

Supporting Mobility for Internet Cars

Zhenkai Zhu and Lixia Zhang, University of California, Los Angeles

Ryuji Wakikawa, Toyota Infotechnology Center

ABSTRACT

In the last few years we have witnessed an increasing number of cars being connected to the Internet. All indicators suggest that this trend will continue, and vehicles will soon become first class citizens on the Internet. Vehicle networking opens the door to a vast new class of applications ranging from car monitoring and diagnosis to passenger assistance, communication, and entertainment. In this article we present a comprehensive survey of the existing mobility support solutions, and discuss various design trade-offs and remaining issues in view of vehicle networking. We hope that these discussions can provide useful inputs to car manufacturers in their design considerations for bringing all cars online.

INTRODUCTION

Thanks to rapid advances in wireless technology, the Internet is becoming increasingly mobile. Not only do smart phones become more affordable and ubiquitous; car manufacturers are also looking into leveraging Internet connectivity to cars to provide advanced applications on car maintenance, such as monitoring and diagnosis, on road assistance such as providing route navigation, weather maps, and automated toll payments, as well as on passenger entertainment including various types of Internet-enriched applications. Although most of today's network-connected cars still rely on telematics systems with low-bandwidth connectivity (e.g., satellite link), which do not meet the needs of emerging new applications, this situation is expected to change quickly. At the time of this writing, several car manufacturers such as General Motors are already offering Internet connectivity for a handful models of cars via the third-generation (3G) network [1]; some other manufacturers are also considering offering Internet-enabled car applications or linking smartphone applications to cars [2]. The current trend suggests that tens of millions of cars will go online in coming years, and innovative car-based Internet applications and services will emerge, which can have a major impact on both manufacturers and passenger experiences.

Generally speaking, vehicle communications can be sorted into three major scenarios: vehicle to vehicle (V2V), vehicle to roadside unit (V2R), and vehicle to infrastructure (V2I). The first two

types of communication scenarios are generally based on 802.11p, which is designed for short-range data exchange between vehicles or between vehicle and roadside units. Example applications of V2V and V2R include safety alert systems, toll collection, and probe data sharing. The third scenario, V2I, can utilize a variety of wireless technologies for a car to communicate with the Internet infrastructure and to enable seamless handover from one communication medium to another. With abundant computing and communication resources from the deployed Internet infrastructure, V2I is expected to offer more advanced applications than the first two scenarios. One critical technology in enabling V2I is mobility support, which is our focus in this article.

What would be the best way to provide global scale IP mobility support for Internet-enabled cars? Multiple answers exist today. In this article we first present a comprehensive survey of existing mobility solutions, and then provide an in-depth analysis of the design trade-offs and remaining issues in view of vehicle networking. We discuss the requirements that are specific to automobile environment and conclude the article.

A BRIEF REVIEW OF MOBILITY SOLUTIONS

The Internet community has been working on mobility support research and standardization since the early '90s. To assess the state of the art, in this section we provide a brief overview of various solutions in mobility support. Interested readers can find a complete and detailed description of all the existing mobility designs in [3].

MOBILE IP AND ITS EXTENSIONS

Mobile IP (MIP) is probably the best known mobility support protocol. It has attracted a lot of attentions from both researchers and industrial practitioners, and has been deployed in a number of commercial systems. Several extensions to MIP have also been developed to adapt MIP to different mobile scenarios.

Mobile IP (MIP) — The Internet Engineering Task Force (IETF) developed the first MIP standard in 1996, and subsequently defined different versions of the MIP standard for IPv4 and IPv6 networks, respectively [3]. Although these

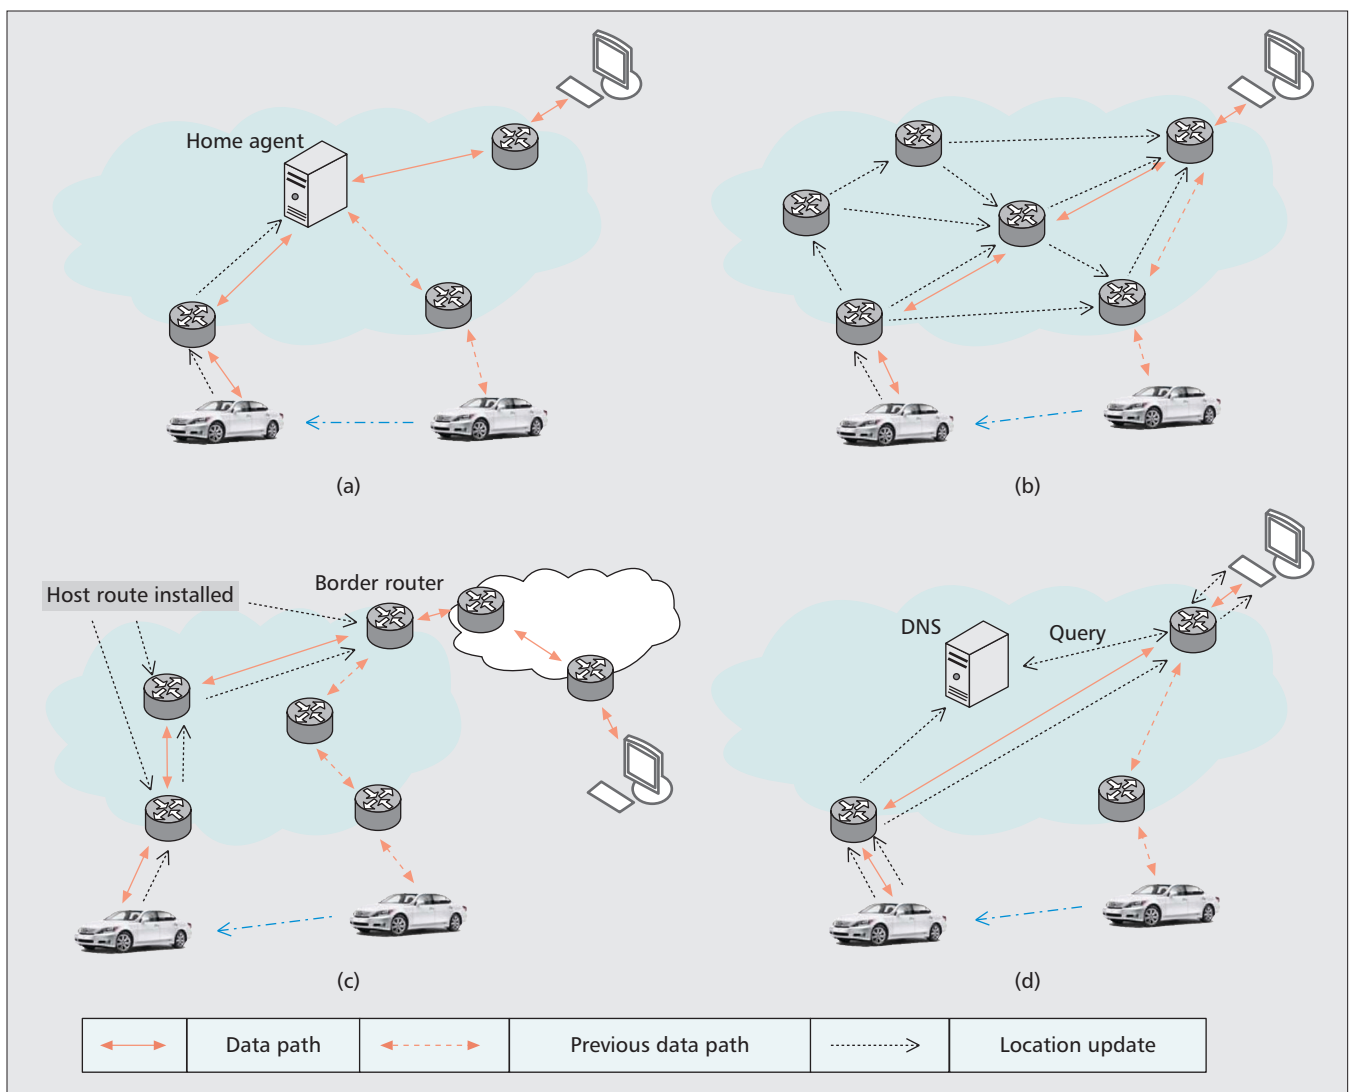


Figure 1. Mobility solutions for vehicles.

standards differ in details, the high-level design is the same. Here we use Mobile IPv6 as an example to illustrate the basic operations.

Figure 1a shows a simple network setting, where the home agent serves as a rendezvous point between a stationary correspondent and a mobile node, in this case a car. MIP assigns each mobile node M a home agent, from which M acquires its home address (HoA) that serves as M 's identifier. As M moves, it also obtains a care-of address (CoA) from its current access router, which can be used to reach M . M notifies its home agent whenever it gets a new CoA. A correspondent node (CN) sends data to the mobile by using the mobile's HoA as the destination IP address. The packets are forwarded to the home agent, which in turn forwards the packets to the mobile by encapsulating the packets using the mobile's CoA as the destination address.

Based on the locations of the mobile, its home agent and the correspondent packets may be forwarded along a triangle path. This problem can be mitigated if the CN is mobility aware and supports route optimization, which works in the following way: whenever the mobile gets a new CoA, it notifies both its home agent and the

CN. This enables the CN to maintain the mapping between the mobile's HoA and CoA, and encapsulate packets to the mobile directly.

NEMO — A group of hosts may move together, as in the case of moving vehicles (e.g., cars, ships, trains, or airplanes), which may have a network with multiple hosts attached. In this case it would be rather inefficient to handle the mobility for each of the hosts separately. Network Mobility (NEMO) [4] was introduced in 2000 as a backward compatible extension to Mobile IP to provide efficient support for network mobility. NEMO assumes that each mobile network has a mobile router which provides IP connectivity to the other hosts on board. Conceptually the mobile router works in a way similar to a mobile node in MIP. However, instead of having a single HoA, it obtains one or more IP prefixes (a.k.a mobile prefixes) from the home agent, and assigns IP addresses to the hosts on board from these address blocks. The mobile router establishes a bidirectional tunnel with its home agent, and the home agent in turn announces the reachability to the mobile prefixes. The mobile router connects to different

PMIP was proposed in 2006 to address the interest of mobile network operators who desire to have direct control over mobility support and to support mobile devices with no mobility support functions.

access routers as the network moves, but does not tell the access routers its mobile prefix(es). Thus, all the packets to and from the mobile network flow through the bidirectional tunnel between the mobile router and the home agent. Mobility is transparent to the hosts on the moving network, but data delivery may go through triangle paths.

Global HAHA — First proposed in 2006 as an extension to MIP, Global HAHA aims to eliminate the triangle routing problem in Mobile IP and NEMO [5]. Global HAHA utilizes multiple home agents that are distributed widely and all announce the same home prefix to the routing system from their different locations, effectively creating a large scale anycast group. Each of all the mobile node is assigned an HoA from the home prefix. A mobile node M can register with any of the home agents closest to it. A home agent H that accepts the binding request from M becomes the primary home agent for M , and notifies all the other home agents of the binding $[M, H]$, so all the home agents keep a synchronized binding information database for all the mobiles. When M moves, it may switch its primary home agent to another one that becomes closest to it.

A correspondent node CN always sends packets to a mobile's HoA. Because of anycast routing, the packets are delivered to the nearest home agent, H_{cn} . H_{cn} then encapsulates the packets to the IP address of the primary home agent H that is currently serving the mobile, and H delivers the packets to the mobile after stripping off the encapsulation header. When many home agents are deployed and distributed widely, H_{cn} can be close enough to CN to effectively eliminate triangle routing in the data delivery path. This makes Global HAHA particularly beneficial for V2I networking, because application servers in the infrastructure do not have to be concerned with mobility support.

Proxy Mobile IP (PMIP) — The three mobility protocols described so far all use home agents to support mobility. These home agents can be provided by end users, network service providers, or even by a third party entity. PMIP [6] was proposed in 2006 to address the interest of mobile network operators who desire to have direct control over mobility support and to support mobile devices with no mobility support functions.

PMIP introduces two new types of nodes, the local mobility anchor (LMA) and mobile access gateway (MAG). Each mobile node M is assigned an LMA within an operator's network. The LMA plays a role similar to that of the home agent in MIP and assigns M a local home network prefix, which serves as M 's identifier in the network. A MAG plays a role similar to that of a foreign agent in MIPv4. However, before providing network access to a mobile M , a MAG first acquires M 's identity and verifies whether M is authorized for access. In addition, a MAG also monitors M 's attaching and detaching events, and generates proxy binding updates to M 's LMA during hand-off. The LMA then updates the mapping between M 's home prefix and the IP address of the MAG

currently serving M . To hide a mobile's roaming from itself, a MAG also advertises each attached mobile's local home network prefix, so the mobile treats MAG as its default router. This enables PMIP to imitate an entire operator's network as a single link to each mobile M , so M does not detect any change with respect to its layer 3 attachment as long as it roams within the network. Other nodes can reach M using its home prefix, and its LMA tunnels packets to M in a way similar to the operation in MIP.

HOST-ROUTE-BASED PROTOCOLS

Another approach to providing network-based mobility support is by setting up a host route for each mobile in a network, as shown in Fig. 1c. A group of existing protocols fall into this category [3]. Here we describe a representative protocol from this group, Cellular IP [7].

Cellular IP is designed as a local mobility solution to work in conjunction with Mobile IP. When entering a Cellular-IP-based network N_{cip} , a mobile reports to its home agent the IP address of N_{cip} 's border router as its CoA, and uses a locally assigned IP address when it roams in N_{cip} . To track the location of each mobile, routers in N_{cip} monitor the packets originating from each mobile and maintain a hop-by-hop reverse path from each router to each mobile. Idle mobiles send dummy packets to the border router with low frequency to help routers maintain the reverse paths to reach them. To keep the overhead low, only a subset of the routers maintain reverse paths for idle mobile nodes. The routers use timeout to remove obsolete path states.

When the border router receives a packet destined to a mobile, it sends out a query using scoped flooding. If a receiving router knows how to reach the destination mobile, it forwards the query to the corresponding interface; otherwise, it forwards the query to all its interfaces except the one from which the query came. Once the mobile receives the query, it sends a route-update message to the border router, setting up a precise reverse path with short timeout value through all the routers along the data path, via which the border router forwards packets to the mobile.

GLOBAL-ROUTING-BASED MOBILITY SUPPORT

Supporting mobility through dynamic routing is conceptually simple: one can broadcast the latest location of a mobile or a mobile network throughout the entire Internet so that all routers know how to reach the mobile (Fig. 1b). Two example designs are described below.

Boeing Connexion Service — Boeing deployed the Connexion service [9] during 2004–2006 that enabled travelers on board of a plane to access the Internet. The design assigns a permanent /24 IP address prefix to each mobile network (in this case an airplane) and uses Border Gateway Protocol (BGP) to propagate the reachability to these prefixes. When an airplane with prefix P moves from access router R_a on the ground to another access router, R_b , R_a withdraws the prefix P and R_b announces P ; thus, the attachment point changes of each airplane are propagated to the rest of the Internet. Conse-

quently, this design raises routing scalability concerns: if the number of mobile networks becomes large, the amount of rapid BGP updates will also increase proportionally, which could lead to severe router overload.

WINMO —WINMO, which stands for Wide-Area IP Network Mobility, aims to address the routing scalability concern of Connexion [10]. Like Connexion, WINMO also assigns each mobile network a stable prefix. However, WINMO can reduce the total number of updates originating from mobile networks by orders of magnitude through two new approaches. First, WINMO uses various heuristics to limit the propagation scope of routing updates caused by mobile movements. As a result, not all routers may know all mobiles' latest locations. Second, WINMO adopts the basic idea of MIP by assigning each mobile network a "home" in the following way. WINMO assigns mobile network prefixes out of a small set of well-known mobile prefixes. These mobile prefixes are announced by a small set of aggregation routers, which also keep track of all mobile networks' current locations. Thus, these aggregation routers play a role similar to home agents in MIP and can be used as the last resort to reach mobile networks.

END-TO-END SOLUTIONS

The protocols described so far, with the exception of MIP's route optimization, share a common assumption that the CN is unaware of the mobile's movement. A different class of solutions have also been developed that let each CN keep track of the mobile's latest location, thus providing mobility support without requiring either intermediates (e.g., home agents in MIP) to forward packets or network routing to track mobiles' locations. We call this class of solutions end-to-end solutions. Typically, they use the existing Domain Name Service (DNS) infrastructure to store the mapping between a mobile's ID and its current location. As shown in Fig. 1d, a CN can query the DNS to obtain the location of mobile M at the beginning of data exchanges, and M is responsible for updating both the CN and DNS server with its latest location. If the CN and M move simultaneously (and hence the former misses the latter's update message), it can always resort back to query DNS again to find M 's latest location.

A well-known example of end-to-end solutions is the Host Identify Protocol (HIP) [11]. HIP assigns to each host an identity made of cryptographic keys, and adds to the protocol stack a new host identity layer between the transport and network layers. Host identities are used to identify mobile nodes, and IP addresses are used for data delivery purposes. In order to reuse the existing IPv6 implementation code, HIP uses the host identity tag (HIT), which is a 128-bit hash value of the host identity, in transport and other upper layer protocols.

HIP can use DNS to store the mappings between mobiles' HITs and IP addresses, or use its own mapping infrastructure made of a set of rendezvous servers (RVSs). Each mobile node M has a designated RVS, which tracks the current location of M . When a CN wants to com-

municate with M , it queries DNS with M 's HIT to obtain the IP address of M 's RVS and sends the first packet to the RVS, which then forwards it to M . After receiving this first packet, M can start communicating with CN directly. When M changes to a new IP address, it notifies CN by sending an HIP UPDATE message with LOCATOR parameter set to its new IP address. It also updates the mapping in RVS.

A MULTI-FACTOR DESIGN SPACE

After a comprehensive overview of the mobility solutions that have been proposed or standardized so far, in this section we contrast and compare different design choices and evaluate their suitability in providing mobility support for V2I communications.

ROUTING-BASED VS. MAPPING-BASED APPROACH

All existing mobility support designs can be broadly classified into two basic approaches. The first is to support mobility through dynamic routing. To deliver packets destined to mobiles to the right place, the network must continuously keep track of each mobile's movement, by either maintaining a path to reach each mobile (e.g., Cellular IP) or reflecting the mobile's current position in the network on the routing table (e.g., Connexion). However, maintaining a path to every mobile host or otherwise informing the whole network of every movement of every mobile incurs a high cost. Thus, this class of solutions is feasible only in small-scale networks or in support of a small number of mobiles. It does not scale in large networks with gigantic numbers of mobiles, which is necessarily the case in supporting Internet-enabled cars.

The second approach to mobility support is to support mobility through providing a mapping between a mobile's stable identifier and its dynamically changing IP address. MIP and HIP serve as typical examples here. When a mobile moves, it only needs to update a single mapping server (e.g., home agent in MIP or RVS in HIP) of its location changes. Another advantage of this class of solutions is the decoupling between mobility support and the networks a mobile may traverse. Take MIP as an example: MIP provides Internet-wide mobility support through the deployment of home agents and requires no changes to any network otherwise. Since home agents can be provided by anyone, MIP offers the flexibility of delegating mobility support to any appropriate parties. This is in sharp contrast with those mobility solutions only mobile providers can offer, as we discuss next.

OPERATOR-CONTROLLED VS. USER-CONTROLLED APPROACHES

By and large the global mobility support today is provided by cellular networks. Different from Mobile IP, cellular networks use a service model that bundles together mobility support with device control and network access control. The huge success of the cellular market clearly indicates that the current cellular service model is

All existing mobility support designs can be broadly classified into two basic approaches. The first is to support mobility through dynamic routing. The second is to support mobility through providing a mapping between a mobile's stable identifier and its dynamically changing IP address.

Generally speaking, mobility support should be global in scope; that is, a mobility protocol should enable communications with a mobile independent of the geographic or topological distance between a correspondent node and the mobile. MIP and HIP are examples of such protocols.

viable and is likely to continue into the foreseeable future. Consequently, there have been many efforts in the IETF in recent years to develop mobility support standards that follow the cellular network model, and PMIP represents one such outcome.

As described earlier, PMIP performs both network access control and mobility support, and the mobility support in PMIP is provided by the combined functions of the LMA and MAG, leaving a mobile completely unaware of its roaming. One main argument for this approach is backward compatibility. By not requiring a mobile to participate in its own mobility signaling, one avoids any changes to legacy mobile nodes while providing them the same level of mobility services as the most advanced mobile devices. According to 3G vendors and operators, this network-provided mobility support is a key aspect of success, as they learned from their deployment experience.

However, PMIP supports mobility within one operator's network only; interprovider roaming requires commercial agreements between providers. Since PMIP makes a mobile unaware of its mobility and relies solely on the network to provide mobility support, the mobile cannot get connected in any network other than its service provider or its provider's roaming partners' networks. Furthermore, the mobile is limited to using only the communication media supported by its provider network. However, to support vehicle networking, it seems unacceptable to have Internet connectivity only when a car stays within a provider's network or to confine its movement within the provider's roaming partners, or to get connected through only one particular medium instead of using any of the communication media it may possess.

Fortunately, most of the mobility support protocols we discuss in this article neither confine a mobile within a specific operator's network nor force the mobile to use a particular access medium. Mobile nodes typically need to update their locations by themselves to the rendezvous points chosen by end users. In these protocols, network access is decoupled from mobility support, and mobiles have the flexibility to connect through any available access channels. Although the mobiles must participate in their own mobility support, in return they obtain flexibility and independence from their service providers. They can choose whatever mobility services are available as long as their software supports that protocol, and they can utilize anyone's network to get connected.

To connect cars to the Internet, manufacturers face choices between network operator controlled and user controlled mobility support. It is likely that both approaches may be adopted at this time. However, as the technologies continue to advance, and mobile devices become more and more powerful, operator-independent mobility support seems to offer a much simpler and more flexible way to move forward.

LOCAL VS. GLOBAL MOBILITY SUPPORT

Discussions on mobility support often involve the concepts of local and global mobility support. Before we proceed to discuss the roles of

local and global mobility support, let us first clarify the definitions of local vs. global mobility.

Generally speaking, mobility support should be global in scope; that is a mobility protocol should enable communications with a mobile independent of the geographic or topological distance between a correspondent node and the mobile. MIP and HIP are examples of such protocols.

A local mobility management protocol, by definition, works within a *local* domain. However, different communities have different interpretations of what defines a local domain. One definition of local domain is a relatively small network or geographic area (e.g., a university campus or campus network). Although MIP or HIP can directly support roaming within a local domain, if some local mobility support is provided, as is the case of a campus network running Cellular IP [7] or Hierarchical MIP [8], not only might a mobile experience reduced handoff delay while roaming inside, but the number of location updates sent by the mobile to its global rendezvous point can also be substantially reduced.

Mobile operators, on the other hand, define local vs. global mobility as whether a mobile moves within a provider's network or across the boundary to a different provider's network, respectively. A mobile may roam across the continental United States and be considered as local mobility if it is still within the same operator's network, or it may be considered to have global mobility even though it can only be connected in three providers' networks. One basic assumption about mobility in cellular networks is that a mobile spends most of the time within the coverage of its home service provider, which may not necessarily hold true for Internet-enabled cars.

MOBILITY AWARENESS

Another dimension to group different mobility solutions is by their assumptions on whether one or both ends of a data exchange are aware of and participate in mobility support.

The first group of solutions hides the mobility from the CN. In this approach, a mobile has a home agent or its equivalent, which keeps track of the mobile's current location and forwards packets to it. The mobile's responsibility is to signal its home agent for location changes, and a CN can be totally unaware of a mobile's movement. This group of mobility solutions can be offered independent of mobile service providers.

The second group of solutions hides the mobility from both the mobile and the CN; thus, the mobile relies solely on the network to provide mobility support. As a result, mobiles are tied to a particular service provider, and lose the freedom to choose the best access method when more than one are available or easily roaming between networks of different providers.

The third group of solutions lets both mobile and the CN participate in mobility support. Consequently the network does not have to be mobility-aware and does not need to install any new functions to support mobility. As mobile devices continue to grow in quantity and variety, and start utilizing all available communication channels, it seems as though this design choice will

become more appealing. One common approach taken by this group of designs is to use DNS to keep track of mobiles' current locations. Given that the DNS infrastructure is ubiquitous and reliable, it makes mobility support easy to deploy, and the data path is always the shortest one.

SPECIAL REQUIREMENTS IN AUTOMOBILE MOBILITY

None of the existing mobility protocols is specifically designed to support V2I communications. In this section we discuss the requirements that arise from automobile mobility support to help put the design trade-offs discussed in the previous section in the context of the automobile environment.

INDEPENDENCE OF NETWORK SERVICE PROVIDERS

Even though today's mobile and wireless technologies are mature and widely deployed, automotive manufacturers still face stagnation in offering networked mobile services. Currently, the dominant mobile services available for cars are provided by cellular networks. A common practice is that a car manufacturer bundles together Internet connectivity service with a car as a package to sell to customers, where the Internet service is provided by a specific cellular operator. Although a cellular network operator can have wide coverage for residential areas, the geographic coverage is often insufficient when cars move nationwide or continent-wide. When a car moves beyond the operator's coverage (e.g., across a country border), it can be connected only if the operator has a contract with the foreign operator and at a much higher cost. Furthermore, cars are frequently driven through underground tunnels or parked in underground garages, where there is no or poor cellular coverage, but possibly WiFi coverage provided by other parties. Therefore, it is important to enable a car to utilize any available access network to establish Internet connectivity, be it Long Term Evolution (LTE), High-Speed Packet Access (HSPA), WiMAX, or WiFi, and regardless of who the provider may be.

The model of binding a car's Internet service to a specific provider may have been inherited from cell phone services. However, unlike smartphones, which have a short life span, a car's average lifetime is about 10 years. Technologies are most likely to advance in significant ways during such a long time span, and the provider market may change during this time as well. Since modification to a car after sale (also known as recall) is a nontrivial operation for car manufacturers, it is highly desirable to have provider-independent Internet service and the ability to update the mobility solution over the lifetime of the car as technology advances.

LEVERAGING WiFi ACCESS

In utilizing multiple access technologies, WiFi access deserves a special attention. WiFi is one of the cheapest solutions for high-speed Internet

access, especially compared to 3G/LTE. It is both challenging and promising to leverage the booming WiFi networking accessible from the vehicle on the road [13]. WiFi coverage has been tremendously increased in recent years and will be even more ubiquitous in the near future, especially in densely populated urban areas. Nonetheless, WiFi access points have limited range considering the high speed of the moving vehicle. Currently, the long connection establishment latency and poor handoff strategy (stick to an access point until it is disconnected) make it difficult for vehicles to benefit from WiFi. Recent research [13, 14] proposed various techniques to overcome these limitations and, together with 3G networks, provide vehicles Internet access with minimal disruption.

We believe that the importance of utilizing ubiquitous WiFi coverage will grow as time passes. The emergence of electric vehicles (EVs) adds further incentives to explore WiFi connectivity. During battery charging, either at a station or at home, it is likely that a car is within WiFi coverage, and this network connectivity can be exploited not only to report charging status notification, but also to update firmware of sensors/actuators or dashboard software since the car is not in operation. This means that a car may need the communication module to run while in parking. In such a scenario, 3G/LTE would be too expensive compared to WiFi, which is most likely provided by parties other than cellular providers.

CONTROL OF MOBILE NETWORK SERVICES

Smartphones and tablet devices are becoming part of daily life, making people connected all the time with each other and their favorite applications. Consequently, users may bring them into cars to enjoy the new digital life everywhere. New smartphone applications have also been developed to enhance users' driving experiences.

On the other hand, car manufacturers desire a certain degree of control over the Internet connectivity and services of the car to fulfill safety requirements when cars are in motion. Car accidents caused by any distractions may result in driver injury or even death. Safety and security are the most important factors for car manufacturers.

Thus, a big question remains open regarding the exact roles of user networked devices in cars given the manufacturers' desire to control all services in the car for safety concerns. Car manufacturers may want to filter content and services on smartphones for safety reasons.

Given the finite resources on a smartphone (e.g., limited screen size and battery lifetime), a second question is whether a car should provide application programming interfaces (APIs) for smartphones so that they can access resources in the car such as power, display, or sound system. Developing such an API can be challenging as there are a large number of smartphone companies manufacturing phones with different hardware and software systems, and at a pace several times faster than that of the car manufacturing industry.

Yet another question concerns network con-

Unlike smartphones, which have a short life span, a car's average lifetime is about 10 years. Technologies are most likely to advance in significant ways during such a long time span, and the provider market may change during this time as well.

Exciting and challenging times are facing car manufacturers. They need to get out of their comfort zones and work with network research community to develop most effective solutions to inter-connecting cars.

nectivity and usage. Generally speaking, the Internet traffic generated from the car has two types: manufacturer-related (diagnostic information, maintenance notifications, etc.) and user-related (e.g., road assistance or entertainment). The manufacturers may pay for the first type of traffic, but may not be willing to pay for the latter. Hence, necessary accounting functions are needed to distinguish the two. Therefore, the coexistence of smartphones and in-vehicle dashboard devices with their own connectivity seems to be a reasonable assumption in the near future.

LOOKING FORWARD

Internet-enabled cars are coming. We have seen fancy commercials of car dashboards with network connectivity. However, those are connected via special providers. Cars are yet to become first class Internet citizens such as laptops or smart phones that can get onto the Internet via any available channels. Based on our study of existing mobility protocols, we envision the solutions to connecting cars to be mapping-based and user-controlled. Furthermore, we also believe that moving vehicles would benefit from the recently initiated National Science Foundation (NSF) expenditure on future Internet architecture [15], which explores new architectures that focus on content and regard mobility as the norm. Car manufacturers face exciting and challenging times. They need to get out of their comfort zones and work with the network research community to develop the most effective solutions for interconnecting cars.

REFERENCES

- [1] <http://www.autonetmobile.com/>, 10-18-2010.
- [2] Lexus ENFORM Mobile, <http://www.lexus.com/enform/mobile/index.html>, 10-18-2010.
- [3] Z. Zhu, R. Wakikawa, and L. Zhang, "A Survey of Mobility Support In the Internet," Internet draft, draft-zhu-mobility-survey-03.txt, 2010. 7
- [4] V. Devarapalli et al., "Network Mobility (NEMO) Basic Support Protocol," RFC 3963, 2005.

- [5] R. Wakikawa, G. Valadon, and J. Murai, "Migrating Home Agents Towards Internet-scale Mobility Deployment," *ACM CoNEXT*, 2006.
- [6] S. Gundavelli et al., "Proxy Mobile IPv6," RFC 5213, 2008.
- [7] A. Valko, "Cellular IP: A New Approach to Internet Host Mobility," *ACM SIGCOMM*, 1999.
- [8] H. Soliman et al., "Hierarchical Mobile IPv6 Mobility Management (HMIPv6)," RFC 4140, 2005.
- [9] L. Andrew, "A Border Gateway Protocol 4 (BGP-4)," Boeing White Paper, 2006.
- [10] X. Hu et al., "Wide-Area IP Network Mobility," *IEEE INFOCOM*, 2008.
- [11] P. Nikander et al., "Host Identity Protocol," RFC 5201, 2008.
- [12] L. Zhang, R. Wakikawa, and Z. Zhu, "Supporting Mobility in the Global Internet," *ACM MOBICOM MICNET Wksp.*, Beijing, China, 2009.
- [13] A. Balasubramanian, R. Mahajan, and A. Venkataramani, "Augmenting Mobile 3G Using WiFi," *ACM MobiSys*, San Francisco, CA, 2010.
- [14] P. Deshpande et al., "Predictive Methods for Improved Vehicular WiFi Access," *ACM MobiSys*, Krakow, Poland, 2009.
- [15] <http://cacm.acm.org/news/98249-nsf-announces-future-internetarchitecture-awards/fulltext>, 8-30-2010.

BIOGRAPHIES

ZHENKAI ZHU (zhenkai@cs.ucla.edu) is a Ph.D. student in the Computer Science Department of the University of California, Los Angeles (UCLA). He received his B.E. in automation from Tsinghua University in 2008. His research interests include mobility support for mobile devices and named data networking.

RYUJI WAKIKAWA (ryuji@us.toyota-itc.com) received his Ph.D. degree in 2003 and his M.E. degree in 2001, both in media and governance from Keio University. He has been a senior researcher at Toyota InfoTechnology Center, USA, Inc. since 2008. His research interests are mobile computing, Mobile IP, mobile ad hoc networks, and mobile networks. He received the 10th Ericsson Young Scientist Award in 2007.

LIXIA ZHANG [M] (lixia@cs.ucla.edu) is a professor in the Computer Science Department of UCLA. She received her Ph.D. in computer science from MIT and was a member of research staff at Xerox PARC before joining UCLA in 1996. In the past she served as the vice chair of ACM SIGCOMM, a member of the editorial board for *IEEE/ACM Transactions on Networking*, a member of the Internet Architecture Board, and co-chair of the Routing Research Group of the Internet Research Task Force. She is a fellow of ACM and the recipient of the 2009 IEEE Internet Award.