Approximating variance in non-product form decomposed models

NIGEL THOMAS
University of Durham, UK

JEREMY BRADLEY
University of Durham, UK

Abstract

A class of models is considered here which 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 state of 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 a system of quasi-separable components is presented and evaluated.

1 Introduction

In the study of stochastic process algebra we wish to consider not only how systems are to be specified, but also how complex systems can be simplified and solved efficiently. One part of this research has been the study of stochastic process algebra models that give rise to product form solutions (see [4] for a review of these results). 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 systems which suffer breakdowns [5], more recently the approach has been applied by Thomas and Gilmore [7] to models specified using the Markovian process algebra PEPA. This characterisation is based on the decomposition of queueing models into submodels each involving one individual queue. By using PEPA a much more general type of component can be considered. In general there is no bound to the number of variables, which describe the state of a component, the only restriction placed is on the nature of the interaction between components. Such a class of models is considered in [6], and it is this characterisation that we here consider also.

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. Not all performance measures of interest can be derived exactly from this decomposition. In particular, whilst the average state of 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. Variance is an extremely important performance measure, knowing how much a system can vary from its mean performance is an essential practical consideration. If the variation of behaviour is large then having only the mean figure for a sojourn is probably not much use for evaluation. Furthermore it has been suggested that it is may be 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].

The structure of the paper is as follows. In Section 2 the concept of quasi-separability is introduced. Steady state average component state and variance are discussed and a simple approximation for variance is presented in Section 3. A characterisation of quasi-separable PEPA models is defined in Section 4. In Section 5 the approach is illustrated through an example, which is used to evaluate to the approximation through numerical experiment, followed by some conclusions.

2 Quasi-Separability

A 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. Clearly, since exact expressions for marginal probabilities can be found, it is possible to derive any performance measure that depends on a single component. 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. Several components operate within this environment such that their behaviour is affected by the state of the environment. 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 is considered to be part of the environment for the purposes of model decomposition.

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.

3 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 $p(\mathbf{S},\mathbf{c}) = p(\mathcal{S} = \mathbf{S}, \mathcal{C}_i = \mathbf{c})$. A 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 $\mathcal{C} = \{\mathcal{C}_1,\dots,\mathcal{C}_N\}$ and $\mathbf{C} = \{\mathbf{C}_1,\dots,\mathbf{C}_N\}$. These probabilities are related in the following way for the submodel involving component i subject to the quasi-separability condition,

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

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:

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

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

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. If we consider the same example on variance the problem is evident.

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

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

$$= \sum_{i=1}^{n} i^2 p(i,.) + \sum_{j=1}^{m} j^2 p(.,j) + \sum_{i=1}^{n} \sum_{j=1}^{m} 2ij p(i,j) - E^2(x,y) \quad (3.1)$$

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 a 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:

$$p(\mathbf{c}) \approx \prod_{i=1}^{N} \sum_{\forall \mathbf{S}} p(\mathbf{S}, \mathbf{c}_i) \tag{3.2}$$

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:

$$p(\mathbf{S}, \mathbf{c}) \approx \frac{\prod_{i=1}^{N} p(\mathbf{S}, \mathbf{c}_i)}{p(\mathbf{S})^{N-1}} \tag{3.3}$$

4 Quasi-separable PEPA components

In this section a characterisation of quasi-separability in PEPA is derived. This characterisation relies on the fact that the model specification in PEPA exactly mirrors the definition of a quasi-separable Markov process given in Section 2. However, in reality there are a great many possible specifications for any given model therefore understandable that this characterisation captures a very limited class of models.

A system of N components expressed in PEPA is now considered. The components are modelled as PEPA components, $C_1, \ldots, C_N$ say. The interactions of these components are co-ordinated by another component S, which is referred to as the scheduler. These components are related to the underlying Markov process of the form discussed in Section 2 such that, S represents all the system state variables $\mathcal{S}$ and C_i represents *at least* all the component state variables $\mathcal{C}_i$, $i = 1, \ldots, N$. That is, the C_i's represent those parts of the system that are considered in isolation and S represents those parts which are not. Note that in general each component C_i may represent more than just $\mathcal{C}_i$ as it may be desirable to include information that restricts the set of possible legal combinations between elements of $\mathcal{S}_i$ and $\mathcal{C}_i$.

No actions are synchronised between the components $C_1, \ldots, C_N$. All actions enabled in C_i are either synchronised between it and S, contained in the set L, or are unique to C_i. Only actions which can be internalised are enabled in the scheduler, S. Any actions in the cooperation set L either do not change S or, if they do, they cannot be blocked by C_i. Thus the system can be described in the following way;

$$(C_1 \| \cdots \| C_N) \underset{L}{\bowtie} S \tag{4.1}$$

The set of synchronised actions L can be rewritten as N subsets L_i, $i = 1, \ldots, N$. The subset L_i contains only those actions that are synchronised between S and C_i i.e. those actions in L and enabled in C_i (and not hidden). Denote by Z_i the set of all actions that are synchronised between S and each C_j such that $j \neq i$, i.e.

$$Z_i = \bigcup_{\forall\, j \neq i} L_j$$

Clearly, $Z_i \subset L$ and $L = (L_i \cup Z_i)$, $i = 1, \ldots, N$. If the actions in subsets Z_i are such that their rates are not influenced by C_i, i.e. $L_i \cap Z_i = \emptyset$, or l is enabled in every derivative of C_i, $1 \leq i \leq N$, with rate $\top$, $\forall\, l \in (L_i \cap Z_i)$.

The complete model specification can then be rewritten in the following way;

$$(C_1 \underset{L_1}{\bowtie} S) \underset{Z_1}{\bowtie} (C_2 \| \cdots \| C_N) \tag{4.2}$$

such that (4.1) and (4.2) are *isomorphic*. Therefore, the subsystem $(C_1 \underset{L_1}{\bowtie} S)$ can now be isolated from the system without altering the individual behaviour of C_1 in any way, subject to the conditions for L_1 and Z_1 given above.

The same approach can be applied for every C_i, L_i and Z_i to give expressions for all N subsystems. A system that can be treated in such a way satisfies the conditions for *quasi-separability* presented in Section 2.

In the characterisation presented above the separation of system state and component variables has already been made. An alternative characterisation is possible whereby the model is expressed as the cooperation of N components.

$$V_1 \underset{Y_1}{\bowtie} V_2 \underset{Y_2}{\bowtie} \cdots \underset{Y_{N-1}}{\bowtie} V_N \;\; where,\; V_i \stackrel{def}{=} S_i \underset{W_i}{\bowtie} C_i \;\; , \;\; 1 \leq i \leq N$$

such that, S_i represents all the system state variables $\mathcal{S}_i$ and C_i represents *at least* all the component state variables $\mathcal{C}_i$, $i = 1, \ldots, N$. There are two possible cases for the synchronisation set W_i for which quasi-separability may hold;

- no action of type $w \in W_i$ causes a change in the behaviour of S_i
- any action of type $w \in W_i$ which causes a change in the behaviour of S_i is enabled in every derivative of C_i with rate $\top$.

An additional pathological case exists whereby an action of type $w \in W_i$ is blocked in every possible combination of S_i and C_i, i.e. when actions of type $w \in W_i$ have no effect within the model. Thus it is assured that the behaviour of S_i is independent of the behaviour of C_i, but the converse is not generally true. Furthermore it is necessary to ensure that the behaviour of S_i is independent of the state of any C_j, $1 \leq j \leq N$. Therefore the cooperation set, Y_i, between components is similarly restricted such that any action of type $y \in Y_i$ which is enabled in C_j, $1 \leq j \leq N$ is enabled in every derivative of C_j with rate $\top$. Hence it is possible to form an expression for S which represents all the system state variables $\mathcal{S}$.

$$S \stackrel{def}{=} S_1 \underset{Y_1}{\bowtie} S_2 \underset{Y_2}{\bowtie} \cdots \underset{Y_{N-1}}{\bowtie} S_N$$

Furthermore, since any shared actions which are enabled in C_i, $1 \leq i \leq N$ are enabled in every derivative of C_i with rate $\top$, there are no shared actions in C_j, $1 \leq j \leq N$, which need the cooperation of C_i. Hence the model can be rewritten as,

$$(C_1 \parallel \cdots \parallel C_N) \underset{W \cup Y'}{\bowtie} S$$

where $W = \bigcup_{i=1}^{N} W_i$ and $Y' \subset \bigcup_{i=1}^{N} Y_i$ such that an action of type $y \in Y'$ is enabled in at least one of the C_i's ($1 \leq i \leq N$). This expression can now be rewritten in the same way as for the previous case to give,

$$C'_j \stackrel{def}{=} (C_1 \parallel \cdots \parallel C_{j-1} \parallel C_{j+1} \parallel \cdots \parallel C_N)$$
$$(C_j \underset{L_j}{\bowtie} S) \underset{Z_j}{\bowtie} C'_j$$

where,

$$\begin{aligned}
L_j &= W_j \cup Y'_j \\
Z_j &= (W_1 \cup \cdots \cup W_{j-1} \cup W_{j+1} \cup \cdots \cup W_N) \cup X_j \\
X_j &= Y'_1 \cup \cdots \cup Y'_{j-1} \cup Y'_{j+1} \cup \cdots \cup Y'_N \\
Y'_j &\subset \bigcup_{i=j-1}^{N} (Y_i \cap Y') \text{ such that an action of type } y \in Y'_j \text{ is enabled} \\
&\quad \text{(and not hidden) in } C_j \ (Y_0 = \emptyset).
\end{aligned}$$

Hence the set W_j denotes the actions shared between C_j and S_j, Z_j denotes the actions shared between the C_i's ($i \neq j$) and S, and Y'_j denotes the actions shared between C_j and S_i ($\forall\, i \neq j$).

Clearly the characterisations presented here are in an idealised form and require further development for more general models to be considered. However, the characterisation of quasi-separable PEPA components given here does not severely restrict the structure of the components in any way other than through the actions that are shared between them.

5 Example

Most known examples of quasi-separability are models of parallel queues. The simplest example of these is the case of two queues 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, which varies independently of the arrival process and the queues. Such a model is illustrated in Figure 1.

$$\begin{array}{rcl}
Queue_{1,0} & \stackrel{def}{=} & (arrival1, \top).Queue_{1,1} \\
Queue_{1,i} & \stackrel{def}{=} & (arrival1, \top).Queue_{1,i+1} + (service1, \mu_1).Queue_{1,i-1} \\
 & & 1 \leq i \leq N-1 \\
Queue_{1,N} & \stackrel{def}{=} & (service1, \mu_1).Queue_{1,N-1} \\
 & & \\
Queue_{2,0} & \stackrel{def}{=} & (arrival2, \top).Queue_{2,1} \\
Queue_{2,j} & \stackrel{def}{=} & (arrival2, \top).Queue_{2,j+1} + (service2, \mu_2).Queue_{2,j-1} \\
 & & 1 \leq j \leq M-1 \\
Queue_{2,M} & \stackrel{def}{=} & (service2, \mu_2).Queue_{2,M-1} \\
 & & \\
Scheduler_1 & \stackrel{def}{=} & (arrival1, \lambda).Scheduler_1 + (toTwo, \eta).Scheduler_2 \\
Scheduler_2 & \stackrel{def}{=} & (arrival2, \lambda).Scheduler_2 + (toOne, \xi).Scheduler_1
\end{array}$$

$$(Queue_{1,0} || Queue_{2,0}) \underset{\{arrival1,\ arrival2\}}{\bowtie} Scheduler_1$$

Figure 1: A PEPA model of two queues with a shared arrival process

This model is not separable because the behaviour of both queues is dependent on the scheduler state. However, the queues do not directly interact with one another and do not affect the state of the scheduler, thus the model may be decomposed into two submodels of the form,

$$(Queue_{1,0} \underset{arrival1}{\bowtie} Scheduler_1) \underset{arrival2}{\bowtie} Queue_{2,0}$$

and,

$$(Queue_{2,0} \underset{arrival2}{\bowtie} Scheduler_1) \underset{arrival1}{\bowtie} Queue_{1,0}$$

The full model as illustrated has $2(N+1)(M+1)$ states, whereas each submodel has $2(N+1)$ and $2(M+1)$ states respectively. The state space saving with this model is clearly relatively modest, but examples considered in previous papers [6, 7] have a much greater saving.

5.1 Numerical Results

This model has been analysed using the PEPA Workbench [2]. For reasons of ease of solution and analysis we have taken the number of places in each queue, N, to be one. The entire model therefore has just 8 states and each submodel has 4 states, so each can be solved very easily, although 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. The PEPA Workbench outputs a hash file, a table file and a Maple input file. The table file represents the mapping states of the system to a matrix representation given in the Maple input file, using identifiers, which are mapped to the named PEPA agents in the hash file. Running Maple with the given input file (matrix) gives the steady state probabilities for the model, we can then derive the measures of interest (in Maple) by finding the states needed in the table file (via the hash file). For example, the mean number of jobs in the first queue is found by summing each steady state probability where the state includes $Queue_{1,1}$.

It is clear that the mean number of jobs in the system is an example of the kind of measure discussed in Section 3, 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 exactly calculate the variance of the number of jobs in the system from the submodels and we have therefore employed the approximations outlined in Section 3. 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.

The calculation of the variance for this model with $N = 1$ is given by (3.1) to be the sum of all probabilities of states including $Queue_{1,1}$ plus the sum of all probabilities of states including $Queue_{2,1}$ plus twice the sum of all states including both $Queue_{1,1}$ and $Queue_{2,1}$ minus the square of the mean. The submodels will give exact values for the sum of all probabilities of states including $Queue_{i,1}$ (where $i = 1,2$), so the only approximation is in calculating the sum of all joint probabilities for the states including both $Queue_{1,1}$ and $Queue_{2,1}$. The approximations for variance employed here are as follows in accordance with the definitions given in Section 3 (3.2, 3.3):

- *Component state approximation*. The sum of all joint probabilities for the states including both $Queue_{1,1}$ and $Queue_{2,1}$ is calculated as the product of the sums of the probabilities of all states including $Queue_{1,1}$ and $Queue_{2,1}$ respectively in the appropriate submodels.

- *System state approximation*. The sum of all joint probabilities for the states including both $Queue_{1,1}$ and $Queue_{2,1}$ is calculated as the sum of all the approximate system state probabilities. Where the approximate system state probabilities are defined for each possible combination of other agents as the product of all states including $Queue_{1,1}$ and that possible combination of other agents and all states including $Queue_{2,1}$ and that possible combination of other agents divided by the sum of the probabilities of all states including that possible combination of agents.

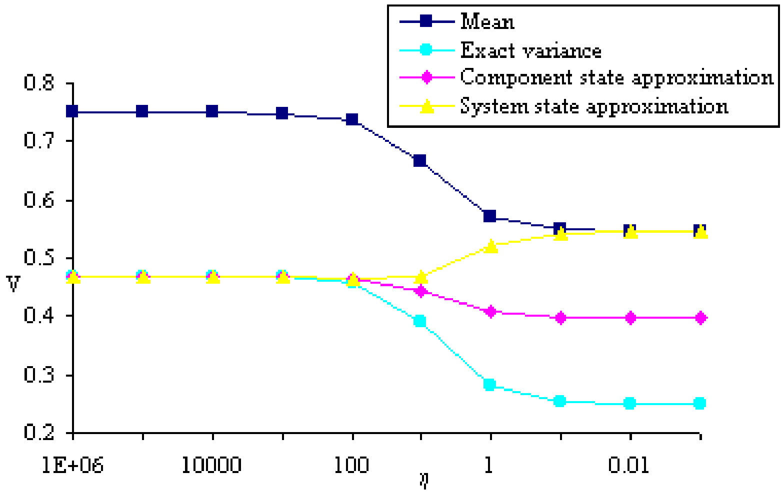

Figure 2: Mean and variance of the total number of jobs against switch rate with constant proportion of jobs to each queue
$\lambda = 12, \xi = \eta, \mu_1 = \mu_2 = 10$

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 as in such cases the scheduler changes state so frequently relative to other actions that its state 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

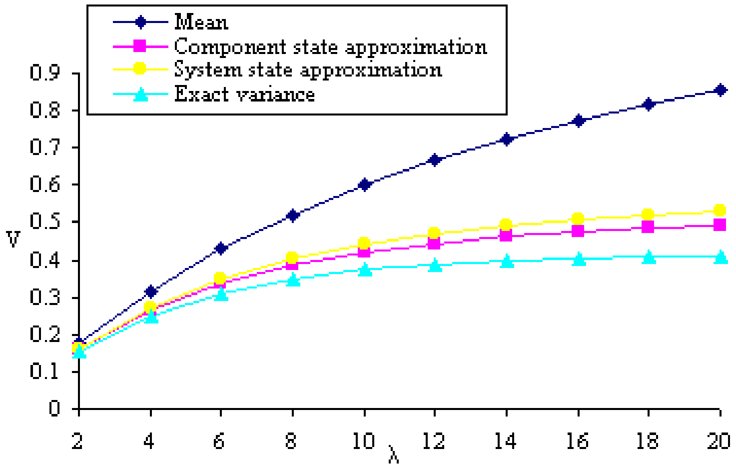

Figure 3: Mean and variance of the total number of jobs against arrival rate $\mu_1 = \mu_2 = 10, \eta = \xi = 10$

clearly work well in this situation, but it would be hard to justify in a more general scenario.

In Figure 3 we note that neither approximation is particularly accurate 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 unsurprising 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.

A final comment concerns a rather surprising result we obtained from another model. In a model of distributed sensors with unreliable communication (see [6]) we found that the component state approximation gave results within 1% of the exact variance consistently over all the values of rates that we tried. This degree of error could easily be counted for in the numerical solution of the matrices, and so it is probable that the approximation is exact for this model. Such a result is not a property of quasi-separability, and so we are investigating the possibility that this particular example has further properties useful in decomposition.

6 Concluding Remarks

Earlier results have justified the study of decomposition techniques to reduce the state space of large models to make their solution feasible. The approximations for variance in submodels without product form solution presented here extends the applicability of those earlier results and adds the confidence that can be placed on performance measures derived in this way. Clearly there are a great many classes of performance model for which measures such as the mean and variance used in this paper are not applicable. However, the approximation used for variance proposed here is actually formed from an approximation for the probability of system state, which is clearly of much wider interest. It should also be pointed out that the applicability of these approximations is limited by the degree of error introduced, although variance is generally employed as a relative performance measure used to determine the degree of confidence in other measures, such as mean here.

Acknowledgements

The authors are supported by the EPSRC funded QWiD project (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] G. Clark, S. Gilmore, J. Hillston and N. Thomas, Experiences with the PEPA performance modelling tools, *IEE Proceedings - Software*, **146**(1), 1999.

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

[4] J. Hillston, Exploiting Structure in Solution: Decomposing Composed Models, in: C. Priami (ed.), *Proceedings of 6th International Workshop on Process Algebra and Performance Modelling*, 1998.

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

[6] 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.

[7] 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.