Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

INCORPORATING TRUNCATED EXPONENTIAL DISTRIBUTIONS IN QUEUEING MODELS WITH ADJUSTABLE SERVICE-RATE CONTROL

Paul Babin

Research and Development
ThyssenKrupp Elevator Innovation GmbH
Essen, GERMANY

Allen Greenwood

Faculty of Engineering Management
Poznan University of Technology
Poznan, POLAND

ABSTRACT

Queueing models with service-rate control provide more realistic simulation results compared to simple M/M/1 models which have too much variability in the queue length. Because the variability in number in queue is increased by the unbounded nature of the exponential distribution, another approach modelers sometimes use is to select a bounded distribution, or to limit the maximum sample value by truncating the exponential distribution. This paper compares the beneficial effect of service-rate control and exponential distribution truncation. Simulation results demonstrate that models incorporating both mechanisms (truncated exponentials and service-rate control) generate the most realistic simulation results when rate adjustment is used to reduce queue-length peaks back to normal.

1 INTRODUCTION

Earlier papers have looked at how adjustable service-rate queueing models can give more realistic results than simple queueing models, especially with regard to the extreme variation observed in the number in queue for an M/M/1 system. The adjustable service-rate feature models the behavior where management adjusts the average service-rate, typically based on the current number in queue. This results in more stability in waiting times and number in queue. Another approach that modelers sometimes use to avoid this high variability of number in queue is to choose a bounded distribution so that extreme service times found in the long right tail of the exponential distribution do not affect the model results. But how much improvement in model stability is obtained by truncating the exponential distribution for the arrival and service times?

This paper studies the performance of an M/M/1 system with truncated tails (now a G/G/1). The performance of the queueing system, especially with regard to the number in queue variability, is studied with various amounts of truncation. A practical implementation in a simulation model is explained, including the need to adjust the mean based on the truncation. Results are compared to the results expected based on Kingman's approximation for the average number in queue for the G/G/1 model.

Then this paper compares the beneficial effect of truncating the distributions to the adjustable service-rate approach, and selects parameters of rate adjustment and of truncation which give similar results with respect to the amount of variability remaining in the observed number in queue. The benefits of including both truncated distributions and service-rate adjustment are demonstrated to give the most realistic results.

This paper provides practical insights for system modelers using discrete-event simulation, and incorporates the recommended practices of selecting appropriate distributions along with modeling the service-rate adjustment. These methods are both relatively straightforward with discrete-event simulation software and this paper provides some insight into their beneficial effects and implementation.

This paper is organized as follows. Section 2 provides some background on adjustable service-rate queueing simulation models and on the truncated exponential distribution. Section 3 uses Kingman's approximation to predict the performance of the queueing system when the arrival and service time distributions are truncated. Section 4 provides simulation results of scenarios selected to illustrate the impact of both mechanisms separately and together. Section 5 provides the conclusions and ideas for future research.

2 BACKGROUND

2.1 Adjustable Service-Rate Queueing Simulation Models

In many real-life systems, there is a mechanism to temporarily adjust the service rate based on the current number in queue. Simple queueing models, like the M/M/1, do not capture this system behavior properly, and lead to simulation results that are not realistic with regard to the extremes seen in the queue length, even if they exactly match the average performance.

With open-loop queueing models, e.g. an 80% utilization M/M/1 system, the number in queue sometimes grows well past the expected value of 3.2, reaching five or more times that value on occasion, before it drops to the average value, and then to zero where the system spends 20% of its time empty and idle.

In real systems, if the number in queue were ever to grow to this extreme, some action would be taken. The service rate would be temporarily adjusted to speed up processing, until the queue length got back to normal. Service-rate adjustment has been studied in the past as described by Hillier and Lieberman (1990). More recent applications include high performance computing (Wu et al. 2005), computer communications (Hellerstein et al. 2004), wireless communications (Ata 2005), dynamic control (George and Harrison 2001), and adaptive workflow modeling (Grabis 2014). A simple mechanism to adjust the service rate to meet some target number in queue was presented in earlier papers by Babin and Greenwood (2007, 2015a, 2015b).

For example, for an 80% utilized M/M/1 queue, the expected number in queue is 3.2. So if the target number in queue is set to 3 (the closest integer value), and "mild" service-rate adjustment is selected, as shown in Figure 1a, the nominal service rate is reduced by a factor of 0.83 when the queue drops to one, and increased by a factor of 1.5 when the queue rises to 9. Such an adjustable rate system will recover from extreme excursions much more quickly, and spend less time idle, as shown in Figures 1b and 1c.

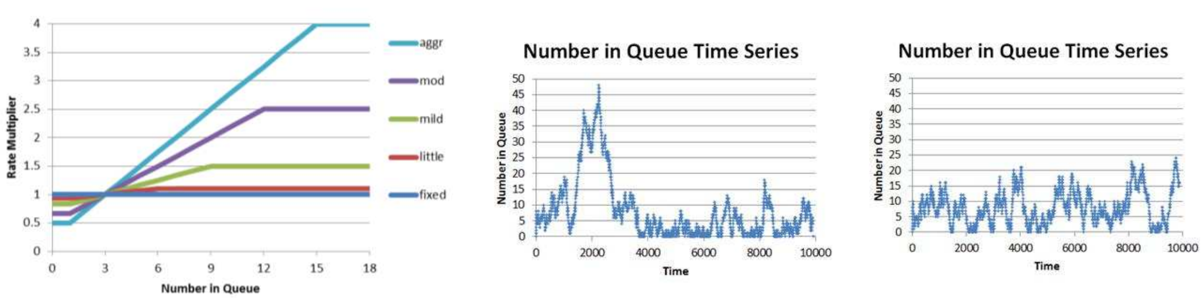

Figure 1: (a) Rate adjustments. (b) M/M/1 performance. (c) Adjustable rate.

2.2 Truncated Exponential Distribution

Truncated exponential distributions are sometimes used by system modelers when they wish to capture the variability in a process, but eliminate the extreme values associated with the infinitely long tail of the exponential distribution. In simulation modeling, including exponential service times will sometimes return unrealistically large values, which may make the replication seem incorrect. One solution is to sample from the exponential distribution and retain only the values that are below some threshold, say a truncation threshold equal to k times the mean value, or TT=k*(1/λ), as shown in Figure 2.

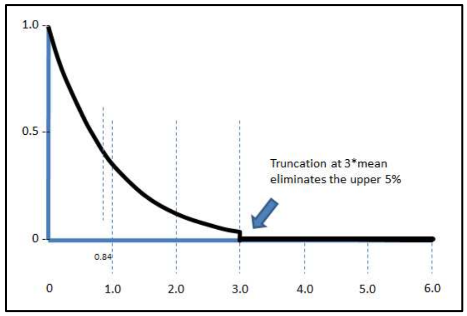

Figure 2: Truncating the exponential distribution.

But in the process of eliminating values from the upper tail, the mean and standard deviation are also somewhat affected. The adjustment to the mean of a truncated exponential distribution is given by Oliver (2015). Thus, if Y is distributed as an exponential random variable with a mean of 1/λ and is truncated at k times the mean, then the expected value of Y and Y^2 is given by

$$E(Y) = \frac{1}{\lambda}\left[\frac{1-(k+1)e^{-k}}{1-e^{-k}}\right] \quad \text{and} \quad E(Y^2) = \frac{2}{\lambda^2}\left[\frac{1-\frac{1}{2}(k^2+2k+2)e^{-k}}{1-e^{-k}}\right].$$

Therefore, using the relationship, $Var(Y) = E(Y^2)-(E(Y))^2$ the mean, standard deviation, and coefficient of variation (CV), at various levels of truncation (for λ = 1), are provided in Table 1.

Table 1: Mean, standard deviation and coefficient of variation for various truncation levels (k) for an exponential distribution with a mean of 1.

k	TT=k*mean	upper tail	E(Y)	Stdev(Y)	CV
2	2	0.1353	0.687	0.5253	0.7647
3	3	0.0498	0.843	0.7097	0.8421
4	4	0.0183	0.925	0.8342	0.9015
5	5	0.0067	0.966	0.9106	0.9426
6	6	0.0025	0.985	0.9541	0.9685
10	10	0.00005	1.000	0.9977	0.9982
100	100	3.72E-44	1	1.0000	1.0000

For example, from this table, if the exponential distribution with λ = 1 is truncated at 3 times its mean, i.e. k=3 and TT=3, then the upper tail consisting of the upper 5th percentile is removed. This has the effect of reducing the mean from 1 to 0.842 – not a small amount. That is why the truncation function should also adjust the samples to still achieve the desired mean value. Notice also that the truncation of k=3 reduces the standard deviation from 1 to 0.71, so the coefficient of variation is reduced from 1 to 0.8421. If the truncation threshold TT is set to a large number, like 100, then the effect disappears so the mean and standard deviation are virtually unchanged.

In Section 3, the effect of this reduction of variability will be predicted using Kingman's equation. Various amounts of truncation from 2 through 10 were also tested. The value of 3 times the mean is selected as a good truncation amount, since that value corresponds to truncating at the 95th percentile, and is not too severe. Some results with a truncation at 2 times the mean are included in Section 3 for comparison, although this severe amount of truncation is not recommended since the shape of the exponential distribution is not retained.

3 PREDICTING THE EFFECT USING KINGMAN'S EQUATION FOR G/G/1 MODELS

The well-known Kingman's equation is used to predict the average time in queue for G/G/1 models, as shown for example in Hopp and Spearman (2000). The average time in queue is a function of the variability, described by the coefficient of variation of the arrival and service times, the utilization, and the average service time.

$$CT_Q = \left(\frac{c_a^2 + c_e^2}{2}\right)\left(\frac{u}{1-u}\right)t_e$$

From the Kingman's estimate of the average cycle time in queue, the average number in queue can be estimated using Little's Law, $L=\lambda W$. Changing to a consistent notation,

$$W_Q = \left(\frac{c_a^2 + c_e^2}{2}\right)\left(\frac{\rho}{1-\rho}\right)\left(\frac{1}{\mu}\right) \quad \text{and} \quad L_Q = \lambda \cdot W_Q \quad \text{and} \quad \rho = \lambda/\mu \quad \text{then}$$

$$L_Q = \lambda W_Q = \left(\frac{c_a^2 + c_e^2}{2}\right)\left(\frac{\rho}{1-\rho}\right)\left(\frac{\lambda}{\mu}\right) = \left(\frac{c_a^2 + c_e^2}{2}\right)\left(\frac{\rho^2}{1-\rho}\right)$$

This approximation gives exact results for the M/M/1 model, where the $c_a^2=c_e^2=1$; i.e., the average number in queue $L_Q = \rho^2/(1-\rho)$. So from Table 1, with a truncation of k=3, the variability is reduced with $c_a^2=c_e^2=0.8421$ and the expected number in queue is reduced from 3.2 to 2.27. The expected values for 80% and 90% utilization scenarios, with truncation at k=2 and k=3 are shown later in Table 2.

While Kingman's approximation gives good results for the average number in queue for G/G/1 systems., further characterizations of the performance (beyond the mean) require higher moments of the distributions be considered. This is shown by Gross and Juttijudata (1997) where they explored the impact of higher moments (beyond mean and variance) of the input distribution on the performance of G/G/1 simulation models. Of course, truncation of the exponential distribution affects not only the mean and standard deviation, but also the higher moments such as skewness.

4 SIMULATION RESULTS

This section provides experimental results obtained from discrete-event simulation. First, the effect of various amounts of truncation is studied and is then compared to values predicted by Kingman's equation. Second, performance measures are compared for four scenarios - M/M/1, truncated exponential, mild rate adjustment, and the combination of a truncated exponential and mild rate adjustment. New performance measures are introduced in order to better illustrate the differences. From this section, it becomes apparent that there is still considerable variability prevalent, even when exponential distributions (interarrival and service times) are truncated. It is shown that service-rate control is more effective at adapting to the rising queue length. As expected, the combined model provides the most "realistic" results, in terms of the observed variability in the number in queue.

4.1 Comparing truncation in arrival times and service times

This section provides the simulation results for truncation amounts of k=2 and k=3 for arrival times, service times, and both. Results are compared to the M/M/1 with 80% utilization as a baseline. For these experiments, 10,000 arrivals are run in each of 30 replications per scenario. Because of the large amount of variability in these models, such large run lengths are recommended to clearly see the trend, although typical industrial applications may have much shorter runs as mentioned in Beaverstock et al. (2012).

The traditional performance measures are shown in Figures 3, 4 and 7, across seven scenarios. The first scenario, labelled "MM1 80 T A0 S0" is for an M/M/1 system with 80% utilization and no arrival or service time distribution truncation. The other scenarios consider truncation levels of the arrival and/or service distributions of either k=2 or k=3. When both the arrival and service distributions are truncated at k=3, as in scenario (T A3 S3), the average number in queue is reduced from 3.23 to 2.25. This is shown both in Figure 3 and Table 2.

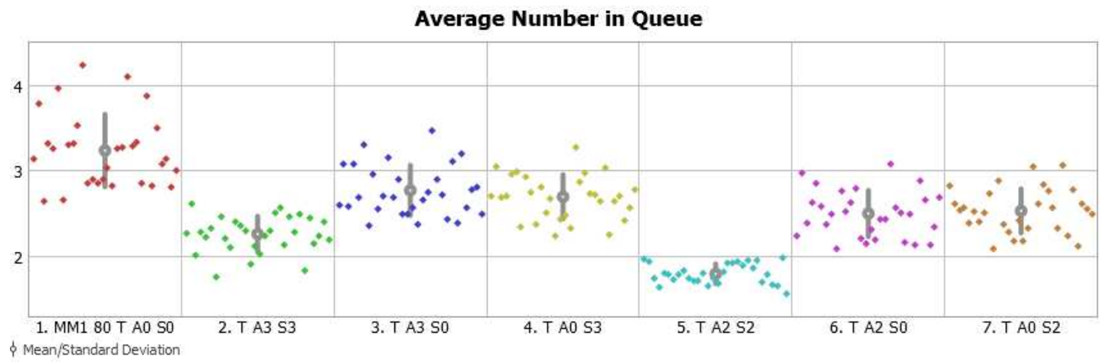

Figure 3: Comparing the average number in queue.

In Figure 4, the standard deviation of the number in queue is also reduced, in a similar manner, when truncation is applied to the arrival and service time distributions.

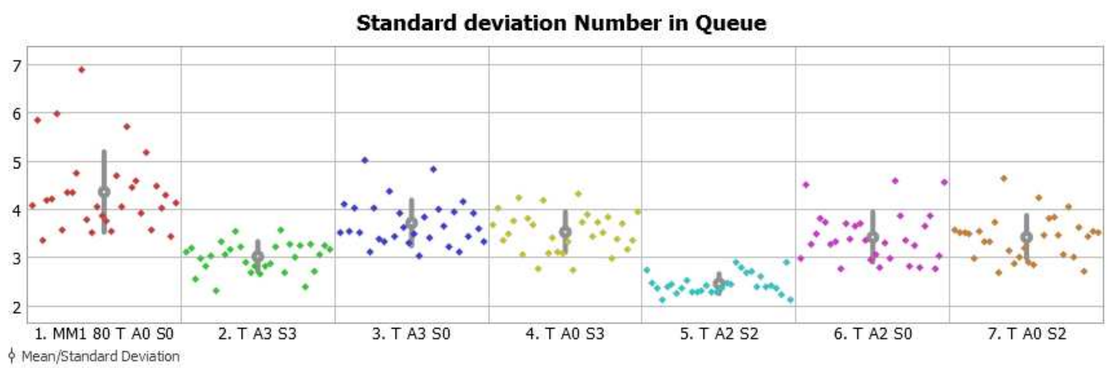

Figure 4: Comparing the standard deviation of the number in queue.

The average number in queue matches well with the expected number in queue predicted by Kingman's equation, taking into account the reduced coefficient of variability that results from the truncation, as shown in Table 2 and Figure 5. The effect of truncating the arrival distribution is similar to truncating the service distribution, and their combined truncation, is as predicted. Lowering the truncation threshold (which increases the amount of upper tail removed) from k=3 to k=2 reduces the average number in queue – just based on the reduction in variability. The mean arrival and service times are properly adjusted so they are not affected – as seen in Figure 7, where the utilization is consistent at the target 80%.

Table 2. Expected and Observed values for Average Number in Queue for 80% and 90% utilization, and arrival and service truncation levels of k=2 and k=3.

Scenario	$1/\lambda$	$1/\mu$	ρ	Ca	Ce	$V=(C_a^2+C_e^2)/2$	$U=\rho^2/(1-\rho)$	Predicted L_Q	Observed L_Q
80T A0 S0	10	8	80%	1	1	1.00	3.2	3.20	3.23
80T A3 S3	10	8	80%	0.8421	0.8421	0.71	3.2	2.27	2.25
80T A3 S0	10	8	80%	0.8421	1	0.85	3.2	2.73	2.76
80T A0 S3	10	8	80%	1	0.8421	0.85	3.2	2.73	2.68
80T A2 S2	10	8	80%	0.7647	0.7647	0.58	3.2	1.87	1.79
80T A2 S0	10	8	80%	0.7647	1	0.79	3.2	2.54	2.49
80T A0 S2	10	8	80%	1	0.7647	0.79	3.2	2.54	2.53
90T A0 S0	10	9	90%	1	1	1.00	8.1	8.10	7.80
90T A3 S3	10	9	90%	0.8421	0.8421	0.71	8.1	5.74	5.66
90T A3 S0	10	9	90%	0.8421	1	0.85	8.1	6.92	6.83
90T A0 S3	10	9	90%	1	0.8421	0.85	8.1	6.92	6.54
90T A2 S2	10	9	90%	0.7647	0.7647	0.58	8.1	4.74	4.41
90T A2 S0	10	9	90%	0.7647	1	0.79	8.1	6.42	6.19
90T A0 S2	10	9	90%	1	0.7647	0.79	8.1	6.42	6.18

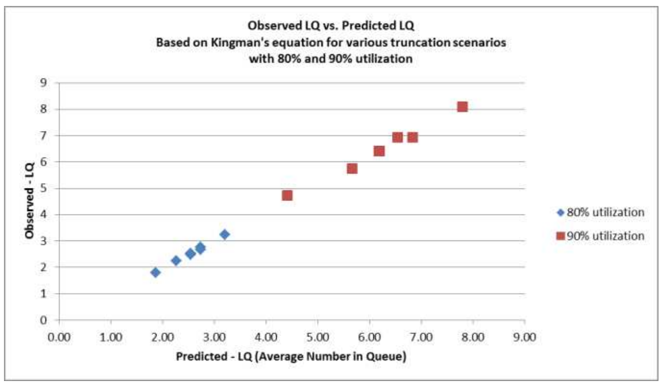

Figure 5: Average number in queue – observed vs. expected.

As predicted by Kingman's equation, the effect of decreasing the truncation threshold is to decrease the variability component. Because the variability term is multiplied by the utilization term to form the expected number in queue, there is an interaction effect of the truncation and the utilization. With higher utilizations, the truncation has a larger effect on the average number in queue, as shown in Figure 6. Utilization has a similar effect for the observed standard deviation of the number in queue from the simulation experiments.

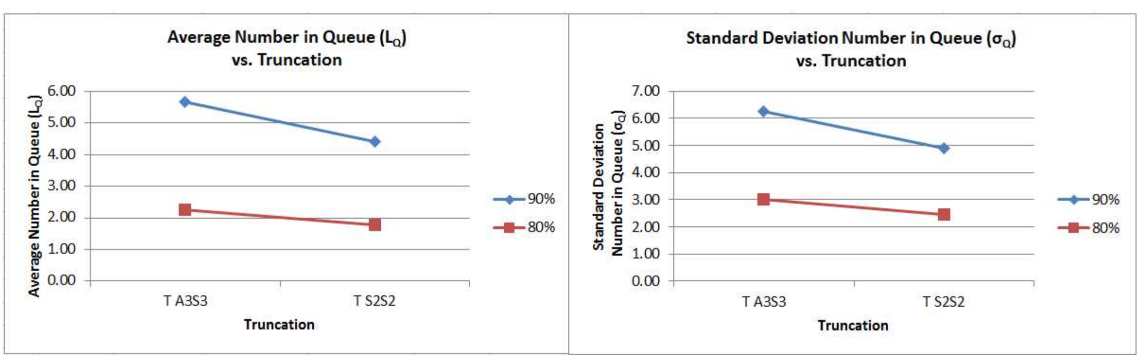

Figure 6: Interaction Plot – the effect of truncation on average number in queue and on the standard deviation of the number in queue depends on the utilization.

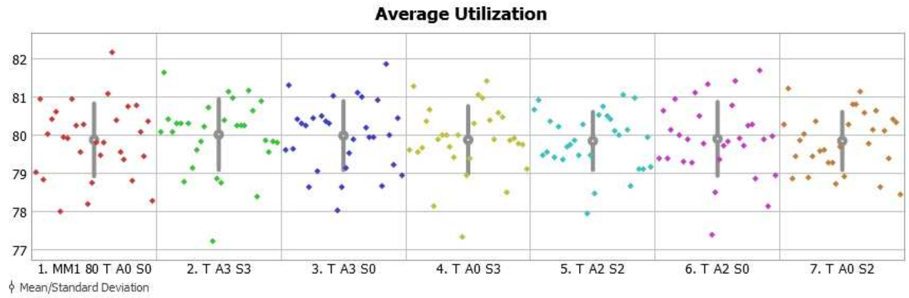

Figure 7: Comparing server utilization.

4.2 Comparing truncated arrival & service with adjustable service rate

This section provides results of the simulation experiments for four selected scenarios. A baseline is established using an M/M/1 with 80% utilization. The arrival and service time distributions are truncated at three times the mean (k=3) forming an *M'/M'/1* (or G/G/1) model. The adjustable service-rate model is then applied to the original un-truncated system, with a "mild" amount of rate adjustment. Finally, the combination of truncation (k=3) for both arrival and service time distributions and the incorporation of a "mild" service-rate adjustment is evaluated.

4.2.1 Number in Queue time series for a single simulation run – for illustration

First, the four selected scenarios are compared using a time-series plot of the number in queue for a single short simulation run of around 250 arrivals (before showing the more comprehensive result with 10,000 arrivals and 30 replications of each scenario in the next section). Figure 8 shows a representative example of what can happen with each scenario.

In the M/M/1 80% utilization panel, notice that even though the expected number in queue is 3.2, sometimes the queue grows to 15, almost 5 times the expected value, before eventually dropping back through the mean and to empty and idle. Notice also that it takes a long time to recover on its own.

When mild rate adjustment is added, the same arrival and service time variability is visible, but the peaks are less extreme, and they recover much more quickly. This shows the essence of the adjustable rate mechanism – if the number in queue exceeds the target value, the nominal service rate is adjusted to bring the queue length down more quickly.

When truncation is applied alone, the peaks are less severe than without truncation, but still present.

Finally, when both rate adjustment and truncation are applied, the peak events are less severe and recover more quickly; however, the variability of the number in queue is still present. Therefore, this combination is the most realistic.

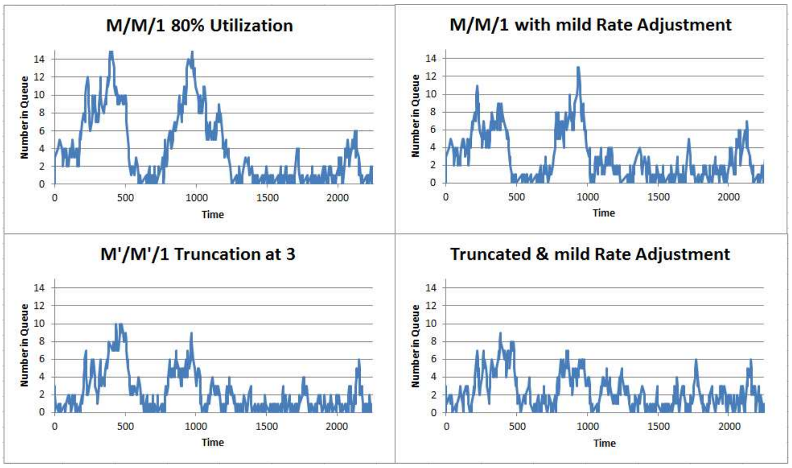

Figure 8: Number in queue time series, for four comparison scenarios.

4.2.2 Comparing scenarios with multiple replications

To draw valid conclusions, a larger run length (10,000 arrivals) and multiple replications (30) were used in the FlexSim experimenter to analyze and compare the four scenarios. Figure 9 shows that the average number in queue is reduced with the truncation approach, as expected. Similarly, in Figure 10, the standard deviation is also reduced for the truncation and rate variation approaches.

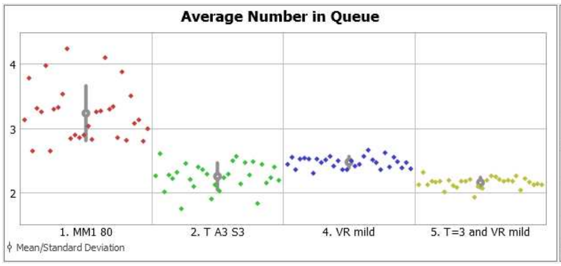

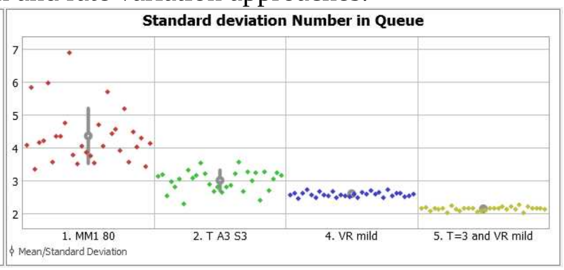

Figure 9: Average number in queue. Figure 10: Standard deviation of the number in queue.

In Figure 11, notice that the server utilization is unaffected when only truncation is employed – it stays at 80% since the mean of both the arrival and service distributions are adjusted to retain the target ratio. When rate adjustment is employed, the utilization is increased, since the server can also slow down a little when the queue is nearly empty. This is easier to see in these results than in the earlier time series, and illustrates another benefit of rate adjustment. Just as the service rate can be increased when the queue rises (to prevent extreme conditions), it can also be slowed to keep the system more stable.

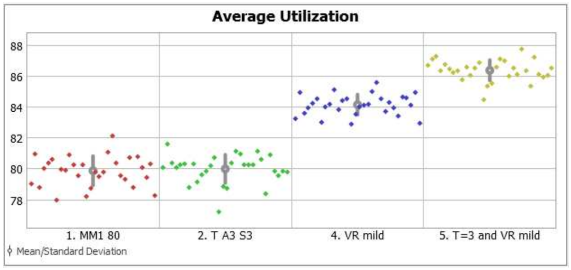

Figure 11: Comparing server utilization.

4.2.3 New performance measures to quantify Peak Excursion behavior

In order to detect the differences, a new set of performance measures is needed. In Figure 12 the number in queue time series is shown along with an illustration of how "peak events" are characterized.

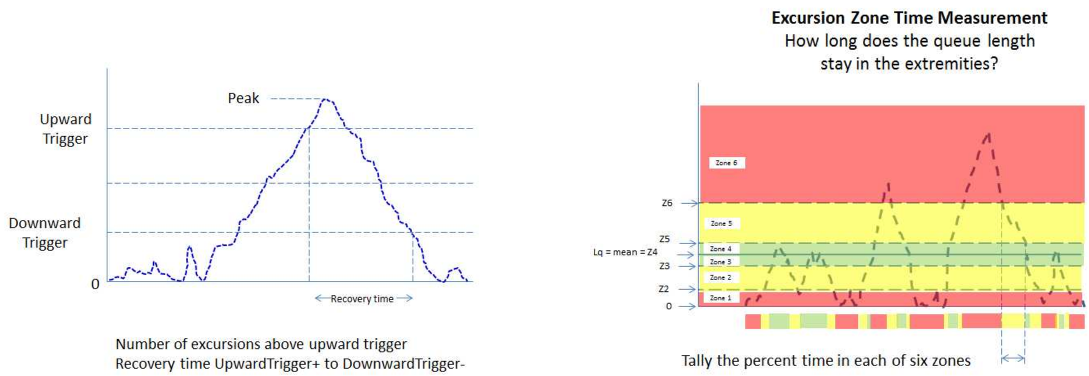

Figure 12: Defining the "peak excursion" performance measures.

To characterize the behavior of the number in queue time series, consider a threshold, an upward trigger, which indicates a 'danger zone' of high queue levels. If the number in queue rises past this trigger level, e.g., set to 5 times the expected number in queue, then a peak event is declared. Of course, the queue will eventually drop back down to its expected value on its own – and this time is considered the "peak recovery time." For the adjustable-rate models, the queue will recover much more quickly down to reasonable levels. Therefore, the additional performance measures considered here are:

- Maximum queue contents. Tracking the maximum observed queue length during a run is highly dependent on the run length, and not an especially good performance measure, but it does give some interesting comparisons.
- Number of peak trigger events. A count of the number of times the queue length ventures above the trigger.
- Average peak recovery time. The average time it takes to recover from peaks, measured from the start of the upward trigger (in this case, above 5 times the mean) to when it drops back to the downward trigger (here set to the expected value).
- Time above the peak trigger limit. (Zone 6 in Figure 12). The percentage of time that the queue length is above the specified threshold.

In Figure 13, notice that for the M/M/1 80% utilization scenario, the queue length rose to maximum values between 21 and 45. Truncating at k=3 reduces the maximum value somewhat, but the average is still above 20. It is interesting to note that even when the maximum service time is truncated at k=3, extreme queue conditions still occur. If a flurry of arrivals occur together, and there is no mechanism to speed up the server to react to that condition, extremes are still seen. Figure 14 counts the number of distinct occurrences of trigger events. Both of these figures indicate that adding truncation along with rate adjustment provide the best reduction of the peak events, to make the model more realistic.

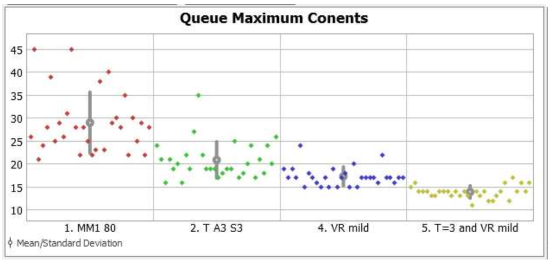

Figure 13: Maximum observed number in queue.

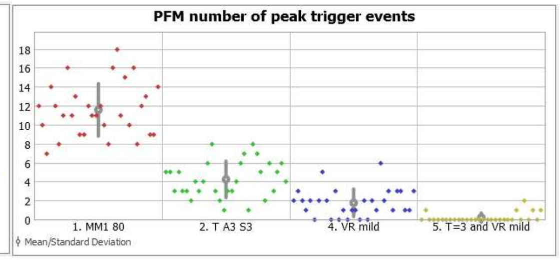

Figure 14: Number of peak trigger events.

In Figure 15 notice that the amount of time required to recover from the peak events is much improved with the rate adjustment scenarios, as expected. Finally, in Figure 16, the percentage of time that the queue spends in the zone above the peak limit is reduced from 3% to nearly zero. With this peak threshold set at 15 or nearly 5 times the expected value of 3.2, the M/M/1 with 80% utilization still spends 3% of its time above this peak. Simply truncating at k=3, the queue still spends an unrealistic amount of time at these extremely high levels.

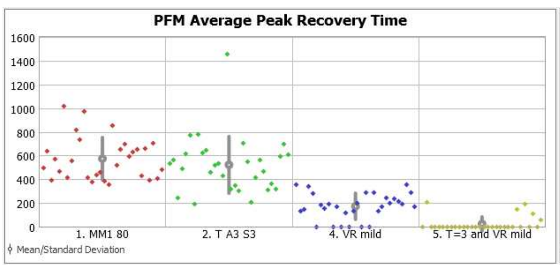

Figure 15: Average peak recovery time.

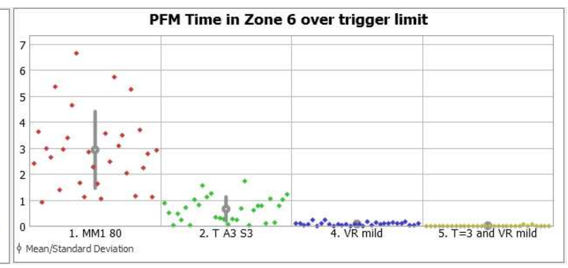

Figure 16: Percentage of time over the trigger limit.

5 CONCLUSIONS AND FUTURE RESEARCH

An important goal of business process modeling is to capture the behavior of a management system, such as adapting to run-time conditions like the work backlog by adjusting the service rate. This paper shows the benefit of incorporating a service-rate adjustment mechanism into queuing simulation models, even when using truncated exponential distribution to limit the extreme service times. Simply truncating the distribution still leaves the system with unrealistic peaks, from management's perspective, of the number of items in queue. Only when rate adjustment is incorporated into the model, does the simulation capture the adaptive behavior of the system – allowing a certain amount of inherent variability, but reacting to the peaks to drive the number in queue down quickly by adjusting the service rate.

For future research, other popular distributions such as the beta or triangular, which are inherently bounded, could also be compared to the exponential and the truncated exponential. More complex rate adjustment mechanisms need to be considered. Rather than taking control feedback based only on the target and current number in queue, the accumulated error from the target and from the rate of change of that error can be used to formulate the rate-adjustment function. Adding this integral term to the controller will drive the steady-state error to zero. Some simple PID controller designs should be evaluated in order to determine how effectively they can be included in simulation models such as in Hellerstein et al. (2004). Similarly, an exponentially weighted moving average (EWMA) could provide a realistic feedback source, as suggested by Box, Luceno and Carmen (2009).

A APPENDIX – TRUNCATED EXPONENTIAL FUNCTION IN FLEXSIM

This appendix provides the truncation function implemented in FlexSim. It illustrates one approach to truncating the exponential distribution while retaining the desired mean values.

```
/**Custom Code*/
// This custom command returns a value from a truncated exponential distribution
// It adjusts the mean based on the truncation, to achieve the correct mean
// However the maximum values will be above the requested threshold
//
double sample;
double desiredMean = parval(1);
double kThreshold = parval(2);
double thresholdValue = kThreshold*desiredMean;
int stream = parval(3);

if (kThreshold > 1)  // prevent lockup with threshold zero, or illogically below
the mean
{
double  truncatedMean  =  desiredMean*(1-(kThreshold+1)*(exp(-kThreshold)))/(1-exp(-
kThreshold));

    sample = exponential(0,desiredMean,stream);
        while (sample > thresholdValue)
        {
            sample = exponential(0,desiredMean,stream);
        }
    sample = sample*desiredMean/truncatedMean; // adjust the sampled value by the
ratio of the truncated to desired means. 12Jan2015
}
else {
    sample = exponential(0,desiredMean,stream);
}
return sample;
```

REFERENCES

Ata, B. 2005. “Dynamic Power Control in a Wireless Static Channel Subject to a Quality-of-Service Constraint.” Operations Research, 53(5), 842-851.

Babin, P., M. Hamilton, and A. Greenwood. 2007. “Modeling Adjustable Rate Processing in Queueing Simulations.” In Proceedings of the Industrial Engineering Research Conference, 1465-1470.

Babin, P., and A. Greenwood. 2015. “Applying Variance Performance Measures for Adjustable Service-Rate Queueing Simulation Models.” In Proceedings of the Industrial Engineering Operations Management Conference.

Babin, P., and A. Greenwood. 2015b. “Applying Cost Measures for Adjustable Service-Rate Queueing Simulation Models.” In Proceedings of the Industrial Engineering Research Conference.

Beaverstock, M., A. Greenwood, E. Lavery, and W. Nordgren. 2012. Applied Simulation: Modeling and Analysis Using FlexSim. Orem, Utah: FlexSim Software Products, Inc.

Box, G. E. P., A. Luceno, and M. Carmen. 2009. Statistical Control by Monitoring and Adjustment, 2nd ed. Hoboken, N.J.: John Wiley & Sons, Inc.

George, J., and J. Harrison. 2001. “Dynamic Control of a Queue with Adjustable Service Rate.” Operations Research, 49:720-731.

Grabis, J. 2014. “Application of Predictive Simulation in Development of Adaptive Workflows.” In Proceedings of the 2014 Winter Simulation Conference, edited by A. Tolk, S. Y. Diallo, I. O.

Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 996-1003. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Gross, D., and M. Juttijudata. 1997. "Sensitivity of Output Performance Measures to Input Distrubutions in Queueing Simulation Modeling." In Proceedings of the 1997 Winter Simulation Conference, edited by S. Andradottir, K. J. Healy, D. H. Withers, and B.L. Nelson, 296-302. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Hellerstein, J., Y. Diao, S. Parekh, and D. Tilbury. 2004. Feedback Control of Computing Systems. Hoboken, N.J.: John Wiley& Sons, Inc.

Hopp, W., and M. Spearman, Factory Physics. 2000. 2nd ed. New York, N.Y.: McGraw Hill, p304.

Hillier, F., and G. Lieberman. 1990. Introduction to Stochastic Models in Operations Research, New York, N.Y.: McGraw Hill.

Olive, D. 2015. Chapter 4 Truncated Distributions. Accessed March 22, 2015. http://lagrange.math.siu.edu/Olive/Ch4.pdf.

Wu, Q., P. Juang, M. Martonosi, and D. Clark. 2005. "Voltage and Frequency Control with Adaptive Reaction Time in Multiple-Clock-Domain Processors." In Proceeding of the 11th International Symposium on High-Performance Computer Architecture.

AUTHOR BIOGRAPHIES

PAUL BABIN is the Global Manager of Test & Qualification for ThyssenKrupp Elevator in Essen, Germany. He is also a Ph.D. student at Mississippi State University in Industrial Engineering where his research interest is in discrete event simulation of controlled queueing systems. He earned a B.S. in Electrical Engineering from Christian Brothers University in Memphis, an M.S. in Electrical Engineering from the University of Tennessee Knoxville, an M.S. in Industrial and Systems Engineering from the University of Memphis, and an M.B.A. from Christian Brothers University in Memphis. He is a member of ASQ, IEEE, and IIE. His email address is Paul.Babin@ThyssenKrupp.com.

ALLEN GREENWOOD is a Professor in the Faculty of Engineering Management at Poznan University of Technology in Poznan, Poland. Previously, for the past 20 years, he was Professor of Industrial and Systems Engineering at Mississippi State University. He received his B.S.I.E from North Carolina State University, M.S.I.E. from the University of Tennessee, and Ph.D. from Virginia Tech. He is a registered Professional Engineer. His teaching and research interests include: simulation modeling, analysis, and optimization; project management; and, process design in production and logistics systems. He is co-author of the textbook Applied Simulation Modeling and Analysis Using FlexSim. He is a member of IIE, ASEE, PMI, and INFORMS. His email address is greenwood@ise.msstate.edu.