Supporting Seamless Mobility in Named Data Networking

Ravishankar Ravindran†, Samantha Lo‡, Xinwen Zhang†, Guoqiang Wang†
†Huawei Research Center, Santa Clara, CA, USA. {ravi.ravindran,xinwen.zhang,gq.wang}@huawei.com
‡Georgia Institute of Technology, Atlanta, GA, USA. samantha@cc.gatech.edu

***Abstract*—Information-Centric Networking (ICN) architectures aim to replace current host-centric IP architecture with an information-centric one for efficient, secure, and reliable dissemination of information. ICN is built on several salient principles such as publish and subscribe paradigm, named content, in-network caching, and security over atomic information objects. These features allow a data chunk to be cached and retrieved from multiple nodes in the network, and can be validated without building a connection with its host. Though these tenets simplify the mobility problem in ICN, seamless mobility for real-time applications still demands a control plane. We propose three cross-layer network-assisted seamless mobility schemes: (1) point of attachment based, (2) rendezvous point based, and (3) multicast based, with the overall design objective of minimizing the loss of interests and data during a handoff scenario. This paper describes these schemes in named data networking (NDN) framework and discusses their trade-offs in terms of control and forwarding plane requirements.**

I. Introduction

Named Data Networking (NDN) [1] is based on content-centric networking (CCN) [2] proposal. The network layer in NDN uses unique names to identify information objects, which are used in the *Interests* expressed by information consumers (or subscribers) and the *Data* response of information producers (or publishers). When an NDN router receives an interest, it first looks for the information object in its own content store (CS). If it is cached, the NDN router responds with the copy in the CS. If not, two tables namely pending interest table (PIT) and forwarding information base (FIB) are used to handle the interest packet. PIT maps information object names to the requesting face(s) of the content, and the FIB stores the mapping of aggregated name prefixes to the next-hop forwarding face(s) towards the information publishers. Furthermore, the PIT aggregates interests and multicasts data to the requesting users, and the FIB forwards interests from the consumers towards longest prefix matching publishers.

Location/identity split is a fundamental requirement in the network layer to handle mobility. Unique and stable identifiers of network hosts and information objects enable uninterrupted application session irrespective of mobile node's (MN) mobility. IP addressing does not satisfy this property, hence another layer of addressing is used to handle mobility. For example, protocols based on Mobile IP [6] have been developed for this purpose and extensively used, but they suffer from issues such as triangular routing, control overhead to manage the routing states between current point of attachment (PoA [1]) and mobility management points such as the home agent, and backward compatibility with handsets that do not support client end protocols for mobility. The principle of location/identity split is an integral part of NDN architecture, as all entities (information objects, hosts, and network nodes) are uniquely named and locations are resolved by routing over these names. NDN allows a MN to retrieve missed data during handoff by re-expressing its interests under a publish/subscribe paradigm. This may meet the requirements of non real-time application such as web browsing, but may not be sufficient for real-time applications, particularly if seamless experience is desired. In this paper we address the problem of handling mobility in an NDN network seamlessly through a control plane framework that enables MN's mobility irrespective of it being a consumer, producer or both.

Towards *seamless mobility* support in NDN, we propose three schemes which leverages the features of NDN including location/identity split, in-network caching, and publish/subscribe paradigm. Our schemes consider the design complexity of control and forwarding plane in NDN, and leverage NDN networking properties, along with L2 capabilities to achieve seamless mobility.

The remainder of this paper is organized as follows. We define the NDN seamless mobility problem and its challenges in Section II, and then present a control plane architecture in Section III to handle seamless mobility. We then illustrate the details of three proposed schemes in Section IV, and discuss further research efforts towards their optimization in Section V. Section VI concludes this paper.

II. Problem and Challenges

Features such as location/identity split, in-network caching, and publish/subscribe paradigm in NDN meet the requirements of elastic applications such as web browsing and e-mail, but may not meet the need of real-time applications such as voice or video sessions with stringent quality of experience (QoE) requirements. Usually these applications have stringent one way delay requirement, e.g., voice in a wireless context has a one way delay requirement of 200-250ms, which is barely met considering voice packetization, de-jitter buffer, network propagation, handoff, and wireless backhaul delay. Furthermore, interests in NDN have to be pipelined to ensure

[1]PoA is L3 network entity (IP or NDN based) which could be an access router or a base station (BS).

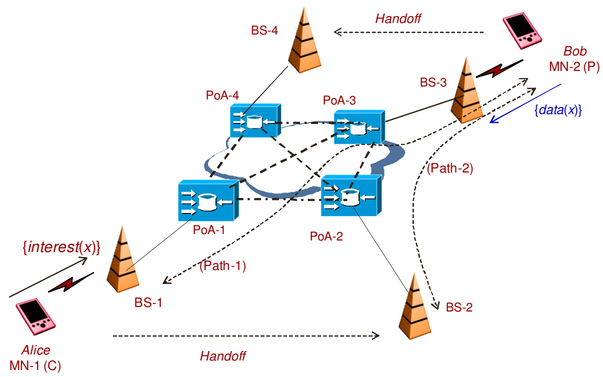

Fig. 1. Seamless mobility in NDN.

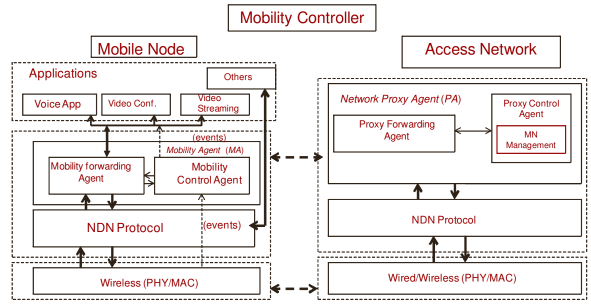

Fig. 2. Components of control plane for seamless mobility in NDN.

constant rate play out of a real-time streamed content at the receiver's end. As shown in Fig. 1, information consumer Alice's MN-1(C) has interests pipelined along Path-1 to the information producer Bob(P) with mobile node (MN-2). When MN-1 handoffs from PoA-1 to PoA-2, it uses Path-2 to re-express interests. The extent to which MN-1 can leverage in-network caching and re-expression of interests to retrieve the pipelined data depends on network topology that the cellular network backhauls to. The efficiency of the retrieval would be high in a hierarchical network, but will be unpredictable in a mesh network. If Path-1 and Path-2 do not intersect or do so after many hops, there is a high probability that the retrieved data will miss the play-out time at the consumer's end leading to poor QoE. Similar problem exists with producer's mobility too. When MN-2 is attached to PoA-3, the interests are routed along Path-1, with application in MN-2 generating the data response. When MN-2 moves to PoA-4, there may still be pending interests in PoA-3 and nodes along Path-1 not responded by Bob, resulting in data loss to MN-1. Furthermore, there has to be a mechanism to track and update Alice of Bob's mobility so that future interests can be routed to its current location. These problems necessitates a framework which ensures minimum loss of interests and data when either consumer, producer or both are mobile.

Mobility in wireless access has two underlying problems. First, the problem of tracking a MN in the network. Here the objective is to enable another MN to resolve it to its current location at any given time. Second, the problem of seamless mobility, where the objective is to minimize the impact of MN's mobility on the applications. One way to solve both the problems is by leveraging the routing control plane ([2] suggests the use of OSPF/ISIS for dynamic routing in CCN) to track MN's mobility. Here, as each MN is uniquely named and if the PoA routers announce the names of the attached MNs into the routing plane, the resulting FIB shall allow interests to be routed to the desired MN's current location. This solution however suffers from issues such as routing convergence, route updates considering high-speed MNs, and in general FIB scalability considering the number of MNs in a domain.

A scalable way to address the two problems is by addressing them separately. The problem of resolving an MN to its current location can be addressed in a centralized manner that manages the mapping of the MN to its current PoA. This is achieved by re-naming an MN when it registers with a domain. The name-format applied is $/\langle PoA\_Attachment\rangle/\langle MN\_ID\rangle$. The first component here changes with MN's location,while the second component is the unique ID of the device. [5] shows how this per-domain centralized mobility framework can be further leveraged to deliver richer mobile services. This approach also scales the routing plane by requiring it to route to fixed network elements in the network. The seamless mobility problem is dealt by orchestrating a re-direction infrastructure to handle active session between two MNs. For this we propose a control plane which uses network based anchor-points to handle seamless mobility. The control plane is particularly aimed to assist conversational applications with pipelined interests to meet real-time application requirements. The number of pipelined interests is a function of the round-trip time (RTT) and application's packet generation rate; that is, the number of outstanding interests in a session increases with increasing RTT or frame rate. Therefore, with respect to a mobile consumer, the objective is to have as many outstanding interests expressed before handoff retrieved efficiently after handoff. From a producer's perspective, the challenge is to re-direct the interests to producer's new location after handoff, at the same time ensure the pending interests from the consumer before handoff are responded to, and data appropriately forwarded to the consumer.

III. Control Plane for Seamless Mobility

Fig. 2 shows the control plane components to assist mobility in NDN. The control plane assumes the participation of the MN during mobility, but this framework can be adapted to the case where the mobility is handled by the network alone, as in PMIPv6 [3]. The control plane comprises of three major components: client mobility agent (MA)in MN, network proxy agent (PA)in the PoA, and mobility controller. Another variant of PA is the rendezvous point agent (RA) which resides on other non-PoA router(s) based on certain mobility optimization objectives. MN's attachment to the network begins by communicating with the PA, then based on the seamless mobility scheme uses the PA or the RA to handle its data flow.

The functions of MA and PA can be further broken down into forwarding and control functions. At the MN's end, the control component implements protocol logic to handle MN's registration and other protocol exchanges during handoff with the access network . The forwarding component takes part in the processing of interests sent out by mobile applications before forwarding them to NDN network layer, and in the reverse path receives data and maps them to applications above it. The MA's control component interfaces with MAC layer to receive events indicating various stages of the handoff process. Applications may also choose to register with the MA to receive such notifications, which enables applications to adapt to various situations. Considering the fact that the seamless mobility support incurs control and forwarding overhead, all applications needn't use the MA. Applications having liberal QoE requirements can rely on its own interest time-outs, re-expression, and in-network caching to recover from connectivity disruptions.

At the network end, a PA manages MNs homed to it and keeps track of their registration and handoff state. Furthermore, based on the seamless mobility strategy, it executes forwarding path logic to achieve seamless mobility. From the control plane design perspective, the objective is to minimize the control and forwarding states the seamless mobility control components introduce, and the delay as a result of processing the interests and data.

The role of the mobility controller is to resolve the location of an MN in its local domain. It manages the mapping between MNs – identified by their unique IDs, and their current location in the network.

IV. Seamless Mobility Schemes

This section presents three different seamless mobility schemes: PoA based, rendezvous point based, and multicast based. In mobile IP, triangular routing results from a dichotomy required to handle location/identity split. This is not a concern in NDN, as MN IDs are fixed and applications bind to these IDs, and the control and forwarding plane handle the problem of resolving the MNs and the session flows to their current locations. Following, we summarize several assumptions upon which we discuss our schemes. We assume that a PoA is a NDN based L3 aggregation router. We further assume that there is standard authorization and authentication process during the network attachment phase. MNs are named relatively to their serving PoAs for name-based routing scalability purpose. We also assume that the expiry of interests and data is set by applications, but network is authorized to modify these settings to aid seamless mobility. In NDN, the interest routing is based on the longest prefix match, hence pre-appending appropriate prefix enables topological aware routing. Although in general an MN is both information consumer and producer during a conversational session, for sake of exposition of our schemes we consider the MN-1 as the consumer (C) and MN-2 as the producer (P).

A. *Point of Attachment(PoA) Based Seamless Mobility*

This scheme leverages the MA and the PA to achieve seamless mobility. A high level view of the setup is shown in Fig. 3. We discuss the various phases of the scheme

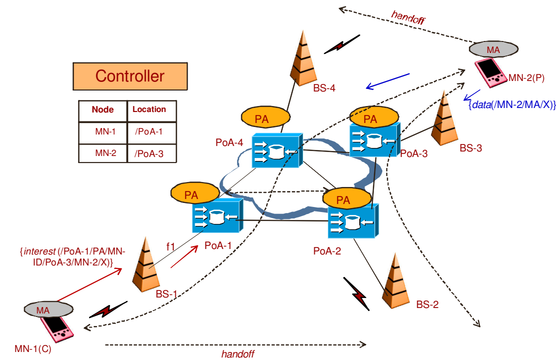

Node	Location
MN-1	/PoA-1
MN-2	/PoA-3

Fig. 3. PoA based seamless mobility.

with respect to normal operation mode and consumer's and producer's mobility.

1) Bootstrap and Normal Operation: We assume the MA in MN-1 in Fig. 3 has a way to discover the PoA(s) in its vicinity and chooses PoA-1 based on its L1 properties. During the discovery phase, the PA notifies the MA with the prefix it can use to correspond with the PA, e.g., /PoA-1/PA/. MN-1's MA then sends a registration message to PoA-1. The minimum information in the interest includes the MN's unique ID (MN-1). The successful registration creates an entry in the PoA-1's FIB in form of /MN-1/MA, which allows the PA to forward any control message or data from PoA-1 to MN-1. This implies that the PoA's FIB has to scale in a linear order of the number of MNs, which we believe is a reasonable engineering requirement. This bi-directional forwarding path between the MA and the PA allows the two entities to maintain forwarding state to manage MN's mobility. The control message is handled by the control agent in the PA and MA, and is distinguished from the application session flows by the sub-components in the prefix [2]. As part of the registration, the PA also updates the controller about the MN-1's presence in the network by binding its ID to its attachment point /PoA-1.

To establish a conversation with MN-2, MN-1 resolves MN-2 by issuing a resolution query to the PA, which resolves the MN-2 through the controller. For our discussion, let MN-2 be resolved to /PoA-3 as in Fig. 3. Here the applications in MN-1 are oblivious of MN-2's location or mobility state, as it is handled by the MA on behalf of the applications.

During the conversation, an interest X from MN-1 to MN-2 is pre-appended by the prefix supplied by the PA (e.g., /PoA-1/PA/MN-1) and the name required to forward it to MN-2. The PA then receives the interest from the MA with prefix /PoA-1/PA/MN-1/PoA-3/MN-2/X, and issues a new interest with prefix /PoA-3/PA/MN-2/X. The PA in PoA-1 manages the mapping between the original interest X and the new interest expressed on behalf of MN-1. The re-expressed interest is routed to PoA-3. When the PA in PoA-3 receives this interest and checks that it contains a contextual detail, i.e., MN-2, it issues a new interest to MN-2 with the prefix of /MN-2/MA/X. Assuming the application at the MN-2 is registered with MA, the MA forwards the interest X to the application.

The prefix setting for data from the MN-2 follows the names corresponding to the PIT state in the reverse path.

2) Information Consumer Mobility: When MN-1(C) is moving to another PoA, the seamless mobility actions are triggered when MN-1's control agent gets a notification from the L2 MAC about the initiation of the handoff. In order to minimize interest packet loss, MN-1's control agent first notifies MA's forwarding agent to stop issuing new application interests, though MA could still continue to receive interests from the application.

Deregistration with old PoA: The MA sends a de-registration interest message to the PA in PoA-1, which consists of the MN's ID and the new PoA (PoA-2) it is registering to. As the PA keeps track of outstanding interests from the attached MNs, it can identify outstanding interests for MN-1. Considering MN-1 is mobile, incoming data for the pending interests are cached with higher expiry as they are expected to be retrieved later. At this point the PA can delete the FIB state created to serve MN-1.

Registering with new PoA: During the registration process, we assume the same form of aforementioned bootstrapping operation, which allows the MA and the PA in PoA-2 to establish communication. After the registration, the MA notifies PoA-2 of its previous attachment with PoA-1. The PA in PoA-2 then signals the PA in PoA-1 to retrieve data corresponding to the pending interests.

Retrieving data from old PoA: The PA in PoA-2 now begins a control plane interaction with the PA in PoA-1. Once the PA in PoA-1 receives such a content retrieval request containing contextual information of MN-1, the PA responds with the set of pending interest prefixes in the context of the MN-1 for which it has received content. The PA in PoA-2 then issues interests to retrieve the cached data for the pending interests on behalf of MN-1.

Re-expressing pending interests: MN-1 starts re-expressing the pending interests as soon as it finishes its registration with the new PoA. The pending interests are marked so that the PA could match it against those for which transfer has been initiated from the previous PoA.

This and the above steps can be fine-tuned by considering several factors such as L2 attachment process, connectivity property, and application requirements.

Updating the controller: PoA-2 then updates the controller about the new location of the MN for future resolution.

3) Information Producer Mobility: Similar to the above protocol, when MN-2(P) moves from its current PoA (PoA-3) to new PoA (PoA-4), the two PoAs coordinate to ensure that data for interests are routed to the new PoA so that these pending interests from the consumer can be satisfied.

Handoff trigger and de-registration: Once the MA in MN-2 is notified of a handoff from L2, it stops issuing data for interests from consumers. The MA then notifies the PA in PoA-3 about its mobility by issuing a de-registration message. As the PA handles incoming interests from the network in the context of the MN-2, it refreshes them locally with higher expiry. The PA can distinguish network interests from those that come from MN-2 based on the face they are arriving on. During the handoff, MN-1 may not be aware of MN-2's mobility, hence continues to issue interests. These interests are cached and held by the PA in PoA-3, and marked as pending interests which MN-2 is not yet aware of. The PA in PoA-3 now waits for further action to be notified by the new PoA.

Registration with new PoA: Once MN-2 attaches to a new PoA (i.e., PoA-4), the PA in PoA-4 contacts PoA-3 by expressing interests for outstanding interests which require data from MN-2. Once the PoA-3 receives this notification, it re-expresses the pending interests from MN-1. The PA in PoA-4 keeps track of these incoming interests, and issues them to the MN-2. The MA in the MN-2 then looks in its own cache to see if there are corresponding data, if so it responds immediately, otherwise it forwards the interests to the application. The PoA-4 then updates the controller about the MN-2's new location. Later, a notification is triggered to MN-1 about MN-2's new location, this can be handled either by the MAs at MN-2 and MN-1 or with the help of their corresponding PAs.

B. Rendezvous Point Based Seamless Mobility

Our second scheme leverages strategically located routers called rendezvous points (RP) across the access network to aid seamless mobility, instead of PoAs. The objective is to minimize the need for multiple anchor-points to handle seamless mobility which shall improve the control plane efficiency. The RPs can be chosen by considering global optimization objectives; these factors include network topology hierarchy, user behavior, current traffic pattern, etc. In line with mobile IP solution such as in [3], RP can be higher level aggregation router managing MN's belonging to multiple PoAs. In this case the seamless mobility of MNs attached to these PoAs can be handled locally by the RP itself. Another approach to choose a RP node is during the MN registration process. In this case the PoA uses local optimization criteria such as network topology, geographical proximity to adjacent PoAs considering MN's mobility characteristics, traffic loading of the RPs, to assist the MN with seamless mobility.

From the control plane perspective, the role of the PA agent in the PoA is now limited to manage MN's network attachment process and assign a RP to the MN. The control and forwarding state of the MNs is maintained by the RP agent (RA). Though this scheme simplifies handling of MN's mobility, the RP should now scale to manage all the MNs attached to the set of PoAs it is currently serving. The controller's role is also enhanced in this scheme to keep track of the mapping of an MN to the RP serving it.

1) Bootstrap and Normal Operation: With respect to Fig. 4, when MN-1 completes its registration with PoA-1, PoA-1 chooses RP-1 to handle its seamless mobility. PoA-1 notifies this assignment in form of prefix /RP/RA to the MA. During MN-2's resolution phase, the request is issued by the MA, which is resolved by the PA at the granularity of the RP and the PoA it is attached to. With respect to Fig. 4, the resolution results in the prefix /RP-3/PoA-3/MN-2. An interest X from

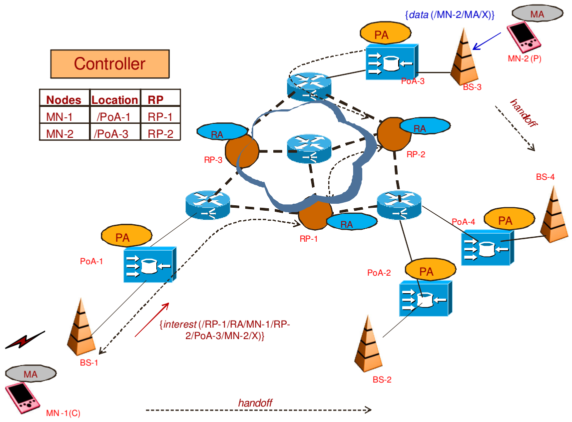

Fig. 4. Randezvous point based seamless mobility scheme.

the MN-1's application is pre-appended with two prefixes, one is MN-1's serving RP information, and the other is the location resolution information of MN-2, i.e., /RP-3/PoA-3/MN-2, resulting in prefix /RP-1/RA/RP-3/PoA-3/MN-2/X.

When the RA in RP-1 receives an interest from the MA in MN-1, it creates a state in the context of MN-1, and re-issues the interest with prefix /RP-3/RA/PoA-3/MN-2/X. Once RA in RP-3 receives this interest, it re-expresses the interest on behalf of the MN-2 with prefix of /PoA-3/MN-2/MA/X. The PoA-3 then routes this interest to the MN-2. The data from MA in MN-2 follows the states created in PITs on its way from MN-1 to MN-2. With this form of forwarding, the RAs in the RP-1 and RP-3 keeps track of the pending interests and data received, therefore serves as anchor-points to assist the mobility of MN-1 and MN-2.

2) Information Consumer Mobility: Following are the stages with respect to consumer's mobility.

Handoff trigger and de-registration: When handoff is triggered, it causes the MA to stop issuing new interests. When the MN-1 de-registers with PoA-1 through a notification containing its own ID, the control information is forwarded to the RA by the PA. As the RA is aware of MN-1, it takes similar actions as those in PoA based scheme, such as to refresh data for the pending interests with higher expiry in the context of MN-1.

Registering with new PoA: When the MN-1 registers with PoA-2, in addition to its own ID, it also includes the RP that is serving it. The PA in PoA-2 then notifies the RA of RP about MN-1's new location.

Re-expressing pending interests: Upon registering with PoA-2, MN-1 starts re-expressing its pending interests to the RA in RP-1. The RP then fetches cached data and generates appropriate responses to the MN-1.

Change of RP: RP-1 may not be optimal with respect to the PoA-2. Hence after the content retrieval process PA could notify MN-1 to use another RP based on the operator's RP selection objectives. Once the PA in PoA-2 signals MN-1 about the change of RP, it notifies the servicing RP-1 about de-registering client MN-1. This can be done gracefully through inter-RP interaction till all the forwarding state for MN-1 has been transferred to the new RP.

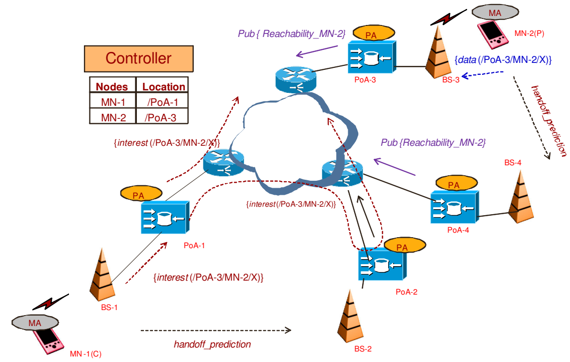

Fig. 5. Multicast based seamless mobility scheme.

3) Information Producer Mobility: When MN-2 moves from its current PoA (PoA-3) to new PoA (PoA-4), MN-2's PA notifies its servicing RP, i.e. RP-3, about its handoff. This causes RP-3 to take actions to ensure that current pending and future interests from the network are saved in the context of the MN-2. The MA also notifies the PA in PoA-4 about its current RP and continues using it during the handoff. When MN-2 registers with PoA-4, the PA of PoA-4 informs RP-3 about MN-2's new location. The RA then issues pending interests to MN-2's new location. Any data from MN-2 is then forwarded to RP-2, which in turn is re-mapped to appropriate prefix and forwarded to RP-1 and then to MN-1. Eventually, if the RP of the MN-2 changes, this information is updated to the controller. Also a notification is triggered to MN-1 about the change of the RP.

C. Multicast Based Seamless Mobility

Our third scheme takes the advantage of two fundamental network properties in NDN to achieve seamless mobility: multi-path interest routing and multi-point content (prefix) publishing. The former supports information consumer's mobility by anycasting interests from an MN along multiple paths, while the latter supports information producer's mobility by publishing prefix publishing from multiple points in the network. In this scheme, the PA doesn't keep session flow state for each MN attached to it, instead its task is to collaborate with PAs in other PoAs to conduct multi-path interest routing and multi-point prefix publishing to assist mobile consumer and the producer with seamless mobility.

1) Bootstrap and Normal Operation: The bootstrap and normal operations are similar to that in previous two schemes. The MA registers with the PA with its unique ID, which in turn creates control plane state to assist it with seamless mobility. Unlike the previous two schemes, here the PA doesn't keep any forwarding state for the MNs. In normal mode of operations, when the MN-1 is not mobile, as in the previous schemes, the scheme first resolves MN-2 through the controller. The MA then modifies application prefix by pre-appending MN-2's reachability prefix before forwarding application interests to the network.

2) Information Consumer Mobility: In this scheme, the PoA takes the help of information available at L2 to predict potential set of PoAs [4] that the MN may attach to. Compared to the previous two schemes, this scheme would require such prediction well in advance, as it leverages default operation of NDN to enable seamless mobility. In the normal mode of handoff operation, this potential neighbor list is compiled by the PoA and provided to MN-1. MN-1 uses its L2 measurement such as signal-to-noise ratio to determine the subset of PoAs that it could potentially handoff to. The objective here is to anycast interests from MN-1 through these potential set of PoAs, so that the information from the producer is cached when MN-1 handoffs to one of them. In normal mode of operation, the interests from MA are forwarded towards the current PoA of MN-2. In order to anycast MN-1's requests during handoff from other PoAs, we suggest the following: when PA receives an interest from MA, for each potential PoA, it pre-appends the corresponding PoA-ID to the interest prefix from MA and forwards it to its corresponding PA. This is in addition to forwarding the interest as per its own FIB policy. For e.g. when PA in PoA-1 receives interest /PoA-1/PA/PoA-3/MN-2/X from MN-1, and with the information that MN might handoff to PoA-2, it not only forwards a regular interest (/PoA-3/MN-2/X) as per its own FIB policy, but also pre-appends /POA-2/PA to the interest prefix /PoA-3/MN-2/X and forwards it. When PA in PoA-2 receives it and observes that it is from another PoA, the PA creates another interest with prefix /PoA-3/MN-2/X and forwards it. This shall enable a copy of the content to be available at PoA-2 when MN-1 handoffs to it.

3) Mobility of Information Producer: In order to track the producer's mobility, the network publishes the reachability information of the MN-2 via multiple PoA's. Once the PoA provides the potential neighbor list and receives feedback from the MN-2 about the potential set of PoAs it may handoff to, the multi-point prefix publishing can be done at the level of the MN or the applications that require seamless mobility support. The decision on the granularity of prefix to be published can be taken by the PA in coordination with the MA, considering user and application requirements and issues related to routing scalability. Once the prefixes to be published are agreed upon, the PA in current PoA coordinates with PAs in other PoAs to publish the prefixes into the content routing plane. This creates forwarding state in the content routers' FIB that enables interests to be multicasted to the potential PoAs to which the producer may handoff to. In Fig. 5, PoA-3 coordinates with PoA-4 to publish the prefixes. This causes interests from the MN-1 to be forwarded to PoA-4 as well. As data forwarding in NDN is loopfree, unlike IP, the requirement for absolute routing convergence is relaxed, e.g., in Fig. 5, the publish action may not be immediately reflected at PoA-1 or PoA-2. However,this technique would work with good probability as long as significant part of the network has updated information about the multi-point publish action. Though, we argued in Section II against using routing plane for user mobility, one way to mitigate the routing churn is by applying it on MNs which are mobile and entering the handoff mode.

V. Discussion

The intent of this paper is to propose various approaches towards seamless mobility in NDN. Further analysis is needed to study the efficiency of these strategies under practical L1-L3 conditions. In the context of PoA based scheme, the phase of transition between a servicing PoA to the new PoA and the retrieval and response of pending interests are the critical parts that require further research. The optimization also depends on L2 handoff mechanism, e.g., information available at L2 and the connectivity property of the MN during handoff. In the case of RP based scheme, the challenge is to assign an RP whenever an MN registers to a PoA. We have proposed two mechanisms: one is static, globally optimized on longer timescale, while the other is dynamic based on local optimization considerations. The multicast based scheme is unique in the sense that it leverages the inherent multi-path interest forwarding and multi-point publishing capability of NDN. While this approach simplifies control and forwarding plane complexity to achieve seamless mobility, it introduces inefficiencies in terms of transient FIB entries, PIT state overhead, and replicated data flows resulting in bandwidth cost. The efficiency of this strategy relies on applying accurate mobility prediction schemes.

VI. Conclusion

This paper proposes three seamless mobility techniques in NDN with the objective of minimizing the loss of pending interests and data between two parties involved in a real-time conversational session. The PoA based scheme and RP based scheme use anchor-points in the network to conduct seamless mobility. The RP based scheme improves the PoA based scheme in terms of providing flexibility in choosing nodes in the network to manage seamless mobility and overall reduce the number of points where forwarding states are managed. The multicast based scheme leverages the multi-path interest forwarding feature of NDN and the ability to publish a prefix through multiple points in the network. We have described the working of these schemes and the requirements and challenges from both control plane and forwarding plane perspectives. We finally discussed several optimization problems under research towards efficient realization of these schemes under practical setting.

References

[1] Named data networking project: www.named-data.org.

[2] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard. Networking named content. In *Proceedings CoNEXT*, 2009.

[3] J. Lei and X. Fu. Benefits of introducing pmipv6 for localized mobiility management. In *Proceedings of IWCMC*, 2008.

[4] P.Pathirana, P.Savkin, and S.Jha. Mobility modelling and trajectory prediction for cellular networks with mobile base stations. In *Proceedings of the MobiHoc*, 2003.

[5] R.Ravindran, X.Zhang, and G.Q.Wang. Towards secure mobile virtual group in information centric networks. In *Proceedings of the IEEE ANTS*, 2011.

[6] Z. Zhu, R. Wakikawa, and L. Zhang. A survey of mobility support in the internet, rfc 6301 (informational), 2011.