2

Specification of Cell Dispersion in ATM Networks

Fabrice Guillemin*, Catherine Rosenberg†, and Aref Meddeb†

* France Telecom, CNET Lannion A, Route de Trégastel, 22300 Lannion, France.
Tel: + (33) 96-05-13-46, Fax: + (33) 96-05-11-98, Email: guillemi@lannion.cnet.fr.
† Ecole Polytechnique de Montreal, C.P. 6079, succ. A, Montreal, Quebec, Canada H3C 3A7.
Tel: + (514)340-4123, Fax: + (514)340-4562, Email: {cath,aref}@comm.polymtl.ca.

Abstract

So far, cell jitter has been investigated within the standardization bodies for Constant Bit Rate connections only and it has been recognized that cell jitter covers two different topics, namely cell clumping and cell dispersion. Whereas the material for the specification of cell clumping seems stable for the time being, the specification of cell dispersion requires further investigations in particular, because we have shown that the algorithm currently under consideration within the ITU-T for specifying cell dispersion fails to provide useful information for the operation of very simple receiving ATM Adaptation Layer (AAL) mechanisms. Moreover, it exhibits an unstable behavior when cell transfer delays become very large and/or cell loss occurs, and the outcome of the algorithm depends on the specific realization of the observed cell stream. Hence, after a detailed description of the issues related to cell dispersion, we propose an algorithm which computes two quantities of interest for specifying cell dispersion.

Keywords

ATM, Cell Delay Variation, ATM Adaptation Layer, cell dispersion

1 INTRODUCTION

A basic feature of telecommunication networks based upon the Asynchronous Transfer Mode (ATM) is that cells progressing along a connection experience random delays, which may be caused for instance by queuing in multiplexing stages. As a consequence, the initial time structure of any cell stream traversing an ATM network is in general altered by stochastic perturbations.

For example, when observing at some point along a connection an initially periodic cell stream, which has passed through several network elements using asynchronous multiplexing, the cell arrival process is no longer periodic but erratic because of random cell transfer delays. This is the phenomenon of cell jitter, also referred to as Cell Delay Variation (CDV) within the standardization bodies (I.356, 1993) and (I.371, 1995). Roughly speaking, cell jitter on a connection is due to the interaction between this connection and other connections when sharing common resources (transmission capacities, buffers, etc.). So far, cell jitter has been investigated within the ITU-T only for cell streams, which

L. Mason et al. (eds.), *Broadband Communications*

are initially periodic (Constant Bit Rate, CBR, connections). In such a case, it has been recognized that cell jitter is composed of two basic complementary phenomenons, namely cell clumping and cell dispersion (Guillemin and Monin, 1992).

To illustrate the cell clumping effect, consider an initially periodic cell stream with period T passing through an ATM network. When this cell stream is observed at some point along the connection (e.g., the T_B reference point, an inter-Network Node Interface, or the receiving S_B interface), it may happen that the distance between two consecutive cells is less than the initial period T because of random delays. This is precisely what is called cell clumping.

Cell clumping is quantified by using a reference algorithm, namely the Virtual Scheduling Algorithm (VSA). A definition of the VSA is given in the ITU-T Recommendations I.356 (1993) and I.371 (1995) Annex 1. The VSA is also known as Generic Cell Rate Algorithm (GCRA) within the ATM Forum (The ATM Forum, 1994).

Consider a periodic cell stream with period T traversing an ATM network and observed at some point along the connection. It may happen that two consecutive cells experience increasing delays so that their inter-arrival time at the observation point is greater than the initial period T, giving rise to a cell gap in the cell stream. This is the phenomenon of cell dispersion, which is critical for connections supporting a circuit emulation (e.g., ATM Adaptation Layer, AAL, of type 1) and/or when re-sequencing (e.g., two-layer video) has to be performed in the receiving terminal.

In Section 2, we present a detailed formulation of the problems related to cell dispersion in ATM networks. A first specification of cell dispersion has been introduced in Gravey and Boyer (1993) and I.356 (1993). However, the proposed algorithm fails to provide useful information for the operation of very simple receiving ATM Adaptation Layer (AAL) mechanisms (see Guillemin, Rosenberg and Meddeb, 1996). We introduce in Section 3 a new method of characterizing cell dispersion. Similarly to cell clumping, which is specified with respect to a reference algorithm (the VSA), we propose a new algorithm to measure cell dispersion.

2 PROBLEM FORMULATION

2.1 Need for Quantifying Cell Dispersion

Consider a CBR connection with peak emission interval T supporting a circuit emulation and assume that this connection traverses one or several ATM networks. As mentioned above, the cell stream arriving at the receiving terminal may be affected by cell dispersion. This phenomenon is especially critical for AAL 1 mechanism. Specifically, such a mechanism must restore as far as possible the periodic structure of the cell stream. For this purpose, arriving cells are stored in a buffer, usually referred to as *elastic buffer* or *playout buffer*.

Several disciplines may be envisaged to remove cells from an elastic buffer. For instance, the service rate may be such that the elastic buffer may be at any instance half-full. In such a case, the service rate adapts to the instantaneous cell arrival rate, which is time-varying because of cell jitter. Another possibility is to store the first arriving cell in the elastic buffer for a certain time (*fixed playback point*) and then, to read out this buffer with rate $1/T$. This is the most simple AAL mechanism, which was used for instance in the early ATM experiment named PRELUDE (Devault, Cochennec and Servel, 1992). This mechanism will be referred to as *reference AAL mechanism* in the following. In some sense, this mechanism is the worst one since it does not adapt to the instantaneous cell arrival rate and relies mainly on one parameter, namely the buffering delay of the first arriving cell.

With regard to the quality of service (QoS) of the circuit emulation, the key problem is to avoid starvation in the elastic buffer (since overflow can be taken care of by means of large buffers). Indeed, because of cell dispersion, it may happen that the elastic buffer empties (e.g., if cells experience over-large delays), leading to an interruption in the cell disassembling process and then to a degradation of the QoS offered to the application (e.g., clipping in a telephone connection). More precisely, after the initial buffering delay, the elastic buffer is explored either periodically with period T in a reference AAL mechanism or quasi-periodically when the service discipline is more sophisticated. If a cell is present in the buffer at a server exploration time, it is removed. Starvation occurs when there is no cell to remove from the elastic buffer at a server exploration time and hence cell stuffing is required to mask the interruption in the cell disassembling process. The stuffing can be more or less sophisticated depending on the application characteristics and on the acceptable level of complexity at the playout.

In any way, in order to dimension the AAL mechanism to offer the right level of QoS, information is required on the cell dispersion affecting the cell stream (i.e., on the characteristics of the cell stream entering the playout). Means are needed to quantify and measure cell dispersion.

2.2 Differences between Cell Clumping and Cell Dispersion

Cell clumping has a direct impact on the network, for instance on UPC/NPC mechanisms and resource allocation when "pick-up" policing mechanisms are used in the UPC/NPC function (Boyer, Guillemin, Servel and Coudreuse, 1992). According to ITU-T Recommendation I.371, cell clumping quantification is required at each T_B and inter-network NNI interfaces (in particular in order to dimension UPC/NPC mechanisms) and is specified in the traffic contract by means of a CDV tolerance value, usually denoted by τ.

Because of the possible impact of cell clumping on resource allocation, there is a direct incentive for a network operator to limit and even eliminate as far as possible cell clumping at the network access point. This can be achieved by performing cell spacing (Boyer, Guillemin, Servel and Coudreuse, 1992).

In the contrary, cell dispersion on a connection is an end-to-end phenomenon which does not have a direct impact on resource allocation. So, there is no direct incentive for a network to limit or eliminate cell dispersion. Cell dispersion could be eliminated (or at least reduced) by buffering the cells of the connection for a certain time somewhere along the connection to absorb as far as possible the random fluctuations in cell transfer delays. Buffering cells for a certain time would greatly increase the QoS parameter Cell Transfer Delay (CTD, I.371, 1995). This should carefully be examined from an application point of view. If one decides in favor of eliminating part of the cell dispersion effect on a connection (e.g., at the output of each network), this should be done in preference within the ATM layer, since it would be highly desirable not to go beyond the ATM layer in the middle of a connection.

As a consequence, the key difference between cell clumping and cell dispersion is that it may be inappropriate to eliminate cell dispersion somewhere along a connection supporting a real time application while the same operation may be very useful for cell clumping.

2.3 How and Where to Quantify and Measure Cell Dispersion ?

As this point, the necessity to quantify and measure cell dispersion appears clearly. Many questions remain open especially concerning how and where to quantify and measure it. In ways similar to the one followed for cell clumping, there is the need to define an *algorithm* to measure cell dispersion.

An ideal situation would be that no interruptions occur in the cell disassembling process in the receiving terminal. However, since cell transfer delays may be very large, this would require huge buffering delays (this may cause echo phenomenons for telephone connections). A more realistic situation is that starvation occurs with a very low probability and that the AAL mechanism recovers from starvation by producing stuffing cells.

This leads to a statistical characterization of cell dispersion and the necessity to define precisely the stuffing function.

Since cell dispersion is definitely a problem related to the AAL level, we should try to quantify it using some information relevant to AAL mechanisms. However, the existence of many possible AAL mechanisms implies that it will be difficult to find some information common to all AAL mechanisms. Furthermore, even if such information could be found, a more fundamental problem, besides the arbitrary choice of not going beyond the ATM Layer in the middle of a connection would forbid the use of information related to the AAL level to quantify cell dispersion. Namely, the algorithms for measuring cell clumping and cell dispersion should be coordinated in the sense that the cell dispersion algorithm should only consider the conforming cells with respect to cell clumping. Indeed, the UPC/NPC mechanism is entitled to discard any cell, which does not conform with the VSA, whose parameter are the declared peak emission interval and cell clumping tolerance. Thus, the UPC/NPC mechanism by discarding non conforming cells may significantly but legally increase the magnitude of cell dispersion. As a matter of fact, if the algorithm for cell dispersion were to take into account all cells, it would measure a cell dispersion effect smaller than the one it would measure by considering conforming cells only. By discarding non conforming cells, the network should not be held responsible for the increase in cell dispersion.

To illustrate this phenomenon, consider an initially periodic cell stream with period T affected by CDV as depicted in Figure 1. Cell clumping at the observation point is quantified by a CDV tolerance parameter, which allows, say, two back to back cells. But, the cell stream is altered by unexpected excessive cell clumping so that cell clumps are composed of four back to back cells and thus, half of the cells are non conforming. Because of the non conformance of some cells, the stream of conforming cells is affected by cell dispersion, exhibited by the presence of cell gaps. Now, with regard to cell dispersion, if all cells were considered as to be conforming, then it is easily checked that it would be possible, from the jittered cell stream, to restore a perfectly periodic cell stream, for example by smoothing cell clumps in a buffer. In such a case, excessive cell clumping masks cell dispersion. Since a UPC/NPC mechanism may discard non-conforming cells, it may severely increase the magnitude of cell dispersion but this is legal and the network should not be considered as responsible for this increase in cell dispersion. As a consequence, the cell dispersion algorithm must handle only cells that are declared as to be conforming by the cell clumping algorithm and thus, both algorithms must be run simultaneously in a coordinated mode.

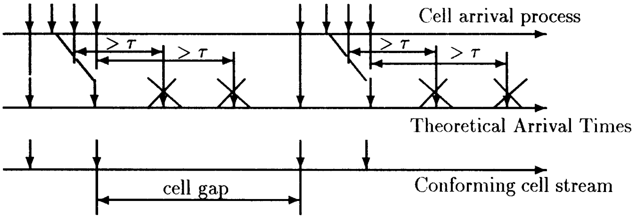

Figure 1 Impact of excessive clumping on cell dispersion.

It follows that cell dispersion should be characterized within the ATM Layer and the algorithms for measuring cell clumping and cell dispersion should be coordinated in the sense defined above.

2.4 Conformance with respect to Cell Dispersion

The problem of cell conformance testing with respect to cell dispersion has not yet been addressed within the standardization bodies but it is believed in this paper that this problem is fundamental, when a connection supports a circuit emulation. Specifically, since the (end-to-end) cell dispersion effect results for a multi-operator connection from the passage through several networks, it will be necessary to create some conformance testing procedure to check that the cell dispersion incurred by the cell stream in each network does not exceed some predefined limits.

As a matter of fact, for a proper operation of the receiving AAL mechanism, the cell dispersion should not exceed some predefined bounds. In the case of a multi-operator connection, a tool, namely a conformance algorithm, should be specified to assess the responsibility in the case of a conflict between a user and the network (at large) with regard to the QoS of a circuit emulation. Roughly speaking, if cells experience over-large transfer delays in a network along the connection, then the QoS of the circuit emulation may be degraded, for instance because of interruptions (or stuffing) due to information starvation in the cell disassembling process in the receiving terminal. In that case, it should be possible, from a *legality* point of view, to determine who is responsible for the QoS degradation, because the user has the right to complain that the different networks involved along the connection have not met their QoS commitments.

Hence, when cells of the connection are taken over from one network to another, the receiving network should be able to check that the delivered cell stream is not too much dispersed. If cell dispersion is too large, then the receiving network should not be considered as responsible for QoS degradation. In addition, each network should behave in such a way that it delivers to the next network a cell stream, which is not too much dispersed.

So even if cell dispersion is an end-to-end problem, the legality question comes down to check that each network behaves correctly with respect to cell dispersion. In fact, cell dispersion conformance testing should be performed at each network access point. Since we have chosen to remain in line with the ATM philosophy regarding the non passage to the AAL level in the middle of a connection, we will have to implement the cell dispersion conformance testing mechanism within the ATM Layer, which precludes the use of AAL information for quantifying the cell dispersion effect.

In summary, the legality problem has two major consequences: the necessity to quantify the cell dispersion effect with information within the ATM Layer and the need for an algorithm able to test conformance of cell dispersion, similar to the VSA used to test cell conformance with respect to cell clumping (I.356, 1993)

2.5 Apportionment of Cell Dispersion

With regard to cell clumping, the user at the T_B interface or the backward network at an inter-network NNI declares a value for the CDV (clumping) tolerance parameter τ that the network can either accept or refuse. An alternative approach would consist in specifying a *reference connection* and in apportioning cell clumping to each component of the connection (the Customer Equipment, each network, etc.). Two reasons why this alternative may be of little interest are that

1. the UPC should be perfectly programmable and thus should allow more flexibility than a reference connection ; in particular, any value for the cell clumping tolerance parameter τ could be declared and may possibly depend on the declared peak cell rate,
2. a network operator can perfectly control cell clumping on a connection via cell spacing (e.g., it can reduce it to that equivalent to one multiplexing stage), and thus, the need for a reference connection to specify cell clumping at each interface is useless ; this last point needs nevertheless further investigations within the standardization bodies since some network operators are still reluctant to perform traffic shaping and in particular cell spacing.

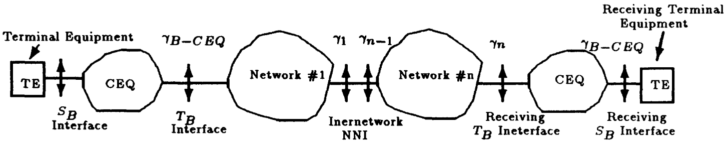

Figure 2 Apportionment of cell dispersion along an ATM connection.

In the contrary, it is of little interest to reduce cell dispersion on a connection. Thus, cell dispersion is mainly a problem to be dealt with at the end-user terminal, hardwired with predetermined AAL characteristics, which cannot be modified on a per connection basis. Since cell dispersion increases as cells progress along a connection and since cell dispersion has an impact on the QoS offered to a circuit emulation, the magnitude of cell dispersion should be kept under appropriate predefined bounds at the end-equipment terminal and thus within each network involved along the connection. A simple method to determine these bounds is to use a reference connection as depicted in Figure 2.

The end-to-end cell dispersion is divided into several components, namely the cell dispersion for the sender Customer Equipment, the cell dispersion for network $\#i$ along the connection, and the cell dispersion for the receiving Customer Equipment, so that the end-to-end dispersion is equal to the sum of all these components.

It is much too early to give an exact definition of a reference connection, given the lack of information on the behavior of switches and AAL mechanisms, in particular on how much cell dispersion is acceptable for guaranteeing a given QoS level. But, under the assumptions that cells are not buffered in the middle of the connection to absorb cell dispersion, we can already assess that the apportionment cannot be linear since cell dispersion is only going to increase along the connection.

2.6 Cell Dispersion Characterization

Besides the points raised above, some questions are still open and will be addressed in Section 3. In particular, the exact definition of the cell dispersion parameters or in other words, the design of a proper algorithm for measuring and testing conformance of the cell dispersion remains an open issue, since we show in Guillemin, Rosenberg and Meddeb (1996) that the algorithm currently under consideration within the ITU-T (I.356, 1993) suffers from many drawbacks.

Before proceeding further, let us summarize the basic points for cell dispersion characterization :

- cell dispersion should not necessarily be absorbed along the connection,
- legality issues imply conformance testing with respect to cell dispersion, which leads to the specification of a conformance algorithm,
- cell dispersion conformance testing requires the quantification of cell dispersion with information within the ATM layer,
- quantification of cell dispersion should be statistical,
- the algorithms for measuring and testing conformance of cell clumping and cell dispersion should be coordinated in the sense that the cell dispersion algorithm should only consider the conforming cells with respect to cell clumping,
- specification of a reference connection and apportionment of cell dispersion are needed.

3 A NEW ALGORITHM FOR THE SPECIFICATION OF CELL DISPERSION

Before addressing the design of an algorithm for the specification of cell dispersion, we should try to better understand which characteristics of cell dispersion are of importance from an application standpoint. As already mentioned, QoS degradation occurs when there is a period of starvation in the playout buffer. Of course, each occurrence of a starvation period may entail a QoS degradation, but the level of the degradation is directly linked to the length of the period.

As a matter of fact, a short starvation period can be masked via cell stuffing, while a long one may well have a disastrous effect on the QoS. In other words, we believe that for a given AAL mechanism, two characteristics of cell dispersion are important from an application standpoint (or that these two characteristics are sufficient to estimate the performances of an AAL mechanism), namely the frequency of starvation periods and the maximum size (or a remote quantile) of a starvation period. If the length of a starvation period is greater than the maximum size, the QoS degradation becomes unacceptable and some drastic actions should be performed by the AAL mechanism to guarantee the QoS of the application.

Roughly speaking, starvation periods of limited size are "natural" in the context of ATM systems and are due to random cell transfer delays through the network. Such starvation periods should be masked by the AAL mechanism (e.g., via cell stuffing) and as long as they do not occur frequently, they have a limited impact on the QoS of the application. This is why the frequency of starvation periods is an important characteristic of cell dispersion. In the contrary, starvation periods of large size may result from error conditions or unexpected traffic configurations within the network, which entail over-large cell transfer delay. Such a starvation period may adversely affect the QoS of the application. For example, the AAL mechanism may be unable to compensate the starvation of cells. For a proper operation of AAL mechanisms, the boundary between starvation periods of small and large size should be specified. This limit is precisely what we have called above maximum size of a starvation period.

If the values of the maximum size of a starvation period and the frequency of starvation periods were known in advance for a particular AAL mechanism, they may be used for dimensioning this AAL mechanism in the receiving terminal equipment. Moreover, the values of these two parameters may depend on the application supported, since they are directly related to the QoS of the application. As a matter of fact, constant bit rate video may have different characteristics and requirements than telephony.

Since a lot of AAL mechanisms are currently studied and implemented, each of them with more or less sophisticated features, we consider in the following the simplest AAL mechanism, namely the reference AAL mechanism introduced in Section 2, which is certainly the most poorly featured mechanism. It is actually expected that the information derived by considering this mechanism is sufficient to make the use of more sophisticated AAL mechanisms possible. Specifically, we are interested in the number of starvation periods N_g as well as in their length in the reference AAL mechanism. Let γ be the maximum allowed size of a starvation period (if the starvation period is greater than γ, the QoS degradation becomes unacceptable).

One question with respect to the algorithm is to decide whether the cell dispersion parameters N_g and γ should be "absolute" (with no closing error) or statistical (i.e., related to some quantile). In the latter case, the magnitude of cell dispersion is specified with some closing error ε. Objective values for ε are typically related to QoS issues. While it is clear that there is no way an application could request to encounter absolutely no starvation period (and thus N_g should be less or equal than a given limit or the ratio $\frac{N_g}{N_c} \leq \varepsilon_g$ where N_c is the number of conforming cells), the size of the starvation period could either be absolutely bounded or statistically bounded by γ. In any case, if a network testing conformance on the dispersion affecting a connection, finds the bound on N_g or the bound on γ exceeded, the network is not responsible for any QoS degradation.

If, the size of the starvation periods is absolutely bounded by γ, i.e., the network cannot create any starvation period of size greater or equal to γ, an algorithm such as the one presently proposed in I. 356 (1993) could check if γ/T (T is the period of the CBR source) is exceeded even if it was not its primary objective and if it has many drawbacks that should preclude its use (Guillemin, Rosenberg and Meddeb, 1996).

We are in favor of statistically bounding the size of the starvation periods, i.e., the proportion of starvation periods of size greater or equal to γ should be less or equal to ε_γ.

Another related issue concerns the ratio of the closing error for cell clumping to that for cell dispersion. For instance, cell clumping could be specified with a 10^{-9} closing error so that a fraction of at most 10^{-9} cells are detected as non-conforming by the UPC/NPC mechanism and cell dispersion could be specified with $\varepsilon_g = 10^{-4}$ and $\varepsilon_\gamma = 10^{-9}$ depending on the admissible QoS, which can be achieved by the receiving AAL mechanism. In this line of investigations on the relationship between peak cell rate control and cell dispersion, another problem, which could be addressed once the algorithm for specifying dispersion has been chosen, is the one concerning cell spacing. Specifically, the general belief that cell spacing increases cell dispersion should be questioned.

So, we expect that, once the reference connection has been defined, that an AAL would have, for this connection, the assurance that the proportion of starvation periods is less than a given ε_g and that for the γ of its choice, the proportion of starvation periods of size greater or equal to γ is less than a given ε_γ.

We propose in this section an algorithm (see Figure 3) for the specification of cell dispersion. This algorithm computes the ratio $\frac{N_g}{N_c}$ where N_c is the number of conforming cells and N_g is the number of starvation periods, and the ratio $\frac{N_\gamma}{N_g}$ where N_γ is the number of starvation periods of size greater than γ.

There are three counters :

- N_c, the number of conforming cells;
- N_g, the number of starvation periods;
- N_γ, the number of starvation periods of size greater than γ.

We have also introduced a variable TST (for Theoretical Service Time) which is the time before which a cell should arrive for the process not to exhibit a starvation period. A cell is conforming and thus taken into account by our algorithm if it is conforming to the VSA(T, τ). For each conforming cell, we compute its TST and compare it with the cell arriving time. If the cell has arrived before its TST (i.e., $TST < t$), there is no starvation period and the algorithm waits for the next conforming cell to arrive. Otherwise there is a starvation period and the algorithm increments the counter N_g, check if the starvation period size is greater than γ (if yes it increments the counter N_γ) and computes in any case the new value of TST. To compute the size of the starvation period and new value of TST, we have introduced the function $\mathrm{TST}(t)$ where $\mathrm{TST}(t) = \mathrm{TST} + \lceil (t - \mathrm{TST})/T \rceil \times T$.

$\mathrm{TST}(t)$ is the theoretical service time for a conforming cell which follows a starvation period. The size of the starvation period is computed as the difference between the previously computed TST and $\mathrm{TST}(t)$.

At the end of the connection, the proportion $\nu_g = N_g/N_c$ and $\nu_\gamma = N_\gamma/N_g$ can be computed and compared respectively to ε_g and ε_γ.

Note that the big difference between this algorithm and the VSA which is used for characterizing clumping is that this one is inherently discrete in nature.

The value of γ is application dependent and thus the network operators would give for the reference connection introduced earlier the apportioned tolerance for different values of γ.

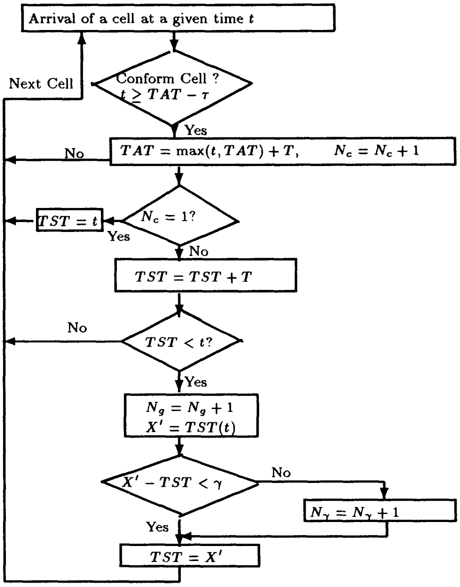

Figure 3 The new algorithm for specifying cell dispersion.

At a given interface, namely the UNI or a NNI, a real time connection is characterized, as far as cell dispersion is concerned, by γ and objective values for ν_g and ν_γ, say, $\varepsilon_g = 10^{-4}$ and $\varepsilon_\gamma = 10^{-9}$, respectively. The connection is said to be conforming if the ratios ν_g and ν_γ computed at the end of the connection for the specified value of γ satisfy $\nu_g \leq \varepsilon_g$ and $\nu_\gamma \leq \varepsilon_\gamma$. In the case of a multiple operator connection, a reference connection should be considered (see Section 2). Each portion of this connection should be assigned an objective for cell dispersion. At each interface, the network operator receiving traffic may check by using the new algorithm whether the connection is conforming with respect to cell dispersion. If the user complains of QoS degradation, the origin of this degradation may be identified.

Note that the algorithm currently considered in ITU-T Recommendation I.356 may be used (see also Guillemin, Rosenberg and Meddeb, 1996), to some extent, as a special case of our algorithm depicted in Figure 3 if $\varepsilon_g = 1$ (i.e, the frequency of cell starvation periods is not considered as a QoS parameter) and $\varepsilon_\gamma = 0$ (i.e., the connection is not conforming as soon as γ is exceeded).

As mentioned in Section 2 and as it can be seen in Figure 3, the cell dispersion algorithm is coordinated with the cell clumping algorithm (only conforming cells are taken into account). Intuitively, cell loss increases the magnitude of cell dispersion. Cell loss may be caused by legitimately discarding cell in the UPC/NPC. In fact, the tolerance τ itself characterizes the magnitude of cell clumping with some closing error ε_c. As a consequence, it may be very unsafe to characterize cell dispersion and cell clumping at an interface with the same accuracy (ε_c and ε_g with the same order of magnitude). It seems more suitable to specify cell dispersion less accurately than cell clumping (say, $\varepsilon_c = 10^{-9}$ and $\varepsilon_g = 10^{-4}$).

4 CONCLUSION

Some issues related to cell dispersion in ATM networks have been discussed in this paper. It turns out that information on cell dispersion is needed for dimensioning AAL mechanisms in receiving terminals. Absorbing cell dispersion in the middle of a connection (e.g., via a circuit emulation) is not suitable and even impossible. Legality problems entail that conformance testing with respect to cell dispersion is necessary. For conformance testing purposes, cell dispersion should be characterized by information within the ATM layer. This characterization should rely on statistical information and should be coupled with that of cell clumping. Namely, only cells, which are conforming with respect to cell clumping, should be considered. In the case of multiple operator connections, cell dispersion should be apportioned by considering a reference connection.

Recognizing that the algorithm currently specified in ITU-T Recommendation I.356 suffers from instability and gives a poor information on cell dispersion, a new algorithm has been proposed in this paper for specifying cell dispersion. Roughly speaking, this algorithm emulates a reference AAL mechanism and counts the occurences of starvation periods and the number of starvation periods whose length is greater than a specified value γ. This two parameters are expected to provide sufficient information for the operation of AAL mechanisms more sophisticated.

This algorithm allows the characterization of cell dispersion. Further investigations are needed to clarify with which accuracy cell clumping and cell dispersion should be specified (respective orders of magnitude of the objectives ε_g and ε_c). Moreover, the general belief that cell spacing adversely increases the magnitude of cell dispersion should be questioned. Namely, the impact of cell spacing on the parameter γ should be clarified. These point will be addressed in forthcoming papers.

REFERENCES

Boyer, P., Guillemin, F., Servel, M., and Coudreuse, J.P. (1992) Spacing cells protects and enhances utilization of ATM network links: *IEEE Communications Magazine*, pages 38–49.

Devault, M., Cochennec, J.Y. and Servel, M. (1988) The PRELUDE experiment : assessments and future prospects: *IEEE Journal on Selected Areas in Communications, Vol 6, n°9.*

Gravey, A. and Boyer, P. (1993) Cell Delay Variation in ATM networks: *Modeling and Performance Evaluation of ATM Technology*, Martinique.

Guillemin, F. and Monin, M. (1992) Management of Cell Delay Variation in ATM networks: *Globecom'92*, Orlando, FL.

Guillemin, F., Rosenberg, C. and Meddeb, A. (1996) Analysis of cell dispersion in ATM networks. *In preparation.*

I.356. ITU-T Recommendation (1993) B-ISDN ATM layer cell transfer performance, Geneva.

I.371. ITU-T Recommendation (1995) Traffic control and congestion control in B-ISDN, Geneva.

Roberts, J. and Guillemin, F. (1992) Jitter in ATM networks and its impact on peak cell rate control: *Performance Evaluation, Special Issue on Modeling of High Speed Telecommunications Systems.*

The ATM Forum (1994). User-Network Interface Specification, Version 3.1.