MULTIMEDIA COMMUNICATIONS TECHNICAL COMMITTEE MMTC

http://www.comsoc.org/socstr/org/operation/techcom/multimedia.html

IEEE COMMUNICATIONS SOCIETY

E-LETTER **Vol. 1, No. 3, December 2004**

CONTENTS

The MultiMedia communications Technical Committee (MMTC) is a volunteer group that examines systems, applications, services and techniques in which two or more media are used in the same session. These media include, but are not restricted to, voice, video, image, music, data, and executable code. The scope of the committee includes conversational, presentational, and transactional applications and the underlying networking systems to support them.

MULTIMEDIA COMMUNICATIONS

- JOIN US -

Joining MMTC is easy. Simply send an email to list@comsoc.org with the following string in the body of the message (NOT the subject line):
join multicomm

The mailing list, multicomm@comsoc.org is the communication channel with the MMTC. To post a message to the list, send e-mail to multicomm@comsoc.org.

You can also navigate through MMTC mailing list archive (since Feb. 2004).
http://barbarian.comsoc.org/comsoc.org/multicomm/

Future MMTC Meetings

CCNC 2005, January 2005, Las Vegas, Nevada, USA
Wednesday, 5 January – 07:00 – 09:00 PM.

ICC 2005, May 2005, Seul, Korea

ICC 2005 MMTC Activities

TECHNICAL SYMPOSIUM (May 16-20, 2005)
ICC 2005: Multimedia Communications and Home Networking Symposium

ICC GLOBECOM IEEE IEEE COMMUNICATIONS SOCIETY

MMTC OFFICERS

A WORD FROM THE E-LETTER E-I-C

Marco Roccetti

The main goal of this electronic publication is to disseminate issues that focus on opinions, initiatives, scientific achievements and perspectives of multimedia with an emphasis on the communication technologies.

The third issue of our E-Letter wishes to communicate to all the MMTCers the most important things happening in our community including information about the MMTC Awards, MMTC Interest Groups and related activities.

The editorial team has been putting great effort on several editorial initiatives that strengthen the reputation of the E-Letter. To this aim, it is worth mentioning that this current E-Letter issue features a new column as well as a new perspective article.

The column was provided by M. Reha Civanlar from the Computer Engineering Department of the Koç University, Istanbul, Turkey. In his column, titled "Content Adaptive Video Coding", Dr. Civanlar discusses issues related to adaptive video coding techniques based on contents.

In his perspective article, titled "Live Streaming via Application-Level Multicast in Content Delivery Networks", Mr. Jian Ni from Yale University, New Haven, USA, discusses ideas and approaches to implement application-level multicast in Content Delivery Networks for live streaming applications.

Finally, a call for contributions per annotated bibliographies has been provided by our colleague J.C. De Martin, Politecnico di Torino, Italy, which concludes this issue.

We renew the invitation to everyone to become regular contributor by submitting proposals for columns, perspective articles and annotated bibliographies. Information for submissions can be found at the MMTC website:
http://www.comsoc.org/~mmc.

Enjoy this issue!

Marco Roccetti Editor-in-Chief
E-Letter

COLUMN *M. Reha Civanlar*

Content Adaptive Video Coding

M. Reha Civanlar
Koç University, Computer Engineering Department
Sarýyer 34450, Istanbul, Turkey

Universal access to multimedia information is widely accepted to be an emerging need of the modern lifestyle. Given the heterogeneity of the network infrastructures and the end devices employed to receive multimedia content, implementing universal access requires presentation of the multimedia data in a variety of formats. These formats should be able to address a wide spectrum of applications ranging from home entertainment centers that require broadcast quality, and may be three-dimensional, video, and high fidelity, multi-channel audio, to low bandwidth, mobile devices, where a story-board type picture and audio or text presentation will be more appropriate. Proper presentation of video, yet the largest bandwidth component of multimedia data, is of vital importance in designing such systems.

Although various techniques for scalable video coding have been developed and some of them standardized, none of these can provide the wide range of media presentations required for the universal access. Most of the existing scalable video coding techniques are based on using signal processing to partition the encoded video data into streams that present the original data at various distortion levels (reduction of spatial or temporal resolutions can be considered as distortions also.) What is common to all of these techniques is that they are based on classical video compression which targets reducing the redundancy and the visual irrelevance in the video to be encoded. Achieving the wide range of scalability needed by universal access, however, is not viable by adjusting only the distortion because, particularly at the low bandwidth of the applications spectrum, the level of distortion can render the received video entirely useless. Such a wide range of scalability can only be implemented by understanding the content of the video data. That is, the visual relevancy concept should be extended to include the “semantic relevancy.”

While visual relevance can be defined independent of the context, “relevance of the content” is a highly context (content domain) dependent concept. For example, in the context of a soccer game, the temporal video segments showing a goal event and the spatial segments around the ball are definitely more important than any other part of the video, but assigning relevance levels to the parts of a video material with unspecified context e.g., a home movie, is not that obvious. On the other hand, there are a variety of important contexts, ranging from sports videos to broadcast news, where the relevance of the content can reasonably be classified. Moreover, there exist techniques for automatically classifying such content.

The target of video content analysis within this application is to partition the input video into spatio-temporal segments each of which can be assigned a semantic relevance level. The semantic relevance levels may be determined based on average choices of the common audience e.g., we may assume that a goal event is the most relevant part of a soccer game for all soccer game audience. It is also possible to tailor the relevance levels for each specific user or for each user’s specific request at a given time, considering e.g., that the same person may want to see the expressions on the players’ faces after a goal when he or she watches the same game for the second time.

Content adaptive coding ideas have been introduced in the literature before. There also are techniques that divide the input video into segments by considering various statistics along these segments that affect the ease of coding without taking into account any relevance issues. For the extremely low bandwidth accesses, several applications that generate a pictorial outline of a video signal have been reported. Content adaptive video coding in general should extend these approaches to efficient coding for video representations in all bandwidths. This way, a wide range of optimally encoded presentations can be obtained.

Given the spatio-temporal, content based segmentation of the input video and the associated relevance levels, the compression algorithm should work such that more resources are made available for more relevant segments. These resources include both the bits allocated to encoding various video segments and to network transmission resources that may include better loss protection or guaranteed quality of service. Within this approach, a scalable coder presents the most relevant data at its base layer and offers less relevant data at the higher layers. The form of presentation for the base layer may have additional levels e.g., starting from a pictorial story book and building up, as well. An explicit co-operation between the video coding and the analysis fields, made possible by the advances in the computational hardware and software, seems to be opening the path to several high impact applications in universal multimedia presentations.

About the Author

M. Reha Civanlar received his B.S., M.S. from METU, and Ph.D. from NCSU, all in EE. He worked in Pixel Machines on parallel computers and in Bell Labs on video communications. Dr. Civanlar became a Department Head in AT&T Research Labs in 1999. He joined Koç University, Istanbul, in 2002. He is also a consultant for Argela Technologies. Dr. Civanlar, an ASSP Senior Award recipient, served as an editor for IEEE Trans. on Communications, IEEE Trans. on Multimedia, EURASIP Image Comm. and JASP. He has numerous publications, standards contributions, and patents.

COSPONSORING / RELATED CONFERENCES AND WORKSHOPS

CCNC 2005

January 3 - 6, 2005
Las Vegas, Nevada, USA
IEEE Consumer Communications and Networking Conference (CCNC) will present the latest approaches and technical solutions in the areas of consumer networking, enabling technologies such as middleware and multimedia, and novel applications and services. CCNC 2005 will include a peer-reviewed program of technical sessions, technology application panels, tutorials, and poster/demo sessions.

ICC 2005

May 16 - 20, 2005
Seoul, Korea
Today, the major trend of telecommunication networks and services is "convergence" and "seamless provision". Reflecting this trend, IEEE International Conference on Communications (ICC 2005) choses "towards the era of ubiquitous networks" as the theme of ICC 2005. Under this theme, ICC 2005 will feature the latest developments in telecommunications from a technical perspective and discuss likely trends with leading technical specialists from all over the world. At the same time, influential business figures will be invited to add business flavor to ICC 2005.

1st International Symposium on Video over Wireless

June 13-15, 2005
Maui, Hawaii, USA
The 1st International Symposium on Video over Wireless will be held at Maui, Hawaii in June 13-15, 2005, in conjunction with the World's Premier International Conference on Wireless Networks, Communications, and Mobile Computing 2005 (WirelessCom'05). The symposium will focus on the state-of-the-art research in various important issues related to emerging technologies and standards on video coding and transmission over wireless networks. The symposium features keynote address, technical program, and outstanding paper awards. Outstanding papers will be invited to extend to full version for a special issue of the Journal of Wireless Communications and Mobile Computing, which is scheduled to be published in early 2007. The deadline for submitting a paper to the symposium is February 15, 2005. For more information, please visit the symposium website: http://www.ece.northwestern.edu/~haohong/wirelesscom05/index.html.

ICME 2005

July 6-8, 2005
Amsterdam, The Netherlands
IEEE International Conference on Multimedia & Expo (ICME) is a major annual international conference organized with the objective of bringing together researchers, developers and practitioners from academia and industry working in all areas of multimedia. ICME serves as a forum for the dissemination of state-of-the-art research, development, and implementations of multimedia systems, technologies, and applications.

CONFERENCE CALENDAR

CONFERENCE	LOCATION	INFORMATION
CCNC 05 IEEE Consumer Communications and Networking Conference	January 3 - 6, 2005, Las Vegas, Nevada, USA	http://www.ieee-ccnc.org/2005/
INFOCOM 05 IEEE Conference on Computer Communication	March 13 - 17, 2005, Miami, Florida, USA	http://www.ieee-infocom.org/2005/
ICC 05 International Conference on Communications	May 16 - 20, 2005, Seul, Corea	http://www.icc05.org/
VTC 05 Spring The 61st IEEE Semiannual Vehicular Technology Conference	May 29 – June 1, 2005, Stockholm, Sweden	http://ewh.ieee.org/soc/vts/conf/vtsconf.html
WirelessCom'05 International Conference on Wireless Networks, Communications, and Mobile Computing	June 13-15, 2005 Maui, Hawaii, USA	http://www.ece.northwestern.edu/~haohong/wirelesscom05/index.html
ICME 05 IEEE International Conference on Multimedia and Expo	July 6 - 8, 2005, Amsterdam, Netherlands	http://www.icme2005.org/

MMTC INTEREST GROUPS

Based on the research interests of MMTC members, five IGs have been initiated led by experts and active researchers in each area. Detailed info about the IG charters, focus areas of each IG, and their activities are announced at
http://www.comsoc.org/~mmc/
and through the reflector. The five IGs are:

(MSIG) Media Streaming
Interim Chair: Pascal Frossard
Interim Vice-chair: Juan Carlos de Martin

(HNIG) Home Networking
Interim Chair: Prof. Madjid Merabti
Interim Vice-chair: Heather Yu

(MobIG) Mobile and Wireless Multimedia
Interim Chair: Prof. R. Chandramouli
Interim Vice-chair: Oliver Wu

(SecIG) Multimedia Security
Interim Chair: Dr. Qibin Sun
Interim Vice-chair: Suba Subbalakshmi

(QoSIG) Quality of Service
Interim Chair: Qian Zhang
Interim Vice-chair: Apostolis Salkintzis

Call for IG Members:

IG Membership: We encourage you to apply for IG membership. IG Membership is free. It is a great networking opportunity. It gives means to contribute to technical activities within the multimedia communications area. Information about how to join each IG will be available at the MMTC Web site. Please stay tuned.

AWARDS

MMTC Distinguished Service Award – Given to a MMTC member with exemplary service to MMTC over a sustained period of time.

Prize
Certificate and plaque

Basis for judging
Exemplary service to MMTC over a sustained period of time

Eligibility
- The nominee must be a MMTC member at the time of nomination
- The nominee must have been a MMTC member for a sustained period of time

Winner of the 2004 ComSoc MMTC Distinguished Service Award
Dr. Charles N. Judice
For his exemplary service to the Multimedia Communications Technical Committee and the multimedia communications community at large.

MMTC Best Paper Award – Given to an outstanding paper in the area of multimedia communications published in any ComSoc magazine, journal, or ComSoc sponsored conference in the previous two calendar years.

Call for Nominations
IEEE Comsoc Multimedia Communications Technical Committee will give a yearly award to the Best Paper in the multimedia communications area. Any paper published in an IEEE Comsoc journal/magazine or in the proceedings of an IEEE Comsoc-sponsored conference/workshop/symposium, in the two years preceding the election, is eligible.

The prize is an IEEE plaque signed by ComSoc President.

Nominations are still solicited for the Best Paper Award 2004. Papers published in 2002 and 2003 will be considered. Exceptional papers published in 2001 could be nominated for the first award. Paper nominations have to be sent by email to MMTCawdcommittee@netscape.net, with subject line 'MMTC-BPA Nomination'.

The nomination should include the complete reference of the paper, author information, a brief supporting statement (maximum one page), the name of the nominator, and an electronic copy of the paper when possible. The hard deadline for paper nomination has been extended to Jan 31st, 2005.

The Best Paper Award 2004 should be presented at ICC 2005, by the MMTC chair, to one of the authors of the best paper. Additional information, and election by-laws are available on the MMTC website.

E-LETTER EDITORIAL BOARD

To all MMTC members: If your postal address, telephone or fax numbers have changed, please update them with the committee secretary. You can review our current records on our web page at http://www.comsoc.org/~mmc/.
If you like to join MMTC Mailing List, the indications how to subscribe/unsubscribe are reported at http://www.comsoc.org/~mmc/membership.html.

PERSPECTIVE ARTICLE

Jian Ni

Live Streaming via Application-Level Multicast in Content Delivery Networks

Jian Ni
Yale University, New Haven,
CT 06520, USA
Email: jian.ni@yale.edu

ABSTRACT

With the increasing popularity of multimedia streaming applications over the Internet, innovative infrastructures and technologies are needed to efficiently distribute the surging amount of multimedia contents. Content delivery networks provide an intermediate layer of infrastructure that helps to deliver the contents from content providers to a large community of geographically distributed users. This article introduces the ideas and approaches to implement application-level multicast in content delivery networks for large-scale live streaming applications.

I. INTRODUCTION

Thanks to the development of multimedia technologies and high-speed networks, traditional WWW sites are changing into multimedia WWW sites and new multimedia networking applications seem to be emerging everyday. Multimedia streaming applications like on-line movies, Internet radio and television, video conferences, distance learning, etc., have now become popular networking applications over the Internet. With the increasing popularity of such applications and the high quality-of-service (QoS) expected by the users, content delivery networks (or content distribution networks, CDNs) have recently been widely deployed to deliver the contents from content providers to a large community of geographically distributed users [1]. CDNs may be deployed by a CDN service provider such as Akamai (http://www.akamai.com) that partners with multiple ISPs. Alternatively, a big ISP like AT&T itself may provide CDN service and deploy CDN servers at the edge of its network. A CDN can achieve scalable content delivery by distributing load among its servers, by serving user requests from servers that are close to the users, and by bypassing congested network paths.

Basically, a CDN is an overlay network constructed from a group of strategically placed and geographically distributed CDN servers. These CDN servers are interconnected by the existing network infrastructure (e.g., the public Internet or transmission lines with dedicated bandwidth). They form a logical overlay network and function as overlay routers to deliver the contents from the origin servers to the users. When a user requests some content, the request is redirected to a nearby CDN server (we call it the user's local CDN server) through certain request redirection mechanism so that content delivery takes place at the edge of the network where bandwidth is abundant. The local CDN server is the entry point for the user to access the CDN. It conducts admission control to either accept or block the request. If the request is accepted, the local CDN server will serve the user if it has the content, otherwise it will perform content routing to locate and then deliver the requested content to the user.

In this article we consider a large class of multimedia streaming applications: streaming of live multimedia contents (live streaming). This class of applications is similar to traditional broadcast radio and television, except that transmission takes place over the Internet. These applications often have many users who are receiving the same live multimedia content synchronously, which can be efficiently accomplished via multicast techniques. Multicast over the Internet was originally proposed at the network layer, referred to as IP multicast. However, after a decade of research, there are still many hurdles in the deployment of IP multicast, such as the lack of higher layer functionalities (e.g., reliable transport; flow, congestion and access control; multicast security, etc.) and scalable inter-domain multicast routing protocols. Therefore, application-level multicast (ALM), or overlay multicast, in which multicast functionality is pushed up to the application layer, has recently been proposed for a number of Internet multicast applications [2], [3], [4].

There are two different ALM architectures, namely, the peer-to-peer (P2P) based architecture and the proxy-based architecture. In the P2P based architecture, the users are considered to be equivalent peers and are organized into an overlay network for multicast data delivery. This architecture is suitable for applications like video conferences or interactive games. In the proxy-based architecture, a group of proxy servers (or multicast service nodes) are organized into an overlay network and provide content delivery service to the users. This architecture is suitable for applications like live streaming. In this article we concentrate on the proxy-based architecture for implementing ALM in CDNs.

A CDN server can aggregate the requests for the same live multimedia content that are redirected to it, and become a multicast group member for that content. A multicast group for a multicast content is a group established for the distribution of that content. The multicast group members for the same content form an overlay multicast tree that delivers the content from the origin server to all multicast group members. If a CDN server that has not yet been a multicast group member for certain content receives a user's request for that content, it initiates multicast (content) routing to add itself to the multicast tree.

We give a brief comparison between IP multicast and ALM when being deployed in CDNs.

1. ***Multicast Group Identifier***. In IP multicast, a Class D address is used to identify a multicast group. In ALM, a URL or other application related key can be used to identify a multicast group.
2. ***Multicast Group Members***. In IP multicast, in order to receive certain multicast content, a user explicitly registers to its directly attached router via the Internet Group Management Protocol (IGMP). The IP routers of the same multicast group form a multicast tree and take responsibility to deliver the multicast content to their registered users. In ALM, the request from a user is redirect to the user's local CDN server. The CDN servers of the same multicast group form a multicast tree or mesh and take responsibility to deliver the multicast content to their users.
3. ***Network Topology***. In IP multicast, the topology of the IP routers exactly reflects the physical network topology. While in ALM, the CDN servers form a logical overlay network on top of the underneath network infrastructure. This overlay network may or may not accomplish a good match with the real physical network.
4. ***Multicast Routing***. Both IP multicast and ALM require certain multicast routing protocol to construct a multicast tree or mesh for delivering the multicast content. IP multicast routing normally relies on the underlying unicast routing protocols that are quasi-static and employ simple routing metrics like the number of hops or delay. On the contrary, as we will discuss in the next section, the implementation of ALM routing is much more flexible.

II. APPLICATION-LEVEL MULTICAST ROUTING

When constructing ALM trees for live streaming applications, first we can choose different routing metrics. One commonly employed metric is delay. Note that ALM for live streaming is unidirectional transmission: a CDN server who is transmitting a live content to another CDN server will not receive the same content back from that CDN server. Hence for each multicast group, a single-source tree is formed where the source is, for example, the origin server. In this case we want to construct a minimum-delay-path spanning tree that minimizes the delay from the source to each of the nodes on the tree.

For live streaming applications that consume a large amount of bandwidth, the available bandwidth of the network path between the sender and the receiver is crucial to maintain the streaming quality. In addition, end-to-end delay normally consists of propagation delay, transmission delay and queueing delay. Once the end-to-end bandwidth is guaranteed, transmission delay and queueing delay can be bounded, only the propagation delay plays an important role. While for streaming applications, a user can tolerate some start-up (propagation) delay to receive the content, but will be annoyed if the streaming is interrupted because of insufficient available bandwidth. Therefore, available bandwidth may be a more suitable routing metric. In this case we want to construct a widest-path spanning tree that maximizes the minimum available bandwidth of the tree links. Here accurately measuring the available bandwidth between the CDN servers without causing intrusive traffic is a challenging task.

[3] and [4] choose the access (interface) bandwidth of the CDN servers as the routing metric. This is based on the assumption that links in the core networks are over-provisioned and therefore are not bottlenecks, while the access bandwidth limits the number of simultaneous multicast sessions that a CDN server can support. In order to avoid overusing a CDN server's access bandwidth, its outgoing degree on the multicast tree should be bounded. In addition to access bandwidth, [3] and [4] also consider delay. For example, [3] considers the minimum-diameter degree-limited spanning tree problem (MDDL), in which the diameter of the tree is defined as the delay of the longest tree path. [4] considers the minimum-average-latency degree-bounded directed spanning tree problem, in which the average latency of the tree is defined as the weighted summation of the latencies of all overlay nodes. The weight of an overlay node is determined by its user population. In general, when multiple routing metrics are considered simultaneously, the idea is similar to that of quality-of-service (QoS) routing, but now this is implemented at the application layer.

Secondly, we can employ different routing strategies.

1. ***Static vs. Dynamic***. In the static or quasi-static approach, a static multicast tree that spans all the CDN servers of a CDN is built beforehand based on the long-term estimation of the network conditions. Every time when a new multicast group is formed, a subtree of the static multicast tree that spans all the multicast group members is used to deliver the content. In the dynamic approach, every time when a new multicast group is formed, a multicast tree is constructed on-demand among the multicast group members based on the current network conditions. Compared with the static approach, the dynamic approach is more adaptive and can achieve better resource utilization, but requires heavier communication and control overhead.
2. ***Centralized vs. Distributed***. In the centralized approach, a central manager/server constructs the multicast tree for every multicast group. It collects the link information (available bandwidth, delay, etc.) and server information (CPU power, residual access bandwidth, etc.) of the entire CDN, and makes decisions based on this global information. In the distributed approach, each CDN server collects the local information of its directly attached links and neighboring CDN servers. The multicast tree is constructed distributively among the multicast group members. The centralized approach can achieve global optimality and works well for small CDNs, while the distributed approach is more robust and scales much better for large CDNs.
3. ***Reconstructive vs. Cumulative***. During the session time of a live streaming, a CDN server may join or leave the

multicast group dynamically, depending on whether there are some users receiving the live content from that CDN server. In the reconstructive approach, every time when a CDN server joins or leaves the multicast group, a new multicast tree is reconstructed. Note that this may change some of the on-going transmissions between CDN servers that still remain in the multicast group. In the cumulative approach, when a CDN server joins the multicast group, a unicast path that connects this CDN server and a node on the multicast tree is added to the multicast tree. When a CDN server leaves the multicast group, if it has no child node on the multicast tree, then it just detaches itself from the tree; otherwise, it remains on the multicast tree until all its child nodes leave the multicast group. Note that in the cumulative approach no on-going transmission between a pair of CDN servers will be interrupted. The reconstructive approach can maintain optimality when changes occur, but it requires smooth switch-over techniques to change on-going transmissions.

III. HIERARCHICAL APPLICATION-LEVEL MULTICAST

It is difficult to construct an optimal or even efficient ALM tree for a multicast group consisting of hundreds or even thousands of CDN servers. This may happen, for example, when Internet broadcasting the World Cup or Olympic Games to users all around the world. We believe that a hierarchical ALM approach is needed for such large-scale live streaming. In order to form a hierarchical overlay network [5], the CDN servers are first grouped into clusters, either by manual configuration, or through some self-organizing scheme. The CDN servers of the same cluster then form a sub-overlay network. Each cluster selects one or more CDN server as the representative of this cluster. Representatives of different clusters form a higher-level overlay network. Thus a hierarchical topology is constructed.

We give an example to illustrate that a hierarchical approach scales much better than the flat approach (in which all CDN servers are in the same level). Assume there are N=1000 CDN servers participating in a large-scale live streaming. Assume the computational complexity of constructing a N-node multicast tree that optimizes certain routing metrics is $O(N^3)$ (e.g., the computational complexity of constructing a minimum-diameter spanning tree is $O(N^3)$). In the flat approach, with the reconstructive routing strategy, every time when a CDN server joins or leaves the multicast group, a new optimal multicast tree is reconstructed, with a computational complexity of $O(10^9)$. While in a three-level hierarchical approach, a cluster at each level has O(10) CDN servers. With the reconstructive routing strategy, when a CDN server joins or leaves the multicast group, under the worst case, a new optimal multicast tree is reconstructed in one cluster at each of the three levels. This only requires a computational complexity of $O(3*10^3)$. Note that the hierarchical approach may not achieve global optimality, but scales much better.

A hierarchical ALM approach is also flexible to implement. At different level of the hierarchical overlay network, we can employ the most appropriate routing strategy and algorithm to construct the ALM trees based on property of that level. For example, at the lowest level, for the CDN servers in the same cluster that are close to each other, we can employ the centralized and reconstructive routing strategy to achieve optimality within that cluster. At the higher level, for the CDN servers that are geographically distributed, we can employ the distributed and cumulative routing strategy in which each CDN server only shares information with its neighbors, and ALM routing is implemented in an ad-hoc distributive way.

IV. SUMMARY

In this article we first introduced the idea of implementing ALM in CDNs for live streaming applications. We then discussed different approaches to construct ALM trees by employing different routing metrics and routing strategies. Finally we demonstrated that a hierarchical ALM approach scales much better than the flat approach to achieve large-scale live streaming. The issues introduced and discussed in this article will help researchers to design practical and efficient ALM protocols in CDNs for large-scale live streaming applications, a task we are currently exploring.

REFERENCES

[1] A. Vakali, G. Pallis, "Content Delivery Networks: Status and Trends," IEEE Internet Computing, Vol 7, Issue 6, pp. 68-74, 2003.

[2] D. Pendarakis, S. Shi, D. Verma, and M. Waldvogel, "ALMI: An Application Level Multicast Infrastructure," Proc. 3rd Usenix Symposium on Internet Technologies and systems (USITS), March 2001.

[3] S. Shi and J. S. Turner, "Routing in Overlay Multicast Networks," Proc. IEEE INFOCOM 2002, June 2002.

[4] S. Banerjee, C. Kommareddy, K. Kar, B. Bhattacharjee, and S. Khuller, "Construction of an Efficient Overlay Multicast Infrastructure for Real-time Applications,"' Proc. IEEE INFOCOM 2003, April 2003.

[5] J. Ni, D. H. K. Tsang, S. H. Yeung, X. Hei, "Hierarchical Content Routing in Large-Scale Multimedia Content Delivery Network," Proc. IEEE ICC 2003, May 2003.

About the Author:

Jian Ni received his B.Eng. degree in Automation from Tsinghua University, Beijing, China, in 2001, and his M.Phil. degree in Electrical and Electronic Engineering from The Hong Kong University of Science and Technology, Hong Kong, in 2003. He is currently working toward his Ph.D. degree in Electrical Engineering at Yale University, New Haven, USA.

CALL FOR CONTRIBUTIONS

Juan Carlos De Martin

Call for Contributions per Annotated Bibliographies for

The Multimedia Communications Technical Committee

E-Letter

Editor in Chief: Marco Roccetti

IEEE Communications Society

The E-letter of the Multimedia Communications Technical Committee of the IEEE Communications Society is an electronic publication that welcomes submissions of annotated bibliographies.

A considerable barrier to entry into a new field of research is to become aware of the existing literature on the topic. The Internet and search engines -such as IEEExplore and, more recently, Google Scholar- have made access conference proceedings and journals immensely easier than it used to be.

However, speed and ease of access, by themselves, do not solve the problem of understanding the state of the art in a given field. Some form of intelligence is needed to filter the raw data represented by the very large number of available publications. Such intelligence may be acquired, in due time, by reading and attending conferences - or it may come from experts already working in the field.

To help fellow engineers and researchers to gain easier access to new fields of activities, the E-Letter of the Multimedia Communications Technical Committee (MMTC) invites multimedia experts to submit annotated bibliographies on topics of their choosing.

It is expected that the annotated bibliographies could be of various kinds – from tutorial level bibliographies on the general field of multimedia communications to bibliographies on very specialized subtopics.

If technically feasible, we will adopt an open approach to bibliographies development. Instruments such as wiki are, in fact, making very easy to build knowledge repositories in a collaborative fashion, as shown, for instance, by the astounding success of wikipedia.org. Initial contributions could, therefore, if the original author agrees, be placed on a MMTC wiki to be integrated by comments and modifications made by the community at large. The E-letter will periodically publish selected annotated bibliographies.

Possible topics for annotated bibliographies include, but are not limited to:

Hardware and Software for Multimedia
Home Networking for Multimedia
Implemented Prototypes
Mathematical Modeling and Simulation for Multimedia
Mobile and Wireless multimedia
Multimedia Communication Systems
Multimedia Security
Multimedia Design
Multimedia Development Tools
Multimedia Networking and Quality of Service
Networked Multimedia Entertainment
Quantitative and Qualitative Studies for Multimedia
Streaming Multimedia
Theoretical/Ergonomic Issues Regarding Multimedia Communications

Annotated bibliographies will be subject to peer review and, upon acceptance, published in an upcoming issue of the E-Letter. All authors should consider the general nature of the E-Letter's readers. Annotated bibliographies should not have been previously published and must not be submitted for publication as well.

Submission guidelines are as follows: length should be no more than 3000 words (four double column pages).

Annotated bibliographies should be submitted in pdf format by e-mail to the E-Letter Assistant Editor J.C. De Martin at demartin@polito.it.

Deadlines:
The next issue of the E-Letter will appear on April 2005. Our deadline for receiving annotated bibliographies articles is 60 days prior to the cover date.