Decomposing 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. It is very important that some degree of confidence in such measures can also be given, however, we show here that it is not generally possible to calculate the variance exactly from the marginal probabilities. In this paper a simple approximation for the variance of the state of a system of quasi-separable components is presented and evaluated.

1 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 [3, 7], more recently the approach has been generalised by Thomas et al [4, 5, 6] using the Markovian process algebra PEPA [2]. 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.

*nigel.thomas@durham.ac.uk
†jeremy.bradley@durham.ac.uk

Not all performance measures of interest can be derived exactly from this decomposition. 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. Variance is an extremely important performance measure, knowing how much a system can vary from its mean performance is an essential practical consideration. Furthermore it has been suggested that, in certain situations, it is more desirable for a system to be reliably predictable (more deterministic), i.e. have a low variance, rather than fast, as might be indicted by a low mean [1].

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.

In Section 2 the model is presented. Quasi-separability is discussed in more detail in Section 3 and in Section 4 we show how mean and variance can be calculated from the marginal queue size probabilities derived. Some numerical results are presented in Section 5 and some concluding remarks are made in Section 6.

2 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.

If, at the time of arrival, a new job finds the scheduler in 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 2^N vectors,

$$\mathbf{q}(\sigma) = [q_1(\sigma), q_2(\sigma), \ldots, q_N(\sigma)] \ , \ \sigma \subset \Omega_N \ , \tag{2.1}$$

such that for every σ,

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

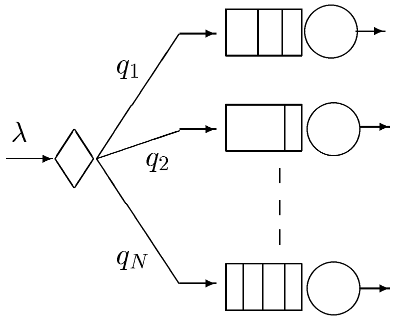

Figure 1: A single source split among N nodes

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 ($k = 1, 2, \dots, N$). 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;

$$W = \frac{1}{\lambda} \sum_{k=1}^{N} L_k \; . \tag{2.2}$$

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, $\{J_k(t)\,,\, t \geq 0\}$ ($k = 1, 2, \dots, N$), are not Markov. However, the performance measures can be determined by studying the stochastic processes $Y_k = \{[I(t), J_k(t)]\,,\, t \geq 0\}$ ($k = 1, 2, \dots, N$), 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.

3 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. Despite the lack of a solution for the joint distribution, many performance measures of interest can still be derived exactly. In addition it is possible to obtain certain whole system performance measures in the form of long run averages, such as the average state of the system and average response time in a queueing network.

A system that is amenable to a quasi-separable solution can be considered informally in the following way. The entire system operates within a single environment, which may be made up of several sub-environments. The state of each component does not alter the state transitions of either the environment or the other components. The behaviour of such components can clearly be studied in isolation from the other components as long as the state of the environment is considered also. The restriction on the behaviour of the components imposed here is unnecessarily strong. We can also consider models where the state space of the components can be separated into that part which does have an impact on state transitions in the environment or other components and that part which has no external influence, not even on the other part of that component. Such a separation requires that the part of a component that influences the state of the environment be considered to be part of the environment for the purposes of model decomposition.

Models such as these have appeared in the literature of the study of queueing systems with breakdowns and rerouting of jobs [3, 7]. In such models the environment is generally made up of the operational state of servers in the system. The routing of jobs to queues is dependent on the operational state of the system i.e. the state of the environment. Such models can generally be decomposed into single queue systems with Markov-modulated arrivals and breakdowns. In the model presented in the previous section; the operational state of each of the servers can be used to determine the scheduler configuration and hence jobs may be directed away from failed nodes. This type of model is conceptually quite simple; there are only two aspects to the state of the components, but in general there may be many aspects of state that must be considered.

Consider an irreducible Markov process, $X(t)$, which consists of N separate components. 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.

4 Deriving mean and 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 $p(\mathbf{S}, \mathbf{c}) = p(\mathcal{S} = \mathbf{S}, \mathcal{C}_i = \mathbf{c})$. These probabilities are related in the following way for the submodel involving component i subject to the quasi-separability condition,

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

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,

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

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

$$E[x] = \sum_{\forall i} E[x_i]$$

Consider, for example, the following case involving just two values:

$$E[x, y] = \sum_{i=1}^{n} \sum_{j=1}^{m} (i + j) p(i, j)$$

$$= \sum_{i=1}^{n} i \sum_{j=1}^{m} p(i, j) + \sum_{j=1}^{m} j \sum_{i=1}^{n} p(i, j)$$

$$= \sum_{i=1}^{n} i p(i, .) + \sum_{j=1}^{m} j p(., j)$$

$$= E[x] + E[y]$$

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.

$$V[x, y] = \sum_{i=1}^{n} \sum_{j=1}^{m} (i + j)^2 p(i, j) - E^2(x, y) \tag{4.1}$$

$$= \sum_{i=1}^{n} \sum_{j=1}^{m} (i^2 + 2ij + j^2) p(i, j) - E^2(x, y)$$

$$= \sum_{i=1}^{n} \sum_{j=1}^{m} i^2 p(i, j) + \sum_{i=1}^{n} \sum_{j=1}^{m} j^2 p(i, j) + \sum_{i=1}^{n} \sum_{j=1}^{m} 2ij p(i, j) - E^2(x, y)$$

$$= \sum_{i=1}^{n} i^2 p(i, .) + \sum_{j=1}^{m} j^2 p(., j) + \sum_{i=1}^{n} \sum_{j=1}^{m} 2ij p(i, j) - E^2(x, y)$$

In this case there is one term involving $p(i, j)$ which cannot be broken down to the marginal probabilities, $p(i, .)$ and $p(., j)$. In the more 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 $p(i,j)$ with $p(i,.)p(.,j)$, i.e. a product based approximation. In the case of quasi-separability the situation is slightly complicated since the submodels give rise to marginal probabilities involving not only component variables (as in the simple example used here), 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:

$$p(\mathbf{c}) \approx \prod_{i=1}^{N} \sum_{\forall \mathbf{S}} p(\mathbf{S}, \mathbf{c}_i) \tag{4.2}$$

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:

$$p(\mathbf{S}, \mathbf{c}) \approx \frac{\prod_{i=1}^{N} p(\mathbf{S}, \mathbf{c}_i)}{p(\mathbf{S})^{N-1}} \tag{4.3}$$

5 Numerical Results

The majority of known examples of quasi-separability are models of parallel queues, such as the model presented in Section 2. 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 $q_1(1) = 1$. At this stage we are not interested in the reason why jobs are routed in this way or why the scheduler configuration changes, but such a situation has been observed in packet switched networks with first generation adaptive routing [8]. The scheduler configuration changes from 0 to 1 and from 1 to 0 according to negative exponential rates of ξ and η respectively.

This model is not separable because the behaviour of both queues is dependent on the scheduler configuration. However, the queues do not directly interact with one another and do not affect the configuration of the scheduler, thus the model may be decomposed into two submodels, each with a single queue. The full model as illustrated has $2(M_0+1)(M_1+1)$ states, whereas each submodel has $2(M_0+1)$ and $2(M_1+1)$ states respectively, where M_i is the capacity of queue i.

For reasons of ease of solution and analysis we have taken the number of places in each queue, M_i, to be one. The entire model therefore has just 8 states and each submodel has 4 states, so each can be solved very easily; of course this will generally not be the case. The process of deriving submodels has not been automated, so the analysis has involved studying three separate models, one for the complete system and one for each of the submodels.

It is clear that the mean number of jobs in the system is an example of the kind of measure discussed in Section 4, where precisely the same value can be found from the solution of the entire model or more simply from the appropriate submodel. We therefore normally calculate this measure from both our solutions in order to give greater confidence that each submodel is correct with respect to the entire model. However, it is not generally possible to calculate exactly the variance of the number of jobs in the system from the submodels and we have therefore employed the approximations outlined in Section 4. The purpose of our numerical experiments is therefore to validate the approximated variance from the submodels with respect to the exact value calculated from the entire model.

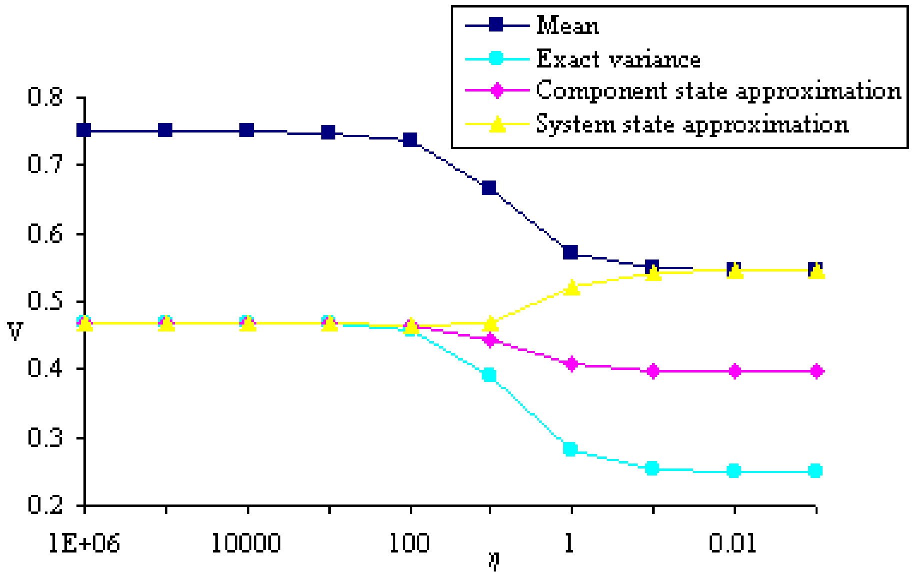

Figure 2: Mean and variance of the total number of jobs against switch rate with constant proportion of jobs to each queue
$\lambda = 12$, $\xi = \eta$, $\mu_1 = \mu_2 = 10$

Figures 2 and 3 show the variance of the total number of jobs in the system against the switch rate at the scheduler and the rate of arrivals respectively. In both cases the queues are identical and the same proportion of jobs is sent to each queue.

In Figure 2 both approximations are extremely good when the rate of switching is high, but both perform relatively badly when the rate of switching is low. The accuracy at high switching rates is easily explained: in such

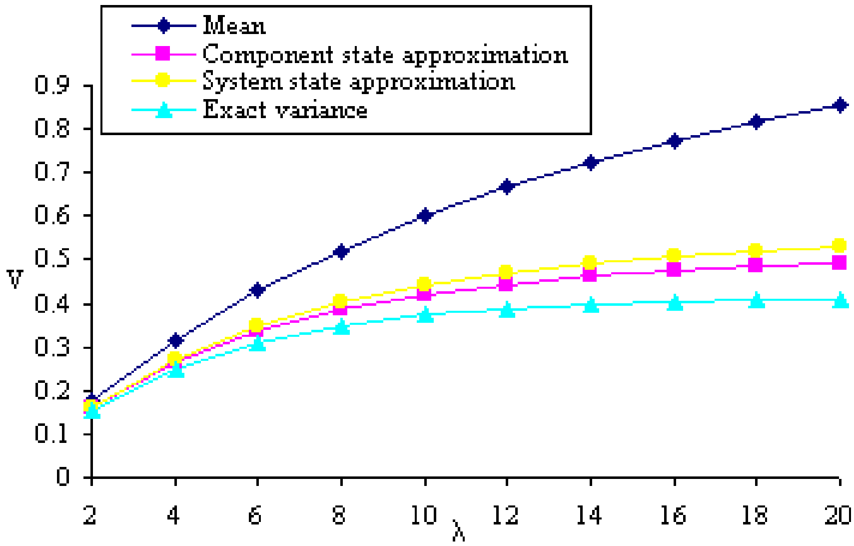

Figure 3: Mean and variance of the total number of jobs against arrival rate $\mu_1 = \mu_2 = 10$, $\eta = \xi = 10$

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. However, both 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 both perform particularly badly in this case. Clearly this suggests a further approximation where the probability of there being a job in both queues is assumed zero. Such an approximation would clearly work well in this situation, but it would be hard to justify in a more general scenario. We have observed in all our experiments with this and other models that the component based approximation forms an upper bound on variance (although we have not attempted to prove this), this further approximation suggests a lower bound, and so the combination might be worth pursuing further.

In Figure 3, we note that neither approximation is particularly accurate (the accuracy around $\lambda = 2$ is an illusion of scale) and also that the system state approximation is less accurate than the component state approximation (also observable on Figure 2). We have generally found this to be the case and, although again we have not proved this yet, this is not surprising as, in the system state approximation, we are attempting to approximate the probability of every state in the system, where as in the component state approximation we are compounding states wherever feasible.

6 Concluding Remarks

In this paper we have illustrated the exploitation of a property known as quasi-separability to decompose a class of finite queueing models. This class of model is similar to the unbounded queueing models studied when developing the notion of quasi-separability. In this paper we consider not only performance measures which it is possible to derive exactly, such as average response time, but also the variance of the number of jobs in the system. Variance is dependent on the joint queue size probabilities which are not derived in the decomposition, thus two approximations for variance are proposed. These approximations are evaluated using a very simple example. Clearly further numerical results need to be obtained in order to evaluate this approach properly.

Acknowledgements

The authors are supported by the EPSRC funded QWiD project (reference GR/N01491) further details of which are available at
http://www.dur.ac.uk/~dcs0nat/qwid.html.

References

[1] J. Bradley and N. Davis, Measuring improved reliability in stochastic systems, in: J. Bradley and N. Davis (eds.), *Proceedings of 15th UK Performance Engineering Workshop*, pp. 121-130, University of Bristol, 1999.

[2] J. Hillston, *A Compositional Approach to Performance Modelling*, Cambridge University Press, 1996.

[3] I. Mitrani and P.E. Wright, Routing in the Presence of Breakdowns, *Performance Evaluation*, **20**, pp. 151-164, 1994.

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

[5] N. Thomas and J. Bradley, Approximating variance in non-product form decomposed models, in: *Proceedings of the 8th International Workshop on Process Algebra and Performance Modelling*, Carleton Scientific Publishers, 2000.

[6] N. Thomas and S. Gilmore, Applying Quasi-Separability to Markovian Process Algebra, in: C. Priami (ed.), *Proceedings of 6th International Workshop on Process Algebra and Performance Modelling*, 1998.

[7] N. Thomas and I. Mitrani, 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*, pp. 248-261, Springer-Verlag, 1995.

[8] J. McQuillan, J. Richer and E. Rosen, The new routing algorithm for the ARPANET, *IEEE Transactions on Communications*, May 1980.