

Multicast with an Application-Oriented Networking (AON) Approach

Xiaohua Tian*, Yu Cheng*, Kui Ren* and Bin Liu†

*Department of Electrical and Computer Engineering, Illinois Institute of Technology
Chicago, IL, USA 60616, email: {xtian3, cheng, kren2}@iit.edu

†Department of Computer Science and Technology, Tsinghua University
Beijing, China, 100084, email: liub@tsinghua.edu.cn

Abstract—This paper proposes an efficient and scalable multicast scheme based on the concept of application-oriented networking (AON). The traditional IP multicast is bandwidth efficient but suffers from the scalability problem. The overlay multicast, proposed in recent decade, manages a data-dissemination tree at the application layer, and only utilizes unicasts among pairs of hosts; the overlay approach, however, usually incurs a considerable amount of redundant traffic. The essence of AON is to integrate application intelligence into the network. For AON-based multicasting, each packet will carry necessary explicit addressing information, instead of an implicit class-D group address, to facilitate the multicast data delivery. Each AON router will leverage the unicast IP routing table to compute necessary multicast copies and next-hop interfaces. The proposed AON multicast eliminates the need for constructing and maintaining the network-layer multicast routing table, while its bandwidth efficiency is very close to that of the IP multicast.

I. INTRODUCTION

The increasing popularity of multimedia applications over the Internet, e.g., on-line multiplayer games, IPTV, video conferencing, file-sharing, software updates, and grids, necessitate an efficient and scalable multicast mechanism to distribute shared data to a group of receivers. The traditional multicast solutions are implemented at the network layer, where the IP routers need to communicate with each other to construct and maintain a tree structure according to a distributed multicast routing algorithm [1]–[4]. Since the active members associated with a group are usually distributed in different administrative domains, the multicast forwarding entries in each router are not easily aggregatable, and grow linearly with the number of multicast groups being supported by the router [5]. Due to the implementation complexity, scalability issue, and some other technical and marketing reasons [5]–[7], the IP multicast has never been widely deployed.

The emergence of overlay networks provides an alternative multicasting approach, where trees or other delivery structures are constructed at the application layer [5], [8], [9]. Each link in the overlay network is an end-to-end logic connection between two end hosts. Overlay multicast is increasingly popular as the underlying unicast infrastructure does not need any modification. Nevertheless, overlay multicast can not perform as efficient as IP multicast in bandwidth utilization. It is not a rare case that separate overlay links pass through common physical links in the underlying transport network, that is, overlay multicast usually incurs redundant traffic.

In this paper, we reinvestigate the multicast problem by examining the fundamental design principles underpinning the existing solutions: while overlay multicast resorts to upper-layer intelligence, IP multicast implies a principle of allowing application specific processing within the network. In fact, enhancing network nodes with application intelligence has become one of the mainstream ideas to design the next-generation Internet [10], [11], stimulated by various applications, including firewalls, Web proxies/caches, mobile gateways, service-oriented architectures, in addition to the multicast. Cisco has already started to produce network devices with application intelligence to enhance the deployment, management, and integration of network applications, termed as *application-oriented networking* (AON) [12], [13]. However, as far as we know, all of the current AON¹ studies are directed to facilitate the upper-layer applications [11], [14], [15]. Our perspective is that AON also provides an opportunity to streamline the design of networking functionalities, with the focus of this paper being on multicast.

We propose an AON-based multicast approach, comprising the components of *membership management* and *forwarding protocol*. For membership management, the border router of a stub autonomous system (AS) domain is selected as a *designated router* (DR). The DR of a receiver-side AS domain is responsible for the membership discovery for the groups active in the domain, and forward the multicast request to corresponding source domains. The DR of a source domain will aggregate the information from those receiver-side DRs and generate the list of receiver-side DRs associated with each group managed by the source DR. The multicast forwarding is facilitated by the AON technique. Each multicast packet output from a source DR will carry the address information of related destination DRs. Each AON router will leverage the existing unicast routes to compute the necessary multicast copies and next-hop interfaces.

The proposed AON multicast has the following properties. 1) It eliminates the need for constructing and maintaining the network-layer multicast routing table. 2) The bandwidth efficiency of AON multicast is very close to that of the IP

¹In this paper, we enrich the denotation of AON compared to Cisco's definition: AON represents any new applications, services, or networking protocols that exploit the application intelligence built into the network; the researches in the AON context are termed as *application-oriented* studies.

multicast; although the DR address information carried in each packet incurs bandwidth overhead, which is significantly lower than that induced in the overlay multicast. 3) The membership management component is completely decoupled with the forwarding component; the complexity of multicast forwarding is totally independent of the number of active groups. We present extensive NS2 simulation results to demonstrate the performance of the proposed AON multicast approach, in comparison with the IP multicast and overlay multicast.

The remainder of this paper is organized as follows. In Section II, we give a more detailed review of the AON concept and present a generic architecture for an AON router. The AON-based multicast approach is described in Section III. Section IV presents the simulation results and Section V gives the conclusion remarks.

II. APPLICATION-ORIENTED NETWORKING

Integration of application intelligence into the network or allow application-specific computation in the network is not a brand new idea, which has been taken as an efficient approach, implicitly or explicitly, to implement some basic networking functionalities, develop new network protocols, or facilitate upper-layer applications. For example, the Domain Name Service (DNS) and Dynamic Host Configuration Protocol (DHCP) are fundamental networking functionalities, but implemented through application-layer processing. In recent decades, along with the global popularity of Internet and the proliferation of various multimedia and security applications, we have seen more and more application-specific nodes, e.g. Web proxies, multimedia gateways, wireless access gates, and firewalls, being inserted into the network. The IP multicast is an example of integrating application intelligence into routers. However, all of these application-oriented solutions had been implemented in an ad hoc manner.

Motivated by the necessity of allowing network to perform customized computations, the *active network* [11], [16] was proposed in the mid-1990s as a generic architecture to provision programmability within the network, instead of those ad hoc approaches. In an active network, packets are replaced with *capsules*, which are program segment (possibly with embedded data) executable by an active network node. The active network has never been widely deployed; the main reasons include the large bandwidth overhead of carrying programs, lack of a common capsule program language, and the security issue due to users' active control capability.

With the Internet evolving into an extremely complex system, developing new Internet applications by coding from scratch becomes inefficient and impractical. The service-oriented architecture (SOA) [17] is being widely adopted as a scalable approach for service creation according to a "find, bind and execute" paradigm. In SOA, complex services or applications can be assembled on demand by combining necessary service components. Such SOA-based service creation is implemented by exchanging messages, described in eXtensible Markup Language (XML), among service requesters, registries, and providers.

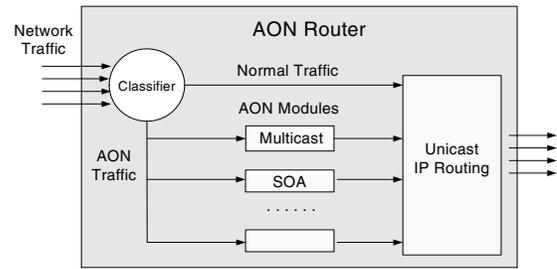


Fig. 1. A generic architecture for an AON router

The efficient implementation of SOA requires a powerful messaging backbone, which is one of the major motivations leading to the Cisco AON technology [12]. An AON-based network can transparently intercept the content and context of application messages, conduct operations on those messages according to business-driven policies and rules. The AON concept is closely related to the active network architecture; an XML message can be interpreted as a capsule capable of activating in-network processing. However, with AON, it is the network itself that determines what application-level processing capabilities to be offered within the network. By limiting users' control (available in the active network), the network-controlled approach not only facilitates the deployment of application-oriented capability in a coordinated manner, but also considerably improves the security level of the whole system.

The current AON studies focus on upper-layer applications [14], [15]. How to systematically exploit the AON capability to enhance the network functionalities is currently obscure in both industry and academia. In this paper, we initiate the study in this area by proposing an AON-based multicast scheme.

Before discussing the details of the multicast scheme, we first present a generic architecture for an AON router, as illustrated in Fig. 1. In an AON router, the incoming traffic will be first classified as *normal traffic*, which does not need application-level processing and is directly forwarded against the IP routing table, and *AON traffic*, which requires application-level processing before forwarded. The AON traffic will be further categorized and dispatched to different application-specific AON modules. We can select 1 bit in the IP header, e.g. one of the Type of Service (TOS) bits in the IPv4 header or one of the Traffic Class (TC) bits in the IPv6 header, to behave as the *normal/AON traffic indicator flag*. The flag is set to "1" for indicating the AON traffic. Although more TOS bits and TC bits may be used to further identify the AON modules, we prefer the fine-grained classification information to be carried in the payload, for higher scalability and flexibility.

III. AON-BASED MULTICAST

In this section, we first explain the multicast membership management scheme, and then present the multicast forwarding protocol. We focus on the single-source multicast for the illustration purpose; it is not difficult to see that the proposed

multicast scheme is also applicable to the case of multiple sources.

A. Membership Management

In the membership management scheme, one of the border routers of a stub AS domain is selected as a *designated router*. For convenience, we use RDR and SDR to denote the DR of a receiver-side AS domain and that of a source-side AS domain, respectively. The membership management functionalities of the RDR and SDR are different.

The RDR basically needs to implement the Internet Group Management Protocol (IGMP) [18] to discover the active groups, which have at least one member host, within its domain. The RDR maintains a *Group Hosts List* (GHL) to store the membership information. Each GHL record is a list of IP addresses of the member hosts associated with a group, indexed by the group address. We assume that there is a multicast address allocation protocol to allocate a unique class-D IP address to each group, which serves as the group ID in our multicast scheme. The group address allocation protocol is out of the scope of this paper. Each RDR will periodically update the membership information to upstream domains in the format as (IP address of RDR: group 1, group 2, ..., group n). Such RDR-group information dissemination may be piggybacked by inter-domain BGP advertisements [19].

The SDR of a source domain will aggregate the RDR-group messages it received and maintain a *Multicast Group List* (MGL). For each group with source node in the domain, the MGL establishes a record in the format as (group ID: RDR 1, RDR 2, ..., RDR n), where each RDR is indicated by its IP address. When the SDR receives a multicast packet (the destination address of which is the multicast group address or group ID) from a certain source node, it will set the normal/AON flag to "1", and insert into the payload the multicast AON-module classifier and the corresponding MGL record. The repacked multicast packets are then forwarded by the SDR to next-hop and further delivered to the member hosts in accordance with the multicast forwarding protocol.

We use an example as shown in Fig. 2 to illustrate the membership management. The active multicast group has a single data source and three active members $R1, R2$ and $R3$. The SDR is router A , and three RDRs are routers D, E and F , respectively. Each RDR maintains a GHL, with the GHL record in RDR F is illustrated in Fig. 2. The RDR-group messages from the three RDRs will be propagated via unicast to the SDR A , where the information is aggregated to a MGL as shown in the figure. In addition, all routers are assumed to be AON routers as shown in Fig. 1.

B. Multicast Forwarding Protocol

The multicast forwarding is facilitated by the AON technique. At each AON router, the normal/AON flag bit and the AON module classifier in the payload will direct the multicast packets to the multicast AON module.

We term the AS domain(s) that connect source and destination stub domains as *transit domain*(s). Fig. 2 illustrates

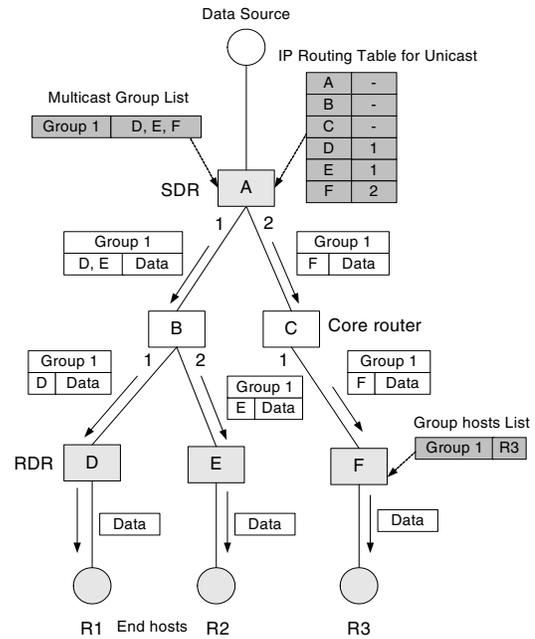


Fig. 2. AON multicast forwarding

the forwarding over a transit domain. For convenience, we converge the DRs with the transit-domain border routers to which they are directly connected. The multicast module of each transit-domain AON router will extract the MGL record from the multicast packet and replicate necessary copies for multicasting. With the list of destination RDRs available from the MGL record, the AON router will check its IP routing table to determine the output interface to each RDR and make necessary aggregation. As in Fig. 2, the IP routing table of the SDR A tells that the output interface 1 is on the path to both RDR D and RDR E , so that only one copy is necessary to be forwarded via interface 1. The IP routing table also shows that another copy should be forwarded via interface 2 to reach RDR F . When the input multicast packet is replicated and put onto each output interface, the MGL record attached with each copy is updated correspondingly to include only the destination RDRs that can be reached via that interface. For example, the MGL record in the packet delivered over A 's interface 1 includes only RDRs D and E . By removing unnecessary addresses from the MGL record, the bandwidth overhead can be reduced. The multicast module at each AON router will execute the same operations of aggregation, replication, and MGL record update, until one multicast packet reaches a RDR.

When an RDR receives a multicast packet, it will replace the MGL record with the GHL record for the corresponding group. As the GHL record includes the full list of IP addresses of the member hosts, the core AON routers in the stub-domain can follow the similar procedure of aggregation, replication, and GHL record update as that executed by the transit-domain AON routers to multicast the packet to each member host, by leveraging the unicast IP routing table.

The AON-based multicast mechanism has the following

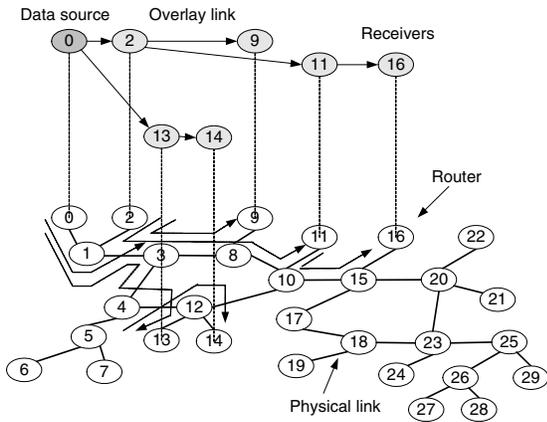


Fig. 3. Simulation topology

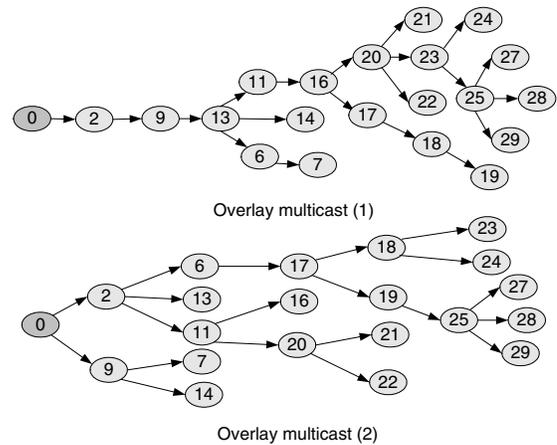


Fig. 4. Two overlay multicast trees

characteristics:

- 1) All the routers involved in the multicast forwarding, other than the DRs, do not need to maintain any status regarding multicasting. The forwarding complexity is totally independent of the number of groups to be supported, resulting in desirable scalability.
- 2) No new multicast routing protocol needs to be introduced. The existing intra-domain and inter-domain IP routing protocols are leveraged to implement the multicast. It is noteworthy that when the proposed multicast forwarding protocol is applied in the inter-domain scenario, redundant traffic may be incurred. The reason is that the path-vector nature of BGP allows a router to compute the shortest path from itself to a set of destinations but not from any source to those destinations [19]. However, the DVMRP protocols can be applied to efficiently suppress such inter-domain redundant traffic.
- 3) The membership management component and the multicast forwarding component are completely decoupled. This property enables the flexibility to develop heterogeneous application-specific membership management schemes over the same multicast forwarding protocol.
- 4) The cost incurred in the AON-based multicast is the bandwidth overhead, due to the AON classifier and the MGL/GHL record carried with each packet. The DR-based membership management scheme in fact implements a two-level addressing hierarchy, which can efficiently reduce the size of the MGL/GHL record. Anyway, further reducing the bandwidth overhead is one of our future topics.

IV. PERFORMANCE EVALUATION

In this section, we present some NS2 simulation results to demonstrate the performance of the proposed AON multicast scheme. For convenience, the illustration figures of the simulations results are generated with Matlab. The network topology for simulation is given in Fig. 3, which is similar to that used in [20]. As the essence of the proposed AON multicast is an

AON-based forwarding protocol, we here focus on examining the forwarding performance, assuming all membership records already established. Thus, all the nodes in Fig. 3 represent routers, and we simulate the multicasting from SDR to RDRs. All the link capacity is set as 1 Mbps, and UDP traffic is used in all the simulations.

We compare the AON multicast with the IP multicast, both in *dense mode* (DM) and *sparse mode* (SM), and the overlay multicast. The IP-DM multicast implements source-based trees with reverse path forwarding and pruning, similar in spirit to DVMRP [2]; the IP-SM multicast employs a core-based approach, constructing a tree rooted in a selected *rendezvous point* (RP) [4]. For overlay multicast, two overlay trees are constructed according to a random scheme [8], [21], as shown in Fig. 4. For each multicast mechanism, we simulate scenarios with different group sizes (i.e., number of RDRs). With node 0 being the SDR, the group of size 4 includes nodes {2, 6, 9, 13}. The node sets {7, 11, 14, 16}, {17, 18, 20, 23}, {19, 21, 24, 25}, and {22, 27, 28, 29} will be added in turn to form the groups of sizes 8, 12, 16, and 20, respectively.

A common measure to demonstrate the multicast bandwidth efficiency is *link cost* [5], which is defined as the total number of physical links that a multicast tree passes for data delivery. For example, the link cost for the overlay multicast tree illustrated in Fig. 3 is $2 + 5 + 4 + 5 + 3 + 2 = 21$. In the following, we further consider the measures of *bandwidth cost percentage*, *group receiving rate*, and *group average delay*, to examine the multicast performance in depth.

A. Bandwidth Cost Percentage

The bandwidth cost percentage (BCP) is defined as

$$BCP = \frac{T}{C \cdot D} \times 100\% \quad (1)$$

where T denotes the total number of bits traversing the physical network, C the total network capacity (i.e., the summation of all link capacities), and D the simulation duration. BCP is a measure of the bandwidth cost of multicast schemes. It is obvious that higher link cost will result in higher BCP.

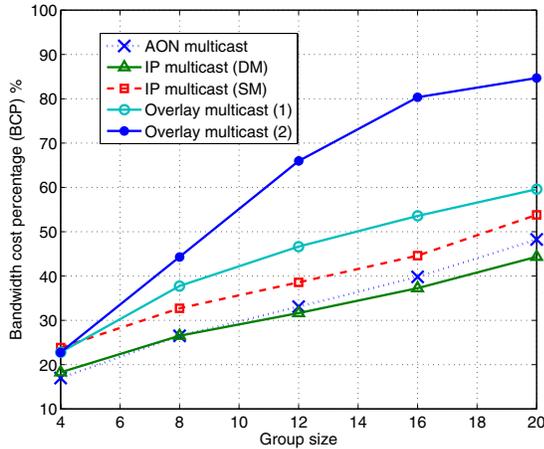


Fig. 5. Bandwidth cost percentage

The BCPs for different multicast schemes are shown in Fig. 5. We can observe that, for all kinds of group size, IP-DM multicast and AON multicast have the similar performance, lower than other schemes. It is not surprising that the overlay multicast incurs the highest bandwidth cost due to redundant traffic. Moreover, we can see that the bandwidth cost of the overlay multicast heavily depends on the structure of the multicast tree. The BCPs of all multicast schemes increase along with the group size, as the multicast will incur more traffic over the network when there are more receivers.

The reason for the bandwidth efficiency of AON multicast is that the AON forwarding protocol, as explained in Section III-B, in fact implicitly establish a source-based tree from the IP routing table, which is the same as the tree constructed by the reverse path forwarding used in the IP-DM multicast. Generally, the BCP of AON multicast is slightly larger than that of the IP-DM multicast, due to the bandwidth overhead induced by carrying AON classifiers and multicast address information within each packet. Moreover, the bandwidth overhead will increase when group size becomes larger, as demonstrated in Fig. 5. The exceptional case is for a very small group size 4, where AON multicast achieves the lowest BCP. The reason is that the tree pruning-messages involved in the IP-DM multicast also lead to bandwidth overhead, which exceeds that generated by the AON multicast.

In the IP-SM multicast scheme, We select node 15 as the RP according to a topology-based RP selection policy [22]; the RP has the smallest average distance (in terms of the number of hops) to other nodes in the network. In IP-SM multicast, data packets are first unicast from the sender to the RP, and then multicast over the tree. The traffic between the sender and the RP can be interpreted as bandwidth cost, leading to a higher BCP compared to IP-DM multicast and the AON multicast.

B. Group Receiving Rate

The group receiving rate (GRR) is defined as the summation of the data receiving rate at each RDR associated with the group, in terms of bits per second. In this experiment, we set

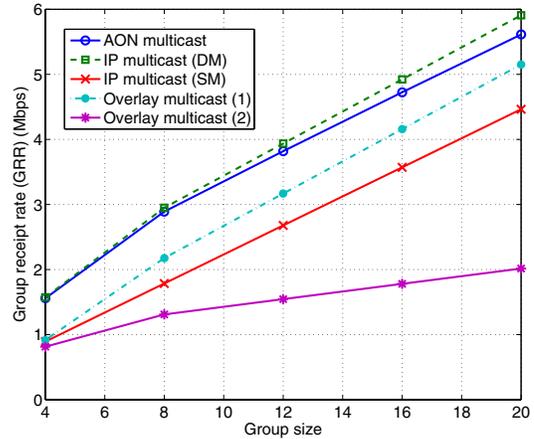


Fig. 6. Group receiving rate

capacities of link (1,3) and (10,15) to be a quarter of other links, with the objective to investigate the impact of bottleneck links on the performance of different multicast schemes.

Fig. 6 depicts the “GRR vs. group size” curves generated from simulations. In all the scenarios, the source node sends out data at the same rate, so the different GRRs achieved by different multicast schemes are due to the packet loss over the bottleneck links. Since the IP-DM multicast and the AON multicast schemes do not generate any redundant traffic, under these two schemes, the bottleneck congestion and associated packet loss are largely avoided, resulting in good throughput. The GRRs of two overlay multicast schemes are much smaller compared to IP-DM and AON schemes, as the redundant traffic incurred by the overlay results in more severe congestions at the two bottleneck links and therefore higher packet loss. One interesting observation is that the IP-SM multicast achieves a GRR even smaller than that of the overlay multicast (1). The reason is that the two bottleneck links are on the path between the source node and RP; therefore, the links seriously impact the throughput of the whole group. It is not difficult to see that such bottleneck impact applying through the RP will take effect independent of the group size; such a fact is indeed demonstrated in Fig. 6, where the GRR of the IP-SM multicast is strictly linear with the group size. This kind of perfect linear relationship is not available in all the other cases, where the packet loss is only due to a portion of the tree branches that pass the bottleneck links.

C. Group Average Delay

The group average delay (GAD) is defined as

$$GAD = \frac{\sum_{i=1}^N d_i}{N} \quad (2)$$

where d_i represents the delay from the data source to each group member (which is a RDR in our simulations), and N is the group size. The GAD can reflect the timing performance of a multicast scheme, which is important to real-time applications.

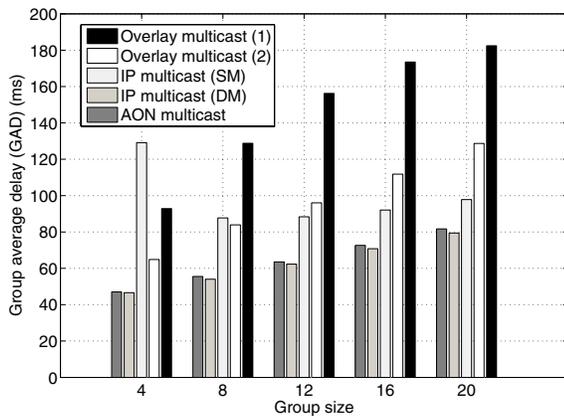


Fig. 7. Group average delay

We show the results for GAD in Fig. 7. The GAD of IP-SM multicast has a different behavior from that of other multicast schemes, which at first decreases versus the group size and then maintains relatively stable. Such a GAD behavior is determined by the position of RP and the distribution of group members. The first 4 group members $\{7, 11, 14, 16\}$ have a relative large distance from the RP node 15, so the GAD is high at first. The later joining groups members are closer around the RD, and they flatten out the GAD according to (2).

All the other multicast schemes adopt a tree rooted at the source node 0. The GAD grows versus the group size, since the later joining nodes are further away from the source node, leading to a larger delay. We can also see from Fig. 7 that bandwidth efficient multicast schemes also have better delay performance. The reason is that if a multicast scheme incurs redundant or overhead traffic, the higher traffic load normally will lead to a longer queue at each output interface. Fig. 7 shows that the bandwidth efficient IP-DM and AON multicast schemes have a significant advantage in the delay performance too. The AON multicast incurs a higher delay over IP-DM. This is because the processing overhead in the AON multicast scheme is longer than that in the IP multicast case; particularly in the current single-group scenario, the multicast routing table has only one entry. We expect that when multiple groups exist, the AON multicast will show advantages over the IP multicast, due to its group-independent forwarding design. One of our ongoing work is to examine the performance of the AON multicast scheme over a practical topology supporting a large number of groups.

V. CONCLUSION AND FUTURE WORK

In this paper, we propose a scalable and efficient multicast mechanism, by exploiting the network-embedded application intelligence or application-specific computation provisioned by the application-oriented networking architecture. By carrying necessary multicast address information within the payload, the AON router can leverage the existing IP routing information to deliver the data to member hosts without incurring any

redundancy. While the current AON research focus on upper-layer applications, our multicast study presented in this paper aims at attracting attentions to reinvestigate the fundamental networking functionality design, under the new application-oriented perspective.

REFERENCES

- [1] S. Deering, D. Estrin, D. Faranacci, V. Jacobson, C. Liu, and L. Wei, "The PIM architecture for wide area multicasting," *IEEE/ACM Trans. Networking*, vol. 4, no. 2, pp. 153–162, Apr. 1996.
- [2] D. Waitzman, C. Partridge, S. Deering, "Distance vector multicasting routing protocol," IETF RFC 1075, Nov. 1988.
- [3] A. Ballardie, "Core based trees (CBT) multicast routing architecture," IETF RFC 2201, Sept. 1997
- [4] D. Estrin *et al.*, "Protocol independent multicast-sparse mode (PIM-SM): protocol specification," IETF RFC 2362, Jun. 1998.
- [5] S. Fahmy and M. Kwon, "Characterizing overlay multicast networks and their costs," *IEEE/ACM Trans. Networking*, vol. 15, no. 2, pp.373–386, April 2007.
- [6] H. W. Holbrook and D. R. Cheriton, "IP multicast channels: express support for single-source multicast applications," in *Proc. ACM SIGCOMM*, pp. 65–78, Aug. 1999.
- [7] C. Diot, B. Levine, B. Lyles, H. Kassem, and D. Balensiefen, "Deployment issues for IP multicast service and architecture," *IEEE Network*, vol. 14, no. 1, pp. 78–88, Jan.-Feb. 2000.
- [8] Y. Chu, S. Rao, S. Seshan, and H. Zhang, "Enabling conferencing applications on the Internet using an overlay multicast architecture," in *Proc. ACM SIGCOMM*, pp. 55-67, Aug. 2001.
- [9] S. Y. Shi and J. S. Turner, "Multicast routing and bandwidth dimensioning in overlay networks," *IEEE J. Select. Areas Commun.*, vol. 20, no. 8, pp. 1444–1455, Oct. 2002.
- [10] D. Clark, *et al.* "New Arch: Future Generation Internet Architecture," Technical Report, 2003. [Online:] <http://www.isi.edu/newarch/iDOCS/final.finalreport.pdf>
- [11] D. L. Tennenhouse and D. J. Wetherall, "Towards an Active Network Architecture," *Proc. DARPA Active Networks Conference and Exposition*, pp. 2–15, 2002.
- [12] Cisco Systems, "Cisco application-oriented networking," white paper. [Online:] http://www.cisco.com/application/pdf/en/us/guest/products/ps6438/c1650/cdcont_0900aecd802c1f9c.pdf
- [13] T. Anthias and K. Sankar, "The network's new role," *ACM Queue*, vol. 4, May 2006. [Online:] <http://acmqueue.com/modules.php?name=Content&pa=showpage&pid=391>
- [14] Cisco Systems, "Cisco Application-oriented networking streamlines financial market-data and trade-order latency," Solution Overview. [Online:] http://www.cisco.com/en/US/products/ps6480/prod_brochure0900aecd804b0abe.html
- [15] Cisco Systems, "Cisco application-oriented networking facilitates intelligent radio frequency identification processing at edge," Solution Overview. [Online:] http://www.cisco.com/en/US/products/ps6480/prod_brochure0900aecd8032811f.html
- [16] D. L. Tennenhouse, J.M. Smith, W.D. Sincoskie, D.J. Wetherall, and G.J. Minden, "A survey of active network research," *IEEE Commun. Mag.*, vol. 35, no. 1, pp. 80–86, Jan. 1997.
- [17] J. Pasley, "How BPEL and SOA are changing Web services development," *IEEE Internet Computing*, vol. 9, pp. 60–67, May-Jun. 2005.
- [18] B. Cain, S. Deering, I. Kouvelas, B. Fenner, A. Thyagarajan, "Internet group management protocol, version 3," IETF RFC 3376, Oct. 2002.
- [19] S. Ratnasamy, A. Ermolinskiy and S. Shenker "Revisiting IP Multicast," in *Proc. ACM SIGCOMM*, September 11-15, 2006.
- [20] S. Egger and T. Braun, "Multicast for Small Conferences: A Scalable Multicast Mechanism Based on IPv6," in *IEEE Communication Magazine*, Jan. 2004.
- [21] Y. Chu, S. Rao and H. Zhang, "A case for end system multicast," in *Proc. ACM SIGCOMM*, pp. 1–12, Jun. 2000.
- [22] H. Wang, T. Sun, "A New RP Selection Algorithm Based on Delay in DiffServ Networks," in *Proc. of IEEE International Conference on Communication Technology*, pp. 1-4, Nov. 2006.