Abstractions and Solutions to Support Smart-Objects in the Future Internet *

Francesco Bronzino
WINLAB, Rutgers University
North Brunswick, NJ 08902
bronzino@winlab.rutgers.edu

Dipankar Raychaudhuri
WINLAB, Rutgers University
North Brunswick, NJ 08902
ray@winlab.rutgers.edu

ABSTRACT

As the number of devices that gain connectivity and join the category of smart-objects increases every year reaching unprecedented numbers, new challenges are imposed on our networks. While specialized solutions for certain use cases have been proposed, more flexible and scalable new approaches to networking will be required to deal with billions or trillions of smart objects connected to the Internet. With this paper, we take a step back looking at the set of basic problems that are posed by this group of devices. In order to develop an analysis on how these issues could be approached, we define which fundamental abstractions might help solving or at least reducing their impact on the network by offering support for fundamental matters such as mobility, group based delivery and support for distributed computing resources. Based on the concept of named-objects, we propose a set of solutions that network and show how this approach can address both scalability and functional requirements. Finally, we describe a comprehensive clean-slate network architecture (MobiityFirst) which attempts to realize the proposed capabilities.

CCS Concepts

•Networks → Naming and addressing; •Human-centered computing → Ubiquitous and mobile computing systems and tools; Ubiquitous and mobile devices;

Keywords

Internet; Mobility; Smart-Objects; Network Services; Network Abstractions; Named Objects; Future Internet Architecture; Virtual Networks

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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SMARTOBJECTS'16, October 03-07, 2016, New York City, NY, USA

© 2016 ACM. ISBN 978-1-4503-4254-4/16/10...$15.00

DOI: http://dx.doi.org/10.1145/2980147.2980153

1. INTRODUCTION

The smart-objects category covers a variety of different devices and models of communication, each one with multiple characteristics and requirements. Consider two classic use cases that belong to the category: a) a connected vehicle and b) an Internet-of-Things (IoT) sensor. While both fall under the same conceptual umbrella, they are characterized by very different models of communication: the first one is highly mobile, while the second one might be more static; the sensor has very limited computational and energy resources, the vehicle is not constrained in this sense. The car requires certain messages to have minimal latency to alert neighboring devices, the sensor can sustain higher delays if this provides gains in meeting its requirements. Without going into more details, it is clear that these devices are extremely multi-faceted and varied. Even with this in mind, it is possible to identify a series of common challenges and requirements that span across some, if not most, of the smart-object classes.

First, more and more devices are introducing new levels of mobility that characterize how networks are accessed today; mobile wireless devices have now outnumbered fixed end hosts and even service end-points. Smart-objects have also followed this trend and are now one of the largest contributors to the total number of portable/mobile devices connected to the mobile Internet. Moreover, their mobility patterns vary a lot as exemplified at the beginning of this introduction. Even though support for seamless mobility is a growing requirement for the Internet as a whole, past proposals and current solutions are either only applicable within limited environments (e.g., cellular [22]) or are inefficient when applied to the Internet (e.g., MobileIP [1]). A few recent scalable approaches to support mobility have been proposed, but these are not standalone and require changes to the routing plane and/or protocol stack defined in TCP/IP [12, 8].

Efficient transmission has been identified as a key requirement for different scenarios, including IoT communications, where efficiency stands for limited overhead imposed on the network and consumption of the battery of the devices, and vehicular scenarios, where efficiency stands for extremely low latency aimed at security scenarios. In order to reach the set goals, improvements to cellular access technologies for lower latency and reduced overhead have been identified as requirements. Current 5G pre-standards discussions are focusing on new technological advances in the radio department, to achieve ultra low latency [17] or drop the connection oriented approach of cellular networks [10]. While these approaches are valuable, a fundamental re-thinking of how data is delivered to these devices may be expected to provide important architectural benefits that go beyond simple performance improvements. For example, because most of the smart-object categories heavily rely on group based communications (e.g. IoT devices), developing novel and powerful schemes that, by means of multicast delivery will reduce overhead and transmission costs, will be extremely beneficial. Also, new context based communications are important for future services involving smart objects with awareness of time, location, network state, etc.

To tackle the different service requirements, in the last few years new cloud based technologies are being proposed. In particular, the continuing growth of mobile driven services has lead researchers to investigate more distributed solutions to better support newly introduced scenarios. An example of this is the edge cloud technology (sometimes referred to as "fog computing") [4], currently viewed as the next logical step in the evolution of mobile data services. While these solutions have high potential of succeeding, they still lack adequate network support to fully achieve their potentil through implementations that could reduce the deployment cost and decrease the access barrier for players interested in utilizing them.

To further complicate the matter, these three fundamental challenges are required to be solved at scale with the expected number of connected devices that will populate the Internet in the future. Current forecasts [3], for a not so far future, envision scenarios composed of tens of billions devices, which can be connected to networks or other devices, or whose meta data or statuses are handled in applications software.

While different ad-hoc solutions have been studied over the years focusing on one of these problems, or on a specific class of smart-devices, we believe that a comprehensive analysis of the common issues they pose is required to drive the evolution of the network architecture and allowing future users to fully achieve the potential of this growing category. With an understanding of these challenges in mind, this paper aims to provide an understanding of the potential abstractions that will be required to support the growing population of smart-objects in the network. The paper is structured as follows: Section 2 introduces the reader to the proposed abstractions and the impact they would have on the network communications between the interested devices. Section 3 then analyzes which approach could be taken to introduce new solutions to provide the abstractions to the network users. Section 4, finally, introduces MobilityFirst, a name-based architecture that can natively support such abstractions laying the ground for the future of the Internet communications. Section 5 concludes the paper.

2. ABSTRACTIONS

The ability to provide advanced communication abstractions will be a fundamental step to tackle the challenges introduced in the first section of the paper. The following paragraphs will list the set of abstractions that we believe can be the base for the right set of tools that will support smart-objects in the future Internet.

Objects based communications: Host-centric abstractions to network services were a solid building block during the conception and early advancement of the Internet. Though, other important principals such as content have emerged since then. While content and services are two established principals (besides hosts), others such as sensors and actuators, as also more abstract ones such as context are quickly gaining traction. Since few foresaw today's usage of the Internet with any accuracy, allowing for a broader definition should be the path forward. In that vein, direct addressability for all principals eliminates any unnecessary bindings of one principal with another. For instance, content should be addressable both directly and independently of where it may be located physically. Obviously, current networks are not able to sustain this type of abstraction, either due to scarcity of available names (IPv4) or due to their hierarchical nature (both IPv4 and IPv6). Allowing for direct addressability for not only content and services, but also other emerging first class entities requires that the name space be practically inexhaustible and to be independent of the network topology.

Mobility as the norm: Communication with a mobile end-point should be no different than that with a fixed end-point. The current "practical" approach results in undesirable asymmetry, where mobile end-points are always responsible for re-establishing connections. The situation is doubly vexing when both end-points are mobile. A basic service abstraction that allows addressing a network end-point by its unique name and not its current location will establish a uniform approach to dealing with fixed and mobile end-points alike, enabling seamless mobility. To support a name-based network service that seamlessly handles mobility, the network requires native support for dynamic and fine-grain location resolution. Some have proposed protocol interposition approaches that dynamically substitute local addresses for end-points with dynamically resolved ones [12]. Efficiency extensions in MobileIPv6 attempt to signal end-points with address updates [2] to avoid triangular routing through fixed home agents. However, we think that besides end-hosts, the network routing fabric must also be able to dynamically resolve and re-bind in-flight packets.

Group based delivery: Certain Internet devices are usually naturally part of groups that for different reasons might need to efficiently communicate between them. Consider for example Internet applications like video streaming, online gaming and social networks. Smart devices are no exception to this. Recent increases in network traffic associated with the growth of mobile devices, IoT devices, smart wearables and connected vehicles, motivate the need for efficient multicast delivery, a service that is not well-addressed through overlay solutions or by IP based solutions due to concerns of scalability which commonly cause them to be left disabled on network elements. Internet applications therefore regularly resort to multiple unicast packets to address groups. Providing streamlined ways for networks to support the creation and management of groups of member identities is a desirable abstraction. To support such a service abstraction, we must provide a means for end-hosts to specify the requested delivery service - i.e. multicast- that's interpreted by the network elements. The network should then dynamically handle the delivery, without requiring overly complicated maintenance schemes.

Distributed computational frameworks: Cloud com-

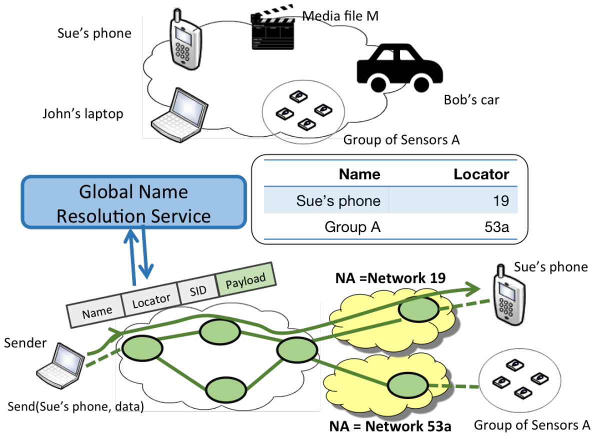

Figure 1: Name-Objects through name resolution

puting has revealed its usefulness over the years, providing access to vaste computational resources to devices that would be very limited otherwise. Thanks to this, more an more devices have increased their "smartness" gaining access to the smart-objects category, e.g. wearables and home automation devices. As the number of these devices grows and their mobility patterns become more and more of impact on their user experienced performance, more distributed solutions have been proposed. The use of cloudlets and edge cloud computing, sometimes referred to with other names, e.g. fog computing, has now become a core objective for many companies and researchers [4]. While these solutions are definitely attractive, they usually lack deep network support to understand fine-grained dynamics affecting network performance lacking access to routing information or acquire it through over the top solutions.

Allowing for additional control over routing decisions based on application information has been demonstrated to be a desirable tool [21, 24]. SDN has proven capable of handling similar problems but has failed to scale to multiple domains. Enabling network fabric through new interfaces to exploit such information would be highly beneficial, providing not only benefits on the applications, but on the network too [21]. Two core abstractions will then be required: 1) the ability to refer to a service as a single entity, allowing the network to handle the delivery through an "anycast" like paradigm and 2) the ability to exploit such information in the fabric itself, without overloading its resources.

3. NAMED-OBJECT BASED SOLUTIONS

Starting from the presented abstractions, we present in this section a set of solutions that are aimed at providing the foundation for three fundamental techniques that are vital toward the support of a large networking environment populated of smart devices: mobility, group based communications and edge cloud based QoS enhancements. At the base of our solution lies a fundamental rethinking of how hosts, devices and network elements in general are identified and communicate in the network. In contrast to the current IP based network, which tends to conflate both names and addresses, we create a new level of separation: names are flat globally unique identifiers (GUIDs) that are large enough to create a name space practically inexhaustible. The location

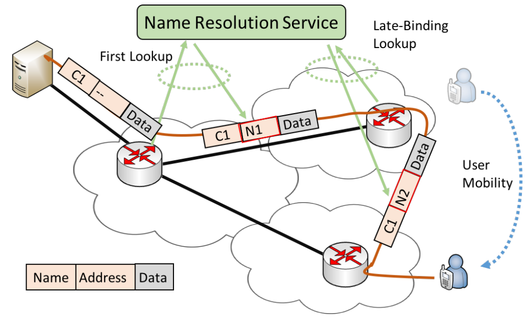

Figure 2: Mobility support through name resolution and late binding

of these objects is then resolved through a Name Resolution Service (NRS) through a dedicated API. While this idea is not completely new and relies on previous work on name/address separation [12, 8, 18, 19], we take it further down the line, proposing support for the so called "name-objects": names that can be used to represent many different Internet objects; for example, a cell-phone, a person, or a group of devices, as shown in Figure 1.

Once we move away from the host-centric nature of the IP world, new and different delivery services can be supported, where routing decisions can be diversified based on the nature of the referenced object. In order to do so, network elements have to support hybrid routing schemes, where a service ID (SID in the Figure) can be located in the networking header to identify the required service. As we will show with the proposed solutions, we believe that this concept, together with appropriate support in the different network entities, can be at the center of the evolution of sustainable communication techniques for the large population of smart-object devices.

3.1 Handling Mobility

Assigning permanent names to mobile objects and decoupling names from addresses has inherent benefits over IP avoiding the need of relying on triangular routing to solve the hierarchical nature of Mobile-IP routing schemes and support host mobility [1]. As end-points talk to each other through names, routers can map these names to the most current locators of the devices, routing packets directly to them and avoiding the need to first route through a previous location persistent element (the home agent). End-points are solely responsible for updating the resolution system with their current locator. Each packet is then routed to the most up-do-date location, as any other device can query the same system to obtain the current locator. When occasional inflight packets reach the previous location of a moving device, routing components can notice that this has moved and perform a new query to obtain the new location. We call this technique "late-binding" and an example is shown in Figure 2.

Avoiding the use of current rendezvous-point based techniques brings two key advantages: first, traffic performance is increased as edge networks are not overloaded by transient traffic and the distance travelled by a single packed is reduced; second, the system can expect higher resilience against failure and fault tolerance compared to the employment of rendezvous servers that can become a single point

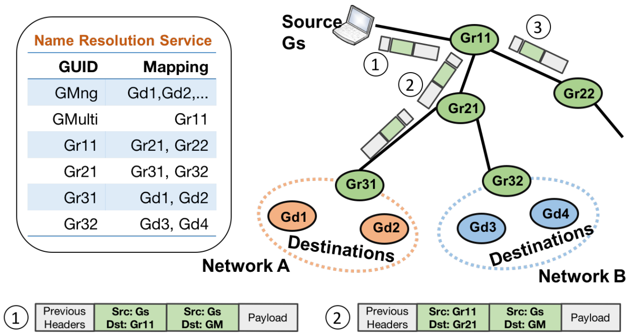

Figure 3: Multicast tree management

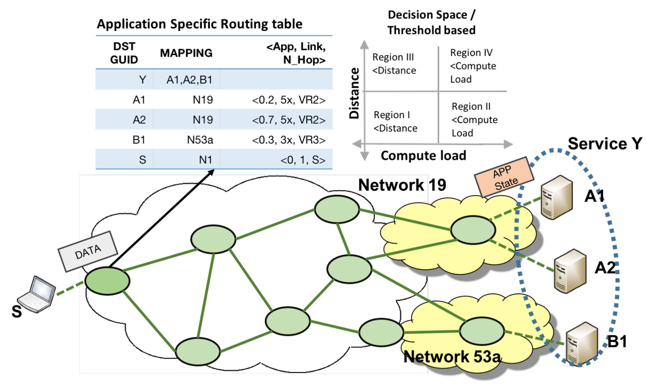

Figure 4: Application specific routing

of failure.

3.2 Group Based Delivery

Smart-objects can often be grouped into sets of devices that require to efficiently communicate to some or all of them at the same time. Consider for example IoT based messaging scenarios: a typical query involves sending short messages to hundreds or thousands of users or application agents, so that scalability becomes an issue, as multiple unicast messages through an overlay service can easily overload the network. Mobility of end-devices results in additional complexity, especially for dynamic environments such as vehicular communications. For example, if a single warning message needs to be pushed to hundreds of cars and pedestrians in a given area, multicast groups would need to be maintained across a large number of access networks in order to efficiently support such one-to-many communication. Using appropriate multicast routing solutions would help improving network efficiency, while reducing the complexity and cost of deploying such applications. However, existing network-layer multicast solutions (e.g., PIM-SM [9], MOSPF [13]) have not been widely adopted due to fundamental problems that are a by-product of the original Internet design geared toward static host-centric communication.

Using named-objects, we propose a solution that exploits the presence of a globally accessible name resolution service, to simplify and streamline the creation and management of multicast groups. Names are used as follows: first, each multicast group is identified by a unique name across all domains, thus separating routing logic from group management. Then, using name recursion, we store the tree topology into the name resolution service. This is achieved by mapping unique names assigned to participating routers to their children nodes, as shown in Fig. 3. Data forwarding can then be performed using tunnels between participating nodes; end-to-end information is preserved within the packet, while the information globally available in the name resolution service is used to identify next hops in the distribution path allowing for quick branching and replicating decisions. Finally, dynamicity of mobile environments is handled by decoupling the participants name from their location through the resolution service and periodically recomputing the multicast tree; the system first needs to translate the name into a list of host names participating in the multicast group. The routable address (locator) of each host (whether mobile or static) can then be identified by a subsequent query to the name resolution service.

3.3 Edge Cloud Support

Edge clouds and fog computing are becoming more and more a relevant topic as the necessity to support mobile devices becomes a stringent requirement for networks and services. Unique technical challenges associated to highly distributed scenarios, such as IoT, emerge. These include the necessity of scaling the Internet architecture to support a very large number of objects while still achieving fast response and stringent QoS requirements. While edge clouds look like a good solution thanks to their highly distributed nature capable of "following" the devices bringing its smartness with them, they usually lack deep network support to understand fine-grained dynamics affecting performance to delivery processing requests to the best replicas (which might not be simply the nearest one in networking terms). Offloading these decisions to the network through an "anycast" like service capable of exploiting both network information and application state would be highly desirable.

We use name-based communications to implement this abstraction. Two core pieces of technology are introduced: first, the ability to aggregate multiple service instances under a single name. This is done by offloading the list of participant locations under a single name into the name resolution service. Second, the ability to make cloud nodes participate in the routing protocol by sharing their application state. This could be either implemented through a new interface in the participating routers, requiring though the introduction of new schemes to identify participants of the protocol, or by offloading this information to the name resolutions service. An example of how this application state information could be used in a routing algorithm is shown in Figure 4, where thresholds based on service load and distance are merged into a single decision process.

4. THE MOBILITYFIRST FUTURE INTERNET ARCHITECTURE

With the above stated network service abstractions in consideration, we have designed [19] and prototyped [6] a network architecture that addresses the principal goals of supporting at-scale and seamless *mobility*, along with *trustworthiness* in the future Internet. While for obvious reasons of space, we will not go deep into details, we will highlight how such architecture, called MobilityFirst, exploits such abstractions to create efficient solutions to the presented

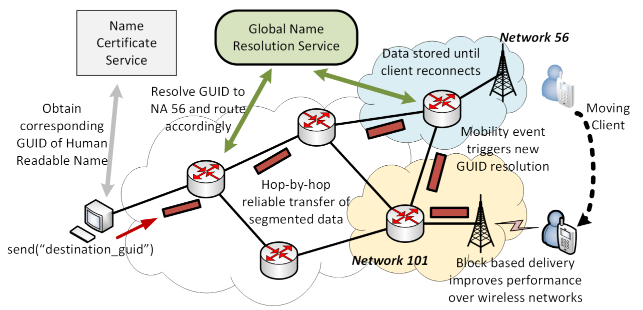

Figure 5: MobilityFirst architecture design.

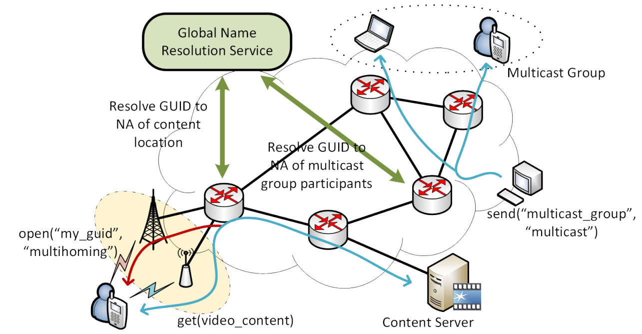

Figure 6: Service abstractions provided via the client API.

problems. Figures 5 and 6 show the main components of the architecture which centers around the concept of self-certifying GUIDs as names for all network principals.

Naming and dynamic resolution. At the crux of the MobilityFirst architecture is a new name-based service layer which serves as the "narrow waist" of the protocol stack. The name-based service layer uses flat GUIDs for all principals or network-attached objects including hosts, content, and services, making each a first-class network object. This resolution is enabled by a globally accessible name resolution service (GNRS), which is used by objects to both announce their latest location/address and lookup end points they wish to communicate with. While a variety of incarnations of the GNRS are possible, we have validated 2 alternate designs that both meet our low resolution latency goals of less than 100ms on average for lookup operations [23, 20]. Additional work is ongoing to improve performance for more extreme scenarios which might require latency upper bounds of 10ms or less.

Name based network interfaces. As all network-attached objects in the MobilityFirst architecture enjoy direct addressability through long lasting unique network names or identifiers, a new GUID-centric network service API [5] is designed to offer network primitives for basic messaging (*send, recv*) and content operations (*get and post*) while supporting several delivery modes innately supported by the MF network such as multihoming, multicast, anycast and DTN delivery. Combined with the GUID indirection and grouping (GUID mapped to one or more other GUIDs) concepts supported by the naming services, the new communication API can produce novel addressing and delivery capabilities.

Storage-informed segmented transport, edge-aware routing. In contrast to end-to-end transports which perform poorly in wireless conditions, MobilityFirst employs a segmented transport to reliably progress data hop-by-hop. Data is segmented into large blocks that are cached at each hop, if storage is available, to enable in-network retransmission under losses [11]. Within a domain, a generalized storage-aware routing (GSTAR) combines link-state routing with DTN elements, and flexibly expands connectivity across wired and wireless segments, as also occasionally connected partitions [16]. Conditions at the wireless-edge are taken into account by adopting an edge-aware inter-domain routing (EIR) approach that scalably gathers (using *telescoping* or aggregation of updates) and utilizes capacity and load conditions at edge networks to instrument effective multi-path and multi-home delivery.

Advanced delivery methods. Building on top of the available name resolution service, MobilityFirst integrates native multicast routing [14]. using two forms of name indirection. A first unique name is assigned to perform the task of node membership; all entities interested in receiving data from the multicast flow, can request to join by inserting their own unique name into the corresponding mapping in the table. This information is then exploited close to the source by a multicast service manager, which builds an efficient tree based on the available resources and the size of the required tree. On the flight multicast is also permitted, using look-ahead longest-common path (LA-LCP) algorithm without requiring building the entire tree. Moreover, context based communications are supported by exploiting in-network computing to keep track of participating objects based on contextual information [7]

Edge cloud support. Using novel network virtualization techniques that exploit the "named-object" abstraction provided, MobilityFirst supports cloud service addressability and advanced anycast delivery capabilities [15]. Through name based communications, MobilityFirst provides native support of virtual networks as the named-object abstraction makes it possible to define virtual networks and store the corresponding topology directly in the GNRS, including the list of participating routers. Such list contains unique names that only belong to the VN logic. This indirection allows for a clean separation between logical and physical layer allowing for easier management and maintenance procedures, e.g. migration. A concept called Application Specific Routing (ASR) is then used to provide fine grained access to distributed edge clouds. ASR allows routing decisions to be based both on network and application metrics. As services attach to the virtual network, they can expose custom information to be employed by the VN routing layer, that will then not only consider classic L3 metrics, but also application layer ones, such as cloud workload/latency.

5. CONCLUSIONS

This paper presents our vision for a set of abstractions and solutions that can serve as a foundation to support communications and advanced services for smart-objects in the future Internet. Using the concept of named-object based networking, we proposed solutions to solve three base communication paradigms that are of fundamental importance: mobility, group based delivery and support for distributed computing resources.

These solutions are among the key features of the MobilityFirst future Internet architecture. By exploiting the work that has been put over the last few years to build a fully working prototype of the architecture, we aim to provide a thorough analysis of the conveniency of such abstractions showing what levels of improvements in efficiency and flexibility tehy can provide.

6. REFERENCES

[1] Ip mobility support for ipv4. http://tools.ietf.org/html/rfc3344.

[2] Mobility support in ipv6. http://www.ietf.org/rfc/rfc3775.txt.

[3] Cisco visual networking index: Global mobile data traffic forecast update, 2015-âĂŞ2020. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html, 2015.

[4] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli. Fog computing and its role in the internet of things. In *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*, pages 13–16. ACM, 2012.

[5] F. Bronzino, K. Nagaraja, I. Seskar, and D. Raychaudhuri. Network service abstractions for a mobility-centric future internet architecture. In *Proceedings of the eighth ACM international workshop on Mobility in the evolving internet architecture*, pages 5–10. ACM, 2013.

[6] F. Bronzino, D. Raychaudhuri, and I. Seskar. Experiences with testbed evaluation of the mobilityfirst future internet architecture. In *2015 European Conference on Networks and Communications (EuCNC)*.

[7] F. Bronzino, D. Raychaudhuri, and I. Seskar. Demonstrating context-aware services in the mobilityfirst future internet architecture. In *Proceedings of the First International Conference in Networking Science and Practice (ITC) 28*. IEEE, 2016.

[8] D. Farinacci, D. Lewis, D. Meyer, and V. Fuller. The locator/id separation protocol (lisp). 2013.

[9] D. Farinacci, C. Liu, S. Deering, D. Estrin, M. Handley, V. Jacobson, L. Wei, P. Sharma, D. Thaler, and A. Helmy. Protocol independent multicast-sparse mode (pim-sm): Protocol specification. 1998.

[10] R. P. Jover and I. Murynets. Connection-less communication of iot devices over lte mobile networks. In *Sensing, Communication, and Networking (SECON), 2015 12th Annual IEEE International Conference on*, pages 247–255. IEEE, 2015.

[11] M. Li, D. Agrawal, D. Ganesan, and A. Venkataramani. Block-switched networks: a new paradigm for wireless transport. In *Proc. of NSDI*, 2009.

[12] R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson. Host identity protocol. *RFC 5201, April*, 2008.

[13] J. Moy. Rfc 1584–multicast extensions to ospf. *SRI Network Information Center*, 1994.

[14] S. Mukherjee, F. Bronzino, S. Srinivasan, J. Chen, and D. Raychaudhuri. Achieving scalable push multicast services using global name resolution, 2016.

[15] K. Nakauchi, F. Bronzino, Y. Shoji, I. Seskar, and D. Raychaudhuri. vmcn: Virtual mobile cloud network for realizing scalable, real-time cyber physical systems. In *Proceedings of the Workshop on Distributed Cloud Computing (DCC)*. ACM, 2016.

[16] S. C. Nelson, G. Bhanage, and D. Raychaudhuri. GSTAR: generalized storage-aware routing for mobilityfirst in the future mobile internet. In *Proc. of MobiArch*, pages 19–24. ACM, 2011.

[17] A. Osseiran, F. Boccardi, V. Braun, K. Kusume, P. Marsch, M. Maternia, O. Queseth, M. Schellmann, H. Schotten, H. Taoka, et al. Scenarios for 5g mobile and wireless communications: the vision of the metis project. *IEEE Communications Magazine*, 52(5):26–35, 2014.

[18] J. Pan, R. Jain, S. Paul, and C. So-In. Milsa: A new evolutionary architecture for scalability, mobility, and multihoming in the future internet. *IEEE journal on selected areas in communications*, 28(8):1344–1362, 2010.

[19] D. Raychaudhuri, K. Nagaraja, and A. Venkataramani. Mobilityfirst: a robust and trustworthy mobility-centric architecture for the future internet. *ACM SIGMOBILE Mobile Computing and Communications Review*, 2012.

[20] A. Sharma, X. Tie, H. Uppal, A. Venkataramani, D. Westbrook, and A. Yadav. A global name service for a highly mobile internetwork. In *ACM SIGCOMM Computer Communication Review*, volume 44, pages 247–258. ACM, 2014.

[21] V. Valancius, N. Feamster, J. Rexford, and A. Nakao. Wide-area route control for distributed services. In *USENIX Annual Technical Conference*, 2010.

[22] A. G. Valkó. Cellular ip: a new approach to internet host mobility. *ACM SIGCOMM Computer Communication Review*, 29(1):50–65, 1999.

[23] T. Vu, A. Baid, Y. Zhang, T. D. Nguyen, J. Fukuyama, R. P. Martin, and D. Raychaudhuri. Dmap: A shared hosting scheme for dynamic identifier to locator mappings in the global internet. In *Distributed Computing Systems (ICDCS), 2012 IEEE 32nd International Conference on*, pages 698–707. IEEE, 2012.

[24] X. Wu and J. Griffioen. Supporting application-based route selection. In *Computer Communication and Networks (ICCCN), 2014 23rd International Conference on*, pages 1–8. IEEE, 2014.