Optimal Measurement-based Pricing for an M/M/1 Queue

Yezekael Hayel[1] and Bruno Tuffin[1]

IRISA-INRIA Rennes, Campus universitaire de Beaulieu
35042 Rennes Cedex - FRANCE
{yhayel,btuffin}@irisa.fr
tel: (+33)299847134
fax: (+33)299842529

Abstract. In this paper, we consider a system modelled as an M/M/1 queue. Several users send jobs to the queue and are characterized by a delay cost per unit of time and a demand function. Our goal is to design an optimal pricing scheme for the queue, where the total charge depends on both the mean delay at the queue and arrival rate of each customer. We also assume that those two values have to be (statistically) measured, introducing errors on the total charge that might avert users from using the system, and then decrease demand. This model can be applied in telecommunication networks, where pricing can be used to control congestion, and the network can be characterized by a single bottleneck queue; the throughput of each user would be determined through passive measurements while the delay would be determined through active measurements.

1 Introduction

In many situations and systems, controlling quality of service is an important task. Indeed, when resource is limited and demand is high, congestion occurs and delay for completion of service might increase to an unacceptable level. By using a suitable pricing scheme, the facility manager can control congestion, offer satisfactory quality of service and properly allocate resource. An application of special attention for the authors is telecommunication networks, where congestion control and service differentiation have become a concern of increasing interest; pricing has therefore been seen as a solution to tackle the problem [3]. See also [4] for recent surveys on pricing in telecommunication networks. Pricing is a topic of argument in the telecommunication community, where some believe that it will be difficult to leave the current flat-rate scheme adopted by service providers, especially with the introduction of optical fiber for which congestion is unlikely to occur. Though, switching to optical fiber is expensive in access networks, and the radio spectrum will remain limited in wireless networks. We thus follow the line of argument that a controlling procedure will have to be applied in these cases, especially with the increasing demand in terms of bandwidth and quality of service.

We consider in this paper an M/M/1 queue (often used to represent the bottleneck queue in a telecommunication network), see for instance [1] and users which adapt their demand to the price charged by the system manager as well as the mean delay at the queue.

This kind of problem has already been studied in [5, 14] where, in a multi-class framework, social welfare optimal prices have been determined. This work has been extended in [17] to non-linear waiting costs, in [12] to a queuing network with dynamic pricing and in [7] to customer-chosen service requirements.

In all those models, users' mean delays and throughputs are assumed perfectly known or observable. Though, in many contexts those values cannot be known exactly and have to be statistically measured, introducing errors. This is for instance the case in the Internet where, for scalability reasons, the throughput of each user is often estimated by *passive measurements*, that is non-intrusive tools, by counting statistically the number of sent packets [8, 10, 11], while end-to-end delay is measured using *active measurements*, that is by sending probes in the network at random times [2, 6, 16, 18].

The main contribution of this paper is the introduction of measurements to determine delay and users' throughput. We consider a single class M/M/1/FIFO queue with heterogeneous customers. Each user is assumed to have his own demand function and delay cost per unit of time. The goal of the facility manager is to find out prices optimizing the social welfare. With respect to the literature, the main difference is then that the standard deviation between the actual and theoretical prices is taken into account in demand functions (averting customers from using the facility, as errors in price computation introduce dissatisfaction), as well as the cost of performing measurements. It also introduces an additional problem that is the determination of sampling frequency parameters: counting more frequently the number of jobs (or packets) at the queue improves the precision but increases management costs, while sending too many probes would increase the processing delay of jobs.

The paper is organized as follows. The basic model is defined in Section 2. Section 3 expresses how throughput and delay can be measured. The optimization problem is described in Section 4, and several algorithms are provided in Section 5, depending on the importance we fix to the last measurements if we wish to incorporate potential changes of parameter values during time. Finally, our concluding remarks can be found in Section 6.

2 Basic framework

We consider an M/M/1/FIFO queue and a population of J users sending jobs (or packets) to the queue. For sake of simplicity, we assume that all jobs (packets) have the same size. Let d be the mean delay by job at this queue. Each user j ($1 \leq j \leq J$) is characterized by his (aggregated) value function $V_j(\lambda_j)$ which specifies the gross value gained when sending jobs to the system at rate λ_j. The value function $V_j(\cdot)$ is assumed to be monotone increasing, continuously differentiable and strictly concave, for all $j \in \{1, \ldots, J\}$. This function is highly

related to demand function in the following way [14]: let z represent the full price, meaning the *charge* plus the *felt* cost. The corresponding demand is $\lambda_j = D_j(z) = (1 - H_j(z))\Lambda_j$ where Λ_j is the maximum potential arrival rate of class-i jobs and $H_j(\cdot)$ is the distribution function of the service valuation. Inverting this function, we have $V_j'(\lambda_j) = D_j^{-1}(\lambda_j)$.
The total (per job) cost perceived by user j is made of p_j, the per job price, $v_j d$, the linear delay cost (with v_j delay cost per unit of time), and $\beta_j \sigma_j$ the cost of aversion for the error in the total charge computation, with β_j cost per unit of error and $\sigma_j = \sqrt{\mathbb{V}(\hat{p}_j)}$, the standard deviation of the price computation, due to measurements, where $\widehat{p}_j$ is the estimated price (depending on measured throughput and delay; see below).

The *demand relationship* is thus given by

$$V_j'(\lambda_j) = p_j + v_j d + \beta_j \sigma_j. \tag{1}$$

In this relation, the delay cost $v_j d$ depends on the *perceived* delay, i.e. the one that is really experienced, and not the estimated one which is used for computing the price. Note again that this delay cost is a *felt* cost, not a charged one.

Finally, define $\lambda = \sum_{j=1}^{J} \lambda_j$, the total arrival rate and $\underline{\lambda} = (\lambda_1, \cdots, \lambda_J)$ as notations that will be helpful in the remaining of the paper.

3 Measurements

We assume here that delay and throughputs are unknown and have to be determined in practice. Based on what is done in telecommunication networks [6], we assume that delay is estimated by sending probes into the system. Probes are special jobs served like standard jobs but designed to observe the response time. Thus, the mean response time over the probes will be used to estimate the average delay. Probes are assumed to be send according to a Poisson process with rate γ. The larger γ, the better the estimation, but at the expense of an increased delay (since those probes have to be served).

Based on similar telecommunication networks arguments, we assume that each job passing through the system cannot be counted for scalability reason, due to management and storage requirements. We thus assume that only a fixed and small fraction ε of traffic is observed [11], each job being selected according to a Bernoulli law of mean ε (independent between jobs). Nevertheless, we assume that probes are automatically detected so that they are not sampled.

We separate time into slots of length T during which we assume that the number J of users is fixed, as well as their sending rates λ_j. Measurements are performed during each slot, and the total charge is computed from those measurements.
Let N be the number of jobs arriving during a measurement slot. Let also N_s be the number of jobs sampled during that period and X_{ij} ($1 \le j \le J$, $1 \le i \le N_s$) be a Bernoulli random variable equal to 1 if the i-th sampled job is a job of user j, and 0 otherwise. The estimation of λ_j is carried out as follows.

Proposition 1. *An unbiased estimator of* λ_j *is*

$$\hat{\lambda}_j = \frac{\sum_{i=1}^{N_s} X_{ij}}{\varepsilon T}. \tag{2}$$

Its variance is

$$\mathbb{V}[\hat{\lambda}_j] = \frac{\lambda_j}{\varepsilon T}.$$

The proof of this proposition is provided in Appendix A.

The mean response time is estimated by sending probes/jobs in the system according to a Poisson process with rate γ. The purpose is only to get an estimation $\hat{d}$ in each measurement slot of length T. The total arrival rate is then $\lambda + \gamma$. Again, for notions on active (intrusive) measurements in communication networks, the reader is referred to [15]. ¿From [9], the response time of each job in such an M/M/1 queue is known to follow an exponential distribution with rate $\mu - \lambda - \gamma$. We assume here that $\gamma << \lambda$ so that the measured response times of probes can be considered independent. This assumption seems relevant since, even if we wish to obtain a precise estimation, it cannot be at the cost of a large increase of delay for actual jobs. Finally, let N_a be the (random) number of probes sent during a slot (of length T) and d_k be the measured response time for the k-th probe. We then have the following estimator of mean response time.

Proposition 2. *An unbiased estimator of mean response time is*

$$\hat{d} = \frac{\sum_{k=1}^{N_a} d_k}{\gamma T}, \tag{3}$$

and its variance (under the assumption of independence between delays of probes) is given by

$$\mathbb{V}(\hat{d}) = \frac{2}{\gamma T(\mu - \lambda - \gamma)}.$$

The proposition is proved in Appendix B. Thanks to those estimations, optimal prices are computed in the next section.

4 Optimal prices

The goal of this paper, like in [13, 14], is to design a pricing scheme optimizing the total expected net value of the users $\sum_{j=1}^{J}(V_j(\lambda_j) - v_j\lambda_j d)$, minus the cost of processing measurements for the system $\alpha\varepsilon\lambda$, where α is the cost per passive measurement, in terms of memory and management requirements. Sending probes to measure delay is not supposed to have any direct cost for the system, but only an indirect one by an increased delay. As well as prices, sampling parameters γ and ε can be properly chosen. The optimization problem can then be formulated as finding out

$$(\underline{\lambda}^*, \gamma^*, \varepsilon^*) = \arg\max_{\underline{\lambda},\gamma,\varepsilon}\{\sum_{j=1}^{J}(V_j(\lambda_j) - v_j\lambda_j d) - \alpha\varepsilon\lambda\}, \tag{4}$$

given that prices p_i and rates λ_i depend one on each other by the demand relationships (1). We recall here that delay y depends on λ.

For given values of ε and γ, the optimal pricing scheme is given by the following theorem.

Theorem 1 *The optimal job price for user j, for $j \in \{1, \ldots, J\}$, is*

$$p_j^* = d^2 \sum_{k=1}^{J} v_k \lambda_k^* - \beta_j \sigma_j + \alpha \epsilon,$$

where $\underline{\lambda} = \underline{\lambda}^*$ is the arrival-rate vector maximizing (4), and with

$$\sigma_j^2 = \frac{d^4}{T} \left[\left(1 + \frac{12}{\gamma T} + \frac{36}{\gamma^2 T^2} + \frac{24}{\gamma^3 T^3}\right) \sum_{j=1}^{J} v_j^2 \frac{\lambda_j^*}{\varepsilon} + \frac{8}{\gamma} \left(1 + \frac{4}{\gamma T} + \frac{3}{\gamma^2 T^2}\right) (\sum_{j=1}^{J} v_j \lambda_j^*)^2 \right].$$

The proof of this theorem is given in Appendix C. Prices $\underline{p}^* = (p_1^*, \cdots, p_J^*)$ induce the optimal arrival rate λ_j^* for each user j and consequently constitutes an optimal price schedule (since, thanks to the demand relationships, with theses prices, users will choose the arrival rates maximizing (4)).
Optimal values of sampling parameters ε and γ depend on $\underline{\lambda}^*$, that is functions V_j. Also, the optimal values $\underline{\lambda}^*$ depend on ε and γ from Theorem 1. In the next section, we propose to dynamically adapt ε and γ measurement slot per measurement slot.

5 Numerical algorithms

We describe here how parameters γ, ε and the λ_i can adapt themselves dynamically over time (discretized in measurement slots), depending on the importance we assign to the last measurement slot in the overall estimations. We thus use index or superscript t to indicate the t-th slot.

5.1 Algorithms

Consider that we are at the end of the t-th slot, with given value of ε_t, γ_t and λ_i^t, so that the values for the $(t+1)$-th slot are to be determined.

- From (2) and (3), we get estimations $\widehat{\lambda}_j^t$, $\forall j = 1, \ldots, J$, and $\widehat{d}^t$ of real consumption/throughputs and delay at t-th slot.
- Estimations from previous slots are combined with the current one by, $\forall j \in \{1, \ldots, I\}$,

$$\tilde{\lambda}_j^t = r_\lambda^t \frac{\sum_{k=1}^{t-1} \hat{\lambda}_j^k}{t-1} + (1 - r_\lambda^t) \hat{\lambda}_j^t,$$

  and,

$$\tilde{d}^t = r_d^t \frac{\sum_{k=1}^{t-1} \hat{d}^k}{t-1} + (1 - r_d^t) \hat{d}^t,$$

where r_λ^t and r_d^t are positive real numbers less than or equal to one corresponding to the proportion we assign to the estimations from previous slots.

- Prices for the next slot are computed based on Theorem 1:

$$\widehat{p}_j^{t+1} = (\tilde{d}^t)^2 \sum_{j=1}^{J} v_j \tilde{\lambda_j}^t - \beta_j \sigma_j(\underline{\tilde{\lambda}}^t, \tilde{d}^t, \gamma_t, \epsilon_t) + \alpha \epsilon_t,$$

  with

$$\sigma_j^2(\underline{\lambda}, d, \gamma, \epsilon) = \frac{d^4}{T}\left(\left(1+\frac{12}{\gamma T}+\frac{36}{\gamma^2 T^2}+\frac{24}{\gamma^3 T^3}\right)\sum_{j=1}^{J} v_j^2 \frac{\lambda_j}{\varepsilon} + \frac{8}{\gamma}\left(1+\frac{4}{\gamma T}+\frac{3}{\gamma^2 T^2}\right)(\sum_{j=1}^{J} v_j \lambda_j)^2\right).$$

- The system manager sets parameters γ_{t+1} and ϵ_{t+1} in order to maximize the social welfare of the system for the $(t+1)$-th slot:

$$(\gamma_{t+1}, \epsilon_{t+1}) = \arg\max_{\gamma,\epsilon} \sum_{j=1}^{J}\left(V_j(\lambda_j^{t+1}) - \frac{v_j \lambda_j^{t+1}}{\mu - \sum_{j=1}^{J} \lambda_j^{t+1} - \gamma} - \alpha\epsilon\lambda_j^{t+1}\right),$$

  extrapolating future demand (depending on those parameters) by:

$$\lambda_j^{t+1}(\gamma, \epsilon) = \Lambda_j \overline{\Phi}_j\left(\hat{p}_j^{t+1} + v_j \tilde{d}^t + \beta_j \sigma_j(\tilde{\lambda}^t, \tilde{d}^t, \gamma, \epsilon)\right).$$

Then estimations then are computed at the $(t+1)$-th slot, and so on.
We consider in this paper three different policies for the proportions r_λ^t and r_d^t assigned to previous measurement slots in the estimations:

- *last*: only the last slot is taken into account in the estimations so that $r_d^t = r_\lambda^t = 0\ \forall t$. This choice is valid in the case where demand varies extremely between slots, meaning that previous estimations are useless for the current one.
- *int*: intermediate scheme where a fixed proportion is assigned to the last slot, i.e., $r_d^t = r_d$ and $r_\lambda^t = r_\lambda\ \forall t$. Thus with respect to *last* policy, estimations are smoothed but the algorithm still reacts to changing conditions in the system.
- *id*: each past slot is assigned the same weight so that $r_d^t = r_\lambda^t = (t-1)/t\ \forall t$. This policy is valid when the system conditions (number J of users, demand functions) are constant over time. We then have

$$\tilde{\lambda}_j^t = r_\lambda^t \frac{\sum_{k=1}^{t-1} \hat{\lambda}_j^k}{t-1} + (1-r_\lambda^t)\hat{\lambda}_j^t = \frac{\sum_{k=1}^{t} \hat{\lambda}_j^k}{t}.$$

  Applying the strong law of large numbers, as $\hat{\lambda}_j^k$ are independent random variables, we obtain

$$\tilde{\lambda}_j^t = \frac{\sum_{k=1}^{t} \hat{\lambda}_j^k}{t} \xrightarrow[t\to+\infty]{} I\!E(\hat{\lambda}_j) = \lambda_j.$$

  The same convergence result stands for the delay estimation.

5.2 Illustrations

As an illustration, consider a system with $\mu = 10$, $T = 1000$, and a number of users $J = 3$.

Case of fixed demand We consider that users have fixed demand, expressed by:

$$V_1(\lambda_1) = 15\lambda_1 - 3\lambda_1^2,$$
$$V_2(\lambda_2) = 17\lambda_2 - 4\lambda_2^2,$$
$$V_3(\lambda_3) = 16\lambda_3 - 5\lambda_3^2.$$

Other parameters are arbitrarily taken as $v = [1; 1.5; 2]$, $\beta = [2; 1; 0.5]$ and $\alpha = 0.005$. At the first time slot, γ_0 is set to 10% of the smallest demand among users while ε_0 is taken from a uniform distribution over $[0, 0.1]$.

We first look at the evolution of throughputs λ_i ($i = 1, 2, 3$) and active sampling rate γ. Figure 1 displays the results for the *last* policy, Figure 2 for the *int* policy with $r_d = r_\lambda = 0.9$ and Figure 3 for the *id* policy. As expected, when the results from previous slots are not used (*last* policy), the estimations experience a large variability since only the last slot is used. On the other hand, *int* policy smoothes the estimations, even if a small variability is kept, while convergence is illustrated for *id* policy.

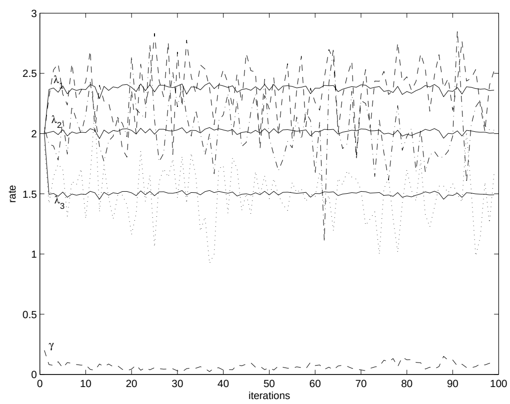

Fig. 1. Evolution of λ_1,λ_2,λ_3 and γ for *last* policy when the number of slots increases (fixed demand). The dashed lines are for the estimations while the plain lines are for actual values of λ_i.

Figure 4 presents the error for the delay estimation in the case of each policy. Again, *last policy* shows a larger variability in the estimation, while *id* policy converges. Note also that for all three cases, $\gamma << \lambda$ so that the independent

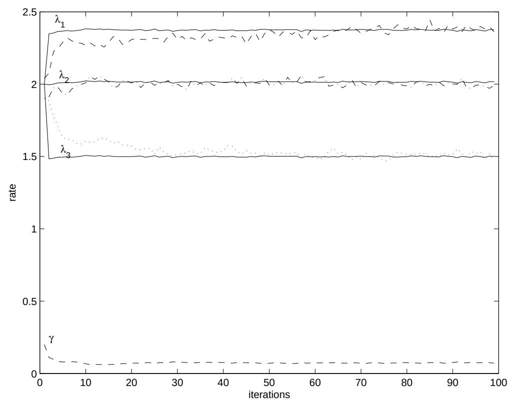

Fig. 2. Evolution of λ_1,λ_2,λ_3 and γ for *int* policy when the number of slots increases (fixed demand). The dashed lines are for the estimations while the plain lines are for actual values of λ_i.

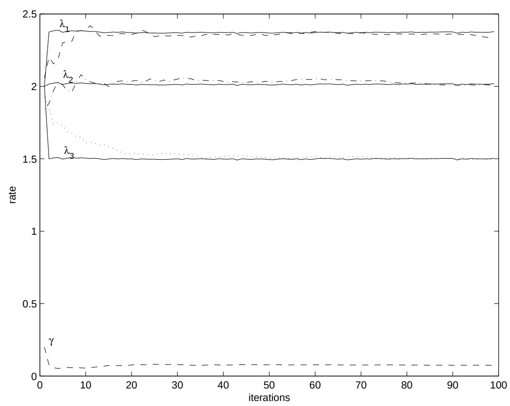

Fig. 3. Evolution of λ_1,λ_2,λ_3 and γ for *id* policy when the number of slots increases (fixed demand). The dashed lines are for the estimations while the plain lines are for actual values of λ_i.

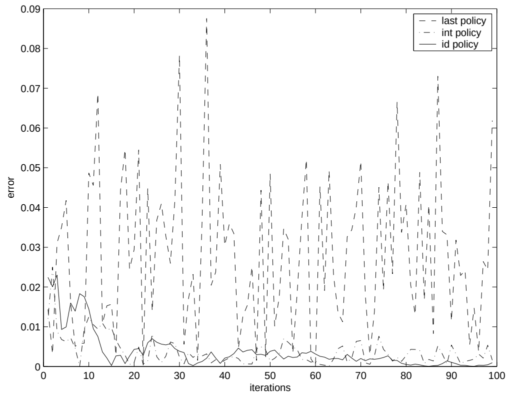

Fig. 4. Evolution of the error between real mean delay and the estimator using the three policies (fixed demand).

approximation between probes is valid. Exactly the same kind of remarks apply for the evolution of the passive measurements parameter ϵ displayed on Figure 5 for each policy.

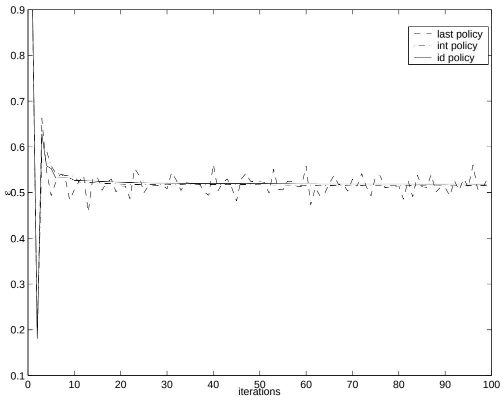

Fig. 5. Evolution of the passive measurement parameter ϵ for each policy (fixed demand).

Case of variable demand Look now at the case where demand varies for two users, and is kept constant for the third one. Demand variations are expressed by varying value functions V_j with time slot t. We use:

$$V_1(\lambda_1, t) = 15\lambda_1 - 3\lambda_1^2,$$
$$V_2(\lambda_2, t) = 17\lambda_2 - 4\lambda_2^2,$$
$$V_3(\lambda_3, t) = (16\lambda_3 - 5\lambda_3^2)\frac{t}{200}.$$

Figures 6, 7 and 8 display the evolution of actual throughputs, their estimation, and the active measurement rate γ for respectively *last*, *int* and *id* policies. We can still observe that variability is the highest for *last* and the lowest for *id*. For the *int* policy, we use $r_\lambda = r_d = 0.35$ in order to reduce the variability in the estimation of parameters with respect to *last*. *id* policy is shown not to stick to changes of demand as well as *last* or *int*. Therefore, a small variability, by using

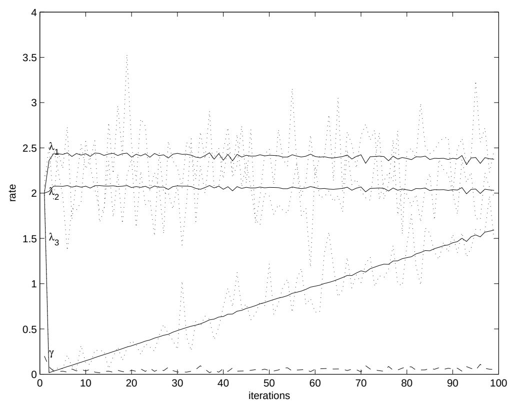

Fig. 6. Evolution of λ_1,λ_2,λ_3 and γ for *last* policy when the number of slots increases (variable demand). The dashed lines are for the estimations while the plain lines are for actual values of λ_i.

past measurements with different parameters, is at the cost of a larger bias in the estimation, like for *id* policy. *int* policy might then be seen here as a good trade-off between variance and bias.

Figure 9 displays errors in the average delay estimation for the three policies. Due to its bias in case of variable demand, the estimation provided by *id* policy degrades as time increases. Here also *int* policy appears a good trade-off.

Figure 10 presents the evolution of passive measurement parameter ϵ. A slightly increasing behavior is observed, due to the increase in total demand.

6 Conclusions

In this paper, we have studied an optimal pricing scheme for a system with heterogeneous users, meaning users with different demand patterns and delay

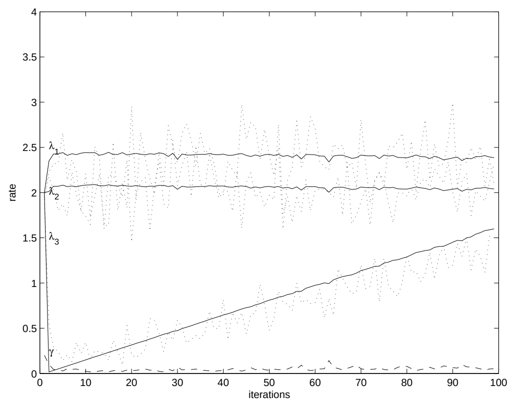

Fig. 7. Evolution of λ_1,λ_2,λ_3 and γ for *int* policy when the number of slots increases (variable demand). The dashed lines are for the estimations while the plain lines are for actual values of λ_i. We consider $r_\lambda = r_d = 0.35$.

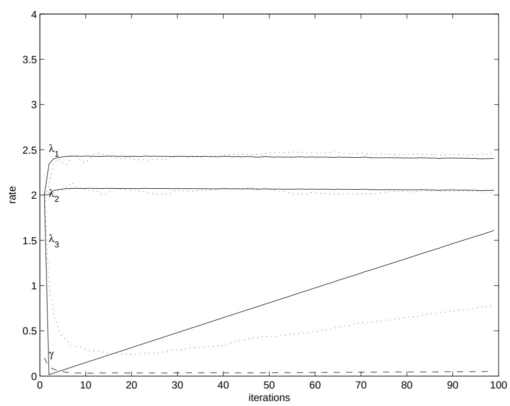

Fig. 8. Evolution of λ_1,λ_2,λ_3 and γ for *id* policy when the number of slots increases (variable demand). The dashed lines are for the estimations while the plain lines are for actual values of λ_i.

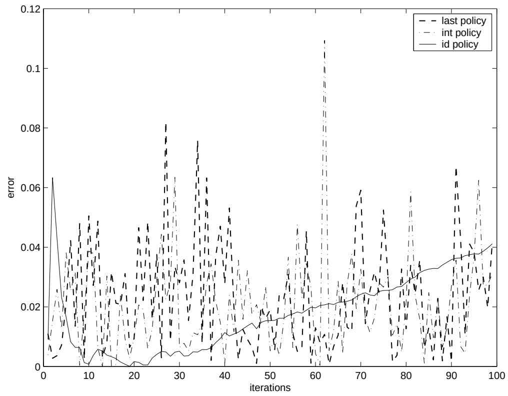

Fig. 9. Evolution of the error between real mean delay and the estimator using the three policies (variable demand).

cost. In our model, the system performance (delay) and users throughput have to be estimated by respectively active and passive measurements. This problem has implications in telecommunication networks for instance. We also have introduced several algorithms to adapt the system to potential modifications in parameter values. The algorithm that uses the *last* policy, where a fixed proportion of past estimations is assigned to overall estimations seems to be a good choice for systems experiencing variations over time in their parameters. Nevertheless, the algorithm based on the *id* policy is well adapted to the case where users' consumption is stable over time.

As an extension, we currently work on a multi-class model where users have to choose between several priority classes. We also plan to study the effect of the granularity parameter T on the measurements parameters. Another extension of our model is to consider heterogeneous traffic, with different packet sizes.

References

1. E. Altman, C. Barakat, and V. Ramos. Queuing Analysis of Simple FEC Schemes over IP. *Computer Networks*, 39(2), 2002.
2. J.C. Bolot. Characterizing End-to-End Packet Delay and Loss in the Internet. *Journal of High Speed Networks*, 2, 1993.
3. C. Courcoubetis and R. Weber. *Pricing Communication Networks—Economics, Technology and Modelling.* Wiley, 2003.
4. L.A. DaSilva. Pricing of QoS-Enabled Networks: A Survey. *IEEE Communications Surveys & Tutorials*, 3(2), 2000.
5. S. Dewan and H. Mendelson. User delay costs and internal pricing for a service facility. *Management Science*, 36(12):1502–1517, 1990.
6. A.B. Downey. Using pathchar to estimate internet link characteristics. In *Proceedings of ACM SIGCOMM*, 1999.

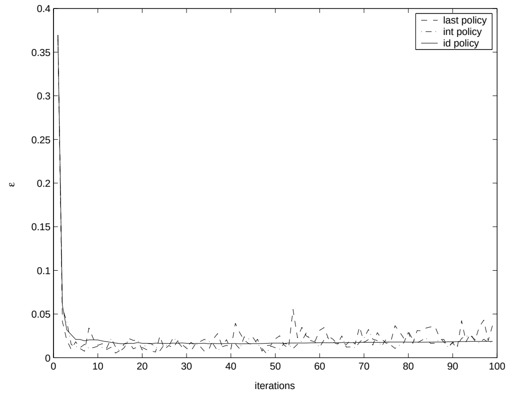

Fig. 10. Evolution of the passive measurement parameter ϵ for each policy (variable demand).

7. A.Y. Ha. Optimal pricing that coordinates queues with customer-chosen service requirements. *Management Science*, 47(7):915–930, 2001.
8. S. Jaiswal, G. Iannaccone, C. Diot, J. Kurose, and D. Towsley. Inferring TCP Connection Characteristics Through Passive Measurements. In *Proceedings of IEEE INFOCOM*, 2004.
9. L. Kleinrock. *Queuing Systems: Theory*, volume 1. J. Wiley & Sons, 1975.
10. M. Kodialam, T. Lakshman, and S. Mohanty. Runs bAsed Traffi Estimator (RATE): A Simple, Memory Efficient Scheme for Per-Flow Rate Estimation. In *Proceedings of IEEE INFOCOM*, 2004.
11. K. Lai and M. Baker. Nettimer : a tool for measuring bottleneck link bandwidth. In *3rd USENIX Symposium on Internet Technologies and Systems*, pages 122–133, San Francisco, CA, March 2001.
12. Y. Masuda and S. Whang. Capacity Management in Decentralized Networks. *Management Science*, 48(12):1628–1634, 2002.
13. J.K. McKie-Mason and H.R. Varian. Pricing Congestible Network Resources. *IEEE Journal on Selected Areas in Communications*, 13(7):1141–1149, 1995.
14. H. Mendelson and S. Whang. Optimal incentive-compatible priority pricing for the M/M/1 queue. *Operations Research*, 38(5):870–883, 1990.
15. V. Paxson. *Measurements and Analysis of End-to-End Internet Dynamics.* PhD thesis, University of California Berkeley, 1997.
16. V. Paxson. End-to-end Internet Packet Dynamics. *IEEE/ACM Transactions on Networking*, 7(3), 1999.
17. J.A. Van Mieghem. Price and Service Discrimination in Queuing Systems: Incentive Compatibility of Gcμ Scheduling. *Management Science*, 46(9):1249–1267, 2000.
18. R. Whitner, G. Pollock, and C. Cook. On Active Measurements in QoS-Enabled IP Networks. In *Proceedings of Passive and Active Measurement (PAM) Workshop*, 2002.

A Proof of Proposition 1

The number N of jobs arriving during a measurement slot of length T follows a Poisson distribution with rate λT. Given N, N_s follows a binomial distribution with N trials and probability of success ε.

The generating function of random variable $\varepsilon T\hat{\lambda}_j$ is

$$\begin{aligned} G_{\varepsilon T\hat{\lambda}_j}(s) = \mathbb{E}[\mathbb{E}[s^{\sum_{k=1}^{N_p} X_{jk}}|N_p]] &= \mathbb{E}[((1-\rho_j)+s\rho_j)^{N_p}] \\ &= \mathbb{E}[\mathbb{E}[((1-\rho_j)+s\rho_j)^{N_p}|N]] \\ &= \mathbb{E}[((1-\varepsilon)+((1-\rho_j)+s\rho_j)\varepsilon)^N] \\ &= \exp(-\lambda T((1-\varepsilon)+((1-\rho_j)+s\rho_j)\varepsilon-1)) \\ &= \exp(\lambda_j T\varepsilon(s-1)). \end{aligned}$$

The random variable $\varepsilon T\hat{\lambda}_j$ thus follows a Poisson distribution with parameter $\lambda_j T\varepsilon$. Taking the expectation, it can be easily seen that $\hat{\lambda}_j$ is a unbiased estimator of λ_j. Also, we get

$$\mathbb{V}[\hat{\lambda}_j] = \frac{1}{\varepsilon^2 T^2}\mathbb{V}[\varepsilon T\hat{\lambda}_j] = \frac{\lambda_j T\varepsilon}{\varepsilon^2 T^2} = \frac{\lambda_j}{\varepsilon T}.$$

B Proof of Proposition 2

The number N_a of probes sent during a slot of length T follows a Poisson distribution with parameter γT. Since we have assumed that response times of probes are statistically independent, the moments of $\hat{d}$ are

$$\mathbb{E}((\hat{d})^m) = \frac{1}{\gamma^m T^m(\mu-\lambda-\gamma)^m}\mathbb{E}\left(\frac{(N_a+m-1)!}{(N_a-1)!}\right). \tag{5}$$

For $m=1$, we obtain $\mathbb{E}(\hat{d}) = (\mu-\lambda-\gamma)^{-1} = d$. The result for the variance follows similarly.

C Proof of Theorem 1

The first order condition over λ_j for maximizing (4) is

$$V_j'(\lambda_k^*) = v_j d + \sum_{k=1}^{J} v_k \lambda_k \frac{\partial d}{\partial \lambda_j}(\underline{\lambda}) + \alpha\epsilon.$$

Furthermore, the demand relationship for user j is

$$V_j'(\lambda_j) = p_j + v_j d + \beta_j \sigma_j.$$

Combining the above equations, we obtain at $\underline{\lambda} = \underline{\lambda}^*$

$$p_j^* = \sum_{k=1}^{J} v_k \lambda_k^* \frac{\partial d}{\partial \lambda_j}(\underline{\lambda}^*) - \beta_j \sigma_j + \alpha \epsilon. \tag{6}$$

Since we consider an M/M/1/FIFO queue, $d = 1/(\mu - \lambda - \gamma)$, so that $\frac{\partial d}{\partial \lambda_j} = d^2$ giving the first part of the theorem.

Also,

$$\sigma_j^2 = Var(\widehat{p}_j) = Var(\widehat{d}^2 \sum_{k=1}^{J} v_k \widehat{\lambda}_k) = \mathbb{E}(\widehat{d}^4)\mathbb{V}(\sum_{k=1}^{J} v_k \widehat{\lambda}_k) + \mathbb{E}(\sum_{k=1}^{J} v_k \widehat{\lambda}_k)^2 \mathbb{V}(\widehat{d}^2).$$

From (5) and the independence between estimations, we get

$$\begin{aligned}
\mathbb{E}(\widehat{d}^4) &= \frac{1}{\gamma^3 T^3 (\mu - \lambda - \gamma)^4} \left(\gamma^3 T^3 + 12\gamma^2 T^2 + 36\gamma T + 24\right), \\
\mathbb{V}(\sum_{k=1}^{J} v_k \widehat{\lambda}_k) &= \sum_{k=1}^{J} v_k^2 \frac{\lambda_k}{\varepsilon T}, \\
\mathbb{E}(\sum_{k=1}^{J} v_k \widehat{\lambda}_k) &= \sum_{k=1}^{J} v_k \lambda_k, \\
\mathbb{V}(\widehat{d}^2) &= \frac{8}{\gamma T (\mu - \lambda - \gamma)^4} \left(1 + \frac{4}{\gamma T} + \frac{3}{\gamma^2 T^2}\right),
\end{aligned}$$

providing the last part of the theorem.