INFORMS Journal on Computing

Online Supplement for "Developing Effective Service Policies for Multiclass Queues with Abandonment: Asymptotic Optimality and Approximate Policy Improvement"

Terry James
STOR-i Doctoral Training Centre, Lancaster University, t.james@lancaster.ac.uk

Kevin Glazebrook
Department of Management Science, Lancaster University, k.glazebrook@lancaster.ac.uk

Kyle Lin
Operations Research Department, Naval Postgraduate School, kylin@nps.edu

---

Appendix A: Proof of Theorem 1

Consider the k-class queueing system under stationary nonidling service control policy π under an assumption that $\sum_{j=1}^{k} \lambda_j/\mu_j < 1$ in which the system is stable under nonidling policies in the absence of abandonments. We shall consider the system in the limit as $\theta \to 0$.

Under policy π and with abandonment parameter θ write $W^{\pi}_{j,\theta}$ for the waiting time of a class j customer in steady state, to be understood as follows: $W^{\pi}_{j,\theta}$ is the time needed for a class j customer arriving at the system in steady state and with zero personal abandonment rate to complete its service. An arriving class j customer in steady state will actually complete its service if this waiting time is no greater than $Y_j \backsim \exp(\theta_j)$, the time available to the customer in the system prior to her abandonment. The quantity $E[\exp(-\theta_j W^{\pi}_{j,\theta})] = E[P(W^{\pi}_{j,\theta} < Y_j)]$ is the long run proportion of class j customers who achieve service completion under policy π. We can now write the reward rate achieved under π as

$$R^{\pi}(\theta) = \sum_{j=1}^{k} \lambda_j R_j E[\exp(-\theta \nu_j W^{\pi}_{j,\theta})]. \tag{1}$$

Now consider any *priority policy* ϖ, namely any policy which operates a fixed priority ordering among the customer classes. The $R\mu\theta$ and $R\mu$ rules are such policies. We shall assume without loss of generality that ϖ chooses individual customers from the chosen class for service in a first-come-first-served fashion. Now write W^{ϖ}_{j} for the waiting time (time to achieve completed service) of a class j job in steady state under priority policy ϖ for the no abandonment case with $\theta = 0$. It is clear from a simple argument based on realisations of the system that the total workload (uncompleted service) in the system from customers in the l classes which have top priority under ϖ (for any $1 \le l \le k$) when $\theta = 0$ stochastically bounds above the corresponding quantity when $\theta > 0$. Since epochs at which any class j job enters service are those at which the workload from higher priority classes is zero, it follows straightforwardly that for priority policies ϖ we have $W_j^{\varpi} \ge_{\text{ST}} W_{j,\theta}^{\varpi}$, and hence from (1) that

$$\begin{aligned} R^{\varpi}(\theta) &\ge \sum_{j=1}^{k} \lambda_j R_j E[\exp(-\theta \nu_j W_j^{\varpi})] \\ &= \sum_{j=1}^{k} \lambda_j R_j - \theta \sum_{j=1}^{k} \lambda_j R_j \nu_j E[W_j^{\varpi}] + O\left(\theta^2\right). \end{aligned} \tag{2}$$

In (2), we have established a lower bound for the reward rate $R^{\varpi}(\theta)$ for any priority policy ϖ. We now develop an upper bound for $R^{\pi}(\theta)$ for any π. To achieve this, we consider first a realisation of the system under nonidling policy π and with no abandonments ($\theta = 0$). This realisation will be determined by π and a given set of arrival times $\mathbf{A}$ and service durations $\mathbf{S}$. We write the lengths of successive busy periods for this realisation as B_n, and the number of customers served in successive busy periods as $M_n, n \in \mathbb{N}$. Write l_{nj} for the number of class j customers in period n, and $W_{jl}^{\pi(n)}$ for their waiting times, for $1 \le l \le l_{nj}, 1 \le j \le k$, where $\sum_j l_{nj} = M_n$.

We now apply abandonment to this realisation. Hence we consider the stochastic process generated when the realisation determined by $\pi, \mathbf{A}$, and $\mathbf{S}$ is modified by random abandonments with class-specific rates $\theta \nu_j$, $1 \le j \le k$. It is trivial that at all epochs at which the system is empty for the no abandonment realisation, it will also be empty when abandonments are applied. Expressed differently, the busy periods for the process without abandonments contain (one or more) busy periods for any generated process with abandonments. When abandonments are applied to the realisation with waiting times $W_{jl}^{\pi(n)}, 1 \le l \le l_{nj}, 1 \le j \le k, n \in \mathbb{N}$, then it is easy to show that the probability that none of the customers served in busy period n is abandoned is bounded below by $\exp(-\theta \nu^* B_n M_n) \ge 1 - \theta \nu^* B_n M_n$ where $\nu^* = \max_j \nu_j$.

Use W_{j-}^{π} for the collection of waiting times for class j customers of the non-abandonment realisation, namely $W_{j-}^{\pi} := \{W_{jl}^{\pi(n)}, 1 \le l \le l_{nj}, n \in \mathbb{N}\}$. We now seek a lower bound for the conditional expectation $E[W_{j,\theta}^{\pi} | W_{j-}^{\pi}]$, which is the mean class j waiting time (after abandonments) conditional on this no-abandonment realisation. To compute this conditional expectation we use $\mathbf{X}^{(n)}$ for the collection of exponential random variables which determine the available lifetimes (deadlines) for the customers concerned with busy period n. We write $\{W_{jl}^{\pi(n)}\left(\mathbf{X}^{(n)}\right), 1 \le l \le l_{nj}, n \in \mathbb{N}\}$ for the new (random) class j waiting times which result from the abandonment process when applied to successive busy periods of the non-abandonment process. By the above argument

$$P(W_{jl}^{\pi(n)}(\mathbf{X}^{(n)}) = W_{jl}^{\pi(n)}, 1 \le l \le l_{nj}) \ge \exp(-\theta \nu^* B_n M_n)$$

from which it follows that

$$E[W_{j,\theta}^{\pi} | W_{j-}^{\pi}] = \lim_{N \to \infty} \frac{\sum_{n=1}^{N} \sum_{l=1}^{l_{nj}} E[W_{jl}^{\pi(n)}(\mathbf{X}^{(n)})]}{\sum_{n=1}^{N} \sum_{l=1}^{l_{nj}} l_{nj}}$$

$$
\begin{aligned}
&\geq \lim_{N\to\infty} \frac{\sum_{n=1}^{N}\sum_{l=1}^{l_{nj}} W_{jl}^{\pi(n)} \exp\left(-\theta\nu^* B_n M_n\right)}{\sum_{n=1}^{N}\sum_{l=1}^{l_{nj}} l_{nj}} \\
&\geq \lim_{N\to\infty} \frac{\sum_{n=1}^{N}\sum_{l=1}^{l_{nj}} W_{jl}^{\pi(n)}}{\sum_{n=1}^{N}\sum_{l=1}^{l_{nj}} l_{nj}} \left(1-O\left(\theta\right)\right) \\
&= E[W_j^\pi] - O(\theta),
\end{aligned}
$$

where the equality follows from the ergodicity of the system, and the lower bound $E[W_j^\pi]-O(\theta)$ is a uniform one that does not depend on W_{j-}^π. Unconditioning, we infer that

$$E[W_{j,\theta}^\pi] \geq E[W_j^\pi] - O(\theta).$$

Combining this with (1), we have that

$$R^\pi(\theta) \leq \sum_{j=1}^{k}\lambda_j R_j - \theta\sum_{j=1}^{k}\lambda_j R_j \nu_j E[W_j^\pi] + O(\theta^2). \tag{3}$$

Please note that in (2) and (3), all the $O\left(\theta^2\right)$ terms, upon division by θ^2 involve expectations which are uniformly bounded as ϖ, π range over their respective policy classes.

Applying (3), we infer that

$$\max_\pi \left(\sum_{j=1}^{k}\lambda_j R_j - \theta\sum_{j=1}^{k}\lambda_j R_j \nu_j E[W_j^\pi] + O\left(\theta^2\right)\right) \geq \max_\pi R^\pi\left(\theta\right), \tag{4}$$

where the maxima in (4) are over all policies π. Making the $R\mu\theta$-rule the choice of priority policy ϖ in (2), using the fact that $\max_\pi R^\pi(\theta) \geq R^{R\mu\theta}(\theta)$, and using (4) it follows that

$$\max_\pi R^\pi\left(\theta\right) - R^{R\mu\theta}\left(\theta\right) \leq \theta\left\{\sum_{j=1}^{k}\lambda_j R_j\nu_j E[W_j^{R\mu\theta}] - \min_\pi \sum_{j=1}^{k}\lambda_j R_j \nu_j E[W_j^\pi]\right\} + O\left(\theta^2\right). \tag{5}$$

We finally observe that the minimisation in (5) can alternatively be written, using Little's Law, as

$$\min_\pi \sum_{j=1}^{k} R_j\nu_j E[N_j^\pi], \tag{6}$$

with N_j^π the number of class j customers in the system (without abandonments) in steady state under policy π. The minimsation in (6) is of a holding cost rate for the system, with cost $R_j\nu_j$ incurred per class j customer and per unit of time. A classical queueing control result (the $c\mu$-rule) asserts that for the no abandonment system, this holding cost rate is minimised by the $R\mu\nu$-rule which provides service according to a priority policy with class ordering determined by (decreasing) values of $R_j\mu_j\nu_j$. Upon multiplication of these values by θ it is clear that this is our $R\mu\theta$-rule. We infer from this fact and from (5) that

$$\max_\pi R^\pi(\theta) - R^{R\mu\theta}(\theta) \leq O(\theta^2)$$

as required, which concludes the proof.

Appendix B: Extensions to Theorem 1

Section 3.1 in the main paper makes the point that the proof of Theorem 1 makes little use of the stochastic structure of the service mechanism and that extensions of Theorem 1 are available to other systems for which a priority policy optimises a holding-cost type objective in the absence of abandonments. We now give some examples.

Consider a multiclass $M/GI/1$ queueing network with Bernoulli feedback, known as a *Klimov Network*; see Klimov (1974) and Klimov (1978). Exogenous arrivals to the system form independent Poisson streams, with λ_i the rate for class $i, 1 \leq i \leq k$. With each class i is associated a collection J_i of service stations with $S_{ij} \backsim G_{ij}$ a generic class i service time at station $ij, 1 \leq j \leq J_i$. All service times are mutually independent and are assumed to have finite second moment. Each class i customer begins service at station $i1$ and is thereafter routed for further service according to the Markovian routing matrix P^i or exits the system. Hence the sequence of stations visited by each class i customer forms a Markov Chain with departure from the system represented by entry into an absorbing state. A single server is available to provide service at all service stations, namely those in the collection $\cup_{1\leq i\leq k} \cup_{j\in J_i} \{ij\}$. This service is provided nonpreemptively in the case of general service times, which is the case we now consider.

We write $S_i \backsim G_i$ for the *total* service requirement of a class i customer, namely the aggregate of all individual service times until the system is exited. We suppose that the $\sum_{i=1}^{k} \lambda_i E[S_i] < 1$ and hence that the system is stable under nonidling service. If we write N_i for the total number of class i customers present in the system in steady state then a holding cost objective $E[\sum_{i=1}^{k} C_i N_i]$ is minimised by a service policy which imposes a priority ordering $KR(\mathbf{C})$ among the stations, where $\mathbf{C} = (C_1, C_2, \ldots, C_k)$. See Klimov (1974) and Klimov (1978) for details.

We modify the above Klimov Network by imposing customer abandonment. Hence all class i customers have their sojourn in the system terminated at a time after entry which has an exponential distribution with rate $\theta_i = \theta\nu_i$, unless the customer has already exited the system upon completion of all service. We suppose that each class i customer who completes all service prior to abandonment earns a reward R_i. As before we write $R^{\pi}(\theta)$ for the reward rate achieved under service policy π. We write $\mathbf{R}\theta = (R_1\theta_1, R_2\theta_2, \ldots, R_k\theta_k)$ and $KR(\mathbf{R}\theta)$ for the Klimov ordering determined by $\mathbf{R}\theta$. The proof of the following result is in all essentials unchanged from that of Theorem 1.

COROLLARY 1. *For the above Klimov Network, if* $\sum_{i=1}^{k} \lambda_i E[S_i] < 1$, *then*

$$\max_{\pi} R^{\pi}(\theta) - R^{KR(\mathbf{R}\theta)}(\theta) \leq O\left(\theta^2\right).$$

A version of the above result also holds for Markovian Klimov Networks in which all individual service times are exponentially distributed and priorities between customers are imposed preemptively. Please also note that trivially the above also provides an analysis for a conventional multiclass $M/GI/1$ queueing system, ie, with no feedback.

One application of the above network structure has all customers of class i needing an initial period of service of $\exp(\mu_{i1})$ duration. This service is conclusive with probability α_i. Should the first phase of service not prove conclusive, a second phase of service of $\exp(\mu_{i2})$ duration is provided prior to exiting the system. Priorities are imposed preemptively. Following the approach taken in Glazebrook (1996), in this case the Klimov Rule $KR(\mathbf{R}\theta)$ operates as follows: station $i1$ has an associated *Klimov Index* η_{i1} given by

$$\eta_{i1} = \max\left\{R_i\mu_{i1}\theta_i\alpha_i; \frac{R_i\theta_i}{\mu_{i1}^{-1} + (1-\alpha_i)\mu_{i2}^{-1}}\right\},$$

while the station $i2$ has an associated index

$$\eta_{i2} = R_i\theta_i\mu_{i2}.$$

In this case the Klimov Rule orders the stations according to the values of the indices $\{(\eta_{i1}, \eta_{i2}), 1 \le i \le k\}$, with highest priority accorded to stations of highest index.

Appendix C: Numerical Comparison of $R\mu$ and $R\mu\theta$ Rules

We illustrate the convergence of the $R\mu\theta$ and $R\mu$ rules in Table 1. Reward rates for the $R\mu\theta$ and $R\mu$ rules to a given accuracy are obtained by truncating the state space and using the uniformisation technique to facilitate the deployment of value iteration. Truncation levels are set so that the resulting finite-state model provides a sufficiently good approximation to the original model. We use $N_i = 100$ for each class i. Problems were randomly generated with respect to assumptions on system parameter values. In light of the discussion in Section 3.2 of the main paper, care was taken to ensure that the problems generated were such that $R\mu\theta$ and $R\mu$ rules were distinct. As in the preamble to Theorem 1 in the main paper, abandonment rates are expressed as a multiple of some underlying abandonment rate, $\theta_j = \theta\nu_j$ and all problems studied are such that $\rho = \sum_j \lambda_j/\mu_j < 1$. Problems were randomly generated as follows:

$$\mu_j \sim U[0.2, 5] \quad \text{(all cases);} \tag{7a}$$

$$\lambda_j \sim U[0.2, 5] \quad \text{(all cases);} \tag{7b}$$

$$\rho \in [0.5, 0.9] \quad \text{(light traffic);} \tag{7c}$$

$$R_j \sim U[1, 3] \quad \text{(all } k=2 \text{ system cases);} \tag{7d}$$

$$R_j \sim U[1, 5] \quad \text{(all } k=3 \text{ system cases);} \tag{7e}$$

$$\nu_j \sim U[1, 3] \quad \text{(all cases);} \tag{7f}$$

In the parameter generation, the μ_j and λ_j were generated according to (7a) and (7b) by means of a rejection algorithm until the desired ρ condition (7c) was met. For each system, 100 problems were generated at random according to (7a) to (7f) for a given set of abandonment rates. For each problem, value iteration was used to compute the gains of the $R\mu\theta$ and $R\mu$ rules and an optimal policy.

As seen in Table 1, the percentage suboptimality of both policies go to zero in the limit $\theta \to 0$. As stated in the main paper, this would indeed be the case for any priority policy. However, it is evident that this convergence is much more rapid (of order θ^2) in the case of $R\mu\theta$, where it is already the case at $\theta = 0.1$ that the median percentage suboptimality of $R\mu\theta$ is zero (to 2 decimal points) for both $k = 2$ and $k = 3$. The much slower $O(\theta)$ convergence of the percentage suboptimality of the $R\mu$ rule is particularly clear from the 'median' columns of the $R\mu$ part of Table 1.

Table 1 Percentage suboptimalities of the $R\mu\theta$ rule and $R\mu$ rule in sets of 100 randomly generated problems in $k=2$ and $k=3$ class systems. Median and 90th percentiles are shown.

	$R\mu\theta$				$R\mu$			
	$k=2$		$k=3$		$k=2$		$k=3$	
θ	Median	90th	Median	90th	Median	90th	Median	90th
5	0.28	1.26	0.10	0.78	0.00	0.06	0.00	0.14
2.5	0.23	1.37	0.06	0.87	0.00	0.39	0.02	0.36
1	0.16	1.36	0.02	0.67	0.00	0.79	0.06	0.68
0.5	0.06	0.97	0.01	0.76	0.07	1.06	0.15	0.94
0.1	0.00	0.32	0.00	0.27	0.22	1.09	0.18	0.84
0.05	0.00	0.17	0.00	0.14	0.20	1.08	0.15	0.72
0.025	0.00	0.07	0.00	0.04	0.15	0.80	0.12	0.56
0.01	0.00	0.01	0.00	0.00	0.09	0.47	0.05	0.39
0.005	0.00	0.00	0.00	0.00	0.05	0.28	0.04	0.20
0.0025	0.00	0.00	0.00	0.00	0.03	0.16	0.02	0.12
0.001	0.00	0.00	0.00	0.00	0.01	0.07	0.01	0.06
0.0005	0.00	0.00	0.00	0.00	0.01	0.04	0.01	0.03
0.00025	0.00	0.00	0.00	0.00	0.00	0.02	0.00	0.01
0.0001	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.01

References

Glazebrook, K. D. 1996. On the undiscounted tax problem with precedence constraints. *Advances in Applied Probability* **28** 1123–1144.

Klimov, G. P. 1974. Time-sharing service systems I. *Theory of Probability & Its Applications* **19** 532–551.

Klimov, G. P. 1978. Time-sharing service systems II. *Theory of Probability & Its Applications* **23** 314–321.