A Protocol for Supporting the ABT/DT Capability

Pierre Crégut, Fabrice Guillemin,
France Telecom, CNET Lannion-A
2, Avenue Pierre Marzin
22 300 Lannion, France
Pierre.Cregut,Fabrice.Guillemin@lannion.cnet.fr

Barbara Heyd
CRIN/CNRS, INRIA Lorraine
BP 239
54506 Vandoeuvre lès Nancy, France
Barbara.Heyd@loria.fr

Abstract

With the ATM Block Transfer with Delayed Transmission (ABT/DT) capability, the amount of bandwidth allocated to the different cell flows of a connection (namely the CLP=0+1 and possibly the OAM cell flows) is constant over an ATM block and then referred to as block cell rate (BCR). Successive BCR values are dynamically negotiated between the user and the network via the exchange of Resource Management cells. The BCR negotiation for a given ATM block may be initiated either by the source or the destination. As a basic principle, before transmitting an ATM block, the source should receive an explicit authorization from the network. This paper gives the existence proof of a protocol for supporting the ABT/DT capability as specified in the May 1996 version of ITU-T Recommendation I.371.

1. Introduction

While the standardization process of ATM and first ATM experiments (e.g., the European ATM Pilot [10]), were dominated up to a few years ago by classic telecommunications concepts (e.g., connections at the peak cell rate), the steadily increasing need for multimedia applications and integrated services generating variable bit rate traffic has urged the standardization bodies to turn to the problem of statistical multiplexing. The ATM Forum, which successively introduced the concepts of sustainable cell rate [1] and elastic applications supported by the available bit rate service [2], has played a predominant role in the development of statistical multiplexing schemes for ATM networks.

Taking benefit of the studies carried out within the ATM Forum, the ITU-T completed in May 1996 a new release of ITU-T Recommendation I.371 on congestion and traffic control in B-ISDN [8]. This new release specifically provides the high level description of three statistical ATM transfer capabilities, namely the statistical bit rate, the available bit rate, and the ATM block transfer (ABT) capabilities. Connections assigned a static amount of network resources computed on the basis of the peak cell rate are supported by the so-called deterministic bit rate capability.

ABT [7, 8] is in fact the capability derived from the fast reservation protocols (FRPs) proposed by Boyer in [3] and equivalently by Doshi in [6]. Since their introduction, these protocols have been progressively refined, notably to account for the elasticity property of some applications. The exchange of Resource Management (RM) cells used in FRP to reserve/release bandwidth within the network, has notably been formalized to give rise to a new ATM layer entity, referred to as ATM block. An ATM block is specifically a group of consecutive cells of an ATM connection delineated by two RM cells.

With ABT, the peak cell rate of a cell flow of a connection (namely the CLP=0+1 and possibly the OAM cell flow) is constant over the duration of an ATM block and is as a consequence referred to as block cell rate (BCR). The option of dynamically managing the bandwidth of the end-to-end OAM cell flow is chosen to enable further developments of OAM functionalities as discussed in [7].

Like FRP protocols, two transmission modes exist for the transfer of ATM blocks, namely the delayed [4] and the immediate [5] transmission modes, giving rise to the ABT/DT and the ABT/IT transfer capabilities, respectively. The basic characteristic of ABT/DT, as opposed to ABT/IT, is that the source should receive an explicit authorization from the network before starting the transmission of an ATM block. This paper provides the existence proof of a protocol to support the ABT/DT capability with all the options (bidirectional BCR negotiation, traffic control and maintenance procedures, etc.) specified in the May 1996 version of ITU-T Recommendation I.371.

2. ABT/DT Principles

To simplify the discussion, only point-to-point bidirectional communications limited to a single network are considered in this paper (see Figure 1). In the following, vec-

0-8186-7453-9/96 $05.00

tor Λ denotes the couple of cell rate values $\Lambda(0+1)$ and $\Lambda(OAM)$ for the CLP=0+1 (including end-to-end OAM) and OAM cell flows, respectively.

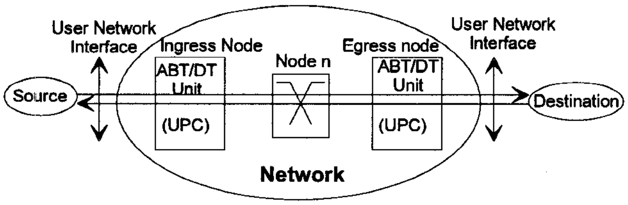

Figure 1. Reference configuration for ABT/DT.

2.1. Traffic Contract and Notation

At the establishment of an ABT/DT connection, a traffic contract is negotiated between the originating entity (say, the source without loss of generality) and the network. This traffic contract notably specifies: the peak cell rate values Λ_{PCR}, the associated CDV tolerances, and the maximum frequency of user BCR requests via the specification of the peak cell rate $1/T_{RM}$ together with the associated CDV tolerance τ_{RM} of the BCR request RM cells submitted by the user to the network.

The BCR request RM cell flow issued by the source has hence to conform to the generic cell rate algorithm (GCRA) [1] with parameters (T_{RM}, τ_{RM}) at the interface. Note that a given user BCR request RM cell flow includes the request RM cells sent by the user for BCR negotiation in the forward direction (the user is then the source) and the backward direction (the user is the destination).

In the following, the respective BCR values currently allocated to the CLP=0+1 and OAM cell flows are denoted by Λ_{Current} at network ingress and λ_{Current} at network egress. These two pairs of BCR values are in principle equal as long as no BCR negotiation is in progress for the connection. Λ_n denotes the corresponding BCR values allocated in network node #n. The BCR values at network ingress (resp. egress), which are reserved via ABT/DT but still not allocated, are denoted by $\Lambda'_{\text{Current}}$ (resp. $\lambda'_{\text{Current}}$).

2.2. RM Cell Content

The ABT/DT RM cell format is given in ITU-T Recommendation I.371 [8]. Specifically, an ABT/DT RM cell M, identified by a protocol identifier set equal to 2, is composed of the following elements:

- **M.Type**: this one bit field (known as Req/Ack bit) indicates whether the ABT/DT RM cell is a bandwidth request or acknowledgement message. When M.Type=0 (resp. 1), the RM cell is a bandwidth request (resp. acknowledgement) RM cell, denoted by BWReq (resp. BWAck) cell for short.
- **M.Dir**: this one bit field (called Direction, Dir, bit) indicates the direction the BCR negotiation applies to. The BCR negotiation applies to the forward direction if Dir=0 and to the backward direction if Dir=1.
- **M.Elastic**: this is the Elastic/Rigid bit, which is set equal to 0 by either user or the network to indicate that the source can adapt its transmission rate to the bandwidth available within the network and then that the network may optionally overwrite the requested BCR values and allocate less resources than requested (elastic mode). If Elastic=1, the BCR negotiation is said to be in the rigid mode.
- **M.CI**: in the rigid mode, this bit (termed Congestion Indication, CI) indicates whether the BCR negotiation has succeeded (CI=0) or failed (CI=1). In the elastic mode, this bit may be set by the network in the case of a BCR increase to indicate that the connection has not at least received the bandwidth it would have received in the case of fair share.
- **M.Maintenance**: the Maintenance bit is set to distinguish RM cells for BCR negotiation purposes by the users and possibly by the network, identified by Maintenance=0, and RM cells used for maintenance of the ABT/DT procedure and identified by Maintenance=1. In the latter case, the RM cell is intended to align the BCR values all along the connection or to retry a BCR negotiation.
- **M.Trf**: this one bit field (named Traffic Management, Trf, bit) distinguishes a usual ABT/DT RM cell used by either user (Trf=0) from an ABT/DT RM cell generated by the network for traffic management purposes (Trf=1).
- **M.Rate**: this 2x2 octets field is the vector composed of the BCR values of the CLP=0+1 and OAM cell flows, denoted by M.Rate(0+1) and M.Rate(OAM), respectively. These BCR values are coded by using the coding scheme given in ITU-T Recommendation I.371.
- **M.Size**: this 4 octet field (referred to as Block Size) may convey the size of an ATM block but it is not used in this paper.
- **M.Number**: this 4 octet field (called Sequence Number, SN) conveys the sequence number of the RM cell and is used to identify the BCR transaction.

A BWReq cell when sent by a given entity (either the source, the destination, or some traffic regulation functions located at network ingress or egress) should convey the respective requested BCR values for the CLP=0+1 and OAM cell flows. These values are denoted by Λ_{Req} for a BCR negotiation initiated by the source and by λ_{Req} for a BCR negotiation initiated by the destination. When the source (resp. the destination) is the originating entity of a BCR negotiation, a BWReq or BWAck cell received by the source or the destination should convey the respective allocated BCR values Λ_{Alloc} (resp. λ_{Alloc}) for the CLP=0+1 and OAM cell flows.

2.3. User BCR Negotiation Principles

ATM BLOCK DELINEATION. Any BCR negotiation (BCR increase or decrease) by either user is initiated by sending BWReq cell into the network and by waiting for a response RM cell from the network (BWAck cell). The only exception is when a BCR decrease is initiated by the source itself. In that case, the BCR decrease is immediately taken into account by the network and not acknowledged.

If positive, a response RM cell received from the network should be acknowledged by the source by sending a BWAck cell. Any BCR modification by the destination should also be acknowledged by the source. The transmission of a new ATM block by the source can start only after the source has sent the corresponding BWAck cell.

As a consequence, the RM cells delineating a given ATM block (see Figure 2) are specifically those BWReq cells indicating a BCR decrease initiated by the source or those BWAck cells sent by the source in response to a BCR increase initiated by the source itself or to a BCR modification initiated by the destination or the network.

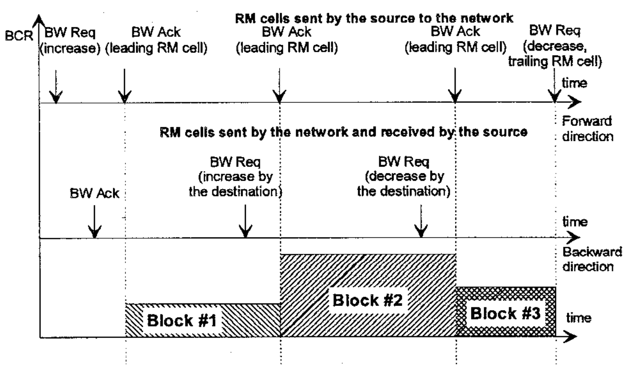

Figure 2. RM cells delineating an ATM block in ABT/DT.

PARAMETER CONTROL. The BCR values conveyed by a BWReq cell sent by the source should satisfy:

$$\begin{cases} \Lambda_{\text{Req}}(0+1) \leq \Lambda_{\text{PCR}}(0+1), \\ \Lambda_{\text{Req}}(\text{OAM}) \leq \Lambda_{\text{PCR}}(\text{OAM}). \end{cases} \quad (1)$$

Otherwise, for a BCR negotiation initiated by the source, the BCR values conveyed by the BWReq cell are forced at network ingress to:

$$\begin{cases} \Lambda_{\text{Req}}(0+1) = \min\{\Lambda_{\text{Req}}(0+1), \Lambda_{\text{PCR}}(0+1)\}, \\ \Lambda_{\text{Req}}(\text{OAM}) = \min\{\Lambda_{\text{Req}}(\text{OAM}), \Lambda_{\text{PCR}}(\text{OAM}), \\ \qquad\qquad \Lambda_{\text{Req}}(0+1)\}. \end{cases} \quad (2)$$

A similar procedure is run at network egress by replacing Λ_{Req} with λ_{Req}.

Similarly, a BWAck cell sent by the source should convey correct BCR values. To protect the network from misbehaved users, the BCR Λ_{Alloc} conveyed by a BWAck cell sent by the source are forced at network ingress to:

$$\begin{cases} \Lambda_{\text{Alloc}}(0+1) = \min\{\Lambda_{\text{Alloc}}(0+1), \Lambda'_{\text{Current}}(0+1)\}, \\ \Lambda_{\text{Alloc}}(\text{OAM}) = \min\{\Lambda_{\text{Alloc}}(\text{OAM}), \Lambda'_{\text{Current}}(\text{OAM}), \\ \qquad\qquad \Lambda_{\text{Alloc}}(0+1)\}. \end{cases} \quad (3)$$

BCR NEGOTIATION LIMITATIONS. Only one BCR negotiation by a given entity (the source or the destination) can be in progress within the network in a given direction. If the same entity initiates a BCR negotiation in a given direction while another one is still in progress within the network, then this new BCR negotiation RM cell is not processed by the network. If the corresponding RM cell is conforming to the traffic contract on the RM cell flow, the network sends a specific RM cell back to the originating entity (namely a Not_Ready message as described below).

COLLISION. Since user BCR negotiations for a same connection may be initiated by either the source or the destination, two BCR negotiations may collide within the network. Specifically, two BCR negotiations collide if the BWReq cell issued by an entity arrives at a network node, which has received a BWReq cell from the other entity and not the corresponding BWAck cell. To solve this problem, BCR negotiations initiated by the destination have priority over those initiated by the source (backward priority principle).

ABT/DT USER MAINTENANCE. In addition to BCR negotiation messages, either user may send a BWReq cell to align the BCR values allocated all along the connection or to know the BCR values currently allocated within the network. This is a Status_Enquiry procedure and the corresponding messages are called maintenance messages, identified by the Maintenance bit set equal to 1. The backward priority principle applies to maintenance messages.

2.4. Traffic Management Principles

TRAFFIC CONTROL. For traffic regulation purposes (e.g., traffic balance functions, block non conformance [7], etc.), the network may need to generate some specific RM cells to incite a source to reduce its transmission rate or to interrupt the transmission of an ATM block in progress. These

RM cells, referred to as traffic control messages, are identified by the Trf and Maintenance bits set equal to 1 and 0, respectively.

ABT/DT NETWORK MAINTENANCE. In some conditions (e.g., expiration of some timers, etc.), the network may need to retry some BCR negotiations on its own initiative. For this purpose, some specific maintenance messages, namely RM cells with Maintenance=1 and Trf=1, are generated by the network. Their role is to clear spurious bandwidth reservations, which have not been acknowledged because of RM cell loss or the absence of response from either user, and to retry the BCR negotiation, which is pending.

COLLISION. A traffic management procedure is identical to a usual BCR negotiation except that the originating entity is the network. In particular, during such a procedure, any conforming BCR negotiation initiated by the source is denied by the network and any BCR negotiation initiated by the destination fails or is denied by the network. Traffic management messages have priority over BCR negotiations initiated by either user. Moreover, traffic management messages issued at network ingress have priority over those initiated at network egress (forward priority principle).

2.5. Basic ABT/DT Protocol Principles

A bandwidth decrease is initiated by the source by sending a BWReq cell into the network with requested BCR values less than those currently allocated in the network. A BCR decrease is immediately taken into account by the network without acknowledgement. Any other BCR negotiation by the source is performed in two phases. The first phase is a BCR reservation phase, where each node checks whether the requested resources are available, possibly modifies the requested BCR values and reserves some resources. The second one is a BCR validation phase: if the reservation phase has been successful, the reserved resources are allocated during the BCR validation phase. The BCR validation phase is triggered by the BWAck cell, which is sent by the source and which delineates the ATM block. Each phase corresponds to the round trip of some RM cells through the network. For BCR modifications initiated by the destination, the BCR reservation and allocation phases are performed in a single RM cell round trip.

Each RM cell round trip is protected by a timer. In this version of the ABT/DT protocol, no messages are discarded in network nodes. Thus, when a message is sent by an ABT/DT unit, it is received by the other ABT/DT unit unless it is lost. An ABT/DT unit cannot discard a message that should be looped back unless it has sent another message with a higher priority and is waiting for a response.

ABT/DT transactions are identified by a sequence number. The source should use an even SN value and the destination an odd SN value in the BWReq cell. The SN value of subsequent BWReq cells issued by a given user with a given Dir bit should then be incremented by two (modulo 2^{32}). Those ABT/RM cells generated by the network for traffic management purposes are also identified by an SN value. A given ABT/DT unit increments by one the SN value of the successive RM cells in a given direction. The SN value is used to test whether the traffic mangement procedure is correctly executed and to eliminate spurious messages.

3. User BCR Negotiation Procedures

3.1. User BCR Negotiations

BCR DECREASE BY THE SOURCE. We assume in a first step that there is no collision between some BCR negotiations initiated by the source and the destination. To decrease the BCR values currently allocated within the network, the source should send a BW Req cell with Elastic=1 (as a convention a bandwidth decrease is in the rigid mode) and the requested BCR values for the CLP=0+1 and OAM cell flows satisfying after application of eq. (2) the inequality $\Lambda_{\text{Req}}(0+1) \leq \Lambda_{\text{Current}}(0+1)$. Immediately after issuing such an RM cell, the source should adapt its transmission rate in order to conform to the new BCR values.

After application of eq. (2), if as expected $\Lambda_{\text{Req}}(0+1) \leq \Lambda_{\text{Current}}(0+1)$, then Λ_{Current} is updated to the new BCR values and the BWReq cell is forwarded to the network nodes, which immediately take into account the bandwidth decrease. A BCR decrease is not acknowledged by the network as shown in Figure 3.

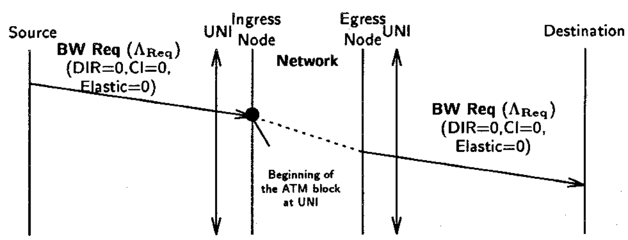

Figure 3. BCR decrease by the source without collision.

BCR INCREASE BY THE SOURCE. If after application of eq. (2), $\Lambda_{\text{Req}}(0+1) > \Lambda_{\text{Current}}(0+1)$, the BCR request is a BCR increase. The requested BCR values are examined in each network node, which accepts or denies the request according to a specific resource allocation policy and to the value of the Elastic bit, and updates the CI bit and BCR fields (elastic mode), if necessary. Upon reception of a BWReq cell with CI=1 in the rigid mode, a network node should ignore the bandwidth request and forward it to the next node, if any. Figure 4 shows the exchange of the messages across the User Network Interface (UNI) for a BCR increase initiated by the source in the rigid and elastic modes.

The BWReq cell with the BCR values to be allocated to the connection is forwarded at network egress to the destination for information if CI=0 or Elastic=0 and is looped back in the form of a BWAck on the backward direction towards the source. If Elastic=1 and CI=1, the resources reserved in a network node are released upon reception of this BWAck cell. If Elastic=0, the reserved resources are adjusted to the BCR values conveyed by the BWAck cell.

The BWAck cell is forwarded at network ingress to the source. Unless CI=1 and Elastic=1, the source should send in turn a BWAck cell on the forward direction with BCR values less than or equal to those communicated by the network. This BWAck cell sent by the source is the leading RM cell of a new ATM block. Upon reception of this RM cell, the BCR values are controlled at network ingress and updated if necessary as in eq. (3). The resources reserved in the network nodes are allocated to the connection upon reception of this BWAck cell propagating on the forward direction.

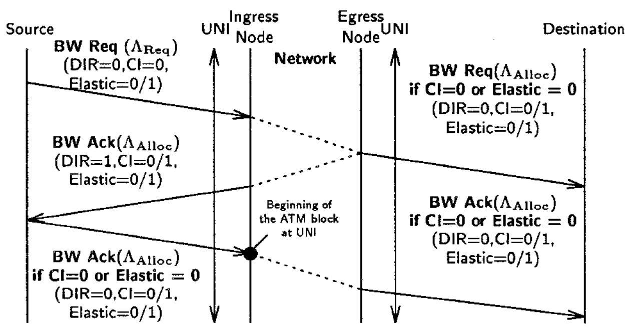

Figure 4. BCR increase initiated by the source with no collision.

BCR MODIFICATION BY THE DESTINATION. As an improvement on FRP/DT, ABT/DT allows the destination to perform BCR modifications (in the rigid as well as elastic modes) applicable to the traffic generated by the source (see Figure 5). To achieve this, the destination may send to the network a BWReq cell with Dir=1. Upon reception of this RM cell by the network, $\lambda_{\text{Req}}(0+1)$ and $\lambda_{\text{Req}}(\text{OAM})$ are updated as in eq. (2), where Λ_{Req} is replaced with λ_{Req}. If $\lambda_{\text{Req}}(0+1) < \Lambda_{\text{Current}}(0+1)$, the BWReq cell indicates a bandwidth decrease and the Elastic bit is set equal to 1.

In each network node, the requested BCR values are examined. In the rigid mode, if CI=0 and the requested resources are available according to the given resource allocation policy, they are reserved and the BWReq cell is forwarded to the next node; otherwise, no resources are reserved, the CI bit is set, if necessary and the RM cell is forwarded to the next network node. In the elastic mode, BCR values are reserved and the CI bit is set according to the bandwidth fair share algorithm for elastic BCR negotiations; the BCR values may be modified.

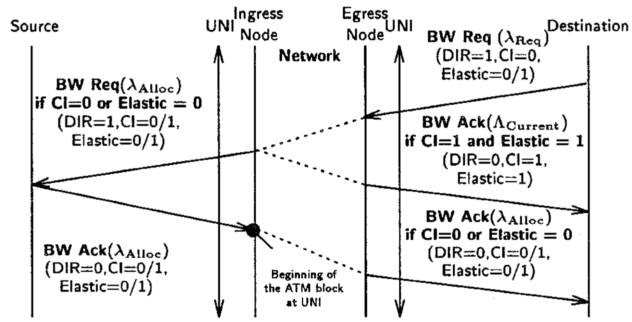

Figure 5. BCR modification initiated by the destination with no collision.

Unless CI=1 and Elastic=1, the BWReq cell is forwarded at network ingress to the source. This RM cell conveys the BCR values, that are to be allocated to the connection. Upon reception of this RM cell, the source should send a BWAck cell and the negotiation proceeds as for a BCR modification initiated by the source.

In the rigid mode, when the BWReq cell arrives at network ingress with CI=1, this RM cell is not forwarded to the source but looped back to the destination in the form of a BWAck cell with CI=1 and the current BCR values Λ_{Current} at network ingress. Upon reception of this message in a network node, the reserved resources, if any are released.

COLLISION BETWEEN BCR NEGOTIATIONS. So far, we have assumed that there is no collision between the BCR negotiations initiated by the source and the destination. In fact, the procedure presented above for BCR negotiations initiated by the destination is still valid, except that the BWReq cell due to the destination is sent at network ingress to the source even if Elastic=1 and CI=1 (see Figure 6).

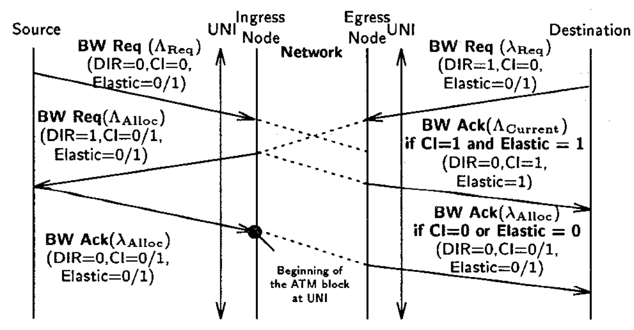

Figure 6. BCR modification initiated by the destination with collision.

Problems due to collision between two BCR negotiations are solved by applying the backward priority principle presented in Section 2. Upon reception at a network node of a BWReq cell due to the destination, any BCR reservation by the source is canceled. Any BWReq cell issued by the source and arriving at a network node while some resources have been reserved by a Bandwidth Request RM cell issued by the destination is forwarded without being processed and the CI bit is set. Upon reception at network egress, it is discarded and no acknowledgement is sent to the source.

3.2. User Maintenance Procedure

STATUS_ENQUIRY PROCEDURE. A Status_Enquiry procedure is identical to a usual BCR negotiation in the rigid mode (Elastic bit set equal to 1) except that the Maintenance bit is set equal to 1. Moreover, Status_Enquiry messages have the CI and Trf bits systematically set equal to 0. Finally, the BCR values conveyed by such a message issued by either user are irrelevant and may be set equal to 0. Upon reception of a Status_Enquiry message at network ingress (resp. egress), the BCR fields are updated to $\Lambda_{\text{Current}}(0+1)$ and $\Lambda_{\text{Current}}(\text{OAM})$ (resp. $\lambda_{\text{Current}}(0+1)$ and $\lambda_{\text{Current}}(\text{OAM})$).

Upon reception of a Status_Enquiry message in a network node, the requested BCR are examined. If they are available according to the resource allocation policy for rigid BCR negotiations, then the BCR value $\Lambda_n(0+1)$ is updated to $\Lambda_{\text{Req}}(0+1)$. Otherwise, the requested BCR is updated to $\Lambda_n(0+1)$. In any case, the CI bit is left unchanged. As for a successful BCR negotiation in the rigid mode, the BCR values communicated by the network should be acknowledged by the source.

NOT_READY MESSAGES. Not_Ready messages are generated by the network when either user tries to initiate a new BCR negotiation when another one initiated by this same user in the same direction is still in progress within the network or when the network has initiated a traffic management procedure. Note that there is no need to send such a message to the source when the network waits for a response from the source.

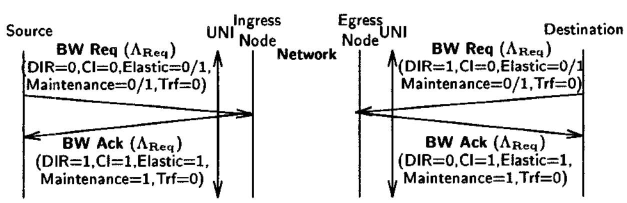

Figure 7. Not_Ready messages.

A Not_Ready message at network ingress is a BWAck cell with CI=1, Maintenance=1, Trf=0, Elastic=1, Dir=1, and the BCR values Λ_{Current}. At network egress, Dir=0 and the BCR values conveyed by such a messages are λ_{Current}. Not_Ready messages are shown in Figures 7.

4. Traffic Management

4.1. Traffic Control Procedures

For traffic control purposes, ABT/DT units located at network ingress and egress as shown in Figure 1 may generate traffic control messages identified by Trf=1 and Maintenance=0. These messages are identical to usual BCR negotiation messages (in the rigid or elastic modes) except that the originating entities are not the source or the destination but network ABT/DT units. Moreover, traffic control messages may be generated at any time, even if a user BCR negotiation is in progress within the network. As mentioned in Section 2, traffic control messages have priority over BCR negotiations initiated by either user and the forward priority principle applies between traffic control messages.

To initiate a traffic control procedure in the rigid mode, the ingress (resp. egress) ABT/DT unit sends a BWReq cell on the forward (resp. backward) direction with the BCR values Λ_{Current} (resp. λ_{Current}) currently allocated at network ingress (resp. egress). Upon reception of this message in a network node, the requested resources are examined. If they are available according to the resource allocation policy for rigid BCR negotiations, the requested BCR is reserved in the network node if it is greater than $\Lambda_n(0+1)$ currently allocated to the CLP=0+1 cell flow. Otherwise, the BCR value is updated to the BCR value currently allocated in the network node and the CI bit is set if necessary.

At network ingress, some BCR values are communicated to the source. If CI=1, then the current amount of bandwidth can no longer be allocated within the network and the network may incite the source to modify its transmission rate. In such a case, the BCR values $\Lambda_{\text{Alloc}}(0+1)$ and $\Lambda_{\text{Alloc}}(\text{OAM})$ actually communicated to the source depend on the traffic policing strategy adopted by the network. For instance, $\Lambda_{\text{Alloc}}(0+1)$ may be set equal to 0 or to a minimum BCR value. $\Lambda_{\text{Alloc}}(\text{OAM})$ may be set to the minimum between $\Lambda_{\text{Current}}(\text{OAM})$ and $\Lambda_{\text{Alloc}}(0+1)$. The source should acknowledge the BCR values communicated by the network.

For a traffic control procedure in the elastic mode, the BCR values contained in the BWReq cell issued at network ingress or egress are equal to the PCR values Λ_{PCR}. Upon reception of this BWReq cell, some resources are reserved in the network nodes by applying the bandwidth fair share algorithm as for a usual BCR negotiation in the elastic mode. The BCR values computed all along the connection are then systematically communicated to the source. For a traffic control procedure initiated at network ingress, the BCR values reserved in network nodes are adjusted, if necessary upon reception of the BWAck cell propagating on the backward direction as for usual BCR negotiations.

A traffic control procedure initiated at network ingress (resp. egress) is depicted in Figure 8 (resp. Figure 9). To illustrate the collision principles presented in Section 2, it is assumed in these figures that a BCR negotiation has been initiated by the source or the destination. If a user initiated BCR negotiation is still in progress in the network while the traffic control procedure gets started, then the BWAck cell corresponding to the user request is discarded, if necessary by the ABT/DT unit, which has initiated the traffic control procedure.

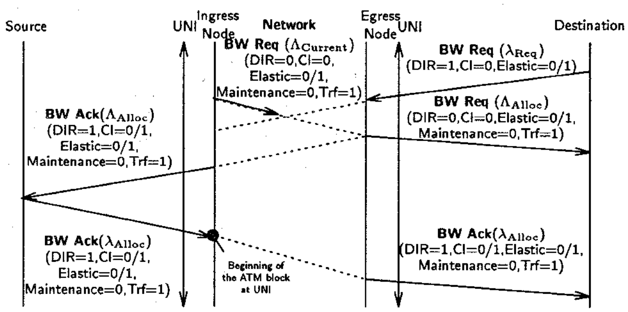

Figure 8. Traffic control procedure initiated at network ingress.

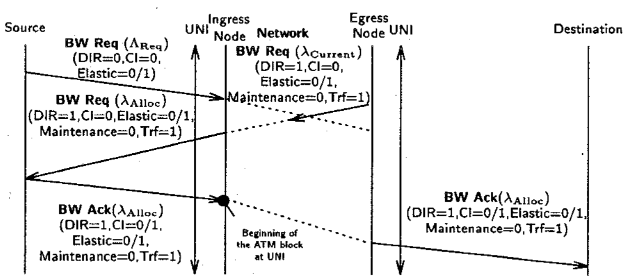

Figure 9. Traffic control procedure initiated at network egress.

4.2. BCR negotiation Control - Network ABT/DT Maintenance Procedures

TIMERS ASSOCIATED WITH BCR NEGOTIATIONS. Since some RM cells may be lost within the network, it may happen that after initiating a BCR negotiation, a given entity (the source, the destination, etc.) does not get any response from the network or either user, or receives RM cells with incorrect SN values. To recover from such error conditions, some timers are introduced as follows:

1. W_{Source} and $W_{Destination}$: after initiating a BCR negotiation, the source or the destination should get a response from the network within a fixed time interval. After that delay the source or the destination may reattempt the BCR negotiation. The same timer is used by the destination to check that it gets a BWAck cell from the network after it has received a BWReq cell.

2. $W_{Ingress}$, W_{Egress}: after sending a BWReq cell into the network (or looping back a BWAck cell at network egress), the ingress (resp. the egress) ABT/DT unit should get a response from the network within a fixed time interval as shown in Figure 10 (resp. Figure 11). After that delay the ingress (resp. the egress) it may initiate a network maintenance procedure (see below).

Note that a BWReq cell is considered as an answer from the network if it has a higher priority than the negotiation initiated by that entity.

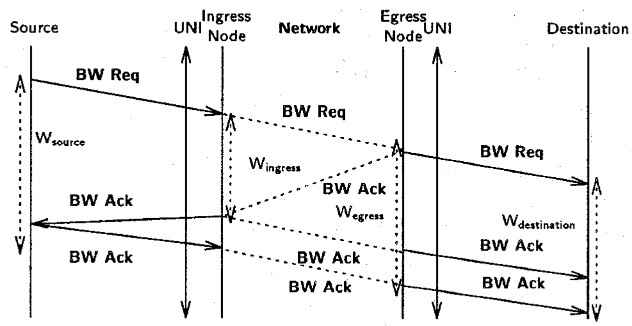

Figure 10. Timer associated with a BCR negotiation at network ingress.

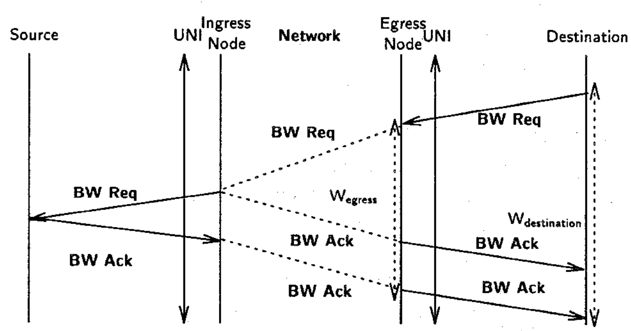

Figure 11. Timers associated with a BCR negotiation processed at network egress.

ABT/DT NETWORK MAINTENANCE. Upon expiration of timer $W_{Ingress}$ (resp. W_{Egress}), a BWReq cell with Maintenance=1, Trf=1, and an appropriate sequence number is sent by the ingress (resp. egress) ABT/DT unit into the network on the forward (resp. backward) direction with

1. the BCR Λ_{Current} (resp. λ_{Current}) if the reserved BCR $\Lambda'_{\text{Current}}(0+1) = 0$ (resp. $\lambda'_{\text{Current}}(0+1) = 0$) or $\Lambda'_{\text{Current}}$ (resp. $\lambda'_{\text{Current}}$) if $\Lambda'_{\text{Current}}(0+1) > 0$ (resp. $\lambda'_{\text{Current}}(0+1) > 0$),

2. the Elastic/Rigid bit of the BCR transaction, which is pending.

This is intended to retry an incomplete BCR transaction.

A network maintenance ABT/DT RM cell is handled in network nodes as a traffic control message. In particular, the requested resources are examined by using the resource allocation policy corresponding to the value of the Elastic bit of the message. Upon reception at network ingress of the BWAck (resp. BWReq) cell corresponding to expiration of timer W_{Ingress} (resp. W_{Egress}), the RM cell is forwarded to the source, which should acknowledge this message by sending a BWAck cell on the forward direction. The BCR values in network nodes are updated upon reception of this cell.

Expiration of timers at ingress node (resp. egress node) and the ensuing maintenance procedure is depicted in Figure 12 (resp. Figure 13).

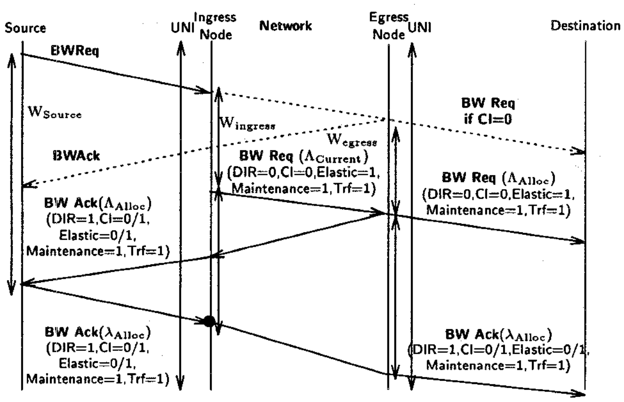

Figure 12. Expiration of timer W_{Ingress}.

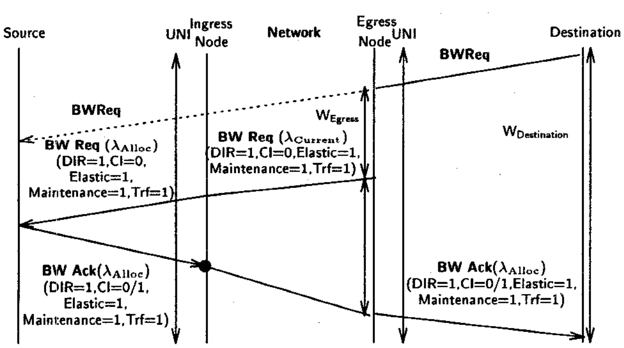

Figure 13. Expiration of timer W_{Egress}.

5 Conclusion

This paper presents a protocol intended to support the ABT/DT capability. It notably takes into account all the features of ABT/DT specified in the latest version of ITU-T Recommendation I.371. The complexity of the protocol may however raise implementation issues. In the early experiments of ABT/DT, the procedure could be implemented in the core network with a limited number of primitives. BCR negotiations may for example be periodically initiated by the network in order to avoid the implementation of the timers. The source would thus have just to acknowledge the BCR values communicated by the network.

In fact, elastic ABT/DT is very close to explicit rate indication (ERI) procedures introduced by Jain in [9]. The major difference, however, is that ERI procedures are intended to estimate the amount of available network resources without making any bandwidth reservation while ABT/DT procedures are precisely meant to reserve bandwidth. This is why in ERI, unlike ABT/DT, several RM cells are simultaneously present within the network.

References

[1] The ATM Forum. UNI Spec 3.1. Boston, September 1993.

[2] The ATM Forum. UNI Spec 4.0-Traffic management specification. December 1995.

[3] P. Boyer. A congestion control for ATM. In *Proc. Teletraffic Expert Seminar*, Morristown, October 1990.

[4] P. Boyer and D. Tranchier. A reservation principle with applications to the ATM traffic control. *Computer Networks and ISDN Systems* 24, pp. 321-334, 1992.

[5] J. Boyer. Statistical multiplexing of data ATM connections. *Proc. 2nd International Conference on Telecommunication Systems Modeling and Analysis*, Nashville, March 1994.

[6] B. Doshi and S. Dravida. Congestion control for bursty data in high-speed wide area packet network. In *Proc. Teletraffic Expert Seminar*, Morristown, October 1990.

[7] F. Guillemin and P. Boyer. ATM block transfer capabilities: the special case of ABT/DT. In *Proc. Globecom'96*, London, November 1996.

[8] ITU-T Recommendation I.371. Traffic control and congestion control in B-ISDN. Geneva, May 1996.

[9] R. Jain. Congestion control and traffic management in ATM networks: recent advances and a survey. Submitted, 1995.

[10] M. Parker et al. The European ATM pilot. In *Proc. ISS'95*, Berlin, April 1995.