TWO SOLUTION METHODS FOR MODELS OF PARALLEL QUEUES

Nigel Thomas* and Jeremy Bradley

Research Institute in Software Evolution
University of Durham, UK

ABSTRACT

A class of queueing models is considered here which in general do not give rise to a product form solution but can nevertheless be decomposed into their components, subject to a property referred to as quasi-separability. Such a decomposition gives rise to expressions for marginal probabilities which may be used to derive potentially interesting system performance measures, such as the average number of jobs in the system. In this paper a simple approximation for the variance of the state of a system of quasi-separable components is presented and compared with an alternative method of approximate solution.

INTRODUCTION

Systems of Markovian queues which give rise to product form solutions have been widely studied in the past. In this paper an alternative method of model decomposition is considered that can be found in the queueing network literature, *quasi-separability.* Quasi-separability was developed in the study of queueing systems which suffer breakdowns (Thomas and Mitrani 1995), more recently the approach has been generalised by Thomas et al (Thomas 1999; Thomas and Bradley 2000a) using a Markovian process algebra PEPA. Decompositions of this kind are extremely useful when tackling models with large state spaces, especially when the state space grows exponentially with the addition of further components. Quasi-separability can be applied to a range of models to derive numerical results very efficiently. While it does not generally give rise to expressions for joint probability distributions it does provide exact results for many performance measures, possibly negating the need for more complex numerical analysis. As such it is a very useful means of reducing the state space of large models.

Not all performance measures of interest can be derived exactly from this decomposition. In particular, whilst the average number of jobs in the system may be calculated exactly, in general its variance cannot. It is clearly advantageous however, to gain some confidence in the calculated mean as a useful performance measure without having to solve a much more complicated model. Our proposed solution to this problem is to approximate the variance of the system state by two distinct methods which we compare by numerical experiment. The first method is to estimate the joint queue size distributions directly from exact calculations of the marginal queue size distributions. The second method is to approximate the behaviour of a number of separate queues as a single multi-server queue and use this to calculate the total number of jobs in the system at any time.

In this paper we consider a class of models consisting of a number of nodes in parallel which share a source of jobs. Each node consists of a finite length queue and one or more servers. Jobs are shared amongst the nodes on an a priori basis according to a routing vector which is dependent on the configuration of a scheduler. The scheduler configuration may change independently or in response to changes in the behaviour of the nodes. We show that if the scheduler configuration is not dependent on the number of jobs in the queues, then the system may be decomposed such that each node may be studied in isolation.

THE MODEL

Jobs arrive into the system in a Poisson stream with rate λ. There are N nodes, each consisting of one or more servers with an associated bounded queue. All jobs arrive at a scheduler which directs jobs to a particular node according to its current state. Jobs sent to a queue which is full are lost. The system model is illustrated in Figure 1.

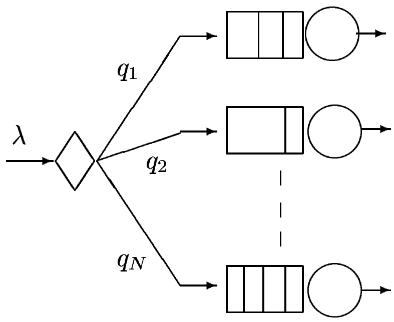

Figure 1: A single source split among N nodes

If, at the time of arrival, a new job finds the scheduler in

*nigel.thomas@durham.ac.uk

configuration σ, then it is directed to node k with probability $q_k(\sigma)$. These decisions are independent of each other, of past history and of the sizes of the various queues. Thus, a routing policy is defined by specifying M vectors, where M is the number of possible configurations which the scheduler may take.

$$\mathbf{q}(\sigma) = [q_1(\sigma), q_2(\sigma), \ldots, q_N(\sigma)] \ \ , \ \ \sigma \subset \Omega_N \ \ , \qquad (1)$$

such that for every σ,

$$\sum_{k=1}^{N} q_k(\sigma) = 1 \ \ . \qquad (2)$$

The system state at time t is specified by the pair $[I(t), \mathbf{J}(t)]$, where $I(t)$ indicates the current scheduler configuration and $\mathbf{J}(t)$ is an integer vector whose kth element, $J_k(t)$, is the number of jobs in queue k ($k = 1, 2, \ldots, N$). Under the assumptions that have been made, $X = \{[I(t), \mathbf{J}(t)], t \geq 0\}$ is an irreducible Markov process.

When the routing probabilities depend on the system configuration, the process X is not separable (i.e., it does not have a product-form solution). As the capacity of the system becomes large, i.e. each queue has a large bound and N is also large, the direct solution of the joint queue size probabilities becomes increasingly costly, although never intractable. The quantities of principal interest are expressed in terms of averages only; they are the steady-state mean queue sizes, L_k, and the overall average response time, W, given by;

$$W = \frac{1}{\lambda} \sum_{k=1}^{N} L_k \ . \qquad (3)$$

To determine these performance measures, it is not necessary to know the joint distribution of all queue sizes; the marginal distributions of the N queues in isolation are sufficient. Unfortunately, the isolated queue processes, $\{J_k(t) \ , \ t \geq 0\}$ ($k = 1, 2, \ldots, N$), are not Markov. However, the performance measures can be determined by studying the stochastic processes $Y_k = \{[I(t), J_k(t)] \, , \, t \geq 0\}$ ($k = 1, 2, \ldots, N$), which model the joint behaviour of the system configuration and the size of an individual queue. The state space of Y_k is dependent only on the capacity of the queue at node K, which simplifies the solution considerably and makes it feasible for reasonably large values of N. The important observation here is that Y_k is an irreducible Markov process, for every k. This is because the arrivals into, and departures from queue k during a small interval $(t \, , t + \Delta t)$ depend only on the system configuration and the size of queue k at time t, and not on the sizes of the other queues.

QUASI-SEPARABILITY

The model presented in the previous section has a property which has become termed *quasi-separability*. Decomposition based on quasi-separability allows expressions to be derived for marginal distributions just as with a product form solution, however unlike product form these marginal distributions cannot, in general, be combined to form the joint distribution for the whole model.

Consider an irreducible Markov process, $X(t)$, which consists of N separate components. The state of each component i can be described by a set of K_i separate variables. Denote by $\mathcal{V}_i$ the set of K_i variables which describe the state of component i. If it is possible to analyse the behaviour of each component, i, of the system exactly by only considering those variables that describe it, i.e. $\mathcal{V}_i$, then the system is said to be *separable*. In this case all the components are statistically independent and a product form solution exists.

For the system to be *quasi-separable* it is necessary only that it is possible to analyse the behaviour of each component, i, of the system exactly by only considering those variables that describe it, $\mathcal{V}_i$, and a subset of the variables from all the other components. Thus the elements of $\mathcal{V}_i$ can be classified into the subsets of either system state variables, $\mathcal{S}_i$ or component state variables $\mathcal{C}_i$, such that:

- the state of $c(t) \in \mathcal{C}_i$ changes at a rate which is independent of the state of any variable $v(t) \in \mathcal{C}_j$, $\forall \ j$ s.t. $j \neq i$.
- the state of $s(t) \in \mathcal{S}_i$ changes at a rate which is independent of the state of any variable $v(t) \in \mathcal{C}_j$, $1 \leq j \leq N$.

If $\mathcal{C}_i \neq \emptyset$, $\forall \ i$, the system can be decomposed into N submodels such that the submodel of the system with respect to the behaviour of component i specifies the changes in the system state variables $\mathcal{S} = \bigcup_{i=1}^{N} \mathcal{S}_i$ and the component state variables $\mathcal{C}_i$. In general the analysis of these submodels gives rise to expressions for their steady-state marginal probabilities if the submodels have stationary distributions with state spaces which are infinite in at most one dimension. As stated above, these marginal probabilities do not, in general, give rise to expressions for the joint probability of the whole system, i.e. no product form solution exists. For quasi-separability to be useful the state space of the submodels should be significantly smaller than the state space of the entire model.

DERIVING VARIANCE FROM MARGINAL PROBABILITIES

If the state space of a model is being reduced then the available information is also reduced unless a product form solution exists. The submodels consist of the system state variables $\mathcal{S} = \bigcup_{i=1}^{N} \mathcal{S}_i$ and the component state variables $\mathcal{C}_i$, hence the steady-state solution of such a system gives probabilities of the form

$p(\mathbf{S}, \mathbf{c}) = p(\mathcal{S} = \mathbf{S}, \mathcal{C}_i = \mathbf{c})$. The solution of the entire model would give rise to probabilities of the form $p(\mathbf{S}, \mathbf{C}) = p(\mathcal{S} = \mathbf{S}, \mathcal{C} = \mathbf{C})$, where $\mathcal{C} = \{\mathcal{C}_1, \ldots, \mathcal{C}_N\}$ and $\mathbf{C} = \{\mathbf{C}_1, \ldots, \mathbf{C}_N\}$. These probabilities are related in the following way for the submodel involving component i subject to the quasi-separability condition,

$$p(\mathcal{S} = \mathbf{S}, \mathcal{C}_i = \mathbf{c}) = \sum_{\forall \mathbf{C} s.t. \mathbf{C}_i = \mathbf{c}} p(\mathcal{S} = \mathbf{S}, \mathcal{C} = \mathbf{C}) \quad (4)$$

If it is possible to associate a value, x_{ij} with each state of a component i then the average state of the component can easily be found. In addition the average of the sum of all components can be found exactly. Thus,

$$E[x_i] = \sum_{\forall j} \sum_{\forall \mathbf{S}} x_{ij} p(\mathcal{S} = \mathbf{S}, \mathcal{C}_i \equiv x_{ij}) \quad (5)$$

Gives the average state of the component, which can be used to derive the average sum,

$$E[x] = \sum_{\forall i} E[x_i] \quad (6)$$

Clearly it is an advantageous property to be able to derive system performance measures from marginal probabilities when they can be found. However, the mean is a special case as the sum of the values is trivially separated. If we consider the same example on variance the problem is evident, in this case there is always at least one term involving the joint probabilities which cannot be broken down to the marginal probabilities. In the general case where there are N components, there will be N terms involving just the marginal probabilities, but $(N-1)!$ terms involving the joint distribution. Clearly then it is not possible to calculate the variance exactly except when a product form solution exists.

The obvious (traditional) solution to this problem is to generate an approximate solution to variance by substituting a product based approximation for the joint distribution. In the case of quasi-separability the situation is slightly complicated since the submodels give rise to marginal probabilities involving not only component variables but also system state variables. The simplest solution (henceforth referred to as the *component state approximation*) would be to eliminate the system state variables by summing over all possible values:

$$p(\mathbf{c}) \approx \prod_{i=1}^{N} \sum_{\forall \mathbf{S}} p(\mathbf{S}, \mathbf{c}_i) \quad (7)$$

where $\mathbf{c} = \{\mathbf{c_1}, \ldots, \mathbf{c_N}\}$. An alternative approach (henceforth referred to as the *system state approximation*) is to attempt to derive approximations for every possible system state:

$$p(\mathbf{S}, \mathbf{c}) \approx \frac{\prod_{i=1}^{N} p(\mathbf{S}, \mathbf{c}_i)}{p(\mathbf{S})^{N-1}} \quad (8)$$

In our earlier study (Thomas and Bradley 2000b) we observed that the component state approximation consistently performs better than the system state approximation. However, the latter has some value as when the two approximations give close values they are observed to be accurate, but when they are diverse they are observed to be somewhat less accurate.

AN ALTERNATIVE APPROXIMATION

It has been shown that in some situations these approximations for the joint probability distribution based on quasi-separability worked well in some situations and less well in others (Thomas and Bradley 2000b). In general we were unable to predict the situations where the approximations worked well, but concluded that if the two approximations gave very close values then they were generally accurate. Clearly there are many other methods which could be applied to this class of models and it is therefore worth considering if any of these can improve on the results we have already for more complex metrics, such as the variance of the number of jobs in the system. The approximation we present here is inspired by an approach presented by Gribaudi and Sereno (Gribaudi and Sereno 2000). It is based on the observation that, for the measures of interest, it is not generally necessary to know the joint queue size distribution, but only the distribution of the total number of jobs in the system.

The quasi-separability decomposition presented in this paper relies on considering each of the queues in isolation to derive the marginal queue size distributions, as such, the information necessary to derive measures based on the total number of jobs in the system is not available. However, if we consider an alternative whereby an individual queue is studied in a model with another queue representing all the remaining queues in the system, we not only derive the marginal queue size distributions exactly, but also the distribution of the total number of jobs in the system approximately. The approximating queue is an amalgamation of all the queues in the system except the one represented explicitly and has a maximum size equal to the sum of all those queues. The total number of available servers are associated with the amalgamated queue and an arrival rate equal to the total effective arrival rate into those queues, i.e. the blocking probability is taken into account. The approximated model is shown in Figure 2.

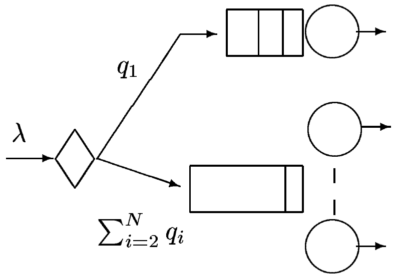

Figure 2: An approximate model of a single source split among N nodes

The entire model has $(K+1)^N M$ states, whereas the quasi-separability solution involves solving N separate models each with $(K+1)M$ states, where K is the maximum size of each queue. This new approximation involves a model with $(K^2(N-1)+NK+1)M$ states, one queue with $K+1$ places and one with $(N-1)K+1$ places. Thus, if K is large and N is small, this new approximation is far more complicated to solve than the quasi-separability case, for example, if $M = 2^N$, $K = 999$ and $N = 3$ then the entire model has 8 thousand million states, the sub-models under quasi-separability have 8 thousand states and the approximate model has a little under 16 million states. If, however, N is large and K less so, then the difference is less pronounced but still significant, e.g. $M = 2^N$, $K = 9$ and $N = 10$ gives the entire model 10240 thousand million states, the sub-models 10240 states, and the approximate model 839680 states.

The accuracy of this approximation is dependent on how accurately the amalgamated queue mimics the flow of jobs through the system. A naive application of this model would involve using the combined arrival and service rates. However, this leads to a distortion as the larger queue size will generally accept more jobs than the sum of the individual queues. We can potentially reduce this problem by matching the sum of the effective arrival rates (the arrival rate of successful jobs) in the complete model with the effective arrival rate in the amalgamated queue. This means resetting the arrival rate in each configuration, σ, and for each number of jobs in the amalgamated queue. As yet we do not have an efficient mechanism for calculating the effective arrival rate into the approximating queue without first deriving the marginal queue size probabilities using quasi-separability and hence finding the blocking probabilities, $p_i(\sigma, K)$, for each queue, i, and every possible scheduler state, σ.

In the infinite buffer case this approximation will always form a lower bound on the total number of jobs, i.e. there will always be no more jobs in the amalgamated queue than in the constituent queues. This is because in the approximation there will never be less service capacity available to a particular job, therefore the server utilisation will always be greater in the approximation, except where utilisation is one for all servers when the approximation becomes exact. This would suggest that if we were to perform this approximated solution considering each queue separately in turn, then the best approximation will be found by selecting a marginal distribution on each scheduler state σ such that the average number of jobs in that state is maximised (since the approximation is always a lower bound). Hence the best approximation will generally be found by combining, on σ, parts of solutions of several the possible approximate models. We could additionally attempt to adjust the offered service rate to approximate the amount of service capacity available, however we have not evaluated either of those approaches as yet.

NUMERICAL RESULTS

For the purposes of numerical comparison we are principally interested in the effectiveness of the amalgamated queue. Hence the results we present are a comparison between a quasi-separability solution and a single amalgamated queue, rather than the complete approximation presented above. The calculation we use for the variance of the number of jobs in the amalgamated queue makes use of the exact calculation for the marginal queue size distributions where possible and only uses the approximation to calculate the unresolvable terms involving the joint queue size distributions.

The majority of known examples of quasi-separability are models of parallel queues, such as the model presented in this paper. The simplest example of these is the case of two queues, 0 and 1, in parallel where the arrival process is controlled by a scheduler which directs jobs to one or other of the queues according to its own internal state, σ, which varies independently of the arrival process and the state of the queues. Hence, $\sigma = 0, 1$, $q_0(0) = 1$ and $q_1(1) = 1$. The scheduler configuration changes from 0 to 1 and from 1 to 0 according to negative exponential rates of ξ and η respectively. Such a model has previously been studied to evaluate the accuracy of the quasi-separability approximation for variance under various conditions (Thomas and Bradley 2000b).

In Figures 3 and 4 we present some numerical results of this same scenario. For reasons of simplicity in

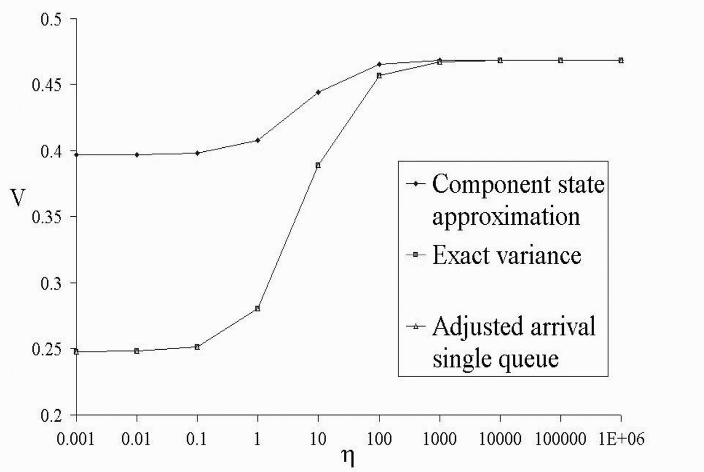

Figure 3: Mean and variance of the total number of jobs against switch rate
with constant proportion of jobs to each queue
$\lambda = 12$, $\xi = \eta$, $\mu_1 = \mu_2 = 10$

solution we have taken the maximum size of each queue to be one, giving the entire model just 8 states. Although this gives rise to a trivial model the comparison is still valid. In Figure 3 we show the three approxima-

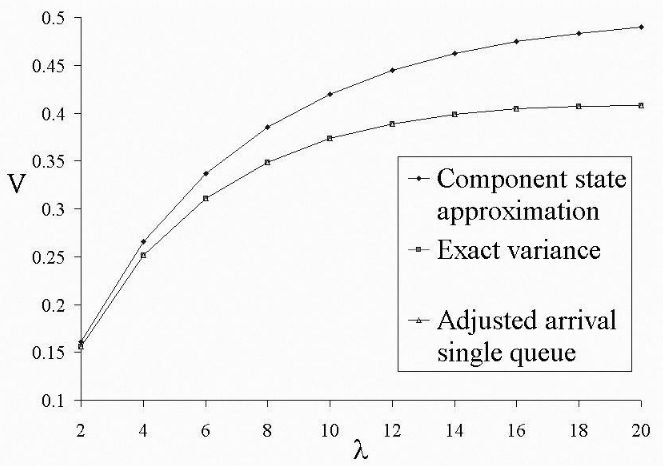

Figure 4: Mean and variance of the total number of jobs against arrival rate
$\mu_1 = \mu_2 = 10$, $\eta = \xi = 10$

tions and exact calculation of variance plotted against the switching rate. In this scenario if the switching rate is high then the model is *almost* product form: in such cases the scheduler changes configuration so frequently relative to other actions that its configuration is seemingly irrelevant and the system is analogous to one having two independent Poisson arrival streams (or a priori splitting). At the other extreme, when the switching rate is very low, the probability that there is a job in both queues is very small, since any job left in the queue after the scheduler has switched away from it will be served relatively quickly. One queue will see a great many jobs before the other queue receives another job. The quasi-separability based approximations work from the premise that the probability of there being a job in a queue is independent of the probability that there is a job in the other, and so the component based approximation shown here performs well when the switching rate is high and poorly when the switching rate is slow.

In Figure 4 we show the various calculations of variance plotted against arrival rate when the switching rate is constant. We can clearly see that the quasi-separability based approximation diverges from the exact result as the arrival rate increases. Again this is due to the fact that this approach assumes that the number of jobs in each queue to be independent of the other, so as the probability of a job being present increases, so does the error. In both cases the single queue approximation performs exceptionally well, indeed any error observed is mainly due to numerical rounding in calculating the effective arrival rate. In this instance we have take the arrival rate when the queue is empty to be λ and calculated the arrival rate when one job is present such that the effective arrival rate in the single queue approximation matches that derived from the quasi-separability solution.

CONCLUDING REMARKS

In this paper we have continued a line of research we have been pursuing for a number of years. Earlier work had identified some potential methods for approximating variance in models exhibiting the quasi-separable property and shown that those methods might give an upper bound to variance. In this paper we have compared those earlier results with a further method whereby the behaviour of the entire model is approximated and the total number of jobs in the system is calculated.

The results we present, based on the evaluation of a very simple model, are exceedingly promising and give a strong indication that this approach is one worth pursuing further. Clearly though, there is much that remains to be done to convince us that the approximation presented here is accurate for more complex models. In particular we need to develop a reliable and efficient method for calculating the effective arrival rates. As indicated by our earlier work we have also been applying these technique to much more general models specified using a Markovian process algebra and we are working towards a general model of state space reduction and model approximation for models subject to the quasi-separability conditions.

REFERENCES

Gribaudo M. and M. Sereno. 2000. “Approximation technique of finite queueing networks exploiting Petri net anlysis”, in: D. Kouvatsos (ed.), *Proceedings of the Fourth International Workshop on Queueing Networks with Finite Capacity*, Ilkley, UK.

Thomas N. and I. Mitrani. 1995. “Routing Among Different Nodes Where Servers Break Down Without Losing Jobs”, in: F. Baccelli, A. Jean-Marie and I. Mitrani (eds.), *Quantitative Methods in Parallel Systems*, Springer-Verlag, 248-261.

Thomas N. 1999. “Extending Quasi-separability”, in: J. Bradley and N. Davies (eds.), *Proceedings of 15th UK Performance Engineering Workshop*, University of Bristol, 131-140.

Thomas N. and J. Bradley. 2000. “Approximating variance in non-product form decomposed models”, in: R. Gorrieri (ed.) *Proceedings of the 8th International Workshop on Process Algebra and Performance Modelling*, Carleton Scientific Publishers.

Thomas N. and J. Bradley. 2000. “Decomposing models of parallel queues”, in: D. Kouvatsos (ed.), *Proceedings of the Fourth International Workshop on Queueing Networks with Finite Capacity*, Ilkley, UK.