Title
Efficient Multi-Source Multicasting in Information Centric Networks

Author
Garcia-Luna-Aceves, J.J.

Publication Date
2015-01-09

Efficient Multi-Source Multicasting in Information Centric Networks

J.J. Garcia-Luna-Aceves
Palo Alto Research Center, Palo Alto, CA 94304
Computer Engineering Department, University of California at Santa Cruz, CA 95064

***Abstract*—Information Centric Multicasting (ICM) is introduced for the support of multicast groups with multiple sources in information centric networks. ICM supports two modalities: source-initiated multicasting and receiver-initiated multicasting. In contrast to all prior multi-source multicast approaches, routers do not have to flood the network from each multicast source, or know about the core or rendezvous point of a given multicast group. A multi-instantiated destination spanning tree (MIDST) is built to interconnect all sources or all receivers of a given multicast group. To disseminate content in a multicast group, receivers can send Interests to the nearest known source, or sources send content to the nearest receiver; either way, content flows over the MIDST to reach all receivers.**

I. INTRODUCTION

Several information centric network (ICN) architectures have been proposed [1], [3], [19] as alternatives to the current Internet architecture in response to the shift in Internet usage patterns from host-oriented communication to peer-to-peer and user-generated content.The goal of ICN architectures is to enable access to content and services by name, independently of their location, in order to improve system performance and end-user experience.

Interest-based ICN architectures [11], [17] integrate name resolution and content routing, and support content delivery based on Interests. Some of the routers (producers or caching sites) advertise the existence of local copies of named data objects (NDO) or name prefixes denoting a set of objects with names sharing a common prefix, and routes to them are established. Consumers of content issue content requests (called Interests) that are forwarded along the routes to the routers that issued the NDO or name prefix advertisements.

It has been argued that Interest-based ICN architectures [11], [17] provide "native support" for multicasting. However, this is the case only for single-source multicasting if the names in Interests denote the source of a multicast group. In some other ICN architectures, the names of data objects are mapped into addresses by means of directory servers, and then address-based routing is used for content delivery. Support for multicasting in these architectures is done as in the IP Internet.

Section II summarizes the prior work on multicasting in general and ICNs in particular. This summary reveals that very limited work has been reported on multi-source multicasting in ICNs to date. The proposals so far have been either based on the same approaches used in IP multicast or focus on single-source multicast.

Section III introduces *Information Centric Multicasting* (ICM), the first solution for efficient multi-source multicasting in ICNs based on name-based content routing.

ICM supports sender-initiated multicasting (SIM) or receiver-initiated multicasting (RIM) modalities. In the SIM modality, sources of a multicast group announce their presence and receivers of the group join the multicast group by sending Interests towards the nearest known source. In the RIM modality, receivers announce their presence and sources send content towards the nearest known receivers. In both cases a *multi-instantiated destination spanning tree* (MIDST) is used to ensure that all receivers of a multicast group receive content from all multicast group sources. ICM requires the use of signaling that is much more efficient than the signaling introduced in the past for shared multicast trees (e.g., [2], [13], [18]) or the spanning-tree approach for publish-subscribe signaling introduced for content-based networking (CBN) [4].

II. RELATED WORK

McQuillan [15] proposed the first link-state routing approach to support multicasting. In essence, each node floods link-state advertisements (LSA) stating the sate of adjacent links and the existence of receivers for different multicast groups. Given this information, each node can compute shortest routes to all receivers of each multicast group. Multicast OSPF [16] constitutes a more recent example of this approach.

Deering [7] described an approach to multicasting similar to what McQuillan introduced, and also introduced a sender-initiated approach to multicast routing based on distances. Various approaches can be used for flooding signaling packets from each source and for nodes with no receivers of a multicast group to send prune messages towards sources. Examples of this approach are ODMRP [12] and PIM dense mode [8].

The first approach for multicast routing that avoided the flooding from each multicast source or the flooding of information about those routers with attached multicast group receivers was Core Based Trees (CBT) introduced by Ballardi, Francis and Crowcroft [2], and many subsequent examples have been proposed (e.g., [8], [13], [18]). In CBT, a pre-defined node serves as the intermediary of a multicast group and is called the core of the group. Nodes maintain routes to all network nodes and hence to all possible cores, and learn the mapping from the multicast group address to the address of the core by some external means. Each receiver of a multicast group sends join requests towards the core of the group to establish a shared multicast tree spanning all the receivers and the core. Sources simply send data packets towards the core, and data packets are sent to all receivers of the multicast group over the multicast tree. Protocol Independent Multicast (PIM) [8] is similar to CBT, but the multicast trees are unidirectional; hence, sources must send multicast data to the intermediary node, called rendezvous point (RP), which then floods the data over the multicast tree.

The ICN architectures proposed recently have adopted two basic approaches to support multicasting: pull and push.

The pull-based approach is adopted in Interest-based ICN architectures [5], [17], in which paths from content consumers to content providers are established by sending content requests (called Interests) towards the content. Content Centric Networking [11] and named data networking [17] are the two main examples of Interest-based ICN approaches. Routers direct Interests towards content producers or caches base don forwarding information bases (FIB) populated by routing protocols. Routers along the paths traversed by Interests from consumers to content store the ingress and egress interfaces in which Interests were received in pending Interest tables (PIT), and data objects requested in the Interests are sent over the reverse paths traversed by the Interest. A router forwards an Interest for the same content only once and stores the ingress interfaces from which additional Interests for the same content are received.

It has been argued that NDN and CCN [11], [17] provide native support for multicasting, given that Interests are aggregated along trees directed towards a site with requested content. However, CCN and NDN can provide efficient support for single-source multicasting only, assuming that Interests state the name of the requested multicast source and that the presence of the multicast group source is advertised throughout the network. However, this approach does not work well for the case of large multi-source multicasting, because each multicast source must be known in the network and a tree to each such source needs to be maintained, which does not scale.

A number of ICN architectures that are not based on Interests adopt a push-baed approach to multicasting using mechanisms that are much the same as those introduced for PIM-SM [8]. A good example of this case is COPSS [6]. Users subscribe to content on a content descriptor (CD), which can be any legal content name, and each CD is associated with a Rendezvous Point (RP). The number of RPs may be as large as the number of ICN nodes. Routers maintain CD-based subscription tables to provide the same functionality as IP multicast, and COPSS supports sparse mode multicasting at the content layer. The RPS receive content from one or more publishers and send it over the multicast trees established by routers for the multicast groups.

III. ICM: INFORMATION CENTRIC MULTICASTING

The motivation for ICM is the lack of support for multi-source multicasting in Interest-based ICN architectures, and the inherent inefficiencies of supporting multi-source multicasting using the PIM-SM approach in push-based schemes for multicasting content in ICNs.

The operation of ICM assumes that: (a) each network node is assigned a name with a flat or hierarchical structure; (b) each multicast group can be requested by means of a unique name; (c) multicast group names (MGN) can be denoted using either flat or hierarchical naming, and the same naming convention is used for the entire system; and (d) a routing protocol operates in the network to provide each router with at least one route to the nearest instance of each multicast group advertised in the network. ICM establishes a MIDST (multi-instantiated destination spanning tree) for each name prefix that denotes a multicast group. Examples of routing protocol that provide the services required by ICM are DCR [9] and NLSR [14].

An instance of a multicast group can be a receiver or a source of the group. A router that advertises having an instance of a multicast group locally available is called an *anchor* of the group. Each anchor of a group originates routing updates for the multicast group periodically, and the update states the MGN, the name of the anchor, a distance to the multicast group, and a sequence number that only the anchor is allowed to change.

ICM can operate in two modalities: sender-initiated multicasting (SIM) and receiver-initiated multicasting (RIM). We assume that the name of the multicast group specifies the modality used for the group. The two components of ICM consist of the establishment and maintenance of a multi-instantiated destination spanning tree (MIDST) that connects all the anchors of a multicast group, and the use of the MIDST for each source to reach all receivers in SIM, or each receiver to be reached by all sources in RIM.

A. Information Stored and Exchanged

For convenience, we refer to a multicast group name as a name prefix or prefix in a number of cases, given that the routing protocol (e.g., DCR) maintains routes to name prefixes that can be multicast groups or not. We denote the lexicographic value of a name i by $|i|$, the set containing router i and its neighbor routers by N^i, and the set of next hops of router i for prefix j by S^i_j. The link from router i to router k is denoted by (i, k) and its cost is denoted by l^i_k. The cost of the link (i, k) is assumed to be a positive number that can be a function of administrative constraints and performance measurements made by router i for the link. The specific mechanism used to update l^i_k is outside the scope of this paper.

A router i uses the following information provided by the content routing protocol running in the ICN: (a) a *link cost table* (LT^i) listing the cost of the link from router i to each of its neighbors; (b) a *neighbor table* (NT^i) stating routing information reported by each neighboring router for each prefix or a topology table (TT^i) storing a link state for each physical link in the network and virtual links from the anchor of a prefix to the prefix itself ; and (c) a *routing table* (RT^i) that stores routing information for each known prefix. ICM builds a *multipoint routing table* (MRT^i) that stores routing information about routing trees created for multicast groups.

The entry in LT^i for link (i, k) consists of the name of neighbor k and the cost of the link to it (l^i_k).

The information stored in NT^i for each router $k \in N^i$ regarding prefix j is denoted by NT^i_{jk}, and consists of routing information for the nearest anchor and the root anchor of the prefix. The routing information for the nearest anchor reported by k consists of: the distance from neighbor k to j (d^i_{jk}); an anchor (a^i_{jk}) storing j; and the sequence number created by a^i_{jk} for j (sn^i_{jk}). The routing information for the root anchor of the prefix consists of: a root anchor (ra^i_{jk}); the distance from neighbor k to that anchor (rd^i_{jk}); and the sequence number created by ra^i_{jk} for j (rsn^i_{jk}). If prefix j is locally available at router i, then $a^i_{ji} = i$ and $d^i_{ji} = 0$. In this case, router i is its own nearest anchor for prefix j, but need not be the root anchor for j.

For simplicity of exposition, it is assumed that the information stated for NT^i is derived from TT^i in the case that the network runs a content routing protocol based on link-state information (e.g., NLSR).

The row for prefix j in RT^i specifies: (a) the name of the prefix (j); (b) the routing update information for prefix j (RUI^i_j); (c) the set of neighbors that are valid next hops (S^i_j); (d) a neighbor that offers the shortest distance to j ($s^i_j \in S^i_j$); and (e) an anchor list (A^i_j) that stores a tuple for each different valid anchor reported by any next-hop neighbor. Each tuple $[m, sn(m)] \in A^i_j$ states the name of an anchor m and the sequence number $sn(m)$ reported by that anchor.

RUI^i_j states: (a) the current distance from i to j (d^i_j); (b) the anchor of j that has the smallest name among those that offer the shortest distance to j (a^i_j); and (c) the sequence number created by a^i_j for j (sn^i_j).

The entry for prefix j in MRT^i specifies: (a) the name of prefix j; (b) the multipoint update information for prefix j (MUI^i_j); and (c) the list of neighbor routers that have joined the MIDST for the prefix ($MIDST^i_j$). MUI^i_j states the root anchor of j (ra^i_j), the distance to the root anchor (rd^i_j), and the sequence number created by ra^i_j for prefix j (rsn^i_j).

An update message sent by router i to neighbor m consists of the name of router i; a message sequence number (msn^i) used to identify the message; and a list of updates, one for each prefix that needs updating. An update for prefix j sent by router i is denoted by U^i_j and states: the name of the prefix j; the distance to j (ud^i_j); an anchor (ua^i_j); and the sequence number created by ua^i_j for prefix j (usn^i_j).

The entry for prefix j in the update message received from neighbor k by router i is denoted by U^i_{jk}. It states the prefix name j, a distance to it (ud^i_{jk}), an anchor for the prefix (ua^i_{jk}), and the sequence number assigned to the prefix by the anchor (usn^i_{jk}).

Lastly, router i updates the entry for j in MRT^i based on updates received from its neighbors and signaling messages exchanged among routers to join the MIDST of j.

B. *Establishing Multi-instantiated Destination Spanning Trees*

ICM supports routing to multicast groups by means of *multi-instantiated destination spanning trees* (MIDST). All the anchors of a given prefix corresponding to a multicast group are connected with one another through the MIDST for the prefix, which is rooted at the anchor of the prefix with the smallest name, which we call the *root anchor* of the prefix. The MIDST is established using routing updates exchanged only by routers located between the root anchor and other anchors of the same group.

The distance from router i to the root anchor ra^i_j is $rd^i_j = rd^i_{js} + l^i_s$, where $s \neq i$ is the next hop to ra^i_j selected by router i. If $i = ra^i_j$ then $rd^i_j = 0$. To build the MIDST for prefix j, routers select the root anchor of the prefix to be that anchor of the prefix that has the lexicographically smallest name; therefore, at each router i and for any neighbor $k \in N^i$, $|ra^i_j| \leq |ra^i_{jk}|$ and $|ra^i_j| \leq |a^i_{jk}|$.

The MIDST is established in a distributed manner. A router that knows about multiple anchors for a prefix other than the anchor it considers to be the root anchor sends updates about the root anchor along the preferred path to each of the other anchors it knows. Routers that receive updates about the root anchor send their own updates to their preferred next hops to each other anchor they know. This way, distance updates about the root anchor propagate to all other anchors of the same prefix. Updates about the root anchor propagate only to those routers in preferred paths between the root anchor and other anchors. If router i changes its routing information for the root anchor of prefix j, it schedules an update about its root anchor to each neighbor that satisfies the following condition.

Root-Anchor Notification Condition (RNC):
Router i sends an update with the tuple [ra^i_j, rd^i_j, rsn^i_j] to each router $k \in N^i - \{i\}$ for whom the following statements are true:

$$|a^i_{jk}| > |ra^i_j| \;\vee\; |ra^i_{jk}| > |ra^i_j| \tag{1}$$

$$\forall v \in N^i(a^i_{jv} = i) \;\vee\; \forall v \in N^i - \{k\}(\, a^i_{jk} \neq a^i_{jv} \;\vee \tag{2}$$
$$(\, d^i_{jk} + l^i_k < d^i_{jv} + l^i_v \;\vee\; [\, d^i_{jk} + l^i_k = d^i_{jv} + l^i_v \wedge |k| < |v| \,]\,)\,)$$

Eq. (1) states that k has not reported as its anchor or root anchor the same root anchor adopted by i. Eq. (2) states that i forwards the update about the root anchor to k if either i is an anchor and all its neighbors report i as their chosen anchor, or k is the lexicographically smallest next hop to an anchor that is not the root anchor.

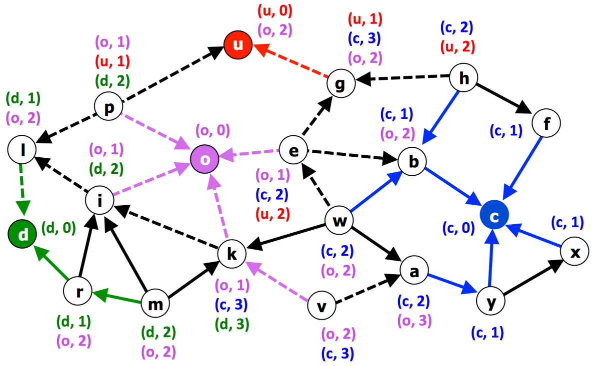

Fig. 1: Propagating root anchor information

Fig. 1 shows how information about the root anchor of a multicast group is propagated. The example assumes that DCR [9] is used in the ICN. There are four anchors for a given multicast group mg in this example, namely routers c,d, o, and u. These routers are anchors of mg, because they have attached sources or attached receivers of mg, depending on the multicasting modality of mg.

The arrowheads in Fig. 1 indicate the direction of interfaces between two routers according to their FIBs. An arrowhead with the color of an anchor indicates that the link is the preferred next hop (i.e., the next hop that is lexicographically smallest) for the anchor (e.g., the link (w, b) colored blue shows that b is the preferred next hop to c for w; and the link (m, r) colored green indicates that r is the preferred next hop to d for m). Black arrowheads indicate next hops to anchors that are not preferred next hops to anchors.

In the example, router c, which is shown in blue color, has the smallest name among all the anchors of the multicast group mg. Dashed arrowheads and blue arrowheads indicate those links over which updates about c being the root anchor of mg propagate in the direction away from c. For example, router b propagates a root anchor update to e because that neighbor is the best choice for b towards anchor o, and router g propagates a root anchor update to u because it is the best choice for anchor u. Router o propagates an update stating that c is the root anchor of mg to all its neighbors, because it is an anchor with all its neighbors reporting d as their preferred anchor.

Because of RNC, updates about the root anchor of a prefix reach all the other anchors of the prefix (see proof in [10]). However, we observe in Fig. 1 that routers (e.g., r and m) need not participate in the propagation of updates about c being the root anchor of mg, and may not even receive updates about mg with c as an anchor.

To allow anchors and relay routers to join the MIDST of a prefix (denoting a multicast group in our case), routers use the following condition to select their next hops towards the root anchor, and forward their requests to those neighbors. Eq. (3) states that the root anchor reported by k has the smallest name among all anchors of j known to i, and also reports an up-to-date sequence number from such an anchor. Eq. (4) states that k must offer the shortest distance to the root anchor among all neighbors. Eq. (5) orders router i with its selected next hop to the root anchor based on the distance to the anchor and the sequence number created by the anchor.

Root-Anchor Ordering Condition (ROC):
Router i can select neighbor $k \in N^i$ as its next hop to its root anchor for prefix j if the following statements are true:

$$|ra^i_{jk}| \le |ra^i_j| \ \wedge \ rsn^i_{jk} \ge rsn^i_j \tag{3}$$

$$\forall m \in N^i \ \left(\ rd^i_{jk} + l^i_k \le rd^i_{jm} + l^i_m \ \right) \tag{4}$$

$$rsn^i_j < rsn^i_{jk} \ \vee \ [\ rsn^i_j = rsn^i_{jk} \ \wedge \ rd^i_{jk} < rd^i_j \ ] \tag{5}$$

Based on ROC, much fewer routers and links are used in the signaling needed to establish the MIDST of group mg. Other than the links that are part of preferred paths to anchor c as the nearest instance of mg, only those routers along shortest paths between the root anchor and another anchor of mg may participate in the propagation of updates regarding the root anchor of mg. This is more efficient than the traditional approach to building shared multicast trees [2], [13], [18] or RP-based multicast trees [8], in which all routers must have routes to the root c, which also needs to be pre-defined.

To establish the MIDST of a multicast group, each anchor of a multicast group originates a join request and sends it to its lexicographically smallest next hop to the root anchor of the group. The join request is identified by the name of the multicast group and the name of the root anchor. Each router receiving and forwarding a join request stores an entry for the request denoting the neighbor from which it was received for a finite period of time. The join request traverses the path towards the root anchor of the prefix, until it reaches the root anchor or a router x that is already part of the MIDST of the multicast group. The response to the request traverses the reverse path of the join request and makes each router processing the response become part of the MIDST.

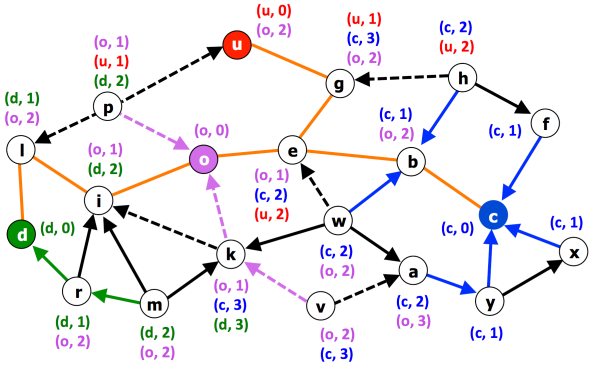

Fig. 2: The MIDST of a prefix

Fig. 2 shows the resulting MIDST for multicast group mg in the same example of Fig. 1. Anchors u, r, and o send their join requests towards d to join the MIDST of mg. The links that constitute the MIDST are indicated by solid orange lines in the figure, and multipoint data traffic for the prefix can flow in both directions of those links.

How the MIDST is used depends on the multicast modality adopted for the multicast group, which we describe next.

C. SIM: Source-Initiated Multicasting

In source-initiated multicast (SIM), the sources of a multicast group advertise their presence in the network and receivers attempt to join the group of sources. The anchors of a multicast group are the sources of the group.

Each source of a multicast group advertises its presence to an attached router using a Multicast Group Management Protocol (MGMP) implemented using Interest-based signaling or push-based signaling, depending on the ICN architecture in which it is used.

The MGMP message sent from a multicast source to its attached router states the name of the multicast group, the request to carry out SIM with the requester acting as a source of the group, and optional attributes. The attributes may include a delete timer informing the router the length of time that the source is to remain active. The name of the multicast group indicates the fact that the multicast group operates in the SIM modality. Given that the sources of a multicast group make attached routers advertise the presence of the group by name, the routers attached to multicast sources are the anchors of the multicast group.

Routers use a content routing protocol like DCR [9] to provide routing to the nearest sources of the multicast group, and ICM builds and maintains the MIDST of the multicast group as explained in the previous section. The resulting MIDST connects all sources of the multicast group.

To receive content from all the multicast sources of the group, a receiver sends simply an Interest to its attached router stating the name of the multicast group. In turn, the router sends an Interest towards the nearest known anchor of the multicast group. Content from the multicast group is delivered over the reverse path traversed by the Interest. A router attached to a source of the multicast group forwards content to each neighbor that submitted an Interest for the group, as well as each neighbor that is a member of the MIDST created for the multicast group. As a result, receivers are able to obtain content from all sources of a multicast group for which they stated an Interest.

The SIM modality of ICM can be used to implement a simple extension of the single-source multicast support available in NDN and CCN. Instead of each receiver having to send Interests for each multicast source, a receiver simply sends Interests requesting content from all the sources of multicast group, which is denoted by name in such Interests.

Requiring receivers to submit Interests for multicast content is not as efficient as defining a push-based approach [6]. ICM could be used more efficiently in CCN and NDN by defining long-term Interests that elicit multiple data objects.

D. RIM: Receiver-Initiated Multicasting

In receiver-initiated multicast (RIM), the receivers of a multicast group advertise their presence in the network and sources send content towards the nearest receivers of the group. The anchors of a multicast group in RIM are the receivers of the group.

Each receiver of a multicast group advertises its presence to an attached router using a Multicast Group Management Protocol (MGMP). The MGMP message sent from a multicast receiver to its attached router states the name of the multicast group, the request to carry out the RIM modality with the requester acting as a receiver of the group, and optional attributes. The attributes may include a delete timer informing the router the length of time that the source is to remain active. The name of the multicast group indicates the fact that the multicast group operates using the RIM modality. The routers attached to multicast receivers are the anchors of the multicast group.

Routers use a content routing protocol like DCR [9] to provide routing to the nearest sources of the multicast group, and ICM builds and maintains the MIDST of the multicast group using the nearest-instance routing information provided by the routing protocol as explained in Section III-B. The resulting MIDST connects all receivers of the multicast group.

To send content to the multicast receivers of the group, a source simply sends the content to its attached router, who in turn sends the content to the nearest anchor of the group based on the nearest-instance routing information. The first anchor of the group or router in the MIDST of the multicast group that receives the content broadcasts it over the MIDST. As a result, all receivers of the multicast group obtain the content from any one source.

IV. CONCLUSIONS

Existing Interest-based ICN architectures (NDN and CCN) do not support multi-source multicasting efficiently. We introduced the first approach to multi-source multicasting for ICNs that is not based on the traditional approach for IP multicasting, namely PIM-SM. The signaling required in ICM is more efficient than that of PIM-SM, and ICM can support different modalities to accommodate the nature of the multicast group.

ICM can be adopted in any ICN architecture in which content routing is supported. The two modalities of ICM can be implemented in NDN and CCN; however, introducing content-push mechanisms in such ICN architectures can make ICM far more efficient by eliminating the limitations discussed in [6]. In the RIM modality, long-term Interests can be used for receivers to subscribe to content from sources. In the SIM modality, multicast content could be sent from sources to the nearest anchor, without the need for Interests.

REFERENCES

[1] B. Ahlgren et al., “A Survey of Information-centric Networking,” *IEEE Commun. Magazine*, July 2012, pp. 26–36.

[2] T. Ballardie, P. Francis, and J. Crowcroft, “Core Based Trees (CBT),” *Proc. ACM SIGCOMM 93*, Oct. 1993.

[3] M.F. Bari et al., “A Survey of Naming and Routing in Information-Centric Networks,” *IEEE Commun. Magazine*, July 2012, pp. 44–53.

[4] A.Carzaniga et al., “A Routing Scheme for Content-Based Networking,” *Proc. IEEE Infocom ’04*, March 2004.

[5] Content Centric Networking Project (CCN) [online]. http://www.ccnx.org/releases/latest/doc/technical/

[6] J. Chen et al., “COPSS: An Efficient Content Oriented Publish/Subscribe System,” *Proc. ACM/IEEE ANCS 2011*, Oct. 2011.

[7] S. Deering, “Multicast Routing in Internetworks and Extended LANs,” *Proc. ACM SIGCOMM ‘88*, Aug. 1988.

[8] S. Deering et al., “The PIM architecture for wide-area multicast routing,” *IEEE/ACM Trans. on Networking*, Vol. 4, No. 2, April 1996.

[9] J.J. Garcia-Luna-Aceves, “Name-Based Content Routing in Information Centric Networks Using Distance Information,” *Proc. ACM ICN 2014*, Sept. 2014.

[10] J. J. Garcia-Luna-Aceves, “Routing to Multi-Instantiated Destinations: Principles and Applications,” *Proc. IEEE ICNP 2014*, Oct. 2014.

[11] V. Jacobson et al., “Networking Named Content,” *Proc. IEEE CoNEXT ’09*, Dec. 2009.

[12] S.J. Lee, M. Gerla, and C. Chiang, “On-demand Multicast Routing Protocol in Multihop Wireless Mobile Networks,” *Mobile Networks and Applications*, Vol. 7, No. 6, 2002

[13] E.L. Madruga and J.J. Garcia-Luna-Aceves, “Core Assisted Mesh Protocol for Multicast Routing in Ad-Hoc Networks,” U.S. Patent 6,917,985, 12 July 2005.

[14] A.K.M. Mahmudul-Hoque et al., “NSLR: Named-Data Link State Routing Protocol,” *Proc. ACM ICN ’13*, 2013.

[15] J.M. McQuillan, “Enhanced Message Addressing Capabilities for Computer Networks”, *Proceedings of the IEEE* , Nov. 1978.

[16] D. Medhi and K. Ramasamy, *Network Routing: Algorithms, Protocols and Architectures*, Morgan Kaufmann, 2007.

[17] NDN Project [online]. http://www.named-data.net/

[18] M. Parsa and J.J. Garcia-Luna-Aceves, “A Protocol for Scalable Loop-free Multicast Routing,” *IEEE JSAC*, April 1997.

[19] G. Xylomenos et al., “A Survey of Information-centric Networking Research,” *IEEE Communication Surveys and Tutorials*, July 2013.