Supporting Mobility in a Publish Subscribe Internetwork Architecture

Varvara Giannaki, Xenofon Vasilakos, Charilaos Stais, George C. Polyzos, George Xylomenos
Mobile Multimedia Laboratory
Department of Informatics
Athens University of Economics and Business
Athens, Greece
{bgiannaki, xvas, stais, polyzos, xgeorge}@aueb.gr

Abstract — **Information-Centric Future Internet architectures are constantly gaining momentum within the Future Internet research community. PURSUIT is an EU-FP7 research project developing a clean-slate Pub/Sub Internetwork (PSI) architecture. PSI (or Ψ) envisions a pure information-centric Internet architecture that has integrated seamless mobility support. The novel information-centric mechanisms supported in Ψ, along with smartly placed in-network caches, enable the architecture to handle both mobile and fixed devices in a uniform way. This paper presents a blueprint for optimizing mobility support in Ψ without applying any modifications or add-on solutions. We demonstrate a micro-mobility scenario that describes the functionality of Ψ core components, and sketch our plans on future work and a proper assessment of our ideas.**

Keywords-component; Future Internet, Information-Centric Networks, Mobility

I. INTRODUCTION

The current Internet architecture is based on End-to-End point (E2E) communication, covering the needs of the time when it was designed. The former is widely considered as the root cause of many of its limitations and inefficiencies. Add-on solutions such as NATs, CDNs or Mobile IP were introduced in order to mitigate problems, even thought they tend to violate aspects of the Internet architecture. Publish/Subscribe (Pub/Sub) continuously gains popularity within the research community as a clean slate Future Internet (FI) architecture due to its inborn information-centrism that manages to decouple information from its location. Research projects such as PURSUIT [1, 3], PSIRP [2, 3], and CCNx [4] investigate Pub/Sub as a FI architectural alternative. Unlike the current Internet, Pub/Sub fulfills the core need of users who are interested in the information itself rather than where the information resides. Further advantages include its multicast nature, its provision of anonymity and its inherent asynchrony, constituting Pub/Sub as an ideal paradigm for mobile environments [12]. First, anonymity and asynchrony enable the quick adaptation to the continuous attachment and detachment of *Mobile Nodes* or *Agents (MNs or MAs)* to *Access-Points (AP)* that reside at the edge of a mobile network. Note that APs are not necessarily wireless; they may refer to any kind of a physical or software mobile layer, including wireless environments such as IEEE 802.11 or UMTS, or to a software MA that migrates to another physical host. Second, the multicast nature of Pub/Sub can deal with large scales of MNs and their tendency to change locations continuously. Last, given the limited capabilities of MNs such as short battery life or bandwidth, multicast helps systems to better utilize their resources, e.g., by performing fewer (re)transmissions that imply energy savings.

This paper presents an on-going effort for enhancing the inborn mobility support of the Pub/Sub Intenetworking architecture (PSI or simply Ψ) [3]. Our work is within the context of the PURSUIT project, an EU-FP7 funded research project for a clean-slate, information-oriented FI architecture. We first discuss the related work; then we present an overview of Ψ and proceed by explaining the insights of our proposed enhancement for seamless mobility support in Ψ. Finally, we conclude and refer to our future plans.

II. RELATED WORK

There are several notable studies on mobility within the context of information-centric architectures. i3 [5] is an overlay network that conceals E2E communication by applying a rendezvous-based model on top of IP. Senders send packets to a specific *rendezvous point* and receivers issue *triggers* on specific packet identifiers. i3 is mobile friendly, it protects location privacy and is shown to perform better than Mobile IP in terms of stretch and fault tolerance [6]. Ψ uses similar concepts through the rendezvous and topology formation processes, but unlike i3, Ψ is native to the network. ROFL [7] and DONA [8] use information identifiers rather than location identifiers or triggers. DONA's discovery mechanisms provide a facility for caches, for forming multicast delivery trees and anycast delivery, i.e., after locating the nearest information source. ROFL uses Distributed Hash Tables and their hierarchical versions. Node "labels" are location-independent identifiers, matching the nature of MNs that are specially treated as ephemeral hosts, participating as little as possible to the overall architecture. This is a major difference from Ψ, for we do not distinguish between stable and MNs in an effort to apply a universal property to our architecture. Assuming another solution, CCN supports mobility by caching and allowing nodes to transmit "interest" packets simultaneously from multiple interfaces, assuming an unreliable transport delivery service [9]. Nonetheless, there is no notion of a rendezvous scheme as is in Ψ; the unstructured organization of information advertised by flooding, poses questions on the potential impact on mobility support.

Authors in [10] describe the most relevant study to ours. Their mobility solution for Pub/Sub networks uses a *neighbor graph* per broker, i.e., a data structure that contains candidate brokers for caching subscription context. Brokers convey subscriptions to other brokers within a single-hop distance. Contrary to authors' statements, forming the neighbor graph requires topological knowledge as subscriptions carry the address of the originating broker of the MN. We prefer to take a different approach by using a Ψ compatible mechanism which allows MN to move without suffering any loss in data and without the need of any control messages.

III. Overview of Ψ and Mobility Support in Ψ

Pub/Sub architectures involve three major entities: the *publisher*, the *subscriber* and the *event notification service*. A Publisher advertises the availability of specific information items by issuing *publication messages*. Subscribers express their interest for consuming specific information items by issuing *subscription messages*. The event notification service locates the publishers who provide desired information items by matching the consumers' subscriptions; then it initializes the forwarding information content from publishers to respective subscribers. Publications and the subscriptions *do not* have to *be issued in sync*. Moreover, publishers and subscribers do not know the identities of each other (anonymity property), nor do they need to be concurrently connected to the mobile network (asynchrony property).

A Rendezvous Network (RENE) incarnates the event notification service. It is composed of several *rendezvous nodes (RNs)*, each of which is responsible for a set of publications. A Rendezvous Point (RVP) of a specific publication is the RN that is responsible for the latter. As in [14], Ψ uses (statistically) unique *labels* for each discrete information item, composed by a pair of flat, semantically free identifiers: the *Rendezvous Id (Rid)* and the *Scope Id (Sid)*. The RId derives from an application specific function applied to the advertised information item, while the SId corresponds to a scope. *Scoping mechanisms* in Ψ are used to control access to information and to enforce policies and strategies on information published within a certain scope. Scopes employ a hierarchical structure where sibling and parent-to-children relationships exist. A subscriber expresses her interest for a specific publication by issuing a subscription message which the RENE delivers to a RVP. Upon receiving and matching a subscription, the RVP initiates the process of constructing a forwarding path from a publisher to one or more subscribers. A *Forwarding Identifier (FId)* is a bloom filter based structure that includes all the identifiers of the links in the forwarding path that data need to traverse in order to be delivered to the subscriber(s) [11]. As a result, publishers can start sending data without knowing who the subscribers are or where they reside.

The Ψ architecture is ready to support MNs since all of its functionalities are location-independent and location-agnostic. We introduce the concept of *Smart Caches (SC)* that assist and enhance mobility management. Given a publication, SCs are smartly placed *in-network caches* designated by the RVP with respect to the (1) topology and the (2) current and expected future attachment positions of subscribers. A SC acts as an intermediate node that apart from caching, it mediates between the data exchanging parties. Therefore it is treated as both a publisher and a subscriber. In the following, we describe a micro-mobility scenario which involves the usage of SCs, and comment on the behavior of the basic network components of the architecture.

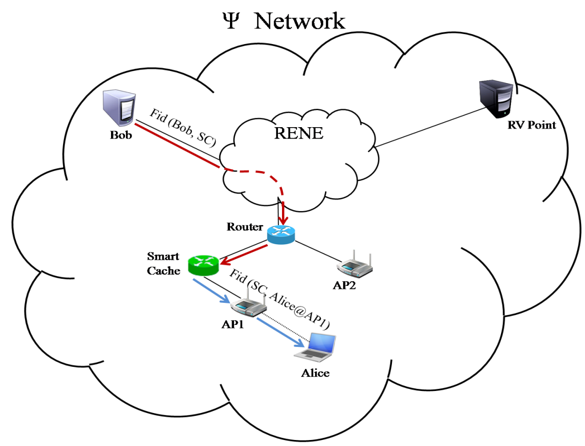

Figure 1: Bob sends data to the SC; the SC resubmits data to Alice .

A. A Simple Use-Case Scenario

Bob is a *fixed* -node that issues a publication <*Bob_Sid, Bob_Rid, [metadata]*>. Alice is a MN[1] currently attached to AP1, who issues a subscription for <*Bob_Sid, Bob_Rid*>. Both the publication and the subscription find their way to RVP through the RENE. The earlier designates a SC based on the topological knowledge of the topology management function [13]. The goal is to facilitate the delivery of data to Alice and any other prospective mobile subscribers in her network proximity. Upon matching a subscription, the RVP issues two forwarding paths: (1) Fid(Bob, SC) that defines the path from Bob to the SC ; (2) Fid(SC, Alice@AP1) that defines the path from the SC to the AP1 where Alice is currently attached to. As shown in figure 1, Bob starts sending data to the SC, which in turn resubmits the data to Alice using FId(SC, Alice@AP1). SC gets data as a subscriber and caches them for at least as much time as it is required for Alice to move to a near-by AP such as AP2. Since this is a micro-mobility scenario, the RVP can guess the APs that Alice might attach herself to next, estimate the caching period and smartly place the SC in order

[1] Our idea treats MNs and software MAs in the same way, thus referring to MNs is equivalent to referring to MAs and vice versa for the purposes of this paper.

to accommodate not only Alice now, but also in some future possible attachment points for Alice or other mobile subscribers.

At some point Alice decides to move to AP2 while data are still in transit to AP1 (figure 2). Upon performing the handoff to AP2, Alice re-subscribes to <Bob_Sid, Bob_RId>. Note that the RVP has currently recorded two different publishers whose publications match this subscription; the one is Bob and the other one is the SC. The RVP makes an anycast choice between the two and thus instructs the *best suited publisher* with respect to the current location of Alice. SC is likely going to be the chosen anycast source of data, for it was originally designated in order to accommodate future possible attachment points for Alice. Note that (1) data are still sent to AP1 via a multicast tree until all past or future subscriptions submitted from AP1 expire and that (2) data that were sent to AP1 during the hand-off phase to AP2 are not lost; they are cached in SC and later forwarded to AP2. Alice is now attached to AP2 and the RVP can either select a new SC for Alice or instruct Bob to directly send the data to Alice in AP2.The fact that no data in transit are lost while Alice moves is particularly important for streaming and real-time applications where SCs can be used to redirect such data to the current position of Alice

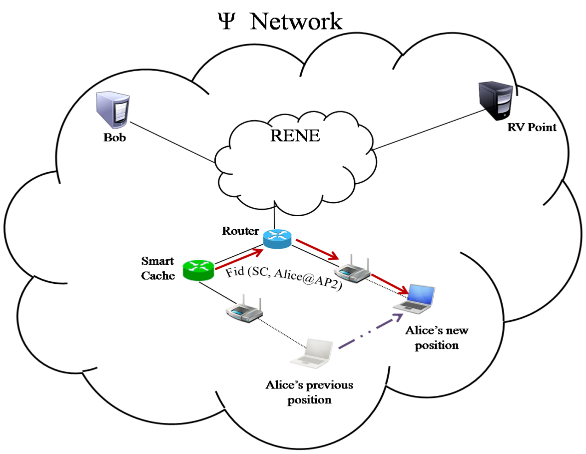

Figure 2: Alice moves while data in transit. Data are sent to both AP1 and AP2 via a multicast tree

SCs *are neither* a new type of component nor an add-on to the Ψ architecture which has integrated support for mobility [12]. SCs are more likely to be existing ordinary caches for frequently requested content; only we "promote" a simple cache to perform as another publisher. In cases where a cache is absent from the network, the RVP delegates the role of a SC to any node that is capable of caching the content. Such nodes could be routers with caching capabilities. As a byproduct, SCs also enhance anonymity by mediating between publishers and subscribers.

B. *Smart Cache Selection Mechanisms*

We consider two different selection mechanisms for delegating the role of a SC: *RVP forecasting (RVPf)* and *AP based (APb)*. In both cases, the RVP treats the SC as another publisher and subscriber for information data. The two mechanisms differ with respect to who initiates this procedure and when this action is performed. RVPf assumes that the RVP has an up-to-date knowledge of the current topological organization of the network via a topology-manager function [13]. Consequently, RVPs can forecast the candidate future position of MNs in micro-mobility handoffs. Thus it is feasible to designate SCs before the subscribed MNs move. Any data that were originally sent to the AP are also sent via a multicast tree to the new SCs.

Unlike that, the underlying idea in APb is that APs in wireless environments can detect a reduction of the signal strength when a MN moves away of it. The currently used AP as well as a prospective AP may infer the tendency of the MN to move. The currently used AP sends a control message to the RVP, requesting the reconsideration of the current SCs and perhaps the designation of new SCs. Once the RVP receives this control message, the procedure followed is identical to what is described in the previous sections (adding a new SC as a publisher to the RVP, create the suitable FIds to send data via new SCs, etc.). This selection mechanism adds some extra control messaging for every hand-over, which is a small overhead, compensated by taking load off the RVP. This is quite important, since the RVP performs the matching of subscriptions to publications; overloading it could cause delays to rendezvous and consequently to data delivery. Moreover, APb allows the RVP to take action only *when the MN actually moves* to another AP without the need to run the topology function for every MN in order to see its possible next positions and delegate SCs. Unfortunately, APb is not suitable for non-wireless mobile networks while AVPf is.

IV. Conclusions and Future Work

This paper presents an enhancement concept for mobility in Ψ, based on a smart placement of caching elements. Such SCs guarantee a seamless delivery of data to mobile subscribers. Our work is part of on-going research towards a uniform Pub/Sub Internetworking architecture for both mobile and fixed environments. We plan to investigate the correctness and effectiveness of our proposed solution, and to study any further possible uses of SCs, e.g., as transport protocol mediators enhancing anonymity of principals and facilitating transport mechanisms. Last, we will proceed in evaluating the performance of our ideas via simulations of Ψ architectures that incorporate SCs. This will help us verify the suitability of our suggestion and further study scenarios that involve mobile publishers.

References

[1] PURSUIT: Web site, 2010. http://www.fp7-pursuit.eu/

[2] PSIRP : Web site, 2010. http://www.psirp .org

[3] Fotiou, N., Nikander, P., Trossen, D., and Polyzos, C., G: Developing Information Networking Further: From PSIRP to PURSUIT". Proc. 7th International ICST Conference on Broadband Communications, Networks, and Systems, 2010

[4] CCNx: Web site, 2010. http://www.ccnx.org

[5] Stoica, I., Adkins, D., Zhuang, S., Shenker, S., Surana, S: Internet indirection infrastructure. IEEE/ACM Trans. Netw. 12(2), 205-218, 2004

[6] Zhuang, S., Lai, K., Stoica, I., Katz, R., Shenker, S: Host mobility using an internet indirection infrastructure. Wirel. Netw. 11(6), 741-756, 2005

[7] Caesar, M., Condie, T., Kannan, J., Lakshminarayanan, K.,Stoica, I: Rofl: routing on flat labels. In: SIGCOMM '06: Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 363-374. ACM, New York, NY, USA, 2006

[8] Koponen, T., Chawla, M., Chun, B.G., Ermolinskiy, A., Kim, K.H., Shenker, S., Stoica, I: A data-oriented (and beyond) network architecture. In: SIGCOMM '07: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 181-192. ACM, New York, NY, USA, 2007

[9] Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggs, N.H., Braynard, R.L: Networking named content. In: CoNEXT '09: Proceedings of the 5th international conference on Emerging networking experiments and technologies, pp. 1-12. ACM, New York, NY, USA, 2009

[10] Gaddah, A. Kunz, T., “Extending mobility to Pub/Sub systems using a proactive caching approach”. Mobile Information Systems 6 (2010) 293-324, 2010

[11] Jokela, P., Zahemszky, A., Esteve Rothenberg, C., Arianfar, S., Nikander, P: Lipsin: line speed Pub/Sub internetworking. In: SIGCOMM '09: Proceedings of the ACM SIGCOMM 2009 conference on Data communication, pp. 195-206. ACM, New York, NY, USA, 2009

[12] Katsaros, K., Fotiou, N., Polyzos, C., G., Xylomenos, G: "Overlay Multicast Assisted Mobility for Future Pub/Sub Networks," ICT Mobile Summit, Santander, Spain, June 2009

[13] Särelä, M., Rinta-aho, T., Tarkoma, S., "RTFM: Pub/Sub Internetworking Architecture", In ICT Mobile Summit, 2008

[14] Tsilopoulos, C., Makris, D., Xylomenos, G., “Bootstrapping a Pub/Sub Information Centric Network”, to appear in Future Network & Mobile Summit, 2011