

Wireless Access Considerations for the MobilityFirst Future Internet Architecture^{*}

Akash Baid, Dipankar Raychaudhuri
WINLAB, Rutgers University, {baid, ray}@winlab.rutgers.edu

Abstract—This paper presents an overview of wireless access considerations behind the design of the MobilityFirst clean-slate future Internet architecture. The MobilityFirst architecture is motivated by a historic shift of the Internet from the fixed host-server model to one in which access from mobile platforms becomes the norm. This implies the need for a future Internet architecture designed to handle the special needs of mobile/wireless access efficiently and at large scale. A number of key wireless access network requirements including user/network mobility, varying wireless link quality and disconnection, multi-homing, ad hoc networking, flexible autonomous system boundaries and spectrum coordination are identified along with a brief discussion of the implications for protocol design. This is followed by a summary of the MobilityFirst protocol stack based on separation of names and locators, global name resolution service (GNRS), storage-aware routing with hop-by-hop transport, integrated spectrum management, along with an enhanced edge-aware interdomain routing framework. Selected results from ongoing protocol design and evaluation work are given for key components such as the GNRS, storage-aware routing and spectrum coordination protocol.

I. INTRODUCTION

The MobilityFirst architecture project is founded on the premise that the Internet is approaching a historic inflection point, with mobile platforms and applications poised to replace the fixed-host/server model that has dominated the Internet since its inception [1]. This predictable, yet fundamental, shift presents a unique opportunity to design a next generation Internet in which mobile devices, and applications, and the consequent changes in service, trustworthiness, and management are primary drivers of a new architecture. The major design goals of our proposed architecture are: mobility as the norm with dynamic host and network mobility at scale; robustness with respect to intrinsic properties of wireless medium; trustworthiness in the form of enhanced security and privacy for both mobile networks and wired infrastructure; usability features such as support for context-aware pervasive mobile services, evolvable network services, manageability and economic viability. The design is also informed by technology factors such as radio spectrum scarcity, wired bandwidth abundance, continuing Moore’s law improvements to computing, and energy constraints in mobile and sensor devices.

This paper presents the wireless/mobile edge network perspective behind the MobilityFirst (MF) design. In particular, we discuss emerging mobile/wireless access network usage

scenarios and identify some of the resulting key protocol features needed to address these requirements. A broad range of emerging wireless access network scenarios have been identified in [2]. In addition to ongoing cellular/Internet integration, important new use cases include heterogeneous network access, mobile content delivery, wireless peer-to-peer (P2P) networking, vehicular (V2V) networking, sensor and machine-to-machine (M2M) applications, and so on. Each of these usage scenarios maps to a certain set of key protocol requirements in terms of naming, routing, transport, security and privacy. These protocol design requirements can then be extracted to serve as inputs to the design of next-generation Internet protocols now being considered by the research community [3].

The MF architecture [4] has been designed with this approach, with a strong emphasis on meeting emerging wireless access requirements while at the same time improving robustness, flexibility and security in the core wired network as well. Although this is not yet well understood, there is an important interplay between edge network requirements and core network design, which have historically been treated as separate problems. In the rest of this paper, we start by identifying some of the important wireless access requirements and provide a discussion of associated protocol design considerations. This is followed by an outline of the MF architecture, showing how the proposed new design attempts to meet many if not all the identified requirements. Selected results from ongoing protocol and design work are also provided for key components where available.

II. WIRELESS ACCESS CHALLENGES & REQUIREMENTS

A. Mobility of devices, networks, content and context

Mobility in networks has been extensively studied for over 20 years in conjunction with protocols and standards such as mobile IP and 3GPP [5], [6]. In the MF architecture, we generalize the concept of mobility support to provide seamless connectivity to more general kinds of network-attached object, which can be devices, hosts, users, content files or even context-defined data. Ideally, seamless mobility can be supported by providing a unique name for any network-attached objects and then providing a mechanism for dynamic binding of the name to the current points of attachment.

These functional requirements can be translated to the following protocol design requirements (as outlined in Fig. 1):

- 1) Support for named network-attached objects.

^{*}Research supported by NSF Future Internet Architecture (FIA) grant CNS-1040735




Fig. 1. A server in Network C is sending packets to a host moving from Network A to B. Seamless delivery of packets can be achieved if packets are destined for the device rather than the current network address of the device.

- 2) Dynamic binding of names to network addresses/locators.
- 3) Redirection of in-transit packets through the routing layer, with delay tolerance to deal with disconnection.

B. Varying wireless access link quality and disconnection

Fluctuations in access link quality is a fundamental property of the wireless medium - see for example the sample trace of downlink throughput in an experimental 4G network showing bit-rate variations greater than 3:1 during just tens of seconds (Fig. 2). In addition, complete disconnection due to mobility and/or insufficient signal strength is also a common occurrence in both 3G/4G and Wi-Fi networks. While these variations are usually handled at the PHY and MAC layers, they invalidate some implicit assumptions in the control loop algorithms used in the Internet. For example, TCP congestion control treats wireless link errors as congestion losses and performs poorly in high variation wireless channels [7].

Given the increasing dominance of wireless medium as the last hop for Internet access, such link quality variations need to be natively supported at different layers of the Internet architecture. This leads to the following requirements:

- 1) Increased visibility of edge link quality to distinguish between wireless errors from other network anomalies.
- 2) Link quality awareness at both the intradomain and inter-domain routing layer to allow for robust packet delivery strategies.
- 3) Disconnection tolerant routing and transport protocols which are capable of temporarily storing packets during disconnections and rerouting to the new point of attachment.




Fig. 2. Variations in downlink throughput measured for a client connected to a GENI WiMax base station

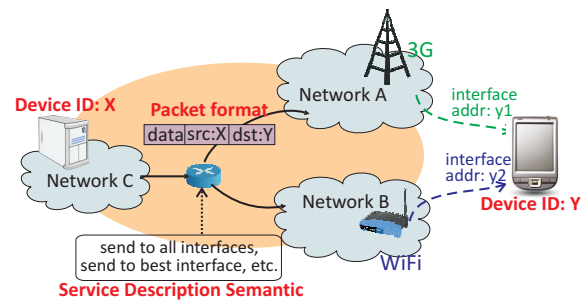


Fig. 3. Qualitative illustration of multi-homing support in the Internet

C. Accessing multiple networks

Most mobile devices today support multiple network interfaces such as 3G/4G cellular and Wi-Fi. A typical wireless device in an urban area might see 3-5 cellular networks and 8-10 Wi-Fi access points, but is currently constrained to access only one of these due to both technical and business model constraints. Currently, a single device can only have one IP address at a time, restricting multi-homing to scenarios where both interfaces belong to the same network operator (autonomous system). In the future Internet, it would be desirable to have improved support for multi-homing in order to allow a device to simultaneously access $k > 1$ networks in a seamless manner.

Efficient support for host multi-homing induces the following key requirements (see Fig. 3):

- 1) Separation of names for network-attached objects and their locators (network addresses (NAs)), allowing a single name to have multiple NAs.
- 2) Routing protocols (both intra and inter domain) with some visibility of link quality along the available alternate paths.
- 3) Service semantics to support policies for choosing between multiple paths (e.g. send to “all” interfaces or “best” interface)

D. Ad hoc network support

Wireless ad hoc and peer-to-peer networks are important for infrastructure less scenarios such as vehicle-to-vehicle (V2V) communications. Applications such as photo/video sharing, last-mile connectivity, local social networking, multi-player gaming and dissemination of traffic information inherently fit the P2P communication model. However support for such wireless P2P scenarios is currently limited to re-use of overlay applications or relegated to lower layer solutions such as Wi-Fi Direct [8]. Integration of such networks within the framework of a future Internet design presents a unique set of challenges. These include:

- 1) Fast service discovery and ad hoc network formation.
- 2) Name resolution, authentication, routing and security services capable of disconnected operation.
- 3) Inter-domain routing support for large-scale migration of networks.

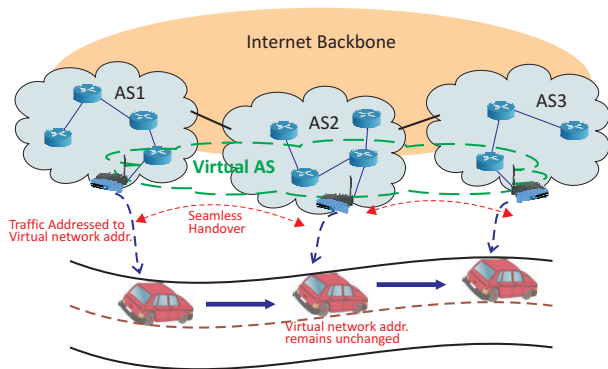


Fig. 4. A mobile client connects to a series of Wi-Fi APs residing in different ASes. A virtual AS encompassing the involved routers can obviate the need for layer 2 handovers.

E. Flexible AS formation in edge networks

Unlike cellular networks where a single entity owns a set of base-stations deployed over a large area, mobile ad hoc networks discussed in Sec. II-D motivate the need for flexible autonomous system (AS) formation in the future Internet. A second motivation for flexible ASs is the emergence of unlicensed band radios which can be used to form “unlicensed networks” which provide an alternative to licensed cellular networks. The current definition of an AS is relatively static and corresponds to an edge or core network fully owned by a single enterprise or operator. Future networks should be able to form ASs in a more dynamic fashion and be able to incorporate components such as unlicensed access points (APs) or ad hoc networks which are not necessarily physically contiguous. In order to transition from unlicensed radios to unlicensed networks, dynamic formations of *Virtual ASs* is required.

A vehicular use-case is illustrated in Fig. 4, where a moving host connects to a series of Wi-Fi APs each residing in different ASs. The set of routers servicing these APs, having ascertained their proximity on a commonly traversed path can form a virtual AS used for handling highly mobile clients. Packets destined to the client, in such a setting, can be addressed to a single virtual network address which the client retains throughout the journey.

F. Network-assisted spectrum coordination

End hosts in the future Internet will predominantly use wireless links as the last-hop connection medium [1], with an increasing proportion expected to operate in unlicensed spectrum or recently allocated TV white space bands. Coordination of the unlicensed spectrum usage is thus a key requirement for ensuring high rate, low delay Internet services in the future. Efficient management of the limited spectrum resource will be especially required in dense usage scenarios. For example, Fig. 5 shows a typical case of thousands of Wi-Fi APs being deployed in a small geographical area. With such a large number of transmitters contending for channel time, the CSMA overhead substantially reduces the throughput that any device can achieve and thus severely degrades the QoS



Fig. 5. Estimated Wi-Fi AP locations in a 0.4x0.5 sq.mile area in Manhattan, NY. Red dots indicate single citation in the crowd sourced war-driving database. Green dots indicate multiple citations. Source:WiGLE [9]

for all Internet services. Coordination of unlicensed spectrum access requires a network management layer capability that can assimilate the local interference measurements at different end-points and disseminate aggregated information such that individual devices can select optimal access parameters. Further a framework for geographical region based message forwarding is required to enable such spectrum information dissemination.

G. Other design requirements

Although we do not focus on the security aspects in this paper, the requirements of location privacy, strong authentication of ownership, mechanisms against mobility spoofing attacks and fast authentication mechanisms must be taken into account for a mobile centric future Internet architecture. In addition, mobile M2M applications such as environment sensing, traffic monitoring, smart grids and inventory tracking [10] require new kinds of services involving geographic scope, content and context - see [11] for further details.

III. MOBILITYFIRST FUTURE INTERNET ARCHITECTURE

The MobilityFirst architecture [4] is based on the idea of separating “names” of end-users or other network-connected objects, and their routable addresses or locators. Separation of names and addresses makes it possible for mobile devices to have a permanent, location independent name or globally unique identifier (GUID) which can then be mapped to a set of routable network addresses (NAs) corresponding to the current point(s) of attachment. This concept of separating names from addresses has been proposed in earlier work, such as [12] but is usually viewed as an overlay service above the network similar in spirit to DNS. The MF architecture aims to integrate a global name resolution service (GNRS) as a basic network-layer service which can be efficiently accessed both by end-user devices and in-network routers, base stations and access points [13].

This concept is illustrated in Fig. 6 which shows the layering of functionality in the proposed MobilityFirst architecture. The design consists of a set of application specific ‘name assignment’ services which translate human readable names such as ‘sensor@xyz’ or ‘John’s laptop’ to GUIDs. This framework also supports the concept of context-based descriptors such as ‘taxi in New Brunswick’ which can be resolved by a

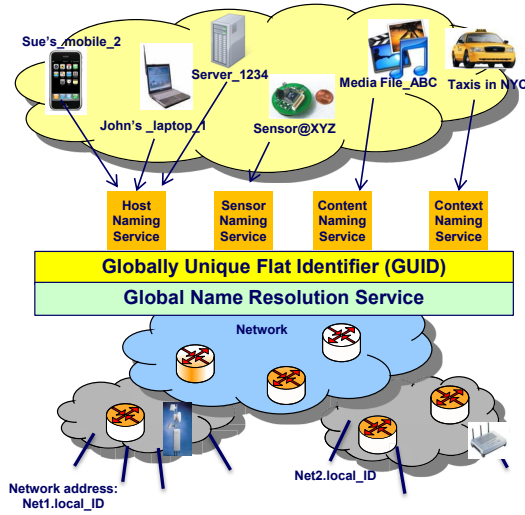


Fig. 6. Naming and name resolution in MobilityFirst

context naming service to a particular GUID which serves as a dynamic multicast group for all taxis currently in that area. The GUID is then assigned to the mobile device (or other network-connected object) and entered into the network-level GNRS service shown in the figure. The GNRS is a distributed network service which is responsible for maintaining the current bindings between the GUID and network address(es). Mobile devices (or routers at their point of attachment) update the GNRS with their current network address(es) resulting in a table entry such as $\langle \text{GUID: Net1, Net2, optional properties} \rangle$. In the following subsections, we discuss some features of the MobilityFirst architecture that target the requirements outlined in Section II.

A. Dynamic Name-Address Bindings

The GUID-based protocol stack described above handles host and network mobility through fast dynamic binding of identifiers to locators. That is, when a user sends packets directed to a particular identifier (GUID), the networking protocol must quickly ascertain the set of locators (NAs) attached to the GUID and route the packets correspondingly. We address the challenge of providing a fast global name resolution service at Internet scale through a router DHT-based *Direct Mapping (DMap)* scheme for achieving a good balance between scalability, low update/query latency, consistency, availability and incremental deployment [13]. In order to perform the name resolution for a given GUID, DMap distributes the GUID \rightarrow NA mappings amongst Internet routers using an in-network single-hop hashing technique which derives the address of the mapping router directly from the GUID. Through a detailed simulation study described in [13], we have shown that DMap achieves a 95th percentile round trip query response time of below 100ms (Fig. 7 presents the key query response time result), considered more than adequate for current and future mobility services [14], [15]. The dynamic binding of GUIDs to network addresses thus helps meet mobility and multi-homing requirements.




Fig. 7. CDF of round trip query time from a measurement driven Internet scale simulation of 1 Million name resolution queries passing through a realistic Internet model. K represents the number of replicas of each mapping and provides a tradeoff between response time and storage load.

B. Storage-aware and Delay Tolerant Routing

MobilityFirst uses a generalized storage-aware routing (GSTAR) approach in order to support delay and disconnection tolerance in the routing layer. The GSTAR approach [16], [17] is based on the principle of in-network storage integrated into the routing later in order to provide improved performance in presence varying link quality and disconnection. The basic idea in GSTAR is for each router to have some storage capability and make forward vs. store decisions based on both short-term and long-term path quality metrics. In addition, packets along paths which become disconnected are handled by a delay-tolerant network (DTN) mode of the protocol with delayed delivery and replication features.

In particular, each router maintains two types of topology information: (i) The intra-partition graph is formed by collecting flooded linkstate advertisements (F-LSAs) which carry fine-grained, time-sensitive information about the links in the network; (ii) The second, termed the DTN graph, is maintained via epidemically disseminated link-state advertisements (DL-SAs) which carry connection probabilities between all nodes in the network. Recent results in [17] indicate that by intelligently utilizing in-network storage, GSTAR outperforms traditional and storage-augmented link-state protocols in both wired and wireless network environments. Fig. 8 shows an example of the gains for a ORBIT testbed implementation of GSTAR with two wireless links. This directly targets the requirements arising from link quality variations, network disconnections and ad hoc networks.

In addition to GSTAR intradomain routing, the MF architecture includes an edge-aware interdomain routing protocol which allows networks to exchange information on multiple path options and path quality so that the sending network can make reasonable routing decisions taking into account wireless access network properties. This interdomain framework also supports efficient multi-homing, multicast and multipath which are all useful features for wireless services.

C. Spectrum Access Coordination

The MF management plane facilitates unlicensed band spectrum coordination through dissemination of spectrum usage




Fig. 8. GSTAR performance in a fluctuating link: data rate is 6 Mbps for half the time period and 54 Mbps for the other half.

information to networks within radio interference range of each other. In this architecture, the routers which directly connect to the base stations/home APs run an evolved flavor of geocast routing [18] which stores the information about the region of operation of each network that they support. As illustrated in Fig. 9, the source X of any spectrum management message, signs it using $\{L_x, r_x\}$ where L_x is the geo-location of X and r_x is the radius of operation obtained by equating: $PL_x(r) = P_{x,max} + G_x - S_{x,min} - N$, where $PL_x(\cdot)$ is the appropriate indoor/outdoor pathloss model used, $P_{x,max}$ is the maximum transmit power of X , $S_{x,min}$ is the minimum received power required for operation and N is the noise floor. Each router stores the list of $\{L_i, r_i\}$ pair for each of the network that it supports either directly or through a child router. Upon receiving this message, the router checks to see if the source region in the message overlaps with any of its networks and passes the message to all overlapping networks. It further routes the message to its parent router (using IP tunneling if there are other routers on the way that do not support this feature) which then sends it to other routers connected to it using a similar overlap search.

The resulting distributed spectrum management service allows for co-operation between independent access networks. For example, two virtual access networks using physically overlapping sets of WiFi AP's can coordinate their spectrum usage in order to achieve improved performance. An example evaluation for overlapping WiFi grids (as in stadiums or dense urban areas) shows 150-200% throughput improvements for clients most affected by the interference [19].

IV. CONCLUDING REMARKS

In this paper we have presented a brief overview of future Internet design considerations driven by emerging wireless access and mobility scenarios. Several key protocol requirements have been identified including name/address separation, robustness with respect to link quality variation and disconnection, multi-homing, ad hoc network formation, flexible interdomain boundaries and spectrum coordination. Key design features of the MobilityFirst protocol stack have been outlined and shown to address some of these requirements. Comprehensive coverage of all the design goals and protocol features is beyond the scope of this paper, but we have

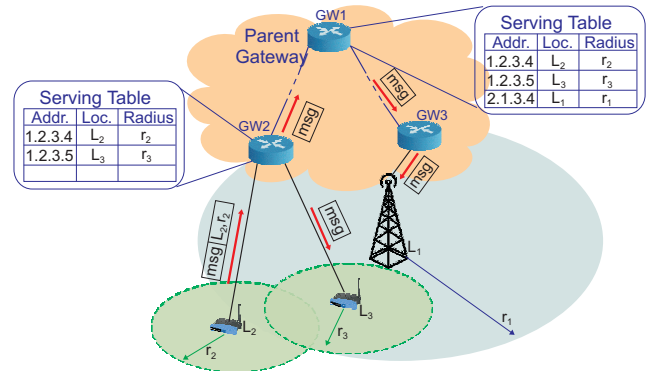


Fig. 9. Network Support for Neighbor Selection

attempted to provide a general understanding of the design approach and key features of the MF stack. Protocol design and evaluation work is ongoing, and a more complete coverage of this topic along with large-scale prototype validations will be reported in future papers on this topic.

REFERENCES

- [1] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2011-2016," Cisco White Paper, Feb. 2012.
- [2] D. Raychaudhuri and M. Gerla, *Emerging Wireless Technologies and the Future Mobile Internet*. Cambridge University Press, 2011.
- [3] NSF Future Internet Architecture Project, <http://www.nets-fia.net/>.
- [4] MobilityFirst Future Internet Architecture Project, <http://mobilityfirst.winlab.rutgers.edu/>.
- [5] C. Perkins, "Mobile IP," *Communications Magazine, IEEE*, vol. 35, no. 5, pp. 84–99, May 1997.
- [6] 3GPP: Mobile Broadband Standard, <http://www.3gpp.org/>.
- [7] G. Xylomenos, G. Polyzos, P. Mahonen, and M. Saaranen, "TCP performance issues over wireless links," *Communications Magazine, IEEE*, vol. 39, no. 4, pp. 52–58, April 2001.
- [8] Wi-Fi Direct: Personal, portable Wi-Fi, <http://www.wi-fi.org/discover-and-learn/wi-fi-direct>.
- [9] WiGLE: Wireless geographic logging engine, <http://www.wigle.net/>.
- [10] G. Wu, S. Talwar, K. Johansson, N. Himayat, and K. Johnson, "M2M: From Mobile to Embedded Internet," *IEEE Communications Magazine*, vol. 49, no. April, pp. 36–43, 2011.
- [11] J. Li, Y. Zhang, K. Nagaraja, and D. Raychaudhuri, "Supporting efficient machine-to-machine communications in the future mobile internet," in *IEEE WCNC 2012, Workshop on IoT-ET*.
- [12] D. Farinacci, D. Fuller, D. Oran, and D. Meyer, *Locator/ID separation protocol (LISP)*, IETF Internet Draft, draft-farinacci-lisp-12.txt, Sep. 2009.
- [13] T. Vu et al., "DMap: A Shared Hosting Scheme for Dynamic Identifier to Locator Mappings in the Global Internet," in *Proceedings of ICDCS, 2012*.
- [14] V. Vassiliou and Z. Zinonos, "An Analysis of the Handover Latency Components in Mobile IPv6," *Journal of Internet Engineering*, vol. 3, no. 1, pp. 230–240, Dec 2009.
- [15] A. Mishra, M. Shin, and W. Arbaugh, "An empirical analysis of the IEEE 802.11 MAC layer handoff process," *SIGCOMM Comput. Commun. Rev.*, vol. 33, pp. 93–102, April 2003.
- [16] S. C. Nelson, G. Bhanage, and D. Raychaudhuri, "GSTAR: generalized storage-aware routing for mobilityfirst in the future mobile internet," in *Proceedings of MobiArch, 2011*, pp. 19–24.
- [17] N. Somani, A. Chanda, S. C. Nelson, and D. Raychaudhuri, "Storage-Aware Routing for Robust and Efficient Services in the Future Mobile Internet," in *Proceedings of ICC FutureNet V, 2012*.
- [18] M. Mauve, A. Widmer, and H. Hartenstein, "A survey on position-based routing in mobile ad hoc networks," *Network, IEEE*, vol. 15, no. 6, pp. 30–39, August 2002.
- [19] A. Baid, M. Schapira, I. Seskar, J. Rexford, and D. Raychaudhuri, "Network Cooperation for Client-AP Association Optimization," in *WiOpt 2012, RAWNET workshop*.