
DOM: A Scalable Multicast Protocol for Next-Generation Internet

Xiaohua Tian and Yu Cheng, Illinois Institute of Technology
Xuemin (Sherman) Shen, University of Waterloo

Abstract

In this article we propose a scalable and efficient destination-oriented multicast protocol for next-generation Internet. With DOM, each packet carries explicit destinations information, instead of an implicit group address, to facilitate the multicast data delivery; and each router leverages the unicast IP routing table to determine necessary multicast copies and next-hop interfaces. The fundamental issue is that we need to limit the bandwidth overhead for such explicit addressing, since it is impractical to attach all the destination addresses to each packet. We resort to the Bloom filter technique to encode the destination information carried by each packet for bandwidth efficiency, with elaborated design to accommodate the features of the practical Internet including longest-prefix matching, route aggregation, and asymmetric interdomain routing. Moreover, DOM enables a fast group joining approach to minimize the joining delay perceived by receivers. The scalability and efficiency of DOM are demonstrated with simulation results.

The next-generation Internet will accommodate various multimedia applications over a common IP-based transport infrastructure. Many multimedia services (e.g., online multiplayer games, IPTV, videoconferencing) require a scalable and efficient multicast mechanism capable of disseminating the shared data to widely distributed destinations and quick in responding to subscribing requests from end users. Traditional IP multicast delivers data for each group along a tree structure at the network layer, which is constructed by a distributed multicast routing algorithm. IP multicast is bandwidth efficient in data delivery but suffers from the scalability issue [1], since the messaging overhead and memory cost grow linearly with the number of multicast groups supported by the router. The more recent overlay multicast establishes the data dissemination structure at the application layer. Each link in the overlay network is an end-to-end logic connection between two end hosts. Although overlay multicast requires no modification of the unicast infrastructure, it induces redundant traffic at the network layer [2]: it is common that separate overlay links pass through common physical links in the underlying transport network.

Multicast is a generic service paradigm commonly used by a wide variety of multimedia applications. A scalable and efficient multicast mechanism should be one of the infrastructural functionalities provisioned by the next-generation Internet [3]. Our perspective is that the rapid development of hardware and software technologies enable networking devices to execute more complex computing and more intelligent processing, which provides an opportunity to streamline the design of networking functionalities for next-generation Internet. Extensive efforts are being made in both academia [3, 4] and industry [5] on incorporating application-oriented intelligence into the network layer.

In fact, enhancing network nodes with more intelligence

has become a mainstream idea to design the modern multicast protocols. REUNITE [6] stores the destination addresses of a multicast group at network routers located at the intersection points along a multicast tree to facilitate multicast forwarding; however, it needs to maintain group-specific information through a soft-state mechanism, which still incurs considerable memory and message overheads. In Xcast [7] each packet carries a list of destination addresses, and the network router can process the destinations list for multicasting. Although Xcast takes a similar service model as adopted in this article, it is not scalable because the limited size of the Xcast header constrains the protocol to be applicable only in a scenario with a small number of multicast receivers. The FRM scheme [8], most closely related to our study, encodes the multicasting tree into the packet using a Bloom filter, and computes packet copies and output interfaces by checking the router's neighboring edges against the embedded tree. Our method instead only encodes receiver destinations in the packet, which can reduce bandwidth overhead and facilitate the group joining process. The LIPSIN scheme [3] uses a similar methodology as FRM with enhanced techniques to reduce false positives incurred by the Bloom filter processing, but requires an extra protocol to assign and manage ID for each link of the network. In MAD [9] it is proposed to adaptively switch between IP and overlay multicast modes, depending on the service scenarios. Such a scheme can mitigate memory overhead when there are a large number of groups with infrequent data traffic, but cannot fundamentally solve the scalability issue.

This article proposes destination-oriented multicast (DOM), following the philosophy that network routers can have enhanced intelligence in addition to basic routing and forwarding. With DOM, each packet carries explicit destination information instead of an implicit group address to facilitate multicast data delivery; and each router leverages the

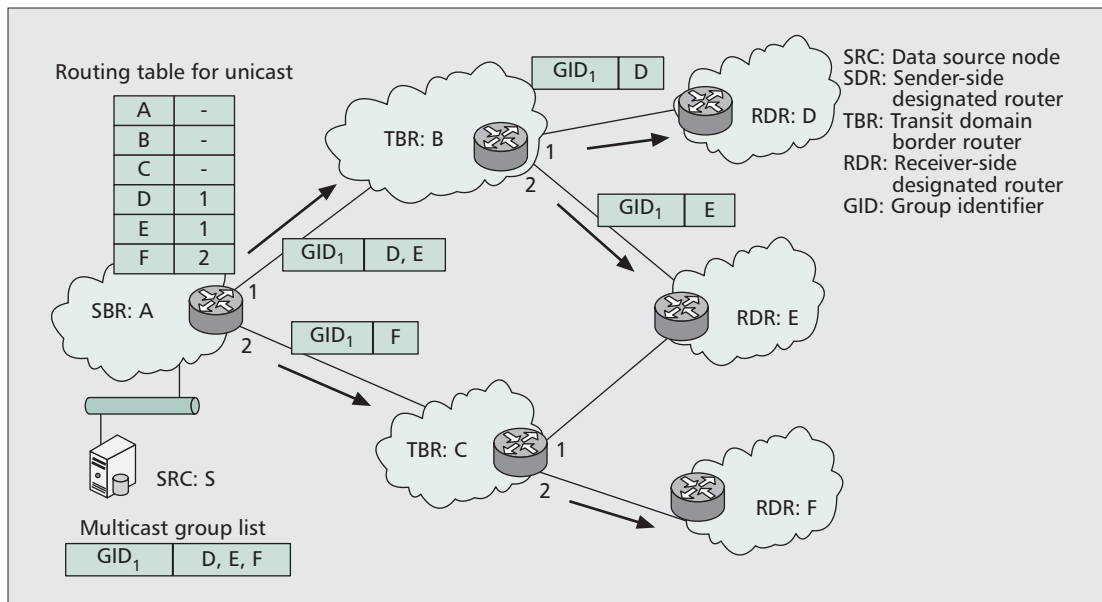


Figure 1. DOM service model.

unicast IP routing table to determine necessary multicast copies and next-hop interfaces. The fundamental issue is that we want to limit the bandwidth overhead for such explicit addressing since it is impractical to attach all the destination addresses to each packet. In our design the in-packet destination information is encoded into the Bloom filter to achieve bandwidth efficiency. The Bloom-filter-based design is further elaborated to accommodate the features of the practical Internet including the longest-prefix matching, route aggregation, and asymmetric interdomain routing. We focus on presenting the concept and design methodology of DOM in this article, while the implementation details of the protocol are available in [10]. Moreover, this article presents a new border gateway protocol (BGP)-view-based joining scheme to address the asymmetric routing issue for interdomain multicasting and a fast group joining approach to minimize the joining delay perceived by receivers.

The proposed DOM has the following properties:

- The memory and messaging overheads incurred at each router for DOM are independent of the number of groups being supported.
- The forwarding bandwidth efficiency of DOM is very close to that of IP multicast.
- DOM enables establishing the reverse shortest path tree (SPT) even with asymmetric routing, and minimizes the joining delay perceived by the receivers.
- DOM decouples the group identifier (ID) from multicast routing and forwarding, and enables a scalable source-specific localized group ID allocation scheme.

In the remainder of this article we first describe the DOM service model and discuss associated practical design issues. We then present the Bloom-filter-based design to realize the DOM protocol in the practical Internet. Performance analysis and simulation results are demonstrated to verify the scalability and efficiency of DOM. Future work is briefly discussed in the concluding remarks.

DOM

Service Model

In the DOM service model, each receiver domain sends a join message to the source node of the groups in which it is interested; based on the join messages, each source node can know the members for each group it provisions. Each multicast

packet will carry the destination addresses of all the members so that each router can retrieve the addresses and leverage the unicast routing table to compute necessary copies and output interfaces. DOM provides a uniformed intra-/interdomain multicasting service, but we focus on the interdomain case in this article for convenience of demonstration.

Membership Management — For membership management, a border router of a stub autonomous system (AS) domain is selected as the designated router (DR). We use RDR (SDR) to denote the DR of a receiver-side (source-side) AS domain. When multicast routing/forwarding is considered, RDR also represents the prefix associated with the corresponding receiver domain. The meaning of RDR will be clear in the context. The transit domain border router is denoted TBR.

The RDR basically needs to implement the Internet Group Management Protocol (IGMP) to discover the active groups within its domain. When new groups are activated, the RDR is triggered to send membership updating messages (MUMs) to the data source node (SRC) in the format as (RDR: GID₁, GID₂, ..., GID_n), where RDR represents a domain prefix and GID represents the group ID. Note that the MUM messages will be forwarded along the shortest path between the RDR and the SRC according to the unicast routing information.

The SRC aggregates the MUM messages it has received and maintains a multicast group list (MGL). For each group provisioned by the SRC, the MGL establishes a record in the format as (GID: RDR₁, RDR₂, ..., RDR_n), where each RDR again indicates a domain prefix. The MGL lets the SRC trace the active receiver domains for each group it provisions. Figure 1 shows an example where the record of group 1 in the MGL is illustrated. When the SRC multicasts data over a certain group, it inserts the corresponding MGL into the packet as the destination information in the format of a shim header. The multicast packets are then forwarded to the SDR for interdomain multicasting.

Multicast Forwarding Protocol — When receiving a multicast packet, each router will extract the MGL record from the packet, and compute packet copies and corresponding output interfaces. Specifically, a DOM router performs the following processing. First, check the unicast routing table to determine the output interface for each destination listed in the MGL of the packet, and aggregate destinations with the same output interface into a set. Second, replicate the packet for each

unique interface found in the first step. Third, update the MGL of each packet copy with the aggregated set yielded in the first step so that the packet copy for a given interface contains only the destinations that can be reached via this interface. By removing unnecessary destinations from the MGL record, the downstream router will not generate unnecessary packet copies for those destinations that have been delivered over other sibling subtrees. Each router will execute the same operations of aggregation, replication, and MGL record updating, until one multicast packet reaches an RDR. Figure 1 depicts the forwarding process of group 1, where the destinations list includes *D*, *E*, and *F*.

We consider the case that all routers (including SDRs, RDRs, and TBRs) are aware of DOM. How to deploy DOM in a compatible manner with the legacy routers will be studied in the future work.

Practical Design Issues

In the prototype DOM service model all the routers involved in the multicast forwarding (other than the DRs) do not need to maintain any state regarding multicasting. The forwarding complexity is totally independent of the number of groups to be supported, resulting in desirable scalability. Nevertheless, considerable bandwidth overhead could be incurred when there are a large number of receivers (RDRs) for each group: the MGL in the packet becomes impractically long, and the number of receivers that can be supported is constrained by the packet header size.

A possible solution is resorting to the Bloom filter technique [8]. In DOM, the MGL can be encoded into a Bloom filter which saves much space in the packet (e.g., about 40 percent in our simulation). However, the Bloom-filter-based design needs to support the particular features of Internet. Normally, Internet routers apply the longest-prefix matching and route aggregation schemes to control the size of the unicast routing table, by which the same destination network may be represented with different network prefixes in different routers. Since the Bloom filter only supports exact query, it is possible that the destination RDR prefixes encoded in the Bloom filter cannot match any forwarding entry stored in a SDR/TBR. Instead of directly utilizing the unicast routing table, there is a need to establish the forwarding states that can recognize the Bloom-filter-formatted MGL along the data delivery path.

Most of the existing multicast protocols [1] establish the forwarding states when the joining request is delivered from the receiver to the source node (or rendezvous point), and then forward the data packets along the path that is reverse to the joining path. Such a reverse path forwarding (RPF)-based approach, however, requires the symmetric routing environment: the path from the source to a receiver follows the same path used to go from the receiver to the source. Unfortunately, the interdomain routing is usually asymmetric for administrative reasons [6]. When designing the DOM, we also need to consider the effect of asymmetric routing on the protocol.

Implementation of DOM Protocol

This section presents a Bloom-filter-based design for DOM, where the practical design issues mentioned above are solved with reasonable cost.

Bloom Filter Based Design

We are to describe the Bloom-filter-based DOM design according to the upstream procedure (i.e., states establishment) and the downstream procedure (i.e., data forwarding),

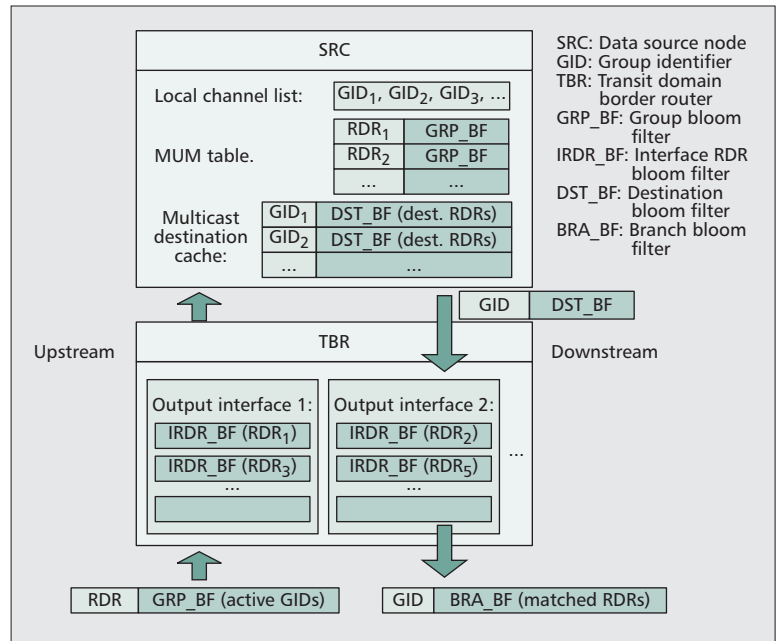


Figure 2. Bloom-filter-based design of DOM.

as illustrated in Fig. 2, where Bloom filters are illustrated as shadowed areas.

The left side of Fig. 2 shows how forwarding states are established by MUM messages. To reduce the bandwidth overhead for membership updating, the list of active groups in the MUM message is encoded with a *group Bloom filter* (GRP_BF). When an MUM message reaches an upstream TBR/SDR router, the router will retrieve the RDR prefix, and store it as a local forwarding state that leads to a reverse path of the MUM incoming interface. By continuously observing the MUMs, each related interface of the TBR/SDR will memorize all the destination domains that can be reached through it, and the reverse SPT from the SRC to subscribing RDRs can then be constructed. At an output interface, each RDR is stored as a separate Bloom filter, called an *interface RDR Bloom filter* (IRDR_BF), which will be used to facilitate multicast forwarding.

The upstream MUM messages will finally reach the SRC node, and each message will be stored as a record of the MUM table. The SRC node should have a *local channel list* indicating the multicast groups it provisions. By checking each GID against the MUM table and identifying the matched GRP_BF, the SRC can detect the destination prefixes for a given group. The destinations information under the group ID will be encoded into a *destination Bloom filter* (DST_BF) and stored in the multicast destination cache. Note that the DST_BF in fact encodes the MGL according to the DOM service model.

The right side of Fig. 2 illustrates how multicast packets are forwarded. At the SRC node, the DST_BF for a group will be inserted as the destination information into each multicast packet. In the downstream data forwarding process, each router generally executes the same operations of aggregation, replication, and MGL record updating introduced in the DOM service model. The only difference is that these operations are conducted with Bloom filters in both the packet and the router. Specifically, each TBR/SDR compares the packet's DST_BF with IRDR_BFs at each interface. A packet replica is generated and dispatched along the interface, if the DST_BF and IRDR_BFs installed at the interface have any element matched. The subset of matched prefixes associated with each output interface is then re-encoded into the *branch Bloom filter* (BRA_BF). The BRA_BF will be inserted into

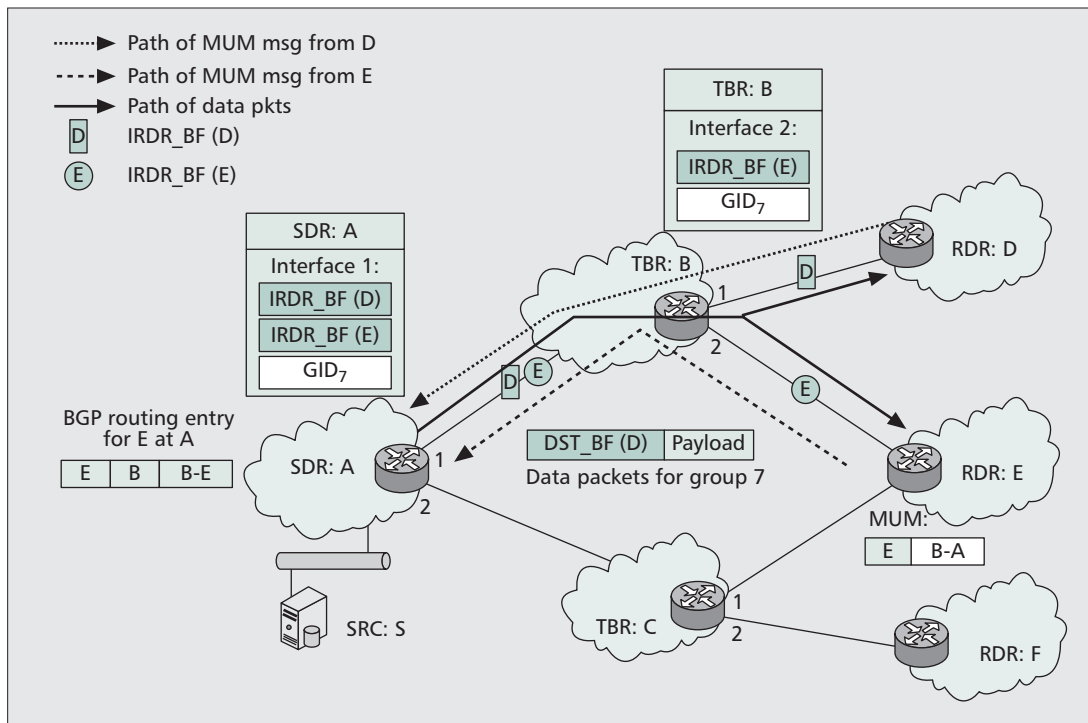


Figure 3. DOM with asymmetric routing and fast group joining.

the packet replica delivered through that interface, serving as the destination information `DST_BF` for further downstream forwarding.

The downstream procedure actually follows the RPF concept: the data packet is forwarded along the path that is reverse to the MUM joining path. The states installed by MUM messages explicitly construct a reverse SPT rooted at the SRC node, and the forwarding procedure avoids directly checking the unicast routing table to decide output interfaces. With such RPF-based multicasting, DOM will not be impacted by the issues of longest-prefix matching and route aggregation.

DOM with Asymmetric Routing

In Fig. 1 it is possible that the MUM sent by *E* takes the path *E-C-A* to reach *S*, while the downstream data path is *A-B-E*; thus, the forwarding states cannot be installed following the RPF concept. To address such an issue, an option is to leverage the Multiprotocol Extensions to BGP-4 (MBGP) [1], which can announce different unicast- and multicast-capable routes to help the MUM messages take the correct joining path to the SRC; however, MBGP incurs high complexity.

We thus propose a low-complexity solution to address the asymmetric routing issue, the BGP-view-based joining scheme. We have two assumptions: the physical links of the data delivery path from the SDR to a RDR are bidirectional, and the interdomain routing policy always allows control messages (e.g., MUMs) traveling along the path reverse to the data delivery path. These two assumptions are very reasonable in practice. Most of the physical links are bidirectional to save the cost in deployment, and a more flexible policy to multicast control messages is preferable considering the remarkable benefit from a scalable multicast design.

The BGP-view-based joining scheme could construct the reverse SPT for multicasting even with asymmetric interdomain routing, as illustrated in Fig. 3. The service provider designates a BGP-speaking SDR, which knows the shortest paths from itself to any possible receivers. The information is stored in the local BGP routing table, where each table entry repre-

sents the local routing view for a given destination network prefix. For instance, the BGP routing entry for the network associated with *E* shows that *E* can be reached through the next-hop *B* and the path vector *B-E* in Fig. 3. The BGP routing entry will be passed to the corresponding RDR so that the receiver side knows the actual routing view the sender side can see. Then the MUM is forwarded along the reverse path indicated by the BGP path vector with source routing (e.g., *E-B-A* in Fig. 3). In this manner the MUM can correctly install forwarding states at corresponding routers along the indicated joining path to the SRC, the dashed line path in Fig. 3.

A natural question is: how can the BGP view seen by the SDR be passed to a RDR? The key observation is that DOM adopts the *source-based* service model [1], where each receiver knows the SRC information (SRC IP address, channel number, etc.) before subscribing to a channel. A number of techniques can be used to transport the BGP routing entry from the SDR to a RDR, for example, via web pages or sessions announcement applications. We note that there should be an interface between the SDR and multicasting applications at SRC to retrieve BGP views considering the BGP routing selection rules; the implementation details are out of the scope of this article.

Fast Group Joining

The DOM design has the benefit of enabling a fast group joining procedure: it is possible for the RDR to start receiving requested packets before its MUM message arrives at the SRC. The DOM protocol needs some modification to support the fast group joining. Along the path that the MUM messages travel, the routers not only add the `IRDR_BF` information but also temporarily store the group IDs. While the regular TBR/SDR forwarding process only compares the `DST_BF` against the `IRDR_BFs`, a `GID`-based forwarding is added as an assistant process to facilitate the fast join scheme.

An example of fast group joining is illustrated in Fig. 3. Consider that a multicast tree has been established between the SRC and the subscriber RDRs after an initiation period. A subsequent MUM from *E* can follow the same path marked

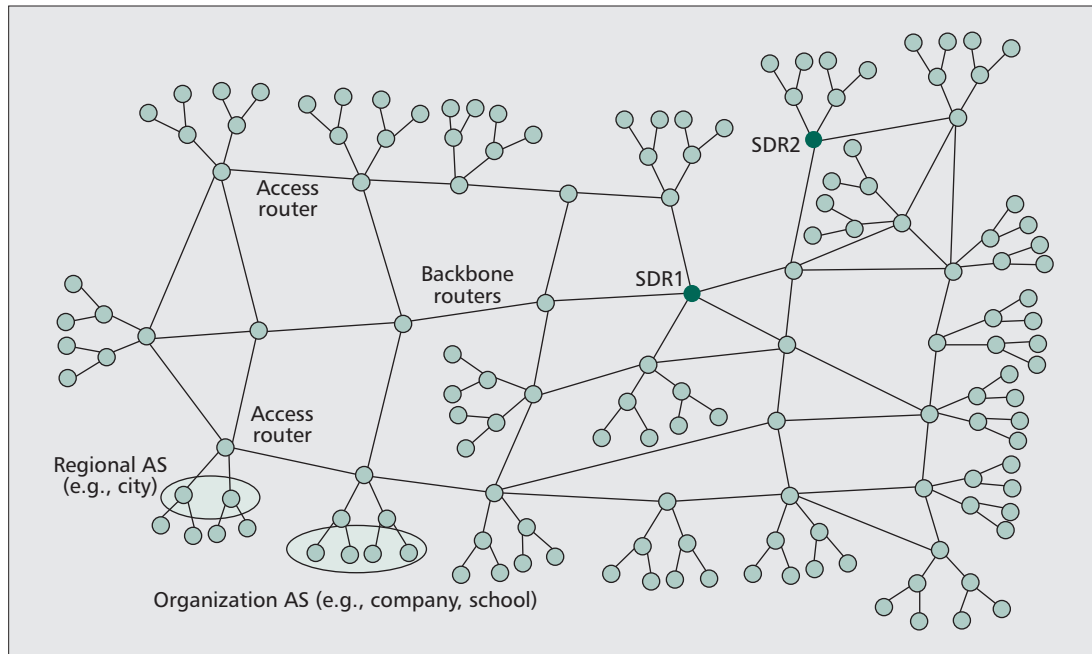


Figure 4. Simulation topology.

by the dashed line to subscribe to a newly activated group within its domain, say group 7. Upon receiving the new MUM, the TBR router *B* will temporarily store the group ID, GID_7 , at its interface 2, and then forward the MUM further to the next hop. If *B* has already been forwarding data packets of group 7 to *D* along its interface 1 when the MUM arrives, it will find that the group 7 data packets being forwarded via interface 1 match the GID_7 labeled at interface 2. Node *B* will then immediately forward the same data packets via interface 2 to *E*, although $DST_BF(D)$ and $IRDR_BF(E)$ do not match. Thus, *E* can receive the requested data packets before the MUM arrives at *S*.

We emphasize that the GIDs need to be stored in the router only temporarily. When the regular forwarding process later confirms that the packet with a certain GID should be dispatched through a certain interface, the GID-based forwarding process is then stopped, and the corresponding GID label on that interface will be removed. Consider the group 7 packet at SDR *A* in Fig. 3; as $DST_BF(D)$ matches the $IRDR_BF(D)$, the regular forwarding process confirms that the group 7 packet should be dispatched through *A*'s interface 1, and GID_7 can be deleted immediately. In TBR *B*, however, the GID_7 can be deleted only after the first packet with the updated $DST_BF(D,E)$ arrives. At that time, the regular forwarding process will confirm that group 7 packets should be forwarded via interface 2, and GID_7 can be removed. The temporal GID labels are the trade-off cost for fast group joining. With the label removing mechanism, the number of forwarding states at each router is still independent of the number of groups in the long run.

Performance Analysis

Scalability

Compared with IP multicast, DOM installs fewer forwarding states at routers on the multicasting tree. This is because DOM stores only one state on each related node for each subscribing domain. In contrast, each subscribing domain may join in tens of thousands of groups, and each group needs a state on each related node under IP multicast. The number of forwarding states per node for DOM is independent of the number of groups being supported by the node. Moreover,

DOM decouples the membership management component from the multicast forwarding component. The forwarding component at a router just requires RDR related information. Group IDs are only used for labeling groups at the SRC and RDRs to establish the service relationship; thus, the group IDs can be allocated at the SRC locally in the form of a two-tuple (source node address, source-specific channel ID) [6], which breaks the address space limitation of IPv4 Class-D addresses.

Bandwidth Efficiency and False Positive

DOM strikes a balance between bandwidth efficiency and small Bloom filter false positive rate, in comparison with the closest related work, FRM. The fundamental difference between DOM and FRM is that the DOM encodes only destination prefixes in the packet shim header while the FRM encodes multicasting tree branches. The Bloom filter incurs false positive, which means the element not encoded in the Bloom filter might be falsely detected. For a fixed-length bloom filter, the more elements are encoded, the higher the false positive rate can be. In DOM and FRM, when the number of receivers/branches exceeds the capacity of a single shim header, multiple packets are sent to cover all destinations, which are counted as redundant traffic. Since covering the same number of destinations normally requires more branches, DOM can generate less redundant traffic than FRM does if they keep the same false positive rate [10].

Joining Delay

In DOM receivers could perceive shorter joining delays than in FRM. This is because the DOM fast joining mechanism can efficiently respond to join requests at the intermediate router already on the multicasting tree, while the joining process of FRM has to be completed at the SRC. In a practical scenario there are a large number of groups running over the Internet, so when an RDR wants to join in a new group discovered within its domain, it is very possible that the join request can meet some intermediate router which has been forwarding the requested data to other members already in the group. Thus, the requested data can be steered to the RDR immediately instead of waiting until the join request arrives at the SRC.




Figure 5. CDF of the number of forwarding entries per node ($N = 80$).

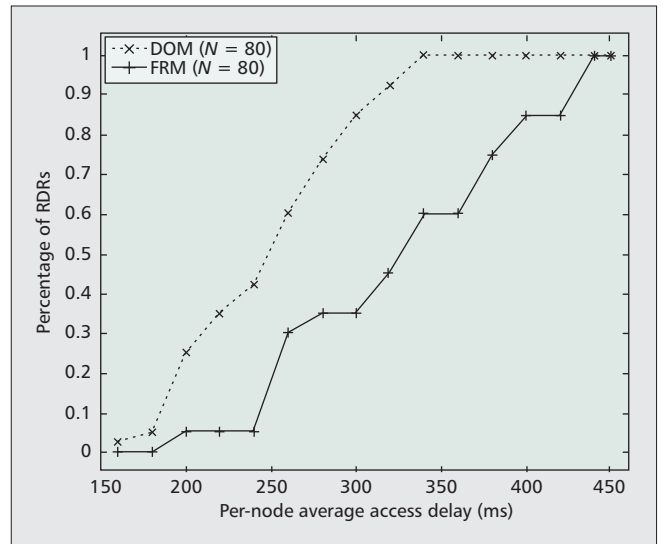


Figure 7. CDF of per-node average access delay ($N = 80$).

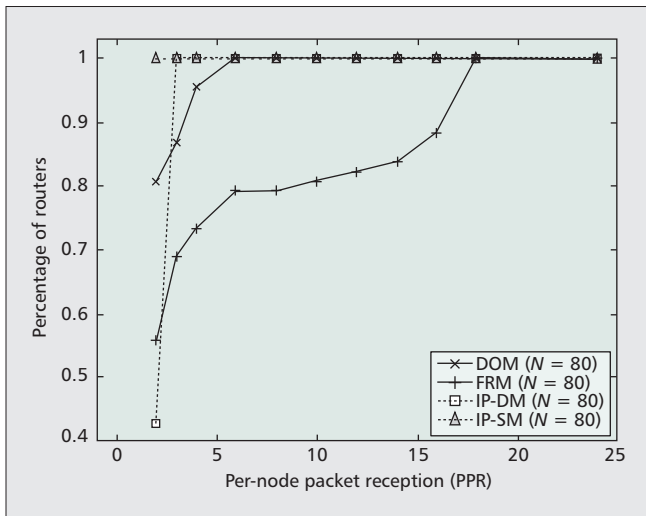


Figure 6. CDF of per-node packet reception ($N = 80$).

Simulation Results

In this section we present some NS-2 simulation results to compare the performance of DOM with other reference multicast schemes including IP multicast Dense Mode (DM), Sparse Mode (SM), and FRM. The network topology for simulation is given in Fig. 4, which is widely used in the literature as a hypothetical U.S. backbone network. There are two SRCs associated with the two SDRs, as illustrated in Fig. 4, and RDRs are placed at organization AS domains. In the simulation each RDR joins in $G = 200$ groups evenly provided by the two SRCs.

For memory overhead, we count the number of forwarding entries at intermediate routers of the multicasting tree for each scheme under study. Figure 5 shows the cumulative distribution function (CDF) of the number of forwarding entries per node in each scheme, where there are $N = 80$ RDRs. DOM can significantly reduce the number of forwarding states stored at each node compared to IP multicast because DOM stores destination-specific states at the router. The number of states per node for DOM is independent of the number of groups passing through, in contrast to IP multicast. The low memory overhead could facilitate high-speed for-

warding on the line card, which has limited memory capacity. The forwarding states in the FRM router are in fact the router's AS neighbor edges, which are used to compare with the tree branches encoded in the packet shim header to determine output interfaces. Thus, the number of forwarding entries per node in FRM is the AS degree of this node.

We examine the bandwidth overhead by counting the number of packets received at a given node when multicasting a single packet from each SRC to all receivers. Figure 6 shows that the number of packets each router receives under DOM is close to that under IP multicast and less than that under FRM, which means that the bandwidth overhead of DOM is close to that of IP multicast and less than that of FRM. The reason is that DOM encodes only destination information in the packet, with the intermediate routers maintaining the forwarding states, while FRM encodes the multicasting tree branches in the packet. It usually takes more tree branches to cover the same number of destinations; therefore, FRM generates more redundant traffic than DOM.

Although DOM and FRM both adopt the source-based joining operation that has to be finished at the SRC, DOM performs better than FRM in terms of joining delay with the fast group joining scheme. We compare CDFs for the per-node access delay averaging the participated 200 groups under DOM and FRM, as shown in Fig. 7. The access delay is the time for the data to start arriving at the RDR since sending the join request. It is obvious that RDRs experience less access delay under DOM than under FRM.

Conclusions and Future Work

In this article we have proposed a scalable and efficient multicast scheme, DOM, for next-generation Internet. By carrying necessary receiver addresses information within the packet shim header, the proposed DOM can leverage the existing unicast IP routing information to deliver the data to receivers. In the Bloom-filter-based implementation of DOM, all the management and addressing information traversing the network is encoded with Bloom filters for memory and bandwidth efficiency. Moreover, the DOM design allows seamless integration with practical Internet features including longest-prefix matching, route aggregation, and asymmetric interdomain routing. In addition, DOM enables a fast group joining approach to minimize the joining delay perceived by receivers. For future work, we will study how to incrementally deploy the DOM with backward compatibility with legacy routers.

References

- [1] K. C. Almeroth, "The Evolution of Multicast: from the Mbone to Interdomain Multicast to Internet2 Deployment," *IEEE Network*, vol. 14, no. 1, Jan./Feb. 2000, pp. 10–20.
- [2] S. Fahmy and M. Kwon, "Characterizing Overlay Multicast Networks and their Costs," *IEEE/ACM Trans. Net.*, vol. 15, no. 2, Apr. 2007, pp. 373–86.
- [3] P. Jokela *et al.*, "LIPSIN: Line Speed Publish/Subscribe Inter-Networking," *Proc. ACM SIGCOMM*, Aug. 2009, pp. 195–205.
- [4] A. Feldmann, "Internet Clean-Slate Design: What and Why?" *SIGCOMM Comp. Commun. Rev.*, vol. 37, no. 3, July 2007, pp. 59–64.
- [5] Cisco Systems, "Cisco Application-Oriented Networking"; <http://www.cisco.com/go/aon>
- [6] I. Stoica, T. S. E. Ng, and H. Zhang, "REUNITE: A Recursive Unicast Approach to Multicast," *Proc. IEEE INFOCOM*, vol. 3, Mar. 2000, pp. 1644–53.
- [7] R. Boivie *et al.*, "Explicit Multicast (Xcast) Basic Specification," Internet draft, Mar. 2001.
- [8] S. Ratnasamy, A. Ermolinskiy, and S. Shenker, "Revisiting IP Multicast," *Proc. ACM SIGCOMM*, Aug. 2006, pp. 15–26.
- [9] T. W. Cho *et al.*, "Enabling Content Dissemination using Efficient and Scalable Multicast," *Proc. IEEE INFOCOM*, Mar. 2009, pp. 1980–88.
- [10] X. Tian, Y. Cheng, and B. Liu, "Design of a Scalable Multicast Scheme with an Application-Network Cross-Layer Approach," *IEEE Trans. Multimedia*, vol. 11, no. 6, Oct. 2009, pp. 1160–69.

Biographies

XIAOHUA TIAN [S'08] (xtian3@iit.edu) received his B.E. and M.E. degrees in communication engineering from Northwestern Polytechnical University, Xi'an, China, in 2003 and 2006, respectively. He is currently working toward a Ph.D. degree in the Department of Electrical and Computer Engineering, Illinois Institute of Technology, Chicago. His current research interests include application-oriented networks, multicast protocols, cross-layer design for multimedia networking, and peer-to-peer networks.

YU CHENG [S'01, M'04, SM'09] (cheng@iit.edu) received his Ph.D. degree in electrical and computer engineering from the University of Waterloo in 2003. Since August 2006 he has been an assistant professor in the Department of Elec-

trical and Computer Engineering, Illinois Institute of Technology. His research interests include application-oriented networking, Internet performance analysis, wireless networks, and network security. He received a Postdoctoral Fellowship Award from the Natural Sciences and Engineering Research Council of Canada (NSERC) in 2004, and a best paper award from the International Conference on Heterogeneous Networking for Quality, Reliability, Security, and Robustness 2007 (QShine '07). He served as a Technical Program Co-Chair for the Wireless Networking Symposium of IEEE ICC 2009. He is an Associate Editor for *IEEE Transactions on Vehicular Technology* and Area Editor for *Computer Networks*.

XUEMIN (SHERMAN) SHEN [M'97, SM'02, F'09] (xshen@bcr.uwaterloo.ca) received his B.Sc. (1982) degree from Dalian Maritime University, China, and his M.Sc. (1987) and Ph.D. degrees (1990) from Rutgers University, New Jersey, all in electrical engineering. He is a professor and University Research Chair, Department of Electrical and Computer Engineering, University of Waterloo, Canada. His research focuses on mobility and resource management in interconnected wireless/wired networks, UWB wireless communications networks, wireless network security, wireless body area networks, and vehicular ad hoc and sensor networks. He is a co-author of three books, and has published more than 400 papers and book chapters in wireless communications and networks. He has served/serves as Technical Program Committee Chair for IEEE VTC '10 Fall, Symposia Chair for IEEE ICC '10, Tutorial Chair for IEEE ICC '08, Technical Program Committee Chair for IEEE GLOBECOM '07, General Co-Chair for CHINACOM '07 and QShine '06, and Founding Chair of the IEEE Communications Society Technical Committee on P2P Communications and Networking. He also has served/serves as a Founding Area Editor for *IEEE Transactions on Wireless Communications*, Editor-in-Chief for *Peer-to-Peer Networking and Application*, and Associate Editor for *IEEE Transactions on Vehicular Technology*, *Computer Networks*, and *ACM/Wireless Networks*. He has also served as Guest Editor for *IEEE Journal on Selected Areas in Communications*, *IEEE Wireless Communications*, and *IEEE Communications Magazine*. He received the Excellent Graduate Supervision Award in 2006, the Outstanding Performance Award in 2004 and 2008 from the University of Waterloo, the Premier's Research Excellence Award (PREA) in 2003 from the Province of Ontario, Canada, and the Distinguished Performance Award in 2002 and 2007 from the Faculty of Engineering, University of Waterloo. He is a registered Professional Engineer of Ontario, Canada, and a Distinguished Lecturer of the IEEE Communications Society.