

Efficient Group Coordination in Multicast Trees

H.-P. DOMMEL

Computer Engineering Department
Santa Clara University
Santa Clara, CA 95053, USA
(hpdommel@cse.scu.edu)

and J. J. GARCIA-LUNA-ACEVES

Computer Engineering Department
School of Engineering
University of California
Santa Cruz, CA 95064, USA
(jj@cse.ucsc.edu)

Abstract. The majority of today's Internet applications relies on point-to-point communication. In recent years, however, multipoint communication support has become the foundation for such applications as multiparty video conferencing, distributed interactive simulations, and collaborative systems. We describe a novel protocol to coordinate multipoint groupwork within the IP-multicast framework. The protocol supports Internet-wide coordination for large and highly-interactive groupwork, relying on the dissemination of coordination directives among group members across a shared end-to-end multicast tree. We also describe how addressing extensions to IP multicast can be used for our multisite coordination mechanism.

Keywords: group coordination, tree-based multicast, process coordination, multiparty collaboration

1. Introduction

Internet-based computing is gradually migrating from the standard unicast transmission model to multicasting. In the IP multicasting model [2], a source needs to send a packet only once to the network interface, and multicast routers replicate the packet on its transmission path to multiple receivers. The Internet Group Management Protocol (IGMP) allows a host to join a multicast group by informing its local router to forward multicast traffic for this group to the leaf subnetwork where the host resides. Protocols such as DVMRP, MOSPF, PIM-SM/DM, MBGP, and MSDP [2] perform the construction of intradomain and interdomain multicast delivery trees and enable packet forwarding among routers. However, the standard IP multicast model provides no guarantees for reliable or order-preserving delivery of packets, and every message is delivered to all members of a multicast group on a best-effort basis. A variety of protocols for multicast routing and



© 2001 Kluwer Academic Publishers. Printed in the Netherlands.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 2001		2. REPORT TYPE		3. DATES COVERED 00-00-2001 to 00-00-2001	
4. TITLE AND SUBTITLE Efficient Group Coordination in Multicast Trees				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California at Santa Cruz, Department of Computer Engineering, Santa Cruz, CA, 95064				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 14	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

reliable multicasting have been proposed in recent years as solutions to the numerous inherent architectural and deployment issues [4]. In this article we investigate supporting coordination among users and processes taking advantage of IP multicasting. We view group coordination among nodes in the IP multicast service model as complementary to group membership and information dissemination within the group, which are handled by the Internet Group Membership Protocol (IGMP) and multicast routing protocols (e.g., PIM, CBT, and the like), respectively. As the following sections describe, any distributed applications exerting multisite control and building on coordination of events and activities may benefit from tapping into group coordination services.

The problem of coordination of processes and their resources in network-centric computing has received considerable attention in recent years. In our context, group coordination services entail activity coordination among interacting hosts, including flow synchronization from different sources, ordered delivery of distributed event information, and the concurrent use of and access to shared resources, also referred to as floor control [5].

The origins of group coordination can be traced back to protection and access control mechanisms in operating systems [11], mutual exclusion in distributed systems [15], and concurrency control in database systems [3]. None of these mechanisms caters to continuous multimedia content or host interactivity. Studies in the context of computer supported cooperative work [6, 9] have shown that social aspects such as lack of eye contact as well as technical factors, for example limited bandwidth and network delays, contribute to coordination problems in network-mediated collaborative work. These problems can materialize in greedy user behavior, contention for resource usage, and a lack of group focus in joint task completion. Many proposed mechanisms and systems for online collaboration [1, 8, 17–19] are proprietary, sparsely documented, and limited both in scope and scale to local area networks and few users.

We address the problem of how to allow group-centric collaborative applications take advantage of the underlying multicast routing service and node geometry to coordinate interactions among the group members efficiently, without restrictions in scope and scale. To this end, we describe a novel tree-based group coordination mechanism called the Aggregated Coordination Protocol (ACP). ACP operates on a shared multicast tree as a shadow of the underlying multicast routing tree to store and forward coordination primitives among hosts. Multiple parties in different multicast on-tree groups are supported by ACP to coordinate their distributed activities via message passing. ACP can be used for floor control by establishing ephemeral resource usage rights

through coordination primitives, rather than operating system locks, and its semantics extends to continuous media flows as well as discrete data. Raymond [15] described an ancestor protocol of ACP, where a mutual exclusion algorithm operates on a static logical propagation tree to determine the node gaining exclusive access to a critical section.

The rest of this paper is organized as follows: Section 2 presents the system model and assumptions. Section 3 describes ACP and addresses its correctness, robustness and performance. Section 4 summarizes the benefits of tree-based group coordination.

2. Model and Definitions

We define a coordination session $C_S = (P, H, L)$ in a computer network to consist of processes P coordinating a set of interactive hosts H over links $L \subset H \times H$. The entities involved in a coordination process are users (agents) and resources at the end-hosts, and the ACP packets coordinating them flowing across links. Communication is by message passing only, and we assume that messages eventually arrive correctly. We refer to ACP packets as control primitives (CPs) and each host in a session is client and server for CPs to other hosts. CPs synchronize the hosts' joint tasks, implement causal or total ordering in distributed events, and mitigate access to shared, but exclusive resources. Coordination management must be aligned with membership operations such as joining, leaving or splitting groups.

We will illustrate group coordination with floor control, which entails mechanisms and policies to mitigate mutually exclusive access on a set of shared resources as the coordination objective. Resources can be located at a specific end host (camera), in replicated form at every host of a multicast group handling the same resource (telepointer), or in the network (voice channel). We distinguish between generic CPs (“grant”, “release”, “open”) and resource or media-specific CPs (“rotate left”, “zoom in”). The temporary privilege to work with a multimedia resource is also called *floor*. CPs are exchanged among hosts to resolve contention for a specific resource floor, and to preserve a global activity order. CPs transfer a sender id, the receiver ids, the time of creation, a control directive, the allowed duration, and an optional priority value. Users may assume social roles (moderator, panelist, student etc.), or control roles, which include the *floor controller* (FC) of a resource R as the host managing access and interactions for R , and the *floor holder* (FH) as the host temporarily permitted to use the resource.

Control over shared resources can be centralized, distributed or a hybrid of both. It can be performed successively by individual session

members, partitioned, where various session hosts contribute different control functions, or democratic, where consensus is achieved by negotiation, yielding a new consistent control state. A host holds a floor in each turn for a time interval T , which can be preset or timed out. In our model, we assume that the FC role is associated with a specific host, but it may infrequently rove among hosts. The FH changes at every turn. We assume that each host in a session runs the same coordination protocol, serving requests from other hosts and transmitting requests for resources placed by users.

3. Aggregated Coordination Protocol

3.1. DESCRIPTION

The Aggregated Coordination Protocol (ACP) operates on a control tree, consisting of three main types of nodes: *holder nodes* host the FH, operating on a resource and being transmission sources; *control nodes* host the FC for a specific resource, and are addressed by other nodes asking for a floor; and *target nodes* receive updates of the operations by a FH. Nodes on the path from a holder to its targets are referred to as hop nodes. Leaf nodes terminate the downward forwarding of control information in the tree. ACP packets are assumed to be transferred reliably, which is implicitly guaranteed by the underlying reliable multicast protocol. If ACP performs independently, it needs to supply its own reliability mechanisms.

Each CP contains a version field V ; a flag field for protection bits; a TTL field to scope the CP; a sequence number field SN and timestamp TS to order CPs and uniquely identify them in the shared event space; an identifier $CPid$ to uniquely tag the CP in the session; a checksum CS ; a command field Cmd for specific directives; a type field Typ to classify the operation; an options field Opt for the command; and the source address SA and target address TA , which can be a multicast group address. The 24 bits for CPids are sufficient to characterize fine-grained resource coordination in large sessions and the 16-bit prefix labeling used for SA and TA is sufficient for sessions up to 64K hosts. Figure 1 depicts this structure of a CP packet.

CPs are disseminated across a single shared tree connecting all hosts in a session. The tree corresponds to a single shared acknowledgment tree for concurrent multicasting, allowing multiple sources instead of building separate dissemination trees per source and multicast group. We outline tree maintenance for the case that group coordination is deployed in a network in conjunction with an existing underlying reliable multicast tree. The coordination tree can be a working copy of




Figure 1. Packet header fields for coordination primitives (CPs).

the reliable multicast tree prepared by a protocol such as Lorax [13]. Various mechanisms for initiation, joining, and advertising of collaborative sessions exist. We assume that one host, representing itself or a multicast group, initiates a session and advertises the session description, multicast address and media used in a session directory [7]. The directory serves as a rendezvous interface, which allows other hosts to join via a call-up mechanism.

The control tree is grown from the initiating node as the root, and other hosts joining the session are considered first off-tree, unicasting request-to-join messages to the inviting root node based on addressing information provided by the session directory. A TTL (time-to-live) field in the join packet restricts the session scope. A successful adoption of a child node to the control tree is confirmed with a bind message. Each newly joined host, as the root of its subtree, locally advertises invitation-to-join messages. It retrieves the current control state for resources of concern by polling its parent node. It may also be the case that separate subtrees may fission together to create a joint control tree. Each node in the tree has a maximum degree D , which must be high enough to reflect the session structure, but not exceed the capacity of a host to efficiently serve its children. Furthermore, if the tree also serves as a media mixing hierarchy [14], the permissible height must satisfy the end-to-end delay constraints.

Open sessions with dynamic membership may incur frequent joining and leaving, or accidental withdrawal of hosts from the control tree. When the root leaves the control tree, the eldest child in the subtree becomes the designated new root. Age may be determined by location, joining time, or address labels. Although the initial source and root will change during the lifetime of a session, the branching properties of the positional tree will not change, when the tree is virtually rehung with shifting of control roles. A hop node, which lost contact with its parent for a timeout larger than possible congestion delays, must rejoin the tree as described. CPs from a node identified as lost are held at the

processing host until a timeout, and are discarded if the host does not reappear. A lost and rejoined host must resend its pending CPs.

Routing of CPs in the tree is performed as follows: if a target node (FC or FH) is in the subtree of a node, the CP is routed downward the subtree branch where the target resides, otherwise it is sent upward to its parent node. The exact location of the target node is not needed since CPs are forwarded on the basis of directionality. This approach is similar in principle to Raymond's mutual exclusion solution [15], where a logical control tree is used to direct critical section requests for one resource at a time towards the node determining critical section entry. In contrast, our protocol handles multiple resources and routes are chosen based on the underlying actual multicast tree. Due to a higher degree of interactivity, where many nodes may assume FC or FH roles during session time, we introduce a more expressive addressing mechanism with self-routing semantics based on prefix matching of node labels.

The idea of using binary labels has first been used for routing of instructions in multiprocessor systems. In networking, it has been used for IP address lookup mechanisms in CIDR [16], promoting hierarchical address aggregation in small forwarding tables using prefix matching, and for router acceleration in multicast routing and reliable multicasting [12]. Prefix matching is also an elegant solution for subcasting and anycasting.¹

ACP assigns unique prefix labels recursively and top-down to each node joining the control tree, such that a child node label contains the label of its parent as a prefix. For example, using a binary alphabet, the root receives label **1**, its children are numbered **10** and **11**, etc. Labels allow nodes to be addressed individually, even when belonging to more than one potentially overlapping multicast groups. Also, hosts can be placed in the tree independent of their membership in different multicast groups, in contrast to other approaches, which demand that the tree organization closely reflects group membership [10].

CPs such as requests for the floor on a resource are propagated across the tree using aggregation. This mechanism corresponds to solutions for the Ack implosion problem in reliable multicast, reducing the feedback of receivers to a source on lost or corrupt packets [13]. Aggregation constrains the control process to local groups, i.e., CPs with requests that can be satisfied within the neighborhood of the issuing node, will not have to be forwarded all the way to the target controller node. The controller node is hence liberated from the need to communicate with

¹ Subcasting delivers or retrieves data between a source and select members of a multicast group, and anycasting transfers data to any one member of a group, for example the nearest proxy in a group of servers.

every host in the session, and deals only with requests which could not be satisfied along the feedback path. If a hop node receives the same CP from different nodes within a timeout, it aggregates them into one CP. It checks then if a response to these requests can be satisfied locally by polling its own state and the state of neighboring nodes. Otherwise the composite CP is self-routed up or down in the tree toward the target node(s) as outlined before. This implies that the number of CPs required to coordinate nodes decreases as the request activity increases, because requests are not sent further, if a hop node is reached that already processed the CP.

Hosts need to maintain the following state locally: the resource ids shared locally with remote hosts, and the remote resource ids accessed locally, together with their CPids; a state table indicating which resources are locally available or held remotely, together with the id of the remote FH; and a request queue *ReqQ* which collects successive CPs from different hosts (the queue is limited by the number of nodes in the session). In addition, each node maintains a FIFO queue of pending CPids, identified by the senders' labels. A hop node receiving a CP compares the label of the target node with the head of the queue, and elects itself as receiver if its own label matches the head, or forwards it according to the routing procedure outlined previously. A control node receiving a request responds immediately to the request by sending back a response to the requester, if its local queue is empty, or it appends it to its queue. When control shifts from one node to the next, the pending request queue is transferred to the new control node and its new address is multicast to all groups sharing the associated resource.

3.2. EXAMPLE

Consider a scenario where three hosts from three different multicast groups MG1 - MG3 must coordinate access to a shared resource they are contending for. Figure 2 depicts a snapshot of the protocol operation in a ternary tree.

Assume that hosts **12**, **100** and **11** all request the floor held by FC at location **101**. All request packets need to be routed along branches of the tree to **101**. The prefix property of the labels allows self-routing of these packets. For example, assume that all hosts are informed that host **101** is FC. Host **12** compares its label with the target label. Its prefix matches (**1**), but the second identifier indicates that the FC is on subtree 0. The request packet is hence sent upward to host **1**, which compares its label with the target, and, detecting that **101** is one of its children, sends the packet to host **10**, whose label matches the prefix of **101**. This node performs the same comparison and the

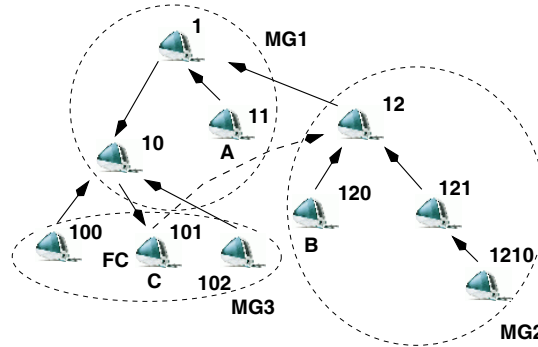


Figure 2. Snapshot of ACP operation among multicast groups MG1 - MG3.

packet ultimately arrives at FC, which finally grants the floor to node 12.

3.3. CORRECTNESS

We outline correctness and resilience properties of ACP. Assuming that CPs are sent reliably, a node can receive the privilege associated with a CP only if sends a request to the node controlling the associated resource. A message enacting the CP (e.g., grant-floor) can only be sent by a controlling node. The addressing labels in the tree are unique and a CP is only issued for a specific requesting node. Provided that *a priori* there is a single privileged FH per resource, we need to show that ACP will continue to forward and interpret CPs such that the single FH property per resource is maintained.

If several nodes send a CP and another node holds the privilege associated with it for an infinite time, a deadlock exists. Reasons may be that the wrong node or no node holds the privilege, it is unreachable, or the propagation of a request fails prematurely. Assume a tree with N hosts. In the worst case, each CP has to traverse $2N$ hops for a node to get its CP request satisfied, where $2N$ is the length of the longest path in the tree. Let us now assume that ACP operates correctly for $N + 1$ nodes in a tree of height $H = 2$ with an additional root node. Using induction over H , we assume that the liveness argument holds for any level L with $2 \leq L < H$. A tree of height $H + 1$ has one additional level, from where CPs can be sent and must be replied to. Assume that a node in the additional level of the tree (the root, or a leaf level of a subtree) sends a CP. The prefix matching semantics of the protocol ensures that messages are also correctly self-routed to the target address at this tree level and will eventually reach the target node. The acyclic tree geometry prevents that a CP request can be outrun by a moving CP response (e.g., a grant message.) Consequently, a host with

a pending CP request cannot starve indefinitely and ultimately gets its CP satisfied. If we consider buffering at nodes, the finite FIFO nature of request queues at each node guarantees that every CP will eventually be served. Thus, ACP operates live and correct at any tree height.

3.4. FAIRNESS AND RESILIENCE

Fairness refers to the average frequency and duration, by which nodes acquire a privilege for a given period. This period is for example the session lifetime, or the lifetime of a shared resource in a session. Network latency, geographic distance and location of nodes, or varying host capabilities can all be factors in causing uneven dissemination patterns for CPs and consequently unfair allocation of resource access privileges. Leaf nodes take more time to propagate their requests across the root to a node on the other side of a tree, than nodes just below the root. Shared trees also do not provide shortest paths between a source and its receiver set [2], which may cause increased latency in CP transfer. It is hence important to establish service policies, which counteract these factors. One simple solution, a “least-recently-served” policy, can be enacted by letting each node maintain a local record of the most recent CPs and their originating nodes. Those nodes are serviced first, which do not appear on the list, or appear least in duration or frequency of service.

Previously, we assumed that transfer of CPs and accounting of control information among nodes is failure free. Even we if assume reliable multicasting to ensure that CPs are eventually transferred, the control apparatus may need additional recovery mechanisms to ensure consistency. This applies to regular node failure, control node failure, link failure, or control token loss or duplication. Such exceptions can be avoided by disseminating status information among active nodes, or by systematic detection of loss and recovery. Regular node or controller failures are typically detected via timeout and recovered with an election protocol, with neighbor nodes providing state updates. Continuation of a split session is possible if a quorum of the members in each partition agrees to continue. One method to deal with the case that a CP is lost completely or reaches only a subset of nodes is to multicast a `CP_probe` message from a control node to the group.

3.5. PERFORMANCE

Attaching positional labels to nodes in a D -ary tree implies an additional storage cost of $\log_2 D$ bits per level in a positional tree of N receivers and height $\log_D N$, i.e., $\lg N$ bits are needed. With 16-bit labels for designating sources and targets in message headers, up

to 2^{16} hosts can be accommodated. Prefix comparison is cheaper for nodes close to the root due to shorter labels. Serving a CP costs $C_{CP} = c_{req} + c_{resp} + c_{upd}$, comprising the cost to send a request to a control node, receive a response, and multicast an update of the new state. The response time t_r to receive a response is bound by the maximum time for the probe to traverse the longest link, t_{max} , plus the time t_{ack} for the receiver nodes to send a positive or negative acknowledgment, $t_r = 2t_{max} + t_{ack}$. If the CP is diagnosed as lost, the controller node for the respective floor can regenerate the token and send an update to the session.




Figure 3. Coordination cost for various communication models.

For a simple comparison of the coordination overhead incurred with unicast, multicast, and aggregated multicast, measured by the number of CPs to be exchanged, we assume full load (each node sends a CP) and a normalized processing time for request, response and update packets. The average path length between nodes is assumed to be the same for all models. λ represents the individual processing, packetization and transmission delay for each type of packet. In unicast, the coordination delay incurs $(N - 1)$ requests, replies from control nodes, and updates, where N is the current session size, i.e., $CD_{uc} = 3(N - 1)\lambda$. In multicast, $(N - 1)$ nodes send requests, and the control node multicasts one reply and one update back to the session, i.e., $CD_{mc} = (N + 1)\lambda$. In aggregated multicast, CPs are handled within multicast groups and only the root of a group forwards a composite request to its parent, or responds to group-local requests, if it holds the information locally. With K

groups we have on the average $G = \lceil \frac{N}{K} \rceil$ members per group, and per group there are G requests inside a group, K aggregated requests sent to a control node from all groups, and one multicast response and update, i.e., $CD_{amc} = ((G - 1) + (K - 1) + 2)\lambda = (G + K)\lambda$. Figure 3 shows the average cost to coordinate hosts in sessions up to size $N = 1000$, clustered into $K = N/10$ groups, and with normalized transmission delay λ . The benefits of aggregation in conjunction with multicast are apparent.

4. Conclusions

The IP multicast model lacks refined support for intergroup coordination. In the wake of a new generation of interactive and collaborative applications, we investigated the possibility to support group-coordinative systems with network layer functionality. To that end, we have outlined the main issues in establishing a coordination mechanism based on the IP multicast service model. The discussed protocol, ACP, operates on a logical control tree and scales to large groups and wide scope. Prefix matching previously used in CIDR addressing and multicast routing has been put to use for our multisite coordination mechanism to effectively channel the exchange of coordination messages. The proposed approach eliminates the need to build a separate control structure for tracking, routing, withholding, or forwarding control directives. A side benefit of prefix matching is the trivial provision of subcasting to and within multicast groups. We have described ACP as a mechanism strongly coupled with the operations in an underlying routing and end-to-end delivery tree. However, ACP may implement its own tree management at layer 4+ in an overlay network aside from multicast routing and thus can operate independently of any layer 3 protocol. While IP multicast is still plagued by deployment and standardization problems, our protocol can be viewed as a first solution towards coordination middleware components for emerging collaborative service infrastructures in the next-generation Internet.

Acknowledgements

This work was supported by the Defense Advanced Research Projects Agency (DARPA) under Grant F19628-96-C-0038.

References

1. Aguilar, L., J. J. Garcia-Luna-Aceves, D. Moran, E. J. Craighill, and R. Brungard: 1986, 'Architecture for a Multimedia Teleconferencing System'. In: *Proc. ACM SIGCOMM*. pp. 126–136.
2. Almeroth, K.: 2000, 'The Evolution of Multicast: From the MBone to Interdomain Multicast to Internet2 Deployment'. *IEEE Network* **14**(1), 10–20.
3. Barghouti, N. S. and G. E. Kaiser: 1991, 'Concurrency Control in Advanced Database Applications'. *ACM Computing Surveys* **23**(3), 269–317.
4. Diot, C., B. N. Levine, B. Lyles, H. Kassem, and D. Balensiefen: 2000, 'Deployment Issues for the IP Multicast Service and Architecture'. *IEEE Network* **14**(1), 78–88.
5. Dommel, H.-P. and J. J. Garcia-Luna-Aceves: 1997, 'Floor Control for Multimedia Conferencing and Collaboration'. *Multimedia Systems J. (ACM/Springer)* **5**(1), 23–38.
6. Ellis, C. A. and S. J. Gibbs: 1989, 'Concurrency Control in Groupware Systems'. In: *Proc. ACM SIGMOD Int. Conf. on Management of Data*. Portland, OR, pp. 399–407.
7. Handley, M. and V. Jacobson: 1998, 'SDP: Session Description Protocol'. RFC 2327.
8. International Telecommunication Union: 1996, 'Recommendation T.120 on Data protocols for multimedia conferencing'. <http://www.itu.int>.
9. Isaacs, E. A., T. Morris, T. K. Rodriguez, and J. C. Tang: 1995, 'A Comparison of Face-To-Face and Distributed Presentations'. In: *Proc. ACM CHI*. Denver, CO.
10. Jia, X.: 1995, 'A Total Ordering Multicast Protocol Using Propagation Trees'. *IEEE Trans. Parallel and Distrib. Sys.* **6**(6), 617–627.
11. Lampson, B.: 1974, 'Protection'. *Proc. 5th Princeton Conf. on Information Sciences and Systems* **8**(1), 18–24.
12. Levine, B. N. and J. J. Garcia-Luna-Aceves: 1997, 'Improving Internet Multicast with Routing Labels'. In: *Proc. IEEE Int. Conf. on Network Protocols*. Atlanta, GA, pp. 241–250.
13. Levine, B. N., D. Lavo, and J. J. Garcia-Luna-Aceves: 1996, 'The Case for Reliable Concurrent Multicasting Using Shared Ack Trees'. In: *Proc. ACM Multimedia*. Boston, MA, pp. 365–376.
14. Rangan, P. V., H. M. Vin, and S. Ramanathan: 1993, 'Communication Architectures and Algorithms for Media Mixing in Multimedia Conferences'. *IEEE/ACM Trans. on Networking* **1**(1), 20–30.
15. Raymond, K.: 1989, 'A Tree-based Algorithm for Distributed Mutual Exclusion'. *ACM Trans. on Comp. Sys.* **7**(1), 61–77.
16. Ruiz-Sanchez, M. A., E. W. Biersack, and W. Dabbous: 2001, 'Survey and Taxonomy of IP Address Lookup Algorithms'. *IEEE Network* **15**(2), 8–23.
17. Vin, H. M., P. V. Rangan, and S. Ramanathan: 1991, 'Hierarchical Conferencing Architectures for Inter-Group Multimedia Collaboration'. In: *ACM SIGOIS Bull., Proc. Org. Comp. Sys.* Atlanta, GA, pp. 43–54.
18. Yavatkar, R. and K. Lakshman: 1994, 'Communication support for distributed collaborative applications'. *Multimedia Systems J.* **2**(2), 74–88.
19. Ziegler, C., G. Weiss, and E. Friedman: 1989, 'Implementation mechanisms for packet switched voice conferencing'. *IEEE J. on Sel. Areas in Comm.* **7**(5), 698–706.

Authors' Vitae

Hans-Peter Dommel

is currently Assistant Professor of Computer Engineering at Santa Clara University, California. In 1990, he received a Master's degree in computer science and theoretical linguistics from the Technical University of Munich, Germany. Prior to coming to the United States in 1992 on a Fulbright Scholarship, he worked for IBM, Siemens and other companies in Germany and Switzerland. His current research interests include networked multimedia systems, group-centric computing, and web technology. He is a member of the ACM, the IEEE Computer Society, and the Gesellschaft for Informatik, Germany.

J. J. Garcia-Luna-Aceves

received the B.S. degree in electrical engineering from the Universidad Iberoamericana, Mexico City, Mexico, in 1977, and the M.S. and Ph.D. degrees in electrical engineering from the University of Hawaii, Honolulu, HI, in 1980 and 1983, respectively. He is Professor of Computer Engineering at the University of California, Santa Cruz (UCSC) and consults for NOKIA wireless routers on protocol design. He held the position of Visiting Professor at Sun Laboratories in Menlo Park, California in 1999. Prior to joining UCSC in 1993, he was a Center Director at SRI International (SRI) in Menlo Park, California. He first joined SRI as an SRI International Fellow in 1982. His current research interest is the analysis and design of algorithms and protocols for computer communication. At UCSC, he leads a number of research projects sponsored by DARPA, ONR, Air Force, and industry that focus on wireless networks and internetworking. Dr. Garcia-Luna-Aceves has published a book and more than 190 refereed papers in journals and conferences. He has been on the editorial boards of the IEEE/ACM Transactions on Networking, the ACM Multimedia Systems Journal, and the Journal of High Speed Networks. Dr. Garcia-Luna-Aceves was Program Co-Chair of the ACM/IEEE Mobicom 2000 Conference; and has been Chair of the ACM special interest group on multimedia, General Chair of the first ACM conference on multimedia: ACM MULTIMEDIA '93, Program Chair of the IEEE MULTIMEDIA '92 Workshop, General Chair of the ACM SIGCOMM '88 Symposium, and Program Chair of the ACM SIGCOMM '87 Workshop and the ACM SIGCOMM '86 Symposium. He has also been program committee member for numerous IFIP 6.5, ACM, IEEE, and SPIE conferences. He received the SRI International Exceptional-Achievement Award in 1985 for his work on multimedia communications, and again in 1989 for his work on adaptive routing algorithms. He is a member of the ACM and the IEEE.

