

Achieving Scalable Push Multicast Services Using Global Name Resolution*

Shreyasee Mukherjee, Francesco Bronzino, Suja Srinivasan, Jiachen Chen and Dipankar Raychaudhuri

WINLAB, Rutgers University, North Brunswick, NJ 08902, USA

Email: {shreya, bronzino, sujirini, jiachen, ray}@winlab.rutgers.edu

Abstract—This paper presents a novel approach to achieving scalable push multicast services using the distributed global name resolution service associated with emerging name-based network architectures. The proposed named-object multicast (NOMA) scheme employs unique names to identify multicast groups while using the global name resolution service (GNRS) to store the tree structure and maintain current mappings to mobile end-user addresses. The NOMA scheme achieves improved scalability and performance over conventional multicast protocols such as PIM-SM and MDSP by taking advantage of the GNRS to simplify tree management and limit control overhead. Performance evaluation results including comparisons with IP multicast are given using a combination of analysis and NS-3 simulation. The results show good scalability properties along with low control overhead for medium to large multicast groups. In addition, NOMA seamlessly handles mobility for end-hosts subscribed to a group, avoiding data losses upon mobility events. Results further demonstrate how separating names from addresses enables NOMA to dynamically forward traffic to mobile users. In conclusion, we describe a proof-of-concept prototype developed for further experimental validation of the proposed NOMA multicast routing scheme.

I. INTRODUCTION

Internet applications like video streaming, online gaming and social networks, e.g. Twitter, often require dissemination of the same piece of information to multiple consumers at the same time. While multicast routing protocols have long been available, most of these applications rely on unicast based solutions that exploit overlay networks aimed at improving the efficiency of pushing the required data without support from the network. Recent increases in network traffic associated with the growth of mobile devices, Internet-of-Things (IoT) devices, smart wearables and connected vehicles, motivate the need for efficient push multicast, a service that is not well-addressed through overlay solutions. Consider for example IoT based messaging scenarios: a typical query involves sending short messages to hundreds or thousands of users or application agents, so that scalability becomes an issue, as multiple unicast messages through an overlay service can easily overload the network. Mobility of end-devices results in additional complexity, especially for dynamic environments such as vehicular communications. For example, if a single warning message needs to be pushed to hundreds of cars and pedestrians in a given area, multicast groups would need to be maintained across a large number of access networks in order to efficiently support such one-to-many communication.

*Research supported under NSF Future Internet Architecture - Next Phase (FIA-NP) grant CNS-1345295

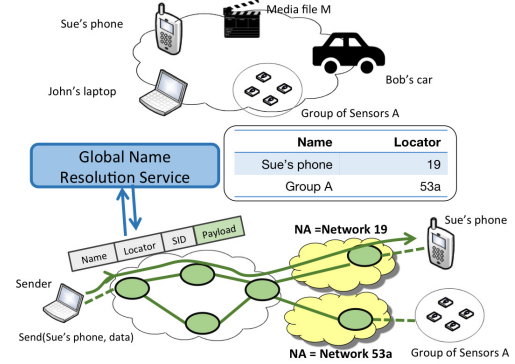


Fig. 1. Named object abstraction with clean separation of naming and addressing

Using appropriate multicast routing solutions would help, however, existing network-layer multicast solutions (e.g., PIM-SM [1], MOSPF [2]) have not been widely adopted. Firstly, these solutions cause packet loss during end-host mobility and large amount of distributed control traffic is generated to modify the multicast tree structure. The problem becomes particularly acute for applications like Twitter where each end-host might have more than 100 groups to join each time it moves. Secondly, extending these protocols to inter-domain has achieved mixed results, with issues of scalability and coordination across domains [3]. For example, in the case of IP multicast based on PIM-SM [4] domains are often unwilling to have rendezvous points (RPs) for their local groups to be maintained in other domains. Multicast group address assignment may require a separate protocol altogether, such as the Multicast Address-Set Claim (MASC) protocol used in conjunction with BGMP [5]. All of these problems have further negative consequences for highly dynamic environments. For example, in the vehicular use-case previously described, group membership changes rapidly with vehicular mobility along with the context of data-delivery. An accident or traffic-alert push-notification to a group of cars in NJ Turnpike is such an example. Table I describes a sample set of application scenarios that require efficient multicast primitives and their characteristics.

Application layer solutions for multicast have also been explored in this context; works like SCRIBE [6] and ZIGZAG [7] sought to find scalable and efficient solutions by building an overlay among the receivers. These solutions do address mobility and inter-domain management issues, but due to the lack of topology awareness, they may incur high levels of network

Application	Multicast Type	Group Size	Group Flux	Group Longevity	Data Flow Size
IoT commands	Push	1000's	Hours	Days	KB-MB
Accident notification	Push	100's	Seconds	Minutes	KB
Twitter	Pull	100's of 1000	Minutes	Months	KB-MB
IPTV	Pull	1000's	Relatively static	Months	GB
Multiplayer games	Push/Pull	100's	Hours	Hours	GB

TABLE I
EMERGING MULTICAST APPLICATION AND THEIR CHARACTERISTICS

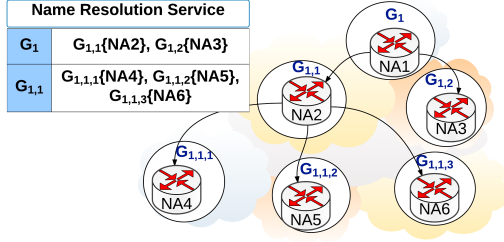


Fig. 2. Hierarchical tree structure maintained in a name resolution service, with names of tree nodes recursively mapping to routable addresses

traffic. In addition, forcing the end hosts to replicate packets, instead of dedicated routers results in heavy workload on the end hosts, which may have intrinsic power and computation constraints.

Based on the above considerations, a native network layer multicasting solution is identified as an important goal for future networks which are increasingly required to support many-to-many communication modes. We propose a solution based on named objects and a dynamic name-resolution service for mapping names to routable network entities. Separating names (identities) from addresses has been advocated by the research community [8], [9] for quite some time and has inherent benefits in handling mobility and dynamism for one-to-one communication. But they also provide additional advantages by facilitating creation of new service abstractions that can be used to design solutions for multicast services. Names can be used to represent many different Internet objects; for example, a cell-phone, a person, or a group of devices, as shown in Fig. 1; the latter perfectly applies in the context of multicast to define participation of end-hosts. Moreover, new entities can be integrated within these names, not being constrained to end points; through this, we gain the ability to directly name network entities, such as routers that implement the multicast routing protocols.

We exploit names to design a Named Object Multicast (NOMA) solution which relies on separation of names and addresses using a globally distributed, logically centralized name resolution service, similar in spirit to an evolved DNS. In NOMA each multicast group is identified by a unique name across all domains, thus separating routing logic from group management. NOMA takes advantage of the dynamic name resolution service to manage the tree, using name recursion, to store the tree topology. This is achieved by mapping unique names assigned to participating routers to their children nodes, as shown in Fig. 2. Data forwarding is then performed using tunnels between participating nodes; end-to-end information

is preserved within the packet, while the information globally available in the name resolution service is used to identify next hops allowing for quick branching and replicating decisions. Finally, dynamicity of mobile environments is handled by decoupling the participants name from their location through the resolution service and periodically recomputing the multicast tree.

The remainder of the paper provides the details of our design and performance evaluation of the proposed scheme which include:

- The design of NOMA architecture that leverages use of names and global name resolution service to manage multicast routing protocols;
- An efficient centralized tree-construction mechanism that minimizes the network traffic with relatively low computational overhead; and,
- Large-scale simulations to demonstrate the reliability, efficiency and scalability of NOMA design even when there is node mobility.

II. NOMA DESIGN

NOMA aims to achieve efficient multicast communications through the employment of a logically centralized, globally distributed name resolution service associated with name based communications. A Global Name Resolution Service (GNRS) is utilized as a network-wide entity that provides an API for inserting and querying mappings between unique name identifiers and a set of values which can include network addresses, other name identifiers and related parameters – e.g. node properties, past locations and more. In spirit, this service is very similar to current Internet's DNS, which has already been effectively applied for new service functions such as load balancing and service replication. Even more interesting services can be realized with the next generation of global name resolution services such as DMap [10] and Auspice [11] introduced recently.

NOMA's design, as proposed here, is based on MobilityFirst (MF [9]), which is a clean-slate network architecture for the next-generation mobile network where DMap [10] is used to provide resolution of names, that are Globally Unique Identifiers (GUIDs), into routable network addresses (NAs). Moreover, MF incorporates a hybrid name-address forwarding scheme, in which routing components use availability of both names and addresses in packet headers to perform forwarding decisions. Note that even though NOMA is based on MF, the same design concept can be applied to IP extensions (such as

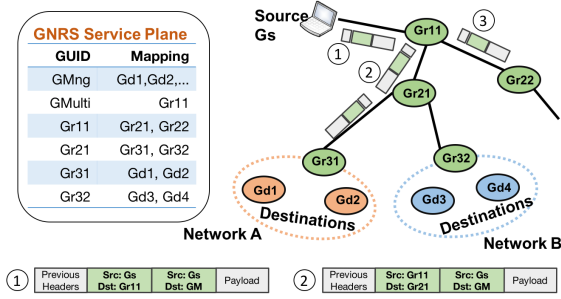


Fig. 3. Multicast architecture overview

HIP [8]), overlay protocols (such as SCRIBE [6]), or clean-slate ICN protocols such as NDN [12] and XIA [13] through the use of a similarly designed name resolution service.

A. Multicast Tree Management

Multicast management consists of two core operations: membership of destination nodes and building and management of multicast trees. Both operations can be streamlined by exploiting the name resolution service (GNRS); in particular by using two forms of name indirection. First, a unique name (*GMng* in Fig. 3) is assigned to perform the task of node membership; all entities interested in receiving data from the multicast flow, can request to join by inserting their own unique name into the corresponding mapping in the table. This information is then exploited close to the source by a multicast service manager, which builds an efficient tree based on the available resources and the size of the required tree. Recursive mappings are then used to express the tree graph: by assigning to each branching router a name that exclusively identifies it in the context of the given multicast flow, we recursively follow the tree structure. For example, in Fig. 3, the root of this tree is identified by mapping the unique name of the multicast group to the first branching router ($GMng \rightarrow Gr11$); this router then maps to its children in the tree ($Gr11 \rightarrow \{Gr21, Gr22\}$) and so on until the leaves of the tree are reached, where we identify the leaves as the destination nodes. As time progresses and destinations join or leave the multicast group, the service manager can rebuild the tree and update the information in the GNRS.

One of the novelties of NOMA is that it can support push mode of multicast, where a source can send a single packet of multicast data, without the knowledge of the tree and this can happen even before the tree has been built. On receiving a multicast packet, for a group *GMulti*, the gateway router at the source domain, acting as the multicast service manager, will do a membership query to the GNRS. GNRS supports recursive queries that return the host GUIDs along with the NAs of the domains they are currently connected to. Having the service manager on the gateway enables the tree computation to be topology-aware, as unicast path information of the NAs is available at the gateway, which is then used to build the tree. Once a tree is computed, it is updated in the GNRS such that downstream nodes do not need to recompute the tree again. This is quite different than distributed tree management techniques used in IP multicast since NOMA does not require

flooding of multicast control messages (for example, source active (SA) or Join messages in PIM-SM and MSDP [4]) across domains. The latter limits the scalability for traditional multicasting techniques to small to medium groups, as shown later in Sec. III. Also, using unique names to represent a group, members of the group, and, the multicast tree eliminates the need of a separate address allocation protocol, similar to MASC required for BGMP [14]. For evaluation purposes, we focused on two categories of multicast tree computation algorithms, i.e. shortest path trees (SPTs) and Steiner trees. A constraint of having centralized computation of trees is complexity and hence we opted for SPT and its modifications, even though our design is not limited to any specific algorithm.

B. Data Forwarding

Once the multicast tree is established, data forwarding can exploit the information contained in the GNRS to efficiently flow through edges between the nodes of the tree. In order to do so, we exploit address encapsulation, where two pieces of information are carried in data packets at the same time: Internally (i.e. second field in the green packets in Fig. 3), the encapsulated information carries the source and destination of the multicast flow, providing valuable information usable by all nodes along the path to easily identify data streams. Externally, routing information to perform hop-by-hop forwarding from one branching node to the next is placed. At each branching node participating in the multicast, forwarding decisions are performed by querying the GNRS to obtain information on how many next hops it has to forward to, generating required duplicates and replacing the external routing information with the new hop source and destination; this process is exemplified in the figure, where node *Gr21* generates 2 duplicates for its two children, replacing headers accordingly. Intermediate nodes along the path forward encapsulated packets based on normal unicast rules. This reduces complexity of multicast packet processing to only a subset of nodes of the tree. To reduce the need of continuously involving the GNRS in the forwarding procedure, mappings can be cached at each hop, avoiding traffic and computational overhead. The tradeoff for this approach comes at the cost of slower reaction times to tree change events. More details on how to handle tree restructuring and end points mobility is provided in the following section.

C. Handling Mobility:

End host mobility support has been a challenging problem for both unicast and multicast delivery. For the latter, the situation is further aggravated by the fact that an end-host mobility can significantly alter the multicast tree and hence its efficiency of delivery to other connected end-hosts. Without a clean separation of names and addresses, the onus of *re-booting* an ongoing *session* falls on to the mobile end-host. For an inter-domain multicast delivery, this means that every time an end-host moves and changes its point of association, it needs to send an explicit join at the new point of connectivity. The router at the new domain will then need to join the multicast tree, before the end-host can receive any data.

Meanwhile, following a best-effort delivery policy, all the data received at the previous point of association will be lost.

For NOMA on the other hand, failure in delivery to a downstream node results in temporary storage of data packets (MF routers are storage-capable [15]) and re-querying the GNRS for an up-to-date downstream node name (GUID) to its address (NA) mapping. This is specially relevant for the leaves of the tree which could be mobile end-hosts. As mentioned earlier for a long-lived flow tree, restructuring takes place periodically and any mobility that happens at a faster time-scale than tree re-computation will suffer. In order to ensure that end-hosts do not lose packets while moving, NOMA supports encapsulated ‘repair’ packets to be sent to the moving client. This again is enabled by the GNRS that maintains the up-to-date location (end-host GUID to NA mapping) as it moves. As shown in Fig. 4, when a end-host *D1* moves from *NA14* to *NA11*, which is not part of the multicast distribution tree, the tree does not change immediately. However, failure to deliver at the edge, causes the gateway router at *NA14* to query the GNRS for up-to-date mapping of *D1*. Following association at *NA11*, the gateway at *NA14* can encapsulate the pending data and send it as unicast repair to *NA11* as shown. In contrast to multicasting, the repair procedure is transparent to an end-host or application and does not require explicit re-joining from the client side. However this is only a short-term mechanism to counter moderate mobility of a subset of destinations. With increase in the number of devices and mobility, the frequency of tree updates should increase proportionally.

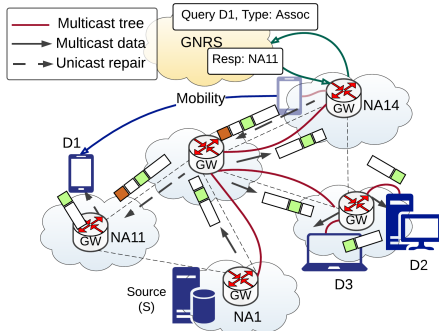


Fig. 4. Device mobility handling through unicast repair messages

III. EVALUATION

In this section we present detailed performance evaluation based on a combination of large scale analytical modeling and fine-grained packet-level simulation on network simulator (ns-3).

A. Tree Generation Algorithms

NOMA provides a framework for managing and deploying multicast communications, independently from the tree generation algorithm employed. While this is a valuable feature of the design, it is necessary to study different algorithms and heuristics in the context of choosing one that can effectively utilize unicast routes, and is lightweight enough to be potentially run at a single router. We looked at two main

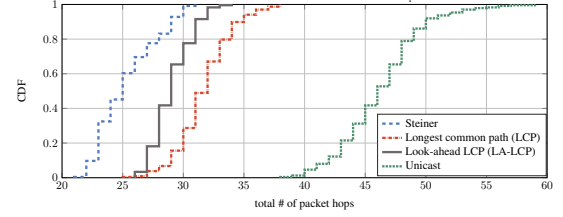


Fig. 5. CDF of packet hops for a 100 node random graph with 20 randomly placed destination nodes

categories of algorithms for building multicast trees, namely shortest path trees (SPTs) and Steiner trees. Although Steiner trees provide an optimal solution in terms of overall network resource utilization, they are NP-hard to compute. Several Steiner heuristics have been proposed over the years to provide near-optimal solutions [16], with relatively high computation cost. However, computational complexity is a key constraint for our design, since the tree computation is *centralized*. We instead opt for the SPT algorithm that uses inter-domain unicast route information and require no further computation, but is less efficient compared to a Steiner tree. In SPT, packets are forwarded along the *longest-common path* (LCP) to all the destinations, as a single copy, until the branching point is reached, where the packet is copied and delivered to multiple destinations. This allows all destinations to receive multicast packets across the shortest path from the source. We also analyzed other heuristics that aimed to further minimize the overall network traffic with moderate computation. One of these heuristics is the *look-ahead longest-common path* (LA-LCP) algorithm. Unlike LCP, which branches whenever there is a divergence of shortest paths to multiple destinations, LA-LCP, compares the overall network cost of branching from the current node and branching from each of the possible next hops, and decides to branch downstream if the cost is lower for the latter, thereby deviating from the SPT. This reduces the overall packet hops in the network, with slight increase in computation complexity.

Fig. 5 plots the CDF of total packet hops to reach 20 randomly placed destinations from a single source on a 100 node *Erdős-Rényi* random graph for each of these algorithms. As seen from the plot, all the multicast algorithms are much more efficient than unicast. Although Steiner provides the most efficient trees, it is computationally intensive. In comparison, LA-LCP provides *reasonable* performance with lower overall network overhead compared to traditional longest common path.

B. Comparison to IP multicast

In this section we compare pull-based multicast of NOMA with IP based inter-domain multicast, namely, PIM-SM standard coupled with MSDP [4]. Through the results we highlight two key benefits of using NOMA, namely, 1) lower control overhead for maintaining a multicast group, and 2) better handling of mobility for data forwarding. Note that BGMP [14] is another prominent inter-domain IP multicast scheme, however, it is not well-suited for applications that involve dynamism and fast changes in the tree, and hence has not been a focus of

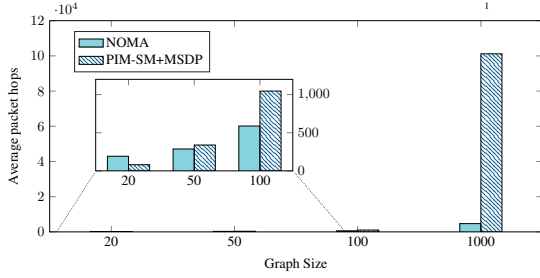


Fig. 6. Control packet overhead for tree setup for varying graph sizes

our evaluation. BGMP allows multicast route updates to be carried along with inter-domain BGP messages and therefore tree changes occur at a much slower time-scale than PIM-SM/MSDP (typical BGP updates take about 100 seconds to propagate throughout the network [17]).

- Control overhead:** The advantage of using unicast routes to build the tree is that no multicast specific control overhead needs to be exchanged across networks. This is crucial for inter-domain settings where flooding periodic multicast tree update messages is not tractable. In Fig. 6 we plot the multicast specific messages exchanged for setting up a tree and forwarding packets for increasing graph sizes, with the topology being an *Erdős-Rényi* random graph, and 50% of the nodes being randomly chosen to have destination clients part of the multicast group. For NOMA the overhead includes 1) GNRS insert messages from each of the destination networks for joining a particular multicast group; 2) GNRS insert from the gateway at the source domain to insert the generated multicast tree; and, 3) GNRS query and responses during data forwarding at the branching nodes. The GNRS is implemented as a distributed hashmap, following the DMap design [10], with the same mapping stored at multiple locations. For evaluation purposes, 3 GNRS instances were maintained, therefore each insert incurred 3 unicast messages to 3 specific nodes (determined by a hash function), whereas each query was *anycasted* to the nearest of the 3. In comparison, for PIM-SM+MSDP the overhead numbers comprise of, 1) Source-Active (SA) messages from the source domain to all other domains; and, 2) Join messages from the domains which have destinations nodes interested in the multicast group. As seen from the plot, maintaining a multicast tree in the GNRS has higher overhead for smaller sized graphs (for example, for a 20 node topology, shown in the zoomed in section of Fig. 6), but it scales elegantly with size. Using PIM-SM+MSDP, on the other hand, becomes intractable as the number of nodes increases. With more than 40 thousand ASes in the Internet today, if every domain was multicast enabled, the cost becomes too high to maintain a distributed tree.

- Handling mobility:** NOMA seamlessly handles client mobility and the dynamism in tree-changes thereof, by periodically recomputing the tree and updating the corresponding GNRS entries. In addition, to counter packet-loss due to mobility, NOMA supports unicast ‘repair’ packets to be sent from the previous to the current point of attachment of a mobile client, until a tree update restructures the tree. We performed detailed packet level simulations in network-simulator (ns-

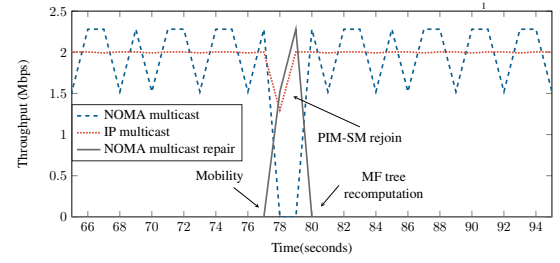


Fig. 7. Average multicast throughput received at a client with mobility

3) on a 20 domain random topology with randomly placed mobile and static clients, for both NOMA and an IP multicast implementation of PIM-SM + MSDP. Fig. 7 plots the fluctuation in received throughput at a client receiving a multicast stream of 2Mbps on the event of mobility. A mobility event is characterized by disconnection of a client from its attachment point and re-association to another node, following a period of association (uniform random variable $U(0,1)$ seconds), as highlighted in the figure at about 77 seconds. NOMA periodically restructures the multicast tree every 10 seconds for this scenario, whereas, IP multicast restructures following the client explicitly joining the tree at the new point of association. Therefore, multicast traffic for NOMA falls to 0, until tree is restructured at $t = 80$ seconds. However, repair packets are delivered to counter packet loss and reordering, highlighted by the black trajectory in the figure. Note that NOMA is based on MobilityFirst transport, that uses reliable hop-by-hop delivery of large chunks, and the throughput received by the client is therefore in steps with the average being 2Mbps. In comparison, for IP multicast, data throughput falls following temporary disconnection and re-connection, as shown by the red dotted trajectory.

Mobility not only affects the instantaneous throughput at a client, but it also leads to loss of packets during the interval of disconnection, re-association of the client, re-joining and re-structuring of the multicast tree. Additionally, in a practical setting, for IP multicast, the mobile client will spend a significant amount of time for new IP address allocation through DHCP, which has not been taken into account for this evaluation. This packet loss and reduction in overall throughput is highlighted in Fig. 8 where we plot the aggregate throughput at a mobile client for increasing rates of mobility, that moves randomly with exponential random mean mobility interval of 50, 20 and 10 seconds. As seen from the plot, aggregate throughput for NOMA does not change with mobility, primarily due to native features of MF such as hop-by-hop reliable delivery and storage-capable routers to handle temporary disconnections. In comparison, IP multicast throughput significantly worsens with increasing mobility speeds.

C. Prototype Description

To validate the implementation feasibility of NOMA, we built a Click software router prototype, based on the one described in [18] and tested it on a small scale topology on the ORBIT testbed [19]. Fig. 9 highlights the key router and multicast host components that were built for the prototype.

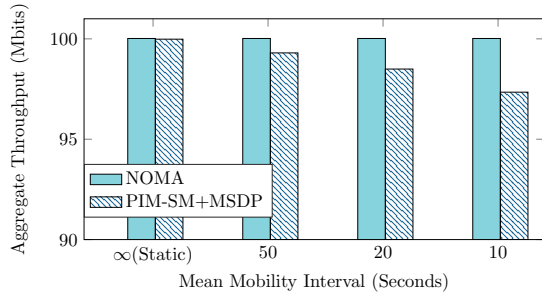


Fig. 8. Aggregate throughput at a mobile client, with increasing mean mobility rates; mobility event is an exponential random variable

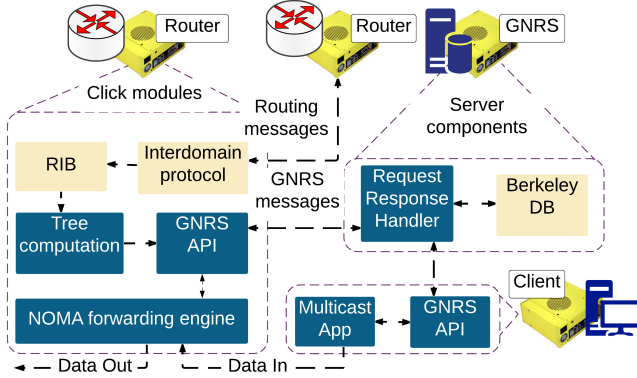


Fig. 9. Components of the NOMA router prototype, GNRS and client implementation, with developed modules shown in blue

In addition, existing DMap based GNRS APIs were modified to allow multicast tree insertion and queries. Ongoing work includes detailed evaluation on larger topologies to validate scalability of NOMA in realistic scenarios.

IV. RELATED WORK

Work on multicast started in the early 1990's when Deering proposed DVMRP [20]. This was followed by multicast extensions to OSPF called MOSPF [2]. However, most of the early work on multicast concentrated on flood-and-prune based techniques that were difficult to scale. Core-based trees [21] and PIM-SM [1] were thus proposed that introduced shared trees and rendezvous points (RPs). However, even these were mostly deployed within a single domain. Challenges in choosing appropriate RPs for inter-domain settings, led to the multicast source discovery protocol (MSDP) that allowed RP's in different networks to coordinate [4]. MSDP suffered from scaling issues that limited its deployability. BGMP is the multicast extension for BGP for inter-domain routing [5], but requires its own address assignment and address allocation protocols for multicast group management [14]. In this context, overlay solutions such as SCRIBE [6] or ZIGZAG [7] offer better scalability. However, such application layer solutions overload end-hosts with multicast tree computation and active management and also result in higher network traffic load. Clean slate network layer solutions have also been proposed, notably COPSS [22] for content delivery. However, the focus of COPSS is on pub/sub based systems that require pull based multicast. In comparison, NOMA allows both push and pull mechanisms to support a wide variety of applications.

V. CONCLUSION

In this paper, we have proposed a name based inter-domain multicast approach, leveraging on a distributed name resolution service for membership and tree management. The proposed NOMA framework scales reasonably well to medium to large scale trees and handles client mobility with disconnections. Large scale analytical results for management overhead and fine-grained packet-level simulations for mobility scenarios were provided. In addition, we presented a proof-of-concept prototype with small-scale experiments as feasibility studies. Future work includes further feasibility studies and deploying NOMA on the GENI large scale testbed to evaluate performance in more realistic inter-domain network scenarios.

REFERENCES

- [1] D. Farinacci, C. Liu, S. Deering *et al.*, "Protocol independent multicast-sparse mode (PIM-SM): Protocol specification," 1998.
- [2] J. Moy, "RFC 1584-Multicast Extensions to OSPF," *SRI Network Information Center*, 1994.
- [3] C. Diot, B. Levine, B. Lyles *et al.*, "Deployment issues for the IP multicast service and architecture," *Network, IEEE*, vol. 14, no. 1, pp. 78–88, 2000.
- [4] D. Meyer and B. Fenner, "Multicast source discovery protocol (MSDP)," 2003.
- [5] S. Kumar, P. Radoslavov, D. Thaler *et al.*, "The MASC/BGMP architecture for inter-domain multicast routing," *ACM SIGCOMM CCR*, vol. 28, no. 4, pp. 93–104, 1998.
- [6] M. Castro, P. Druschel, A. Kermarrec, and A. Rowstron, "SCRIBE: A Large-Scale and Decentralized Application-Level Multicast Infrastructure," *JSAC*, pp. 1489–1499, 2002.
- [7] D. Tran, K. Hua, and T. Do, "ZIGZAG: An Efficient Peer-to-Peer Scheme for Media Streaming," in *INFOCOM*, 2003.
- [8] R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson, "RFC 5201–Host Identity Protocol," 2008.
- [9] I. Seskar, K. Nagaraja, S. Nelson, and D. Raychaudhuri, "MobilityFirst Future Internet Architecture Project," in *ACM AINTEC 2011*.
- [10] T. Vu, A. Baid, Y. Zhang *et al.*, "Dmap: A shared hosting scheme for dynamic identifier to locator mappings in the global internet," in *IEEE ICDCS*, 2012.
- [11] A. Sharma, X. Tie, H. Uppal *et al.*, "A global name service for a highly mobile internetwork," in *ACM SIGCOMM*, 2014.
- [12] V. Jacobson, D. Smetters, J. Thornton *et al.*, "Networking Named Content," in *ACM CoNEXT*, 2009.
- [13] A. Anand, F. Dogar, D. Han *et al.*, "XIA: An Architecture for an Evolvable and Trustworthy Internet," in *HotNets*, 2011.
- [14] P. Radoslavov, D. Estrin, R. Govindan *et al.*, "The multicast address-set claim (MASC) protocol, RFC-2909," Tech. Rep., 2000.
- [15] S. C. Nelson, G. Bhanage, and D. Raychaudhuri, "GSTAR: generalized storage-aware routing for mobilityfirst in the future mobile internet," in *ACM MobiArch*, 2011.
- [16] B. Kadaba and J. Jaffe, "Routing to multiple destinations in computer networks," *Communications, IEEE Transactions on*, 1983.
- [17] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed internet routing convergence," *ACM SIGCOMM CCR*, 2000.
- [18] F. Bronzino, D. Raychaudhuri, and I. Seskar, "Experiences with testbed evaluation of the mobilityfirst future internet architecture," in *Networks and Communications (EuCNC), 2015 European Conference on*. IEEE, 2015, pp. 507–511.
- [19] D. Raychaudhuri *et al.*, "Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols," in *IEEE WCNC*, 2005.
- [20] S. Deering and D. Cheriton, "Multicast routing in datagram internetworks and extended lans," *ACM TOCS*, vol. 8, no. 2, pp. 85–110, 1990.
- [21] A. Ballardie, J. Crowcroft, and P. Francis, "Core Based Tree (CBT)–An Architecture for Scalable Inter-Domain Routing Protocol," in *ACM SIGCOM*, 1993.
- [22] J. Chen, M. Arumathurai, L. Jiao *et al.*, "COPSS: An Efficient Content Oriented Pub/Sub System," in *ACM ANCS*, 2011.