Asynchronous Iterative Solution for State-Based Performance Metrics

Douglas V. de Jager   Jeremy T. Bradley

Department of Computing
Imperial College London
180 Queen's Gate, London SW7 2AZ
United Kingdom
{dvd03,jb}@doc.ic.ac.uk

ABSTRACT

Solution of large sparse fixed-point problems, $\mathbf{M}\overline{\mathbf{x}} = \overline{\mathbf{x}}$ and $\mathbf{M}\overline{\mathbf{x}} + \overline{\mathbf{b}} = \overline{\mathbf{x}}$, may be seen as underpinning many important performance analysis calculations. These calculations include steady-state, passage-time and transient-time calculations in discrete-time Markov chains, continuous-time Markov chains and semi-Markov chains. In recent years, much work has been done to extend the application of asynchronous iterative fixed-point solution methods to many different contexts. This work has been motivated by the potential for faster solution, more efficient use of the communication channel and/or access to memory, and simplification of task management and programming. In this paper, we present theoretical developments which allow us to extend the application of asynchronous iterative solution methods to solve for the key performance metrics mentioned above—such that we may employ the full breadth of Chazan and Miranker's classes of asynchronous iterations.

Categories and Subject Descriptors

G.3 [**Probability and Statistics**]: Markov processes; G.1.3 [**Numerical Linear Algebra**]: [Eigenvalues and eigenvectors, sparse, structured and very large systems]

General Terms

Performance, algorithms, theory

Keywords

Asynchronous iterations, performance analysis, dominant eigenvectors, Perron–Frobenius, matrix-vector splitting

1. INTRODUCTION

Extracting performance measures from performance models is a computationally intensive process. Typically a performance model will consist of a large Markov chain in excess of 10^9 states, and producing even a simple steady-state vector for a utilisation metric will require considerable amounts of computing resources.

In this paper, we present theoretical developments which show that steady-state, passage-time and transient-time performance measures—when reformulated as linear fixed-point problems—are amenable to asynchronous iterative solution. This type of solution allows each computational node to perform calculation iterations without having to synchronise frequently with other nodes and, thereby, to create communication bottlenecks.

SIGMETRICS'07, June 12–16, 2007, San Diego, California, USA.
ACM 978-1-59593-639-4/07/0006.

2. ASYNCHRONOUS ITERATIONS FOR LINEAR FIXED-POINT PROBLEMS

The theory of asynchronous iterations to solve linear fixed-point problems was originally presented by Chazan and Miranker. [4] It may be outlined as follows.

Consider the linear fixed-point problem:

$$\overline{\mathbf{x}} = \mathbf{A}\overline{\mathbf{x}} + \overline{\mathbf{b}}, \text{ where } \mathbf{A} \in \mathbb{C}^{n\times n}, \overline{\mathbf{x}}, \overline{\mathbf{b}} \in \mathbb{C}^n. \qquad (1)$$

To solve for the unique fixed-point, $\overline{\mathbf{x}}$, we generate a convergent sequence of iterates from the corresponding *class of asynchronous iterations.*

The class of asynchronous iterations is given by the class of sequences, $<\overline{\mathbf{x}}^{(t)}>$, of column vectors in $\mathbb{C}^n$ recursively defined:

$$x_i^{(t+1)} := \begin{cases} \sum_{j=1}^n A_{ij} x_j^{(t-d(t,i,j))} + b_i & \text{if} \quad i \in U(t) \\ x_i^{(t)} & \text{if} \quad i \notin U(t) \end{cases} \qquad (2)$$

where U, d are functions satisfying a set of conditions. Function $U : \mathbb{N} \to \mathbb{P}\{1, \ldots, n\}$ gives the set of vector entries to be updated at each step (where $\mathbb{P}\{1, \ldots, n\}$ denotes the power set of {1,. . . ,n}); function $d : \mathbb{N}\times\{1, \ldots, n\}\times\{1, \ldots, n\} \to \mathbb{N}$ gives the relative "age" of the entries used in the update of each vector entry at each step. These two functions satisfy the following two conditions:

($\mathcal{CM}$1) Each vector entry, i ($1 \le i \le n$), features in an infinite number of update sets.

($\mathcal{CM}$2) For each pair of vector entries, i, j ($1 \le i, j \le n$), we have that both $(t - d(t, i, j)) \to \infty$ as $t \to \infty$ and $\forall t(d(t, i, j) \le t)$.

Generalising minimally the fundamental result from [4], we get the following convergence guarantee for generated sequences. It allows us to know pre-generation whether sequences are convergent.

THEOREM 1. *Consider the class of asynchronous iterations corresponding to equation* (1). *If $\rho(|\mathbf{A}|) < 1$, then every sequence of iterates within the class converges to the unique fixed-point, $\overline{\mathbf{x}}$, as $t \to \infty$; if $\rho(|\mathbf{A}|) \ge 1$, then some non-convergent sequence of iterates exists within the class.*

In this theorem, $|\mathbf{A}|$ denotes the modulus matrix corresponding to $\mathbf{A}$; $\rho(|\mathbf{A}|)$ denotes the spectral radius of $|\mathbf{A}|$.

3. ASYNCHRONOUS ITERATIONS FOR STEADY-STATE VECTORS

It is a standard result that steady-state analysis of both continuous-time Markov chains (CTMCs) and semi-Markov chains (SMPs) can be expressed in terms of the steady-state analysis of their respective embedded discrete-time Markov chains (DTMCs). This has the familiar eigenvector form, $\overline{\mathbf{x}} = \mathbf{M}\overline{\mathbf{x}}$, where $\mathbf{M} \in \mathbb{R}^{n\times n}$ is a column-stochastic square matrix. This DTMC problem is of course a special case of the dominant eigenvector problem:

$$\rho(\mathbf{M})\overline{\mathbf{x}} = \mathbf{M}\overline{\mathbf{x}} \tag{3}$$

where $\mathbf{M} \in \mathbb{C}^{n\times n}$. Without loss of generality we may assume that $\rho(\mathbf{M}) = 1$.

If we are to employ the corresponding asynchronous iterations to solve for this vector, $\overline{\mathbf{x}}$, then we require a convergence guarantee for the iterations. Unfortunately, Theorem 1 is of little immediate help in the case of the dominant eigenvector problem, since from [5] we have:

$$\rho(|\mathbf{M}|) \geq \rho(\mathbf{M}) = 1 \tag{4}$$

However, Theorem 1 may be employed indirectly. This is made possible by proof of the following theorem:

THEOREM 2. *Let irreducible* $\mathbf{M} \in \mathbb{C}^{n\times n}$ *have* $\rho(|\mathbf{M}|) = 1$. *Let* $\mathbf{V}^{(K)} \in \mathbb{C}^{n\times n}$ $(K \subseteq \{1, 2, \ldots, n\})$ *be defined such that* $V^{(K)}_{ij} := v_i$ *if* $j \in K$ *and* 0 *if* $j \notin K$, *where* $\overline{\mathbf{v}} \in \mathbb{C}^n$. *Let* $\overline{\mathbf{v}}$ *and* K *be such that either* $0 \leq Re(V^{(K)}_{ij}) \leq Re(M_{ij})$ *or* $Re(M_{ij}) \leq Re(V^{(K)}_{ij}) \leq 0$, *for every* i, j *pair; and analogously for the imaginary components. Let* $V^{(K)}_{ij} \neq 0$ *for some* i, j *pair.*
Then, the following holds true:

(C1) $\rho(|\mathbf{M} - \mathbf{V}^{(K)}|) < 1$;

(C2) *If 1 is an eigenvalue of* $\mathbf{M}$, *then it is a simple eigenvalue of* $\mathbf{M}$ *and there is a corresponding right eigenvector,* $\overline{\mathbf{x}}$, *of* $\mathbf{M}$, *which comprises no zero entries, and is the unique fixed-point of*

$$\overline{\mathbf{x}} = (\mathbf{M} - \mathbf{V}^{(K)})\overline{\mathbf{x}} + \overline{\mathbf{v}}. \tag{5}$$

Beyond opening the door to unrestricted asynchronous iterative solution, Theorem 2 is significant for at least two reasons. It provides a complex extension of sorts to the Perron–Frobenius theorem: it informs us that if $\mathbf{M}$ comprises negative entries or entries with imaginary parts, and if there is some suitable corresponding dominant eigenvector, $\overline{\mathbf{x}}$, then this eigenvector, $\overline{\mathbf{x}}$, is necessarily unique (up to scalar multiples). Secondly, it offers the potential for improved sparsity patterns and/or reduced spectral radius for the principal matrix—and, thereby, the potential for accelerated convergence using traditional iterative techniques. To this end, an example application involves removing dangling node and teleportation entries from the principal PageRank matrix [6].

4. ASYNCHRONOUS ITERATIONS FOR TRANSIENT AND PASSAGE TIMES

It can be shown that the calculation of passage-time and transient distributions in CTMCs and SMPs ([3] and [2]) is a fixed-point problem of the sort:

$$\overline{\mathbf{x}} = \mathbf{U}'\overline{\mathbf{x}} + \overline{\mathbf{b}} \tag{6}$$

where the entries of $\mathbf{U}'$ are derived from the Laplace transform of the sojourn time distribution transition matrix of the stochastic process. This is soluble by the method of asynchronous iterations as given by Theorem 1 if $\rho(|\mathbf{U}'|) < 1$. We need to prove that this inequality holds and we do so as follows. From [1], we know that $\rho(|\mathbf{U}'|) < 1$ if we can construct a matrix, $\mathbf{M}$, from $\mathbf{U}'$ by replacing some of the zeros in $\mathbf{U}'$ with nonzeros, where the newly constructed $\mathbf{M}$ satisfies all three of the following:

($\mathcal{PT}$1) $\|\mathbf{M}\|_\infty \leq 1$;
($\mathcal{PT}$2) $\exists k \sum_j |M_{kj}| < 1$;
($\mathcal{PT}$3) $\mathbf{M}$ is irreducible.

If this matrix can be constructed then we can deduce—using the Perron–Frobenius theorem and spectral-radius continuity with respect to matrix entries—that:

$$\rho(|\mathbf{U}'|) \leq \rho(|\mathbf{M}|) < 1 \tag{7}$$

and we are done.

We can show that the only cases where we cannot construct such a matrix $\mathbf{M}$ are: if all the transitions in the SMP are immediate—in which case all passage times are 0 without calculation and the transient calculation is degenerate; or there is an absorbing strongly-connected subcomponent of the state-space, which comprises states all of which have immediate out-transitions. In the latter case, the SMP can be aggregated to remove the set of states in a way that has no effect on the passage-time or transient calculation.

In summary, in all non-trivial cases, $\mathbf{U}'$ has $\rho(|\mathbf{U}'|) < 1$, or a suitable $\mathbf{U}'$ is easily found by state aggregation.

5. REFERENCES

[1] AXELSSON, O. *Iterative solution methods.* Cambridge University Press, New York, NY, USA, 1994.

[2] BRADLEY, J. T., DINGLE, N. J., HARRISON, P. G., AND KNOTTENBELT, W. J. Distributed computation of transient state distributions and passage time quantiles in large semi-Markov models. *Future Generation Computer Systems 22*, 7 (August 2006), 828–837.

[3] BRADLEY, J. T., DINGLE, N. J., KNOTTENBELT, W. J., AND WILSON, H. J. Hypergraph-based parallel computation of passage time densities in large semi-Markov models. *Journal of Linear Algebra and Applications 386* (July 2004), 311–334.

[4] CHAZAN, D., AND MIRANKER, W. L. Chaotic relaxation. *Linear Algebra and Its Applications 2* (1969), 199–222.

[5] HORN, R. A., AND JOHNSON, C. R. *Matrix analysis.* Cambridge University Press, New York, 1986.

[6] SEPANDAR D. KAMVAR, TAHER H. HAVELIWALA, C. D. M., AND GOLUB, G. H. Exprapolation methods for accelerating PageRank computations. In *Proceedings of the 12th Int. World Wide Web Conference* (2003).