Fluid Queue Models of Energy Storage and Distribution

Gareth L. Jones, Peter G. Harrison, Uli Harder, Anthony J. Field
Imperial College London,
180 Queen's Gate, London, SW7 2RH, United Kingdom
{gljones,pgh,uh,ajf}@doc.ic.ac.uk

July 2, 2012

Abstract

In this paper we present an approximation method for networks of fluid queues. This approximation method can be applied to a very large class of fluid networks, a significantly broader class than that previously proposed by some of the authors of this paper.

Previously the approximation required that every input to a queueing node was at a rate greater than the service rate of that node. We relax this restriction and allow inputs to be at any rate.

We consider a network which describes energy distribution and storage, an increasingly important problem to understand as renewable (and less reliable) energy generation becomes more common around the world.

We apply our new approximation method to some example networks which could not have been analysed with the previously proposed algorithm and show encouraging results.

1 Introduction

This paper is concerned with networks of fluid queues. Results for the stationary distribution of single fluid queues are well known (see, for example Kulkarni's survey [14]) and the busy period distribution is also known for infinite buffers fed by Markov modulated sources [5, 9].

However, surprisingly little progress has been made on performance metrics for networks of these queues. Two queues in tandem with on/off arrivals at

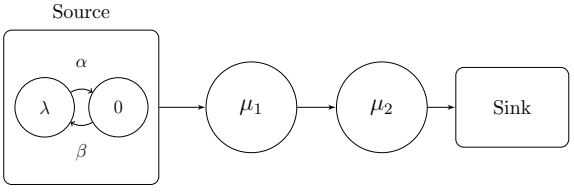


Figure 1: Simple tandem network as considered in [7] shown not to have a product form stationary distribution.

the first node—the simplest sort of network as shown in Figure 1—have been studied [13, 16, 17, 7] and shown not to exhibit a product form solution. In slightly more complicated networks like Figure 3 (with $\mu > 45$), computation of the moments of the busy period at the third node is an open problem (because this node is fed by non-exponential sources each feeding at a rate less than the service rate of the node).

Therefore to analyse fluid networks, we must embrace some form of approximation algorithm. In this paper an approximation which can be applied to a very general class of networks of fluid queues is introduced. Previously the approximation considered only a single input [7] or a multiple inputs [8] with the restriction that every input in to every node was at a rate faster than the service rate of that node. Examples of networks that satisfied this property were studied in these previous publications, but they are a very restrictive class. Suppose a node is fed by 3 other nodes. Then to fit in this restriction each of the feeding nodes must be active less than $\frac{1}{3}$ of the time on average and if each of the source nodes is in turn fed by a further three sources then each of these original sources must be active less than $\frac{1}{9}$ of the time on average.

1.1 Modelling renewable power networks

Globally a movement to renewable power generation is underway. One recent study suggests that currently available commercial technologies are sufficient to supply 80% of US electricity generation [3]. However, the generation of and demand for renewable power generation do not always coincide. Indeed, Texas has invested heavily in wind power and last summer found itself having to buy electricity on the spot market at up to 30 times the regular price. Wind power generates most power at night, but demand peaks during the hot summer day with air conditioning use [2].

Energy storage is one of the ways to tackle this problem. Currently 99% of the world's installed grid connected storage is in the form of pumped hydroelectric storage where water is pumped backwards through a hydroelectric plant. The technology is about 75% efficient, but has limited scope for further development because of limited number of suitable sites.

Significant investment will be needed in storage capacity, and planning efficient energy and storage distribution networks will be important if renewable generation goals are to be met. We envisage networks of fluid queues as being helpful in this direction. The typical capacity factors for wind turbines are 20–40% [1]. As we will see later the new approximation method introduced here is essential for evaluating the kind of networks we are likely to see in renewable networks, networks the previous approximation method could not handle.

In this paper we introduce a framework for modelling such systems using networks of fluid queues. We show how an approximation algorithm can be used to compute metrics of interest (busy period and fluid level moments) in directed acyclic networks.

2 Fluid queues

A fluid queue is a bivariate stochastic process (J_t, X_t) defined by a generator matrix Q, input rate vector λ, service rate scalar μ and initial conditions J_0 and

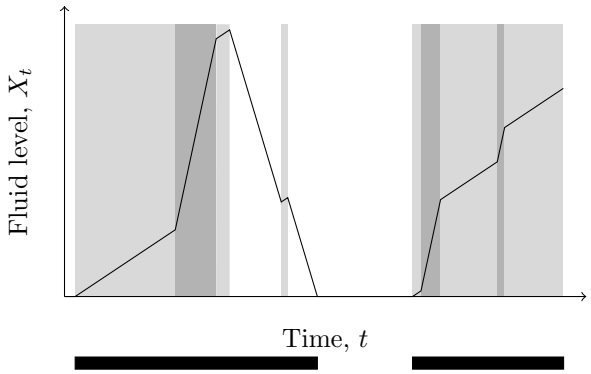


Figure 2: A sample trace from the fluid level at node 1 in Figure 3. Busy periods are highlighted with a thick black line.

X_0. J_t is a continuous time Markov chain with initial state J_0 and generator matrix Q. X_t is a stochastic process with X_0 given and X_t such that

$$\frac{\mathrm{d}X_t}{\mathrm{d}t} = \begin{cases} \lambda_{J_t} - \mu & \text{if } X_t > 0 \\ \max(\lambda_{J_t} - \mu, 0) & \text{if } X_t = 0. \end{cases}$$

A sample trace for node 1 in Figure 3 is shown in Figure 2. There are two sources feeding node 1. The background highlighting represents the state of the Markov process J_t: white represents periods when both sources are off, light grey when only one source is on and darker grey when both sources are on.

The *busy period* of a fluid queueing node is the time period for which the fluid level in the buffer at the particular node is greater than zero. Busy periods are highlighted in Figure 2 with a thick black line.

3 Networks of fluid queues

Networks of fluid queues consist of sources, nodes and sinks. Sources produce fluid at specified rates for stochastically determined periods of time. This fluid is send to service nodes, which process fluid at a specified rate.

Fluid leaving a node is fed in to other nodes in the network. When a buffer level is greater than zero the output of that buffer is at the service rate of the buffer. When a buffer level is equal to zero, the output of the buffer is equal to the current input rate to the buffer. After service at one or more nodes, fluid is routed to a sink where it leaves the network.

When a fluid flow is branched probabilistically the output rate is split according to the probabilities between the receiving nodes. In Figure 7 the output from node 2 is split between nodes 4 and 5 according to a $\frac{3}{4} : \frac{1}{4}$ split, so when the output of node 2 is at rate 30, $\frac{90}{4}$ is sent to node 4 and $\frac{30}{4}$ to node 5.

For the analysis of long term averages (the performance metrics we seek to compute) we require that each node and the entire network is stable. That is, the utilization of each service node is less than 1. My We also require that for each service node there is some input state where the buffer fills. This is an

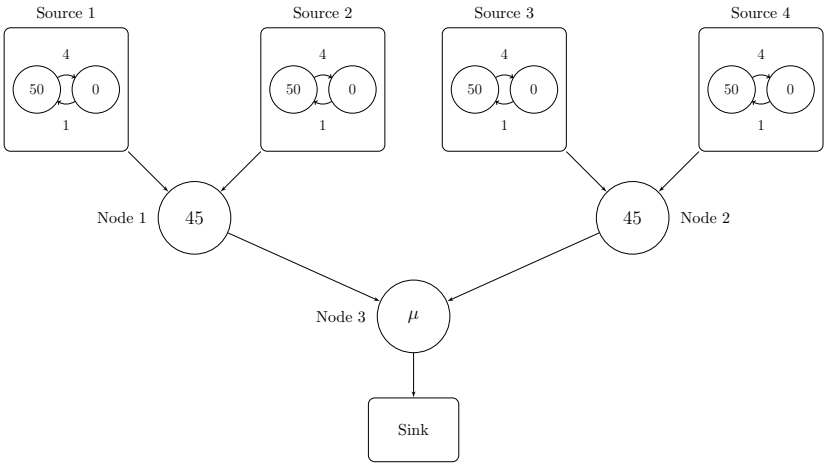


Figure 3: Topology of the first example network considered. For $40 < \mu < 45$ the busy period moments at node 3 can be computed using Boxma and Dumas' result [5], but for $\mu > 45$ we require the new approximation method.

intuitively reasonable requirement, otherwise the service node is defunct and all fluid passes immediately through as though the node did not exist.

The algorithm proposed here allows for the (approximate) computation of moments of both busy periods and mean fluid levels at all nodes in the network. Another approach to fluid levels in networks of fluid queues using martingales can be found in [11, 12].

4 Example 1

We first consider a network with three storage nodes. Two are locally connected to intermittent sources, each of which are then connected to a central storage which feeds a consumer. We represent the demand of consumers by the rate at node 3 (μ) so the busy period at node 3 represents the time period for which node 3 has their demand constraint satisfied. We seek to approximate the mean busy period at node 3 accurately. (A single value was computed for a similar model by the authors of this paper previously [10].)

For stability, we require $\mu > 40$. The busy period and fluid level moments can be computed analytically at nodes 1 and 2 using known results. Similarly, for $40 \leq \mu \leq 45$ the moments can be computed at node 3. For $\mu > 45$, the Boxma and Dumas result [5] does not apply, and indeed the computation of busy period moments in such a system is an open problem. We approximate the output of nodes 1 and 2 by continuous time Markov chain-modulated flows and then compute the busy period for a fluid queueing node subject to these two (approximate) inputs using a recent result [9].

Output from node 1 (and similarly node 2) is either at rate 45, or at rate 0. The time spent in the 0 state is exponentially distributed with parameter 1. The time spent in the 45 state is not exponentially distributed, so we approximate with a phase type distribution.

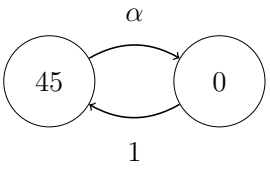


Figure 4: Two state model for output.

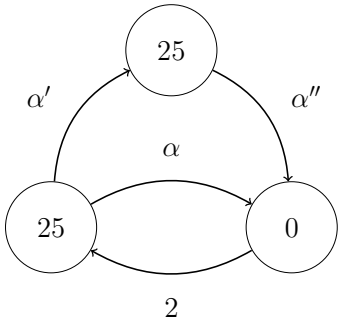


Figure 5: Three state model for output.

4.1 Two state approximation

Using a single exponentially distributed state to estimate the length of on period for output from node 1, we set $\alpha = \frac{45\times 5}{50} - 1$ to conserve average throughput of the node. Node 3 is then fed by two on/off sources. The input process to node 3 has four possible states: when both nodes are off, when only one is one, when only the other is on and when both are on. Computation of the busy period here can be done using the ideas in [9]. (Due to repeated eigenvalues which are described in [9], but have not yet been implemented we perturb the input rates to 45.1 and 44.9.)

4.2 Three state approximation

We fit a chain with state space as shown in Figure 5 (as in [7]). We choose the parameters α, α' and α'' to minimize the least squares difference between the approximation busy period moments and actual busy period moments which are known for nodes 1 and 2.

The input process to node 3 now has 9 states. The busy period moments can be computed for this process [9]. (Again, we perturb input rates to 45.1 and 44.9 to avoid repeated eigenvalues.)

4.3 Numerical results

In Figure 6 the mean busy period is plotted against μ. The solid line is the exact result from simulation (confidence intervals too small to display) and the dotted lines are the two and three state approximation values.

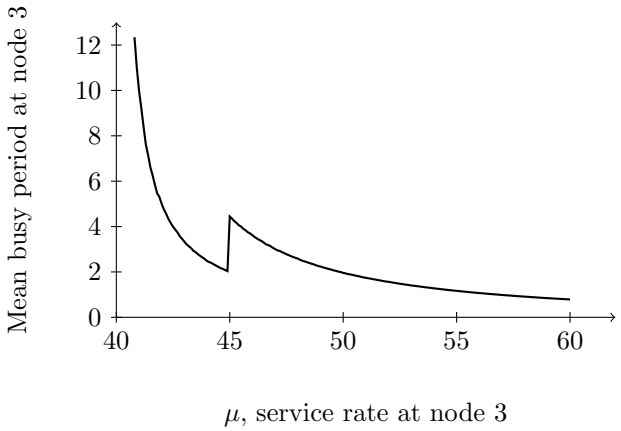


Figure 6: Mean busy period at node 3 in the network shown in Figure 3 for μ between 40 and 60. Simulation plot is the solid line. The two dashed lines are two and three state approximations.

The jump observed in the figure is the same as seen previously in Field and Harrison [9] and is caused by the number of net emptying/filling states changing from 3 to 1 as μ increases past 45.

5 Example 2

A more complicated network is shown in Figure 7. Three on/off sources feed the network. These are each have a capacity factor of 33%, typical of wind turbines (generally 20–40% [1]). These three independent sources feed in to local storage buffers at nodes 1, 2 and 3. Fluid from nodes 1 is sent to node 4 and fluid from node 3 to node 5. Fluid leaving node 2 is split $\frac{3}{4} : \frac{1}{4}$ between nodes 4 and 5. Fluid leaving nodes 4 and 5 is collected in node 6, where it is routed to a sink.

The rate at node 6 again represents the demand of consumers and the busy period the length of time for which this demand is met. We use our algorithm to approximate the busy period at node 6. We will ignore dependencies between flows in the network, so expect the approximation not to perform as well in this situation.

To demonstrate the effect that the dependencies have an independent network which has flows as the approximation sees them is shown in Figure 11 and results for this network shown on our graphs.

5.1 Stability

The variable parameter in the model shown in Figure 7 is μ, the service rate at node 5. The constraints on μ are that

- For node 6 to have a filling state we require that $\mu > 10$.
- For node 5 to have a filling state we require that $\mu < 22.5$.

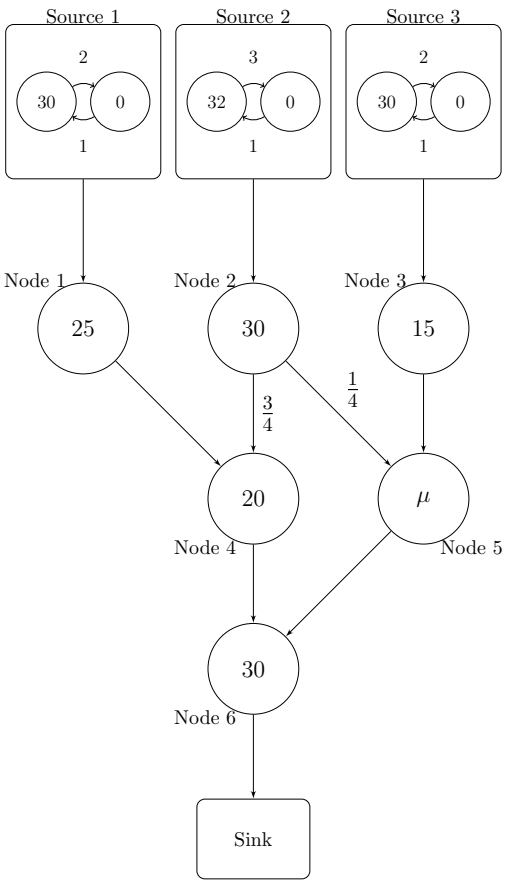


Figure 7: Topology of the second example network considered in this paper.

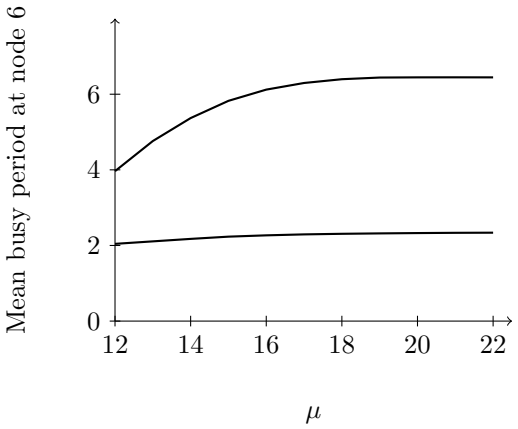


Figure 8: Mean busy period at node 6 for μ between 12 and 22.5. Simulation plots are the solid line. The upper line corresponds to the network in Figure 7, while the lower line corresponds to the independent version of the network as drawn in Figure 11.

- For node 5 to be stable we require that $\mu > 12$.

So the region of interest is $12 < \mu < 22.5$.

Again, there are no analytical results to compute the busy period moments at node 6 so we use our approximation algorithm.

5.2 Approximation algorithm

Computing the mean fluid level and busy period moments at the first three nodes is straightforward. Node 4 is of the class which Boxma and Dumas' result can handle [5], but both node 5 and 6 require the more recent Field and Harrison result [9] to compute. At node 5 either both or one of the sources feed at a rate faster than the service rate, while at node 6 no single source feeds at a rate higher than the service rate.

5.3 Numerical results

Our approximation method ignores dependencies between flows. In Figure 8 the thick solid line represents results for the actual network and the thinner solid line results for an independent version of the network. Our approximation agrees well with the independent version of the network, but the disagreement between the independent and true network shows that future work is needed to incorporate dependencies in to our algorithm.

In Figures 10 and 9 we plot the mean fluid levels at nodes 5 and 6 in the network to show how the approximation performs at approximating mean fluid level.

The average busy period against different values of μ is plotted in figure 8.

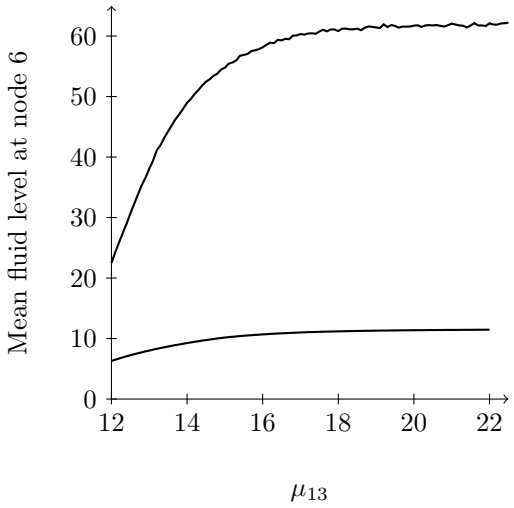


Figure 9: Mean fluid level at node 6 for μ between 12 and 22.5. Simulation plots are the solid lines. The upper line corresponds to the network in Figure 7, while the lower line corresponds to Figure 11.

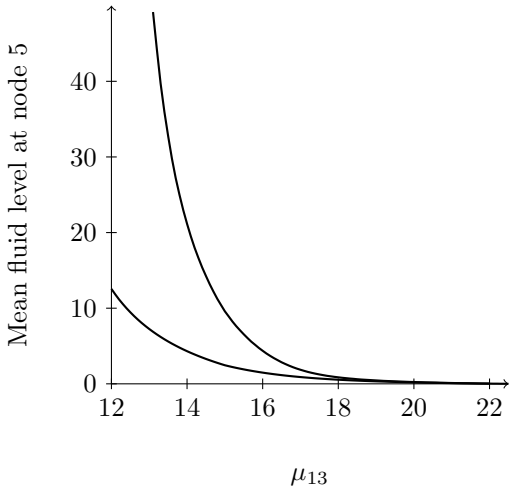


Figure 10: Mean fluid level at node 5 for μ between 12 and 22.5. Simulation plots are the solid lines. The upper line corresponds to Figure 7 and the lower line to Figure 11.

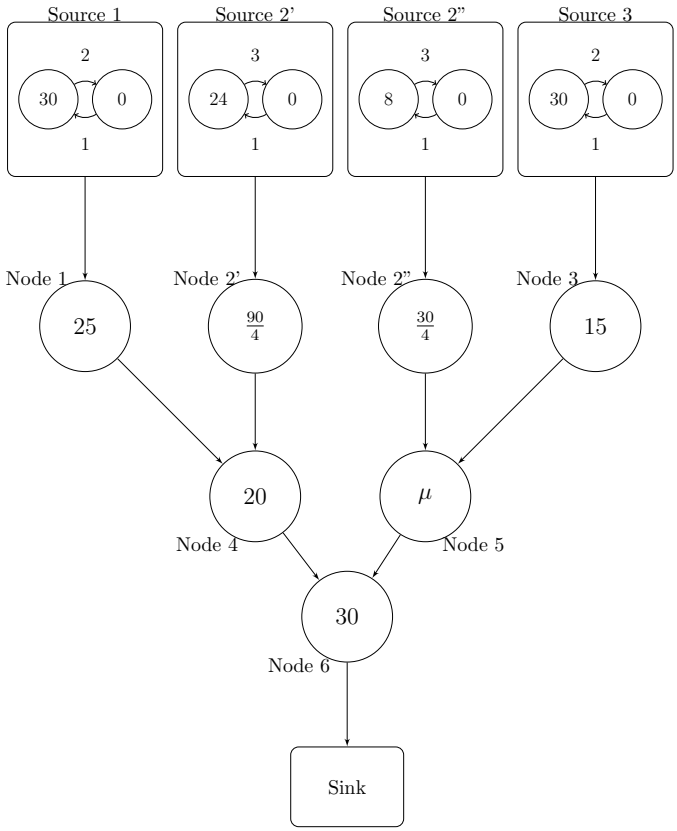


Figure 11: Altered topology of the second example network considered in this paper to show the network without dependencies. Our approximation performs well on this network.

6 Approximation method

In this section we explain how to apply our approximation method to a general fluid queue network (with acyclic directed topology). It is required that each node have some filling state (the buffer does not remain empty forever) and that at each node the throughput is less than the service rate (so each node is stable).

In a network of fluid queues, the input processes are modulated by Markov chains. However, the output from queueing nodes has non-exponentially distributed holding times. We approximate the state holding times with (one and two state) phase type distributions and assume all inputs are independent.

1. Input the following
   - A list of sources, nodes, sinks indexed as
   $$\underbrace{1, 2, \ldots, l}_{\text{sources}}, \underbrace{l+1, \ldots, m}_{\text{queueing nodes}}, \underbrace{m+1, \ldots, n}_{\text{sinks}}.$$
   - An n dimensional routing matrix P which describes the flows between sources, nodes and sinks. The routing matrix P must be upper triangular with zeroes on the diagonal. Elements p_{ij} of P describe the proportion of fluid leaving node i that is routed to node j.
   - For each Markov modulated source node $i = 1, \ldots, l$
     - o_i a vector of rates at which input arrives
     - Q_i a generator matrix determining the switching between different states
   - μ an $m-l$ length vector describing service rates at node with elements indexed $\mu_{l+1}, \mu_{l+2}, \ldots, \mu_m$.
2. For $i = l + 1$ to m:
   (a) Make a list of the input nodes to source i, called S_i.
   (b) Compute the combined inputs for node i by taking the independent product of the sources in S_i. Write U_i for the combined generator matrix and v_i for the vector of net input rates. These quantities are computed as the Kronecker sum of the Q_is and o_is as
   $$U_i = \bigoplus_{j \in \mathcal{S}_i} Q_j$$
   $$v_i = \text{Diagonal}\left(\bigoplus_{j \in \mathcal{S}_i} \text{DiagonalMatrix}(p_{ji} o_j) \right)$$
   (c) Compute the fluid and busy period moments at node i which serves at rate μ_i. (These results are exact for the given inputs.)
   (d) Approximate the output of the node using either a two state or a three state approximation. Write o_i for the output vector of node i and Q_i for the Markov chain which modulates this output. (Use one of the methods described below.)
3. Return the list of fluid level and busy period moments at each of the nodes.

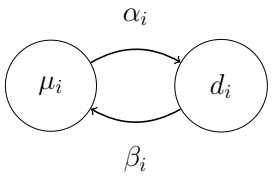


Figure 12: General two state model for output.

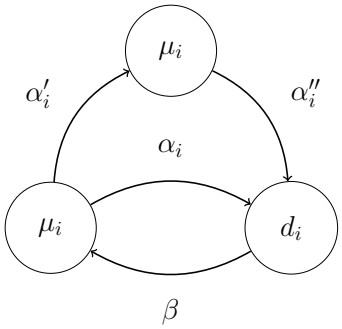


Figure 13: General three state model for output.

6.1 Two state output approximation

We fix the format of the output to be as in Figure 12 with rate μ_i and d_i. d_i is chosen to represent any flows that are below the service capacity of the node (including zero flow). In example 1, all non-full rate outputs were zero. In example 2, node 5 has input rates which are less than the service rate.

μ_i is already fixed for the model

d_i is chosen to describe the average non-full output rate. We compute the stationary distribution of the input chain π, select only emptying states (π^{E}) and normalise this vector, multiplying by the net output rate in each state (so $d_i = \frac{\pi^{\mathrm{E}} \cdot v_i^{\mathrm{E}}}{|\pi^{\mathrm{E}}|}$)

β_i is chosen

α_i is chosen to match throughput (we know the throughput of the node, and require the output $\frac{\mu_i\beta_i + d_i\alpha_i}{\alpha_i + \beta_i}$ match it)

6.2 Three state output approximation

In the 3-state approximation we set the output in the form shown in Figure 13.

μ_i is already fixed for the model

d_i is chosen

β_i is chosen

α_i α_i' and α_i'' are chosen (by least squares minimisation) to match moments of the busy period

7 Conclusions & Future Work

In this paper an improved approximation has been offered. The improvements offered in this paper are three fold

- using the recent Field and Harrison [9] result has allowed a much broader class of networks to be evaluated
- we have included a state to approximate small flow rates
- the result of [9] has been computed for significantly larger models than those first presented with the result.

The first of these is a particularly significant improvement. Networks where the previous approximation could be applied were a very limited class. For any node with more than two or three sources the inputs would need to have very short on periods for the node to remain stable.

Further improvements to the approximation are possible. We have neglected information about the correlation between sources and assumed all inputs to be independent. The authors of this paper have considered some results for fluid queues where the holding times states in filling states were not exponentially distributed [9]. Using results of this type with parameters chosen to accurately represent real data is an area for future work.

An important feature of power distribution we have not discussed in this paper is loss. While the traditional power distribution network is reasonably efficient (losses are about 7% [4]), energy storage is much less efficient. Pumped-hydroelectric storage, for example, is around 75% efficient [6]. Including these costs for buffer use is another area of future work.

We have not yet applied the approximation to a larger network (10+ nodes), but would expect the approximation of heavy tailed distributions by phase-type distributions at each node to give poor results in larger networks.

We want to include in our modelling

- Finite buffers as real storage can become full. We will need to use the busy period result of Lee [15] for finite buffers. And possibly extend that result [[how general is that result?]]. How to deal with excess fluid will be a problem in this model. Initially the easiest thing to do would be for excess fluid to overflow and be lost from the network. A smarter thing would be to re-route fluid if spare storage capacity remains elsewhere.
- Implement the ideas in [9] where $m_s \neq 1$ (because of symmetry in the network) and use them to re-work the results in this paper accurately.
- Buffers with thresholds where power is used at different rates [[is this realistic?]]
- Buffers with losses due to inefficiencies as most storage is only 75% efficient this is an important cost of storing energy

- Variable demand rates at different nodes as demand is very diurnal
- At the last stage (node 3 in the first example) numerically invert the transform of busy period to find the 95% point so we can meet service level agreeements in terms of percentage of busy period rather than just mean busy period.

8 Acknowledgements

This work was partly funded by the EPSRC grant *Fluid Approximations for Quantitative Analysis* EP/F048726/1.

References

[1] *Wind Power: Capacity Factor, Intermittency, and what happens when the wind doesn't blow?* University of Massachusetts, 2009.

[2] Packing some power. *The Economist*, 3 March 2012. `http://www.economist.com/node/21548495`.

[3] Renewable electricity futures study, June 2012. `http://www.nrel.gov/analysis/re_futures/`.

[4] U. E. I. Adminstration. How much electricity is lost in transmission and distribution in the united states?, June 2011.

[5] O. J. Boxma and V. Dumas. The busy period in the fluid queue. In *Proceedings of the 1998 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*, SIGMETRICS '98/PERFORMANCE '98, pages 100–110, New York, NY, USA, 1998. ACM.

[6] J. M. Carrasco, L. G. Franquelo, J. T. Bialasiewicz, E. Galván, R. C. P. Guisado, M. Ángeles, M. Prats, J. I. León, and N. Moreno-alfonso. Power-Electronic Systems for the Grid Integration of Renewable Energy Sources: A Survey. 53(4):1002–1016, 2006.

[7] A. J. Field and P. G. Harrison. An Approximate Compositional Approach to the Analysis of Fluid Queue Networks. *Performance Evaluation*, 64(9–12):1137–1152, October 2007.

[8] A. J. Field and P. G. Harrison. Approximate Analysis of a Network of Fluid Queues. In *Workshop on Mathematical performance Modeling and Analysis (MAMA), June 2007*, volume 35 of *Performance Evaluation Review*, pages 30–32, August 2007.

[9] T. Field and P. G. Harrison. Busy periods in fluid queues with multiple emptying input states. *Journal of Applied Probability*, 47:474–497, 2010.

[10] G. L. Jones and P. G. Harrison. Using Fluid Queues to Model Energy Storage and Distribution. In *Imperial College Energy and Performance Colloquium (ICEP)*, May 2012.

[11] O. Kella. Markov-modulated feedforward fluid networks. *Queueing Syst. Theory Appl.*, 37(1/3):141–161, Mar. 2001.

[12] O. Kella and W. Whitt. Useful martingales for stochastic storage processes with lévy input. *Journal of Applied Probability*, 29(2):396–403, June 1992.

[13] D. P. Kroese and W. R. W. Scheinhardt. Joint distributions for interacting fluid queues. *Queueing Syst.*, pages 99–139, 2001.

[14] V. G. Kulkarni. Fluid models for single buffer systems. *Frontiers in Queueing: Models and Applications in Science and Engineering*, pages 321–338, 1997.

[15] E. Y. Lee and K. K. J. Kinateder. The expected wet period of finite dam with exponential inputs. *Stochastic Processes and their Applications*, 90(1):175–180, Nov. 2000.

[16] L. Rabehasaina. Moments of a Markov-modulated, irreducible network of fluid queues. *Journal of Applied Probability*, 43(2):510–522, 2006.

[17] W. R. W. Scheinhardt and B. Zwart. A tandem fluid queue with gradual input. *Probab. Eng. Inf. Sci.*, 16(1):29–45, Jan. 2002.