

Extending the Functionality of RTP/RTCP Implementation in Network Simulator (NS-2) to support TCP friendly congestion control

Christos Bouras
Research Academic Computer
Technology Institute and University
of Patras

N. Kazantzaki Str., University of
Patras, 26500 Rion, Greece
+30 2610 960375

bouras@cti.gr

Apostolos Gkamas
Research Academic Computer
Technology Institute and University
of Patras

N. Kazantzaki Str., University of
Patras, 26500 Rion, Greece
+30 2610 960465

gkamas@cti.gr

Georgios Kioumourtzis
University of Patras, Department of
Computer Engineering and
Informatics

University of Patras, 26500, Rion,
Patras, Greece
+30 2610 960316

gkioumou@ceid.upatras.gr

ABSTRACT

In this paper, we present a modification of the ns2 code for the RTP/RTCP protocols. The legacy RTP/RTCP code in ns2 has not yet been validated but it provides a framework of the protocol's specification for experimental use. We have modified the code by adding all the RTP/RTCP protocol's attributes that are defined in RFC 3550 and related to QoS metrics. We have also implemented additional algorithms and functions in order to enhance our modified code with TCP friendly bandwidth share behavior. Our protocol, named RTPUP ("UP" stands for the University of Patras), is offered as a package and is fully documented so that it can be used for simulations and research within the ns2 simulation environment.

Categories and Subject Descriptors

I.6 [Simulation and Modeling]: Applications, Model Validation and Analysis, Model Development, Simulation Output Analysis, Miscellaneous.

General Terms

Algorithms, Performance, Design, Standardization, Verification.

Keywords

Network Simulator (NS-2), RTP/RTCP protocol, Multimedia transmission, TCP Friendly.

1. INTRODUCTION

Real time multimedia applications have enjoyed the global interest over the last years. These applications are characterized by three main properties: the demand for high data transmission rate (bandwidth-consuming applications), the sensitiveness to packet delays (latency and jitter), and last the tolerance to packet losses (packet-loss tolerant applications), when compared to other kind of applications. The Transmission Control Protocol (TCP) is the dominant and most widely used protocol at the

transport layer. However, there are three characteristics of this protocol that makes it insufficient for real time data delivery:

- TCP has a built-in retransmission mechanism that may be useless for delay-sensitive applications.
- TCP does not carry any time related information, which are needed by real time applications, and lastly.
- TCP employs a "strict" congestion control mechanism that reacts even in the light of a single packet loss event.

Similarly, the User Datagram Protocol (UDP) does not provide any support for multimedia applications. Therefore, the need of a new protocol led the research community to design the Real Time Protocol (RTP) and the associated RTP Control Protocol (RTCP) [1], in order to support multimedia applications. The RTP protocol constitutes a new de facto standard and is the dominant transport protocol for multimedia data transmission.

The implementation of the RTP in NS2 [2] is very generic. It only provides the main functions of a "common" transport protocol and runs on top of UDP. In this work, we extend the functionality of the RTP and RTCP code in NS2 to include:

- The feedback functions that are described [1] and related to QoS metrics.
- TCP friendly behavior with the meaning that the transmitted flow consumes no more bandwidth than a TCP connection, which is traversing the same path with the transmitted flow.

With these new feedback functions any multimedia application can employ the internal mechanisms of the RTP and RTCP for Quality of Service (QoS) measurements. The TCP friendly bandwidth share mechanism is based on the TCP Friendly Rate Control (TFRC) protocol presented in [3]. Our motivation is to use the RTP modified code for simulations of multimedia data transmission from a server to a number of receivers, through multicasting and different multicast RTP streams. The ns2 code provides the framework for these simulation scenarios. However, one has to extend the code to support these scenarios because the RTP code in ns2 cannot support multiple RTP streams running in one network node. The rest of this paper is organized as follows: Section 2 discusses the Algorithmic aspects. The extensions made to RTP code in ns2 are presented in section 3, as section 4

presents the performance evaluation of our modified code. Conclusions and future work are discussed in section 5.

2. ALGORITHMIC ASPECTS

In this section we describe the algorithm to estimate a TCP friendly bandwidth share. Then we explain how we estimate the packet loss ratio and the Round trip Time (RTT) that are used for the TCP friendly bandwidth calculations. Finally, we present the inter-arrival jitter delay estimations, which are based on the RFC 3550 recommendations.

2.1 TCP Friendly Bandwidth Share Estimations

The subject of transmission of TCP friendly flows over networks has engaged researchers all over the world, [4], [5] and [6]. In this work we use the analytical model of TCP presented in [4] in order to estimate a TCP friendly bandwidth share.

2.2 Packet Loss Rate Estimation

Every receiver that joins the RTP session can measure the packet loss rate based on RTP packet sequence numbers. In order to prevent a single spurious packet loss having an excessive effect on the packet loss estimation, the receivers smooth the values of packet loss rate using the filter presented in [6], which computes the weighted average of the m most recent loss rate values. The authors of [6] have also evaluated this filter and the results are very positive.

2.3 RTT Estimations

When a receiver i receives a RTP packet from a sender, it uses the following algorithm to estimate the RTT between the sender and the receiver:

$$\begin{aligned} & \text{if no feedback has been received before} \\ & \text{RTT} = \text{sqr}t(\text{effective_RTT}) \\ & \text{else} \\ & \text{RTT} = q * \text{RTT} + (1-q) * \text{effective_RTT} \end{aligned} \quad (1)$$

where, q has a default value of 0.9

2.4 Inter-arrival Jitter Estimations

Our implementation for delay jitter calculations is based on the algorithm defined in RFC 3550.

3. EXTENSIONS TO RTP CODE

In this section we describe the extensions made to RTP code in ns2. Our work is mainly divided into two main areas:

- Providing the RTP code the additional functionality defined in RFC 3550 and related to QoS metrics.
- Employing TCP friendly bandwidth share mechanisms for experimental use.

The extensions made in the ns2-2.30 version, on a Linux platform running Fedora 6 operating system.

3.1 Software Architecture

We present the structure of the RTPUP code using the UML diagram in figure 1. First of all we have renamed the RTP packet header from “*hdr_rtp*” to “*hdr_rtpup*” (“up” stands for

University of Patras) to distinguish our code from the legacy code in ns2 in order to avoid confusion within the ns2 users community. We also defined new data structures named *server_report* and *receiver_report* to store the fields of the RTCP SR and RR, respectively. A new class named *RTPUPReceiver* was declared to hold the fields that are used by the receiving Agents for QoS measurements. Every new instance of the *RTPUPSession* class creates two instances of the *RTPUPSource* and one instance of the *RTPUPReceiver* classes, accordingly. The *RTPUPSessionClass* is called by the TCL script and in turn two new Agents (*RTPUPAgent* and *RTCPUPAgent*) are assigned to every node in the network that participates in the multicast stream. The *RTPUPAgent* holds all the functionality for sending and receiving RTPUP packets, whereas the *RTCPUPAgent* is responsible for transmission and reception of the RTCPUP sender and receiver reports. We have implemented a one-to-many scheme of the RTP/RTCP protocol, in which one sender transmits a multicast stream to a set of receivers. It is however, easy and quite straightforward to extend the code so that a node can be a receiver and at the same time an active sender. This applies to VoIP applications in which the sender is also a receiver during the VoIP session. Last, new functions are also used for the implementation of the algorithms described in the previous section.

3.2 Modified and New Functions

In our RTPUP code we distinguish three major functions/modules.

3.2.1 Send and Receive RTPUP Packets

RTPUP packets are generated based on a timeout event of the *RTPUPTimer*. The RTPUP Agent creates a new RTPUP packet by calling the send function:

```
void RTPUPAgent::sendpkt() {}
```

The send function invokes the make packet function, which creates the new RTPUP packet and adds the following fields in the packet header:

```
void RTPUPAgent::makepkt(Packet* p) {}
```

- the sequence number of the RTP packet.
`rh->seqno() = seqno_++;`
- the source id of the sending source
`rh->srcid() = session_->srcid();`
- the timestamp
`rh->timestamp() = timestamp_;`
- the receivers which this sender serves with the receiver source id field and the effective RTT

```
rh->receivers_ = session_->receivers_;
```

in which the effective RTT is defined by:

$$\text{eff_rtt} = A - t_{LSR} - t_{DLSR} \quad (2)$$

where, t_{LSR} is the time during which the receiver received the last SR, t_{DLSR} is the time elapsed between the reception of the SR last report and the generation of a new RR report, and A stands for the current time of the reception of the RR. We will see later how the calculation of the effective RTT is done by the sender.

3.2.2 Build RTCPUP Sender and Receiver Report Function

The build function is called by the *RTCPUPAgent* as a result of an *RTCPUPTimer* time-out event. The sender generates a new *SR* if it has sent *RTPUP* packets since the previous *SR*. When this condition is met the sender sets the *we_sent* flag to 1 and generates a sender report (SR). Next lines present the declaration and construction of the SR:

```
//add sender report
sender_report* sr;
//fill in the report
sr = new sender_report;
//assign the sender's id
sr->sender_srcid()= localsrc_->srcid();
//assign the RTPUP packets sent
sr->pkts_sent() = localsrc_->np();
//assign the total bytes sent
sr->octets_sent() = localsrc_->nbytes();
//include the receivers served
sr->rcvr_ = receivers_;
//store the report
rh_->sr_ = sr;
```

The sender includes the total number of RTPUP packets and the total number of bytes that has sent since the beginning of the session. It also includes the receivers that this source serves.

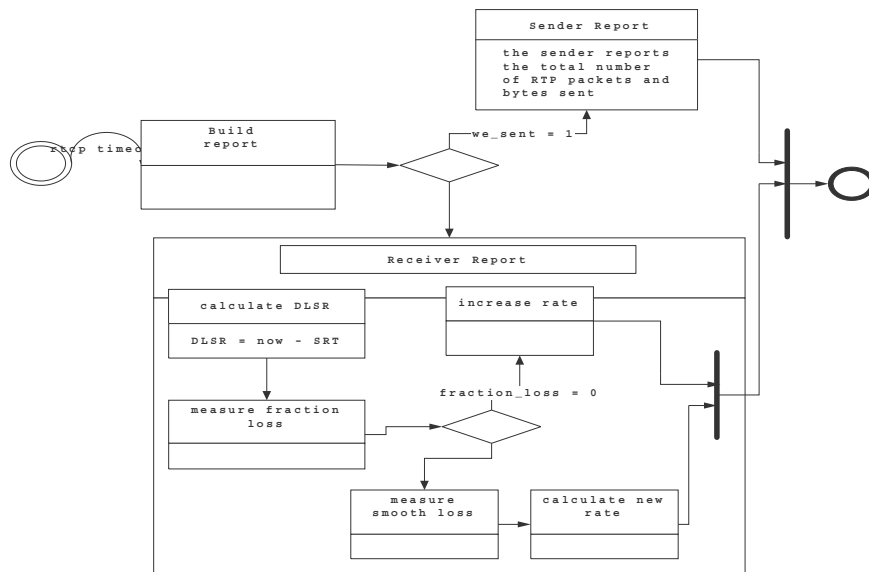


Figure 2. State chart of the build report function

3.2.3 Receive Control RTCPUP Packet Function

We have seen so far how the receivers access the RTPUP packets and how both sender and receivers build the SR and RR reports. We have also explained how the receivers perform the various calculations in order to provide the sender with QoS measurements. In this subsection we will describe what the actions are from the sender side in order to adjust its transmission rate. Therefore, the receive control function is the “merging” function in which the results of the program are presented and actions take place.

Upon the reception of a new RTCPUP report sender and receivers perform different functions. The sender first evaluates if the originator of this RR does exist in its receiver’s list. At this point it has to be mentioned that in the legacy ns2 code the *allsrcs_* field for the sending source is empty as long as it does not received any RTPUP packets from any source. That was the reason that led us to define the *RTUPReceiver* class, so that the sending source could be able to keep a list of all the receivers in

the session it serves. Therefore, if the condition is not met (the originator of the RR has not “heard” by the sender) the sender adds the originator to his receivers’ list. We use a similar function to the ns2 legacy code for constructing the receiver’s list:

```
enter_rcv(RTCPUPReceiver* s)
```

The sender processes the RR and calculates the effective RTT time as follows:

```
eff_rtt = alpha - rh->rr_->LSR() -
rh->rr_->DLSR();
where alpha is the current clock time
```

The TCP receiver’s estimation is kept in a separate data structure. We use for it an instance of the class *list* in which its size is dynamically updated with the number of its elements so that we can hold a fair large amount of receivers. In addition, the class *list* offers a number of built-in functions that are very convenient for accessing and sorting its elements. Every time the sender receives a new report from the RTCPUP Agents in the

multicast session it adjusts its transmission rate. The sender takes into account the minimum bandwidth estimations from the receiver set according to the algorithm below:

$$new_rate = \min(r_{r_tcp}^1, \dots, r_{r_tcp}^i) \quad (3)$$

where, $r_{r_tcp}^i$ is the bandwidth estimation of receiver r_r^i . At this point and in order to prevent oscillations we use a *noFeedbackTimer* to check whether or not the sender has received feedback reports from all the receivers within a feedback interval. This feedback interval is defined as:

$$feedback_interval = 2 * ps / tx_rate_$$

where ps is the packet size of the RTP packet and $tx_rate_$ is the current transmission rate. When the sender does not receive an expected RTCPUP report from a receiver within the feedback interval it cuts its sending rate to half. This is a congestion avoidance mechanism because a lost RTCPUP receiver report indicates a congested path. It has been noticed in our experimental simulations that this mechanism increases the overall performance of the protocol.

4. PERFORMANCE EVALUATION

We evaluate our model with simulations performed with the ns2 simulation software. Our main objective is first to verify that the RTPUP works properly and second that it has indeed TCP friendly TCP behavior.

4.1 Simulation Environment and Network Topology Setup

Our benchmark for the evaluation of the RTPUP protocol is a Local Area Network (LAN), which consists of one multimedia server and six heterogeneous receivers. The heterogeneity of the receivers lays in the variation of the link capacity, which connects the receivers with the LAN. We have intentionally created a “bottleneck” between routers 2 and 3 to create two different sets of wired receivers. The first set of receivers (Nodes 1, 2, 3 “fast receivers”) is able to receive at higher bit rates than the second set (Nodes 4, 5, 6 “slow receivers”). We run a simple simulation scenario in which the multimedia server transmits RTPUP traffic at an initial rate of 256Kb/s. The RTCP transmission interval is set to 500 msec. At the same time a File Transfer Protocol application (FTP) is transmitting TCP packets through the same pipe with the RTPUP traffic from Node 7 (TCP Agent) to Node 8 (TCP Sink). We run two different simulation sets to investigate:

- The behavior of our proposed protocol towards the TCP traffic
- The behavior of the TFRC implementation in ns2 towards the same TCP traffic
- Pros and cons between our implementation and the TFRC code in ns2.

Figure 3 depicts the network topology for the simulated scenarios.




Figure 3. Simulated network topology

4.2 First Simulation: Transmission of RTPUP Multicast Stream with Background TCP Traffic

We initially set the bandwidth capacity of the RTPUP traffic to 256Kb/s and the RTCPUP reporting interval to 500 msec. However, these two parameters do not remain unchangeable and adopt their values according to network conditions. As for the TCP protocol we use the standard TCP Reno version in ns2.

The transmission of both RTPUP and TCP traffic starts in the beginning of the simulation. We run our simulation for 200 seconds. In the chart presenting the simulation results (Figure 4) we can see the receiving rates from two representing nodes of the two different groups (Node 1, “fast receiver” and Node 4, “slow receiver”) and also the TCP receiving rate named as “TCP Sink”. We extract the following conclusions of the simulation results:

- The RTPUP protocol presents the characteristics of the TCP congestion control mechanism, in which the protocol increases its sending rate as long as the end-to-end path is congestion-free. This is a direct result of the TCP friendly algorithm that has been implemented in our code.
- The RTPUP protocol has also the characteristics of a multiplicative-decrease protocol, similar to TCP protocol. This is also the direct result of the implementation of the TCP analytical model in our code. However we have not implemented all the characteristics of the TCP protocol as our intention is to enrich the RTPUP with TCP friendly behavior and not to replicate the legacy TCP code.
- Another important conclusion is that our modified RTPUP protocol does have TCP friendly behavior. The TCP traffic is being delivered from the source to the destination node, although the path is heavily congested by the RTPUP traffic.
- RTPUP presents the same oscillations with the TCP protocol due to the implementation of the congestion control mechanism. We observe that when the TCP transmission rate increases, the RTUP transmission rate decreases and vice versa.
- One last important observation is that RTPUP has similar delivery ratio to both “slow” and “fast” receivers. This is a desired attribute of the protocol as it ensures a fair delivery

ratio, which most times is above 100 Kb/s. We will see in the next simulation test whether or not the TFRC implementation in ns2 is able to keep an equal delivery ratio to the whole set of receivers.




Figure 4. RTPUP versus TCP traffic




Figure 5. Delay jitter measurements

4.3 Delay Jitter Measurement

The delay jitter measurement is straightforward and is being done with a procedure that is defined in the TCL *Session/RTPUP* class. The results below (Figure 5) present the measurements of the “slow” receiver (Node 4) in contrast to the delay jitter of the “fast” receiver (Node 1). These results were measured during the previous simulation and are presented in the same chart, although the delay jitter values are in different scales. We present the results from Node 4 on the left Y-axis and the results from Node 1 on the right Y-axis. All the results are represented in seconds. We extract the following observations:

- Fast receivers enjoy minimum values of delay jitter; the highest observed delay jitter value throughout the simulation time is 2 msecs. We regard this as a good performance metric for the RTPUP protocol as the simulation scenario was set up in such way to challenge the protocol’s performance.

- Slow receivers present delay jitter values between 10 and 15 msec and in general one-way jitter up to 150 msec is considered to be acceptable even for VoIP applications.

4.4 Packet Loss Rate Measurement

For the packet loss rate measurement we have also defined a new procedure in the TCL *Session/RTPUP* class. In this way we can get directly this metric from our simulation script. We measure the loss rate as the ratio of packets lost over the packets received during the sample interval. This sample interval is the time elapsed between two consequent Receiver Reports, (RR).

$$ploss_rate = plost / prcv * 100 \quad (4)$$

We can observe from the simulation results (figure 6) that lost events occur mainly when the network is heavily congested and this happens only for a very short period. We present only the packet loss ratio from a “slow” receiver (Node 4) as we have not observed any packet losses from fast receivers. This is a desired attribute of our RTPUP implementation as we have a multicast protocol that is able to transmit at high bit rates in a congested network, with low delay jitter and minimal packet losses. In the next simulation we will see how our implementation outperforms the TFRC implementation in ns2.

4.5 Second Simulation: Comparison with the TFRC Implementation in ns2

In the last simulation we compare the TFRC implementation in ns2 against our RTP/RTCP with the TCP friendly enhancements. The TFRC code in ns2 has been used for simulation by a number of researchers and provides an acceptable implementation of the TFRC specification.

The simulation scenario has exact the same network attributes with our previous simulation in order to achieve a fair comparison. In this case, RTP traffic is transmitted to the same set of receivers and the congestion control is left to TFRC protocol. We transmit also the same TCP traffic across the network from Node 7 to Node 8. Figure 7 depicts the simulation results.




Figure 6. Loss rate measurements




Figure 7. TFRC in ns2 versus TCP traffic

- We observe from the above results that the *ns2* TFRC implementation has “smoother” oscillations than our implementation, which is a desired attribute especially for video transmission. The TCP friendly behavior is also stable except for some cases, in which TCP traffic is reduced to zero. In our implementation the TCP traffic has always equal or higher values when compared to the initial transmission rate.
- A second observation is that although Node 1 (“fast receiver”) enjoys high receiving rates, Node 4 (“slow receiver”) has very low receiving rates. However, it has to be mentioned that the TFRC code in *ns2* is used for unicast transmission. Thus, the sender transmits different unicast streams to each one of the receivers and adjusts its transmission rate accordingly.
- Our final conclusion is that our RTP/RTCP implementation introduces very good characteristics when we have multicast video stream that is transmitted via a congested path. The code and the implementation complexity of our implementation are very low when compared to the TFRC module in *ns2*.

5. CONCLUSIONS/FUTURE WORK

We present in this work an extension of the RTP code in *ns2*. Our motivation was to enrich the functionality of the existing code by including all the RTP/RTCP protocol’s specification in RFC 3550, which are related to QoS metrics. We also extended our code to enhance it with TFRC mechanisms for research and experimental use. Our effort was to keep the functions and the data fields of the original *ns2* code, to modify existing functions and to define only the necessary functions for the implementation of the new algorithms. We tried also to keep the “code style” of the *ns2*, document our code and offer it as a package for easier integration into *ns2* libraries. There were several simulation runs and tests along with those that are presented in this work in order to verify that we get the correct QoS measurements. Simulation results show that the RTPUP performance has certain advantages for multicast transmission of delay sensitive data, such as VoIP and video streaming. In our future work we will extend the RTUP code to support simultaneous RTPUP multicast streams in one node for experimental use. Finally, simulation examples,

sources and documentation are available in the following URL: http://ru6.cti.gr/ru6/ns_rtp_home.php

6. ACKNOWLEDGEMENT

We thank the anonymous SIMUTools 2008 reviewers for their helpful comments.

7. REFERENCES

- [1] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, “A Transport Protocol for Real-Time Applications” RFC 3550, July 2003
- [2] <http://www.isi.edu/nsnam/ns/>
- [3] Handley, M.; Floyd, S.; Padhye, J.; Widmer, J. “TCP Friendly Rate Control (TFRC): Protocol Specification” Request for Comments (RFC) 3448, The Internet Society, January 2003
- [4] J. Padhye, J. Kurose, D. Towsley, R. Koodli, “A model based TCP-friendly rate control protocol”, Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV), Basking Ridge, NJ, June 1999.
- [5] D. Sisalem, A. Wolisz, “MLDA: A TCP-friendly congestion control framework for heterogeneous multicast environments”, in Eighth International Workshop on Quality of Service (IWQoS 2000), Pittsburgh, PA, June 2000.
- [6] L. Vicisiano, L. Rizzo, J. Crowcroft, “TCP - like congestion control for layered multicast data transfer”, in IEEE INFOCOM, March 1998, pp. 996 - 1003.
- [7] C. Bouras, A. Gkamas, G. Kioumourtzis, “A Framework for Cross Layer Adaptation for Multimedia Transmission over Wired and Wireless Networks”, The 2007 International Conference on Internet Computing (ICOMP’07), Las Vegas, Nevada, USA, 25 - 28 June 2007.