Optimal Adaptive Feedback Control of a Network Buffer.

V. Guffens, G. Bastin
Centre for Systems Engineering and Applied
Mechanics (CESAME)
Université Catholique de Louvain
Belgium, 1348 Louvain-La-Neuve
{guffens,bastin}@auto.ucl.ac.be

Abstract— **A general fluid flow model for a network buffer with tail drop queueing policy is introduced and shown to be suitable for representing a large class of queueing system. An on-line identification scheme for this model is presented and an optimal control strategy is developed using the minimal principle of Pontryagin. Experimental results show that the implementation of this control policy is nearly optimal for a large range of experimental conditions.**

I. INTRODUCTION

Optimal control theory, based on the minimum principle of Pontryagin[2], is a common framework to produce attractive control laws for dynamical systems. While the set of conditions to be satisfied under this principle gives rise to an extremal trajectory, the obtained control law is open loop, model sensitive and is therefore not suitable for an actual implementation. However, the optimal control analysis also suggests a heuristic control design that can be implemented in closed loop using further information on the system to obtain a robust quasi optimal adaptive control law. Such an approach is used for instance in [7] for the control of biochemical reactors. In this paper, the methodology described above is applied in the context of congestion control of packet switched communication networks. The considered problem is the problem of finding the optimal control of a network buffer that gives the best trade-off between the number of lost packets and the average queue length.

In Section II the fluid flow model used throughout the text is described and its relevance in the context of stochastic queueing systems is illustrated. In Section III, the optimal control problem is solved and the optimal control strategy is shown to be of the bang-bang type with the possibility of a singular arc. It is worth mentioning that the same conclusion was already obtained in [6] using a very different approach. The minimisation of an similar cost function is also considered in [3]. However, the simplicity of our model allows for a much simpler solution of the problem. In Section IV a closed loop approximation of the optimal control is proposed. The needed parameters are identified on-line based on very simple network measurements : the average packet retention time and the number of output packets during the considered time interval. These parameters are always available, regardless of the traffic characteristics. In Section V, experimental results show a comparison between the proposed strategy and a tail-drop policy. It is shown that our strategy is nearly optimal and its adaptive feature makes it very attractive in the case of a time varying input stream.

This paper presents research results of the Belgian Programme on Interuniversity Attraction Poles, initiated by the Belgian Federal Science Policy Office. The scientific responsibility rests with its author(s).

II. FLUID FLOW MODEL OF A NETWORK BUFFER

Consider the system depicted in Fig. 1 showing a typical management scheme of a network buffer. In this setup, a quantity $w(t)$ of packets per second arrive in the system. Depending on the number $x(t)$ of packets already waiting in the buffer and and on the service rate μ of the network server, it may be desirable to drop a fraction of this incoming flow to prevent network congestion. The incoming flow $w(t)$ is therefore split into two flows: a flow $d(t)$ of dropped packets and a flow $u(t)$ actually fed to the network buffer.

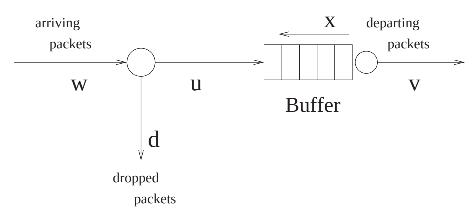

Fig. 1. A simple network buffer. Packets arrive at a rate $w(t)$, a quantity $d(t)$ of packets per second is lost and the remaining flow $u(t)$ is fed to the buffer.

The time evolution of the buffer load may be written :

$$x(t) = x(0) + U(t) - V(t)$$

where $U(t)$ and $V(t)$ are piece-wise constant (from $\mathbb{R} \to \mathbb{N}$) non-decreasing functions representing respectively the number of incoming and outgoing packets over the time interval $[0, t]$. In the fluid-flow paradigm of communication networks, counting processes such as $U(t)$ and $V(t)$ are replaced by a C^1 approximation (from $\mathbb{R} \to \mathbb{R}$) of the corresponding discrete processes. Therefore, the different rate functions $w(t), d(t), u(t)$ and $v(t)$ introduced above may be defined as the time derivative of their respective fluid flow counting processes. The dynamics of the buffer load is then rewritten as :

$$\dot{x} = u(t) - v(t) \tag{1}$$

In order to get a control system, we have obviously to describe the modelling of the flow rates v in terms of the state variable x and an appropriate control input u. To achieve this objective, we need to introduce the concept of processing rate, denoted $r(x)$:

$$v(t) = r(x(t)) \tag{2}$$

The construction of a suitable processing rate function is detailed in the next subsection.

A. The processing rate function

A number of processing rate functions may be found in the literature. For instance in [3] and [1], a relationship summarised as:

$$r(x(t), u(t)) = \begin{cases} \mu & \text{if } x > 0 \\ \min(\mu, u(t)) & \text{if } x = 0 \end{cases} \tag{3}$$

is proposed. If the buffer is fed with a constant average input rate $u(t) = \lambda$ smaller than μ, the system (1)-(2)-(3) has a single equilibrium at $x = 0$, regardless of the value λ.

Let us now, consider the following experiment : A network buffer with average service rate $\mu = 40$ packets per second [pps] is fed by a source with exponentially distributed inter-packet arrival time whose average rate jumps every $T = 10$ seconds from a value of 10 [pps] to a value of 15 [pps]. This experiment is carried out on a discrete event queue simulator and repeated five thousand times so as to obtain the ensemble average of the buffer load at regular time intervals. The result is shown in Fig. 2(A) (light curve). It is clear that the average buffer load is non-zero and increases when the average input rate increases. If the fluid flow input rate $u(t)$ is chosen as :

$$u(t) = \begin{cases} 10 & kT \leqslant t < (k+1)T \\ 15 & (k+1)T \leqslant t < (k+2)T \quad k = 0, 2, \ldots \end{cases} \tag{4}$$

then, the fluid model (1)-(2)-(3) produces a constant buffer load always equal to zero which is clearly not the observed behaviour.

Therefore we consider the following alternative family of processing rate functions [4]

$$r(x) = \frac{x}{\theta(x)} \quad \text{[pps]}$$

where $\theta(x)$ is defined as the average packet residence time. If we select a linear relationship between θ and the buffer level x, we obtain :

$$r(x) = \frac{\mu x}{a + x} \tag{5}$$

with $a > 0$, a parameter. For a constant input rate $\lambda = u(t) < \mu$, the system (1)-(2)-(5) has a single equilibrium given by :

$$\lambda = \frac{\mu x}{a + x} \quad \text{or} \quad x = \frac{a\lambda}{\mu - \lambda}$$

For $a = 1$, we then recover the classical formula of queueing theory for M/M/1 systems and μ can therefore

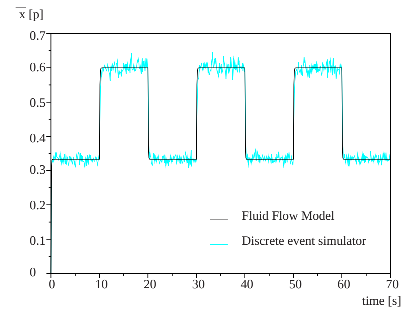

(A)

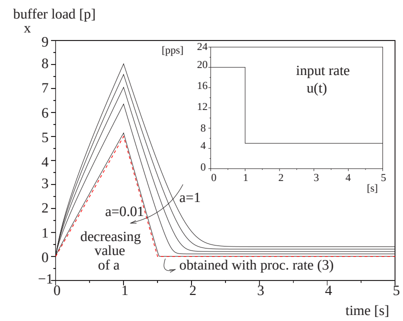

(B)

Fig. 2. (A) Comparison between a discrete event queue simulator and the integration of model (1)-(2)-(5) (B)Average buffer queue length for different value of the parameter a when the input rate is greater then the link capacity during one second.

be interpreted as the average service time. Integrating the system (1)-(2)-(5) with $a = 1$ and the input rate given by eq. (4) yields the curve shown in Fig. 2(A) (dark curve) which is compared with the curve obtained with the discrete event queue simulator (light curve). It is clear that our fluid model reproduces the desired behaviour.

Fig. 2(B) shows the comparison between the processing rate functions (3) and (5). Both models are integrated with $\mu = 15$ and the input u equal to

$$u(t) = \begin{cases} 20 & t < 1 \\ 5 & t \geqslant 1 \end{cases}$$

for different values of the parameter a. For a tending to zero, it may be observed that model (1)-(2)-(5) gives results which converge to the result of model (1)-(2)-(3).

Even though the model (1)-(2)-(5) does not have a discontinuous processing function, results presented in Fig. 2 show that this model can be made arbitrarily close to a model with a discontinuous processing rate function which are often used in the literature. Furthermore, results from the queueing theory suggest that the relationship, at equilibrium, between the queue length and the output rate should be of the form (2) rather than (3). Therefore, if fluid flow measures such as the average queue length or the average retention time are used, a model of the form (1)-(2)-(5) seems more appropriate than a model with a discontinuous processing rate function.

B. *Cost function*

The separation of the input stream $w(t)$ (Fig. 1) between a stream of dropped packets $d(t)$ and a stream $u(t)$ actually fed to the buffer is mainly motivated by two conflicting criteria : On one hand, the number of dropped packets in a network should be kept as small as possible to avoid retransmission but on the other hand, the transmission of large bursts requires large queues which increases the transmission delay. This trade-off is captured in the following function :

$$\begin{aligned}\mathcal{L}(x,t,u) &= x(t)+Rd(t)\\ &= x(t)+R(w(t)-u(t))\end{aligned}$$

with $R > 0$, a positive weight on the dropping rate. Therefore, we consider the cost function :

$$J(x,t_f,u) = \int_0^{t_f} \mathcal{L}(x,t,u)dt \quad (6)$$

The minimisation of this cost function is considered in the next section.

III. Optimal control

The optimal control problem can be stated as follows : Given the system

$$\dot{x} = f(x,t) = u(t) - \frac{\mu x}{a+x} \quad 0 \leqslant u(t) \leqslant w \quad (7)$$

with $d(t) = w - u(t)$
find the optimal control $u^*(t)$ which minimises the cost function (6) along the trajectories of system (7).

The solution of this minimisation problem is a direct application of the minimum principle of Pontryagin. The Hamiltonian (with costate p) is as follows :

$$\begin{aligned}\mathcal{H}(x,t,u) &= \mathcal{L}(x,t,u)+pf(x,t)\\ &= x(t)+R\Big(w-u(t)\Big)\\ &\quad +p\Big(u(t)-\frac{\mu x}{a+x}\Big)\end{aligned}$$

The optimal control u^* then satisfies the following conditions :

$$\begin{aligned}u^* &= arg.min_{0\le u(t)\le w}\mathcal{H}(x^*,t,u)\\ \dot{p} &= -1+p\frac{a\mu}{(a+x)^2}\\ \dot{x} &= f(x,t)\end{aligned}$$

A. *Minimisation of the Hamiltonian*

Minimising the Hamiltonian with respect to u yields

$$u^* = \begin{cases} 0 & p > R\\ w & p < R\\ \text{singular} & p = R\end{cases}$$

As already stated in [6], the optimal control of a fluid flow buffer is of the bang-bang type with the possibility of a singular arc. Defining $s(t) = p(t) - R$, a singular arc $[t_1, t_2]$ will be obtained if

$$\frac{d^i s}{dt^i} = 0 \quad \forall i \in \mathbb{N} \quad \forall t \in [t_1, t_2]$$

- $\frac{ds}{dt} = \frac{dp}{dt} = -1 + p\frac{a\mu}{(a+x)^2}$

$$\frac{ds}{dt}(x^*,p^*)_{p=R} = 0 \iff x^* = \sqrt{aR\mu} - a$$

- $\frac{d^2s}{dt^2} = \dot{p}\frac{a\mu}{(a+x)^2} + p\frac{-2a\mu}{(a+x)^3}\dot{x}$

$$\begin{aligned}\frac{d^2s}{dt^2}(x^*,p^*)_{p=R} = 0 &\iff \dot{x}_{x=x^*} = 0\\ &\iff u_{sing} = \frac{\mu x^*}{a+x^*}\\ &\qquad u_{sing} = \mu(1-\sqrt{\frac{a}{R\mu}})\end{aligned}$$

Therefore, the singular arc (of order 2) is characterised by

$$\dot{p} = \dot{x} = 0 \quad x_{sing} = \sqrt{aR\mu} - a \quad u_{sing} = \frac{\mu x_{sing}}{a + x_{sing}} \quad (8)$$

B. *Boundary conditions*

The formulation of the optimisation should be completed with suitable boundary conditions. These conditions are stated here with the idea that we are looking for the optimal control path when the input rate $w(t)$ represents a burst, that is to say a function of the type :

$$w(t) = \begin{cases} w & t \le t_{burst}\\ 0 & \text{otherwise}\end{cases} \quad (9)$$

The control period $[t_0, t_f]$ is therefore envisioned to be larger than the burst duration. The initial condition $x_0 = x(t_0)$ is given. Additional conditions depend on the formulation of the problem :

1) *Final time t_f fixed, $x(t_f)$ free*: In this situation, the additional condition is on the final value of the costate :

$$p(t_f) = 0$$

Given the shape of the input function $u(t)$ and the fact that t_f is greater than the burst duration t_{burst}, the final value of the state $x(t_f)$ is expected to be small. Therefore, the following alternative formulation may also be considered.

2) *Fixed final state value $x(t_f)$ with $x(t_f)$ small, t_f free.* The final time t_f is given by :

$$\mathcal{H}(x(t_f), t_f, u(t_f)) = 0$$

The integration of the dynamics of the system yields the final time t_f and the above relationship may then be used to determine the final value of the costate.

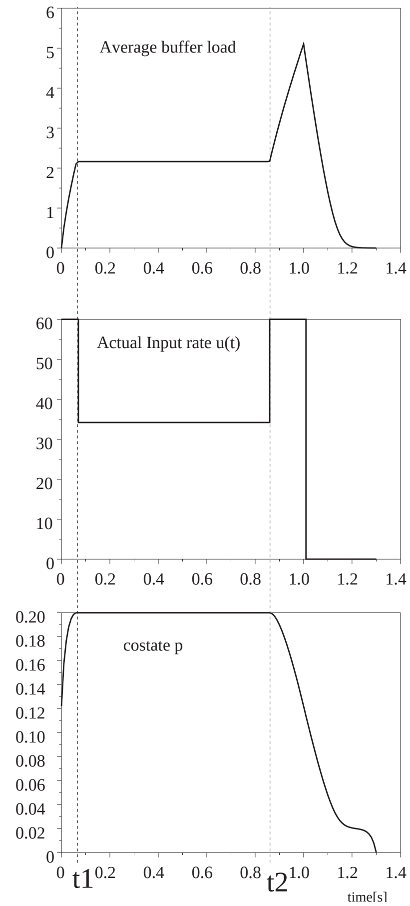

Fig. 3. Optimal control in the presence of a burst.

C. Example

Let us now consider an illustrative example of an M/M/1 queue fed at a rate $w(t)$ given by eq. (9) with $w = 60$ and $t_{burst} = 1$. The average service rate is set to $\mu = 50$. The optimisation problem is considered with a fixed time interval $[0, t_f = 1.3]$. The weight R is set to 0.2.

Results are shown in Fig. 3 where the average buffer length x, the optimal control u^* and the costate p are displayed. The optimal control presents three distinct stages :

In the time interval $[t_0, t_1]$, the control u is set to its maximal value to allow the state variable x to reach its singular value x^* as quickly as possible. Once the singular arc is reached, the system is controlled so that $\dot{p} = 0$ and $\dot{x} = 0$ and the system therefore stays on this arc during the time interval $[t_1, t_2]$. The singular arc must eventually be left in order to satisfy the boundary condition $p(t_f) = 0$. The time instant t_2 may be determined iteratively by "shooting" from selected initial instants.

With this profile, the optimal cost calculated with eq. (6) is 6.49. This number may be compared in Fig. 4 with different costs obtained with a tail-drop policy for different value of the threshold. With this policy the minimum

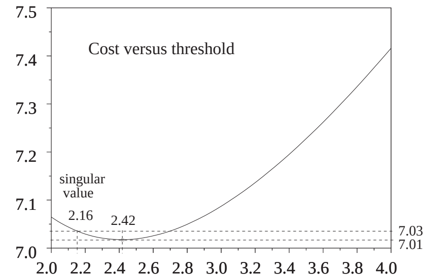

Fig. 4. Cost for different values of the threshold under tail-drop policy. The optimal profile gives an optimal cost of 6.49.

cost is given for a threshold set at 2.42 and is 7.01. A substantial improvement is obtained using the optimal profile. However, it must be noted that the computation of the time instant t_2 requires the a priori knowledge of the burst duration which is not known in practice. This problem is discussed further in the next section.

IV. Implementation of the optimal control

The computation of the optimal control requires the integration of the costate equation and the system dynamics. As pointed out in the previous section, the burst duration itself must be known to calculate the optimal profile. In order to obtain a useful control law, the optimal strategy should be described as a feedback of some state measurement which would make it robust to model uncertainties and should not rely on a-priori knowledge on the traffic characteristics.

A closed loop control is easily obtained if the final interval $[t_2, t_f]$ is neglected. In this case, the optimal strategy becomes identical to the tail-drop policy which only requires the measurement of the state x. The control law therefore reduces to the tracking of the singular value x_{sing} given by eq. (8). For the example given above, $x_{sing} = 2.16$. The effect on the cost may be observed in Fig. 4. It can be seen that this value yields a suboptimal cost which lies in the vicinity of the best cost obtained under tail-drop policy.

An idea of the importance of the neglected interval is obtained if we consider a burst with t_{burst} greater than t_f. In that case the processing rate function (5) is approximated by $r(x) = \mu$ and the costate evolution by $\dot{p} = -1$. An explicit value for t_2 is $t_2 = t_f - R$ which indicates that the singular arc should be left earlier if the weight on the dropping rate is increased. If the last interval is to be neglected, the parameter R should therefore be small compared to a typical burst duration. Note that, if short burst occur and if an important weight is set on the dropping rate, a sensible policy would be to accumulate the entire burst in the buffer and prevent any overflow to occur.

Our sub-optimal control can now be stated :

1) Obtain fluid-flow measures of needed variables:
   - $\hat{x}$: Estimate of the average buffer length

- $\hat{\lambda}$: Estimate of the rate of packets through the buffer

2) Obtain an estimate $\hat{a}$ of the parameter a
3) Compute the singular values: $x_{sing} = \sqrt{R\mu\hat{a}} - \hat{a}$ and $u_{sing} = \mu(1 - \sqrt{\frac{\hat{a}}{R\mu}})$
4) if $\hat{\lambda} > u_{sing}$, drop packets so as to control $\hat{x}$ at its singular value x_{sing}

Each of these steps are now described in the following subsections

A. Fluid flow measures

Fluid-flow variables are measured with a sliding window of length Δ [sec]. During these Δ seconds, the number of outgoing packets N and the sum τ of the seconds spent by each packet in the system is recorded. Variables are estimated as follow :

$$\begin{aligned} \hat{\lambda} &= \frac{N}{\Delta} && \text{: average rate} \\ \hat{T} &= \frac{\tau}{N} && \text{: average retention time} \end{aligned}$$

The average buffer length is then calculated using the Little's formula[5] :

$$\hat{x} = \hat{\lambda}\hat{T} = \frac{\tau}{\Delta}$$

B. On-line model identification

The computation of the singular value x_{sing} requires the knowledge of the parameter a which depends on the stochastic properties of the traffic. It is therefore necessary to estimate this parameter on the basis on some traffic measurements.

Given a set of K measurement vectors $p_k = (\hat{x}_k, \hat{\lambda}_k), k = 1, \ldots K$, it is easy to show that the value a_{est} that best fits the processing rate (5) in the sense

$$a_{est} = arg.min_a \sum_{i=1}^{K} \big(\frac{\mu\hat{x}_i}{a + \hat{x}_i} - \hat{\lambda}_i\big)^2$$

is given by :

$$\begin{aligned} a_{est} &= \frac{\mu \sum_{i=1}^{K} \hat{x}_i - \sum_{i=1}^{K} \hat{x}_i\hat{\lambda}_i}{\sum_{i=1}^{K} \hat{\lambda}_i} \\ &= \frac{\mu\Psi_x^K - \Psi_{x\lambda}^K}{\Psi_\lambda^K} \end{aligned}$$

where the variable $\Psi_x^K, \Psi_{x\lambda}^K, \Psi_\lambda^K$ are introduced for convenience. Therefore, the following method for the estimation of the parameter a is proposed :
Every Δ seconds compute

$$\begin{aligned} \Psi_x^K &= \Psi_x^K + \hat{x} \\ \Psi_{x\lambda}^K &= \Psi_{x\lambda}^K + \hat{x}\hat{\lambda} \\ \Psi_\lambda^K &= \Psi_\lambda^K + \hat{\lambda} \\ a_{est} &= (\mu\Psi_x^K - \Psi_{x\lambda})/\Psi_\lambda \\ \hat{a} &= \hat{a} + h(a_{est} - \hat{a}) \end{aligned}$$

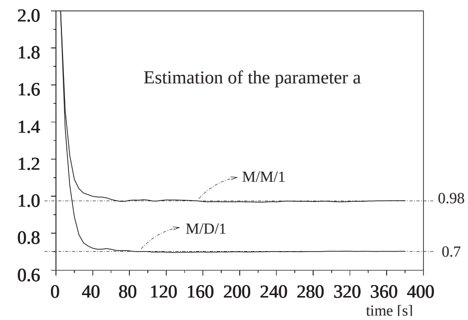

Fig. 5. On-line estimation of the parameter a of the proc. rate (5) in the M/M/1 and M/D/1 case. Initial estimation is set to $\hat{a}(0) = 2$

The successive estimate a_{est} are filtered with a gain $h < 1$ to avoid rapid change between estimations of a. The parameter a is evaluated every Δ which is the value used for the estimation of $\hat{x}$ and $\hat{\lambda}$. This is not required but minimises the number of parameters used for the algorithm. This method is illustrated in Fig. 5 showing the convergence of the estimated value of a toward its theoretical value. This figure is obtained with the OMNET++ [8] simulator with the following parameters :

$$\Delta = 5[s], \mu = 500[pps], h = 0.5$$

In order to obtain a suitable excitation of the system during the identification, the input rate is changed every 10 seconds. The following two values are used recursively:

$$1/0.003, 1/0.005$$

In the M/M/1 case, the estimation $\hat{a}$ converges toward 0.98 which is close to the theoretical value 1. In the M/D/1 case, the relationship that links the number of customers in the system as a function of the rate λ is given by

$$\frac{\lambda}{\mu} = (1 + x) - \sqrt{1 + x^2} \qquad (10)$$

It is easy to verify numerically that the value of a that minimises the distance between (5) and (10) is close to 0.63. The convergence of $\hat{a}$ toward 0.7 indicates that the system has been successfully self-adapted to the M/D/1 situation.

C. Adaptive threshold

In order to regulate the average queue length at its singular value, an adaptive threshold c is used. This threshold is used at every packet arrival to decide if the packet should be enqueue or dropped.

The parameter c is updated, once again, every Δ seconds, so as to keep the average buffer length within ten percent of the tracked value : the operation is as follows :

if $(\hat{\lambda} > u_{sing})$

$$\begin{aligned} &\text{if } \hat{x} > 1.1x_{sing} : \quad c = c - 1 \\ &\text{if } \hat{x} < 0.9x_{sing} : \quad c = c + 1 \end{aligned}$$

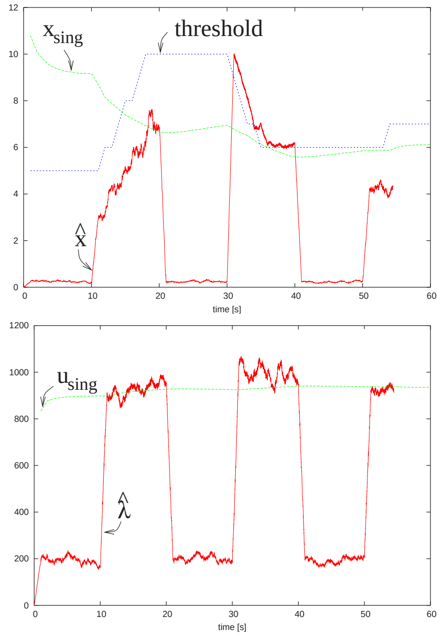

Fig. 6. top:Average buffer length $\hat{x}$, tracked value x_{sing} and adaptive threshold c. Bottom: Estimated throughput $\hat{\lambda}$ and singular value u_{sing}. The adaptive threshold is modified so as to bring $\hat{x}$ within ten percent of the tracked value.

With this algorithm, the control strategy is triggered only when the rate through the buffer is larger than the singular value. The adaptive threshold will maintain the average buffer load around the singular value x_{sing}.

V. Simulation results

Experimental results are obtained with the OMNET++ simulator. The source uses exponentially distributed inter-packet delay (Poisson source) whose average changes periodically, every $T = 10$ seconds. The values 0.005,0.0009, 0.005 and 0.0005 are used recursively to obtain some rates of 200, 1111, 200 and 2000 [pps].

The service time is also exponentially distributed, its average is set to 0.001 ($\mu = 1000$ [pps]). The fluid-flow variables are updated every $\Delta = 1$ [s]. The weight on the dropping rate is set to $R = 0.1$ and the filter gain is $h = 0.5$.

The average buffer load $\hat{x}$, the singular value x_{sing} and the adaptive threshold c obtained with the control strategy presented in this paper are displayed in Fig. V. The estimated throughput $\hat{\lambda}$ is also displayed with the calculated singular value u_{sing} showing the time interval where the control law is active.

It can be verify that the threshold is successfully adapted so as to bring the average buffer load in the vicinity of the tracked value.

Fig. 7 offer a comparison between the cost obtained with the adaptive threshold and the different costs obtained with a constant threshold. As expected, our adaptive control law operates the system so as to yield a cost approaching the best cost obtained with a constant threshold.

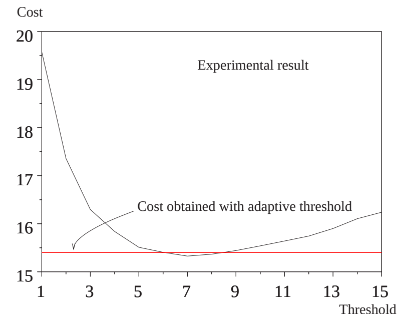

Fig. 7. Cost obtain with the adaptive threshold control compared to the costs obtained with a constant threshold for different values of this threshold.

VI. Conclusion

A new fluid-flow model has been presented. This model has the particularity to represent a wide class of stochastic queueing systems depending on the value of a single parameter. An on-line identification algorithm has been described in order to adjust the value of this parameter. This model has then been used to derive an optimal control policy that minimises the sum of the retention time and the number of dropped packets over a finite time interval.

This strategy has then been shown to be implementable in closed loop by measuring two easily observable parameters: the sum of the outgoing packets over the given time interval and the average retention time.

Experimental results have shown that the implementation of this strategy, although suboptimal, is nevertheless nearly optimal for a wide range of experimental conditions.

References

[1] Jean-Chrysostome Bolot and A. Udaya Shankar. Analysis of a fluid approximation to flow control dynamics. In *INFOCOM (3)*, pages 2398–2407, 1992.

[2] A. Bryson and Y-C Ho. *Applied Optimal Control*. Blaisdell, 1969.

[3] C. G. Cassandras, Y. Wardi, B. Melamed, G. Sun, and C. G. Panayiotou. Perturbation analysis for online control and optimization of stochastic fluid models. *IEEE transaction on automatic control*, 47(8):1234–1248, 2002.

[4] V. Guffens, G. Bastin, and H. Mounier. Using token leaky buckets for congestion feedback control in packets switched networks with guaranteed boundedness of buffer queues. In *proc. of ECC03 Cambridge, UK*, 2003.

[5] L. Kleinrock. *Queueing Systems, Volume 1: Theory*. Wiley & Sons, 1975.

[6] S. Rajagopal, V.G. Kulkarni, and S. Stidham. Optimal flow control of a stochastic fluid-flow system. *IEEE journal on selected areas in communications*, 13(7):1219–1228, 1995.

[7] I. Smets, G. Bastin, and J. Van Impe. Feedback stabilisation of fed-batch bioreactors for microbial conversion processes with non-monotonic kinetics. *Biotechnology Progress*, 18(5):1116 – 1125, 2002.

[8] Andras Varga. Omnet++ simulator. Available at http://www.omnetpp.org/.