A Fuzzy Decision Maker for Source Traffic Control in High Speed Networks

V. Catania, G. Ficili, S. Palazzo, D. Panno
Istituto di Informatica e Telecomunicazioni, University of Catania
V.le A. Doria 6, 95125 Catania - ITALY

Abstract

Most of the policing techniques proposed so far in high speed networks using the ATM technique are based on conventional approaches which use a crisp decision making logic. They don't meet the selectivity and responsivenes requirements needed to make a policing efficient. In this paper we propose a policing mechanism based on fuzzy logic. The main characteristics of the proposed fuzzy policer are simplicity and the capacity to combine a high degree of responsiveness with a selectivity close to that of an ideal policer. Moreover, it can easily be implemented in hardware, thus enhancing both cost and processing performance. The simulation results reported show that the performance of our fuzzy policer is much better than that of conventional policing mechanisms.

1. Introduction

Policing in high speed networks based on ATM technology is entitled to ensure that each source conforms to its negotiated parameters[1]. The policing function can be defined as the set of actions taken by the network, during the entire phase of the call, to monitor and control the offered traffic, with the purpose of protecting network resources from malicious as well as unintentional misbehaviour which can affect the QoS of already established connections, by detecting violations of negotiated parameters and taking appropriate action. This action can be either cell-dropping, cell-marking, or shaping the source rate.

The policing function should fulfill the following basic requirements:

1) High selectivity with respect to the traffic monitored (that is, the capability of detecting any illegal traffic situation and transparency for connections that respect the parameter values negotiated, on whose cells no policing action need be taken);
2) high responsiveness, that is, low response time to parameter violations;
3) simplicity of implementation and cost effectiveness.

Several problems arise in defining an efficient policing mechanism. A key issue is defining a traffic enforcement mechanism which can meet requirements 1, 2 and 3 at the same time.

In literature several mechanisms such as the leaky bucket and window mechanisms have been proposed, an extensive review of which can be found in [2]. However, for none of them has it been possible to achieve the ideal policing behaviour, but only a tradeoff between the above-mentioned conflicting requirements.

The difficulty of characterizing a policer accurately if "traditional" methods and models are used led us to explore alternative solutions based on soft computing techniques, specifically in the field of fuzzy systems.

Fuzzy systems are suitable for approximate reasoning, above all in systems for which it is difficult, if not impossible, to derive an accurate mathematical model. Imprecision or uncertainty can, for instance, affect the input values or parameters of the system, as well as the inference rules which characterize the control algorithm. In such cases, fuzzy logic is a powerful tool which allows us to represent qualitatively expressed control rules quite naturally, often on the basis of a simple linguistic description.

In addition, when applied to appropriate problems, fuzzy systems have a faster and smoother response than conventional systems, also thanks to the fact that fuzzy control rules are often simpler and do not require great computational complexity. The latter aspect, along with the spread of VLSI hardware structures dedicated to fuzzy computation, makes fuzzy systems cost-effective [3-4].

In this paper we propose a policing mechanism based on fuzzy logic, the main characteristics of which are a performance level which is decidedly better than that obtainable with conventional mechanisms and a great flexibility in maintaining its efficiency even when the source statistical characteristics vary.

In defining the fuzzy policer the aim was to detect violations of the parameters negotiated, regardless of the action that could subsequently be taken against cells which turn out to be excessive.

0-8186-7216-1/95 $04.00

The philosophy on which the mechanism is based exploits the fuzzy logic capability to deduce a system model on the basis of linguistic variables, fuzzy sets and fuzzy inferences. This allowed the rules of behaviour of a policer - expressed in approximate terms but at the same time corresponding to an expert description - to be translated into a rigorous fuzzy inferential system, the performance of which is much better than that of most methods proposed so far.

Another advantage of the approach proposed is that the inferential system which describes the policer is very simple and can be implemented in hardware, thus improving both costs and processing speed.

The paper is organized as follows. In Sect. 2 an overview of the most significant policing mechanisms already proposed in literature is presented. In Sect. 3 we introduce a proposal for a fuzzy logic based policing mechanism. Sect. 4 contains a performance evaluation of the proposed solution and a comparison with other mechanisms. Evaluation and implementation suggestions are given in Sect. 5. Finally, in Sect. 6 some conclusions are drawn.

2. Related work

In literature, policing mechanisms have mainly been evaluated against their capability to enforce such traffic parameters as mean and peak bit rate. Policing of the peak rate is generally not complex and can be achieved, for example by using a cell spacer or other mechanisms. Enforcement of the mean rate is more problematic, since short-term statistical fluctuations of the source traffic are admissible as long as the source respects the average value negotiated, λ_n, in the long term. With respect to this, comparative performance studies have been carried out, referring to the ON-OFF or bursty source model, which is widely accepted as the worst-case traffic pattern for this purpose. This source model is very suitable for the representation of packetized voice, still images and interactive data services.

The most popular and simple policing mechanism is the Leaky Bucket (LB) [5-6]. It is based on the concept of pseudoqueue and consists of a counter which is increased on arrival of the cells and decreased, if positive, at a constant frequency (depletion rate of the pseudoqueue server, λ_e). When the counter exceeds a pre-established threshold N (length of the pseudoqueue or counter limit), the cells are detected as excessive and the policing action agreed on is taken.

The parameters for the sizing of the LB are the threshold N and the depletion rate, λ_e. In principle, when enforcing the mean cell rate negotiated, it should be assumed that $\lambda_e=\lambda_n$. The choice of N plays a very important role. As certain statistical fluctuations around the average value negotiated are allowed in the cell rate, N has to be long enough to reduce the false alarm probability, that is, the probability of detecting some cells of a non-violating source as "excessive". This requirement is met when N values are high, but the reaction time of the mechanism grows excessively. From the analysis in [2, 7-8] it has emerged that in order to achieve greater flexibility in size and reduce the probability of false alarms it is necessary to introduce an overdimensioning factor C ($C\geq 1$) between the negotiated cell rate, λ_n, and that which is really policed, λ_p; it follows that $\lambda_e=\lambda_p=C\lambda_n$. On the other hand, this artifice reduces the capacity to detect violation over a long term. In the extreme case of a deterministic source, the traffic generated may even exceed the negotiated cell rate up to a factor C without any cell being detected as excessive. In spite of its pitfalls, the LB mechanism is still regarded as particularly attractive due to its simplicity of implementation.

Other control methods are the window-based ones [2, 9]. Among of these the Exponentially Weighted Moving Average (EWMA) exhibits the best performance. In it the number, N_i, of cells allowed in a fixed time interval, T (window), is dynamically updated as a function of the average value N of cells allowed per window and a sum of terms which take the past into account, i.e. the number of cells accepted in the previous intervals. More specifically, the contribution of each term decreases exponentially as the windows get further away in time. The rule for calculation of N_i in the i-th window is:
$N_i=(N-\gamma S_{i-1})/(1-\gamma)$, $0<\gamma<1$, where $N=\lambda_p T=C\lambda_n T$, $S_{i-1}=(1-\gamma)\cdot x_{i-1}+\gamma\cdot S_{i-2}$, and x_{i-1} is the number of cells accepted at the (i-1)-th window.

The γ parameter is a constant weighting factor which makes the mechanism more or less flexible with respect to the burstiness of the traffic. If $\gamma=0$, N_i is constant and the algorithm is thus the traditional Jumping Window. A value of $\gamma > 0$ makes the mechanism more tolerant of statistical fluctuations (larger bursts are allowed), thus reducing the false alarm probability; however, when γ increases, responsiveness decreases. EWMA performance greatly depends on the value chosen for γ: for 'respectful' sources a high value is needed; vice versa, for violating sources it has to be low. The limit of this mechanism is therefore the static, a priori choice of the value of γ. Once a tradeoff value has been chosen for γ, the false alarm probability can only be reduced by increasing T. This, however, would further reduce responsiveness. The problem can be avoided by introducing, as with the Leaky Bucket, the overdimensioning factor C which, with a fixed pair of values for T and γ, gives lower false alarm probability values. The EWMA mechanism is more complex than the previous one.

None of the above mechanisms are able to cope efficiently with the conflicting requirements of ideal policing, that is, a low false alarm probability and high

responsiveness. Their limits are due to the fact that they control stochastic magnitudes, which are allowed to fluctuate around an average value, by means of crisp-threshold mechanisms (i.e., the counter limit value, the γ parameter value).

A challenging alternative is to use innovative solutions based on artificial intelligence techniques.

In [10] two Artificial Neural Networks (ANNs) are trained to predict the count process of arriving cells. One network is trained using source patterns which do not violate the parameters negotiated, while for the other both violating and non-violating sources are used. In this method prediction of the count process of the next window is based on the values it took in the M previous windows. This means that each of the two networks requires M inputs. By comparing the outputs of the two networks, an error signal is determined which is null if the source being examined is not violating the parameters negotiated and positive if it is.

With a view to finding solutions based on artificial intelligence and taking advantage of the merits of soft-computing, a novel solution based on the use of fuzzy logic is investigated in this paper.

We were led to use fuzzy logic by analysis of the limits of traditional mechanisms. These are based on conventional algorithms which, from a decisional point of view, are crisp decision makers: the decision to consider arriving cells as excessive or not is, in fact, a crisp logic evaluation over a set of fixed thresholds. This leads to a poor dynamic action on the part of the control mechanism, which in turn can cause extremely nonlinear behaviour. The adoption of fuzzy logic as the decison-making logic, on the other hand, allows the use of algorithms which are soft decision-making: the evaluation of whether a the source respects the parameters negotiated or not is represented by a truth value which is not restricted to either false (the truth value is zero) or true (the truth value is one), but in a continuum of [0,1]. The softness of truth values, which is inherent in the concept of fuzzy sets, along with the power of expression of fuzzy inferential systems, promises a more appropriate representation of the decision process which a policing mechanism has to feature.

3. Fuzzy Logic Applied to Policing

Fuzzy logic [11] is based on the concepts of linguistic variables and fuzzy sets. A fuzzy set in a Universe of Discourse U is characterized by a membership function m_f which assumes values in the interval [0,1]. A fuzzy set F is represented as a set of ordered pairs, each made up of a generic element $u \in U$ and its degree of membership $m_f(u)$.

A linguistic variable x in a Universe of Discourse U is characterized by a set $W(x)=(W_{1x}, \ldots\ldots W_{nx})$ and a set $M(x)=(M_{1x},\ldots,M_{nx})$, where W(x) is the term-set, i.e. the set of names the linguistic variable x can assume, and W_{ix} is a fuzzy set whose membership function is M_{ix}. If, for instance, x indicates a temperature, W(x) could be the set W(x)=(Low, Medium, High), each element of which is associated with a membership function.

The rules governing a fuzzy system are often written using linguistic expressions which formalize the empirical rules by means of which a human operator is able to describe the process in question using his own experience. If x and y are taken to be two linguistic variables, fuzzy logic allows these variables to be related by means of *fuzzy conditional rules* of the following type:

"IF (x is A) THEN (y is B)"

where *(x is A)* is the *premise* of the rule, while *(y is B)* is the *conclusion*. This rule makes it possible to deduce, using specific inferential methodologies, a fuzzy set for y for each input value of x, whether it is associated with a fuzzy set or assumes a numerical value (*crisp*).

The fuzzy policer proposed is a window-based control mechanism in which the maximum number N_i of cells that can be accepted in the i-th window is dynamically updated by inference rules based on fuzzy logic.

The aim is to make a generic source respect the average cell rate negotiated, λ_n, i.e. to ensure that on average the source transmits N cells per window with $N=T\lambda_n$. We assume that the peak cell rate is separately controlled.

The philosophy on which our mechanism is based is one of granting credit to a source which in the past has respected the parameters negotiated by increasing its control threshold N_i as long as it perseveres with non-violating behaviour. Viceversa, if the behaviour of the source is violating or risky, the mechanism reduces its credit by decreasing the threshold value.

The parameters describing the behaviour of the source and the policing control variables are made up of linguistic variables and fuzzy sets, while control action is expressed by a set of fuzzy conditional rules which reflect the cognitive processes that an expert in the field would apply.

The source descriptor parameters used are: the average number of cell arrivals per window since the start of the connection, A_{oi}, and the number of cell arrivals in the last window, A_i. The first gives an indication of the long-term trend of the source; the second indicates its current behaviour. A third parameter, the value of N_i in the last window, was also introduced to indicate the current degree of control (degree of permissiveness) the mechanism has over the source. These parameters are the three linguistic variables which make up the fuzzy policer input.

The output chosen was the linguistic variable ΔN_{i+1} which represents the threshold variation to be made in the next window.

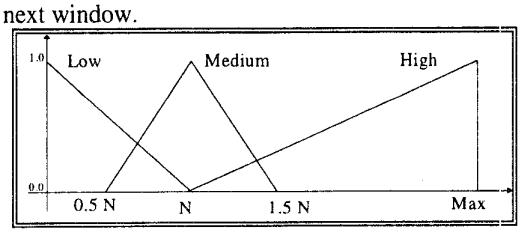

Fig. 1: Membership Functions for the A_{oi} and A_i input variables.

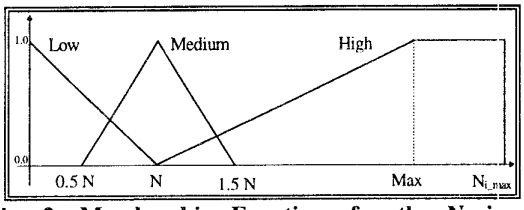

Fig. 2: Membership Functions for the N_i input variable.

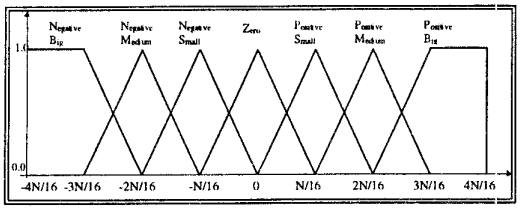

Fig. 3: Membership Functions for the ΔN_{i+1} output variable.

The model of the fuzzy system, comprising the control rules and the term sets of the variables with their related fuzzy sets, was obtained through a tuning process which started from a set of initial insight considerations and progressively modified the parameters of the system until it reached a level of performance considered to be adequate.

In particular, the term sets of input variables have the following fuzzy names: Low (L), Medium (M) and High (H).

The term set of the output variable is composed of seven fuzzy sets with the following fuzzy names: Zero (Z), Positive Small (PS), Positive Medium (PM), Positive Big (PB), Negative Small (NS), Negative Medium (NM), Negative Big (NB).

The membership functions chosen for the fuzzy sets are shown in Figs. 1, 2 and 3.

MAX represents the maximum value between 1.5 N and the maximum number of cells that can arrive in a window (T/t_c), where t_c is the cell interarrival time during a burst. N_{i_max} indicates the upper bound value for the N_i variable.

	A_{oi}	N_i	A_i	ΔN_{i+1}
1	L	H	L	PB
2	L	H	M	PS
3	L	H	H	Z
4	M	M	L	PB
5	M	M	M	PS
6	M	M	H	Z
7	M	H	L	PB
8	M	H	M	Z
9	M	H	H	NB
10	H	L	L	PB
11	H	L	M	PM
12	H	L	H	PS
13	H	M	L	PB
14	H	M	M	PM
15	H	M	H	Z
16	H	H	L	NS
17	H	H	M	NM
18	H	H	H	NB

Table 1: Fuzzy Policer Rules

Table 1 shows the knowledge base of our fuzzy policer. By way of illustration, the Rule 1 in Table 1 has to be read as: *If (A_{oi} is low) and (N_i is high) and (A_i is low) then (ΔN_{i+1} is positive big).*

To make the fuzzy policer's knowledge base easy to understand, below we discuss the three cases in which the source is fully respectful (*A_{oi} is low*), moderately respectful (*A_{oi} is medium*) and violating (*A_{oi} is high*), respectively.

1) N_i is necessarily high due to the fact that the source has gained credit. Thus if the number of cells which arrived in the last window is medium or high, that is, the source continues non-violating behaviour, its credit is increased (rules 1, 2); vice versa, if *A_i is high*, a sign of a possible beginning of violation on the part of the source or an admissible short-term statistical fluctuation, the threshold value remains unchanged (rule 3);
2) We distinguish between two subcases:
   - N_i is *medium*: the choice of ΔN_{i+1} is based on the same logic as before (rules 4-6).
   - N_i is *high:* This indicates a steady-state situation due to a respectful source or a transient situation due to a source which is starting to violate; the choice of ΔN_{i+1} is greatly influenced by the variable A_i as can be seen in rules 7, 8 and 9.
3) We distinguish between three subcases:
   - N_i *is low*: the threshold must be brought back immediately to values close to N to avoid excessively rigorous policing which would raise the false alarm probability (rules 10-12);
   - N_i *is medium*: without doubt this is a steady-state situation in which the source is violating and the threshold has therefore settled around N. Here again it may make sense to increase the source credit (rules 13, 14) to be able to cope with the

situation correctly if the arrivals in the current window are medium-low;

- N_i *is high*: this situation occurs when the source starts to violate in the initial stages of the connection. The threshold value has to be lowered immediately and the choice of the consequents is therefore clear (rules 16-18).

In order to determine N_I, the value for N_i in the first window, and $N_{i\ max}$ tuning was performed, by means of several simulation runs, which led to the following choice: N_I=3.5N and $N_{i\ max}$=9N. It should be noted that the fuzzy policer model is parametric with respect to the values MAX, N_{i_max} and N_I which are a function of N, t_c and T. This allows us to use the same model for bursty sources with different statistical properties, as will be shown in the next section.

4. Performance Evaluation

In this section we assess the efficiency of the fuzzy policer proposed.

We will compare the performance of the fuzzy policer with that of the LB and the EWMA as these mechanisms have been recognized as better than others [2]. The figures of merit considered are the selectivity and responsiveness of the mechanisms.

To assess the performance of the three mechanisms software simulators were developed in C language. More specifically, to evaluate the fuzzy inferences an inferential engine was implemented using a discrete Universe of Discourse with 256 points, MAX-MIN as the inference method and Center of Gravity as the defuzzification technique. The simulation results given were obtained with a 95 per cent confidence interval of the true value. This interval is not given in the figures as it is so slight, and omission makes the values easier to read.

To be able to compare the three mechanisms fairly, we referred to the bursty source characteristics studied by Rathgeb [2]. Therefore, we assumed that the number of cells per burst has a geometric distribution with a mean of E[x]=5 cells; the duration of the idle phase has a exponential distribution with a mean of E[s]=0.14772 sec; and the intercell time during a burst is t_c=0.016 sec. So, the cell arrival rate negotiated is λ_n=22 cell/s. In [2], policing mechanisms were devised to enforce the mean rate by selecting N=45; in addition, to achieve a low false alarm probability with this value of N, C=1.42 for the LB and C=1.42 with γ=0.8 or C=1.27 with γ=0.91 for the EWMA were derived.

Let us note that for the EWMA mechanism with C=1,42 and γ=0.8, the window size is $T=N/C\lambda_n$=1.44 s. Our fuzzy policer was dimensioned by assuming a window size of T=1.44 s; considering that in our case C=1 can be assumed, the N value is equal to 32 cells. This window size corresponds, on average, to six burst/silence patterns.

First, let us focus our attention on selectivity. If we indicate with P_d the probability that the policing mechanism detects a cell as excessive, the ideal behaviour would be that P_d is zero with the mean cell rate up to the nominal one (the false alarm probability is zero), and $P_d = (\sigma - 1)/\sigma$ for $\sigma > 1$, where σ is the long-term actual mean cell rate of the source normalized to the negotiated mean cell rate.

In order to obtain the curve P_d versus σ, we assume that a variation in the cell rate is due to a change in the average number of cells per burst, while the average silence time is assumed to be constant (E[s]=0.14772 sec). As shown in Fig. 4, our fuzzy policer presents a null false alarm probability ($\sigma \leq 1$) and, in the case of violating sources ($\sigma > 1$), a probability of detection of violation very close to ideal and certainly much greater than that of the other policing methods. For example with σ=1.1 there is an improvement for P_d of over one order of magnitude as compared with the other mechanisms.

We also compared the dynamic behaviour of the mechanisms, i.e. we evaluated their responsiveness in terms of the average number of cells emitted by a violating source before the policer takes control action. We considered sources with an average rate 50% higher than the negotiated rate, i.e. σ=1,5.

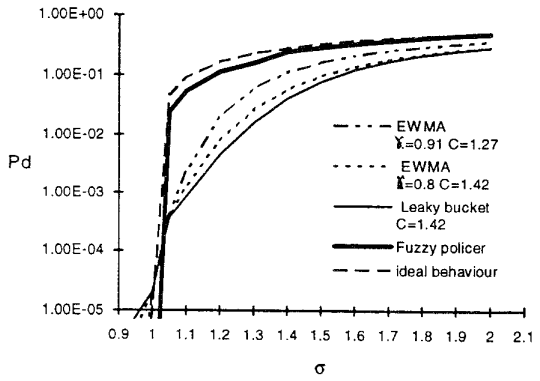

Fig. 4: Selectivity performance versus cell rate variations.

From Figs. 4 and 5, a comparison of traditional mechanisms shows that the mechanism which has the best behaviour towards long-term violations (namely EWMA with γ=0.91 and C=1.27), is the worst as far as responsiveness performance is concerned; the opposite holds for the LB. This confirms the fact that traditional mechanisms are not able to cope efficiently with the conflicting requirements of ideal policing, that is, a low false alarm probability and high responsiveness.

This problem is not encountered with our fuzzy policer. In fact, a trend very close to the ideal curve in the steady

state corresponds to decidedly better dynamics than those of the other mechanisms. More specifically, although the LB starts detecting violation after only 150 cells, the percentage of cells detected as excessive is very low, in the range of 5%, as compared with an ideal detection probability of about 33%. Although our mechanism takes control action after about 450 cells, its detection probability grows very fast, and, after only 1500 cells it reaches a value of about 0.20, thus showing a marked improvement over the other policing methods.

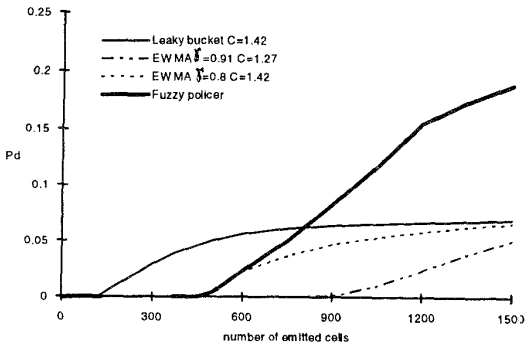

Fig. 5: Dynamic Behaviour.

In order to evaluate the fuzzy policer's capability to react to different kinds of violations, we made the mean cell rate vary by increasing the average silence time E[s], and keeping the number of cells per burst constant (E[x]=5 cells).

As the results given in Figs. 6-a and 6-b show, this kind of violation is harder to detect with traditional policing methods, as was also pointed out by Rathgeb [2]. Our mechanism, on the other hand, is robust and efficient regardless of the kind of violation.

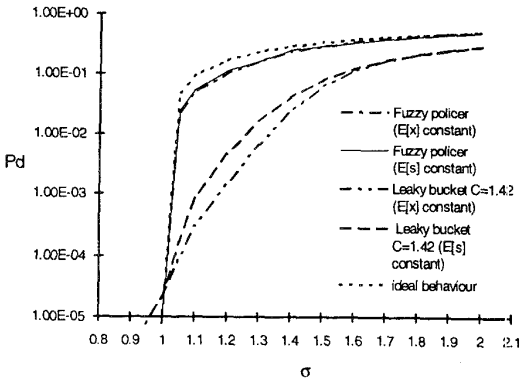

Fig. 6-a: Selectivity performance versus different kinds of cell rate variations

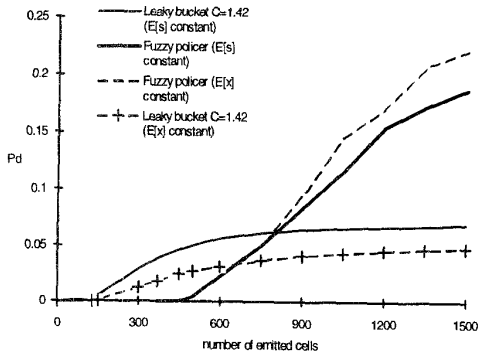

Fig. 6-b: Dynamic Behaviour with different kinds of cell rate variations.

To illustrate the field of application of the mechanism, we assessed its performance in the case of a real bursty source, namely a packetized voice source. Considering a peak bit-rate of 32 Kbit/s and an ATM cell size of 53 bytes, we have a cell interarrival time t_c of 12 ms. In addition, the nominal values of the traffic parameters are typically: E[x]=29 cells, E[s]=650 ms.

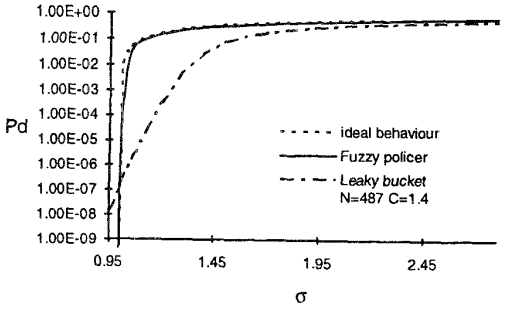

Fig. 7-a: Selectivity in control of a voice source.

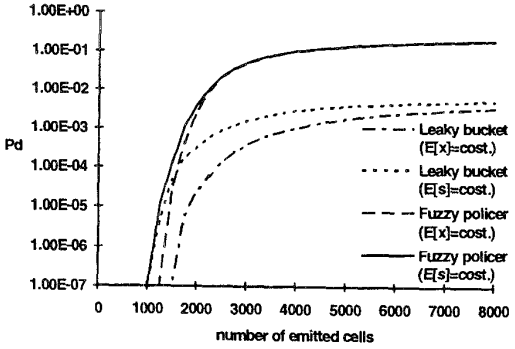

Fig. 7-b: Dynamic behaviour in policing of a violating voice source with σ=1.5

It is worth pointing out that for this kind of source traditional policing methods proved to be inefficient [2].

To achieve a low false alarm probability, it is necessary to have either a high value for the counter limit, N, which means a poor dynamic response, or a high value for the overdimensioning factor, C>2, which reduces the capacity to detect a violation. For example, we have compared the fuzzy policer with a LB, sizing both mechanisms for guaranteeing a given false alarm probability $P_d=10^{-7}$ for $\sigma=1$. For our mechanism we were able to choose T equal to three burst/silence patterns.

The simulation results, depicted in Fig. 7-a and 7-b, show that the fuzzy policer keeps a good preformance level for both selectivity and responsiveness in control of real sources.

In conclusion, the performance results obtained have shown that:

- The fuzzy policer offers performance levels which are decidedly better than those obtainable with conventional mechanisms. In addition, it maintains its efficiency even when the source characteristics vary, managing - unlike other mechanisms - to achieve a selectivity close to the ideal.
- The parametric model adopted for the fuzzy policer gives it a highly general nature, thus making application to different types of bursty sources quite straightforward.

5. Implementation Issues

One of the requirements a policing mechanism has to meet is simplicity, a fundamental property to ensure feasible, cost-effective implementation.

In this perspective, we will now give some elements to evaluate the complexity and costs our mechanism entails, when implemented in hardware, if a certain level of performance is to be guaranteed.

Recently there has been a spread of solutions on VLSI chips which allow fuzzy inferences to be hardware-computed [12]. There are also several architectural proposals which allow the hardware resources to be implemented on chips to be chosen in such a way as to obtain a certain cost/performance ratio [3, 4, 13].

In our evaluation we refer to the architecture proposed in [4], as illustrated in Fig. 8. A set of Fuzzy Computation Units (FCUs) operates in parallel to process the fuzzy application rules stored in the Rule Memory. The Knowledge Base Memory stores all the membership functions relating to the term sets of the input and output variables of the rules. The architecture is designed in such a way that the number of FCUs can be chosen according to the inference processing speed required.

The Control Unit maps the rules to be processed on the FCUs and also synchronizes the various blocks in the architecture. Lastly, the Defuzzifier block, fed by the single FCUs, deals with defuzzification. To estimate the cost of implementing the fuzzy policer proposed on the architecture shown in Fig. 8, it is necessary to estimate the cost of memorizing the rules and the Knowledge Base, and the number of FCUs needed to obtain a certain processing speed. The memory cost can be estimated assuming a storage technique like the one described in [4], in 1 mm^2 for the Rule Memory and in 8 mm^2 for the Knowledge Base Memory.

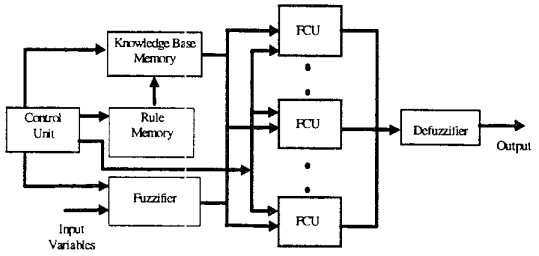

Fig. 8: Scalable architecture

The processing speed depends on the statistical characteristics of the source to be policed. We consider a bursty source and indicate the cell interarrival time as t_c. This is the maximum time limit by which the fuzzy policer has to make a decision, that is, infer the output ΔN_{i+1} by processing the 18 rules in the Knowledge Base. This constraint derives from analysis of the typical timing of the policing strategy. The fuzzy policer can, in fact, only begin processing rules when the input variables, namely A_{oi}, A_i and N_i, are all available; this does not happen until the end of the current window. If we indicate the latency introduced by the policer with t_L, the latter cannot exceed the value of t_c, thus preventing a cell arriving at the beginning of the new window from escaping control action.

If we want to express the time constraint in terms of FLIPS (Fuzzy Logic Inferences Per Second), i.e. inference execution speed, the fuzzy policer has to give a minimum performance of $1/t_c$ FLIPS. In the case of a packetized voice source as considered in the previous section, if t_c=12 ms the performance required is 84 FLIPS. The architecture we are taking as our reference has a minimum performance of 77 KFLIPS if a single FCU is operating; it therefore amply meets the speed required to police a single source. On the basis of the results given in [4], and considering the application memory requirements and a single FCU, the cost of implementation (expressed in terms of silicon area occupied) is estimated to be only 15 mm^2. With large production volumes this translates into a monetary cost of a few dollars per chip. It should also be pointed out that, in view of the performance level offered, the same chip could be used to police sources with more severe dynamic characteristics. In the case of still picture sources with a peak bit rate of 2 Mbit/s, the processing speed would be much lower than the 77 KFLIPS of the fuzzy processor.

An improvement in the exploitation of hardware

resources can be obtained if the fuzzy processor is used to police a set of sources in time sharing. For example, in the case of voice sources it is capable of policing almost a thousand sources. This would lead to further decreasing the cost per source.

6. Conclusions

In this paper we have introduced a policing mechanism for ATM networks which is based on fuzzy logic. It is a window control mechanism in which the number of cells that can be accepted per window is dynamically updated in accordance with the degree of compliance of the source with the negotiated parameter, namely the average cell rate. The mechanism detects the arriving cells as excessive or not according to a soft decision-making logic using truth values which are not restricted to either false (truth value "zero") or true (truth value "one"), but range in a continuum of [0,1].

The parameters describing both the source behaviour and the policing control actions are expressed by linguistic variables and fuzzy sets. The control strategy is described through a set of fuzzy inferences which emulate the knowledge base that is typical of human expertise.

Thanks to the softness of fuzzy logic, the proposed mechanism exhibits good dynamic behaviour, high flexibility and a fair capacity to meet the statistical properties of a wide range of traffic sources.

The performance of our mechanism has been evaluated through several simulations and compared with that of some of the most popular policing methods, such as the Leaky Bucket and EWMA. The results obtained have shown that the performance of our fuzzy policer is much better than that of conventional policing mechanisms, in terms of both responsiveness and selectivity, and approaches that of an ideal policer.

The fuzzy policer can easily be implemented in hardware, thus enhancing both cost and processing performance.

Some final remarks should be made about the prospects of the approach proposed. Although a problem still to be solved concerns the choice of parameters which describe the statistical properties of the various kinds of traffic and lend themselves to being policed, our mechanism is much more efficient than others when the parameter to be policed is the long-term average bit rate. This supports the hypothesis that an approach based on soft computing techniques may prove to be more appropriate for the policing of other parameters which characterize the source.

References

[1] R. O. Onvural: "Asynchronous Transfer Mode networks: performance issues", *Artech House* 1994.

[2] E. Rathgeb: "Modeling and performance comparison of policing mechanisms for ATM networks", *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 3, pp. 325-334, April 1991.

[3] V. Catania, A. Puliafito, M. Russo, and L. Vita: "A VLSI fuzzy inference processor based on a discrete analog approach", *IEEE Transactions on Fuzzy Systems*, Vol. 2, No. 2, May 1994.

[4] G. Ascia, and V. Catania: "A VLSI parallel architecture for fuzzy expert systems", *International Journal of Pattern Recognition and Artificial Intelligence*, Vol. 9, No. 2, 1995.

[5] J. S. Turner: "New directions in communications (or which way to the information age?)", *IEEE Communications Magazine*, Vol. 24, No. 10, pp. 8-15, Oct. 1986.

[6] G. Gallassi, G. Rigolio, and L. Fratta: "ATM: bandwidth assignment and bandwidth enforcement policies", *Proc. GLOBECOM '89*, Dallas, Nov. 1989.

[7] J. Monteiro, M. Gerla, and L. Fratta: "Leaky Bucket input rate control in ATM networks", *Proc. ICCC 90*, New Delhi, Oct. 1990.

[8] M. Buttò, E. Cavallero, and A. Tonietti: "Effectiveness of the Leaky Bucket policing mechanism in ATM networks", *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 3, April 1991.

[9] L. Dittmann, S.B. Jacobsen, and K. Moth: "Flow enforcement algorithms for ATM networks", *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 3, pp. 343-350, April 1991.

[10] A. Tarraf, I. Ibrahim, and T. Saadawi: "A novel neural network traffic enforcement mechanism for ATM networks", *IEEE Journal on Selected Areas in Communications*, Vol. 12, No. 6, pp. 1088-1096, August 1994.

[11] L.A. Zadeh: "Fuzzy sets", *Inf. Contr.*, Vol. 8, pp. 338-353,1965

[12] K. Nakamura, N. Sakashita, N. Nitta, K. Shimomura, and T.Tokuda: "Fuzzy inference and fuzzy inference processor", *IEEE Micro*, Vol. 13, No. 5, pp. 37-48, Oct. 1993.

[13] H.Watanabe, J.R.Symon, W.D.Detloff, and K.E.Yount: "VLSI fuzzy chip and inference accelerator board system", *Fuzzy Logic for Management of Uncertainty*, pp. 211- 243, John Wiley & Sons Inc., 1992.