Mobility Work Re-Visited Not Considered Harmful

Dirk Trossen*, Xenofon Vasilakos†, Paris Flegkas‡§, Vasilis Sourlas‡§ and George Parisis*

* Computer Laboratory, Computer & Communication University of Cambridge, United Kingdom.
† Dept. of Informatics, Athens University of Economics and Business, Greece.
‡ Department of Computer & Communication Engineering. University of Thessaly, Greece.
§ Centre for Research and Technology Hellas - ITI.
Email: dirk.trossen@cl.cam.ac.uk, xvas@aueb.gr, {pflegkas, vsourlas}@inf.uth.gr, georgios.parisis@cl.cam.ac.uk

Abstract—**The Internet is a universal source of information for its users. However, obtaining desired content has become the main focus of interest rather than the communication with machine endpoints. The increased availability of mobile devices makes such desired access continuous in time and space, adding an extra dimension and level of complexity for satisfying users' needs. Together with the almost ubiquitous availability of computing and storage resources, this has led to efforts that focus on disseminating information to a possibly changing set of consumers. In this paper, we revisit several strands of IP mobility work with respect to the benefit when being seen from an information-centric networking (ICN) angle. We present a set of architecture invariants together with a functional model, both of which allow for formulating the various concepts of mobility from an ICN angle. We furthermore present and evaluate two approaches to support mobility in ICN, a proactive one based on prefetching subscriptions to possible reconnection points of a mobile node and a reactive one based on buffer relocation. Finally, we also outline various challenges in evaluating such approaches, these challenges caused by the paradigm change that ICN constitutes.**

Index Terms—**Mobility, Information-Centric Networking, Publish-Subscribe.**

I. INTRODUCTION

Since the inception of the Internet more than 30 years ago, there have been significant advances in technologies. One of the more remarkable developments is the increase in available storage and computing resources, fueled by Moore's Law. A modern mobile phone can easily hold many gigabytes of memory while processing data at GHz speeds. This is a significant departure from the times of the inception of the Internet during which computing and memory resources were extremely scarce, usually assembled in a dedicated computing centre. Nowadays thought, users have access to personal computers or laptops which possess more computing power than many of these.

These changes have an impact on the design of future communication systems. With information potentially available in many storage places, caches, mobile phone memories or even embedded devices, the exact location of the "server" to be contacted moves in the background. The focus of interest becomes that of information, with the communication substrate being chartered with finding the best provider of the desired information.

Recent research efforts, such as NDN [1], PURSUIT [2], SAIL [3] and others, have made this starting point their central principle, i.e., they declare information a first class citizen on the networking level, while building a transient communication relationship between providers and consumers of information at any point in time.

In addition to the exponential increase in computing and memory resources, another development has marked communication in recent years: mobility. People, and increasingly objects, move about while providing their resources within a communication system that needs to cater to the various patterns of mobility that are inhibited by its participants. While cellular systems use specialized approaches for micro-mobility, many attempts have been undertaken to introduce mobility support in IP. Originally without any dedicated mobility support, approaches like Mobile IP introduce various techniques for supporting IP endpoints moving about in different networks.

More advanced techniques include the relocation of transfer buffers from one access point (AP) to another, the transfer of transport contextual information [4] or the identification of possible (IP-level) handover candidates [5] in the vicinity of the mobile node (MN) that is about to move. In the hay day of these activities, the Internet Engineering Task Force (IETF) devoted a dedicated working group to the aspect of seamless (IP) mobility from 2000 to 2003. Many of the employed techniques in these works are relevant for mobility approaches in information-centric networking (ICN), too, since they deal with identifying the regionalized neighborhood and relocating information in preparation of a handover or pre-establish context of communication for the purpose of uninterrupted communication.

However, there is a fundamental difference when looking at these approaches from an ICN angle. An information-centric approach is concerned with disseminating information to a possibly changing set of consumers, while the endpoint-centric approach of IP is concerned with establishing a communication relation between two dedicated endpoints with an opaque transfer of content over the established communication channel. This difference poses a challenge when revisiting this prior work. For instance, handover efficiency in IP is largely centered on the successful re-establishment of communication between endpoints, a procedure that hardly contributes to any efficiency increase in the future. On the

978-963-8111-76-0

contrary, an information-centric angle to this problem can increase efficiency for future handover procedures by utilizing dissemination of information that occurred in the past, i.e., utilizing cached information.

In this paper, we extend on this example for the different angle of how to look at previous IP mobility research. We start with categorizing previous (IP-level) mobility research in Section II. We then outline in Section III a Strawman information-centric architecture that allows for translating the various concepts onto the field of ICN. In Section IV we address the mobility concepts within the functional model, presented in Section III, and we present two approaches to support mobility in ICN. In Section V we evaluate the proposed approaches, while in Section VI we provide a list of challenges when moving from an endpoint-centric angle of mobility to an information-centric one. Finally, we conclude the paper in Section VII and provide insights for future work.

II. Re-visiting Seamless Mobility

Let us briefly describe the work regarding mobility in the current IP-based Internet carried out mostly by the relevant IETF working groups in their hay day throughout 2001 to 2003, as well as to present the related work in the area of mobility in overlay application layer publish-subscribe systems.

A. Regionalization

In order to enable seamless IP-layer handover, a MN should be able to discover its region, i.e., the identities and capabilities of Candidate Access Routers (CARs) to which a MN can be handed over. The concluded IETF SEAMOBY working group specified a CAR Discovery (CARD) protocol [5] that enables the acquisition of information about access routers that are candidates for the MN's handover. CARD involves identifying a CAR's IP address, resolving a Layer 2 identity to an IP address, and the capabilities that might influence the handover decision by MN, such as the available bandwidth. If there are more than one CARs, the MN chooses the one whose capabilities best match its requirements. However, the decision algorithm was left out of the scope of this work item.

B. Dissemination

Various protocols have been proposed for mobility management in the IP-based Internet, the most dominant being mobile IP [6]. Mobile IP supports mobility of IP hosts by using a home address that represents the fixed address of the node, and a care-of address (CoA) that changes with the IP subnet the MN is currently attached to. In Mobile IPv4 [7], these mappings are exclusively handled by home agents (HA). A correspondent node (CN), i.e., a node that wants to send traffic to some MN, sends packets to the home address of the MN. The HA 'intercepts the packets and tunnels them either directly to the MN itself, or to a foreign agent with which the MN shares a direct link.

The MIP6 IETF working group enhanced IPv6 so that HAs do not exclusively deal with address mappings; each CN can have its own 'binding cache' for storing home address to CoA mappings. In order to keep these mappings up-to-date, a MN has to signal any corresponding changes to its HA and/or CNs when performing a handoff to another IP subnet. MNs that move quickly or far away from their respective home domain or CNs, produce significant signaling traffic; this, results in suffering from handoff latency and packet losses when no extension to the baseline Mobile IP protocol is used.

Hierarchical Mobile IPv6 (HMIPv6) [9], developed by the MIPSHOP IETF working group, is a localized mobility management proposal that aims at reducing the signaling load due to user mobility. For this, it employs a hierarchy of foreign agents (FAs) to locally handle Mobile IP registration. According to this protocol, MNs send mobile IP registration messages with appropriate extensions to update their respective location information. Registration messages establish tunnels between neighboring FAs along the path from the mobile host to a gateway FA. Packets addressed to MNs travel in this network of tunnels which can be viewed as a separate routing network overlay on top of IP routing.

C. Establishing Context

The protocols described above are responsible for changing the routing paths in order to deliver traffic destined for a particular MN to its new point of access. Nonetheless, the success of real time services in a mobile environment, e.g., VoIP telephony and video, depends heavily upon minimizing the impact of this traffic redirection. For example, there may be serious impact on delay-sensitive real time traffic in case a node that has negotiated a successful AAA (Authentication, Authorization and Accountability) exchange with an access server, is required to re-establish this service when moving to a new location. For this problem, the IETF SEAMOBY working group has specified the Context Transfer Protocol (CXTP) [4] which focuses on re-establishing context without requiring that the MN explicitly performs all protocol flows for those services from scratch. When a MN decides to change its point of network attachment, CXTP undertakes to move the MN's current context from the old access router to the new one. The context may include header compression data, authentication credentials, or accounting records.

D. Mobility in Overlay Publish-Subscribe Systems

Systems Mobility solutions have not only found their entry in the IP layer itself, but also in overlay networks. In the following, we outline the most important solutions for supporting mobility in publish-subscribe systems. The first publish-subscribe overlay system that supported mobile clients was JEDI [8]. In JEDI, clients use two functions, namely move-out and move-in, to explicitly detach from the network and reconnect to it, possibly via a different node. In [10], authors implement a mobility support service that is independent from the underlying publish-subscribe overlay and manages active subscriptions and incoming messages transparently when a client detaches from one node until it reattaches at another. They use mobile service proxies, which

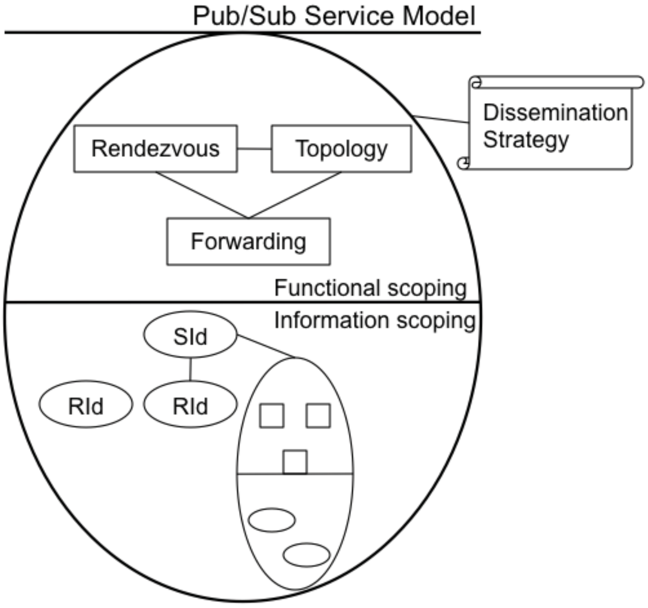

Fig. 1. Functional Model.

are independent stationary components that run at the edges of the overlay network, i.e., where clients exist. Service proxies are responsible for gathering the published events that match the interests of the mobile client, and deliver them when the client reconnects to the network. The former approach is similar to Mobile IP with the addition of caching the published events during the transition period.

Fiege et al. [11] modify the REBECA publish-subscribe middleware in order to support physical mobility of nodes as well as location-dependent subscriptions. Huang et al. [12] show how the anonymity, the dynamism, the asynchrony, as well as the multicast nature of publish-subscribe architectures make them ideal for wireless and mobile environments. Moreover, they discuss how various publish-subscribe models can be adapted to a mobile environment. Authors in [13] explore publisher mobility and propose various solutions to ameliorate multicast tree reconstructions and publication delivery times when the publisher changes location, as well as to limit delivery latency. These solutions include the use of a data prefetching mechanism based on publisher mobility prediction, the use of proxies that deliver data on behalf of publishers as well as the delayed teardown of a multicast tree when publishers change location. COMAN [14] is a protocol for organizing the nodes of a MANET in a tree-shaped network that is resilient to frequent topological reconfigurations. It was designed to minimize the number of nodes whose routing information is affected by topological changes.

Last, the authors in [15] present a mobility approach for publish-subscribe systems that envisions regions of dissemination called bubbles. As we examine in the following section, our ICN approach borrows from this approach while extending the recursive layering with a functional model and the notion of dissemination strategies.

III. AN INFORMATION CENTRIC DESIGN APPROACH

We now present a design approach for information-centric architectures that is based on a simple functional model from which we can assemble complex information-centric networking structures. For that, we outline the invariants for this functional model and position our concepts of mobility within this approach.

We formulate four invariants as the basis for our functional model. The first one is that of referencing information items, therefore establishing identities for information rather than the hosting endpoints. An information item is any form of data that is relevant within a particularly given context, i.e., it can represent a principal of a transaction, an identifier to a service being invoked, a policy rule acting on another piece of data or a pointer to some other data. Each information item is identified using a (statistically unique) flat label identifier. These identifiers are self-generated and the associated semantic of the information is only known to the implementation generating said identifier.

The second invariant is that of scoping. Scopes define the set of information that is being disseminated in a given context. With that, each scope represents the boundaries of a defined dissemination strategy for the information it contains. Hence, we achieve scoping on functional and information level. Each information item is being placed in at least one scope. Being a set of information, scopes are seen as information items themselves. Hence, they can be placed in a scope itself, potentially building complex graphs of information. Like information items, scopes are identified via (statistically) unique identifier. The scope's data contains the identifiers of all underlying information structures, including its sub-scopes.

The third invariant is the service model as that of publishing of and subscribing to individual information items within a scope. The resulting service API is that of creating items (and scopes) and publishing them towards anybody who subscribes to it. We leave out the details of such API in this paper.

The fourth invariant is that of separating the functions for information dissemination within a given scope into three distinct ones:

1) rendezvous concerns the matching of publishers' availability of information and subscribers' interest in it. This results in some form of location information for the set of publishers and subscribers that has been matched.
2) topology management and formation determines a suitable topology for the transfer of the respective information, and
3) forwarding executes this transfer of information. The exact nature and implementation of these functions is defined by the dissemination strategy of the scope for which these functions are realized, establishing a functional scoping through which these distinct functions can be optimized within a given strategy.

The invariants can be put together in a functional model, as shown in Fig. 1. It shows the functional and information scoping that is provided through combining the first, second and fourth invariant. It also exposes the service model, defined in our third invariant, as an API to other functional parts of the system. Through sub-scoping, our functional model enables the recursive assembly of distributed systems, as indicated in the nested structure in Fig. 1, within a given dissemination strategy. This brings us to the assembly of end-to-end solutions. Particularly our functional model can be assembled towards more complex network structures by virtue of the following end-to-end principle:

> *A given problem within a distributed system can be implemented through an assembly of sub-problem solutions, whose individual dissemination strategies are not in conflict with the ones set out by the problem in question.*

In other words, distributed solutions within our architectural framework are realized through a process of reconciling various (possibly conflicting) dissemination strategies. We foresee approaches for such reconciliation being rooted in the processes of engineering solutions in design time, utilizing standardization and specification of functional components in the architecture (e.g., responsible for handover in mobility). But we also see the opportunity to formally specify dissemination strategies as information itself and therefore utilize runtime reconciliation of strategies, e.g., when handing over to new technologies that expose these formal strategy descriptions. Due to the lack of space, we leave out the detailed nature of processes for reconciliation at runtime.

IV. MAPPING OUR MOBILITY CONCEPTS

When addressing the mobility concepts in an IP context, we have to remember that the network is agnostic to the information that is being exchanged by the interested parties. As a result, the only action that needs to be undertaken when nodes move is to efficiently restore the end-to-end communication between the moving nodes. Our information-centric angle changes the perspective of how to handle mobility of participating network nodes by focusing on disseminating information rather than mere end-to-end exchange of an opaque bit stream. With that in mind, we first address in the following paragraphs how the three main mobility concepts described in Section II are specifically supported by our ICN functional model. Then, we briefly describe a reactive and a proactive approach that facilitate node mobility in accordance to this functional model

A. Supporting Mobility Concepts within the Functional Model

Regionalization is addressed through the ability to define a coherent structure of information dissemination. This is in direct support of the ideas in Sollins [16] in which the concept of regions is generalized to include those of mobility-related regions, e.g., locality, as well as regions of trust and security. Such a region in our framework is constituted by a single functional component in our model with its own dissemination strategy that can be optimized for a particular purpose. For instance, the aspect of router discovery can be captured through a localized region across various access technologies. Information within that region traverses the various underlying access technologies through an appropriate dissemination strategy, very much alike the physical multicast overlay solution proposed in [17]. More specifically, the generalized view of a region can be seen as an information structure, which is published by network points where MNs can be attached to or by network points representing them. Then, discovering the region is a matter of subscribing to the appropriate information sub-graph(s), given that the dissemination strategy for this set of information allows it (e.g., when part of the region belongs to other administrational domains). The former also better adapts to software mobility; software-based mobile agents (MAs), such as JADE [18] and Aglets [19], tend to move fast, autonomously and spontaneously, from one region to another.

Dissemination of information across various regions, e.g., for buffer relocation or HA approaches for reachability, can be achieved through reconciling the various dissemination strategies of the involved regions in accordance to our end-to-end principle. Such reconciliation can be optimized, e.g., for information dispersion akin to a buffer relocation in Mobile IP. For instance, when publishers move, the Rendezvous function in coordination with the Topology Formation may instruct other publishers of the same information to publish data to the current set of subscribers.

It is important to understand that a solution to the information dissemination problem is easier to optimize since many (potentially cooperating) parts of the network are aware of the information items exchanged among interested parties. Establishing context is supported by the concept of metadata that accompanies the information structures as defined through our first two invariants. In order to emphasize the information-centric aspect of context transfer in the mobility problem, we quote the definition of “context” from the respective RFC [4]: “The information on the current state of a service required to re-establish the service on a new subnet without having to perform the entire protocol exchange with the mobile host from scratch”. It is obvious that based on our two first invariants the context itself can be represented within the information structure for which a node acts as a publisher or subscriber. When subscribers move within an intra-domain scenario, information structures pertaining to current subscriptions can be transferred to the new AP (if already known) or to a set of APs that act as proxy caches that buffer data while the MN is moving[1]. With that, the publish-subscribe model allows network elements (e.g., APs or fixed publishers and subscribers) as well as the MNs themselves to update this information whenever required by the context. In addition, metadata can be attached to various points in the information structure, such as the super-scope of a particular part of the structure but also individual information items. Furthermore, the reconciliation of dissemination strategies across information structures, e.g., in the context of a handover from one technology towards another, also reconciles the access to the metadata that comes with the reconciled information structure.

B. Solutions in Accordance to the Functional Model

Let us now describe two approaches that are in accordance with the functional model of section III. Both approaches are based on subscription proxies for seamless mobility support of mobile subscribers. The role of a proxy is to handle subscriptions on behalf of MNs and to buffer - i.e., to cache - data that correspond to matched subscriptions when a MN moves to another proxy.

1) Proactive Approach: The first approach [20] falls under the family of *proactive approaches*, i.e., approaches based on *prefetching subscriptions* to neighboring proxies. We use the term *neighbor* to denote a proxy that is one-hop-ahead from the MN. When a MN moves from its current proxy to one of the proxy-neighbors, it immediately receives data that were transmitted while MN was moving. This is an ideal approach for delay-sensitive applications such as real-time and streaming applications; only it comes with costs related to both bandwidth and memory consumption for prefetching subscriptions and buffering data.

Let us consider a MN that is currently connected to some proxy i. Proxy i knows each probability p_{ij}, i.e., each probability that MN disconnects from i and moves to a neighboring proxy j. It is feasible to assume that i has a priori knowledge of such probabilities, for instance, by maintaining a history of handoffs from i to each neighbor proxy j. Alternatively, the work in [23] utilizes knowledge of the user's context (e.g., location, habits, etc) to pre-allocate buffer resources. Based on such knowledge, the current proxy i selects the subset of neighbor proxies $S*$ to which it proactively sends the subscriptions it handles for the MN.

Eq. 1 shows how to compute P_{hit}; P_{hit} is the probability that MN connects to a proxy in a given subset of neighbors S. Amongst all subsets of neighbors, i selects $S*$ which implies the probability $P_{hit}*$ that minimizes the cost function presented in Eq. 2. Note that the resulting $N(P_{hit}*)$, i.e. the size of $S*$, must be equal to the minimum number of neighbors to achieve $P_{hit}*$.

The process of computing $N(P_{hit})$ is straight forward. First, we order all neighbors j with descending probability order p_{ij}. Next we compute a stepwise function $N\big(\sum_{k=1}^{n} P_{ik}\big) = n$ where $n = 1, 2, \ldots, |J|$. Given the descending order of probabilities, each value of $P_{hit} = \sum_{k=1}^{n} p_{ik}$ uniquely determines a subset of neighbors $\{1, 2, \ldots, n\}$. Finally, we compute $P_{hit}* = \sum_{k=1}^{n*} p_{ik}$ (and thus $S* = \{1, \ldots, n*\}$), i.e., the P_{hit} value that minimizes the cost in Eq. 1 by searching over all discrete values of P_{hit}.

$$P_{hit} = \sum_{j \in S} Pij \tag{1}$$

$$P_{hit} \times C_{hit} + (1 - P_{hit}) \times C_{miss} + N(P_{hit}) \times C_{cache} \tag{2}$$

Regarding the individual costs presented in Eq. 2, C_{hit} is the cost for buffering subscriptions and data that match subscriptions at selected neighbors; thus, the probability for such a cost is P_{hit}. On the contrary, C_{miss} represents the cost for receiving data from their original publisher(s) in case the MN moves to a proxy that is not a member of $S*$; therefore, C_{miss} is always higher than C_{hit}. The respective probability for C_{miss} is equal to the probability of the next proxy not to be selected in $S*$, i.e., $1 - P_{hit}$. Unlike the previous costs that can be expressed in terms of latency or bandwidth consumption, C_{cache} represents the memory cost for caching subscriptions and their matching items respectively.

When a MN disconnects from its current proxy i, it publishes a notification to all neighbor proxies $j*$ in $S*$ to start buffering items that match the subscriptions of the MN. Next, i starts to publish to each $j*$ the items it receives while the MN is still disconnected. This ensures that when the MN connects to some new proxy in $S*$, the latter will have already buffered all items transmitted during the disconnection period of the MN. Upon connecting to the new proxy, the MN publishes another notification to both its former proxy and the (rest of the) members in $S*$, informing them to stop buffering items.

2) Reactive Approach: The second solution [21] falls into the category of *reactive approaches*. Particularly using, as in the proactive approach, a portion of each proxy's buffer we allow proxies to manage subscriptions and publications on behalf of the MN, when the MN is disconnected and during the switch over phase (reconnection phase).

When the MN is connected, it receives publications directly from the network. Before detaching, the MN publishes a notification to the proxy, that is attached to, informing it that will detach. The corresponding proxy informs the rendezvous function of the network that he is responsible for the detached MN. Particularly, the proxy sends a subscription message to the rendezvous function with the exact interest with the detached MN. Also informs the rendezvous function that is also a publisher for those publications that will receive. From that moment the proxy stores the publications that match the subscription of the detached MN. The corresponding proxy also forwards those publications to the rest of the connected subscribers with a matching subscription. When the MN reconnects to the network, from a different proxy, issues the same subscription as before. The rendezvous function together with the topology manager, apart from directing to the MN any new matching publication, also notifies the proxy to publish the stored publications to the MN. The proxy also unsubscribes from the rendezvous function the interest of the re-connected MN. It also erases those messages cached for the MN and stops acting as a publisher for them.

In order to make the mobility support robust, we should equip the proxies with a preemption priority mechanism for those publications that are stored for a MN. Using such a mechanism, these messages are stored in a FIFO manner disregarding the rest of the publications, which contend only for the remaining slots. Thus, a message stored for a MN is only dropped from the store when the store is full with such

---

[1] An inter-domain scenario can be implemented by reconciling the strategy for disseminating information according to the inter-domain policies.

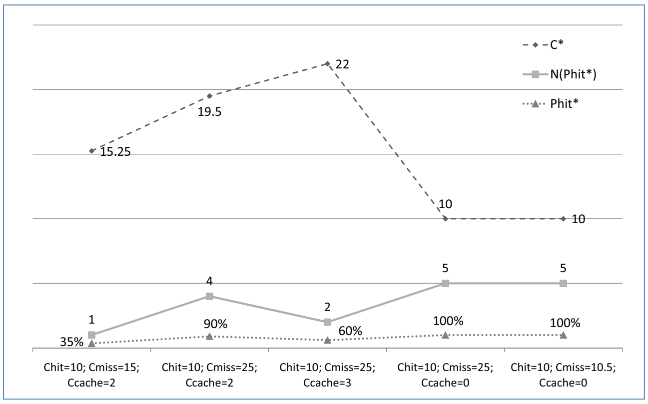

Fig. 2. Proactive approach. Demonstration scenario.

publications, which arrived later than the given one.

V. PERFORMANCE EVALUATION

In this section we provide a practical demonstration and an evaluation of the proactive and the reactive mobility support solutions presented in Sections IV-B1 and IV-B2 respectively. Regarding the proactive approach, let us assume the following scenario for demonstration purposes: The current proxy has five neighbors in descending order of their probabilities: $P_1 = 35\%$, $P_2 = 25\%$, $P_3 = 15\%$, $P_4 = 15\%$, $P_5 = 10\%$. We assume five different configuration scenarios for C_{hit}, C_{miss} and C_{cache}. The diagram of Fig. 2 shows the respective results for $C*$, P_{hit} and $N(P_{hit})$ in each of the five different configurations shown on the x-axis of the diagram. This diagram is interesting due to how and why the results on $C*$, P_{hit} and $N(P_{hit})$ change when moving from one scenario to the next by altering the individual cost values in the x-axis. Comparing the results of the second and the first configuration, we see that a relatively high increase in C_{miss} results in increasing $N(P_{hit}*)$ and consequently $P_{hit}*$. Furthermore, this need for selecting more neighbors in $S*$ increases the optimal cost $C*$ computed. On the contrary, a slight increase in the memory cost in the third configuration, results in a smaller $N(P_{hit}*)$. In this case, caching memory becomes more expensive relatively to the benefit from avoiding the retransmission cost by pre-buffering items to more proxies.

In both the forth and the fifth configuration we assume that memory is “costless” and as a result, all neighbors are selected for $S*$. In such cases, fetching and buffering items in advance is always beneficial for C_{miss} is always higher than C_{hit}; assuming that C_{miss} is less than or equal to C_{hit} is logically incongruous because it would imply that it is cheaper to transfer items from their original publisher rather than a buffer that is closer to the MN. For the same reason, decreasing C_{miss} in the fifth configuration to a value that is close to Chit results again in selecting all neighbors for $S*$ as buffering data is costless.

Regarding the reactive approach we implemented it on top of REDS [22] using five computers. The three of them formed the Information Centric Network (i.e. brokers) while the other two played the roles of the publisher and the mobile subscriber (MN) respectively. In the examined scenario the MN issues one subscription while a series of publications is sent (matching the MN’s subscription) at a rate of 2 publications per second. The whole experiment lasts 50 seconds while the MN disconnects from one proxy and reconnects to another after a fixed interval of 15 seconds. Fig. 3 presents the publish/delivery time trace in the cases where mobility is not supported (left graph) and in the case where the reactive mobility support is enabled (right graph). The x axis corresponds to the publish time of the messages (time left the publisher, set time to zero when the first message is published), while the delivery time (y axis) is when the message is received by the MN. Every point in the figures corresponds to a message received by the MN either through the original publish or after the reconnection process. The part of the figures where there is no message delivery represents the time interval that the MN is disconnected from the overlay network, while the vertical part of delivered messages after the reconnection of the MN represents the messages delivered to the MN upon its reconnection. It is obvious that all the published messages finally delivered to the MN when the mobility support mechanism is enabled, while when it is not enabled 30 messages are lost. This amount of messages corresponds to the 15 seconds that the MN was disconnected and the published messages were not cached for him.

VI. CHALLENGES IN EVALUATION OF MOBILITY IN ICN

There is a range of challenges in evaluating solutions for mobility in ICN. In the following, we address a few of them, defining a research agenda for future work.

The first challenge is that of finding appropriate *traffic models* for evaluation. While it is difficult to predict future applications, it is apparent that a simple flow-based traffic model will not suffice in capturing the characteristics of underlying traffic that is being generated by ICN applications. For instance, event-based communication patterns with possibly very large subscriber groups are more crucial to evaluating the underlying system functions. Furthermore, the publish-subscribe service model allows for decoupling time and space in communication, a situation that any suite of traffic models needs to capture. Such decoupling is particularly important for applications that try to utilize tradeoffs between storage and bandwidth usage.

The second challenge relates to *mobility patterns* that can be observed. While end node mobility still prevails in ICN, with similar issues as in IP mobility, information mobility poses a different set of challenges. Here, publishers and subscribers can move about in a way not associated to (physical) mobility of endpoints. This poses new challenges on regionalization and context transfer since changing regions raises issues upon locality, trust, security and the ability to reconcile different dissemination policies. It even touches upon policies that relate to governance and provenance of information, e.g., in scenarios where information provisioning moves to a different legislative

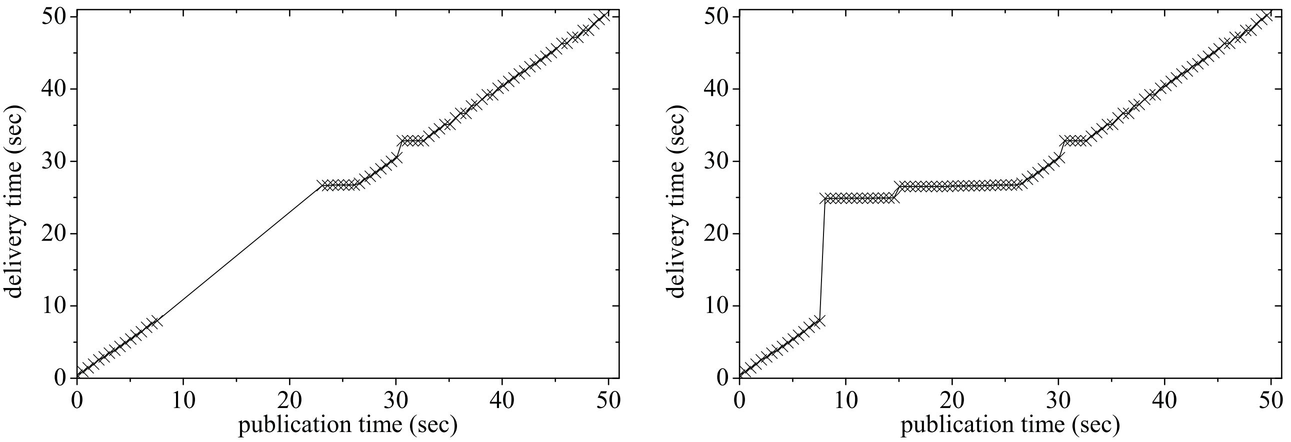

Fig. 3. Reactive approach: no-mobility support (left), mobility support (right).

domain. While caching might alleviate some of these issues, for instance through pre-caching information items within a region due to recognized demand in the population, generally taking into account these issues is a difficult challenge. One area where new mobility patterns are expected is that of cloud computing, for instance.

This leads us to the third challenge, namely that of *defining appropriate metrics* for mobility performance. While it is generally accepted that minimizing handover failures is a desirable goal in IP mobility, such efficiency metric is more complicated to get by in ICN mobility. Take the example of utilizing in-network caching in ICN and assume an equivalent to buffer relocation in IP, as outlined in Section IV-B and evaluated in Section V. Although a handover of a subscriber might fail with respect to, e.g., a real-time constraint, the resulting information transfer into a different cache is possibly aiding future handovers for subscribers to the same content within the region. Hence, the challenge here is that of weighing off the potential of improving overall throughput utilization for an entire network region against the handover optimization of a single subscriber.

Finally, there is the challenge of *evaluating mobility solutions*. While real deployment remains a viable evaluation option, emulation and simulation in large scale is likely to increase in importance when it comes to evaluating mobility in ICN. This is due to the possible non-local mobility patterns that can be observed, for instance in cloud computing scenarios where resources move about on global scale without any attachment to a localized region. With this comes the problem of finding the right level of abstraction required for simulations. Given the potential scale, packet-level simulations are unlikely to perform, requiring inter-domain or network-level abstractions.

VII. CONCLUSIONS

Mobility is an increasingly important aspect of today's networking and even more so in any proposal for the Future Internet. In this paper, we revisit some of the previous IP mobility work under the angle of an information-centric way to networking. To aid our presentation, we outline the crucial concepts of regionalization, information dissemination, and context establishment within network regions. We furthermore outline an approach for designing information-centric networking solutions that help us position the mobility concepts in an ICN context.

Our contribution in this paper aligns the existing IP mobility work with these mobility concepts within the ICN context. This alignment provides an initial step towards a coherent and substantiated theoretical basis for supporting seamless mobility in an information-centric internetworking architecture. Furthermore, our work pulls these mobility concepts down from a publish-subscribe overlay approach onto the network level of a new information-centric internetworking architecture. With this, we see the potential to improve the mobility support within this type of architectures by drawing from potential benefits when shifting from the current endpoint-centric paradigm to an information-centric one.

However, investigating mobility in ICN, in particular with prior knowledge of mobility solutions in IP, is not without challenges. We highlight these challenges, ranging from traffic models over regionalization to efficiency metrics. With that, we define a research agenda that will need addressing in future mobility work to come.

ACKNOWLEDGMENT

The work reported in this paper was supported by the FP7-ICT-2010-257217 PURSUIT program.

REFERENCES

[1] Named Data Networking project, at http://named-data.net, 2011.
[2] PURSUIT project, at http://www.fp7-pursuit.eu, 2011.
[3] SAIL project, at http://www.sail-project.eu/, 2011.
[4] J. Loughney et al. "Context Transfer Protocol (CXTP)", IETF, July 2005, RFC 4067.
[5] M. Liebsch, et al., "Candidate Access Router Discovery (CARD)", IETF, July 2005, RFC 4066.

[6] D. Johnson, C. Perkins, and J. Arkko, “Mobility Support in IPv6”, IETF, Jun 2004, RFC 3775.

[7] C. Perkins, “IP Mobility Support for IPv4”, IETF RFC 3344, August 2002.

[8] G. Cugola, E. Di Nitto, A. Fuggetta, “The JEDI event-based infrastructure and its application to the development of the OPSS WFMS,” in IEEE Trans. On Software Engineering, vol. 27, pp. 827-850, Sept. 2001.

[9] C. Castelluccia, “HMIPv6: A hierarchical mobile IPv6 proposal”, in ACM SIGMOBILE Mobile Computing and Communications Review, v.4 n.1, p.48-59, Jan. 2000.

[10] M. Caporuscio, A. Carzaniga, A. L. Wolf, “Design and evaluation of a support service for mobile, wireless publish/subscribe applications” in IEEE Trans. on Software Engineering, 29(12):1059-1071, December 2003.

[11] L. Fiege, F. Gartner, O. Kasten, A. Zeidler, “Support-ing mobility in content-based publish/subscribe middleware”, in Middleware 2003.

[12] Y. Huang, H. Garcia-Molina, “Publish/subscribe in a mobile environment”, in Wireless Networks 10, pp. 643-65,2 2004.

[13] V. Muthusamy, M. Petrovic, D. Gao and H. Jacobsen “Publisher mobility in distributed publish/subscribe systems”, in 25th IEEE International Conference on Distributed Computing Systems Workshops, pp. 421-427, 2005.

[14] L. Mottola, G. Cugola, G. Picco, “Tree-based overlays for publish-subscribe in mobile ad hoc networks”, Technical report, 2005.

[15] N. Fotiou, G. C. Polyzos, D. Trossen, “Illustrating a Publish- Subscribe Internet Architecture”, in Future Internet Architectures: New Trends in Service Architectures, 2009.

[16] K. Sollins, “Designing for scale and differentiation”, Proc. of the ACM SIGCOMM FDNA, October 2003.

[17] D. Trossen, H. Chaskar, G. Krishnamurthi, “Enabling Ubiquitous Computing in Heterogeneous Environments by Physically Scoped Multicast,” Proc. of IEEE 3G Wireless, May 2002, San Francisco, CA, USA.

[18] JADE, Java Agent Development Framework, at http://jade.tilab.com.

[19] D. Lange and M. Oshima, “Programming and Deploying Java Mobile Agents with Aglets,” Addison-Wesley, 1998.

[20] V. Siris, X. Vasilakos, and G. C. Polyzos, “A Selective Neighbor Caching Approach for Supporting Mobility in Publish/Subscribe Network”, in 5th ERCIM Workshop on eMobility (in conjunction with WWIC 2011),Vilanova i la Geltru’, Catalonia, Spain, 2011.

[21] V. Sourlas, G. S. Paschos, P. Flegkas, and L. Tassiulas. “Mobility support through caching in content-based publish/subscribe networks,” in IEEE/ACM Int’l Conference on Cluster, Cloud and Grid Computing, 2010.

[22] Cugola G. and Picco G., “REDS, A Reconfigurable Dispatching System,” in 6th International workshop on Software Engineering and Middleware, pp. 9-16, Oregon, USA, 2006.

[23] D. Pavel, D. Trossen, “Context-aware Resource Pre-Allocation in Nomadic Applications using a Service Provider Approach,” in IEEE 3G Wireless 2003, May 2003, San Francisco, CA, USA.