

Delay Differentiation by Optimal-Balancing-of-Queue-Lengths Scheduling

Avijit Chakraborty[†], Rajarshi Roy[‡], and Utpal Mukherji[†]

[†] Department of Electrical Communication Engineering
Indian Institute of Science
Bangalore, 560 012 WDIA
080-293-2387

avjitic@ece.iisc.ernet.in, utpal@ece.iisc.ernet.in

[‡] Department of E & ECE
Indian Institute of Technology
Kharagpur 721 302 INDIA
03222-75-5221
royr@ece.iitkgp.ernet.in

Abstract— Scheduling policies adopted for statistical multiplexing should provide delay differentiation between different traffic classes, where each class represents an aggregate traffic of individual applications having the same target queueing delay requirement. We propose scheduling to optimally balance queue lengths as an approach to delay differentiation.

In particular, we assume a discrete-time, two-class, single-server queueing model with unit service time per packet. We find a scheduling policy that we show to be discounted-cost optimal, for Bernoulli packet arrivals using dynamic programming analysis and for i.i.d. batch arrivals using a step-wise cost-dominance analytical approach. We then use cumulative queue length state variables in the one-step cost function of the optimization formulation and obtain the next-stage optimal policy. Simulations show that this policy achieves long-term mean queueing delays closer to the respective target delays than the first policy, and also achieves smaller errors in short-term mean queueing delays than the Coffman-Mitrani policy.

1. INTRODUCTION

In scenarios such as communication via satellite, where signal propagation delays are large, the queueing delay margins available to packet communication applications can be small [1]. **Further**, these margins can be different for different applications. Hence, when packets from different applications are multiplexed statistically and transmitted on a single shared channel, the packets should be classified according to their delay requirements and scheduled for transmission using a suitably designed class-based scheduling policy.

Given the packet arrival statistics for different classes, it is possible to determine whether it is feasible to attain mean class queueing delays within their respective delay margins [2,3]. Having chosen a feasible set of mean class queueing

delays, it remains to determine a scheduling policy that will achieve delays that are close to the target mean delays when averaged over time intervals of sufficient length

The Coffman-Mitrani policy in [2] exactly achieves a feasible set of target long-term mean queueing delays, but the short-term performance may be poor. Subsequent approaches that have been reported either are primarily heuristic in approach [3], or requires somewhat involved computations to obtain policy parameters [4].

Our approach is to formulate a weighted-queue-length-balancing optimization problem in the framework of Markov decision theory, where class weights are set inversely proportional to the respective products of target delays and packet arrival rates. For the sake of simplicity, we assume that the class weights are identical (set to 1), i.e., the target delays are inversely proportional to the arrival rates of the respective classes. If instantaneous queue lengths could be equalized by a scheduling policy then the mean queueing delays achieved by the policy would equal the respective target delays: Although exact instantaneous equalization is not possible, the mean queueing delays achieved by a policy that attempts to equalize cumulative queue lengths may be close to the respective target delays.

The rest of the paper is organized as follows. In Section 2, we describe a multiplexing model and analytically obtain an optimal policy for Bernoulli packet arrivals and for i.i.d. batch arrivals. In this section, we also use cumulative queue length state variables in the optimization problem and obtain the next-stage optimal policy. In Section 3, simulation results are presented for the above policies as well as for the Coffman-Mitrani policy, indicating the errors in the short-term and long-term mean queueing delays relative to the target delays, and the relative values of the coefficients of variation for queueing delays. Section 4 concludes the paper.

2. MODELING AND ANALYSIS

A survey of numerous Markov decision models, proposed for applications in communication networks, can be found in [5]. In [6], a two-class multiplexer has been modeled as a discrete-time single-server two-queue system with unit service time per packet and i.i.d. packet hatch arrivals to each input queue (figure 1).

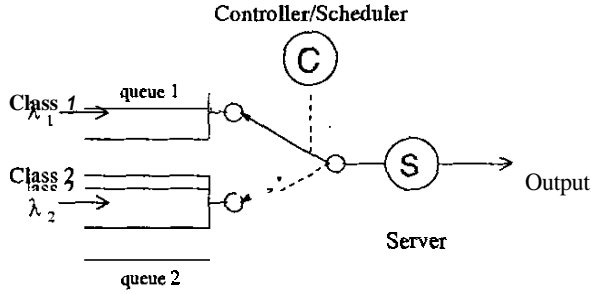


Figure 1. A two-class multiplexer

Bernoulli Arrivals

For the same model specialized to the case of Bernoulli arrivals to each input queue (arrival rate λ_i to queue i), we formulate an optimal scheduling problem with total discounted cost minimization criterion associated with a one-step cost function of the two queue length state variables. The cost function is taken to be symmetric in the two co-ordinates, nondecreasing in each co-ordinate, and increasing with modulus of difference between co-ordinates for fixed sum of co-ordinates. Examples of such one-step cost functions are $C(i, j) = i^p + j^p$ for $p > 1$, $C(i, j) = \max(i, j)$, and $C(i, j) = b^i + b^j$ for $b > 2$.

The total expected discounted cost over N stages under policy π , for discount factor $\beta \in (0, 1]$ and initial state \tilde{x} , is given by the value function

$$V_N^\pi(\tilde{x}) = E^\pi \left[\sum_{k=0}^{N-1} \beta^k C(\tilde{X}(k)) \mid \tilde{X}(0) = \tilde{x} \right], \forall N \geq 1;$$

$$V_0^\pi(\tilde{x}) := 0$$

where $\tilde{X}(k)$ denotes the state at time k

Let Π be the set of all policies. We want to find $\pi^* \in \Pi$, such that $V_N^{\pi^*}(\tilde{x}) = \inf_{\pi \in \Pi} V_N^\pi(\tilde{x})$. This optimal value function is denoted by $V_N(\tilde{x})$.

The Dynamic Programming results for finite action space $\mathcal{U}(\tilde{x})$ [7] yield the following optimality equation:

$$V_N(\tilde{x}) = C(\tilde{x}) + \beta \min_{\tilde{u} \in \mathcal{U}(\tilde{x})} \left[\sum_{\tilde{x}'} P_{\tilde{x}\tilde{x}'}(\tilde{u}) V_{N-1}(\tilde{x}') \right], \forall N \geq 1$$

where $P_{\tilde{x}\tilde{x}'}(\tilde{u})$ is the controlled Markov chain state transition probability.

We solve for an optimal policy by proving some structural properties for $V_N(\cdot, \cdot)$ (Lemma 1, stated in Appendix A.1) and obtain the following result.

Theorem 1: An optimal policy, for each horizon length N , serves at each time n the queue with index

$$\arg \max_{i=1,2} \{X_i(n) + \lambda_i\},$$

i.e., serves the longer queue when queue lengths are unequal, and a queue with the maximum arrival rate when queue lengths are equal.

We also note that, following this policy, the modulus of the difference of the individual queue lengths is at most 2 at any time after the queue lengths are first equal.

1.1.D. batch Arrivals

The model with i.i.d. batch arrivals in successive time units was considered in [6] with specific one-step cost function equal to the sum of the squares of the queue lengths. It was shown in [6] using a step-wise cost-dominance analytical approach, that an optimal stationary policy, for each horizon length, serves the longer queue when queue lengths are unequal, and the queue with the maximum arrival rate when queue lengths are equal and the arrival rates differ.

In the special case of identical arrival statistics for the two classes, we show that an optimal policy serves either one of the queues when both queue lengths are equal.

One-Step Cost Function of Cumulative Queue Lengths

So far, we have used only instantaneous queue lengths as state variables in our optimization problems. Simulations show that the ratios of the resulting long-term average queueing delays for the two queues are not satisfactorily close to the ratios of the target delays. We seek an improvement by assuming a one-step cost function equal to the sum of squares of cumulative queue lengths. The state at time n , $\tilde{X}(n)$, is given by the four-tuple $(X_1(n), X_2(n), \sum_{k=0}^n X_1(k), \sum_{k=0}^n X_2(k))$. Then the one-step cost function at time n is

$$C(\tilde{X}(n)) = \left(\sum_{k=0}^n X_1(k) \right)^2 + \left(\sum_{k=0}^n X_2(k) \right)^2$$

$$= \frac{1}{2} \left[\left(\sum_{k=0}^n X_1(k) - \sum_{k=0}^n X_2(k) \right)^2 + \left(\sum_{k=0}^n X_1(k) + \sum_{k=0}^n X_2(k) \right)^2 \right]$$

We note that the sum $X_1(k) + X_2(k)$, at any time k , is identical for all work-conserving policies. Hence, with the optimization restricted to the space of work-conserving policies, this choice for $C(\tilde{X}(n))$ is an attempt to equalize cumulative queue lengths.

We obtain the following next-stage optimal policy (i.e., with $N = 2$).

Theorem 2: The next-stage optimal work-conserving policy serves, at each time n for which $X_1(n) \neq 0$ and $X_2(n) \neq 0$, the queue with index

$$\arg \max_{i=1,2} \left\{ \sum_{k=0}^n X_i(k) + X_i(n) + \lambda_i \right\},$$

i.e., serves the **queue** with higher value of the sum of cumulative queue length, instantaneous queue length, and arrival rate. If the **sums** for both **queues** are **equal**, serving either **queue** is optimal.

3. SIMULATION RESULTS

In this section, we discuss the performance of the optimal policy using instantaneous queue lengths, the next-stage optimal policy using cumulative queue lengths, and the Coffman-Mitrani policy [2].

We first briefly describe the Coffman-Mitrani policy for our assumed model. Let S_1 and S_2 be two policies giving strict priority to class 1 and class 2 packets respectively. Then a mixing policy is constructed which, at the beginning of each busy period, decides with probability α that the scheduling decisions during that busy period are to be made according to S_1 , and with probability $1 - \alpha$ the decisions are to be made according to S_2 . Given the arrival statistics, the probability α is a constant and is determined solely by the target mean queue length of either class. Let \bar{X}_1^1 and \bar{X}_1^2 be the mean queue lengths of class 1 under policies S_1 and S_2 respectively. For all work-conserving policies, the sum of the mean queue lengths of both the classes is identical and denoted by \bar{X}_{total} . Because of the equal class weight assumption, the target mean queue length of class 1 is $\bar{X}_1^t = \bar{X}_{total}/2$. Then,

$$\alpha = \frac{\bar{X}_1^2 - \bar{X}_1^t}{\bar{X}_1^2 - \bar{X}_1^1}$$

For all the cases below.

- we have simulated a total of 36 uniformly distributed arrival rate pairs at total arrival rates from 0.2 upto 0.9, in increments of 0.1 in the individual class arrival rates,
- long-term average queueing delays are taken over a fixed length of 10^6 time units,
- short-term average queueing delays are taken over 1000 windows of length 1000 time units each,
- average over different arrival rate pairs refers to the arithmetic average over feasible arrival rate pairs alone (i.e., those pairs for which the ratio of the target mean queueing delays, which is determined as the reciprocal of the ratio of arrival rates because of the equal class weight assumption, is feasible given the arrival statistics).

Bernoulli Arrivals

Irrespective of the policy adopted, the number of feasible arrival rate pairs is 23. out of the total of 36 arrival rate pairs.

Using Instantaneous **Queue** Lengths, the ratios of the resulting long-term average queueing delays are within 25% of the ratios of the target delays for 18 out of 23 feasible arrival rate pairs. The average of the errors in the ratio of long-term average delays is approx. 21% and the average of the errors in the ratio of short-term average delays is also approx. 21%. For the arrival rate pair (0.3,0.4), the error in the ratio of long-term average delays is approx. 25%, and the errors in the ratios of short-term average delays are also 25% on average, the maximum value over 1000 windows being approx. 42%. These results are improvements upon the performance of the serve-the-longest-queue policy.

Using Cumulative Queue Lengths, the ratios of the resulting long-term average queueing delays are within 0.71% of the ratios of the target delays for all feasible arrival rate pairs, and within 0.31% of the ratios of the target delays for 18 out of the 23 feasible arrival rate pairs. The average of the errors in the ratio of long-term average delays is approx. 0.2%, while the average of the errors in the ratio of short-term average delays is approx. 6%. For the arrival rate pair (0.4,0.2), the error in the ratio of long-term average delays is approx. 0.31% but the errors in the ratios of short-term average delays are 7% on average, the maximum value over 1000 windows being approx. 26%.

If the Coffman-Mitrani policy is employed, the ratios of the resulting long-term average queueing delays are within 1% of the ratios of the target delays for 18 out of the 23 feasible arrival rate pairs. The average of the errors in the ratio of long-term average delays is approx. 0.7%, but the average of the errors in the ratio of short-term average delays is approx. 25%. For the arrival rate pair (0.1,0.8), the error in the ratio of long-term average delays is approx. 1% but the errors in the ratios of short-term average delays are 21% on average, the maximum value over 1000 windows being approx. 151%.

I.I.D. Batch Arrivals with Poisson Batch Sizes

Irrespective of the policy adopted, the number of feasible arrival rate pairs is 22. out of the total of 36 arrival rate pairs.

Using Instantaneous **Queue** Lengths, the ratios of the resulting long-term average queueing delays are within 30% of the ratios of the target delays for 16 out of 22 feasible arrival rate pairs. The average of the errors in the ratio of long-term average delays is approx. 22% and the average of the errors in the ratio of short-term average delays is approx. 24%. For the arrival rate pair (0.2,0.6), the error in the ratio of long-term average delays is approx. 30%, and the errors in the ratios of short-term average delays are also 30% on average, the maximum value over 1000 windows being approx. 48%.

Using Cumulative **Queue** Lengths, the ratios of the resulting long-term average queueing delays are within 0.59% of the ratios of the target delays for all feasible arrival rate pairs, and within 0.25% of the ratios of the target delays for 16 out of the 22 feasible arrival rate pairs. The average of the errors

in the ratio of long-term average delays is approx. 0.2% while the average of the errors in the ratio of short-term average delays is approx. 10%. For the arrival rate pair (0.4,0.4), the error in the ratio of long-term average delays is approx. 0.25%, but the errors in the ratios of short-term average delays are 6% on average, the maximum value over 1000 windows being approx. 25%.

If the Coffman-Mitrani Policy is employed, the ratios of the resulting long-term average queueing delays are within 1% of the ratios of the target delays for 16 out of the 22 feasible arrival rate pairs. The average of the errors in the ratio of long-term average delays is approx. 0.8% but the average of the errors in the ratio of short-term average delays is approx. 43%. For the arrival rate pair (0.4,0.5), the error in the ratio of long-term average delays is approx. 1%, but the errors in the ratios of short-term average delays are 102% on average, the maximum value over 1000 windows being approx. 1659%.

We note that apart from improved long-term average behaviour, using cumulative queue lengths results in improved short-term average behaviour as well, as shown in figures 2 and 3 (drawn in log-log scale).




Figure 2. Comparison of short-term behaviour for Bernoulli arrivals

Coefficients of Variation of Queueing Delays

Another criterion of importance is the coefficient of variation (ratio of standard deviation to mean) of queueing delay. Let c_i , c_c , and c_m denote the coefficients of variation of the queueing delays of any class for the policies using instantaneous queue lengths, using cumulative queue lengths, and using the Coffman-Mitrani policy respectively.

For feasible Bernoulli arrival rate pairs, the values of c_i , c_c , and c_m are in the respective ranges 0.04 – 0.90, 0.00 – 0.90, and 0.00 – 1.30. And the differences $c_c - c_i$, $c_m - c_i$, and $c_m - c_c$ vary between -0.27 and 0.24, -0.30 and 0.50, and -0.11 and 0.43 respectively.




Figure 3. Comparison of short-term behaviour for i.i.d. batch arrivals with Poisson batch sizes

And for i.i.d. batch arrivals with Poisson batch sizes, the values of c_i , c_c , and c_m are in the respective ranges 0.25 – 1.22, 0.32 – 1.27, and 0.32 – 1.54. And the differences $c_c - c_i$, $c_m - c_i$, and $c_m - c_c$ vary between -0.12 and 0.11, -0.12 and 0.55, and -0.02 and 0.52 respectively.

4. CONCLUSION

The next-stage optimal policy obtained using cumulative queue length state variables in the one-step cost function of the optimization formulation achieves long-term mean queueing delays closer to the respective target delays than by using only instantaneous queue length state variables. This policy using cumulative queue lengths also achieves short-term mean queueing delays closer to the respective target delays than the Coffman-Mitrani policy.

Extensions of this work to arbitrary class weights, larger number of classes, and correlated arrivals, and comparisons with the policies in [3,4], are desirable.

5. APPENDIX

A.1 Lemma 1

The three properties of the one-step cost function assumed for the Bernoulli arrival case in Section 2 are as follows:

- Property A : $C(i, j) = C(j, i)$
- Property B : $C(i + 1, j) \geq C(i, j)$ if $i \geq j$,
 $C(i, j + 1) \geq C(i, j)$ if $j \geq i$
- Property C : if $i_1 + j_1 = i_2 + j_2$, $|i_1 - i_2| = 1$,
and $|i_1 - j_1| > |i_2 - j_2|$,
then $C(i_1, j_1) > C(i_2, j_2)$

Starting with these properties, as $V_1(\cdot, \cdot)$ is $C(\cdot, \cdot)$, we prove the following lemma by induction.

Lemma 1: The N -stage optimal value function $V_N(\cdot, \cdot), \forall N$, has the following structural properties :-

(P1) if $i \geq j + 1$, then $V_N(i, j) \geq V_N(i - 1, j)$ if i

(equality holds if $i = j + 1$);

(P2) if $j \geq i + 1$, then $V_N(i, j) \geq V_N(i + 1, j - 1)$

(equality holds if $j = i + 1$);

(P3) if $i = j$, then $V_N(i + 1, j - 1) \geq \leq V_N(i - 1, j) + 1$ according as $\lambda_1 > = < \lambda_2$;

(P4) $V_N(i, j)$ is non-decreasing in each co-ordinate.

A.2 Proof of Theorem 2

Let the state at time $n, \tilde{X}(n) = (X_1(n), X_2(n), \sum_{k=0}^n X_1(k), \sum_{k=0}^n X_2(k))$, be $(i, j, s, t), i, j \neq 0$.




Figure 4. Timing Diagram

Departure occurs just after a control epoch, and arrivals occur just prior to a control epoch (figure 4).

Then, if queue 1 is served at time $n, \tilde{X}(n) + 1$ will be $(i - 1 + A_1(n) + 1, j + A_2(n) + 1, s + i - 1 + A_1(n) + 1, t + j + A_2(n) + 1)$,

and if queue 2 is served at time $n, \tilde{X}(n) + 1$ will be $(i + A_1(n) + 1, j - 1 + A_2(n) + 1, s + i + A_1(n) + 1, t + j - 1 + A_2(n) + 1)$,

where $A_m(n) + 1$ for $m = 1, 2$ denote the arrivals just prior to time $n + 1$ (figure 2) to queue 1 and queue 2 respectively.

From the choice of the one-step cost function, and denoting the cost incurred at time $n + 1$ when queue m is served at time n by $C^m(n) + 1, s + i + A_1(n) + 1$ by P_1 , and $t + j + A_2(n) + 1$ by P_2 , we have

$$\begin{aligned} & E[C^1(n) + 1 \mid \tilde{X}(n) = (i, j, s, t)] \\ &= E[(P_1 - 1)^2 + P_2^2 \mid \tilde{X}(n) = (i, j, s, t)], \\ & E[C^2(n) + 1 \mid \tilde{X}(n) = (i, j, s, t)] \\ &= E[P_1^2 + (P_2 - 1)^2 \mid \tilde{X}(n) = (i, j, s, t)]. \end{aligned}$$

$$\begin{aligned} \text{Hence, } & E[C^1(n) + 1 - C^2(n) + 1 \mid \tilde{X}(n) = (i, j, s, t)] \\ &= E[\{P_2^2 - (P_2 - 1)^2\} - \{P_1^2 - (P_1 - 1)^2\} \\ & \mid \tilde{X}(n) = (i, j, s, t)] \\ &= 2 E[\{P_2 - P_1\} \mid \tilde{X}(n) = (i, j, s, t)] \\ &= 2 \{t + j + \lambda_2\} - \{s + i + \lambda_1\} \end{aligned}$$

The next-stage optimal policy serves queue 2 / either queue 2 / queue 1 according as

$$E[C^1(n) + 1 - C^2(n) + 1 \mid \tilde{X}(n) = (i, j, s, t)] \geq < 0, \text{ i.e., according as } \{t + j + \lambda_2\} \geq < \{s + i + \lambda_1\}.$$

We note that for Markovian arrival processes, the state of the arrival process has to be included in the state description of the system, and λ_1 and λ_2 will be arrival state dependent.

Replacing i, j, s , and t by $X_1(n), X_2(n), \sum_{k=0}^n X_1(k)$, and $\sum_{k=0}^n X_2(k)$ respectively, Theorem 2 is proved.

REFERENCES

- [1] F. Gargione, T. Iida, F. Valdoni, and F. Vatalaro, "Services, Technologies, and Systems at Ka Band and Beyond - A Survey," *IEEE J. Selected Areas in Communications*, vol. 17, no. 2, pp. 133-144, Feb. 1999.
- [2] E. G. Coffinan. Jr. and I. Mitrani, "A Characterization of Waiting Time Performance Realizable by Single-Server Queues," *Operations Research*, vol. 28, no. 3, part II, pp. 810-821, May-June 1980.
- [3] C. Dovrolis, D. Stiliadis, and P. Ramanathan, "Proportional Differentiated Services: Delay Differentiation and Packet Scheduling," *IEEE/ACM Trans. Networking*, vol. 10, no. 1, pp. 12-26, Feb. 2002.
- [4] M. K. H. Leung, S. C. S. Lui, and D. K. Y. Yau, "Adaptive Proportional Delay Differentiated Services: Characterization and Performance Evaluation," *IEEE/ACM Trans. Networking*, vol. 9, no. 6, pp. 801-817, Dec. 2001.
- [5] E. Altman, "Applications of Markov Decision Processes in Communication Networks," *Handbook of Markov Decision Processes* (E. A. Feinberg and A. Shwartz Ed.), Kluwer Academic, 2002.
- [6] R. R. Pillai, and U. Mukherji, "Optimal Scheduling for Balancing Buffer Occupancy at an ATM Access Point," *Proceedings: International Conference on Information, Communications and Signal Processing*, IEEE, Singapore, pp. 649-653, Sept. 1997.
- [7] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, vol. 1, Athena Scientific. 1995.