Title

Efficient Multicasting in Content-Centric Networks Using Locator-based Forwarding State

Authors

Garcia-Luna-Aceves, J.J.
Mirzazad-Barijough, Maziar

Publication Date

2017

Efficient Multicasting in Content-Centric Networks Using Locator-based Forwarding State

J.J. Garcia-Luna-Aceves[1,2] and Maziar Mirzazad Barijough[2]
[1]Palo Alto Research Center, Palo Alto, CA 94304
[2]Computer Engineering Department, University of California, Santa Cruz, CA 95064
Email: jj@soe.ucsc.edu, maziar@soe.ucsc.edu

Abstract—**The Named Data Networking (NDN) and Content-Centric Networking (CCNx) architectures use a forwarding plane that requires large Forwarding Information Bases (FIB) listing the next hops to name prefixes and Pending Interest Tables (PIT) that maintain per-Interest forwarding state. We introduce CCN-RAMP (Routing to Anchors Matching Prefixes), a new approach to content-centric networking that substitutes the large FIBs and PITs used in NDN and CCNx with small forwarding tables listing anonymous sources of Interests and routers that announce name prefixes being local. The results of simulation experiments comparing NDN with CCN-RAMP based on ndnSIM show that CCN-RAMP requires forwarding state that is orders of magnitude smaller than what NDN requires, and attains smaller end-to-end delays in the dissemination of multicast content to consumers.**

I. INTRODUCTION

Several Information-Centric Networking (ICN) architectures have been proposed [3], [4], [25] aimed at improving the performance of the Internet by means of new ways to integrate name resolution (mapping of names to locations) and routing (establishing paths between locations) functions.

As we state in Section II, most architectures keep name resolution and routing independent of each other; however, the Named Data Networking (NDN) [16] and Content-Centric Networking (CCNx) [5] architectures allow consumers to request content objects (CO) or services by name by merging name resolution with routing. This is accomplished using three tables: A content store (CS) lists the COs that are cached locally; a Pending Interest Table (PIT) keeps forwarding state for each forwarded Interest (i.e., a request for a CO or service) processed by a router; and a forwarding information base (FIB) lists the next hops to known name prefixes. The inherent limitations with this approach are the need to lookup very large FIBs and PITs [7], [18], [21], [22], [24] and vulnerabilities to DDoS (distributed denial of service) attacks introduced by PITs [24].

A major selling point for CCNx and NDN has been that they provide "native" support for multicasting in the data plane with no additional signaling required in the control plane. Multicast receivers send Interests towards the multicast source. As Interests from receivers and previous-hop routers are aggregated in the PITs on their way to the multicast source, a multicast forwarding tree (MFT) is formed and maintained in the data plane. Multicast Interest are forwarded using the same forwarding information base (FIB) entries used for unicast traffic, and multicast data packets are sent using reverse path forwarding (RPF) over the paths traversed by aggregated Interests. Using PITs is appealing in this context, because it eliminates the need for complex multicast routing protocols operating in the control plane (e.g., [8], [17].

Fortunately, the benefits of NDN and CCNx can be attained without the complexity involved in using PITs or FIBs listing name prefixes. We have introduced **CCN-RAMP** (Routing to Anchors Matching Prefixes) [12], an approach to content-centric networking based on small forwarding tables with entries to anonymous sources of Interests and routers that announce name prefixes being local, which we call *anchors* of name prefixes. This paper focuses on multicast content dissemination using CCN-RAMP.

Section IV describes how CCN-RAMP supports multicast traffic in the data plane with no need for per-Interest forwarding state. Using the information disseminated in the name-based routing protocol operating in the control plane, a router builds and maintains two tables: A Forwarding to Anchors Base (FAB) listing the routes to anchors, and a Prefix Resolution Table (PRT) listing the anchors of each name prefix. A router receiving an Interest for content from a given multicast-group name from a local consumer (call it origin router) uses its PRT to bind the CO name to the nearest anchor for the name prefix that is the best match for the multicast-group name. To allow relaying routers to use only their FABs to forward Interests, an Interest states the name of the anchor chosen by the origin router. The origin router and other relaying routers establish multicast forwarding trees (MFT) rooted at the anchors of multicast-group sources, and forward the Interest as needed using their FABs and the anchor name in the Interest.

Section V presents the results of simulation experiments comparing the performance of NDN and CCN-RAMP under multicast traffic. The results show that CCN-RAMP attains even smaller end-to-end latencies than NDN in retrieving content; however, CCN-RAMP requires an average number of forwarding entries per router that is more than 150 times smaller than the number of PIT entries needed in NDN.

II. RELATED WORK

Excellent reviews exist of prior work aimed at making the forwarding planes of ICN architectures efficient [3], [4], [25] and we focus on key aspects of NDN and CCNx.

Such ICN architectures as DONA, PURSUIT, SAIL, COMET, and MobilityFirst implement name resolution and routing as independent functions. Name resolution servers are organized hierarchically, as multi-level DHTs, or along trees spanning the network [4], and consumers and producers contact such servers to publish and subscribe to content in various ways. Consumers obtain the locations of publishers from name resolution servers, and send their content requests to those locations to get the required content or services, and location-based routing is used to establish paths between consumers and subscribers or between resolution servers and subscribers or consumers.

A major limitation of keeping name resolution independent of routing stems from the complexity incurred in keeping name-resolution servers consistent with one another, and allowing consumers and producers to interact with the name-resolution system. Enabling the updates of name-to-address mapping is a non-trivial problem using hierarchical structures, spanning trees, or DHT-based organizations of servers [4].

In contrast to most ICN architectures, NDN and CCNx merge name resolution and routing. This eliminates the complexity of designing and maintaining a network of name-resolution servers. However, given that the name of a CO or service is bound directly to a route on a hop-by-hop basis, each router along the path traversed by an Interest must look up a FIB listing the known name prefixes. To operate at Internet scale, FIB sizes in NDN and CCNx are acknowledged to eventually reach billions of entries [20], which is orders of magnitude larger than the largest FIB size for the IP Internet today and is unattainable without further advances in technology [18]. Furthermore, none of the recent proposals for the reduction of FIBs listing name prefixes [2], [19], [20] ensure that the storage, communication, and time complexities of NDN and CCNx are comparable to or smaller than those of the current Internet forwarding plane.

In addition to the limitations associated with FIBs listing name prefixes, PITs grow linearly with the number of distinct Interests received by routers as consumers pipeline Interests or request more content, or more consumers request content [7], [21]. This is a problem, given that PITs do not deliver substantial benefits and can actually be counter-productive. We have shown [6] that the percentage of aggregated Interests is minuscule when in-network caching is used, even Interests exhibit temporal correlation. We have also shown [10], [11] that Interest aggregation combined with the Interest-loop detection mechanisms used in NDN and CCNx can lead to Interests being aggregated while traversing forwarding loops without such loops being detected. Furthermore, using PITs makes routers vulnerable to Interest-flooding attacks [22], [23], [24] in which malicious users can send malicious Interests aimed at making the size of PITs explode.

III. MULTICAST SUPPORT CCN-RAMP

We make a few assumptions in our description of CCN-RAMP; however, they should not be considered design requirements. A request for content from a local user is assumed to be sent to its local router in the form of an Interest. Routers know which interfaces are neighbor routers and which are local consumers, and forward Interests on a best-effort basis. The name of a multicast group j is denoted by $g(j)$ and the name prefix that is the best match for name $g(j)$ is denoted by $g(j)^*$. The set of neighbors of router i is denoted by N^i.

A. Information Exchanged and Stored

A multicast Interest forwarded by router k regarding group $g(j)$ is denoted by $MI[g(j), a^I(k), D^I(k), mv^I(k)]$, and states the name of the requested group ($g(j)$), the anchor selected by the first router processing the Interest ($a^I(k)$), the distance from k to the requested group ($D^I(k)$), and a multicast-counter value ($mv^I(k)$) used to pace the source of the multicast group.

A multicast data packet sent by router i in response to a multicast Interest is denoted by $MP[g(j), mv^R(i), sp(j)]$, and states, in addition to the CO, the name of the group ($g(j)$), a multicast-counter value ($mv^R(i)$), and a security payload ($sp(j)$) used optionally to validate the CO being sent.

An error message sent by router i in response to a multicast Interest is denoted by $ERR[g(j)$, $mv^R(i)$, $\mathsf{CODE}]$ and states the name of a multicast group ($g(j)$), a multicast-counter value ($mv^R(i)$), and a code (CODE) indicating the reason why the reply is sent. Possible reasons for sending a reply include: an Interest loop is detected, no route is found towards requested group, no content is found, and an upstream link is broken.

Router i uses three tables for multicast packet forwarding: A Prefix Resolution Table (PRT^i), a Forwarding to Anchors Base (FAB^i), and a Multicast Anonymous Routing Table ($MART^i$). If router i has local multicast-group members, it maintains a Group Membership Table (GMT^i), and maintains a Content Store (CS^i) if it provides content caching locally.

PRT^i is indexed by known name prefixes advertised by their anchors. Each entry of the PRT^i states the names of the selected anchors that advertised the prefix. Depending on the specific approach, the list may state the nearest anchors or all the anchors of the name prefix.

FAB^i is indexed by anchor names and each entry in FAB^i states available next hops to the anchor. The distance stored for neighbor q for anchor a in FAB^i is denoted by $D(i, a, q)$. This information is updated by means of a name-based routing protocol running in the control plane.

$MART^i$ maintains forwarding state to the receivers of multicast groups. $MART^i[g(j), a, mv, NH]$ denotes the entry for group $g(j)$ in $MART^i$ and denotes the multicast group name, the anchor a chosen by the origin router for the group, the current multicast-counter value mv for the group, and a list NH of next hops to the group of receivers who have sent multicast Interests for the group. Each item x of the entry in $MART^i$ for group $g(j)$ is denoted by $x^i[g(j)]$.

GMT^i lists the mappings of multicast group names to the lists of local receivers that requested to join the groups. If in-network caching is used as part of multicsting, the entry for group $g(j)$ also states a pointer $p[g(j)]$ to the content that has been cached for the group listing each CO by the value of the multicast counter used to retrieve COs for $g(j)$. The local receivers for group $g(j)$ listed in GMT^i is denoted by $GMT^i[g(j)]$.

B. Avoiding Forwarding Loops

CCN-RAMP eliminates forwarding loops by ordering the routers forwarding Interests based on their distances to anchors. To attain this, each Interest carries the distance to an anchor and FABs list the next hops *and* the distances to anchors. Let S_a^i denote the set of next-hop neighbors of router i for anchor a, router i uses the following rule to forwards Interests:

Multicast Anchor-Based Loop-Free Forwarding (MALF): If router i receives $MI[g(j), a^I(k) = a, D^I(k), mv^I(k)]$ from router k, it can forward $MI[g(j), a^I(i) = a, D^I(i), mv^I(i)]$ if:

$$\exists v \in S_a^i \ (\ D^I(k) > D(i, a, v) \ ) \tag{1}$$

MALF is based on the same approach we have verified to eliminate Interest looping in NDN and CCNx [11]. The proof that a forwarding rule similar to Eq. 1 ensures loop-free Interest forwarding is presented in [12].

C. Multicast Content Dissemination

Multicast content dissemination is based on the forwarding of Interests along multicast forwarding trees (MFT) to the sources of multicast groups, followed by the forwarding of multicast data packets on the reverse paths traversed by Interests. Forwarding multicast Interests is based on the information stored in the FABs maintained by routers.

We assume that routers with local consumers maintain caches of multicast content. The first content object (CO) of a multicast group is labeled by the name of the group and a multicast-counter value equal to one, and an empty entry for a multicast group is initialized with a multicast-counter value equal to zero. We assume that all initial requests to join a group state a multicast-counter value equal to one, and that forwarding state for a group stored in the MART of a router is deleted after a timeout if no Interests are received for the multicast group.

Algorithms 1 to 3 outline the steps taken by routers to process and forward multicast Interests, and return multicast data packets or replies for the case of real-time multicasting. To compare multicasting in CCN-GRAM directly with NDN, we assume pull-based dissemination of real-time multicast content, such that a single CO is sent to multicast receivers in response to an Interest sent to the source of a multicast group over the MFT.

For simplicity, we do not include the steps taken by routers to respond to the failures or additions of interfaces with neighbor routers or local consumers. Furthermore, we assume that Interests and responses to them are transmitted reliably between any two neighboring routers. Forwarding state related to a failed interface is deleted and the corresponding replies with negative acknowledgments are sent to previous next hops to remote receivers or local receivers as needed. Forwarding state associated with new interfaces is instantiated as a result of new Interests being forwarded.

Algorithm 1 shows the steps taken by router i to process requests for multicast content received from local consumers, which are assumed to be Interests stating the name of a group, the name of the consumer, and an empty distance to the content assumed to denote infinite. It is important to observe that name prefix cannot have a next hop in the FAB of a router if no entry for it exists in the PRT of the router.

The same format of data packets and replies used among routers is used to denote the responses a router sends to local consumers. Consumers increase the values of their multicast counters by one to request the next pieces of multicast content from multicast sources.

Router i adds consumer c as a local receiver in group $g(j)$ by adding an entry for $g(j)$ in GMT^i with c as a local receiver for the group, and indicates that it has local receivers in $MART^i$ by adding itself as a next hop towards receivers of the group. Router i forwards a single copy of a multicast Interest requesting more content from a multicast source independently of how many local receivers or neighbor routers send multicast Interests to router i. This is done by means of the multicast-counter value (mv) maintained by each router and multicast receiver, and the multicast-counter value field included in Interests and responses to them.

Algorithm 1 Processing Interest from user c at router i

```
function Interest_Source
INPUT: $GMT^i, PRT^i, FAB^i, MART^i$;
INPUT: $MI[g(j), a^I(c) = nil, D^I(c) = nil, mv^I(c)]$
if $g(j)^* \in PRT^i$ (% There is an anchor for $g(j)$) then
    if $MART^i$ entry for $g(j)$ does not exist then
        $a^i[g(j)]$ = nearest anchor listed in $PRT^i$ for $g(j)^*$;
        $mv^i[g(j)] = 0; NH^i[g(j)] = \emptyset$;
        create entry $MART^i[g(j), a, mv, NH]; GMT^i[g(j)] = \emptyset$;
    end if
    $GMT^i[g(j)] = GMT^i[g(j)] \cup c; NH^i[g(j)] = NH^i[g(j)] \cup i$;
    if $mv^I(c) \neq mv^i[g(j)] + 1$ then
        if $p[g(j)] \neq nil$ then
            retrieve CO for $mv^i[g(j)]$; $mv^R(i) = mv^i[g(j)]$;
            send $MP[g(j), mv^R(i), sp(j)]$ to $c$
        else
            $mv^R(i) = mv^i[g(j)]$; send $ERR[g(j), mv^R(i)$, Interest error$]$ to $c$
        end if
    else
        if $i$ is the source for $g(j)$ then
            $mv^R(i) = mv^i[g(j)]$;
            send $MP[g(j), mv^R(i), sp(j)]$
            to receivers in $GMT^i[g(j)]$ and next hops in $NH^i[g(j)]$
        else
            $mv^I(i) = mv^i[g(j)]; a^I(i) = a^i[g(j)]$;
            $D^I(i) = Min\{D(i, a^I(i), u)$ for $u \in S^i_{a^I(i)}\}$;
            for each $v \in N^i$ by rank in $FAB^i$ do
                if $D(i, a^I(i), v) = D^I(i)$ then
                    send $MI[g(j), a^I(i), D^I(i), mv^I(i)]$ to $v$; return
                end if
            end for
        end if
    end if
else
    $mv^R(i) = mv^i[g(j)]$; send $ERR[g(j), mv^R(i)$, no route$]$ to $c$
end if
```

A content consumer c asks to join a multicast group $g(j)$ as a receiver by sending an Interest $MI[g(j),\ a^I(c) = nil,$ $D^I(c) = nil, mv^I(c) = 1]$. If the value of the multicast counter for the group stored by the router processing the Interest is larger, the router responds with the latest multicast data packet corresponding to the current value of the multicast counter for the multicast group. A router sends a negative acknowledgment to an Interest from a local consumer with a multicast-content value different than the next expected value. This action forces a retransmission by the consumer and keeps all local consumers of the same multicast group using the same current value of the multicast counter. At the same time, it reduces end-to-end latencies incurred in delivering multicast content to consumers far away from group sources. A consumer requests more content from a multicast group by sending a multicast Interest after incrementing the value of the multicast counter for the group.

Using its PRT, a router processing a valid Interest from a local consumer selects the nearest anchor to the source of the multicast group $g(j)$, and forwards Interest $MI[g(j), a^I(i), D^I(i), mv^I(i)]$ towards that source based on the information in its FAB.

Algorithm 2 shows the steps taken by router i to process an Interest received from a neighbor router p. Router i follows similar steps to those in Algorithm 1 to respond to an Interest with a multicast data packet to the neighbor router if the content is local and the multicast counter in the Interest is smaller than the current value of the multicast counter at the router. If the Interest requests the next CO from the group and the group source is local, the multicast data packet is sent to all next hops along the MFT. If the multicast source is remote, router i forwards the Interest towards the anchor of the multicast group stated in the Interest using is FAB and ensuring that no forwarding loops occur according to MALF. The highest-ranked router satisfying MALF is selected as the successor for the Interest and router i. If no neighbor is found that satisfies MALF, an error message is sent stating that a loop was found.

Algorithm 2 Processing multicast Interest from router p at router i

```
function Interest_Forwarding
INPUT: $GMT^i, FAB^i, MART^i, MI[g(j), a^I(p), D^I(p), mv^I(p)]$;
if $a^I(p) \in FAB^i$ (% Route to anchor $a^I(p)$ exists) then
    if $MART^i$ entry for $g(j)$ does not exist then
        $a^i[g(j)] = a^I(p); mv^i[g(j)] = 0; NH^i[g(j)] = \emptyset$;
        create entry $MART^i[g(j), a, mv, NH]; GMT^i[g(j)] = \emptyset$
    end if
    $NH^i[g(j)] = NH^i[g(j)] \cup p$;
    if $mv^I(p) \neq mv^i[g(j)] + 1$ then
        $mv^R(i) = mv^i[g(j)]$;
        if $p[g(j)] \neq nil$ then
            retrieve CO for $mv^i[g(j)]$; send $MP[g(j), mv^R(i), sp(j)]$ to $p$
        else
            send $ERR[g(j), mv^R(i)$, Interest error$]$ to $p$
        end if
    else
        $mv^i[g(j)] = mv^i[g(j)] + 1$;
        if $i$ is the source for $g(j)$ then
            retrieve CO for $mv^i[g(j)]$; $mv^R(i) = mv^i[g(j)]$;
            send $MP[g(j), mv^R(i), sp(j)]$ to
            receivers in $GMT^i[g(j)]$ and next hops in $NH^i[g(j)]$
        else
            $a^I(i) = a^i[g(j)]; mv^I(i) = mv^i[g(j)]$;
            $D^I(i) = Min\{D(i, a^I(i), u)$ for $u \in S^i_{a^I(i)}\}$;

            for each $v \in N^i$ by rank in $FAB^i$ do
                if $D^I(p) > D^I(i)$ (% MALF is satisfied) then
                    send $MI[g(j), a^I(i), D^I(i), mv^I(i)]$ to $v$;
                    return
                end if
            end for
            $mv^R(i) = mv^i[g(j)]$; send $ERR[g(j), mv^R(i)$, loop$]$ to $p$
        end if
    end if
else
    $mv^R(i) = mv^i[g(j)]$; send $ERR[g(j), mv^R(i)$, no route$]$ to $p$
end if
```

Algorithm 3 outlines the processing of multicast data packets. If local consumers requested the content in the data packet, it is sent to those consumers based on the information stored in GMT^i. If the router has neighbor routers that are next hops towards remote receivers of the multicast group, router i forwards the data packet to all neighbors listed for $g(j)$ in $NH^i[g(j)]$ other than router i itself if there are local receivers. Routers take similar steps in the forwarding of replies to multicast Interests when retransmissions are done by consumers, i.e., routers simply forward replies back to the consumers along the MFT created by the forwarding of multicast Interests.

Algorithm 3 Processing multicast data packet from router s at router i

```
function Multicast Data Packet
INPUT: $GMT^i, MART^i, MP[g(j), mv^R(s), sp(j)]$;
[o] verify $sp(j)$;
[o] if verification with $sp(j)$ fails then discard $MP[g(j), mv^R(s), sp(j)]$;
if $NH^i[g(j)] \neq \emptyset$ then
    $mv^R(i) = mv^R(s)$; if $mv^i[g(j)] < mv^R(s)$ then $mv^i[g(j)] = mv^R(s)$;
    if $GMT^i[g(j)] \neq \emptyset$ (% router $i$ has local receivers in group $g(j)$) then
        for each $c \in GMT^i[g(j)$ do
            send $MP[g(j), mv^R(i), sp(j)]$ to $c$
        end for
    end if
    if $NH^i[g(j)] - \{i\} \neq \emptyset$ then
        for each $h \in NH^i[g(j)] - \{i\}$ do
            send $MP[g(j), mv^R(i), sp(j)]$ to $h$
        end for
    end if
    [o] store CO in local storage at $p[g(j)]$ indexed with $mv^R(i)$
else
    drop $MP[g(j), mv^R(s), sp(j)]$
end if
```

IV. EXAMPLE OF MULTICAST DISSEMINATION IN CCN-RAMP

Figure 1 illustrates the forwarding of multicast Interests and multicast data packets in CCN-RAMP. As the figure shows, router i maintains a forwarding table ($MART^i$) specifying the next hops to multicast receivers for each multicast-group name, and a table (GMT^i) listing the local receivers for each multicast-group name. In the figure, the entry for group $g(j)$ in $MART^i$ lists router i as a next hop, which indicates the presence of local receivers; the one local receiver (R_a) for group $g(j)$ is listed in GMT^i.

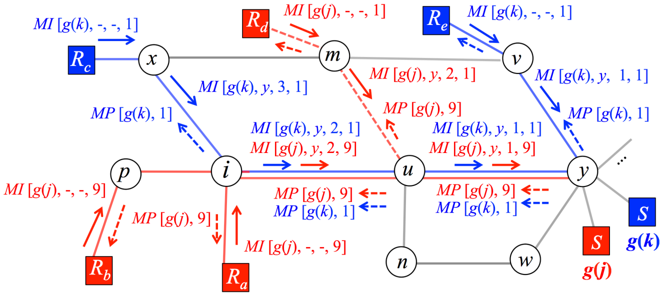

$MART^i$

group name, anchor	mv	next hops
g(j), y	9	i, p
g(k), y	1	x

GMT^i

group name	local receivers
g(j)	R_a
g(k)	-

Fig. 1. Native multicast support in CCN-RAMP

The entries in the MARTs and GMTs maintained by routers define the forwarding multicast trees of all multicast groups created in the network, and are established by the forwarding of multicast Interests, just as in NDN or CCNx. However, the use of multicast counters eliminates the need to maintain per-Interest forwarding state needed in NDN and CCNx. In the figure, the current state for the source of multicast group $g(j)$ is 9, and the state for multicast group $g(j)$ is 1. Receiver R_d joins group $g(j)$ late, and router u sends back the most current multicast packet with $mv = 9$. The application may be satisfied with that reply, or a separate off-line mechanism can be used for the receiver to retrieve the prior COs.

V. PERFORMANCE COMPARISON

We compare the forwarding state and end-to-end delays for multicast traffic in CCN-RAMP and NDN using simulations based on the ndnSIM simulation tool [1]. The implementation of CCN-RAMP is based on the algorithms presented in this paper, and NDN implementation from ndnSIM is used without modification. We use the AT&T network topology, which is considered to be a realistic topology for simulations [13]. This topology includes 153 nodes and 184 point-to-point links with 30 ms delay. To reduce the effects derived from sub-optimal implementations of CCN-RAMP and NDN, we set the data rate of point-to-point links to 10Gbps. Using on-path caching strategy, each router in these experiments can cache up to 1000 content objects. In each simulation scenario, there are multiple multicast groups, and each group contains multiple consumers and one producer. Consumer and producer nodes for each group is selected at random from the routing nodes in the network.

We consider the impact that the Interest request rate, the number of multicast groups, and the size of multicast groups have on performance. We consider minimum download rate of 1.5 Mbps for audio/radio streaming, 5 Mbps for HD video streaming, and 25Mbps for Ultra HD video streaming. Considering the standard packet size of 4KB advocated in NDN, we compared different scenarios with constant rate of 50 to 800 interests per second from each consumer application.

A. Size of Forwarding Tables

Figure 2 shows the results of a simulation experiment that includes 20 multicast groups, each with 20 consumers and one producer.

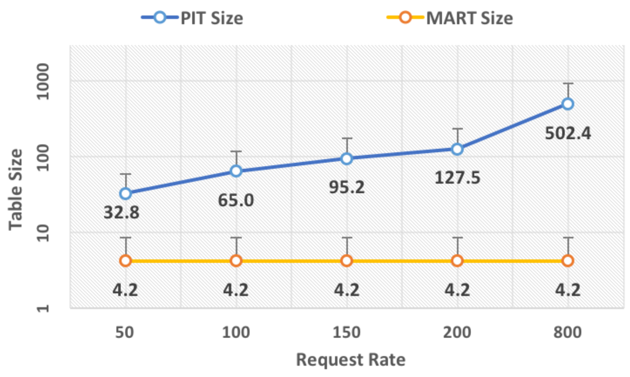

Fig. 2. Average size of forwarding tables for varying request rates

The above figure shows the average size of a forwarding table in logarithmic scale as a function of Interest (request) rates. Given that CCN-RAMP adds a single entry per multicast group, the number of entries in a MART is independent of the rate at which Interests arrive at routers. By contrast, the size of PITs in NDN is a function of the rate at which Interests arrive at routers. The maximum MART size for this scenario is smaller than 10, and the average size of a MART table is 4.2 independently of the request rate. As the figure shows, the number of PIT entries is highly affected by the Interest rates from consumers. The average PIT size for a request rate of 800 Interests per second is 502.4

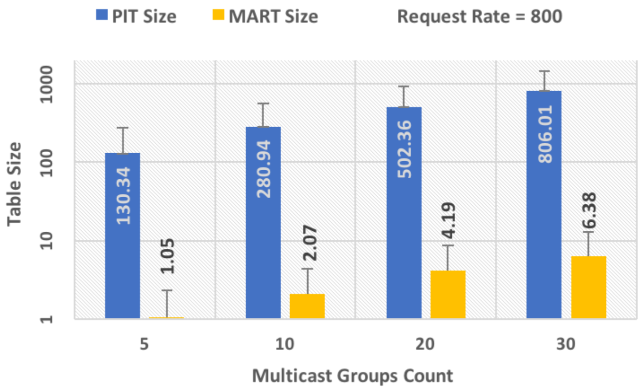

Fig. 3. Average size of forwarding tables vs. number of multicast groups

Figure 3 shows the average size of MARTs and PITs for varying number of multicast groups from 5 to 30 groups. Each multicast group has 20 consumers with Interest (request) rate of 160 Interests per second, which is enough to support HD video streaming with each data packet being 4KB. In CCN-RAMP, the number of entries of MART tables cannot exceed the total number of multicast groups. On the other hand, as the figure shows, the number of entries in the PIT of a router is directly related to the number of interests received by the router, which in turn depends on the number multicast groups and the request rate per group. Accordingly, the average PIT size can grow dramatically.

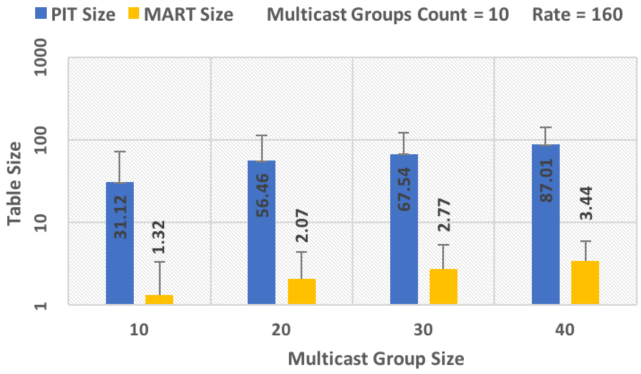

Fig. 4. Average size of forwarding tables for varying size of multicast groups

Figure 4 shows the average size of PITs and MARTs for varying multicast group sizes from 10 to 40 consumers per group. As the size of a multicast group increases, more routers become involved in forwarding multicast Interests and multicast data packets in both NDN and CCN-RAMP. This results in larger average sizes of both PITs and MARTs. However, the grow rate for NDN is higher because of entries are added to PITs on a per Interest basis.

B. Average Delays

As Figure 5 shows, the average delay for CCN-RAMP is shorter than the delays incurred in NDN for the dissemination of multicast content. The delay variance is not shown but is smaller in CCN-RAMP than in NDN.

According to Algorithm 3, the first multicast Interest received by a producer results in the multicasting of data toward current members of the multicast group in CCN-GRAM, even if the Interest from a member or previous-hop relay in the multicast forwarding tree has not been received yet. On the other hand, in NDN, if one consumer node is far from the producer compared to other consumer nodes such that its interests is not aggregated with the same interests from other consumers, request of that node for a multicast data will be processed separate from other group members, which results in lower throughput and higher delays. The operation of NDN could be modified to mimic the way in which CCN-GRAM forwards multicast data, in which case end-to-end latencies could be similar.

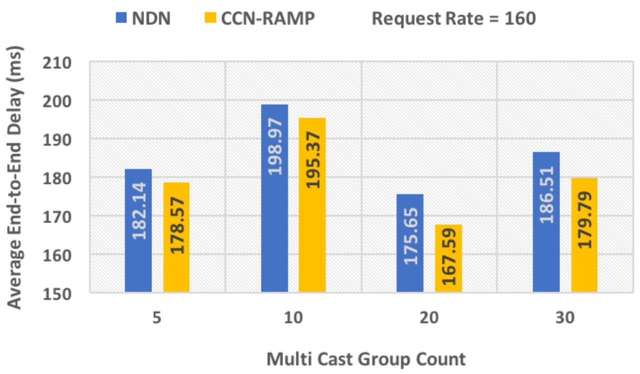

Fig. 5. Average end-to-end delay for varying number of multicast groups

VI. CONCLUSIONS

We presented a multicast approach for content-centric networks based on CCN-RAMP that eliminates the need to maintain per-Interest forwarding state, uses small forwarding tables, and still operates based on the forwarding of Interests in the data plane, without the need for a multicast routing protocol in the control plane.

The results obtained in simulation experiments clearly indicate that maintaining per-Interest forwarding state using PITs to support multicast content dissemination is not needed, and that a far simpler and efficient forwarding strategy can be used based on the anchors of multicast groups. The forwarding tables needed in CCN-RAMP are orders of magnitude smaller than the PITs and FIBs needed in NDN.

Additional work is needed to define the mechanisms needed to enable multicasting from multiple sources over a single multicast forwarding tree, rather than requiring a separate tree for each multicast source as it is currently assumed in NDN and CCNx and described in this paper for CCN-RAMP. More generally, a detailed analysis of push-based and pull-based multicast content dissemination is also needed.

REFERENCES

[1] A. Afanasyev et al., “ndnSIM: NDN simulator for ns-3”, *University of California, Los Angeles, Tech. Rep*, 2012.

[2] A. Afanasyev et al., “SNAMP: Secure Namespace Mapping to Scale NDN Forwarding,” *in Proc. IEEE Global Internet Symposium ‘15*, 2015.

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

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

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

[6] A. Dabirmoghaddam et al., “Characterizing Interest Aggregation in Content-Centric Networks,” *Proc. IFIP Networking 2016*, May 2016.

[7] H. Dai el al., “On Pending Interest Table in Named Data Networking,” *Proc. ACM ANCS ‘12*, Oct. 2012.

[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 ‘14*, Sept. 2014.

[10] J.J. Garcia-Luna-Aceves, “A Fault-Tolerant Forwarding Strategy for Interest-based Information Centric Networks,” *Proc. IFIP Networking 2015*, Toulouse, France, May 2015.

[11] J.J. Garcia-Luna-Aceves and M. Mirzazad-Barijough, “Enabling Correct Interest Forwarding and Retransmissions in a Content Centric Network,” *Proc. ACM ANCS ‘15*, May 2015.

[12] J.J. Garcia-Luna-Aceves and M. Mirzazad-Barijough, “Content-Centric Networking at Internet Scale through The Integration of Name Resolution and Routing,” *ACM ICN ‘16*, Kyoto, Japan, Sept. 26-28, 2016.

[13] O. Heckmann et al., “On Realistic Network Topologies for Simulation,” *Proc. ACM SIGCOMM MoMeTools ‘03*, Aug. 2003.

[14] E. Hemmati and J.J. Garcia-Luna-Aceves, “A New Approach to Name-Based Link-State Routing for Information-Centric Networks,” *Proc. ACM ICN ‘15*, Oct. 2015.

[15] V. Lehman et al., “A Secure Link State Routing Protocol for NDN,” Technical Report NDN-0037, Jan. 2016.

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

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

[18] D. Perino and M. Varvello, “A Reality Check for Content Centric Networking,” *ACM ICN ‘11*, 2011.

[19] T.C. Schmidt et al., “Let’s Collect Names: How PANINI Limits FIB Tables in Name Based Routing,” *Proc. IFIP Networking ‘16*, May 2016.

[20] T. Song et al., “Scalable Name-Based Packet Forwarding: From Millions to Billions,” *Proc. ACM ICN 2015*, Sept. 2015.

[21] M. Varvello et al., “On The Design and Implementation of a Wire-Speed Pending Interest Table,” *Proc. IEEE Infocom NOMEN Workshop ‘13*, April 2013.

[22] M. Virgilio et al., “PIT Overload Analysis in Content Centric Networks,” *Proc. ACM ICN ‘13*, Aug. 2013.

[23] M. Wahlisch et al., “Lessons from the Past: Why Data-driven States Harm Future Information-Centric Networking,” *IFIP Networking ‘13*, May 2013.

[24] M. Wahlisch et al., “Backscatter from The Data Plane? Threats to Stability and Security in Information-Centric Network Infrastructure,” *Computer Networks*, Vol. 57, No. 16, Nov. 2013.

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