QoS based Integration of IP and ATM: Resource Renegotiation

Silvia Giordano Jean-Yves Le Boudec

Institute for Computer Communications and Applications (ICA) ,
EPFL CH-1015 Lausanne

September 29, 1998

Abstract

The dynamic and efficient usage of the resources is one of the fundamental aspects studied in the DIANA project [1]: at first the traffic specification should reflects the real traffic demand, but, at the same time, optimise the resources requested. In applications, which uses RSVP, this can be reached with the introduction of the RVBR Service described in [2], which is based on the renegotiation of the traffic specification. In this paper we present and discuss the RVBR service in detail. Then we describe how it applies to resource reservation for Internet traffic with RSVP, and we show some results from simulation.

1 Introduction

The integration of voice, data, and video services modified the target of networking technologies. Instead of providing a single type of service (e.g. best effort), the network now has also to deal with the integration of services and, related with that, with providing Quality of Service (QoS).

ATM and IP both offer protocols for resources reservations, which provide QoS. However, while applications are mainly IP, there exists a large ATM backbone with high bandwidth. It is therefore straightforward to promote the integration of those two technologies.

Several papers give technical overviews on the competing integrated services network solutions [3], and this is also the topic of the work of the NIG G3 IP ATM Integration Chain Group of the European Community [4].

The ACTS project called DIANA focuses on the integration of ATM with different IP protocols for resources reservation: RSVP [5], Scalable Reservation Protocol (SRP) [6] and Simple Integrated Media Access (SIMA) [7]. The latter two are based on differentiated service architecture. DIANA networking model for RSVP is based on IP as a common network layer and the assumption that end systems are connected to different link layers but the applications request resources in terms of RSVP traffic specification (*Tspec*).

In Section 2 we present the DIANA project, and the architecture proposed in the project. Then we introduce the dynamic reservation aspect of the RSVP

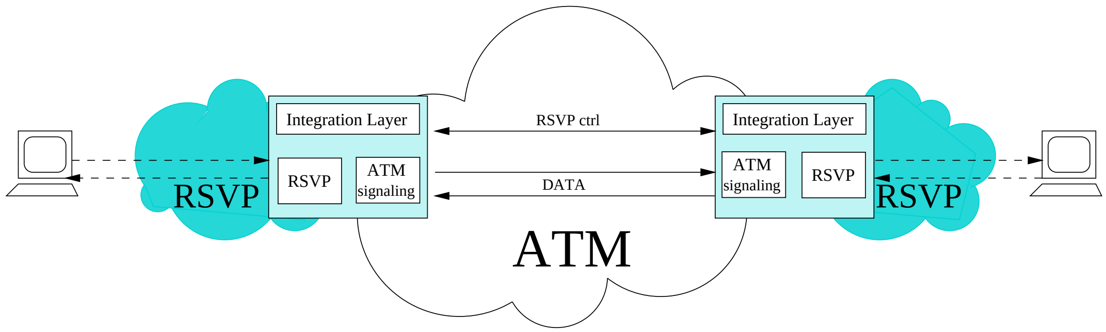

Figure 1: DIANA network scenario

reservation protocol and some application related interworking and traffic control issues between ATM and IP based networks.

In Section 3 we describe the Renegotiated Variable Bit Rate (RVBR) service. This is introduced in [2] in order to specify the complex traffic generated by multimedia applications. RVBR service allows to adjust in a dynamic way the renegotiated traffic parameters to the input traffic. An RVBR source is characterised by a renegotiable leaky bucket specification (with rate r and depth b) plus a fixed size buffer X drained at maximum at renegotiable peak rate p.

In real life examples of this service are traffic shaping done at source sending over VBR connections as defined in [8] and Internet traffic that takes the form of Int-Serv specification (e.g. Controlled Load (CL) [9]) with RSVP reservation.

To the RVBR service is associated with an algorithm, which solve the problem of finding the traffic parameters for the next renegotiation interval when the cost of a set of traffic parameters is represented by a linear function and the input traffic in the next interval is known[1]. Here we present how the RVBR service and this algorithm can be introduced with evident benefits and without remarkable additional cost into an application that uses RSVP to reserve the resources to the network.

In Section 4.1 we consider this case study, and we show, by means of simulations, the benefits of using the RVBR Service for multimedia IP traffic with RSVP.

2 The DIANA Project

DIANA is a project of the European Union 4th Framework Programme ACTS started in March 1998. As its main goal, the DIANA consortium will develop, integrate, validate and demonstrate resource reservation and traffic control functionality to seamlessly interoperate between ATM and IP Int-Serv networks in order to provide guaranteed QoS end-to-end [10]. Although DIANA will mainly focus on RSVP and ATM, the design of the trial platform will be kept flexible enough to allow investigating different solutions, such as SRP or SIMA. DIANA's networking model for RSVP is based on IP as a common network layer and the assumption that end systems are connected to different link layers

[1] In real life we may use approximation

but use RSVP to provide the application with the control capabilities of the respective layer underneath and/or the next RSVP capable network element.

A device named Integration Unit [10] is placed at the boundary between ATM and IP domains to provide the functionality for the translation between IP and ATM reservation protocols, as illustrated in Figure 1. The control plane of this Integration Unit is assigned the key role for prototyping the translation from RSVP and ATM UNI [11] signalling and vice versa, for the mapping of QoS specifications given by the flow descriptors objects with Integrated Services and ATM traffic descriptor information elements respectively, and for the allocation of ATM virtual connections for IP flows. The DIANA architecture will be implemented on the Linux platform of the Flextel 1200 [12].

2.1 QoS in IP Networks

Two opposite directions from the resource reservation protocol point of view are represented by the explicit reservation model and the implicit reservation model [13]. With the explicit reservation model a traffic profile is negotiated between users and the network. The implicit reservation model tries to obviate the difficulty of specifying the traffic parameters by aggregating flows into the network without explicit signaling of flow parameters.

DIANA project is primarily based on explicit reservation protocols: RSVP (and ATM), but also implicit reservation protocols like SIMA and SRP are investigated.

For simplifying either the management of the router status and the signaling procedures, RSVP uses soft state for the reservation. This means when a reservation is made, it must be periodically refreshed. The soft state does not imply that resources are renegotiated, because the traffic parameter specification can be reissued without changes. However, this mechanism allows for expressing dynamic reservation changes in a straightforward way and thus can be easily used to support resource renegotiation, as we illustrate in Section 4.1.

2.2 Application Related Interworking and Traffic Control Issues

The traffic generated by multimedia applications presents a high degree of burstiness that can be hardly described by a static set of traffic parameters.

Since the traffic specification provided by signalling controls the network's traffic management (traffic control, congestion control), the application's traffic specification plays an important role for QoS. The traffic generated by applications must be compliant to its specification and the specification has to reflect its real demand. The way an application generates a traffic specification is implementation and service specific, but can in general either be complex or only a rough approximation. In fact it is not always possible to describe the (complex) traffic in terms of a single traffic descriptor, therefore, if we do not want to use to many resources or to have unacceptable performance, it is necessary to renegotiate it.

In the scenario illustrated in Figure 1 the application describes its traffic in terms of RSVP traffic specification, which is mapped onto ATM traffic specification at the network. A fundamental task is for the application, which has to describe the traffic in form of *Tspec* [14]. As we said above, the application can

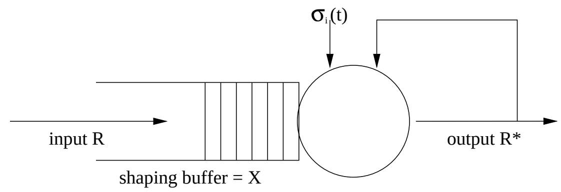

Figure 2: RVBR Service. There is a renegotiable leaky bucket specification (with rate r and depth b) plus a fixed size buffer X drained at maximum at renegotiable peak rate p.

better do it when it can apply renegotiation. The introduction of the RVBR Service at application layer is assumed to simplify and generalise this task. This service allows an application using RSVP protocol with Int-Serv traffic specification, not only to specify the traffic for the initial negotiation, but also to to find the optimal *Tspec* for the next renegotiation. Whenever renegotiation is taking place, the RVBR scheme generates the traffic specification that conforms to the real demand, in order to reallocate the network resources in an optimal way while guaranteeing QoS to the traffic flows.

RVBR service uses the knowledge of the past status of the system and the profile of the traffic expected in the near future, which can be either pre-recorded or known by means of exact prediction. This scheme suits perfectly the dynamics of the traffic generated by multimedia application. Moreover it naturally integrates with the soft state mechanism of RSVP.

3 Resource Renegotiation: RVBR Service

As introduced above, the RSVP protocols offers intrinsic mechanisms for renegotiation. The original role of the soft state mechanism is to simplify the status management in the routers, but it can be easily and without additional costs, used for renegotiating the resources.

The scheme designed for the RVBR service is intended to be integrated with applications that use an explicit reservation protocol, such as RSVP.

The RVBR service is based on a renegotiable VBR traffic specification, and offers a scheme for optimising the traffic specification in the next period of time where this traffic specification is valid.

In the following section we briefly describe the RVBR service in terms of network calculus [15] and the algorithm used to optimise the traffic specification. We then proceed to illustrate how this applies in the case of applications using the RSVP protocol with int-serv traffic specification.

3.1 Overview of RVBR Service

We first recall the characterisation of the RVBR service in terms of input and output functions as given in [2].

The elements of a RVBR source, as illustrated in Figure 2, are a renegotiable leaky bucket specification (with rate r and depth b) plus a fixed size buffer X drained at maximum at renegotiable peak rate p [2].

[2]In [2] the RVBR service is described with two leaky bucket specifications. In the case

The observation time is divided into intervals, and $I_i = (t_i, t_{i+1}]$ represents the i-th interval. Inside each interval the system does not change. The parameters of the RVBR service in I_i are indicated with (p_i, r_i, b_i).

The RVBR service is completely defined by:

- the time instants t_i at which the parameters change
- the RVBR parameters (p_i, r_i, b_i), for each interval I_i
- the fixed shaping buffer capacity X

A RVBR source cannot send more than the traffic specified by the shaping function σ_i, defined as

$$\sigma_i(u) = \min(p_i \cdot u, r_i \cdot u + b_i) \tag{1}$$

Moreover the RVBR service, at the transient times t_i between two adjacent intervals keeps the level of the buckets and restarts from that level at the next interval. The justification of this choice can be found in [2]. Therefore there is another function, resulting from taking into account the bucket level $q(t)$, which limit the traffic in I_i

$$\sigma_i^0(u) = \min(p_i \cdot u, r_i \cdot u + b_i - q(t_i)) \tag{2}$$

If we indicate with the function $R(t) : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ the amount of traffic that has entered in the system in time interval $[0, t]$, the resulting output $R^*(t)$ is given by Proposition 5 of [2]

$$R^*(t) = \min\left(\sigma_i^0(t - t_i) + R^*(t_i), \inf_{t_i < s \leq t}(\sigma_i(t - s) + R(s))\right) \tag{3}$$

3.2 Optimisation of the RVBR parameters

This input-output characterisation of the RVBR service, is further used to solve the problem of finding, at any renegotiation, the optimal σ_i to negotiate with the network. This problem is well know to have no trivial solution. For example some input traffic could be specified from a large r_i and a small b_i as well as from a small r_i and a large b_i.

In [2] the authors introduce an algorithm (*localOptimum*) that finds the optimal solution when the choice of the network is driven by a linear cost function.

The peak p_i results to have an absolute minimum that corresponds to the effective bandwidth of $R(t)$ in this interval [15].

$$p_i = \max\left(\sup_{t,s \in I_i} \frac{R(t) - R(s) - X}{t - s}, \sup_{t \in I_i} \frac{R(t) - R(t_i) - X + w(t_i)}{t - t_i}\right) \tag{4}$$

where $w(t)$ is the backlog of the shaping buffer.

of RSVP the bucket associated to the peak p is the MTU size, hence it is fixed. We further assume it equal to zero to simplify the computation, given that this is not a limitation

The other two parameters r_i, and b_i are derived by solving the following optimisation problem assuming a linear cost function $u \cdot r_i + b_i$ which represent the cost of the traffic specification to the network

$$\text{minimise } u \cdot r_i + b_i \text{ in the region } \begin{cases} 0 \leq r_i \leq r_{max} \\ 0 \leq b_i \leq b_{max} \\ b_i + r_i \cdot s + X - \beta_i(s) \geq 0 \quad \forall s \in I \end{cases} \quad (5)$$

where $I = [0, t_{i+1} - t_i]$, and (r_{max}, b_{max}) represent the maximum values of the rate r and the bucket b, respectively.

$\beta_i(s)$ is a function that computes the maximum amount of traffic sent over the any interval of size s, taking in account the conditions at time t_i

$$\beta_i(s) = \max\left(\sup_{0 \leq v \leq t_{i+1} - t_i - s} (R(v+s) - R(v)), R(s+t_i) - R(t_i) + w(t_i) + q(t_i) \right)$$

The solution to this problem is expressed in terms of the following algorithm [3]

Algorithm 1 localOptimum$(X, (R(t))_{t \in I_i}, b_{max}, r_{max}, u, w(t_i), q(t_i), t_{i+1})$

if $b_{max} < \sup_{s \in I}(\beta_i(s) - r_{max} \cdot s - X)$ **then** there is no feasible solution;
else {
  $p_i = \max\left(\sup_{t,s \in I_i} \frac{R(t) - R(s) - X}{t - s}, \sup_{s \in I_i} \frac{R(s) - R(t_i) - X + w(t_i)}{s - t_i} \right)$;
  if $u \leq 0$ **then** {
    $x_0 = min(r_{max}, p_i)$;
  }
  else {
    $x_0 = \sup_{s \in I} \frac{\beta_i(s) - \beta_i(u)}{s - u}$;
    $x_A = \sup_{s \in I, s > 0} \frac{\beta_i(s) - X - b_{max}}{s}$;
    $x_B = \sup_{s \in I, s > 0} \frac{\beta_i(s) - X}{s}$;
    if $(x_0 > \min(x_B, r_{max}, p_i))$ **then** $x_0 = \min(x_B, r_{max}, p_i)$;
    else if $(x_0 < x_A)$ **then** $x_0 = x_A$;
  }
  $r_i = x_0$;
  $b_i = \sup_{s \in I}(\beta_i(s) - X - s \cdot x_0)$;
}

4 Application of the RVBR Service to RSVP

In this section we describe how we used the previous algorithm to simulate a typical real case: transmission of MPEG2-encoded video using the IntServ Controlled Load service with the RSVP reservation protocol.

In RSVP the sender sends a PATH message with a *Tspec* object which characterises the traffic it is willing to send. If we consider a network that

[3]this algorithm requires that β_i is concave. When this is not true we transform it in a concave function.

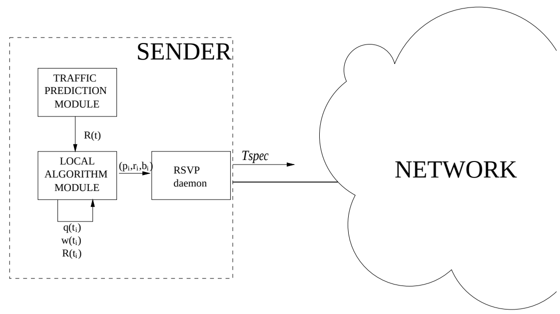

Figure 3: A basic architecture to support the usage of the RVBR service for RSVP with CL service reservation: every 30 seconds $R(t)$ is predicted and used to compute the optimal p, r and b to generate the new *Tspec*. The traffic prediction module is substituted by some other module for pre-recorded traffic.

provides a service as specified for the Controlled Load service (CL) the *Tspec* takes the form of a double bucket specification [16] as given by the RVBR service. There is a peak rate p and a leaky bucket specification with rate r and bucket size b. Additionally there are also a minimum policed unit m and a maximum packet size M. We neglect m and M, which are assumed to be fixed. With RSVP as reservation protocol, the reservation has to be periodically refreshed. The suggested period is 30 seconds. Therefore p, r and b need to be reissued at each renegotiation time. There is no additional signaling cost in applying a *Tspec* renegotiation at that point, even if there exists some computational overhead due to the computation of the new parameters, to the call admission control etc. It is important to note here that, contrary from the negotiation of a new connection, with the renegotiation the reservation is never interrupted.

If the requested traffic specification cannot be supported by the network the old traffic specification is restored, and the network may not be able to accommodate the next traffic. Mechanism to prevent this failure for occurring are still under study. Here we assume that the *Tspec* is accepted all over the network as well as at the destination, such that the the source can transmit conforming to its desired traffic specification.

To apply the RVBR service in this scenario we assume that at any time $t_i = 30 \cdot i$ the application knows (because pre-recorded or predicted) the traffic for the next 30 seconds. We further assume to know the cost of the *Tspec*s to the network (indicated by the cost function $u \cdot r + b$) and b_{max} and r_{max}. The backlog $w(t_i)$ and the bucket level $q(t_i)$ can be measured in the system. Then we use the algorithm *localOptimum* at Section 3.2 for computing the *Tspec* the sender will send at the next renegotiation time. The basic architecture of the sender node is described in Figure 3.

4.1 Simulation results

Here we illustrate and discuss the simulation results obtained in a scenario similar to what we will use in DIANA: IntServ services with RSVP reservation protocol.

In our simulations, we use a 4000 frame-long sequence conforming to the

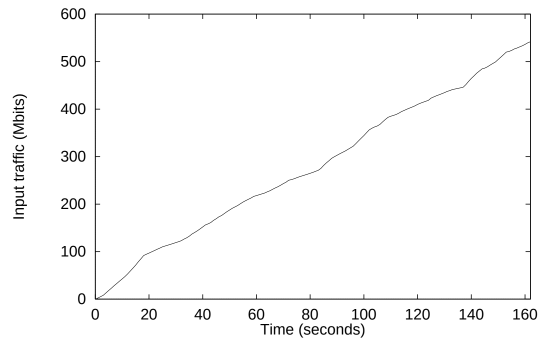

Figure 4: Traffic evolution of the sequence used as input in the simulation.

ITU-R 601 format ($720 * 576$ at 25 fps). The sequence is composed of several video scenes that differ in terms of spatial and temporal complexities. It has been encoded in an open-loop variable bit rate (OL-VBR) mode, as interlaced video, with a structure of 11 images between each pair of I-pictures and 2 B-pictures between every reference picture. For this purpose, the widely accepted TM5 video encoder [17] has been utilised.

The traffic generated by the video is transported by a trunk regulated by a RVBR service (p, r, b) with shaping buffer X. In this context we do not consider any scheduling issues, that is material of ongoing work. Therefore we assume that the video, with a total size of 550 Mbits, is transmitted in 163 seconds (25 frames pro second). The cost function is linear with u. We consider three different scenarios:

Scenario 1: $X = 40$ Mbits, $r_{max} = 5$ Mbps, $b_{max} = 9$ Mbps and $u = 1$

Scenario 2: $X = 30$ Mbps, $r_{max} = 6$ Mbps, $b_{max} = 12$ Mbits and $u = 1$

Scenario 3: $X = 20$ Mbits, $r_{max} = 8$ Mbps, $b_{max} = 10$ Mbps and $u = 6$

The initial conditions are: $q(0) = 0$ and $w(0) = 0$. The file is pre-recorded and, given that we do not enter in scheduling matters, we know $R(t)$ for all t. At time t_i we know $R^*(t)$ for $t \leq t_i$, we measure $w(t_i)$, $q(t_i)$ and compute $\beta_i(t)$. We obtain the optimal shaper parameters by applying the algorithm as described above. The evolution of the input traffic is given in Figure 4.

4.2 Backlog evolution with and without renegotiation

In Figure 5 we plot the backlog for the three scenarios in both cases where we apply the renegotiation and where we do not renegotiate [4]. In order to better distinguish the two approaches, the area of the curve representing the case without renegotiation is coloured.

We observe that in the beginning the curves representing the two approaches have the same behaviour. This is because the traffic is very high in the first 30 seconds, and both traffic specifications conform to this traffic.

After that period the traffic rate decreases. The case without renegotiation has to keep the traffic specification negotiated at time $t = 0$, even if it is no

[4] Even in this case we compute the optimal traffic specification as introduced in [15].

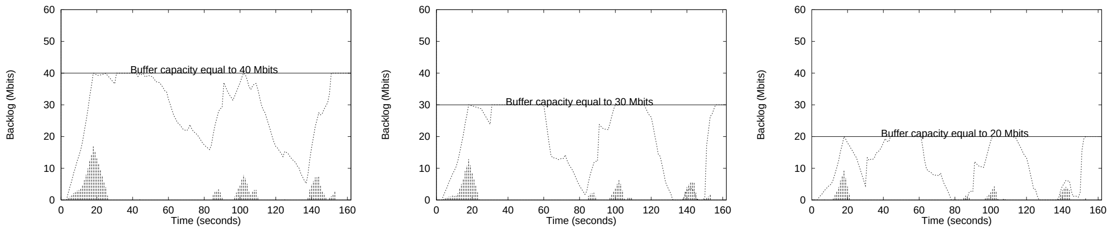

(a) Scenario 1 (b) Scenario 2 (c) Scenario 3

Figure 5: Comparison of the shaping buffer used with renegotiation (white area) and without renegotiation (black area) for the three scenarios

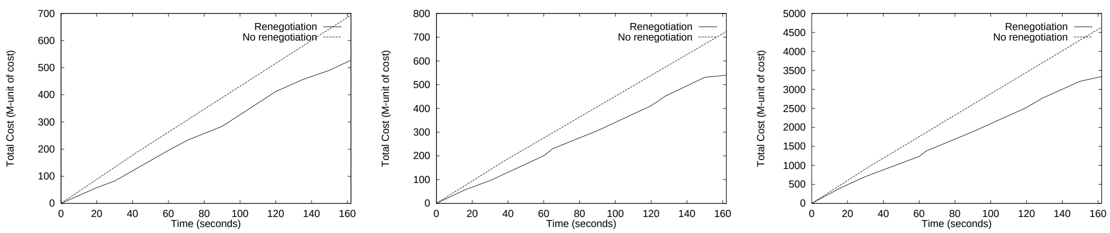

(a) Scenario 1 (b) Scenario 2 (c) Scenario 3

Figure 6: Comparison of the cost of allocating a renegotiated traffic specification and a traffic specification without renegotiation for different scenarios. The cost of the traffic specification is given in “millions of unit of cost” (M-unit of cost) and computed with the linear cost function used for the optimisation.

longer adequate for the demand. The resources allocated in the network are so large that it is possible to empty the buffer, and thereafter the buffer is rarely used.

The curve for the case where we used the RVBR service shows that the buffer is much better utilised, because the traffic specification decreases in the next intervals.

Therefore in the approach, where we apply the renegotiation with the RVBR service the resources in the network are much better used. In fact, when the buffer is almost always filled the output is conforms to the traffic specification, and this means that all the resources in the network are optimally used.

In the first scenario the usage of the buffer without renegotiation is 13%, while with renegotiation it is 58%. In the second scenario the percentages are 59% and 11%; in the last one they are 60% and 11%. In any case we have to remember that the optimisation is done for the worst case, and this explains why, when we do not renegotiate, the buffer never fills completely.

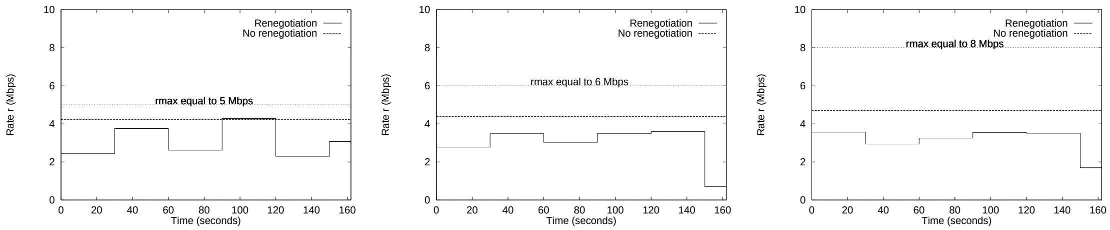

(a) Scenario 1 (b) Scenario 2 (c) Scenario 3

Figure 7: Comparison of the evolution of the rate r with renegotiation and without renegotiation for different scenarios

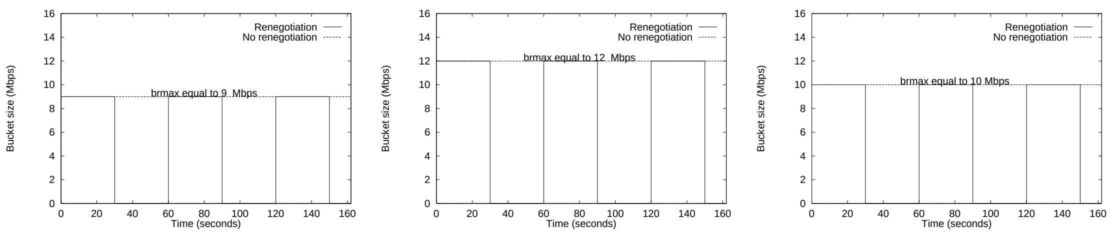

(a) Scenario 1 (b) Scenario 2 (c) Scenario 3

Figure 8: Comparison of the evolution of the bucket b with renegotiation and without renegotiation for different scenarios

4.3 Cost evolution with and without renegotiation

In the graphs in Figure 6 we compare the two approaches in terms of the cost of the traffic specification to the network.

The cost of the traffic specification is given by the linear cost function used by the RVBR service in order to compute the optimal traffic parameters. In the previous section we showed, for the case where we renegotiate the traffic specification, a better utilisation of the shaping buffer, that coincide with a better utilisation of the shaping buffer and consequently of all the resources allocated into the network. The additional result we derive from those other figures is that there is also a substantial advantage from the cost point of view, because the cost of the traffic specifications is in general smaller than or equal to the cost of the one allocated for the case when we do not use the renegotiation. This is even more evident from the figures in the next section.

4.4 Traffic specification parameters evolution with and without renegotiation

Figures 7 and 8 illustrate the fact that with renegotiation we can optimise the resources requested to the network and therefore at the end the total r and b allocated in this case are in general smaller. We also notice that inside an interval it can happen that the RVBR service allocates a *Tspec* that is larger than the one used when not renegotiating. This occurs when the traffic is very bursty and the buffer is full from the previous interval. For scenario 1 this situation occurs also at the forth interval (90 . . . 120 seconds), as illustrated in Figure 7. This happens because the buffer is full and the bucket is not sufficient to absorb the burstiness of the input traffic. It does not take place in scenario 2 and 3, because there is more bucket available and therefore the application can request a larger bucket b.

5 Conclusion

Work in progress in DIANA mainly focuses on the specification, implementation and evaluation of signalling translation, and related with that, traffic and QoS parameter mapping between IP and ATM networks.

One fundamental aspect is to enable the application to renegotiate the traffic specification in order to adjust to the dynamics of the real demand.

When the reservation protocol used is RSVP, the soft state mechanism allows for expressing dynamic reservation changes in a straightforward way and thus can be easily used to support resource renegotiation. As has been pointed out, an application may use the RVBR service to find the optimal traffic specification to renegotiate.

The results of initial simulation suggest that renegotiation allows to better utilise network resources and that in protocols as RSVP, where there is no additional cost for signaling (or so we mainly assume), it is better to renegotiate. Future work on RVBR service includes either the possible integration in a real application and study on the renegotiation period, as well as the integration of the network delay and the application to Guaranteed Service [18].

6 Acknowledgements

We would like to thank Werner Almesberger for many interesting and useful suggestions. We are grateful to Olivier Verscheure, who provided the video used in the simulation, and all the DIANA partners for valuable discussions.

References

[1] DIANA Consortium, *DIANA AC319 home page*, 1998. http://www.telscom.ch/Diana.

[2] S. Giordano, J.-Y. Le Boudec, "Renegotiable Reservations," *work in progress*, 1998.

[3] S. Giordano, R. Schmid, R. Beeler, H. Flinck, J.-Y. Le Boudec, "IP, ATM - current evolution for integrated services ," Technical Report SSC/1998/2, DI-EPFL, CH-1015 Lausanne, Switzerland, January 1998.

[4] NIG-G3 Chain-Group, *NIG-G3: Internet, ATM coexistence Guideline*, 1998. http://gina.iihe.ac.be/nig-g3.

[5] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, *RFC2205: Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification.* IETF, September 1997.

[6] W. Almesberger, T. Ferrari, J.-Y. Le Boudec, "Scalable Reservation Protocol (SRP)," Technical Report 97/234, DI-EPFL, CH-1015 Lausanne, Switzerland, January 1997.

[7] K. Kilkki, J. Ruutu, *Simple Integrated Media Access (SIMA), "draft-kalevi-simple-media-access-02.txt" - work in progress.* IETF - INTERNET-DRAFT, 1998.

[8] ITU Telecommunication Standardization Sector - Study group 13, *ITU-T Recommandation Q.2963.2. Broadband integrated services digital network (B-ISDN) digital subscriber signalling system No. 2 (DSS 2) connection modification - Modification procedure for sustainable cell rate parameters*, 1998.

[9] J. Wroclawski, *RFC2211: Specification of Controlled-Load Network Element Service.* IETF, September 1997.

[10] W. Almesberger, S. Giordano, P. Cremonese, H. Flink, J. Loughney,M. Lorang, "A Framework for the QoS Based Integration of IP and ATM in the DIANA Project," *submitted to: European Telecommunications Transactions*, July 1998.

[11] The ATM Forum, *ATM User-Network Interface (UNI) Signalling Specification, Version 4.0*, 1996. ftp://ftp.atmforum.com/pub/approved-specs/af-sig-0061.000.ps.

[12] Flextel Company, *Flextel home page*, 1998. http://www.flextel.it.

[13] W. Almesberger, S. Giordano, J.-Y. Le Boudec, "Reservation Models: From Arequipa to SRP," Technical Report SSC/1998/029, DI-EPFL, CH-1015 Lausanne, Switzerland, 1998.

[14] J. Wroclawski, *RFC2210: The Use of RSVP with IETF Integrated Services.* IETF, September 1997.

[15] J.-Y. Le Boudec, "Network Calculus, Deterministic Effective Bandwidth, VBR trunks," *IEEE Globecom 97*, November 1997.

[16] S. Shenker, J. Wroclawski, *RFC2216: Network Element Service Specification Template.* IETF, September 1997.

[17] C. Fogg, "mpeg2encode/mpeg2decode," *MPEG Software Simulation Group*, 1996.

[18] S. Shenker, C. Partridge, R. Guerin, *RFC2212: Specification of Guaranteed Quality of Service.* IETF, September 1997.