UNIVERSITY OF SOUTHERN CALIFORNIA

INFORMATION SCIENCES INSTITUTE
4676 Admiralty Way, Marina del Rey, CA 90292

Tera-node Network Technology (TASK 3)

Scalable Personal Telecommunications

Contract: DABT63-95-C-0995
Final Technical Report
submitted by
USC Information Sciences Institute

Sponsored By:

Defense Advanced Research Projects Agency (DARPA)
Information Technology Office
Contracting Agent: Fort Huachuca

REPORT DOCUMENTATION PAGE			*FORM APPROVED* *OMB NO. 0704-0188*

1. AGENCY USE ONLY	2. REPORT DATE March 14, 2000	3. REPORT TYPE AND DATES COVERED Final Report – 7/18/1995 – 7/17/1999

4. TITLE AND SUBTITLE
Tera-node Network Technology (Task 3)
Scalable Personal Telecommunications

5. FUNDING NUMBERS
DABT63-95-C-0095

6. AUTHORS

Bannister, Joseph

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

USC INFORMATION SCIENCES INSTITUTE
4676 ADMIRALTY WAY
MARINA DEL REY, CA 90292-6695

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Directorate of Contracting
Attention: ATZS-DKO-I
Post Office Box 12748
Fort Huachuca, Arizona 85670-2748

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12A. DISTRIBUTION/AVAILABILITY STATEMENT

UNCLASSIFIED/UNLIMITED

12B. DISTRIBUTION CODE

13. ABSTRACT ***(Maximum 200 words)***

The goal of the TNT/SPT project was to develop and demonstrate technologies to support the use of personal conferencing on the Internet. The project has focused on supporting streaming applications. It has also made definitive contributions in the field of end-to-end quality-adaptive streaming applications, TCP-friendly congestion control and avoidance, and rate adaptation mechanisms.

14. SUBJECT TERMS
Multicast Teleconferencing, Streaming Multimedia, Rate and Quality Adaptation, Reliable Multicast.

15. NUMBER OF PAGES
8

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UNLIMITED

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
289-102

Scalable Personal Telecommunications (SPT)

FINAL TECHNICAL REPORT for DABT63-95-C-0995
(covering the period 4/17/1997 - 7/17/1999)

The goal of the TNT/SPT project was to develop and demonstrate technologies to support the use of personal conferencing on the Internet. The project has focused on supporting streaming applications over the Internet and made definitive contributions to the field in the following areas:

1. End-to-end Architecture for Quality-adaptive Streaming Applications over the Internet,
2. End-to-end TCP-friendly Congestion Control and Avoidance,
3. Quality Adaptation,
4. Multimedia Proxy Caching,
5. Experiments with the Rate Adaptation Protocol (RAP)
6. Providing leadership and innovation to the Internet Research Task Force (IRTF) Reliable Multicast Research Group (RMRG)

1. End-to-end Architecture for Quality-adaptive Streaming Applications over the Internet

We identified end-to-end congestion control, quality adaptation, and error control as the three major building blocks for any realtime streaming applications in the Internet. We explored the design space for each one of these components, and within that space, developed an end-to-end architecture suited for playback of layered-encoded stored video streams. The idea is to separate congestion control from error (and quality) control because the former depends on the state of the network while the latter is application specific. Our architecture not only achieves scalability but it also accommodates heterogeneity (in different aspects) among clients. Our architecture reconciles congestion control and quality adaptation which occur on different timescales. In particular, it attempts to establish "peaceful coexistence" with the Transmission Control Protocol (TCP), which is the dominant transport protocol in the Internet (and is expected to remain so for the foreseeable future). It exhibits a TCP-friendly behavior by adopting the Rate Adaptation Protocol (RAP) for end-to-end congestion control. Additionally, it uses a layered framework for quality adaptation with selective retransmission to maximize the quality of the delivered stream as available bandwidth changes.

This architecture can be generalized by replacing the suggested mechanism for each component by another from the same design space as long as all components remain compatible.

We have also devised a novel proxy caching mechanism for multimedia streams that perfectly complements the end-to-end architecture to further improve heterogeneity and scalability.

2. End-to-end TCP-friendly Congestion Control and Avoidance

We designed and developed an end-to-end TCP-friendly Rate Adaptation Protocol (RAP), which employs an additive-increase, multiplicative-decrease (AIMD) algorithm. It is well suited for unicast playback of realtime streams and other semireliable rate-based applications. Its primary goal is to be well-behaved and TCP-friendly while separating network congestion control from application-level reliability.

We studied the dynamics of TCP congestion control and avoidance in different timescales. We evaluated RAP through extensive simulation by the VINT/ns toolset and concluded that bandwidth is usually evenly shared between TCP and RAP traffic. Basic RAP behaves in a TCP-friendly fashion in a wide range of likely conditions, but we also devised a fine-grain rate adaptation mechanism to increase RAP's responsiveness to transient congestion, and to emulate the acknowledgement-clocking property of TCP.

We examined the fine-grain rate adaptation with different granularities. Our results reveal that fine-grain rate adaptation extends the range of scenarios where RAP exhibits TCP-friendly behavior. Furthermore, fine-grain rate adaptation effectively improves RAP's stability. We showed that deploying Random Early Detection (RED) queue management in routers can result in an ideal fairness between TCP and RAP traffic. We also assessed contribution of TCP's burstiness and its internal constraint on interprotocol fairness. We concluded that divergence of TCP's congestion control from AIMD is often the main cause for the unfairness against TCP traffic in special cases. Finally, we examined the self-limiting aspect of RAP and did not observe any sign of instability.

We have simulated the RAP protocol using VINT/ns-2 (Virtual Internet Testbed Network Simulator version 2) and implemented a prototype of the protocol for further experiments over CAIRN (Collaborative Advanced Interagency Research Network) and the Internet to validate of our simulation results with real data.

3. Quality Adaptation

The main challenge for congestion-controlled streaming applications is to cope with random variations in bandwidth while delivering a stream with an acceptable and stable quality. A common approach is to delay slightly the playback time and buffer some data at the client side to absorb transmission-rate variations. The more data that is initially buffered, the wider are the variations that can be absorbed, but a higher startup playback latency will be experienced by the client. The main reason that we targeted playback applications is because they can tolerate this buffering delay. Wide and random variations of the transmission rate in a long-lived session could easily result in client buffer overflow or underflow when there is a mismatch between the stream's consumption rate and long-term average bandwidth. Underflow causes interruption in playback and is very undesirable. Although buffer overflow can be resolved by deploying a flow control mechanism, it then means the received quality is less than the available bandwidth allows for. This implies that streaming applications must be *quality adaptive.*

We adopted a layered framework where the server maintains a hierarchically encoded version of each stream. As more bandwidth becomes available, more layers of the encoding are delivered. If the average bandwidth decreases, the server may then drop some of the active layers. This approach accommodates both heterogeneity and scalability. The main challenge for an layered approach to quality adaptation is to design an add-and-drop mechanism that smoothly adjusts the number of active layers without any information about the future variation in available bandwidth.

We devised an interlayer bandwidth- and buffer-sharing scheme that efficiently maximizes number of delivered layers for any loss pattern. Our scheme is generic, since we do not make any assumption about the loss pattern or available bandwidth. However, the mechanism introduces a smoothing parameter that controls the level of smoothing (i.e., the frequency of adding and dropping a layer) This scheme can be further customized based on the available information about the behavior of background traffic.

The quality adaptation mechanism was implemented and integrated into the ns-2 simulator with help from the ISI-based, DARPA-supported VINT project.

4. Multimedia Proxy Caching

Despite the success of proxy caching in the Web, proxy servers have not been used effectively for the caching of Internet multimedia streams such as audio and video. Explosive growth in demand for web-based streaming applications justifies the need for caching popular streams at a proxy server close to the interested clients. Proxy caching for multimedia streams also provides a good opportunity to support VCR-like functionalities more interactively as the control latencies to the proxy are lower.

The main challenge for proxy caching of Internet multimedia streams has to do with the need for congestion control and quality adaptation. Once a stream is cached, the proxy can replay it from the cache for subsequent requests but it still needs to perform congestion control and quality adaptation during delivery. However, using the variable-quality cached stream to perform quality adaptation for subsequent requests is problematic because there is no correlation between the variation in quality of the cached stream and the required quality for the new session.

Proxy caches perfectly complement our end-to-end architecture. We studied the implications of congestion control on proxy caching mechanisms. We devised a prefetching scheme to smooth out the variations in quality of a cached stream during subsequent playbacks. This enables the proxy to perform quality adaptation more effectively and maximizes the delivered quality. We also extended the semantics of popularity and introduced the idea of the *weighted hit* to capture both the level of interest and the usefulness of a layer for a cached stream. Finally, we presented a fine-grain replacement algorithm for layered-encoded multimedia streams at Internet proxy servers and showed that its interaction with prefetching results in the state of the cache converging to the optimal state such that the quality of a cached stream is proportional to its popularity, and the variations in quality of a cached stream are inversely proportional to its popularity. This implies that after serving several requests for a stream, the proxy can effectively hide low bandwidth paths to the original server from interested clients. Thus the delivered quality of a popular stream is not limited to the available bandwidth from the original server. Instead, the quality of each stream is determined by their popularity and the average bandwidth between the proxy and interested clients.

5. Experiments with RAP and Streaming Media

The SPT project designed and developed a protocol for providing congestion control of a real-time video stream. This Rate Adaptation Protocol (RAP) has been extensively simulated in the ns simulator, and revised based on these simulations, and we verified that it coexists well with TCP in situations likely to be experienced in real networks. RAP is intended to function in an environment where layered video is played out from a multimedia server, and depending on prevailing network conditions, layers are added and dropped as required. The assumption is that the server does not have enough resources to recode the stream live based on feedback from the network, but it can control how the layers are played out. Utilizing receiver buffering, the discrete layering can be played out at a rate determined by RAP, and receiver buffering is used to protect against unexpected congestion events.

The project team members conducted coast-to-coast experiments over CAIRN and the public Internet with RAP to evaluate the protocol's empirical performance and to validate the simulation results. Our results show that RAP's behavior is TCP-friendly, sharing bandwidth with existing Internet traffic in a fair and responsive fashion. The project also conducted simulations and experiments of the effectiveness of repair techniques in packet-based audio transmission. Techniques evaluated included repetition of previous packets during a loss gap and interpolation of missing packets. Simulation results of this work may be found in http://north.east.isi.edu/spt/audio.html.

6. Internet Research Task Force Reliable Multicast Research Group

The Internet Engineering Research Task Force (IRTF) Reliable Multicast Research Group (RMRG) has been chartered to solve some hard problems about reliable multicast transport protocols. Among the critical topics being addressed are:

1. Ways to specialize protocols to their applications given the wide range of the different requirements for reliability (from none to TCP-like byte-streams, with many shades and variations in between).

2. Solving congestion control for reliable multicast transports; understanding the interaction among RM transports, unreliable multicasts, TCP applications, and unicast non-TCP applications.

Projects members Mark Handley and Allison Mankin, both of USC/ISI, cochaired the RMRG during various points of its existence. For the December 1998 meeting, Handley coauthored (with Sally Floyd) the first draft of the congestion control specification for bulk transport reliable multicast.

Mark Handley and Allison Mankin have been actively involved in standardization work in the IETF related to scalable teleconferencing. In particular, this work relates to the initiation of multimedia sessions, multicast address allocation, and infrastructure for multicast.

7. Project Accomplishments

The SPT project has delivered several software components that implement the functionality described in this report. The RAP software runs on UNIX platforms and is available from the project web page. The NTE package is also available. Reference implementation of SDR, SAP and SIP were also produced. An integrated prototype teleconferencing platform is currently in use over the CAIRN testbed: it carries high quality audio and full rate video (at 30 frames per second) and is used in operational mode.

Quality Adaptation for Congestion Controlled Playback Video over the Internet
Reza Rejaie, Mark Handley, Deborah Estrin
To appear in Proceedings of ACM SIGCOMM '99 , Cambridge, September 1999

Architectural Considerations for Playback of Quality Adaptive Video over the Internet
Reza Rejaie, Mark Handley, Deborah Estrin
Technical report 98-686, Computer Science Department, USC.

Proxy Caching Mechanism for Multimedia Playback Streams in the Internet
Reza Rejaie, Mark Handley, Haobo Yu, Deborah Estrin
Proceedings of the 4th International Web Caching Workshop , San Diego, March 1999.

RAP: An End-to-end Rate-based Congestion Control Mechanism for Realtime Streams in the Internet
Reza Rejaie, Mark Handley, Deborah Estrin
Proceedings of IEEE INFOCOM'99 , New York, March 1999.

An End-to-End TCP-friendly Architecture for Realtime Playback Applications over the Internet
Reza Rejaie,
Ph.D. dissertation proposal, Technical report 98-681, Computer Science Department, USC, August 1998.

RTP Payload for Redundant Audio Data
C. Perkins, I. Kouvelas, R. Hardman, M. Handley, J. Bolot, G. Vega-Garcia, J. Fosse-Parisis RFC 2198, Sept 1997

NTE - A Scalable Shared Text Editor for the Mbone
M. Handley,
Proceedings of ACM Sigcomm '97, Cannes, France, September 1997.

The Internet Multimedia Conferencing Architecture
Mark Handley, Jon Crowcroft, Carsten Bormann, Joerg Ott (Internet-Draft, Sept 1997)

SDP: Session Description Protocol
Mark Handley, Van Jacobson (Internet-Draft, Sept 1997)

SAP: Session Announcement Protocol
Mark Handley (Internet-Draft, Nov 1997)

SIP: Session Initiation Protocol
Mark Handley, Henning Schulzrinne, Eve Schooler (Internet-Draft, Nov 1997)

Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification (Version 2)
D. Estrin, D. Farinacci, A. Helmy , D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma , L. Wei (Internet-Draft, Sept 1997)

The Internet Multicast Address Allocation Architecture
M. Handley, D. Thaler, D. Estrin (Dec 1998)

The Multicast Address Set Claim (MASC) Protocol
D. Estrin, M. Handley, D. Thaler, S. Kumar (Internet-Draft, Nov 1997)

The Address Allocation Protocol
M. Handley (Internet-Draft, Dec 1997)

Multicast Scope Zone Announcement Protocol
M. Handley (Internet-Draft, Dec 1997)

Guidelines for writers of RTP payload format specifications
M. Handley (Internet-Draft, Dec 1997)

8. Project Personnel

A. Mankin
M. Maher
R. Rejaie
M. Handley
D. Estrin
L. Gharai

USC
UNIVERSITY
OF SOUTHERN
CALIFORNIA

Information Sciences Institute

Divisions

Advanced Systems

Advanced Technology

Enterprise Integration Systems

High Performance Computing and Communications

Information Processing Center

Intelligent Systems

Silicon Systems (MOSIS)

Software Sciences

March 15,2000

Dr. Douglas Maughan
Program Manager
DARPA/ITO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Subject: TNT Data Submittal
Award No: DABT63-95-C-0095
CDRL No: A001

Dear Dr. Maughan:

Please find enclosed the Final Technical Report, inclusive of the required SF Form 298, for the SPT task funded under the subject contract entitled "Teranode Network Technology" (TNT).

Sincerely,

Beverly Ann Hartmeyer
PostAward Contract Administration

Enclosure

BAH

cc: Sandy Broten, Ft. Huachuca Contract Specialist, w/1 encl.
DARPA/SID, w/ 1 encl.
DTIC, w/ 2 encls.
Roy Peters, Ft. Huachuca COR, w/1 encl.
TACTEC, w/ 1 encl.
USC/DC&G, w/1 encl.
FILE: 53-4540-9192 (TNT) w/1 encl.

University of Southern California
School of Engineering
4676 Admiralty Way
Suite 1001
Marina del Rey, California 90292-6601
Tel: 310 822 1511
Fax: 310 823 6714