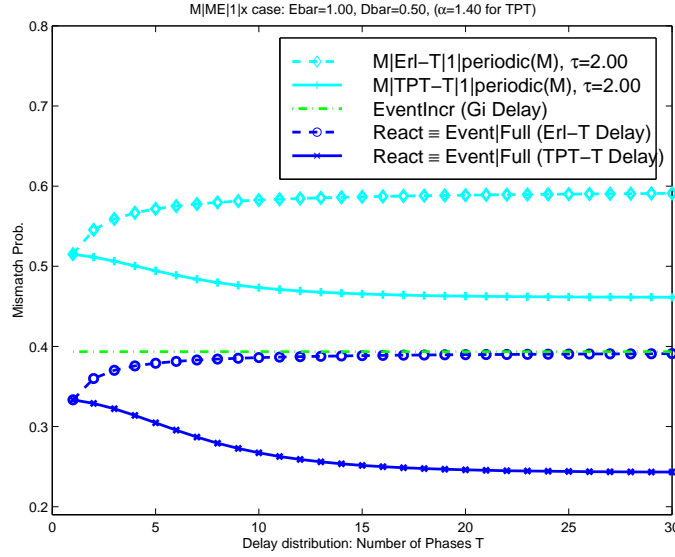


Fig. 6. Mismatch Probability for reactive and the two different proactive strategies for downstream delay process which are renewal processes with matrix-exponential representation: Shown for Erlangian and TPT distributions with increasing number of phases along the x -axis.

Figure 6 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 downstream 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. This decrease in variance actually results in an in-

creased mmPr. The lower curve (marked with 'x') shows the mmPr for a TPT- T distributed network delay. The mmPr decays with increasing T but appears to converge to a value slightly below 25%.

The upper two curves in Figure 6 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, which can also be confirmed by using a deterministic distribution in Equation (12).

The fact that the mmPr increases for deterministic delays, while decreasing for highly varying delays, can be explained in the following way: Consider an event at some point in time which would lead to a mismatch for a deterministic delay, then for the same event there is a probability that the information is matching if it stochastic. For a delay equal, or longer than the deterministic, obviously, there will be a mismatch, but for a smaller delay it may not necessarily be a mismatch. Over time, the randomness leads to a reduction in the mmPr as shown in Figure 6.

6.1.3 Matrix-exponential network delays and event process




Fig. 7. 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.

Similar qualitative behavior as in the previous section is observed when the

distribution of the inter-event times is varied in Figure 7. The change from exponential towards a deterministic distribution (Erlangian with many phases) 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 ($.|D|1|(rect = event|inc = full)$) thereby represent the case of deterministic downstream delays, in which four of the strategies are actually equivalent, see Section 3.3.3.




Fig. 8. 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.

When both processes, the event and the downstream delay, are represented by TPT respectively Erlangian distributions, the impact on the mmPr is strongest, i.e. the worst case in the considered candidate set occurs, when both distributions are Erlangian with large number of phases, i.e. close to the $D|D|1|x$ case, see Figure 8. This case is interesting in the sense that, some combinations of delay and event times may lead to situations where information is always wrong $mmPr = 1$ which happens for example if the delay is larger than the event inter-arrival times, i.e. when we are certain at least one event happens in the downstream period.

6.2 Multiple information providers, non-recurrent event process

For scenarios involving more than one information element, the information 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. Those scenarios were investigated in Sections 3.4, 4.4 and 5.3, for the

proactive event-driven, periodic and reactive access strategy, respectively.

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 (8). 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 (18) allows to compute the mmPr, in most cases involving numerical integration.




Fig. 9. 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.

We show in Figure 9 the numerical results for multiple information providers, $N > 1$, where we restrict ourselves to exponential distributions, i.e. $M|M|N|x$ cases. In the same time we also show the derived limits of the different cases, to show the empirical accuracy of the approximations increased at the benefit of lower complexity of calculating the 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 $N \rightarrow \infty$. For the proactive cases, the difference in using the limits and the exact expressions becomes very small as N increases, giving an opportunity for saving computational effort in e.g. realtime calculations of mmPr in situations with many information sources. Similar with the reactive case, but the effect is less profound as seen.

6.3 Single information provider, recurrent event process

The numerical results so far assumed information elements that can never change back to a previously taken value. Now we remove this assumption, so that the information element can change back to previous values, as described by a continuous time Markov process. We focus our numerical analysis on the reactive case, as described in Section 5.4. For this evaluation 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.




Fig. 10. Mismatch Probability for reactive strategy for an event process which is an ON/OFF process with same average duration of the ON+OFF cycle.

Figure 10 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 MP process (solid) which shows a lower mmPr than a 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 10. Notice that for the ON/OFF element used here, we do not care whether the state is ON or OFF, but focus on whether we get mismatching information or not. This state symmetry leads to $\text{mmPr}(\kappa) = \text{mmPr}(1/\kappa)$.

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 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 greater 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 (12) has been obtained for the scenario, when the instances of sending updates form a Poisson process. Integral-free solutions (14) 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 (25).
- The case of multiple information providers is treated for all strategies and limit results are obtained that allow to identify interesting differences in scaling behavior as well as allow to obtain computationally efficient approximation 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 case of multiple information elements has been analyzed 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.

Finally, the analysis was applied to the case of a binary information element, which toggles between two 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.

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. Finally, the application of the mmPr analysis to the actual use-cases of routing, context-sensitive networking, and replica 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 also like to thank the consortia members for the fruitful discussions. Furthermore, The Telecommunications Research Center Vienna (ftw.) is supported by the Austrian Government and by the City of Vienna within the competence center program COMET. Parts of the work of Martin Bgsted Hansen was carried out at the Department of Mathematical Sciences, Aalborg University.

References

- [1] R.Skehill, M.Barry, W.Kent, M.Ocallaghan, N.Gawley, S.McGrath, The common RRM approach to admission control, *Wireless Communications Magazine*, Volume 14, Issue 2, April 2007 Page(s):48 - 56
- [2] S. Asmussen. *Applied Probability and Queues*. Springer Verlag, New York, NY, 2 edition, 2003.
- [3] P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, NY, 1968.
- [4] M. Bozinovski. *Fault-tolerant Platforms for IP-based Session Control Systems*. PhD thesis, Aalborg University, 2004.
- [5] 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.
- [6] W. Chen, S. Toueg, and M. Kawazoe Aguilera. On the quality of service of failure detectors. *IEEE Transactions on Computers*, 51(5), May 2002.
- [7] D.J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes: Elementary Theory and Methods*. Springer, NY, 2002.
- [8] M. Ghader, R. L. Olsen, M.G. Genet, and R. Tafazolli. Service management platform for personal networks. *Proceedings of IST Summit 2005, Dresden*, 2005.
- [9] A. Graham. *Kronecker products and matrix calculus with applications*. Ellis Horwood, 1981.
- [10] M. Greiner, M. Jobmann, and L. Lipsky. The importance of power-tail distributions for telecommunication traffic models. *Operations Research*, pages 313 - 326, 1999.
- [11] A.A. Helal, A.A. Heddaya, and B.B. Bhargawa. *Replication Techniques in Distributed Systems*. Kluwer Academic Publishers, 1996.
- [12] 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 04, Hong Kong*, 2004.
- [13] L. Lipsky. *Queueing Theory: A Linear Algebraic Approach*. McMillan and Company, New York, 1992.
- [14] A. Seddik-Ghaleb, Y.Ghamri-Doudane, S.-M.Senouci, Effect of Ad Hoc Routing Protocols on TCP Performance within MANETs , *Sensor and Ad Hoc Communications and Networks*, 2006. SECON '06. 2006 3rd Annual IEEE Communications Society on Volume 3, 28-28 Sept. 2006 Page(s):866 - 873
- [15] N. Fujinami, Y.Yokote, Naming and addressing of objects without unique identifiers, *Distributed Computing Systems*, 1992., *Proceedings of the 12th International Conference on 9-12 June 1992* Page(s):581 - 588

- [16] M. F. Neuts. Matrix-Geometric Solutions in Stochastic Models. John Hopkins, University Press, Baltimore, 1981.
- [17] R.L. Olsen, H.P. Schwefel, and M.B. Hansen. Quantitative analysis of access strategies to remote information in network services. In proceedings of Globecom 2006, San Fransisco, California, USA, 2006, 2006.
- [18] Forest Baskett, K. Mani Chandy, Richard R. Muntz, and Fernando G. Palacios. Open, closed, and mixed networks of queues with different classes of customers. Journal of the ACM, 22(2):248-260, April 1975
- [19] G. Harrus and B. Plateau. Queueing analysis of a reordering issue. IEEE Transactions on Software Engineering, 8(2):113-123, 1982.
- [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. Matthiasen, H.P. Schwefel, and R. Prasad. Inconsistency evaluation in a replicated IP-based call control center. In proceedings of ISAS06, Helsinki, Finland, May 2006.
- [22] R.L. Olsen and H.P. Schwefel. Determination of context value in multiple context source scenarios for context management systems. In Proceedings of Wireless Personal Multimedia Communication 07, Jaipur, India, 2007.
- [23] R.L. Olsen, H.P. Schwefel, and M.Bauer. Influence of unreliable information on context aware service discovery. In Proceedings of Third Workshop on Context Awareness for Proactive Systems, University of Surrey, Guildford, England, 2007.
- [24] 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.
- [25] H.P. Schwefel, M.B. Hansen, and R.L. Olsen. Adaptive caching strategies for context management systems. In In Proceedings of PIMRC07, Athens, Greece, 2007.

A 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,

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

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 [10]. The higher the number of phases, T , the later the drop-off occurs. The exponential drop-off is characterized in more detail in [24] 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.1) have to be (see [10]):

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

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

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

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

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

B Shifted delay distribution

In certain scenarios, time shifted delays may need to be considered. For this type of scenario, the stochastic delay variable D_v , is offset by a constant delay d_0 , leading to a total delay D given by

$$D = d_0 + D_v$$

whereby

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

Utilizing such shifted delay distribution, e.g., for the cases $ME|ME|1|event|full$ and $ME|ME|1|react(G)$, Equation (3) yields a mismatch probability of:

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

which in the matrix-exponential case becomes

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