The Intelligent Transportation Society of America (ITS America)

Connected Vehicle Technical Insights

Vehicle Applications and Wireless Interoperability

Heterogeneous Networks, "Multi-Path" Applications and their Impact on Transportation

Steven H. Bayless and Adrian Guan

Technology Scan Series 2011-2015

Contents

Introduction

The GSM Association predicts that the growth of embedded, cellular-based telematics will reach nearly 11 million units per year by 2020, which would encompass nearly every new vehicle that rolls off U.S. assembly lines. Along with cellular, *Wireless Fidelity* (Wi-Fi) will also likely appear embedded in most vehicles to support "local area" connectivity inside the vehicle. Local area connectivity outside the vehicle will be supported by *Dedicated Short Range Communications* (DSRC), a radio access technology similar to Wi-Fi, but is used for time-critical applications such as vehicle-to-vehicle crash avoidance.

These vehicles will connect using multiple radio access technologies, all of which rely upon conventional inter-networking standards such as *Transmission Control Protocol-Internet Protocol* (TCP-IP). DSRC simultaneously supports both conventional TCP-IP connections as well as a specially designed *Wireless Access for Vehicular Environments* (WAVE) *Short Message Protocol* (WSMP) for time-sensitive, safety-critical crash avoidance applications. It is estimated that nearly 80% of overall traffic over the internet relies on TCP-IP.

The foundation for nearly every popular internet-enabled application has been TCP-IP. While IP is more famous – resolving routing between the data source and destination – TCP is critical. TCP manages the connection, providing a reliable, ordered, and error-checked delivery of a stream of bytes (basic units of digital information) between programs running on computers connected to a network. Connections must be properly established in a multi-step handshake process (connection establishment) before entering the data transfer phase. After data transmission is completed, the connection termination closes established virtual circuits. In a connection-oriented, packet-switched data link layer or network layer protocol, all data is sent over *the same path* during a "communications session."

The original designers of TCP-IP never imagined that one day wireless mobile devices may roam, changing location or paths in the middle of a communication session. During roaming, data communications break down and must be reestablished once the user reestablishes a single network pathway – a process completed once a device leaves one wireless area network and connects to another. The original designers also did not imagine that a mobile device might be connected to multiple wired or wireless networks simultaneously. To address this opportunity, the *Multi-Path TCP* standard was established in an ongoing effort of the Internet Engineering Task Force's (IETF) Multi-Path TCP working group, which aims at allowing a TCP connection to use multiple communications paths, known as *inverse multiplexing*.

In the near future, most mobile device operating systems (OS) will support Multi-Path TCP as standard. Multi-Path TCP is currently embedded in Apple iOS to support *Siri*, a personal assistant application, and is the first visible large-scale commercial application of the protocol. Apple's *Siri* relies on Multi-Path TCP to rapidly stream very large audio files, leveraging multiple networks to centralized servicers for

intensive voice-to-text processing. Multi-Path TCP aggregates network capacity by combining Wi-Fi and cellular together, treating them as one network from the perspective of the *Siri* application. Multi-Path TCP enables *Siri* to transmit megabytes of data and still provide millisecond response times for users seeking to receive automated responses to their "voice search" queries. This form of multi-network *aggregation* does not require coordination with the cellular carrier or the provider of Wi-Fi services. Multi-Path TCP not only allows the aggregation of multiple wireless networks into one "network of networks," but also allows them to be used interchangeably on the fly.

Apple iOS' use of Multi-Path TCP allows *Siri* to switch dynamically between wireless media when new network access appears or when old networks disappear – a concept known as "vertical" or wireless media-independent roaming. Such multi-radio network access and interoperability, if widely implemented, may be of great benefit to a number of vehicle/road infrastructure mobility (non-safety critical) applications, in a similar way to how it has benefited users of Apple *Siri*.

However, a number of constraints do still exist, the most significant of which are related to security. Multi-Path TCP by itself cannot address the difficulty of authenticating to multiple networks simultaneously, though this problem is being slowly overcome. Multi-Path TCP must also be monitored carefully to reduce vulnerability to certain denial-of-service cyber-attacks. Over time, such constraints may be overcome as network carriers (e.g. mobile network operators and cable companies) form consortiums to allow, for example, Wi-Fi users to roam with the same security credentials across different operator networks.

This technical insight report suggests that devices and networks may soon incorporate wireless inverse multiplexing standards like Multi-Path TCP-IP and/or possibly other schemes in large numbers. This report describes the momentum behind these technologies, and some of the technical and operational challenges to widespread adoption. The report also suggests that if these challenges are overcome, widespread adoption of standards such as Multi-Path TCP may positively influence the evolution of the vehicle-infrastructure communications architecture as envisioned within the U.S. Department of Transportation's (USDOT) *Connected Vehicle* program. In 2014, USDOT's National Highway Traffic Safety Administration (NHTSA) initiated a rulemaking process that may result in a rulemaking proposal to create a new Federal Motor Vehicle Safety Standard (FMVSS) which would require vehicle-to-vehicle (V2V) communication capability for light vehicles.

Varieties of Wireless Technologies in the Modern Vehicle

For most recent model year vehicles, a reasonably priced mid-market light passenger vehicle might include wireless AM/FM and Satellite Radio, multi-media device USB, Bluetooth, Wi-Fi, and remote direct-access telematics (cellular 2G-4G). The growth of machine-to-machine (M2M) communications in automotive contexts will likely be more rapid than anticipated, as mobile devices enable inexpensive aftermarket devices that utilize the common On-Board Diagnostics (OBD-II) port or other platforms. The expectations of consumers and businesses will result in the introduction of BYOD (Bring Your Own Device) to vehicles, since road users will want to take advantage of useful mobility, logistics, and "infotainment" applications and services. A number of these radio access technologies support TCP-IP directly, or encapsulate it or translate it within a proprietary network. There are two broad categories of vehicular communications: intra- and inter-vehicle *safety-critical* communications, and *non-safety-critical* intra-vehicle and extra-vehicle communications.

Intra-vehicle safety communications are typically high-reliability wired networks but may also include wireless. Next-generation vehicles may also expand the use of wireless networks for intra-vehicle communications using Wireless Controller Area Networks (WCAN), Wireless Keyless Entry, and Tire Pressure Monitoring Systems (TPMS). WCAN allows powertrain, control, chassis, body, and telematics systems within vehicles to communicate without extensive and heavy wire harnessing. WCANs add redundancy and reliability to the traditional wired CAN. WCANs also allow non-critical features to be added without increasing the vehicle's weight, thus increasing the space and reconfiguration of components on the vehicle platform. Instead of TCP-IP, intra-vehicle networks typically use propriety protocols such as CAN and FlexRay.

Applications for intra-vehicle communications have developed over the last decade. WCANs allow reliable monitoring of parameters belonging to moving or rotating parts such as the wheels, and improves the reliability of vehicle systems through the mitigation of common failure modes by means of dissimilar redundancy (that is, WCANs provide alternative connections to braking systems in the event a wire is damaged or severed in the vehicle). Mandated by NHTSA in 2007, TPMSs alert drivers to severe under-inflation of tires, a condition which can reduce the safety and fuel efficiency of vehicles. WCAN may also easily connect rear-view cameras or other sensors that are intended to allow drivers to see completely around the vehicle to avoid obstacles. NHTSA incorporated new rearward visibility requirements into the Federal Motor Vehicle Safety Standards (FMVSS) in 2014.

The next generation of light vehicles may utilize inter-vehicle safety – i.e., Vehicle Area Network (VANET) communications. Dedicated Short Range Communications (DSRC) technology supports both vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) applications. NHTSA announced in 2014 that it intends to begin a rulemaking process to standardize vehicle-to-vehicle communications for short-range crash

avoidance applications. DSRC roadside units (RSUs) running vehicle-to-infrastructure applications may also be deployed to support traffic safety, mobility, and environmental applications. These roadside DSRC "hotspots" may ultimately be located at the nearly 300,000 signalized intersections (and potentially many other locations) across the country. Ultimately, if the deployment of DSRC and Wi-Fi expands, vehicles equipped with telematics systems will likely be connecting to DSRC and/or Wi-Fi networks while they are parked or as they move in and out of smaller coverage areas such as home driveways, parking garages, or even intersections and other confined spaces.

The WAVE/DSRC protocol supports VANET in a way that is similar to the way Wi-Fi supports Wireless Local Area Networking indoors. However, unlike your local Wi-Fi hotspot, vehicles and traffic signals operating WAVE/DSRC transmit safety-critical application data that is "highly local" (e.g., relevant within a short distance, event-driven, close enough that vehicles could potentially crash) and "highly current" (moving at speeds that provide only short response times for drivers). Such applications include vehicle-to-vehicle crash avoidance, through which vehicles provide warnings to other vehicles (e.g. Forward Collision Warning, Cross Traffic Warning), or through which a traffic control device may coordinate vehicle movements, such as at an intersection, traffic light, or tolling gantry. NHTSA intends to begin a rulemaking process on the incorporation of DSRC into vehicles in the next several years; concurrently, the Federal Highway Administration (FHWA) will issue guidance to road operators on the deployment of traffic control and other devices using RSUs. *Diagram 1* summarizes the variety of utilized vehicle communications, navigation and active sensors that will likely be found in most next generation vehicles.

Diagram 1: Vehicle and Infrastructure Communications, Navigation, and Active Sensing Technologies

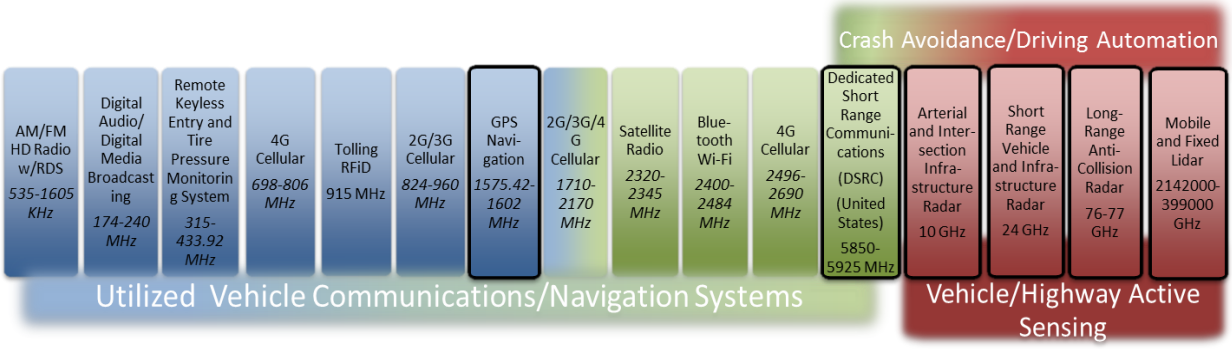

Source: ITS America

Intra- and extra-vehicle communications for non-safety-critical services, such as for infotainment and mobility, are widely deployed and are familiar to most drivers. Two-way cellular is chosen for most telematics mobility applications because of its nationwide, nearly ubiquitous coverage, as well as for the ability of applications using cellular to "roam" across operators and coverage areas without any significant loss of continuous service. Satellite radio also offers continuous geographical coverage over

much of North America and is often utilized as a channel to broadcast mobility data to vehicles. Both satellite and digital terrestrial radio often provide information services, such as weather and traffic, which are universally relevant to drivers within wide geographical areas, such as in large cities or regions.

Intra-vehicle wireless systems include Bluetooth, ZigBee, and Wi-Fi. Bluetooth is primarily used to bridge mobile phones to the head unit of the vehicle to support comfortable and hands-free telephony, infotainment, and other mobility services. Wi-Fi is occasionally used to function in the same way, or to act as a bridge between the cellular telematics broadband connection embedded in the vehicle and mobile devices in the car, such as back seat infotainment systems or tablet computers. Wi-Fi is rarely used to connect the vehicle to a Wi-Fi hotspot, though that would likely serve as a future interface to support connectivity in places such as repair shops or dealerships.

Broadband extra-vehicle connectivity is also used for a number of non-safety-critical applications. For drivers and passengers, mobility information services reduce the cost of searching for information and reduce the uncertainty regarding trip routes, travel times, costs, payment options, and other needs. Mobility services may also include many other broad operational goals such as improving safety or reducing the environmental impact of travel over time. One example is eco-routing, which uses personal navigation devices that minimize trip fuel savings and/or environmental impact, rather than travel time or distance.

There are also a number of vehicle- or infrastructure-oriented services that benefit from extra-vehicle mobile broadband connections. Some information services match underutilized assets, such as unoccupied parking spots, idle taxis or rental vehicles, or empty passenger seats with consumer demand for trips or destination amenities. Vehicle-oriented services are often designed to preserve and maintain vehicle assets from wear and damage through remote monitoring, such as vehicle diagnostics, fleet management, and pay-as-you drive insurance. USDOT has research programs that address a number of these mobility and environmental application areas, especially focusing on the operations of and application to public road agencies.

A given safety, mobility, or environmental application typically utilizes only one vehicle-based radio access technology. Coverage area and spectrum availability typically determine which radio access technology would support which application. For example, short-range radio technologies are usually only used for “vehicle area network” applications, and include Bluetooth, ZigBee, Wi-Fi (for vehicle-to-consumer mobile device or vehicle-to-dealer/mechanic connectivity), and WCAN and DSRC for intra- and inter-vehicle safety-critical functions. DSRC, in particular, was chosen for vehicle crash avoidance applications because of regulatory protection assigned to its spectrum allocation. These regulatory protections were intended to eliminate the risk that radio frequency interference may compromise safety-critical automotive functions, and to incentivize the auto industry and road agencies to invest in new crash avoidance technology.

Aggregation of Wireless Media and HetNets

Fourth Generation LTE represents the complete transition of cellular from a system designed for the unique requirements of voice to a general-purpose system that can manage a number of applications. Wi-Fi is ideal for local area networking with a high density of users within Wi-Fi's very limited (100 square meter) footprint. Often these networks overlap, with modern mobile device users able to connect to both one at a time. However, no standard for simultaneously connecting and aggregating different radio access networks has been widely deployed yet, either by mobile device manufacturers or wireless network operators.

The ability to aggregate capacity across different radio access link services (both homogenous and heterogeneous technologies) has been a problem addressed in standards bodies for a while, but no one standard has been widely adopted until now. Two competing standards, Multi-Path TCP and the Stream Transmission Control Protocol (STCP), both address link aggregation. STCP failed to gain adherents in the marketplace because many firewalls and other network devices did not natively support it. LTE Advanced IP Multi-Media System (IMS), with its Multi-Path Real Time Protocol, also supports Multi-Path TCP for 4G cellular networks, but cellular carriers have not widely deployed LTE IMS and/or all of its features.

Multi-Path TCP, however, did not suffer from such network appliance incompatibility, and the first deployment of the protocol was seen with Apple's iOS 7. Multi-Path TCP is enabled only for *Siri*, with no open interface for other applications to leverage. *Siri* was the first large-scale commercial adoption of Multi-Path TCP. Android- and Linux-based devices are reported to be testing Multi-Path TCP packages, and open interfaces may be available soon. The internet relies heavily on two protocols: *Internet Protocol* (IP) and *Transmission Control Protocol* (TCP). In the network layer, IP provides a "datagram" service and ensures that any host can exchange packets with any other host. TCP operates in the transport layer and offers reliable, ordered, and error-checked delivery of data streams.

Standard TCP delivers packets of the same TCP session over a single path using a sequence of links between a sender and a receiver. This path is usually selected to attain the best transport performance, such as minimal delay and maximal throughput (the average rate of successful data stream delivery). One of the early architectural decisions was to bind IP and TCP at the connection establishment time, which is a frustration to mobile users who may change IP addresses as they move in and out of range of different wireless access points. Despite the growing importance of mobile nodes such as mobile phones and tablet computers, TCP connections cannot move from one IP address to another. When a laptop switches from cellular (indoors and outdoors) to Ethernet or Wi-Fi (mostly indoors) it obtains another IP address. All existing TCP connections must be torn down and new connections initiated.

Multi-Path TCP is a set of extensions to standard TCP which provide a Multi-Path TCP service, enabling a transport connection to operate across multiple paths simultaneously. Since TCP works reasonably well in most cases, it is not uncommon that one link's capacity is completely occupied or exceeded, while another has capacity to spare. Sometimes, a request from a very big TCP session may even surpass one single link's capacity and thus cannot be accommodated. This disparity in network resource allocation, constrained by a single path, sometimes incurs transmission latency and throttles transmission throughput.

Multi-Path TCP lifts the constraint of transmitting via a single path by allowing for data segments of one TCP session to flow over multiple paths, which almost always exist between a packet origin and destination in a sophisticated network topology like the internet or wireless communication networks. By splitting a data stream of the same TCP session into multiple sub-streams or sub-flows, Multi-Path TCP is able to spread the traffic of a TCP session over diverse paths that are all chosen adaptively based on the capacity of these paths. As a result, Multi-Path TCP can maximize throughput in TCP sessions and combat the disparity in network resource allocation.

At the core of Multi-Path TCP lie the establishment, termination, and management of sub-flows. First, unlike standard TCP, Multi-Path TCP does not associate a session with a single IP address of a source. Instead, when a source node has additional IP addresses due to multiple network connections, they are passed to the destination and then added to the source's IP address list. As the result of agile addresses, the sub-flows from IP addresses in the same list are grouped with the same source by the destination node. *Diagram 2* illustrates the concurrent sub-flows within a single TCP session.

Second, when one of the IP addresses in the list becomes unavailable due to the loss of a network connection, the IP address is then removed and this sub-flow is terminated; however, the TCP session may be preserved because of other active sub-flows.

Third, it is desirable that a source with multiple available paths will transfer more traffic using the least congested of the paths, achieving a bundle of links which behaves effectively like one shared link with bigger capacity. This link aggregation increases the overall efficiency of the network as well as its robustness to failure. The so-called "coupled congestion control" is proposed around a parameter of the aggressiveness of the Multi-Path flow, which can be estimated from oft-used congestion control parameters, such as packet loss rate and packet round-trip time.

Diagram 2: Illustration of a Multi-Path TCP Session Over Multiple Radio Access Networks

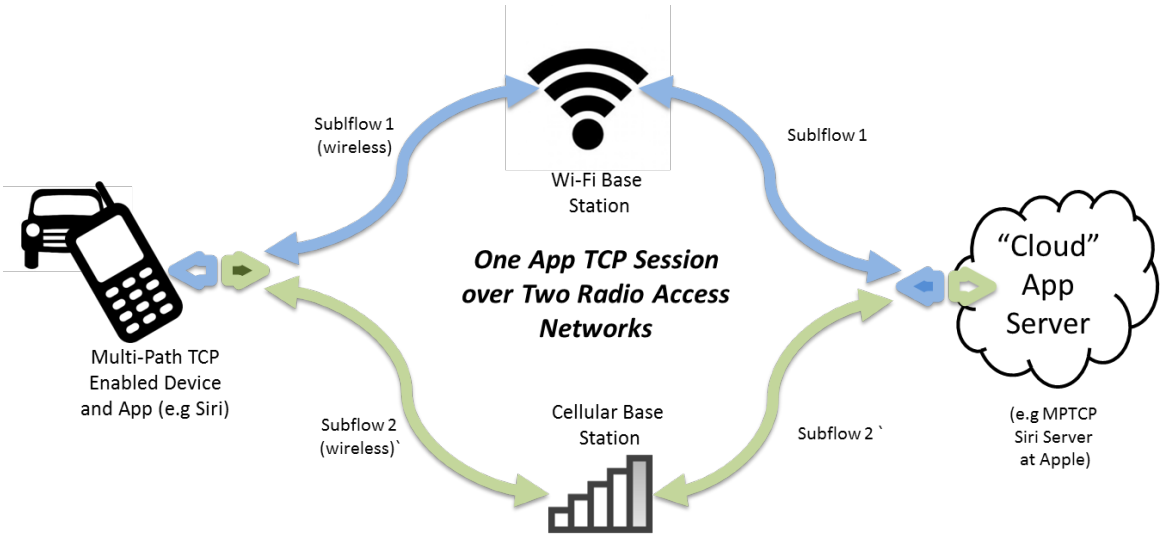

Source: ITS America

Vehicle Mobility and the Challenge of Roaming Across HetNets

As a variant of TCP that is widely supported on almost all network devices, Multi-Path TCP may offer benefits across multiple layers in the Open Systems Interconnection (OSI) model. On the transport layer, a TCP session with multiple paths may adjust transport loads to diverse paths, adaptively based on a network's congestion level. The effect is that the session is more resilient and can thus combat network congestion and session interruption. Communications networks, especially commercial cellular and consumer/commercial Wi-Fi networks, often experience a decline of service associated with congestion (e.g. too many users join and absorb too much network capacity) and may experience system interruption.

On lower layers, including data link and physical layers, Multi-Path TCP's diverse paths enable concurrent access interfaces to diverse networks. Concurrent access, in particular, breeds the potential to ensure seamless handovers in a wireless network – that is, wireless nodes transferring from one communication medium or access point to another. Current cellular systems are unique in that they provide roaming across cells (coverage areas) using a single radio access technology or network – a concept known "horizontal roaming." Multi-Path TCP suggests another layer of roaming may be possible: "vertical roaming." Vertical roaming describes how a user can roam across coverage areas using several different radio access technologies. This concept is also described as "Media Independent Handover"

and is a component of "heterogeneous" networking. *Diagram 3* illustrates heterogeneous inter-networking.

Diagram 3: Illustration of Heterogeneous Inter-networking ("HetNets")

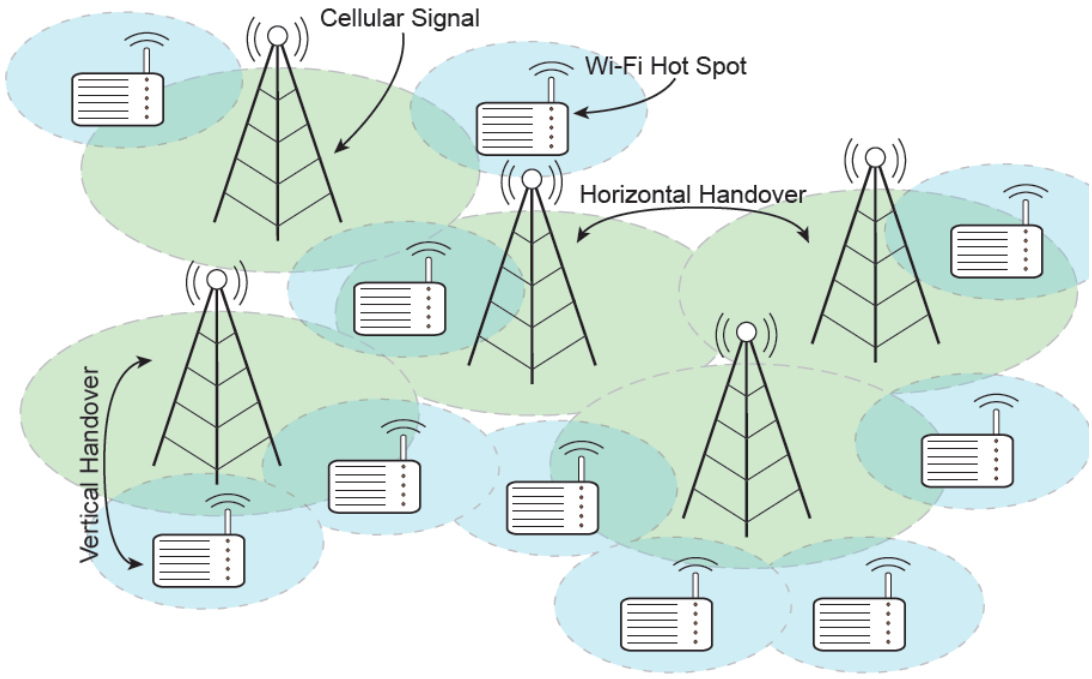

Source: ITS America

A growing number of mobile handsets, including cell phones and tablets, are equipped with interfaces with which to connect to more than one wireless medium: Wi-Fi, 3G, and/or 4G radio access technologies. However, conventional TCP is not able to leverage all these wireless media in one session, since data streams of one transport session can only choose one of these wireless media at a time. When a connection to one wireless medium gets lost or significantly degraded, the transport session suffers from poor performance and, in many cases, may need to be terminated and re-established. Wireless performance at the receiver degrades as signal strength falls – which typically declines as the receiver moves farther and farther away from the base station.

Wireless performance at the base station may also degrade. Degradation typically occurs when too many users connect to a single base station or cell tower, pushing down capacity per user. Capacity per base station is limited by spectral efficiency, which is generally defined as bits/hertz/site for a given radio access technology. When the ratio of users per base station increases, congestion occurs, and application performance either degrades (at best), or outright fails (at worst).

Vertical handovers can be done in two ways: "make-before-break" or "break-before-make." *Make-Before-Break* occurs when one Multi-Path TCP session is associated with more than one wireless medium, so that a session may respond to an interrupted or degraded connection by simply diminishing the traffic load to the decreasingly available medium and increasing the load on another wireless interface in use. In other words, the device predicts that an existing media interface will disappear soon, a prediction usually based on declining signal strength or network congestion, and will switch over to a second medium.

Break-Before-Make handover assumes that an interface is lost suddenly without warning. Multi-Path TCP can react to a failure on the first media interface by enabling a second interface and starting a connection (sub-flow) on it. Once the second connection has been created, the data that was lost as a result of the failure of the first interface can be retransmitted on the new sub-flow, and the connection can continue without interruption.

Throughput fluctuations or service interruptions caused by media switches are also averted, resulting in seamless and transparent vertical handovers. Seamless vertical handovers are crucial to mobile devices due to their highly mobile working environment. During a change of location, mobile devices need to constantly identify and, if needed, switch to the best available wireless media (e.g. the medium with the highest signal strength) without creating friction (e.g. increased latency) in any necessary medium switch. Typically, these media changes are more frequent the faster the mobile device user is moving.

The benefits of Multi-Path TCP over handovers that are handled by a central coordinator (i.e. a wireless network operator) may be exaggerated, however. After all, concurrent paths from sources to destinations with Multi-Path TCP are primarily intended to maximize network throughput, not to enable roaming. Only experimentation and experience with Multi-Path TCP will show how effective it is in enabling vertical roaming. The work of Apple and others in this area will likely be revealing.

Security and Usability Challenges with Aggregated HetNets

There are two security problems that could hold back the promise of Multi-Path TPC if not addressed by network and device manufactures or telecom service providers. The first is authentication and usability. Multi-Path TCP by itself cannot address the difficulty of authenticating to multiple networks simultaneously. Second, Multi-Path TCP must be monitored carefully to reduce vulnerability to some certain denial-of-service cyber-attacks.

Despite the impressive technical design and early commercial adoption of Multi-Path TCP, the agility of the protocol is also one its greatest vulnerabilities. Multi-Path TCP's traffic flow over multiple paths and

agile addresses causes unique security issues, and mobile network providers might limit the scale of its future adoption on mobile devices.

Security Vulnerabilities Inherent in Multi-Path Communications

The first problem is risks related to denial of service. The trust model of conventional TCP is based on unique IP address pairs between packet origins and destinations, but Multi-Path TCP allows for diverse IP pairs within one TCP session. Conventional TCP associates trust between packet origins and destinations through their IP address pairs. Multi-Path TCP has to ensure that additional pairs have the same degree of trust and that the termination of a transmission between original pairs does not compromise the established trust.

By exploiting this agile IP address system, several known types of attacks may be launched. Flooding is one type of attack and is a Denial of Service (DoS) attack that is designed to bring a network or service down by "flooding" it with large amounts of traffic. Flood attacks occur when a network or service becomes so weighed down with packets initiating incomplete connection requests that it can no longer process genuine connection requests. Multi-Path TCP can be leveraged by hackers to use other computers to flood third parties with connection requests.

During this type of "reflective" flooding attack, for example, a hostile attacker at node A ("A" for attacker) may forge and add the IP address of a targeted third party node (node V, "V" for victim) as node A's second available IP address when he sets up Multi-Path TCP. When node A requests data from an unwitting Multi-Path TCP-capable source, like a *Siri* server, this source may be phished to send a sub-flow to the victim's IP address and "flood" the target with unrequested traffic.

This is the real world equivalent of an anonymous scammer (analogous to node A) mail ordering a large quantity of merchandise payable upon delivery to several verified addresses, but then afterwards issuing official "Mail Forwarding Requests" to the Post Office to redirect (phish) all packages destined for the verified address to the address of the victim, someone not willing to pay and accept for delivery. As with a Multi-Path third party flood attack, both the merchant (server), post office (telecom/internet service provider), and the victim (the fake destination of the newly added path) are inconvenienced with the handling of unwanted merchandise.

Although detailed Multi-Path TCP protocol implementation has not been officially adopted in internet standards specifications, the effort of promoting and implementing Multi-Path TCP is well underway. In January 2013, a Request for Comments (RFC) was issued by the Internet Engineering Task Force (IETF) on "TCP Extensions for Multi-Path Operation with Multiple Addresses." In this document, experimental protocol of Multi-Path TCP has been documented for examination, experimental implementation, and evaluation.

The use of forged IP addresses has been a problem in "internetworking" for some time, and mitigating DoS attacks has been a huge drain on resources for most large enterprises or government entities. Multi-Path TCP, however, further muddies the waters by changing a fundamental assumption in internet architecture – specifically that only one IP address is associated with one communicator and communications session. A number of companies such as Apple and Google, as well as the IETF, have been exploring the implications of this fundamental change and will likely continue to experiment.

The Usability Impact of Multi-Network Authentication

The second constraint for the success of Multi-Path TCP is the lack of security credentials that are valid for, or portable across, multiple wireless access points. Authentication to a network refers to the process where an entity's identity is proven, typically by providing evidence that it holds a specific digital identity and the corresponding credential. Even though a user may have Multi-Path TCP on their device and be within communication range of a number of different networks, he or she may not have the security credentials to access any of them.

The most common credentials are usernames and passwords, and others include digital certificates or other unique identifiers. Users that present credentials that are "authentic" are given authorization to access a network. Cellular users are authenticated through their Subscriber Identity Module (SIM), which stores its International Mobile Subscriber Identity (IMSI) and an authentication key that can be presented to any participating network nationwide and internationally. Wi-Fi users, on the other hand, typically must have a separate username and password for every access point hotspot they might encounter.

Two issues that are particularly important for Wi-Fi access are the fact that most Wi-Fi hotspots are relatively small, and that there are many thousands of them. In a densely populated area, a user is often within range of up to 30 different Wi-Fi networks, without knowing whether they are publicly accessible, and if so, how credentials can be obtained.

The performance of enterprise wireless LANs (WLANs) over the past few years, and especially since the introduction of Wi-Fi 802.11n, has evolved to the point where industry analysts now expect Wi-Fi to replace wired Ethernet as the network connection of choice indoors. Outdoor Wi-Fi is appearing as a "bridge" to mobile hotspots, such as from a portable mobile router (i.e. a mobile phone acting as Wi-Fi router using a 3G or 4G cellular network to backhaul to the internet) or a vehicle with an embedded cellular connection, to a mobile device such as a smartphone or tablet computer for passenger use. Google Fiber recently began testing a free outdoor public Wi-Fi system in Kansas City, Missouri, as part of an experiment to build more public Wi-Fi systems in that area.

However, a user must painstakingly inquire in person for credentials to connect to public hotspots that might be found in hotels, airports, retail establishments, or other public areas. This is an inconvenience given that there were 72,156 Wi-Fi hotspots in the U.S. in 2010, with 55.1 percent of public Wi-Fi

locations available at no charge. In 2013, cable operators such as Cablevision, Comcast, Time Warner Cable, Cox Communications, and Bright House Networks declared that they run more than 150,000 hotspots in major cities across the nation.

In an ideal world, a user could access any Wi-Fi hotspot he or she might encounter. The internet was designed for connectivity, not for security, and special care must be taken to secure hotspots. WLAN networks are common attack "vectors," or bases from which to launch anonymous attacks on other systems, or simply to eavesdrop on users connecting to the internet. For most users this is typically only at home or at the workplace, where they have the security credentials that allow them to access their networks.

However, this is the problem space that the Institute Electronics and Electrical Engineering (IEEE) standard 802.11u standard addresses. The 802.11u standard allows a user mobile device to learn more about a network before deciding to connect to it. A hotspot that is 802.11u-enabled can broadcast information about whether it is accessible to the general public and, if so, which "roaming consortia" they belong to and under what conditions they can be used. The 802.11u standard is currently supported by Apple iOS and is incorporated into an effort known as "Hotspot 2.0 / Passpoint." Hotspot 2.0 is a Wi-Fi Alliance certification initiative to provide seamless interoperability for network discovery and authentication, allowing mobile users to roam between Wi-Fi networks without additional credentials.

However, having 802.11u and Hotspot 2.0 installed in mobile devices and Wi-Fi access points still requires a network services operator that can establish and operate one or more roaming consortia. This services operator, known as an *Authentication, Authorization, and Accounting* (AAA) provider, would register users and maintain credentials for them. For example, a cellular mobile device may use its SIM card credentials to connect to an AAA server, with the AAA server then providing Wi-Fi credentials automatically to a user, allowing the user to access both cellular and Wi-Fi at the same time.

Time Warner Cable confirmed its rollout of Hotspot 2.0 in mid-2014. "TWC Wi-Fi-Passpoint" is a national Wi-Fi network that includes Hotspot 2.0 technology on most of its public access points as well as upgraded authentication. Hotspot 2.0 has been slow to take off despite the attractiveness of enabling seamless roaming between Wi-Fi networks. However, as cable providers move to become Wi-Fi network operators, they may provide the impetus for larger scale deployment of Hotspot 2.0. Cable operators first seek to provide value-added, nomadic wireless service for its fixed broadband customers, but may later drive integration with cellular carriers' 4G core networks to allow roaming or offloading between Wi-Fi hotspots and cellular networks. Time Warner Cable, for example, has already established partnership agreements with Verizon Wireless to promote seamlessness between their services.

Implications of HetNets on Mobile Broadband Services

The technological ideal, assuming multiple wireless systems are available, is that any given mobile device can select the most direct, unencumbered, standardized, efficient, secure, and cost-effective wireless path (or paths) through the internet to connect to any other host or terminal. However, even though some of the very big cellular and cable companies' services may develop common authentication and roaming services, there will still be challenges and uncertainties to achieving this ideal.

Multi-Path TCP and similar protocols cannot help applications "choose" which network path is best. Some applications, for example, may require high reliability and capacity, low latency or extra layers of security, such as is typically the case in public safety or crash avoidance communications. Multi-Path TCP applications have little knowledge of the performance, capacity, cost, or security of a given network. Critical applications could certainly benefit from transparency in network operations or common services, such as level of performance, capacity, cost, and security. A mobile device "connectivity manager" that can evaluate multiple networks, and then select the network with resources that best match a given applications needs would be useful, but is still not commercially available.

Fourth-generation cellular, specifically LTE-Advanced, may have a gateway function that can centrally coordinate both vertical roaming and authentication. Within current standards development for LTE-Advanced, the IP Multi-Media System would allow a wireless network operator to "shape" the terminal devices on their network, requiring them to "offload" some application traffic to less congested links such as Wi-Fi networks or even roam to less-congested competitor networks whenever their own network becomes too congested. The goal of the IP Multi-Media System is to expose a large number of cellular operator services as capabilities to application developers, not just support offloading to relieve congestion. However, rollout of the IP Multi-Media System and its features could take years, possibly a decade, to become widely available through 4G cellular carriers.

In the interim, Multi-Path TCP may support the relief of congestion on cellular networks in a way that early versions of 4G LTE alone have been unable to manage. In the United States, consumers used an average of 1.2 gigabytes per month over cellular networks in 2013, nearly doubling in a single year from 690 megabytes per month in 2012. According to Cisco, by 2017, almost 21 exabytes (1 exabyte = 1 billion gigabytes) of mobile data traffic will be offloaded from cellular networks to Wi-Fi and femtocells (small, low-power cellular access points) each month. Without Wi-Fi and femtocell offloading, total mobile data traffic would grow 16-fold, instead of the projected 13-fold growth.

Besides offloading, "virtualization" of mobile network operations and resources may also push the growth of Multi-Path TCP. Mobile virtual network operators (MVNOs) purchase wholesale minutes from traditional mobile network operators and re-sell them to consumers. MVNOs have no network

infrastructure or spectrum, which reduces their cost, but must rely on traditional cellular carriers for connectivity. There are approximately 335 million mobile connections in the U.S, of which around one in every ten is managed through a virtual operator. Most vehicle connections are serviced by MVNOs, for example. (See ITS America's report *Machine-to-Machine Communications - M2M Technology and Potential for the Transportation Sector,* 2011)

Some MVNOs have even turned the cellular services they provide on their head. In order to economize even further, many MVNOs are instead provisioning smartphones to consumers that rely *primarily* on Wi-Fi instead of cellular to provide telephony/SMS and mobile data to their subscribers, using cellular as a backup whenever there is no other alternative. (This service is ideal in urban areas where Wi-Fi coverage is nearly as ubiquitous as cellular.) Rock-bottom Voice-over-IP services such as Google Voice or Vonage are bundled to provide very low-cost mobile telephony/broadband service. Initially, cellular carriers resisted giving MVNOs access to their networks, or provided access only under prohibitive rates or conditions. However, most cellular carriers are generally receptive to virtualization, varying only in the degree of access they accord.

Trends such as Wi-Fi offloading, operator virtualization, and the creation of new categories of mobile apps such as Apple's *Siri* may incentivize mobile device hardware and software providers to provide Multi-Path TCP or similar protocols in all devices when they ship. If technologies such as Multi-Path TCP "virtualize" networks even further, it will be difficult to predict how cellular carriers will react or how disruptive it will be to their business models. However, cellular carriers will find it hard to resist these trends as the demands for connectivity grow.

The Impact of HetNets on Future Connected Vehicle Architecture

Multi-Path TCP or similar protocols could conceivably be adopted to provide vertical roaming between not only cellular and Wi-Fi, but also to support roaming to and from DSRC at intersections and other fixed locations along roadway infrastructure. DSRC could conceivably leverage Multi-Path TCP to allow roaming among overlapping Hotspot 2.0 Wi-Fi access points or even between DSRC roadside units.

Cars stopped at intersections or, more likely, parked within range of a DSRC roadside unit may be able to utilize Multi-Path TCP to connect vehicle applications to the internet. For example, DSRC roadside units could provide local area broadband service the same way Wi-Fi hotspots do and could relieve congestion on other networks (e.g. cellular networks, congested Wi-Fi access points) through offloading. Beyond generic broadband services, however, DSRC could also provide data services for vehicles which would be especially useful for applications that require "bursty" high data rate transmission, such as transmitting recorded or streaming video.

For example, the examination of vehicle telemetry and "video logs" may be one such "bursty" application. Currently, some sophisticated commercial freight fleet managers occasionally stream or transmit recorded video from commercial trucks' forward field-of-view cameras whenever they are alerted to an anomalous event (e.g. detection of sudden braking or steering) by the truck's telematics system. Fleet managers might upload the raw video footage of the event from the vehicle very quickly, using a number of wireless connections simultaneously (e.g. DSRC, Wi-Fi, cellular). Once the video is analyzed at the fleet management center, the fleet manager can call the driver to make sure that the driver is safe and to alert the driver to prevent any further "near crashes."

Other fleets, such as those belonging to first responders (police, fire, and emergency medical), may also be able to take advantage of such services. In the distant future, autonomous vehicles may benefit from fleet managers or telematics service providers who may be able to carefully monitor the performance of safety critical systems. Such Safety Management Systems (SMS) are emerging in the aviation, maritime, and rail domains, and may be adopted by large vehicle fleets as a formal, systematic, explicit, and comprehensive process for managing safety risks.

Some transportation safety and mobility applications may also be enabled by heterogeneous networking, due to its improvement over the performance and reliability of DSRC networks. In particular, Multi-Path TCP may ensure smooth handovers between overlapping DSRC roadside units, especially in dense urban areas. In the highly mobile environment in which vehicles operate, one continuous and lengthy transport session may be challenged by the frequent switching of access points. In USDOT's Vehicle Infrastructure Integration Proof-of-Concept (POC) test in 2008, it was identified that when a vehicle's DSRC on-board unit (OBU) is within the overlapping coverage area of multiple RSUs, issues occurred in the prioritization and service selection across multiple RSUs. In USDOT's Connected Vehicle Safety Pilot of 2010, contiguous RSUs with overlapping coverage were tested again.

Both studies discovered that when a large amount of data needs to be transferred between a vehicle and an RSU, the vehicle generally needs to stay within the range of the RSU until the data transport is finished. However, in many cases, the car left the coverage area of the RSU before the vehicle application could transmit all of the data to the RSU wireless access point. Therefore, moving between RSUs without smooth (in this case horizontal) handovers often terminates an application session.

There are some uncertainties regarding the practicality of using Multi-Path TCP in a highly mobile environment. Although Multi-Path TCP works in "nomadic" mode, where pedestrians with mobile devices are walking (e.g. 1-4 mph) between coverage areas, it is uncertain how well it would perform at vehicle speeds (e.g. 10-55 mph). The establishment and termination of a TCP session incurs additional telecom signaling overhead data traffic (i.e. the messaging that needs to occur to set up the connection), and in a common driving environment, a moving vehicle may remain within the communication zones of multiple RSUs for a time span on the order of ten seconds (e.g. stopped at an intersection). Maintaining

multiple DSRC connections with several overlapping RSUs may be useful for some high-bandwidth applications. However, the vehicle must travel inside the overlapping RSU communications zone for periods longer than approximately ten seconds all the way up to a full minute, which would be rare. Such events may be when road traffic is at a near standstill, or when a vehicle is at a long stoplight or is parked.

There are three potential research tasks related to the use of Multi-Path TCP with DSRC. First, it is uncertain whether Multi-Path TCP may avert such early session terminations for DSRC by keeping multiple communications sessions running over overlapping RSUs, or between RSUs and multiple cellular and Wi-Fi connections. Currently, USDOT is conducting a test at their Southeast Michigan Testbed to determine how Multi-Path TCP might be of utility in supporting handovers among overlapping DSRC hotspots.

Where cars are moving at maximum posted speed limits, the increased communications throughput benefit of connecting simultaneously to multiple RSUs would likely be partially offset by the additional telecom signaling overhead. These trade-offs would need to be examined if a road operator, wireless carrier, or vehicle application service provider were to deploy a DSRC roadside unit with a specific application in mind. The vehicle/infrastructure communications architecture would also need to be reexamined.

Second, it is also uncertain what universe of potential vehicle/infrastructure applications might benefit from the ability to establish concurrent connections to multiple RSUs either while traveling at vehicle speeds or parked. However, few people envisioned the need for Multi-Path TCP either for mobile apps until Apple's *Siri*, so there is nothing that necessarily rules out the possibility that future "Connected Vehicle" application service providers may develop a vehicle app that might leverage Multi-Path TCP.

Conclusion

Big Data describes current and next-generation IT applications where the sheer size of data is the most significant technology challenge: data too big to process conventionally, such as trillions of web pages or petabytes of raw audio/video files. *Stream computing* suggests that not only is *Big Data* too big to process, but in many cases too big to transport wirelessly. The process of upgrading cellular networks from legacy mobile narrowband systems (2G) to broadband 4G Long Term Evolution (LTE) and newer versions of Wi-Fi has been slow and still continues, while the growth in data generated and consumed by mobile devices is explosive (See ITS America's report *Connected Vehicle Insights: Fourth Generation Wireless - Vehicle and Highway Gateways to the Cloud*).

Apple's *Siri* symbolizes a new class of applications, known as "Over-the-Top" or "Big Streaming." In many ways, Multi-Path TCP is a solution to network resource scarcity. Congested networks present problems for special applications like *Siri* that must move very large amounts of data very quickly without interruption. Even after cellular networks have fully upgraded to 4G LTE Advanced, capacity may still be constrained as newer "big streaming" applications demand more communications bandwidth than network operators can provision in short order. Multi-Path TCP would be a useful workaround for any situation in which a radio access network is congested and an application seeks another available network. Apple *CarPlay*, released in 2014, is a version of Apple's iOS designed to function with built-in automobile dashboards; so *Siri*, and by extension the Multi-Path TCP, has already entered the automotive domain.

In the future, DSRC may provide extensive outdoor broadband coverage in urban areas, particularly in city streets heavily trafficked by always-connected vehicles and mobile device-wielding pedestrians (so called vehicle-to-device or V2D). Ultimately, as cellular coverage and capacity expand, as Wi-Fi coverage and capacity grows, and as DSRC begins to be deployed, a connectivity management function that can leverage secure, spare, and underutilized network capacity on the fly would be very useful. LTE-Advanced, Hotspot 2.0, and other telecommunications standards will attempt to address this need to support quality of service, service discovery, selection and security, usage modeling, and pricing from the wireless infrastructure operator side in a more formal fashion. These standards will complement the efforts of Apple and others to provide a bare-bones inverse multiplexing service that the Multi-Path TCP standard provides.

Select Bibliography

1. Bagnulo, M. "RFC 6181 - Threat Analysis for TCP Extensions for Multipath Operation with Multiple Addresses." *RFC 6181 - Threat Analysis for TCP Extensions for Multipath Operation with Multiple Addresses*. Internet Engineering Task Force, Mar. 2011. Web. 7 Dec. 2013. <http://tools.ietf.org/html/rfc6181>.

2. Corbet, Jonathan. "Multipath TCP: An Overview." *[LWN.net]*. N.p., 26 Mar. 2013. Web. 19 Dec. 2013. <http://lwn.net/Articles/544399/>.

3. Chen, Yung-Chih, Yeon-sup Lim, Richard J. Gibbens, Erich M. Nahum, Ramin Khalili, and Don Towsley. "A Measurement-based Study of MultiPath TCP Performance over Wireless Networks." *Association for Computing Machinery (2010): n. pag.* ACM 978-1-4503-1953-9/13/10. ACM, 13 Sept. 2010. Web. 22 Aug. 2014. <http://conferences.sigcomm.org/imc/2013/papers/imc231-chenA.pdf>.

4. *CSWS'12 Proceedings of the ACM Conference on the 2012 Capacity Sharing Workshop*. New York: ACM, 2012. A First Look at Multi-Access Connectivity for Mobile Networking. CSWS'12, December 10, 2012, Nice, France., 10 Dec. 2010. Web. 18 Aug. 2014. <http://conferences.sigcomm.org/co-next/2012/eproceedings/csws/p9.pdf>

5. Dillet, Romain. "IOS 7 By The Numbers: 200 Million Downloads In 5 Days, 64% Adoption Rate, 20 Million ITunes Radio Listeners." *TechCrunch*. N.p., 22 Oct. 2013. Web. 12 Jan. 2014. <http://techcrunch.com/2013/10/22/200-million-devices-running-ios-7-five-days-after-launch-64-of-all-idevices-20-million-itunes-radio-listeners/>.

6. Ford, A., C. Raiciu, M. Handley, and O. Bonaventure. "RFC 6824 - TCP Extensions for Multipath Operation with Multiple Addresses." *RFC 6824 - TCP Extensions for Multipath Operation with Multiple Addresses*. Internet Engineering Task Force, Jan. 2013. Web. 12 Jan. 2014. <http://tools.ietf.org/html/rfc6824>.

7. Paasch, Christoph, et al. "Exploring mobile/WiFi handover with multipath TCP." *Proceedings of the 2012 ACM SIGCOMM workshop on Cellular networks: operations, challenges, and future design*. ACM, 2012.

8. Paasch, Christoph and Olivier Bonaventure. "Multipath TCP." *ACM Queue. Association for Computing Machinery*, 4 Mar. 2014. Web. 22 Aug. 2014.

9. Raiciu, C., M. Handly, and D. Wischik. “RFC 6356 - Coupled Congestion Control for Multipath Transport Protocols.” *RFC 6356 - Coupled Congestion Control for Multipath Transport Protocols. Internet Engineering Task Force*, Oct. 2011. Web. 7 Dec. 2013. <http://tools.ietf.org/html/rfc6356>.

10. United States. Department of Transportation. Research and Innovative Technology Administration. *Vehicle Infrastructure Integration Proof of Concept Technical Description--vehicle Final Report*. By Scott Andrews and Michael Cops. Washington, DC: U.S. Dept. of Transportation, Research and Innovative Technology Administration, 2009. Print.

11. Van Beijnum, Iljitsch. "Multipath TCP Lets Siri Seamlessly Switch between Wi-Fi and 3G/LTE." *Ars Technica*. N.p., 26 Sept. 2013. Web. 4 Oct. 2013. <http://arstechnica.com/apple/2013/09/multipath-tcp-lets-siri-seamlessly-switch-between-wi-fi-and-3glte/>.

About the *Connected Vehicle Technology Scan* Series

Under sponsorship from the *U.S. Department of Transportation* (USDOT) Intelligent Transportation Systems Joint Program Office (ITS-JPO), the *Intelligent Transportation Society of America* (ITS America) is conducting the *Connected Vehicle Technology Scan and Assessment* project.

This scanning series of *Connected Vehicle Insight* reports will assess emerging, converging, and enabling technologies outside the domain of mainstream transportation research. ITS America seeks technologies that will potentially impact state-of-the-art or state-of-the-practice in ITS deployment over the next decade, with an emphasis on the "connected vehicle."

The *Technology Scan Series* notes trends, technologies, and innovations that could influence, or be leveraged as part of, next-generation intelligent transportation systems within the next five to seven years. The series' focus is on developments in applied science and engineering and innovation in data acquisition, dissemination, processing, and management technologies and techniques that can potentially support transportation.

To learn more about the Technology Scan Series, and to see more *Connected Vehicle Insight* reports, please visit http://www.itsa.org/knowledgecenter/technologyscan.

To learn more about USDOT's research in Intelligent Transportation Systems, please visit http://www.its.dot.gov.