

A Survey of Mobility Support in Named Data Networking

Yu Zhang

Harbin Institute of Technology, China

Email: yuzhang@hit.edu.cn

Alexander Afanasyev, Jeff Burke, Lixia Zhang

University of California, Los Angeles

Email: {aa, lixia}@cs.ucla.edu, jburke@remap.ucla.edu

Abstract—The initial Named Data Networking (NDN) architecture design provided consumer mobility support automatically, taking advantage of NDN’s stateful forwarding plane to return data to mobile consumers; at the same time, the support of data producer mobility was left unspecified. During the past few years, a number of NDN producer mobility support schemes have been proposed. This paper provides a clear definition of NDN mobility support, to enable fetching of data produced by mobile users, and then to classify the proposed solutions by their commonalities and to articulate design tradeoffs of different approaches. We identify remaining challenges and discuss future research directions for effective and efficient data producer mobility support in NDN.

I. INTRODUCTION

Named Data Networking (NDN) [1], [2] is a proposed Internet architecture that changes the network communication model from point-to-point packet delivery to named data retrieval, where consumers send interests to fetch data. The NDN architecture naturally supports consumer mobility through its stateful forwarding plane, but the data producer mobility problem remains an active research topic.

In this paper, we aim to understand the design space of NDN mobility support, its dimensions, and the design tradeoffs. More specifically, this paper focuses on mobility support in the context of the globally connected infrastructure. The examination of mobility solutions in locally scoped, highly dynamic, ad hoc, and disruption tolerant environments will be the subject of our next paper.

Given the long history of IP mobility solution development [3], in this paper, we first sketch out a summary of IP mobility support approaches (§II). This is followed by a discussion of the NDN mobility support design space (§III). Our goal is to understand whether any commonalities exist between IP and NDN mobility support despite their architectural difference—location-centricity of IP versus data-centricity in NDN—and how this difference leads to differences in the design spaces for mobility support. We then sort the existing works [4]–[17] into several bins based on their differences in design approaches (§IV). Afterwards, we discuss the design tradeoffs and security challenges (§V), and conclude the paper with our evaluation of the solutions and describe future work (§VI).

II. MOBILITY SUPPORT IN IP

We take previous IP mobility research as a beacon to help us explore the design space of mobility support in general.

A survey through existing IP mobility solutions leads us to the following key observation: the common core of all the solutions is a *rendezvous mechanism* that tracks each mobile node (*MN*) and directs IP packets towards it [18]. In this paper, we let “*rendezvous (RV)*” represent a general concept of indirection. Based on the type of rendezvous mechanisms, we divide the existing solutions into three categories: *routing*, *mapping*, or *tracing*.

The routing-based approach uses the routing plane as the RV. An MN can keep its IP address unchanged while moving, but the MN is required to continuously send routing updates to inform all routers of its whereabouts. Some real deployments (e.g., Connexion service [19]) went in this direction. However, concerns about the routing scalability moved the community focus to more scalable solutions.

The mapping-based solutions scale much better by utilizing stable mapping servers as the RV service. In these solutions, the MN reports its current IP address to the RV whenever its address changes. When the mapping is implemented at the networking layer, such as in Mobile IP [20], the RV (home agent) keeps the mapping of a stable *home address* of the MN to its current care-of IP address. Packets destined to the MN reach the RV, which tunnels them (e.g., using IP-IP encapsulation) to the MN’s current address. ILNP [21] and BTMM [22] represent approaches that perform the mapping above the network layer. In these solutions, each MN has a DNS name and updates its name-to-address mapping in DNS whenever its IP address changes. Thus, one can query DNS to find the MN’s current IP address.

The *tracing-based* approach to IP mobility has received much less attention than the other two. The MN sends signaling messages to the RV after each move, to create a hop-by-hop reverse path from the RV back to itself. In other words, the signaling messages from the MN set up the forwarding entry for the MN at each router between the RV and the MN, so that the RV can forward packets to the MN without using tunneling. Trace-based solutions require that each involved router keeps a per-mobile routing state, which is considered prohibitively expensive. Cellular IP [23] and HAWAII [24] are two solutions falling into this category. Their commonalities are that both use the root gateway in a home/foreign domain as the RV for micro mobility, and utilize Mobile IP for macro mobility. To address the MN within a foreign domain, the former retains the MN’s home address, while the latter assigns

the MN a new care-of address.

III. MOBILITY SUPPORT IN NDN

The data-centric nature of NDN changes the mobility problem from “*delivering packets to an MN*” to “*retrieving data produced by an MN*”. Given the two-way interest/data packet flow in NDN, the NDN mobility can be split into two sub-problems: how the requested data can be returned to a moving consumer (consumer mobility), and how (mobile or not) consumers’ interests can reach the data generated by moving producers (producer mobility).

A. Consumer Mobility

To retrieve the desired data, NDN consumers express interests to the NDN network [1]. An NDN router uses the interest name to look up the content store (CS) first, and then, if there is a match, return the data back to the consumer. Otherwise, the router records the interest in its Pending Interest Table (PIT) and further forwards it based on the forwarding strategy, which takes input from the routing (Forwarding Information Base, FIB) and data (traffic measurements) planes. The state recorded in the PIT of each router traversed by the interest creates a reverse path (“*breadcrumb trail*”) that is used to return the requested data once it is found.

Because data packets are returned by tracing the interest path back to the consumer, the NDN architecture provides the built-in support for consumer mobility. When the consumer moves while the network is retrieving the requested data, the consumer can re-express its pending or timed-out interests to update/recreate the reverse path back to its current location. If the old and new paths cross, re-expressed interests retrieve the previously requested data from a router’s cache or are combined with the previous interest without propagating further.

The consumer mobility support in NDN may look remotely similar to the tracing-based mobility support in IP, however they differ in a fundamental way. The latter performs a routing process separated from user communication, while the former is a byproduct of the stateful data plane built into the architecture. Although the per-packet forwarding state represents a new cost brought to routers by NDN, it is a necessary component that provides multiple much needed functions [25].

B. Producer Mobility

The NDN producer mobility problem may seem similar to IP mobility upon first look. It has, however, a conceptual difference: NDN is about retrieving data and not about delivering packets to a mobile node (*MN*). Therefore, producer mobility support is to make interests rendezvous with data generated by a mobile producer (*MP*).

One direction is using an RV to find out where an MP is, namely *MP chasing*. In this direction, interests are steered towards an MP to retrieve data (the interests may not have to reach the MP if they find requested data in router caches along the way). This is similar to the IP mobility support, and IP’s routing-, mapping- and tracing-based approaches can be adapted to NDN with similar characteristics and tradeoffs. We

exclude routing-based solutions in the rest of this paper, as it is unlikely feasible for MPs to announce the data’s prefix into the routing system on the Internet-scale.

Another direction is to ensure that the mobile-produced data can be easily found, namely *data rendezvous*. This relatively new direction can leverage the data-centricity of NDN in two ways. First, because all data packets in NDN are named and secured, they can be easily detached from the original MP. For example, data generated by an MP can be moved to a stationary and easily reachable location. Second, for applications that generate data at a given location, one can make the data name “stationary” and independent from MPs. Such a name enables consumer interests to be steered toward the location and retrieve data from any MP at or near the location.

IV. PRODUCER MOBILITY SOLUTIONS

In this section, we review the producer mobility solutions proposed in the last few years. To help understand commonalities and differences among the solutions, we group them by the key design choices. We identified two MP-chasing (*mapping* and *tracing*) and two data rendezvous (*data depot* and *data spot*) approaches. Table I summarizes the approaches, their characteristics, and the solutions utilizing each approach. Note that some of them use a combination of multiple mechanisms. The rest of the paper uses terms listed in Table II and conceptual illustrations based on the legend in Figure 1.

TABLE I
NDN PRODUCER MOBILITY APPROACHES

MP chasing			
§IV-A	Mapping	The MP reports to the RV its PoA(s) through which its data can be retrieved	[5]–[10]
§IV-B	Tracing	The MP creates a “breadcrumb trail” that can be followed by interests from the RV to reach the MP	[10]–[14]
Data rendezvous			
§IV-C	Data depot	Data produced by MPs is moved to a known stationary server	[15]
§IV-D	Data spot	Data is produced within, and available from, a stationary region	[16]




Fig. 1. Legend of the conceptual illustrations (red color represents a moved-out MP, an old point of attachment, or an outdated trace)

A. MP Chasing: Mapping-Based Solutions

The mapping-based solutions more or less borrowed the ideas from IP mobility support, i.e., using stable RVs to track MPs. The core idea is to require an MP to report to the RV the mapping from the MP’s “stable” data name to its “temporary”

TABLE II
TERMS AND DEFINITIONS

General terms	
RV name	Globally reachable name of an RV, which may name a service, a network, or a node “/dropbox”, “/att/lte”, or “/timewarner/socal/la/rtr20”
PoA name	Globally reachable name of a point of attachment, which may name a network or a router “/att/lte” or “/timewarner/socal/la/rtr20”
Data name	name of a piece of data, which may or may not be globally reachable “/dropbox/a/family/photos/Paris/I91.jpg” “/a/family/photos/Paris/I91.jpg”
Interest to fetch data whose name is not globally reachable	
Interest with hint	Interest with unreachable data name, and carries a hint field containing a globally reachable name (e.g., PoA name) “/a/family/photos/Paris/I91.jpg” with hint “/att/lte”
Interest with prepending name	Interest with a name N that is a concatenation of a reachable (PoA) name N_{PoA} and an unreachable data name N_D . Note that this interest retrieves an encapsulated data packet (i.e., data with name N_D encapsulated in a data with reachable name $N_{PoA}N_D$). “/att/lte/./a/family/photos/Paris/I91.jpg”
Definitions for tracing-based approaches	
Trace command interest	Command interest towards RV to set up PIT or tFIB state (trace) for data retrieval “/dropbox/...</a/family/photos>”
Trace name	Full name of trace command interest “/dropbox/...</a/family/photos>”
Interest with trace name	Interest with data name, and carries a trace field (trace name, instead of the data name, is used to forward interest) “/a/family/photos/Paris/I91.jpg” with trace name “/dropbox/...</a/family/photos>”

current point of attachment (PoA) name, whenever the MP moves to a different PoA. Near half of the published solutions fall into this category [5]–[10]. Their designs can be roughly split along two dimensions, based on what role the RV plays and how the mapped PoA names are carried in interest packets.

1) *The role of the RV*: In a mapping-based solution, the RV may be a name mapping service [6] that provides the mapping from the name of the data produced by an MP to the MP’s current PoA name; a home agent which “tunnels” interests towards the MPs [5], [9]; or be a hybrid system that includes name mapping and interest tunneling mechanisms [7], [8], [10] (Figure 2).



Fig. 2. Conceptual illustration of mapping-based solutions based on RV role. Cross-out (red) icon represents the previous reachable prefix.

When the RV acts solely as a mapping service, consumers first find from the RV the mapping between the data name and the current PoA name, and then they “tunnel” interests

towards the current PoA. The RV name may or may not be independent from the data name, see Table II for examples. Exactly how to “tunnel” interests is discussed later.

When the RV serves as a home agent of the MP, the data is published under the “stationary” RV name that is globally reachable. A consumer’s interests reach the RV, and the RV “tunnels” the consumer’s interests towards the MP at the PoA based on the current mapping.

In the hybrid mode of operations, the RV serves as a home agent to “tunnel” the first interest to the MP. The returned data packet carries the current PoA name of the MP, which can then be used by the consumer to “tunnel” interests towards the MP directly.

Some of the existing proposals assume the existence of a dedicated home agent that plays the role of the RV [7], [9], [10]. To update the mapping, the MP directly contacts the RV by sending an interest whose name shares the same prefix as the home agent together with an RV-specific command suffix that carries the PoA name of the MP. A few other proposals rely on an infrastructure-supported rendezvous service to play the role of RV [5], [6], [8]. One such example is NDNS [26], which provides a secure distributed mapping service for all parties in an NDN network.

2) *Fetching data using the mapped PoA names (“tunneling” interests and data)*: In order to guide the interest carrying an unreachable name toward the MP’s current PoA, the PoA name needs to be attached to the interest in some form. The existing proposals explored two options: (1) prepending the PoA name to the data name [8]–[10] carried in the interest, namely *interest with prepending name* in Table II and (2) letting the interest carry the PoA name aside in a separate “hint” field of the interest [5]–[8], namely *interest with hint* in Table II.

Given that one cannot alter the name of any given NDN data packet, PoA name prepending requires on-demand creation of a new data packet, which has the compound name (PoA name plus content name) and encapsulates the original data packet as its content. Using hint does not change the original data name, but requires modifications to the interest forwarding process.

B. MP Chasing: Tracing-Based Solutions

The tracing-based solutions leverage and extend NDN’s stateful forwarding plane to create the “breadcrumb trail” to retrieve interests from consumers [10]–[14] (Figure 3). Similar to the mapping-based solution with a home agent, tracing relies on data being published under the globally reachable prefix of the RV, i.e., the RV name shares the prefix with the data name. Whenever the MP moves, it sends *trace command interests* (see Table II) towards the RV, creating or updating a soft-state hop-by-hop reverse path, namely the $MP \leftarrow RV$ trace. Instead of retrieving data, trace command interests solicit interests from consumers to retrieve data. A regular interest from consumers will be forwarded along the trace if one exists, otherwise it is forwarded using the standard forwarding logic of NDN.



Fig. 3. Conceptual illustration of tracing-based solutions: consumer’s interests are forwarded toward the RV, they may meet the MP-trace along the way and move toward the MP without going through the RV.

It is worth noting that the RV in the tracing-based approach does not need to participate in producer/consumer communications; it can only make routing announcements for the data name prefix to attract both consumers’ interests for MP’s data and trace command interest toward itself, so that the former can meet the $MP \leftarrow RV$ trace either on their way toward the RV or otherwise at the RV.

When the MP moves and the new $MP \leftarrow RV$ trace crosses a previously set up trace, the consumers’ interests that are pending along the old trace can be automatically forwarded to the new trace [10], [14]. In [12], [13], in addition to sending trace command interest toward the RV, the MP also sends trace command interest towards its previous PoA to fetch the pending interests and minimize disruption of data fetching.

Two alternatives to store the trace state have been explored: *trace-in-FIB* [10]–[13] and *trace-in-PIT* [14]. The trace-in-FIB proposals introduce a dedicated *tentative FIB* (tFIB), separated from the FIB produced by routing protocols (routing FIB, or FIB). The tFIB is populated and updated whenever a router receives and authenticates the trace command interest. A regular interest will be forwarded along the trace, if there is a match in the tFIB.

The trace-in-PIT proposal [14] re-uses the existing data plane in PITs, extending its function to bring back interests. The trace is created as a side-effect of the interest forwarding. [14] introduces a new `traceName` field in interests which contains the name of command interest creating the $MP \leftarrow RV$ trace, and a `traceable` flag to indicate whether an interest can be traced by other interests. Upon receiving an interest with a `traceName` inside, the router first looks for an exact match of the `traceName` and trace names stored in the PIT. If a match is found, the interest is forwarded via the incoming interface of the matched command interest; otherwise, it is forwarding using the standard mechanisms.

Note that a tFIB entry or a PIT entry created by a trace command interest is soft-state, and it should timeout and remove itself in the absence of an explicit removal request. The tFIB entry can have a properly engineered timeout value, and the PIT entry may timeout in the same way as any other PIT entries. Alternatively, both could be engineered with a different timeout mechanism, e.g., having their lifetime refreshed by interests, so that the entry can last until the data retrieval from the MP is finished.

C. Data Rendezvous: Data Depot

Because NDN enables data packets to be easily detached from the original producers, instead of chasing the MP, one

can solve the producer mobility problem by moving the data produced by an MP to an easily accessible (e.g., stationary) location, which we call *data depot* (Figure 4). An example of data depot is custodian servers proposed in [15]. A data depot looks similar to the home agent in mapping-based solutions discussed in §IV-A, except that data depot takes the full responsibility for hosting the data instead of simply forwarding interests. For example, one can set up a data depot for a collection of family photos taken by multiple family members’ mobile phones. When the depot receives a request for a photo, it will either return the photo if it has been uploaded from a mobile device to the depot, or otherwise try to fetch it on demand using either a mapping- (§IV-A) or tracing-based (§IV-B) solution.



Fig. 4. Conceptual illustration of data depot: consumer interests are forwarded toward the depot to fetch data; they may meet the MP-trace on their way.

Generally speaking, a data depot is really a *name-based rendezvous*: a depot attracts both data and interests to meet at itself. Today’s Internet performs such functions at the application layer. For example, cloud storage today, e.g., Dropbox, largely serves as a name-based rendezvous and data depot for all its users who continuously generate data (text messages, photos, video clips, etc.), potentially while moving. Once data is uploaded to the cloud storage, it can be easily retrieved by others. NDN supports name-based rendezvous at the network layer by letting the data depot announce its name prefix into the routing table. In the data depot version of Dropbox, users will still upload data to the Dropbox servers, but the network will be able to help retrieve this data from wherever the data is available. For example, interests may be satisfied before reaching the Dropbox server by in-network caches, or when consumers and producers are near each other, or do not have connectivity to the infrastructure but have ad hoc connectivity in-between.

With a data depot, one can develop a *mobile producer upload* protocol that can both avoid the need for a third-party RV and hide producer mobility from consumers. For example, if the MP has data to upload, it can express an “upload request” in form of a trace command interest directly to the depot, creating the $MP \leftarrow depot$ trace on demand (§IV-B). The depot can then request the published data from the MP along the trace, and reply with an “ACK” data to satisfy the “upload request” interest after the upload is completed. Or with mapping, the MP can express an “update request” in form of an interest with hint carrying its latest PoA name. The depot can then update the mapping, fetch data from the MP, and reply with an “ACK” data packet to satisfy the “update request” interest afterward. Note that, to be a general mobility solution, the depot should keep up-to-date mapping to the MP’s PoA name to enable it to retrieve data from the MP upon demand.

D. Data Rendezvous: Data Spot

In some applications, data is associated with a specific geographical region and can be generated by any MP “on the spot” (Figure 5). One example is that in the road traffic applications the data about the road conditions at a certain location can be generated by any of the cars currently in this region [16] (“/I405/CA/Westwood/traffic/Jan9,2015-8pm”). Another example is that a car manufacturer collects diagnosis information of a specific car model within a region [17] (“/US/CA/LosAngeles/diagnosis/toyota/prius/2009”). Interest forwarding to retrieve such data can be realized either using a geo-routing or relying on roadside units to announce the region-specific prefixes in the routing system. When an MP in the region receives an interest, it evaluates it, e.g., using current GPS coordinates and associated cartographical information, and, if conditions are met, responds with newly produced data. The MP can also respond to the received interest without additional evaluations, if the matching data is available in cache, e.g., data that was previously generated or requested/overheard from others. Once the MP moves out of the “spot” region, it may no longer receive interests for the region’s data, which can be responded by other MPs in the region. Thus, in the data spot approach, the MP is a producer by chance.



Fig. 5. Conceptual illustration of data spot

V. DISCUSSION

Given the diversity of existing proposals, we discuss the design tradeoffs from several perspectives related to complexity, performance, and security implications.

A. Signalling Overhead

Each of the approaches comes with its own signaling cost. In the mapping-based approaches, the cost is associated with (1) the MP informing the RV when it moves to a new PoA, (2) the mapping information retrieval, and (3) additional overhead in interest packets (i.e., the hint or prepended name). In the tracing-based approaches, the cost is associated with (1) the MP continuously sending trace command interests to keep the trace alive, even if there are no active retrievals of MP’s data; and (2) additional overhead on routers to keep the trace state.

Depending on the mobility and data consumption patterns, the signaling cost can be amortized in different ways. If the MP have relatively high mobility, the mapping update and trace maintenance costs will be similar, with potentially high costs for the mapping information lookup. When MP-produced data is being actively retrieved, the soft state of the trace can be kept alive by using data from the MP as signals, so that the cost could be equalized with the mapping-based approach.

B. Involvement of RVs

With the mapping using a mapping service, consumers can learn the MP’s PoA name from the RV and send interests directly toward the MP. Mapping through home agents removes the need for a global mapping service, however interests and data packets will flow through the RV, potentially creating “triangle paths.”

Trace-based solutions have the potential to create “triangle paths” for data retrieval. Note, however, that in the data depot solution with tracing, depending on the network topology, an interest may not go to the depot to retrieve data as it can be “captured” by the traces and forwarded toward the MP directly. In addition, the data depot provides long-term storage for the MP’s data to satisfy later requests from consumers even when the MP goes offline.

C. Changes to Packet, or to Packet Processing

Interest with hint keeps the data name intact, but requires changes in the forwarding mechanism. Interest with prepending name does not change interest forwarding, but changes the name of requested data.

The data name change leads to several issues. First, if the data can be retrieved from multiple places, one must decide which PoA to use to fetch data, eliminating the flexibility for the network to retrieve data from alternative locations (i.e., the selected prefix pre-defines the direction the interest must go). Second, if different consumers choose different PoA names to retrieve the same data, due to the name change, they will each fetch a copy from the producer instead of being able to utilize router caches. Third, NDN allows consumers to retrieve data using full names (i.e., including the implicit digest), but prepending PoA names to the data name retrieves an encapsulated packet which has a different digest. Fourth, there is also a question of which key to use to sign the encapsulated data packet. We should clarify that none of these issues are unsurmountable roadblocks; they either reduce the network efficiency and/or require additional complexity to handle. We also note that as names are at the center of NDN design, a change to data names may have other impacts not listed above.

D. Security Considerations

As we mentioned earlier, the ability to attach a “hint” field in an interest to carry MP’s PoA name can be abused to direct the interest to a place of an attacker’s choice and bring back forged data which gets stored in router caches. One way to mitigate this denial-of-service through cache poisoning attack is for routers to pair up [interest, hint] with the retrieved data [6], [27]. This way, a data packet retrieved by an interest with a forged link cannot satisfy an interest carrying a different link.

In the tracing-based solutions, precautionary measures are needed to protect against attackers forging and replaying the trace interests, e.g., to masquerade victims and eavesdrop/blackhole interests or poison caches [28]. For example, the solutions may need to ensure that an interest is forwarded along a trace only when explicitly requested, to avoid easy hijacking of regular interests (e.g., if one sends a regular

Interest for “/google/search” data, the attackers should not be able to divert such interest by setting up a trace for “/google/search”). Trace interests may also need to have MP-specific signatures. However, this imposes verification cost at the routers and potentially opens a DDoS possibility against routers with signature verification requests.

The data spot approach requires a reliable way of filtering bogus data, e.g., to prevent malicious vehicles publishing false traffic information. For example, the data publishing module needs to be properly secured and protected from tampering with and, potentially, a voting mechanism may be needed to ensure that multiple MPs agree on the published content.

E. Solution Evaluation

Different solutions to the producer mobility problem bring different sets of tradeoffs. In evaluating these tradeoffs, if one makes “ease of consumer data retrieval” as the first criterium, our analysis suggests that the following two combinations of the proposed solutions could cover a wide spectrum of application and communication patterns: “data depot+mapping” and “data depot+tracing”. A data depot hides producer mobility from consumers by providing permanent or temporary storage to enable detachment of data from MPs and by serving as the RV to let interest/data meet. However a data depot needs to be enhanced with either a mapping or a tracing solution to fetch data from mobile producers in order to support applications that require support for retrieving data on-demand.

VI. CONCLUSION

This paper investigates the design space of mobility support in the NDN architecture. While consumer mobility is naturally supported by NDN’s stateful data plane, producer mobility requires a rendezvous mechanism to let interests meet data generated by mobile producers. We identified four ways: two conventional mechanisms to “chase” the moving producer (mapping and tracing) and two data-centric ways to let interest meet data (data depot and data spot).

For future work, we call for investigations in three directions: 1) quantitative and comparative evaluation of the few proposed approaches that seem promising, such as tracing in FIB and in PIT and mapping; 2) thoroughly examining security aspects of the MP chasing approaches and developing mitigation mechanisms; and 3) developing application scenarios that utilize data rendezvous, to gain further insight on its utility. The first two are crucial for putting NDN mobility support into practice, and the last one is needed to effectively leverage NDN’s benefit of data-centricity.

ACKNOWLEDGMENT

This work is supported by National Natural Science Foundation of China (No. 61202457) and National Science Foundation (CNS-1345318, CNS-1457074).

REFERENCES

- [1] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, K. Claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang, “Named Data Networking,” *ACM Computer Communication Reviews*, June 2014.
- [2] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, “Networking named content,” in *CoNEXT ’09*, 2009.
- [3] Z. Zhu, R. Wakikawa, and L. Zhang, “A survey of mobility support in the Internet,” RFC 6301, 2011.
- [4] M. Meisel, V. Pappas, and L. Zhang, “Ad hoc networking via Named Data,” *ACM MobiArch 2010*, 2010.
- [5] A. Azgin, R. Ravindran, and G. Wang, “Scalable mobility-centric architecture for named data networking,” *IEEE ICCCN ’14*, 2014.
- [6] A. Afanasyev, C. Yi, L. Wang, B. Zhang, and L. Zhang, “SNAMP: Secure namespace mapping to scale NDN forwarding,” in *IEEE Global Internet Symposium ’15*, 2015.
- [7] F. Hermans, E. Ngai, and P. Gunningberg, “Global source mobility in the content-centric networking architecture,” in *NoM ’12*, 2012.
- [8] D. Li and M. C. Cuah, “SCOM: A Scalable Content Centric Network Architecture with Mobility Support,” in *IEEE MSN*, 2013.
- [9] J. Lee, S. Cho, and D. Kim, “Device mobility management in content-centric networking,” *IEEE Commun. Magazine*, 2012.
- [10] D.-h. Kim, J.-h. Kim, Y.-s. Kim, H.-s. Yoon, and I. Yeom, “Mobility support in content centric networks,” in *ICN ’12*, 2012.
- [11] L. Wang, O. Waltari, and J. Kangasharju, “MobiCCN: Mobility support with greedy routing in Content-Centric Networks,” *Globecom*, 2013.
- [12] D. Han, M. Lee, K. Cho, T. T. Kwon, and Y. Choi, “PMC: Publisher Mobility Support for Mobile Broadcasting in Content Centric Networks,” *ASIA Future Internet 2012 Summer School*, 2012.
- [13] J. Augé, G. Carofiglio, G. Grassi, L. Muscariello, and G. Pau, “Anchorless Producer Mobility in ICN,” in *ACM ICN’15*, 2015, pp. 189–190.
- [14] Y. Zhang, H. Zhang, and L. Zhang, “Kite: A mobility support scheme for NDN,” in *ACM ICN’14*, 2014, pp. 179–180.
- [15] V. Jacobson *et al.*, “Custodian-based information sharing,” *IEEE Communications Magazine*, vol. 50, no. 7, pp. 38–43, 2012.
- [16] G. Grassi, D. Pesavento, G. Pau, L. Zhang, R. Vuyyuru, and R. Wakikawa, “VANET via Named Data Networking,” in *NoM*, 2014.
- [17] J. Wang, R. Wakikawa, and L. Zhang, “DMND: Collecting data from mobiles using named data,” 2010.
- [18] L. Zhang, R. Wakikawa, and Z. Zhu, “Support mobility in the global internet,” in *MICNET ’09*, 2009.
- [19] A. L. Dul, “Global ip network mobility using border gateway protocol(bgp),” *Boeing White Paper*, 2006.
- [20] C. E. Perkins, “IP mobility support for IPv4, revised,” RFC 5944, 2010.
- [21] R. Atkinson and S. Bhatti, “Identifier-locator network protocol (ILNP) architectural description,” RFC 6740, November 2012.
- [22] S. Cheshire, Z. Zhu, R. Wakikawa, and L. Zhang, “Understanding apple’s back to my mac (BTMM) service,” RFC 6281, 2011.
- [23] A. G. Valkó, “Cellular IP: A new approach to internet host mobility,” *SIGCOMM CCR*, vol. 29, no. 1, pp. 50–65, 1999.
- [24] R. Ramjee, T. La Porta, S. Thuel, and K. Varadhan, “Hawaii: a domain-based approach for supporting mobility in wide-area wireless networks,” *ICNP ’99*, 1999.
- [25] C. Yi, A. Afanasyev, I. Moiseenko, L. Wang, B. Zhang, and L. Zhang, “A case for stateful forwarding plane,” *Computer Communications*, vol. 36, no. 7, pp. 779–791, 2013.
- [26] A. Afanasyev, “Addressing Operational Challenges in Named Data Networking Through NDNS Distributed Database,” Ph.D. dissertation, ULCA, 2013.
- [27] NFD, “Design mobility with link,” <http://redmine.named-data.net/issues/3000>, accessed January 11, 2016.
- [28] Y. Zhang, H. Zhang, and L. Zhang, “Kite: A mobility support scheme for NDN,” NDN, Technical Report NDN-0020, 2014. [Online]. Available: <http://named-data.net/techreports.html>