



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

LINEAR ALGEBRA
AND ITS
APPLICATIONS

Linear Algebra and its Applications 386 (2004) 311–334

www.elsevier.com/locate/laa

Hypergraph-based parallel computation of passage time densities in large semi-Markov models

Jeremy T. Bradley ^{a,*}, Nicholas J. Dingle ^a,
William J. Knottenbelt ^a, Helen J. Wilson ^b

^aDepartment of Computing, Imperial College London, South Kensington Campus, London SW7 2AZ, UK

^bDepartment of Applied Mathematics, University of Leeds, Woodhouse Lane, Leeds LS2 9JT, UK

Received 12 August 2003; accepted 3 December 2003

Submitted by D. Szyld

Abstract

Passage time densities and quantiles are important performance and quality of service metrics, but their numerical derivation is, in general, computationally expensive. We present an iterative algorithm for the calculation of passage time densities in semi-Markov models, along with a theoretical analysis and empirical measurement of its convergence behaviour. In order to implement the algorithm efficiently in parallel, we use hypergraph partitioning to minimise communication between processors and to balance workloads. This enables the analysis of models with very large state spaces which could not be held within the memory of a single machine. We produce passage time densities and quantiles for very large semi-Markov models with over 15 million states and validate the results against simulation.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Semi-Markov chains; Passage time densities; Hypergraph partitioning

1. Introduction

A rapid response time is an important performance criterion for almost all computer communication and transaction processing systems. Examples of systems with

* Corresponding author.

E-mail addresses: jb@doc.ic.ac.uk (J.T. Bradley), njd200@doc.ic.ac.uk (N.J. Dingle), wjk@doc.ic.ac.uk (W.J. Knottenbelt), h.j.wilson@leeds.ac.uk (H.J. Wilson).

stringent response time requirements include mobile communication systems, stock market trading systems, web servers, database servers, flexible manufacturing systems, communication protocols and communication networks. Typically, response time targets are specified in terms of quantiles—for example “95% of all text messages must be delivered within 3 s”. Response time quantiles are frequently specified as key quality of service metrics in service level agreements and industry standard benchmarks such as TPC.

In the context of the high-level models used by performance analysts (e.g. queuing networks and stochastic Petri nets), response times can be specified as *passage* times—that is, the time taken to enter any one of a set of target states starting from a specified set of source states in the underlying Markov or semi-Markov chain. Traditional steady-state performance analysis of such models is adequate to predict standard resource-based measures such as utilisation and throughput, but is inadequate to predict passage time densities and quantiles.

The focus of the present study is on semi-Markov processes [1], a generalisation of Markov processes which support arbitrary state holding times. Techniques already exist for the practical extraction of passage time densities and quantiles from large purely Markovian systems [2–4]. Until now, however, only relatively small semi-Markov systems have been analysed for passage time quantities [5–7].

In this paper, we present a scalable iterative passage time density extraction algorithm for very large semi-Markov processes (SMPs). This extends our preliminary work on passage time extraction [8] with significant theoretical and empirical convergence results for the central iterative algorithm. Our approach is based on the calculation and subsequent numerical inversion of the Laplace transform [9,10] of the desired passage time density. In [6,7], the time complexity of the numerical derivation of passage time and transient quantities for a semi-Markov system with N states, using the Laplace domain, is $O(N^4)$. This approach is dominated by the complexity of maintaining the Laplace transforms of state holding time distributions in closed form. We solve this problem by characterising a distribution by the samples from its Laplace transform that are ultimately required for the inversion process. The result is an iterative algorithm which calculates arbitrary semi-Markov passage times in $O(N^2r)$ time, for r iterations.

The data partitioning strategy employed is key to the scalability (in terms of both efficiency and capacity) of all parallel algorithms. We exploit recent advances in the application of hypergraph data structures and their partitioning [11] to minimise inter-processor communication while balancing computational load.

The rest of this paper is organised as follows: Section 2 summarises the theory relating to the computation of passage time densities in semi-Markov processes. Our iterative passage time algorithm and its theoretical convergence analysis are detailed in Section 3. Section 4 addresses the problem of general distribution representation, while Section 5 discusses the application of hypergraph partitioning techniques. Section 6 presents a complete parallel passage time analysis pipeline. Section 7 presents

numerical results showing passage time densities extracted from two very large semi-Markov models, as well as scalability and convergence results.

2. Definitions and background theory

2.1. Semi-Markov processes

Consider a Markov renewal process $\{(X_n, T_n) : n \geq 0\}$ where T_n is the time of the n th transition ($T_0 = 0$) and $X_n \in \mathcal{S}$ is the state at the n th transition. Let the kernel of this process be

$$R(n, i, j, t) = \mathbb{P}(X_{n+1} = j, T_{n+1} - T_n \leq t \mid X_n = i) \tag{1}$$

for $i, j \in \mathcal{S}$. The continuous time semi-Markov process (SMP), $\{Z(t), t \geq 0\}$, defined by the kernel R , is related to the Markov renewal process by

$$Z(t) = X_{N(t)}, \tag{2}$$

where $N(t) = \max\{n : T_n \leq t\}$, i.e. the number of state transitions that have taken place by time t . Thus $Z(t)$ represents the state of the system at time t . We consider time-homogeneous SMPs, in which $R(n, i, j, t)$ is independent of any previous state except the last. Thus R becomes independent of n :

$$\begin{aligned} R(i, j, t) &= \mathbb{P}(X_{n+1} = j, T_{n+1} - T_n \leq t \mid X_n = i) \quad \text{for any } n \geq 0 \\ &= p_{ij} H_{ij}(t), \end{aligned} \tag{3}$$

where $p_{ij} = \mathbb{P}(X_{n+1} = j \mid X_n = i)$ is the state transition probability between states i and j and $H_{ij}(t) = \mathbb{P}(T_{n+1} - T_n \leq t \mid X_{n+1} = j, X_n = i)$, is the sojourn time distribution in state i when the next state is j .

2.2. First passage times

Consider a finite, irreducible, continuous-time semi-Markov process with N states $\{1, 2, \dots, N\}$. Recalling that $Z(t)$ denotes the state of the SMP at time t ($t \geq 0$), the first passage time from a source state i at time t into a non-empty set of target states \vec{j} is

$$P_{i\vec{j}}(t) = \inf\{u > 0 : Z(t+u) \in \vec{j}, N(t+u) > N(t), Z(t) = i\}. \tag{4}$$

For a stationary time-homogeneous SMP, $P_{i\vec{j}}(t)$ is independent of t and we have

$$P_{i\vec{j}} = \inf\{u > 0 : Z(u) \in \vec{j}, N(u) > 0, Z(0) = i\}. \tag{5}$$

This formulation of the random variable $P_{i\vec{j}}$ applies to an SMP with no immediate (that is, zero-time) transitions. If zero-time transitions are permitted in the model then the passage time can be stated as

$$P_{i\vec{j}} = \inf \{u > 0 : N(u) \geq M_{i\vec{j}}\}, \tag{6}$$

where $M_{i\vec{j}} = \min\{m \in \mathbb{Z}^+ : X_m \in \vec{j} \mid X_0 = i\}$ is the transition marking the terminating state of the passage.

$P_{i\vec{j}}$ has an associated probability density function $f_{i\vec{j}}(t)$ such that the passage time quantile is given as

$$\mathbb{P}(t_1 < P_{i\vec{j}} < t_2) = \int_{t_1}^{t_2} f_{i\vec{j}}(t) dt. \tag{7}$$

In general, the Laplace transform of $f_{i\vec{j}}$, $L_{i\vec{j}}(s)$, can be computed by solving a set of N linear equations:

$$L_{i\vec{j}}(s) = \sum_{k \notin \vec{j}} r_{ik}^*(s)L_{k\vec{j}}(s) + \sum_{k \in \vec{j}} r_{ik}^*(s) \quad \text{for } 1 \leq i \leq N, \tag{8}$$

where $r_{ik}^*(s)$ is the Laplace–Stieltjes transform (LST) of $R(i, k, t)$ from Section 2.1 and is defined by

$$r_{ik}^*(s) = \int_0^\infty e^{-st} dR(i, k, t). \tag{9}$$

Eq. (8) has a matrix–vector form, $A\vec{x} = \vec{b}$, where the elements of A are general complex functions; care needs to be taken when storing such functions for eventual numerical inversion (see Section 4). For example, when $\vec{j} = \{1\}$, Eq. (8) yields

$$\begin{pmatrix} 1 & -r_{12}^*(s) & \cdots & -r_{1N}^*(s) \\ 0 & 1 - r_{22}^*(s) & \cdots & -r_{2N}^*(s) \\ 0 & -r_{32}^*(s) & \cdots & -r_{3N}^*(s) \\ \vdots & \vdots & \ddots & \vdots \\ 0 & -r_{N2}^*(s) & \cdots & 1 - r_{NN}^*(s) \end{pmatrix} \begin{pmatrix} L_{1\vec{j}}(s) \\ L_{2\vec{j}}(s) \\ L_{3\vec{j}}(s) \\ \vdots \\ L_{N\vec{j}}(s) \end{pmatrix} = \begin{pmatrix} r_{11}^*(s) \\ r_{21}^*(s) \\ r_{31}^*(s) \\ \vdots \\ r_{N1}^*(s) \end{pmatrix}. \tag{10}$$

When there are multiple source states, denoted by the vector \vec{i} , the Laplace transform of the passage time density at steady-state is

$$L_{\vec{i}\vec{j}}(s) = \sum_{k \in \vec{i}} \alpha_k L_{k\vec{j}}(s), \tag{11}$$

where the weight α_k is the probability at equilibrium that the system is in state $k \in \vec{i}$ at the starting instant of the passage. If $\vec{\pi}$ denotes the steady-state vector of the embedded discrete-time Markov chain (DTMC) with one-step transition probability matrix $P = [p_{ij}, 1 \leq i, j \leq N]$, then α_k is given by

$$\alpha_k = \begin{cases} \pi_k / \sum_{j \in \vec{i}} \pi_j & \text{if } k \in \vec{i}, \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

The vector with components α_k is denoted by $\vec{\alpha}$.

3. Iterative passage time analysis

3.1. Overview

In this section, we describe an algorithm for generating passage time densities that creates successively better approximations to the SMP passage time quantity of Eq. (8). Our technique considers the r th transition passage time of the system, $P_{i\vec{j}}^{(r)}$. This is the conditional passage time of the system having reached any of the specified target states within r state-transitions. The unconditioned passage time random variable, $P_{i\vec{j}}$, is then obtained in the limit as $r \rightarrow \infty$. We calculate the Laplace transform of $P_{i\vec{j}}^{(r)}$, $L_{i\vec{j}}^{(r)}(s)$, and pick a sufficiently high value of r to give an approximation to $L_{i\vec{j}}(s)$ to within a specified degree of accuracy. $L_{i\vec{j}}(s)$ can then be numerically inverted to obtain the desired passage time density $f_{i\vec{j}}(t)$.

This iterative method bears a loose resemblance to the well-known *uniformisation* technique [12–14] which can be used to generate transient-state distributions and passage time densities for Markov chains. However, as we are working with semi-Markov systems, there can be no *uniformising* of the general distributions in the SMP. The general distribution information has to be maintained as precisely as possible throughout the process, which we achieve using the representation technique described in Section 4.

3.2. Technical description

Recall the semi-Markov process $Z(t)$ of Section 2.1, where $N(t)$ is the number of state transitions that have taken place by time t . We formally define the r th transition first passage time to be

$$P_{i\vec{j}}^{(r)} = \inf \{ u > 0 : Z(u) \in \vec{j}, 0 < N(u) \leq r, Z(0) = i \}, \tag{13}$$

which is the time taken to enter a state in \vec{j} for the first time having started in state i at time 0 and having undergone up to r state transitions.¹ $P_{i\vec{j}}^{(r)}$ is a random variable with associated Laplace transform $L_{i\vec{j}}^{(r)}(s)$. $L_{i\vec{j}}^{(r)}(s)$ is, in turn, the i th component of the vector

$$\tilde{L}_{\vec{j}}^{(r)}(s) = (L_{1\vec{j}}^{(r)}(s), L_{2\vec{j}}^{(r)}(s), \dots, L_{N\vec{j}}^{(r)}(s)), \tag{14}$$

representing the passage time for terminating in \vec{j} for each possible start state. This vector may be computed as

$$\tilde{L}_{\vec{j}}^{(r)}(s) = U(I + U' + U'^2 + \dots + U'^{(r-1)})\tilde{e}_{\vec{j}}, \tag{15}$$

¹ If we have immediate transitions in our SMP model (as in Eq. (6)) then the r th transition first passage time is $P_{i\vec{j}}^{(r)} = \inf\{u > 0 : M_{i\vec{j}} \leq N(u) \leq r\}$.

where U is a matrix with elements $u_{pq} = r_{pq}^*(s)$ and U' is a modified version of U with elements $u'_{pq} = \delta_{p \notin \vec{j}} u_{pq}$, where states in \vec{j} have been made absorbing. We include the initial multiplication with U in Eq. (15), so as to generate cycle times for cases such as $L_{ii}^{(r)}(s)$ which would otherwise register as 0 if U' were used instead. The column vector $\tilde{e}_{\vec{j}}$ has entries $e_{k\vec{j}} = \delta_{k \in \vec{j}}$.

From Eqs. (5) and (13),

$$P_{i\vec{j}} = P_{i\vec{j}}^{(\infty)} \quad \text{and thus} \quad L_{i\vec{j}}(s) = L_{i\vec{j}}^{(\infty)}(s). \quad (16)$$

We can generalise to multiple source states \vec{i} using the vector $\tilde{\alpha}$ of Eq. (12):

$$\begin{aligned} L_{\vec{i}\vec{j}}^{(r)}(s) &= \tilde{\alpha} \tilde{L}_{\vec{j}}^{(r)}(s) \\ &= \sum_{k=0}^{r-1} \tilde{\alpha} U U'^k \tilde{e}_{\vec{j}}. \end{aligned} \quad (17)$$

The sum of Eq. (17) can be computed efficiently using sparse matrix–vector multiplications with a vector accumulator, $\tilde{\mu}_r = \sum_{k=0}^{r-1} \tilde{\alpha} U U'^k$. At each step, the accumulator (initialised as $\tilde{\mu}_0 = \tilde{\alpha} U$) is updated as $\tilde{\mu}_{r+1} = \tilde{\alpha} U + \tilde{\mu}_r U'$. The worst-case time complexity for this sum is $O(N^2 r)$ versus the $O(N^3)$ of typical matrix inversion techniques. In practice, for a sparse matrix with constant number of non-zeros per row, this can be as low as $O(Nr)$.

3.3. Analytic convergence and truncation error analysis

In the iterative passage time calculation of Eq. (17), we approximate the Laplace transform of the passage time density $L_{\vec{i}\vec{j}}(s)$ by a sum of the form $\sum_{i=0}^k \tilde{x} A^i$. In this section, we demonstrate that this sum converges onto the analytic solution of Eq. (8); also, by using the Power method to provide an expression for the dominant eigenvalue, we derive an approximation for the truncation error after k iterations.

3.3.1. Limit and convergence

Let us consider the finite sum $\sum_{i=0}^k \tilde{x} A^i$ where A has dominant eigenvalue λ_1 . If we post-multiply this expression by $(I - A)$, we obtain

$$\sum_{i=0}^k \tilde{x} A^i (I - A) = \tilde{x} - \tilde{x} A^{k+1}, \quad (18)$$

and so, if the modulus of the dominant eigenvalue of A , $|\lambda_1| < 1$, then $\lim_{k \rightarrow \infty} \tilde{x} A^{k+1} = 0$, and we obtain

$$\sum_{i=0}^{\infty} \tilde{x} A^i (I - A) = \tilde{x} \quad \text{and thus} \quad \sum_{i=0}^{\infty} \tilde{x} A^i = \tilde{x} (I - A)^{-1}. \quad (19)$$

We can relate the above analysis to the iterative passage time algorithm by taking $\tilde{x} = \tilde{\alpha}U$ and $A = U'$. We know empirically that this sum will converge because the sum of Eq. (15) converges as $r \rightarrow \infty$ according to Eq. (16). A formal proof that our iterative scheme will converge in the limit to the analytic solution of Eq. (8) can be found in [15]. It is also shown in [15] that U' has dominant eigenvalue $|\lambda_1| \leq 1$ but the iterations converge even if $|\lambda_1| = 1$.

3.3.2. Truncation error analysis

Let A have eigenvalues λ_j and associated eigenvectors \tilde{v}_j , with dominant eigenvalue λ_1 as above. Assuming non-degenerate eigenvalues, the \tilde{v}_j form a basis and so for an arbitrary vector \tilde{x} :

$$\begin{aligned} \tilde{x} &= \sum_j a_j \tilde{v}_j, \\ \tilde{x}A^i &= \sum_j a_j \lambda_j^i \tilde{v}_j \\ &= a_1 \lambda_1^i \tilde{v}_1 + \sum_{j>1} a_j \lambda_j^i \tilde{v}_j, \\ \tilde{x}A^k &\sim a_1 \lambda_1^k \tilde{v}_1 \quad \text{as } k \rightarrow \infty. \end{aligned} \tag{20}$$

Thus the quantity $\tilde{x}A^k$ is dominated by the largest eigenvalue, with the rate of asymptotic convergence being governed by the ratio of the moduli of the dominant and sub-dominant eigenvalues, $|\lambda_1|/|\lambda_2|$. Eq. (20) yields $\tilde{x}A^k \simeq \lambda_1 \tilde{x}A^{k-1}$ for large k ; so, by right-multiplying both sides by $(\tilde{x}A^k)^*$, we may approximate λ_1 with

$$\lambda_1 \simeq \frac{|\tilde{x}A^k|^2}{(\tilde{x}A^{k-1}) \cdot (\tilde{x}A^k)^*}, \tag{21}$$

where for $\tilde{\omega}$, a vector with complex elements ω_i , the complex conjugate $\tilde{\omega}^*$ has elements ω_i^* .

The error incurred in truncating the sum at the k th term is $\sum_{i=k+1}^{\infty} \tilde{x}A^i$. Using the approximation $\tilde{x}A^i \simeq a_1 \lambda_1^i \tilde{v}_1$ for $i \geq k$ yields:

$$\begin{aligned} \sum_{i=k+1}^{\infty} \tilde{x}A^i &= \sum_{i=1}^{\infty} \tilde{x}A^{k+i} \\ &\simeq \sum_{i=1}^{\infty} \lambda_1^i \tilde{x}A^k \\ &= \frac{\lambda_1}{1 - \lambda_1} \tilde{x}A^k \end{aligned} \tag{22}$$

providing $|\lambda_1| < 1$ as in Section 3.3.1. Thus, in the case where $|\lambda_1| < 1$, an approximate convergence condition is to find the minimum k such that

$$\left| \frac{\lambda_1}{1 - \lambda_1} \tilde{x} A^k \right| < \epsilon, \quad (23)$$

where λ_1 is approximated as in Eq. (21).

3.4. Practical convergence

In practice, convergence of the sum $L_{ij}^{(r)}(s) = \sum_{k=0}^{r-1} \tilde{\alpha} U U^k$ can be said to have occurred if, for a particular r and s -point,

$$\left| \operatorname{Re}(L_{ij}^{(r+1)}(s) - L_{ij}^{(r)}(s)) \right| < \epsilon \quad \text{and} \quad \left| \operatorname{Im}(L_{ij}^{(r+1)}(s) - L_{ij}^{(r)}(s)) \right| < \epsilon, \quad (24)$$

where ϵ is chosen to be a suitably small value, say $\epsilon = 10^{-16}$. We present empirical observations on the convergence behaviour of this technique (i.e. the order of r) in Section 7.

An optimisation which makes the sum converge more quickly (assuming the approximation of Eq. (21) for λ_1 is a good one), is to use the approximate truncation error to improve the accuracy of the calculation. For each iteration we can take

$$L_{ij}^{(r)}(s) = \sum_{k=0}^{r-1} \tilde{\alpha} U U^k + \frac{\lambda_1}{1 - \lambda_1} \tilde{\alpha} U U^{(r-1)}. \quad (25)$$

We can then compare successive $L_{ij}^{(r)}(s)$ in the real and imaginary parts as before. If $|\lambda_1| = |\lambda_i| < 1$ for one or more values of $i \neq 1$ then convergence of $L_{ij}^{(r)}(s)$ will still occur, but the error dynamics are more complicated and we are forced to resort to the stricter notion of convergence above.

4. Distribution representation and Laplace inversion

The key to practical analysis of semi-Markov processes lies in the efficient representation of their general distributions. Without care the structural complexity of the SMP can be recreated within the representation of the distribution functions. This is especially true with the manipulations performed in the iterative passage time calculation of Section 3.

Many techniques have been used for representing arbitrary distributions—two of the most popular being *phase-type distributions* and *vector-of-moments* methods. These methods suffer from, respectively, exploding representation size under composition, and containing insufficient information to produce accurate answers after large amounts of composition.

As all our distribution manipulations take place in Laplace-space, we link our distribution representation to the Laplace inversion technique that we ultimately use. Our tool supports two Laplace transform inversion algorithms, which are briefly out-

lined below: the Euler technique [16] and the Laguerre method [9] with modifications summarised in [3].

Both algorithms work on the same general principle of sampling the transform function $L(s)$ at n points, s_1, s_2, \dots, s_n and generating values of $f(t)$ at m user-specified t -points t_1, t_2, \dots, t_m . In the Euler inversion case $n = km$, where k can vary between 15 and 50, depending on the accuracy of the inversion required. In the modified Laguerre case, $n = 400$ and, crucially, is independent of m (see Section 4.2).

The process of selecting a Laplace transform inversion algorithm is discussed later; however, whichever is chosen, it is important to note that calculating $s_i, 1 \leq i \leq n$ and storing all our distribution transform functions, sampled at these points, will be sufficient to provide a complete inversion. Key to this is the fact that matrix element operations, of the type performed in Eq. (17), (i.e. convolution and weighted sum) do not require any adjustment to the array of domain s -points required. In the case of a convolution, for instance, if $L_1(s)$ and $L_2(s)$ are stored in the form $\{(s_i, L_j(s_i)) : 1 \leq i \leq n\}$, for $j = 1, 2$, then the convolution, $L_1(s)L_2(s)$, can be stored using the same size array and using the same list of domain s -values, $\{(s_i, L_1(s_i)L_2(s_i)) : 1 \leq i \leq n\}$.

Storing our distribution functions in this way has three main advantages. Firstly, the function has constant storage space, independent of the distribution-type. Secondly, each distribution has, therefore, the same constant storage requirement even after composition with other distributions. Finally, the function has sufficient information about a distribution to determine the required passage time (and no more).

4.1. Summary of Euler inversion

The Euler method is based on the Bromwich contour inversion integral, expressing the function $f(t)$ in terms of its Laplace transform $L(s)$. Making the contour a vertical line $s = a$ such that $L(s)$ has no singularities on or to the right of it gives

$$f(t) = \frac{2e^{at}}{\pi} \int_0^\infty \operatorname{Re}(L(a + iu)) \cos(ut) \, du. \tag{26}$$

This integral can be numerically evaluated using the trapezoidal rule with step-size $h = \pi/2t$ and $a = A/2t$ (where A is a constant that controls the discretisation error), which results in the nearly alternating series:

$$f(t) \approx f_h(t) = \frac{e^{A/2}}{2t} \operatorname{Re}(L(A/2t)) + \frac{e^{A/2}}{2t} \sum_{k=1}^\infty (-1)^k \operatorname{Re} \left(L \left(\frac{A + 2k\pi i}{2t} \right) \right). \tag{27}$$

Euler summation is employed to accelerate the convergence of the alternating series infinite sum, so we calculate the sum of the first n terms explicitly and use

Euler summation to calculate the next m . To give an accuracy of 10^{-8} we set $A = 19.1$, $n = 20$ and $m = 12$ (compared with $A = 19.1$, $n = 15$ and $m = 11$ in [16]).

4.2. Summary of Laguerre inversion

The Laguerre method [9] makes use of the Laguerre series representation:

$$f(t) = \sum_{n=0}^{\infty} q_n l_n(t), \quad t \geq 0, \quad (28)$$

where the Laguerre polynomials l_n are given by

$$l_n(t) = \left(\frac{2n-1-t}{n} \right) l_{n-1}(t) - \left(\frac{n-1}{n} \right) l_{n-2}(t), \quad (29)$$

starting with $l_0 = e^{t/2}$ and $l_1 = (1-t)e^{t/2}$, and

$$q_n = \frac{1}{2\pi r^n} \int_0^\pi Q(re^{iu}) e^{-iru} du, \quad (30)$$

where $r = (0.1)^{4/n}$ and $Q(z) = (1-z)^{-1} L((1+z)/2(1-z))$.

The integral in Eq. (30) can be approximated numerically by the trapezoidal rule, giving

$$q_n \approx \bar{q}_n = \frac{1}{2nr^n} \left(Q(r) + (-1)^n Q(-r) + 2 \sum_{j=1}^{n-1} (-1)^j \operatorname{Re}(Q(re^{\pi j i/n})) \right). \quad (31)$$

As described in [3], the Laguerre method can be modified by noting that the Laguerre coefficients q_n are independent of t . This means that if the number of trapezoids used in the evaluation of q_n is fixed to be the same for every q_n (rather than depending on the value of n), values of $Q(z)$ (and hence $L(s)$) can be reused after they have been computed. Typically, we set $n = 200$. In order to achieve this, however, the scaling method described in [9] must be used to ensure that the Laguerre coefficients have decayed to (near) 0 by $n = 200$. If this can be accomplished, the inversion of a passage time density for any number of t -values can be achieved at the fixed cost of calculating 400 truncated summations of the type shown in Eq. (17). This is in contrast to the Euler method, where the number of truncated summations required is a function of the number of points at which the value of $f(t)$ is required.

4.3. Automatic inversion algorithm selection

We have implemented a distributed Laplace transform inverter which automatically selects which inversion algorithm is the most suitable for a given model in the following manner. Firstly, it attempts to scale the Laguerre coefficients (as de-

scribed in the previous section) and if this scaling is successful the modified Laguerre method is used to invert the Laplace transform $L_{ij}^{-1}(s)$ and yield the required passage time density. If, however, it proves impossible to scale the Laguerre coefficients—typically when the density function contains discontinuities—then the tool switches to the Euler inversion method. This ensures that the Laguerre method is used whenever possible as it minimises the amount of computation needed, while the ability to calculate pathological passage time densities using the Euler method is preserved at the cost of requiring more computation.

5. Hypergraph partitioning

As discussed in Section 3.2, the core operation in our algorithm is the repeated sparse matrix–vector multiplication of Eq. (17). In order to exploit the combined processing power and memory capacity of several processors to compute passage time densities for very large systems (with state spaces of the order of 10^7 states or more), it is necessary to partition the sparse matrix U' across the processors. Such a scheme will necessitate the exchange of data (vector elements and possibly partial sums) after every iteration in the solution process. The objective in partitioning the matrix is to minimise the amount of data which needs to be exchanged while balancing the computational load (as given by the number of non-zero elements of U' assigned to each processor).

Hypergraph partitioning is an extension of graph partitioning. Its primary application to date has been in VLSI circuit design, where the objective is to cluster pins of devices such that interconnect is minimised. It can also be applied to the problem of allocating the non-zero elements of sparse matrices across processors in parallel computation [11].

Formally, a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{N})$ is defined by a set of vertices \mathcal{V} and a set of nets (or hyperedges) \mathcal{N} , where each net is a subset of the vertex set \mathcal{V} [11]. In the context of a row-wise decomposition of a sparse matrix A , matrix row i ($1 \leq i \leq n$) is represented by a vertex $v_i \in \mathcal{V}$ while column j ($1 \leq j \leq n$) is represented by net $N_j \in \mathcal{N}$. The vertices contained within net N_j correspond to the row numbers of the non-zero elements within column j , i.e. $v_i \in N_j$ if and only if $a_{ij} \neq 0$. The weight of vertex i is given by the number of non-zero elements in row i , while the weight of a net is its contribution to the edge cut, which is defined as one less than the number of different partitions spanned by that net. The overall objective of a hypergraph sparse matrix partitioning is to minimise the sum of the weights of the cut nets while maintaining a balance criterion. A column-wise decomposition is achieved in an analogous fashion.

The matrix on the right of Fig. 1 shows the result of applying hypergraph-partitioning to the matrix on the left in a four-way row-wise decomposition. Although the number of off-diagonal non-zeros is 18 the number of vector elements which must be transmitted between processors during each matrix–vector multiplication




Fig. 1. A 16×16 non-symmetric sparse matrix (left), with corresponding four-way hypergraph partition (right) and corresponding partitions of the vector.

(the communication cost) is 6. This is because the hypergraph partitioning algorithms not only aim to concentrate the non-zeros on the diagonals but also strive to line up the off-diagonal non-zeros in columns. The edge cut of the decomposition is also 6, and so the hypergraph partitioning edge cut metric exactly quantifies the communication cost. This is a general property and one of the key advantages of using hypergraphs—in contrast to graph partitioning, where the edge cut metric merely approximates communication cost. Optimal hypergraph partitioning is NP-complete but there are a small number of hypergraph partitioning tools which implement fast heuristic algorithms, for example PaToH [11] and hMeTiS [17].

6. Parallel pipeline

The process of calculating a passage time density (shown in Fig. 2) begins with a high-level model specified in an enhanced form of the DNAmaca interface language [18,19]. This language supports the specification of queueing networks, stochastic Petri nets and stochastic process algebras. Next, a probabilistic, hash-based state generator [20] uses the high-level model description to produce the transition probability matrix, P , of the model's embedded Markov chain, the matrices U and U' , and a list of the initial and target states. Normalised weights for the initial states are determined by the solution of $\tilde{\pi} = \tilde{\pi} P$, which is readily done using any of a variety of steady-state solution techniques (e.g. [21,22]). U' is then partitioned using a hypergraph partitioning tool.

Control is then passed to the distributed passage time density calculator, which is implemented in C++ using the Message Passing Interface (MPI) [23] standard. This employs a master–slave architecture with groups of slave processors. The master processor computes in advance the values of s at which it will need to know the

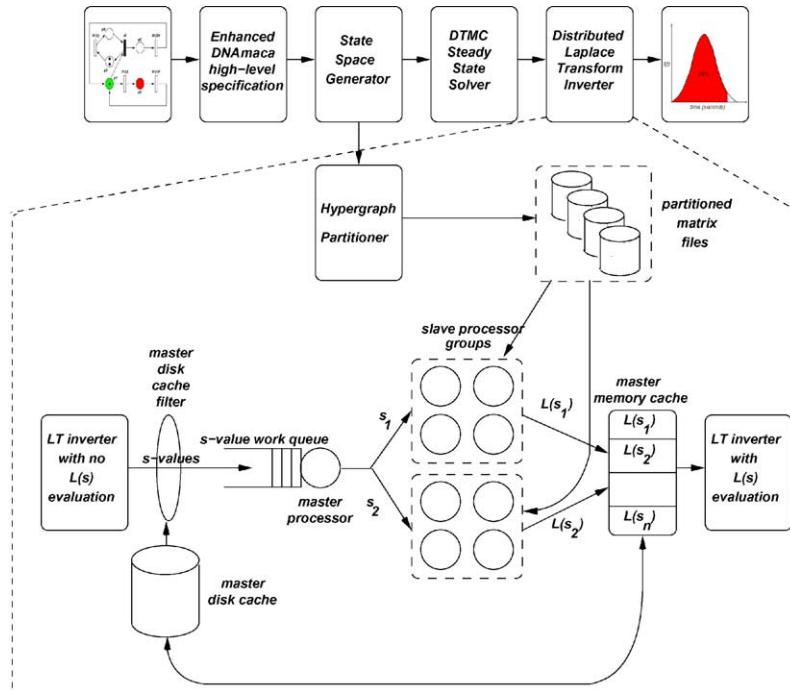


Fig. 2. Parallel hypergraph-based passage time density calculation pipeline.

value of $L_{ij}^{-1}(s)$ in order to perform the inversion. As described in Section 4, this can be done irrespective of the inversion algorithm employed. The s -values are then placed in a global work-queue to which the groups of slave processors make requests.

The highest ranking processor in a group of slaves makes a request to the master for an s -value and is assigned the next one available. This is then broadcast to the other members of the slave group to allow them to construct their columns of the matrix U' for that specific s . Each processor reads in the columns of the matrix U' that correspond to its allocated partition into two types of sparse matrix data structure and also reads in the initial source-state weighting vector $\tilde{\alpha}$. *Local* non-zero elements (i.e. those elements in diagonal matrix blocks that will be multiplied with vector elements stored locally) are stored in a conventional compressed sparse column format. *Remote* non-zero elements (i.e. those elements in off-diagonal matrix blocks that must be multiplied with vector elements received from other processors) are stored in an ultrasparse matrix data structure—one for each remote processor—using a coordinate format. Each processor then determines which vector elements need to be received from and sent to every other processor in the group on each iteration, adjusting the row indices in the ultrasparse matrices so that they index into a vector of received elements. This ensures that a minimum amount of communication

takes place and makes multiplication of off-diagonal blocks with received vector elements efficient.

For each step in our iterative algorithm, each processor begins by using non-blocking communication primitives to send and receive remote vector elements, while calculating the product of local matrix elements with locally stored vector elements. The use of non-blocking operations allows computation and communication to proceed concurrently on parallel machines where dedicated network hardware supports this effectively. The processor then waits for the completion of non-blocking operations (if they have not already completed) before multiplying received remote vector elements with the relevant ultrasparse matrices and adding their contributions to the local vector–matrix product cumulatively.

Once the calculations of a slave group are deemed to have converged, the result is returned to the master by the highest-ranking processor in the group and cached. When all results have been computed and returned for all required values of s , the final Laplace inversion calculations are made by the master, resulting in the required t -points.

7. Numerical results

The results presented in this section were produced on a Beowulf Linux cluster with 64 dual processor nodes. Each node has two Intel Xeon 2.0 GHz processors and 2 GB of RAM. The nodes are connected by a Myrinet network with a peak throughput of 250 Mb/s.

We demonstrate the SMP passage time analysis techniques of the previous sections with large semi-Markov models of a distributed voting system (Fig. 3) and a distributed web server (Fig. 4). The models are specified in a semi-Markov stochastic Petri net formalism (see [24] for full SM-SPN semantics) using an extension of the DNAmaca Markov chain modelling language [18]. Generally distributed transitions, if simultaneously enabled, are selected by probabilistic choice; in this way, we are guaranteed an underlying semi-Markov state space. This probabilistic selection can be seen as an approximation of the more expressive concurrency provided by generalised semi-Markov process (GSMP) models. However, semi-Markov processes are an accurate model in many useful cases; for example when modelling mutual exclusion or probabilistic task switching on a uniprocessor computer system. Concurrent Markovian execution is, however, fully supported (and used in Fig. 4). Transition firing time densities are specified in terms of their Laplace transforms, with macros provided for common distributions (e.g. uniform, gamma, deterministic), and can be made marking dependent by use of the $m(p_i)$ function (which returns the current number of tokens on place p_i). Support for inhibitor arcs is also provided.

Fig. 3 represents a voting system with CC voters, MM polling units and NN central voting servers. In this system, voters cast votes through polling units which in turn register votes with all available central voting units. Both polling units and

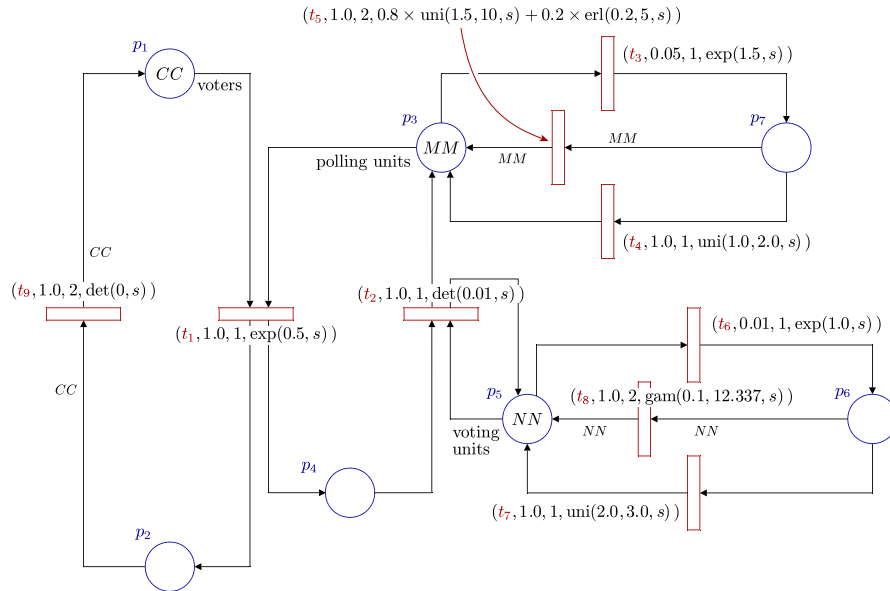


Fig. 3. A semi-Markov stochastic Petri net of a voting system with breakdowns and repairs.

central voting units can suffer breakdowns, from which there is a soft recovery mechanism. If, however, all the polling or voting units fail, then, with high priority, a failure recovery mode is instituted to restore the system to an operational state.

Fig. 4 represents a web server with *RR* clients, *WW* web content authors, *SS* parallel web servers and a write-buffer of *BB* in size. Clients can make read requests to one of the web servers for content (represented by the movement of tokens from p_8 to p_7). Web content authors submit page updates into the write buffer (represented by the movement of tokens from p_1 onto p_2 and p_4), and whenever there are no outstanding read requests all outstanding write requests in the buffer (represented by tokens on p_4) are applied to all functioning web servers (represented by tokens on p_6). Web servers can fail (represented by the movement of tokens from p_6 to p_5) and institute self-recovery unless all servers fail, in which case a high-priority recovery mode is initiated to restore all servers to a fully functional state. Complete reads and updates are represented by tokens on p_9 and p_2 respectively.

For the voting system, Table 1 shows how the size of the underlying SMP varies according to the configuration of the variables *CC*, *MM*, and *NN*. Similarly, for the web server model, Table 2 shows state space sizes in relation to the configuration of the variables *RR*, *WW*, *SS* and *BB*.

In the following, we consider the rate of convergence of the iterative passage time algorithm and the extraction of passage time densities and cumulative distribution functions for the example semi-Markov systems.

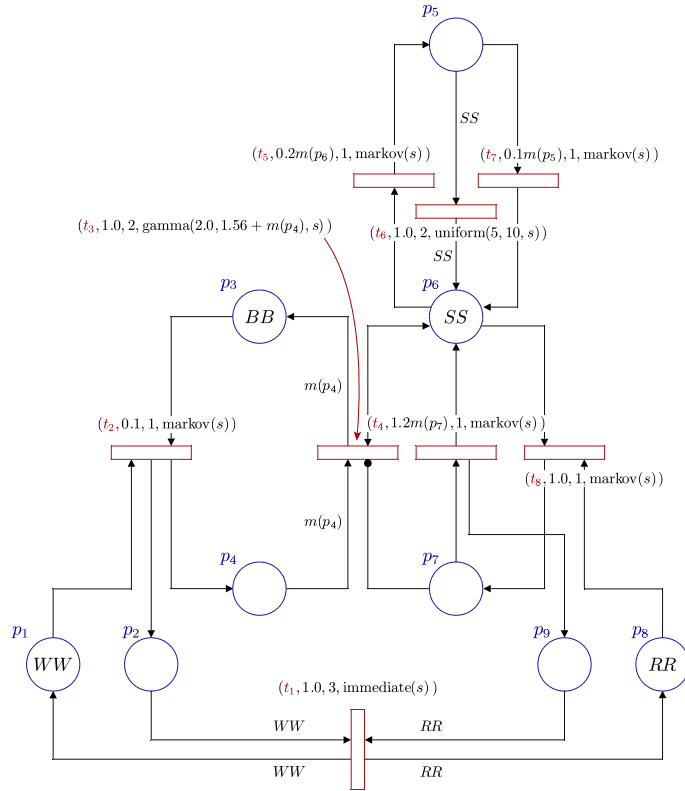


Fig. 4. A semi-Markov Petri net of a parallel web server.

Table 1
Configurations of the voting system as used to present convergence results and passage times

System	CC	MM	NN	States
1	60	25	4	106,540
2	100	30	4	249,760
3	125	40	4	541,280
4	150	40	5	778,850
5	175	45	5	1,140,050
6	300	80	10	10,999,140

7.1. Convergence of iterative passage time algorithm

Our iterative algorithm terminates when two successive iterates are less than ϵ apart (cf. Eq. (24)), for some suitably small value of ϵ . Fig. 5 shows the average number of iterations the algorithm takes to converge per s -point for both models for two different values of ϵ (10^{-8} and 10^{-16}).

Table 2

Configurations of the web server system as used to present convergence results and passage times

System	RR	WW	SS	BB	States
1	45	22	4	8	107,289
2	52	26	5	10	248,585
3	60	30	6	12	517,453
4	65	30	7	13	763,680
5	70	35	7	14	1,044,540
6	100	50	18	20	15,445,919




Fig. 5. Average number of iterations to converge per s point for two different values of ϵ over a range of model sizes.

It is noted that the number of iterations required for convergence as the model size grows is sub-linear; that is, as the model size doubles the number of iterations less than doubles. This suggests the algorithm has good scalability properties. Fig. 6 shows the average amount of time to convergence per s -point, while Fig. 7 shows how the number of iterations per unit time decreases as model size increases. The curves are almost identical for both values of ϵ , suggesting that the time spent per iteration remains constant, irrespective of the number of iterations performed. The rate of computation (iterations per unit time) is $O(1/(N \log(N)))$ for system size N . This gives a time per iteration of $O(N \log(N))$, suggesting an overall practical complexity of better than $O(N^2 \log(N))$ (given the better than $O(N)$ result for the number of iterations required).

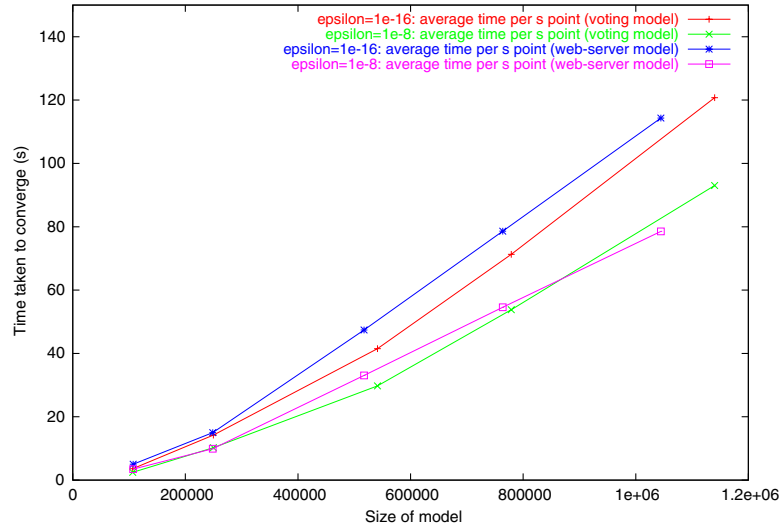


Fig. 6. Average time to convergence per s-point for two different values of ϵ over a range of model sizes.




Fig. 7. Average number of iterations per unit time over a range of model sizes.

7.2. Passage time densities and quantiles

In this section, we display passage time densities produced by our iterative passage time algorithm and validate these results by simulation.




Fig. 8. Analytic and simulated (with 95% confidence intervals) density for the time taken to process 300 voters in the voting model system 6 (10.9 million states).

Fig. 8 shows the density of the time taken to process 300 voters (as given by the passage of 300 tokens from place p_1 to p_2) in system 6 of the voting model. Calculation of the analytical density required 15 h and 7 min using 64 slave processors (in 8 groups of 8) for the 31 t -points plotted. Our algorithm evaluated $L_{i\bar{j}}^{-\bar{j}}(s)$ at 1,023 s -points, each of which involved manipulating sparse matrices of rank 10,999,140. The analytical curve is validated against the combined results from 10 simulations, each of which consisted of 1 billion transition firings. Despite this large simulation effort, we still observe wide confidence intervals (probably because of the rarity of source states).

Fig. 9 is a cumulative distribution for the same passage as Fig. 8 (easily obtained by inverting $L_{i\bar{j}}^{-\bar{j}}(s)/s$ from cached values of $L_{i\bar{j}}^{-\bar{j}}(s)$). It allows us to extract reliability quantiles, for instance:

$$\mathbb{P}(\text{system 6 can process 300 voters in less than 730 s}) = 0.9876.$$

Fig. 10 shows the density of the time taken to perform 100 reads and 50 page updates in the web server model 6. Calculation of the 35 t -points plotted required 2 days, 17 h and 30 min using 64 slave processors (in 8 groups of 8). Our algorithm evaluated $L_{i\bar{j}}^{-\bar{j}}(s)$ at 1,155 s -points, each of which involved manipulating sparse matrices of rank 15,445,919. Again, the analytical result is validated against the combined results from 10 simulations, each of which consisted of 1 billion transition firings. We observe excellent agreement.




Fig. 9. Cumulative distribution function and quantile of the time taken to process 300 voters in the voting model system 6 (10.9 million states).




Fig. 10. Analytic and simulated (with 95% confidence intervals) density for the time taken to process 100 reads and 50 page updates in the web-server model system 6 (15.4 million states).

7.3. Scalability results

Table 3 and Fig. 11 display the performance of the hypergraph-partitioned sparse matrix–vector multiplication operations. They show good scalability with a linear

Table 3
Speedup and efficiency of performing hypergraph-partitioned sparse matrix–vector multiplication across 1–32 processors

Processors	Time (s)	Speedup	Efficiency
1	3968.07	1.00	1.000
2	2199.98	1.80	0.902
4	1122.97	3.53	0.883
8	594.07	6.68	0.835
16	320.19	12.39	0.775
32	188.14	21.09	0.659

Calculated for the 249,760 state voting model for 165 *s*-points.




Fig. 11. Runtime, speedup and efficiency using hypergraph partitioning.

speedup trend, which is unusual in problems of this nature. This is because the hypergraph partitioning minimises the amount of data which must be exchanged between processors. The efficiency is not 100% in all cases, however, as even this reduced amount of inter-node communication imposes an overhead and computational load is not perfectly balanced.

Table 4 and Fig. 12 show the hypergraph scalability in the case where 32 slave processors were divided into various size sub-clusters (32 groups of 1, 16 groups of 2, 8 groups of 4, and so on). This was to measure the benefit to be gained from adding extra groups to draw *s*-points from the global work queue versus doing the computation across larger groups of slave processors (which may be necessary when the state space of the model under analysis is very large). The efficiency decreases as

Table 4
Speedup and efficiency using 32 slave processors divided into various different size sub-clusters

Processors	Time (s)	Speedup	Efficiency
32 × 1	150.13	26.43	0.830
16 × 2	159.55	24.87	0.777
8 × 4	162.13	24.47	0.765
4 × 8	165.24	24.01	0.750
2 × 16	173.76	22.84	0.714
1 × 32	188.14	21.09	0.659




Fig. 12. Runtime, speedup and efficiency for using 32 processors in groups of varying sizes.

the number of groups decreases. Note, however, that with runtimes of between 150 and 188 s, there is still a dramatic improvement over the runtime on a single slave processor (3,968 s) regardless of the group size employed. These results suggest that, given a fixed number of slave processors, it is best to allocate them into the smallest size subgroups (that is, to maximise the number of groups drawing from the global work-queue) subject to the constraints imposed by the size of the model and the memory available on each processor.

8. Conclusion

In this paper, we have derived passage time densities and quantiles from semi-Markov models with over 15 million states. To achieve this, we have developed a parallel iterative algorithm for computing passage time densities. The key to its scalability is the hypergraph data partitioning used to minimise communication and balance load. We represent the general distributions found in SMPs efficiently by using a constant-space representation based on the evaluation demands of a numerical Laplace transform inverter.

We have analysed the truncation error of our algorithm and derived an optimised convergence test that may be applied provided that the dominant eigenvalue of the system can be found. Finally, we have demonstrated the applicability, scalability and capacity of the method on several semi-Markov systems derived from stochastic Petri nets, and have observed practical complexity of better than $O(N^2 \log N)$ processing times for systems with N states.

As future work, we are planning to extend the iterative scheme to generate transient distributions while maintaining a similar complexity profile.

Acknowledgements

The authors would like to thank the Imperial College Parallel Computing Centre for the use of the Viking Beowulf Cluster. We are particularly grateful to Steve Newhouse and Keith Sephton for their help and technical advice.

References

- [1] R. Pyke, Markov renewal processes: definitions and preliminary properties, *Ann. Math. Statist.* 32 (December) (1961) 1231–1242.
- [2] N.J. Dingle, P.G. Harrison, W.J. Knottenbelt, Response time densities in Generalised Stochastic Petri Net models, in: *Proceedings of the 3rd International Workshop on Software and Performance (WOSP'2002)*, Rome, July 2002, pp. 46–54.
- [3] P.G. Harrison, W.J. Knottenbelt, Passage-time distributions in large Markov chains, in: M. Martonosi, E. de Souza e Silva (Eds.), *Proceedings of ACM SIGMETRICS S002*, Marina Del Rey, USA, June 2002, pp. 77–85.
- [4] J.-P. Katoen, M. Kwiatkowska, G. Norman, D. Parker, Faster and symbolic CTMC model checking, in: L. de Alfaro, S. Gilmore (Eds.), *PAPM-PROBMIV 2001, Proceedings of Process Algebra and Probabilistic Methods, Lecture Notes in Computer Science*, vol. 2165, Springer-Verlag, Aachen, 2001, pp. 23–38.
- [5] J.T. Bradley, N.J. Davies, A matrix-based method for analysing stochastic process algebras, in: J.D.P. Rolim (Ed.), *ICALP Workshops 2000, Process Algebra and Performance Modelling Workshop*, University of Waterloo, Ont., Canada, Carleton Scientific, Geneva, 2000, pp. 579–590.
- [6] R. German, D. Logothetis, K.S. Trivedi, Transient analysis of Markov regenerative stochastic Petri nets: a comparison of approaches, in: *PNPM'95, Proceedings of Petri Nets and Performance Models*, Durham, NC, 1995, pp. 103–112.
- [7] G.G. Infante López, H. Hermanns, J.-P. Katoen, Beyond memoryless distributions: model checking semi-Markov chains, in: L. de Alfaro, S. Gilmore (Eds.), *PAPM-PROBMIV 2001, Proceedings of Process Algebra and Probabilistic Methods, Lecture Notes in Computer Science*, vol. 2165, Springer-Verlag, Aachen, 2001, pp. 57–70.
- [8] J.T. Bradley, N.J. Dingle, P.G. Harrison, W.J. Knottenbelt, Distributed computation of passage time quantiles and transient state distributions in large semi-Markov models, in: *Performance Modelling, Evaluation and Optimization of Parallel and Distributed Systems*, IEEE Computer Society Press, Nice, 2003.
- [9] J. Abate, G.L. Choudhury, W. Whitt, On the Laguerre method for numerically inverting Laplace transforms, *INFORMS J. Comput.* 8 (4) (1996) 413–427.
- [10] J. Abate, W. Whitt, The Fourier-series method for inverting transforms of probability distributions, *Queueing Syst.* 10 (1) (1992) 5–88.
- [11] U.V. Catalyürek, C. Aykanat, Hypergraph-partitioning-based decomposition for parallel sparse-matrix vector multiplication, *IEEE Trans. Parallel Distrib. Syst.* 10 (July) (1999) 673–693.
- [12] G. Bolch, S. Greiner, H. de Meer, K.S. Trivedi, *Queueing Networks and Markov Chains*, Wiley, 1998.
- [13] B. Melamed, M. Yadin, Randomization procedures in the computation of cumulative-time distributions over discrete state Markov processes, *Oper. Res.* 32 (July–August) (1984) 926–944.
- [14] J.K. Muppala, K.S. Trivedi, Numerical transient analysis of finite Markovian queueing systems, in: U.N. Bhat, I.V. Basawa (Eds.), *Queueing and Related Models*, Oxford University Press, 1992, pp. 262–284.
- [15] H.J. Wilson, Convergence and correctness of an iterative scheme for calculating passage times in semi-Markov processes, *Tech. Rep.*, 2004/1, Department of Applied Maths, University of Leeds, Woodhouse Lane, Leeds, LS2 9JT, UK, February, 2004, <http://www.doc.ic.ac.uk/~jb/reports/abstract.html#2004-wilson-0>.
- [16] J. Abate, W. Whitt, Numerical inversion of Laplace transforms of probability distributions, *ORSA J. Comput.* 7 (1) (1995) 36–43.
- [17] G. Karypis, V. Kumar, Multilevel k -way hypergraph partitioning, *Tech. Rep.* #98-036, University of Minnesota, 1998.
- [18] W.J. Knottenbelt, Generalised Markovian analysis of timed transitions systems, MSc thesis, University of Cape Town, South Africa, July 1996.

- [19] W.J. Knottenbelt, Parallel performance analysis of large Markov models, PhD thesis, Imperial College, London, UK, February 2000.
- [20] W.J. Knottenbelt, P.G. Harrison, M.S. Mestern, P.S. Kritzing, A probabilistic dynamic technique for the distributed generation of very large state spaces, *Perform. Eval.* 39 (February) (2000) 127–148.
- [21] D.D. Deavours, W.H. Sanders, An efficient disk-based tool for solving very large Markov models, in: *TOOLS 1997, Computer Performance Evaluation: Modelling Techniques and Tools*, Lecture Notes in Computer Science, vol. 1245, Springer-Verlag, St. Malo, 1997, pp. 58–71.
- [22] W.J. Knottenbelt, P.G. Harrison, Distributed disk-based solution techniques for large Markov models, in: *NSMC'99, Proceedings of the 3rd International Conference on the Numerical Solution of Markov Chains*, Zaragoza, September 1999, pp. 58–75.
- [23] W. Gropp, E. Lusk, A. Skjellum, *Using MPI: Portable Parallel Programming with the Message Passing Interface*, MIT Press, Cambridge, MA, 1994.
- [24] J.T. Bradley, N.J. Dingle, W.J. Knottenbelt, P.G. Harrison, Performance queries on semi-Markov stochastic Petri nets with an extended Continuous Stochastic Logic, in: G. Ciardo, W. Sanders (Eds.), *PNPM'03, Proceedings of Petri Nets and Performance Models*, University of Illinois at Urbana-Champaign, IEEE Computer Society, 2003, pp. 62–71.