8-2013

Enabling Public Safety Priority Use of Commercial Wireless Networks

Ryan Hallahan
Carnegie Mellon University

Jon M. Peha
Carnegie Mellon University, peha@andrew.cmu.edu

Published In
Homeland Security Affairs , 9, 13.

Enabling Public Safety Priority Use of Commercial Wireless Networks

Ryan Hallahan and Jon M. Peha

ABSTRACT

By providing public safety users with roaming access to commercial broadband networks on a priority basis, it's possible to increase the capacity, coverage, and reliability beyond what's possible with dedicated public safety networks alone. This article quantifies the advantages with respect to capacity, showing that by establishing multiple arrangements with commercial carriers in every locality, public safety can access an amount of capacity that has been projected for very serious emergencies without seriously compromising quality of service for commercial customers. However, this article also demonstrates some of the issues that must be addressed when crafting roaming agreements between public safety and commercial carriers. LTE technology provides a wide range of capabilities to support priority and roaming, but these must be used in accordance with policies and governance structures that have yet to emerge. It must be decided whether priority and resource allocation decisions are made in an automated way or with human intervention, and if the latter, the locus of control. Moreover, agreements must find ways to accommodate significant technical differences in commercial networks, even though they all comply with a common (LTE) standard, and to support changes in technology and needs over the coming years. This will require a single entity with the expertise and authority to bridge public safety stakeholders, commercial carriers, and technical standards bodies.

INTRODUCTION

Wireless broadband networks present a unique opportunity to revolutionize the way public safety responds to emergencies, bringing a number of new and important applications to first responders who previously had to rely on only narrowband voice.[1] In order to bring this functionality to first responders, public safety agencies around the world may deploy wireless broadband networks.[2] At the same time, commercial operators will continue to deploy and operate their own broadband systems. At least in the United States, leading commercial operators,[3] and the most prominent public safety network which is now known as FirstNet,[4] are all moving towards the same underlying technology: Long Term Evolution (LTE).[5] This has created an extraordinary opportunity for public safety agencies to make use of the services offered by these commercial networks in a way that complements the capabilities of public safety networks; public safety users could roam onto commercial networks seamlessly, and receive priority access when they do. For example, had this been possible on September 11, 2001 when the communications system used by firefighters in the World Trade Center stopped working,[6] firefighters could simply have roamed onto a commercial network and received the order to evacuate the building. In reality, many of the 128 firefighters still inside when the second tower collapsed probably never heard that evacuation order, and lost their lives as a result.

Roaming onto commercial networks for important communications would be a major shift in both policy and technology from the traditional approach to public safety communications, in which all mission-critical communications are carried over networks that are the exclusive domain of public safety agencies. It would also be a controversial shift. Many in public safety believe they can only rely on systems built for their agency alone, and not on systems that others control and that were designed for a consumer market.

Priority roaming for public safety began receiving increased attention in the United States in 2010 when the US National Broadband Plan (NBP) recommended that public safety be able to roam with priority onto commercial networks (FCC 2010a).[7] As should be expected from a high-level 240-word recommendation, most of the details of this policy were left to be worked out at a later date. The idea of priority roaming drew intense debate, as public safety advocates questioned whether priority roaming could meet public safety needs, especially for mission-critical communications that require a higher degree of reliability than consumers typically demand, and consumer advocates wondered about the impact on commercial users. Commercial carriers have explored how to provide roaming services to public safety, as shown by the Motorola-Verizon alliance,[8] although Motorola and Verizon only offered to support non-mission-critical communications, which would make the approach far less useful to public safety. Ultimately, the future level of interest from industry will depend on the financial terms, technical requirements of prioritization and the costs imposed on commercial networks. Thus, this paper quantifies some of the benefits of roaming to public safety, as well as the impact it would have on the commercial sector.

When Congress funded the nationwide public safety network, they maintained the Federal Communications Commission's authority to require priority roaming, although within limits.[9] Among them, requirements could not be imposed on commercial carriers that used technologies that are incompatible with those used by public safety, commercial carriers must be reasonably compensated, and priority mechanisms cannot be preemptive. Whether the FCC uses its authority and how has yet to be determined, and may be contentious.

FirstNet has also the authority to negotiate voluntary priority roaming agreements with one or many commercial networks. In addition to the long-term benefits of such agreements described in this article, priority roaming agreements can also be put in place quickly at relatively low initial cost, which could help FirstNet succeed despite tight constraints on both time and budget in its initial phases.[10]

FirstNet, the FCC, the commercial cellular providers and tens of thousands of state and local public safety agencies will have to make important decisions about priority roaming, beginning with whether to adopt it. This article will inform those decisions by shedding light on whether and how LTE technology can meet public safety needs when emergency responders roam onto commercial networks, the potential benefits to public safety, and the potential risks to commercial carriers and their customers. The article will therefore address a number of questions. For example, what priority mechanisms exist in LTE technology and how can they be mapped to public safety needs? Would priority roaming meet the capacity needs of public safety, even in major disasters, and would doing so be harmful to consumers who also need to communicate? This article also studies the issues and tradeoffs presented by specific technical and operational design decisions where these have implications for future policy and governance decisions. These raise additional questions. For example, what technical issues need to be addressed in roaming agreements, and is there significant benefit of establishing multiple agreements per region? Technical decisions of design and operations are also intertwined with decisions of organizational planning and governance. Should procedures be established that place a person in charge of resource allocation during certain kinds of emergencies, which would require careful negotiations given the large number of local, state, and federal organizations involved, or should those decisions be automated based on agreements made before the actual emergency?

Providing public safety users roaming access to commercial wireless networks can yield a number of benefits including: (1) increased aggregate capacity and possibly increased cell site diversity, (2) increased coverage, and (3) increased resiliency. To reiterate the benefits of priority roaming more specifically, where both commercial wireless service and public safety wireless service are available, public safety will have access to increased aggregate capacity from both sources.[11] This is a significant benefit,

which this paper quantifies in the next section. Moreover, even greater aggregate capacity is available where commercial and public safety cell sites are not co-located, because devices can connect to the closest tower (or more precisely, the source of the strongest signal) regardless of which network it is associated with. Reducing distance between tower and mobile device can enable a much higher data rate per MHz of spectrum. Second, where commercial wireless service is available but public safety service is not, public safety gains the ability to operate by roaming, effectively increasing geographic coverage. Third, where commercial wireless service is available in addition to public safety service, public safety will have access to more resilient and dependable communications, because communications is possible whenever at least one of the networks is functioning.[12] Public safety can realize some of these benefits whether or not they receive priority access. However, with priority, public safety users can rely on commercial networks to a much greater degree.

There has been other work done on prioritizing public safety traffic on LTE-based networks,[13] but these efforts have focused on developing use cases and user requirements, detailing the implementation of specific LTE mechanisms and demonstrating these mechanisms on a dedicated public safety network. And more recently, the FCC's Technical Advisory Board for First Responder Interoperability presented their "minimum technical requirements for interoperability" to FirstNet. That report touched on LTE priority mechanisms, but the focus of was on ensuring a nationwide level of interoperability for the public safety broadband network.[14] This article, unlike the others, focuses on the broader policy challenges of prioritizing public safety traffic on commercial networks, and the implications technical design decisions have for agreements between public safety and commercial operators.

We assume throughout this work that public safety operates its own network, and supplements that network's capabilities with commercial services. Consider Figure 1, which shows the cellular towers belonging to public safety, which presumably operate in one spectrum band, and the towers of a commercial carrier, which operate in another spectrum band. The region is divided up into a set of "cells" or areas, each of which is served by a different commercial cell tower. The region is also divided into a different set of cells, each of which is served by a different public safety tower. Public safety and commercial carriers may share some physical towers, each with their own equipment on the tower operating in their own spectrum. (An alternative outside the scope of this paper is to operate a single network with shared capacity, which can take a variety of forms.[15]) Both public safety and commercial towers will be connected by connections known as "backhaul" back to high-speed centers of their respective, which are known as "packet core networks." (The internal workings of packet core networks are described in greater detail in the appendix) Within this region, there will be areas where the public safety and commercial coverage overlaps; in these areas a device capable of priority roaming could connect to either network.

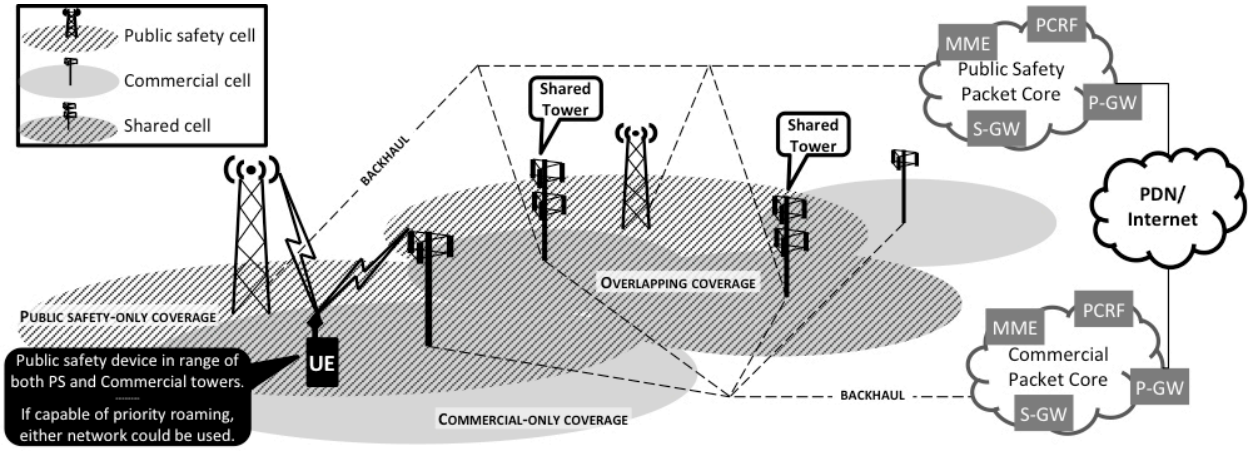

Figure 1. High-level illustration of a public safety network that is deployed in the same region as a commercial cellular network.

Section 2 quantifies some of the benefits of allowing public safety users to roam onto commercial networks. Section 3 discusses the prioritization capabilities that might be desirable to public safety. Section 4 presents the mechanisms available in the LTE standard, and shows that these can be used to meet a wide range of public safety's needs, including protecting critical communications when roaming. Sections 5 and 6, respectively, analyze technical and operational design decisions that must be made, and identify important policy and governance implications. Finally, section 7 presents conclusions.

Many of the observations are rooted in the specifics of LTE technology. The appendix provides an overview of LTE, and highlights the important mechanisms and concepts in the LTE standard that enable preferential treatment of some users and applications.

BENEFITS OF ROAMING AND IMPACT ON COMMERCIAL SECTOR

As discussed in the introduction, one of the primary benefits to public safety of priority roaming is the ability to supplement the dedicated capacity on public safety networks with the capacity available on existing commercial networks. The more commercial networks public safety users have access to, the more capacity that will be available to those users, and the less of an effect public safety roaming traffic will have on commercial subscribers. Thus, to maximize capacity, as well as coverage and resiliency, there is incentive for public safety agencies to craft roaming agreements with as many commercial carriers as possible, although these benefits might be weighed against administrative overhead and transaction costs. This section quantifies the extent to which roaming agreements can give public safety the capacity it needs by estimating utilization levels after serious disasters as the number of roaming agreements and the corresponding amount of spectrum available through roaming is varied. Of course, it also matters what constitutes a serious disaster, so this too is varied. From the carriers' perspective, it is also important to consider how roaming would affect the ability of consumers to communicate during disasters, so this is quantified as well.

To demonstrate how the amount of roaming capacity depends on the number of roaming agreements, Figure 2 shows the amount of commercial capacity per cell sector that is available to public safety in both the uplink and downlink as a function of the

amount of spectrum available for public safety to roam onto. The amount of spectrum available to public safety depends on the number of roaming agreements with commercial carriers, and on the technology and handset design. That is, in LTE there is a predefined set of spectrum bands (i.e., band classes) in which the technology standard can be used, and as the number of band classes supported by a single device increases, the complexity of the device (and thus its cost) increases as well. For this analysis, the *x*-axis in Figure 2 extends to 60 MHz, which is about the amount of spectrum that is available in the three 700MHz band classes (i.e., LTE bands 13, 14, 17).[16] (The 700MHz band is of particular interest in the U.S, in part. because that is where public safety's broadband allocation is located.)

Furthermore, Figure 2 includes two scenarios. There is a base case scenario in which 4.8 Mbps is required in the downlink and 1.5 Mbps in the uplink to support response to a localized emergency. As discussed in Hallahan and Peha,[17] the base case scenario is based on a hypothetical emergency in which there is a chemical and biological terrorist attack in downtown Washington D.C. that causes a substantial number of casualties. This scenario, which was developed by the Spectrum Coalition for Public Safety,[18] was designed to illustrate a localized 'worst-case' capacity scenario in which thousands of first responders participate in an emergency response that is concentrated in a small area, served by a limited number of cell sites. To consider even larger and presumably less likely disasters, we include a second scenario that requires four times as much capacity as the base case.

Figure 2 (as well as Figure 3 discussed below) was generated using a model described in Hallahan and Peha,[19] which calculates capacity requirements and costs with a variety of scenarios and assumptions. The capacity in a given region is proportional to the density of cell sites. In Figure 2 and Figure 3, it is assumed that the density of cell sites in each commercial carrier's network is equal to the density of cell sites in the public safety network. This is a reasonable first-order approximation given that the number of cell sites in each major carrier's network is roughly equal to the number of cell sites recommended by the NBP for a nationwide public safety network (i.e., all four major networks have about 40,000 to 50,000 cell sites nationwide,[20] while the NBP calls for about 44,000 cell sites in a nationwide public safety network.[21] It is assumed that cell sites are co-located, so devices have no opportunity to choose the carrier with the closest tower, which can greatly increase spectral efficiency and total capacity beyond what is shown here. Also, it is assumed in this analysis that the cell sites in both the public safety and commercial networks are divided into sectors in the same way (e.g., three sectors/cell). Additionally, it is assumed that the commercial spectrum is divided equally between the uplink and the downlink and that the spectral efficiency is the same on both the commercial and public safety networks (i.e., 0.5 bps/Hz uplink; 1.5 bps/Hz downlink), which is reasonable given that both sets of infrastructure must support the same LTE technology to facilitate the roaming envisioned in this paper.

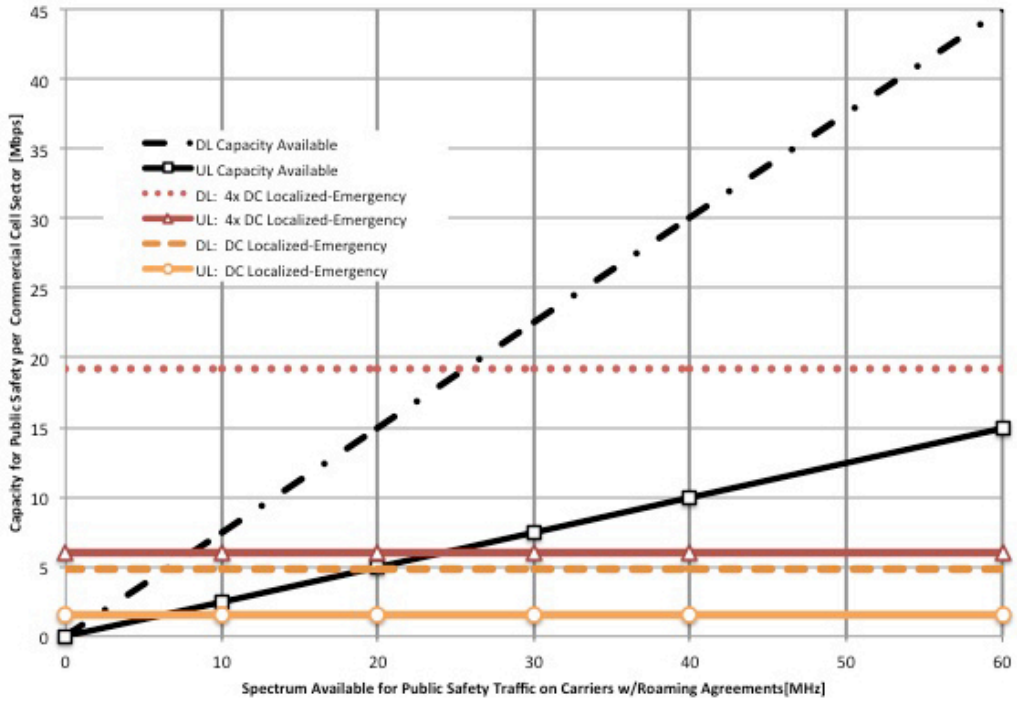

Figure 1. Roaming capacity available for public safety traffic on commercial networks that have roaming agreements with public safety

As shown in Figure 2 as the amount of spectrum on which public safety can roam increases, the magnitude of emergency that can be supported while roaming increases, such that roaming alone can provide sufficient capacity to support an extremely large emergency response. Even if public safety can roam onto just 10 MHz of commercial spectrum, there would be enough capacity to support the response to the terrorist attack in Washington DC studied here without using a dedicated public safety network. If, instead, public safety can roam onto 24 MHz of spectrum (e.g., all of band class 17), there would be enough roaming capacity to support an emergency four times as great as the base case DC-based scenario. Finally, if public safety can roam onto all of the 700MHz band classes (i.e., LTE bands 13, 14, 17), this 60 MHz of spectrum would provide over 60 Mbps of supplemental capacity per sector (45+ Mbps downlink; 15+ Mbps uplink), which is far beyond the capacity per sector expected on the dedicated public safety network envisioned in the National Broadband Plan.[22]

To assess the impact of public safety roaming on commercial subscribers, Figure 3 shows the utilization of a commercial cell sector as a function of the total amount of spectrum used by commercial carriers that have agreements with public safety. Note that this includes both spectrum that public safety can roam onto and spectrum they cannot roam onto, because when a commercial operator makes spectrum at 700MHz available for roaming, its commercial subscribers can be shifted to other spectrum bands outside of 700MHz as needed, even if public safety users can't. Assumptions for Figure 3 are the same as for Figure 2 The *x*-axis in Figure 3 extends to 320 MHz, which, as of 2010, was roughly the amount of spectrum licensed in the United States below 2.5 GHz that could be used for mobile broadband.[23]

Of the commercial spectrum available below 2.5GHz, the two largest nationwide

commercial carriers in the United States (i.e., the tier-1carriers) currently average roughly 100 MHz in the major markets and the next two carriers (i.e., the tier-2 carriers) average about 50 MHz in these markets.[24] Not all of the available spectrum has been put to use yet, and not all of the spectrum in use is used for mobile broadband. (Some is used for voice traffic.) To provide some context, Figure 3 includes labels on the *x*-axis showing the impact of agreements with tier-1 and tier-2 carriers, assuming tier-1 carriers have 60 MHz of spectrum available for mobile broadband and the tier-2 carriers have 30 MHz.[25]

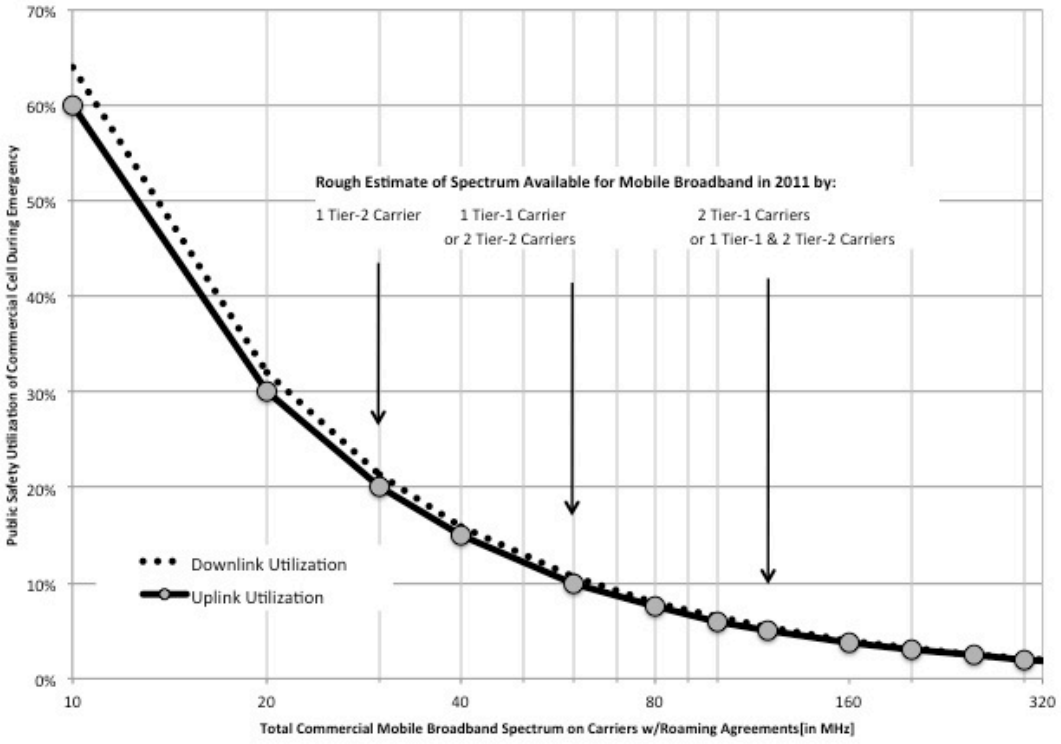

Figure 1. Average commercial network utilization due to public safety roaming during a localized emergency response

Figure 3 demonstrates that the impact of public safety roaming on commercial subscribers is small, and decreases further when the commercial networks on which public safety can roam have more total spectrum. For example, if public safety only has a roaming agreement with one tier-1 carrier, only about 10 percent of that carrier's capacity will be utilized by public safety during the response to the serious localized emergency scenario studied here. Thus, the quality of service observed by commercial subscribers will not be greatly affected. With two tier-1 carriers, utilization due to public safety falls to about 5 percent during the same emergency. Clearly, the more agreements public safety negotiates, the less commercial users will be affected. Moreover, from a given commercial carrier's perspective, the impact of public safety's roaming depends more on the number of carriers with which public safety forms roaming agreements than on how much of that carrier's spectrum public safety users can roam directly onto.

Table 1. Key Factors that may be used to differentiate resource requests on a wireless broadband network, the possible levels of granularity, and specific examples for each distinction

	Distinction	Granularity	Examples
Commercial	User Identity	Type of Subscription	Business vs. Standard
			High Price vs. Low Price
			Post-paid vs. Pre-paid
	Service Type	Real-time vs. Non-real-time app	Voice vs. Web browsing
		Type of real-time application	Voice vs. Video
		Importance of application	911 Voice vs. Other Voice
Public Safety	User Identity	Level of Government	Local vs. State vs. Federal
		Type of Agency	Fire vs. Police vs. EMS
		Rank of User	Officer vs. Chief
	Device Type	Mobility of device	Portable vs. Car-mounted vs. Fixed
		User-issued device type	Handheld vs. Laptop
		Fixed device type	Fixed sensor vs. Fixed Video Camera
	Service Type	Real-time vs. Non-real-time app	Voice vs. Web browsing
		Type of real-time application	Voice vs. Video
		Importance of application	Emergency Voice vs. Routine Voice
	Location of Usage	Within a Jurisdiction	Commercial Building vs. Residential Building vs. Highway
		Within Jurisdiction vs. Outside Jurisdiction	Local Responder vs. Neighboring Jurisdiction Responder
	Time of Usage	Time of Day	Day vs. Night; Busy Hour vs. Off-peak
	Situation of Usage	Type of Event	4-Alarm Fire vs. Hurricane vs. Terrorist Attack
	Network Used	Roaming on other Public Safety Network	Network for Region A vs. Network for Region B
		Roaming on Commercial Networks	Carrier A vs. Carrier B

PRIORITIZATION DECISIONS IN A BROADBAND NETWORK

This section identifies some of the main factors involved in differentiating resource requests in a broadband priority access system that meets public safety needs. Table 1 lists several likely distinctions that commercial and public safety broadband network operators may make to differentiate resource requests, along with some of the possible levels of granularity, and specific examples for each distinction. Some distinctions are static, and easy to accommodate. For example, voice traffic may always be given priority over web browsing, or a dispatcher may always be given priority over other users. Other distinctions are highly dynamic. For example, the same firefighter running the same applications may be given higher priority when responding to an emergency than when running errands, or higher priority when 100 feet from a known fire than when 50 miles from the closest known fire.

While Table 1 provides examples of possible factors used to distinguish resource requests, determining exactly what traffic should be prioritized over other traffic will require extensive input from public safety stakeholders working closely with technologists. The complexity of these decisions is often understated, and there is not yet one organization or group with both the knowledge and authority to address them.

A similar process is needed to determine which traffic remains on the dedicated public safety network, and which traffic should roam. This decision should take into account the fact that public safety and commercial networks typically differ in the types of applications they support, the outdoor signal reliability, the level of in-building coverage, cell-edge data-rates, and latency/call-setup times.[26] For example, given the better signal reliability and indoor coverage that we expect for public safety networks, it may be better to use these networks for critical voice communications, whereas video may be better supported in commercial networks due to their high data rates at edge of cell. (Or if cells are not co-located, unicast video streams might be carried on the system with the closest tower.)

LTE MECHANISMS USEFUL TO PUBLIC SAFETY

Some in the public safety community have expressed concern about whether they can count on priority mechanisms. This concern comes in part from experience with the Wireless Priority Service (WPS),[27] through which public safety users currently get priority access when making cellular telephone calls. Relatively speaking, WPS is a blunt instrument. As discussed in greater detail in Hallahan and Peha,[28] there are six levels of priority in the WPS system: the lowest is that of the general public, with five levels above that reserved for authorized emergency personnel. Higher priority calls are served first when calls are queued,[29] but high-priority users may still have to wait for lower-priority calls to complete when all voice channels are occupied (i.e., there's no preemption).[30] However, quality of service (QoS) is more complicated in a packet-switched network, and LTE provides a rich multifaceted set of mechanisms that can be used in a variety of ways to meet public safety's needs on a commercial network. This section will briefly introduce some of these mechanisms. For more details, see the appendix.

The coarsest form of priority in LTE is *Radio Admission Control,* which utilizes a mechanism called *Access Class Barring* to prevent entire classes of devices from connecting to one or more cells in the cellular network.[31] This allows a network operator to decide whether or not a given cell tower will connect to user classes such as commercial customers, police and other security services, public utilities, and firefighters and emergency medical services.

Once devices are allowed to connect, the network uses a variety of methods to allocate resources where they are most needed. First, to give one set of traffic better treatment than another, it must be possible to tell them apart. As the previous section shows, there are many different criteria that could matter when determining which traffic should get preferential treatment. LTE has mechanisms to support this. In LTE, it is possible to differentiate one user's traffic from another (e.g., a public safety user from a commercial user) and one application's traffic from

another (e.g., a video stream from a web browsing session) based on the *bearer channel* (as described in the appendix). A bearer channel is a form of virtual channel between endpoints. A single device may be exchanging traffic over multiple bearers at any given time. A *Traffic Flow Template* is used to sort packets into the appropriate channels in both the upstream and downstream directions using fields such as source, destination, and port number (which is usually an indicator of application type). Bearer channels are therefore useful instruments to separate traffic associated with applications that have different QoS requirements (e.g., real-time voice traffic vs. email) and traffic to or from users for which QoS during an emergency may be more or less important. All packets in the same bearer channel have similar quality of service objectives, and should receive similar treatment, while network algorithms can give packets from one bearer preferential treatment over packets from another bearer.

In LTE, preferential treatment during periods of congestion can take many forms. It is possible to block, drop (i.e., preempt), or reduce the QoS of the communications of one user or application that is considered less important. At the same time, it is possible to ensure more important packet streams are prioritized over others such that predetermined QoS characteristics are met, that established sessions are guaranteed a minimum bit rate, and that individual users do not use more than a preset amount of network resources.

Allocation and Retention Priority (ARP) mechanisms work at the granularity of bearers, as opposed to individual packets. Three ARP parameters are associated with every bearer: a scalar reflecting the importance of the stream, and two flags that indicate the preemption capability and vulnerability of the bearer. (See the appendix.) Through ARP, lower priority requests for resources can be blocked if enough resources are being used for higher priority sessions, and lower priority communication sessions can even be preempted by higher priority sessions. Session dropping and/or blocking can be used to give preference to one user over another, or for a given user, give preference to one application over another. For example, during a severe disaster, bearers with lower ARP priority level values and preemption vulnerability may simply be dropped to free up capacity for higher priority bearers. Alternatively, an operator could map the voice component and the video component of the same video telephony session to separate bearers with different ARP parameters.[32] During times of congestion, the video component could then be dropped without affecting the bearer carrying voice, allowing voice communications to continue during times of severe congestion.

Additional mechanisms operate on a packet-by-packet basis to protect the QoS of those bearers that are admitted/allowed to continue, and these may also provide preferential treatment to some traffic. For example, there are two different classes of bearer: *guaranteed bit rate* (GBR) and non-guaranteed bit rate (non-GBR). For GBR bearers, the network will ensure that sufficient resources are available to meet or exceed this rate, even if this means blocking new bearers (of lower ARP). Conversely, the network may impose a maximum data rate on a bearer, or collectively on all bearers from a given device, by establishing a *maximum bit rate* (MBR) for GBR bearers or an *aggregate maximum bit rate* (AMBR) for non-GBR bearers. This prevents a given user or a given application from generating too much traffic, thereby endangering the QoS of other data streams.

At even finer granularity, the *QoS Class Identifier* (QCI) mechanism ensures that the treatment a packet receives at each node in the network is tied to specific QoS objectives (e.g., packet delay budget and packet error loss rate) and is subject to varying levels of prioritization accordingly. Using QCI values, the desired QoS characteristics of each bearer are understood by the network without the need to specify individual QoS characteristics (e.g., a packet delay budget) for each bearer. For example, one QCI value indicates a QoS that would be well suited for interactive voice communications, and another QCI value would fit video streaming.

Finally, in LTE it is possible to receive preferential treatment even when roaming onto a new network. However, the question to be debated is whether the home network

or the visited network should decide how traffic to and from the roaming device will be handled. As discussed further in the appendix, LTE supports both possibilities. For example, in a *home-routed* configuration, all traffic from the roaming device is routed through that device's home network, and it is the *Policy and Charging Rules Function* (PCRF) module in the home network (i.e., H-PCRF) that makes the QoS policy decisions. On the other hand, with *local breakout*, there is no requirement to route traffic through the home network, and it is the visited network's PCRF (i.e., V-PCRF) that makes the QoS policy decisions.

Thus, it is technically possible under the LTE standard to provide several priority-related capabilities that are likely to be important to public safety users. These capabilities mean that during times of congestion, a LTE-based priority system could meet the QoS requirements of the individual users or applications that public safety deems most important, whether or not public safety is roaming on commercial networks or using their own dedicated systems. However, crafting the agreements and designing the rules that would govern such a system raises a number of challenges, as will be discussed in the next two sections.

TECHNICAL DESIGN DECISIONS: IMPLICATIONS FOR AGREEMENTS

Public safety use of commercial networks will require entirely new forms of agreements between public safety agencies and commercial carriers. Among other things, these agreements must reflect technical design issues.

Some people in policy circles believe that simply by specifying a standard, which in this case would presumably be LTE, all interoperability and quality of service issues associated with priority roaming will be settled. This is not the case. It is all too common that two systems that both comply with a given standard do not work together, or at least not at a quality of service that users expect or demand.

In the case of the LTE standard, not all network elements are required for an LTE deployment, and not all the features of LTE are supported by every network configuration. Agreements should be structured such that they ensure public safety's needs are met even when vendors and/or operators may make different decisions in their implementations of an LTE network, perhaps by making the agreements independent of those vendor/operator specific decisions (e.g., by specifying performance requirements and letting vendors and operators decide how to meet them) or by including vendor-specific breakout pieces in agreement guidelines.

As discussed in the previous section, two networks that both comply with the LTE standard can have different ways of handling roaming. For example, if public safety representatives conclude that they require direct control over the QoS that their users experience when roaming onto commercial networks, then the home network(s) operated by public safety must include a PCRF and the commercial network must support 'home-routed' traffic for roaming public safety users. Both of these are optional implementations within the LTE standard,[33] thus requiring coordination and agreement between commercial and public safety network operators to ensure the desired functionality is properly supported.

The previous section also shows the importance of the values assigned to QCI and ARP parameters. However, there is no guarantee the defined values will match the needs of all important public safety applications.[34] Worse yet, there is no guarantee the QoS experienced on one commercial network with a given set of QCI and ARP values will be the same as the QoS on another network. For example, one commercial carrier may give its best business customers ARP values even higher than public safety's, while another carrier doesn't, so public safety blocking probabilities in the first network would be worse even with the same ARP values. Alternatively, two different carriers may simply adopt a different *scheduling algorithm*, which is the algorithm that determines which queued packet should be transmitted next.[35] LTE standardizes the QoS parameters that serve as inputs to a scheduling algorithm, but the algorithms themselves are not standardized, and different vendors could easily take different approaches.[36]

Some degree of commonality across agreements based on meaningful guidelines will likely reduce the need for thousands of different agreements to be negotiated separately, while still allowing for divergence between agreements. One challenge with creating commonality is that there are many commercial carriers across the country and even more public safety agencies. (In the United States, there are more than 50,000 local, state, and federal agencies.[37]) Thus, identifying (or creating) a single entity to be responsible for soliciting input from all interested agencies and establishing appropriate guidelines for agreements will reduce the burden on individual agencies while ensuring a more consistent roaming experience for first responders in all localities. There is an inevitable tradeoff in such guidelines. The more specific the agreement, the more likely the agreement can ensure the needs of public safety will be met. However, overly specific agreements can have a stifling effect on innovation over time and diversity of product and service offerings across providers. Thus, the goal is to identify the minimum level of specificity that can meet public safety's needs, while enabling innovation and evolution to occur.

Finally, note that agreements must be constructed based on current technologies and standards, but the needs of both commercial and public safety users are likely to change in the years ahead. For example, the LTE standard currently recognizes 15 values of ARP priority and 9 QCI values.[38] It is possible that a new and important public safety application will emerge that is not well served by the currently defined values. If public safety's needs evolve in the same way as commercial needs, then standards bodies like 3GPP will probably meet public safety needs without any prompting. However, it is also possible that public safety's needs could diverge from those of commercial users. This raises two important issues for policy and governance. First, some entity with the ability to solicit input from public safety agencies and the authority to speak for those agencies should actively participate in the standards process. Second, it must be possible to update the guidelines on agreements between public safety and commercial carriers to reflect changes in the standard (including changes that may be optional). For example, imagine that after years of effort, 3GPP defines a new QCI value that perfectly fits an important public safety application. Some commercial carriers may have little incentive to upgrade their networks to support this new QCI. When such a network observes a QCI value it doesn't understand, it may simply interpret this as a QCI value that it does understand,[39] potentially giving public safety users an inconsistent QoS experience as they roam. (As discussed by Hallahan and Peha, recognizing a new QCI may require nothing more than software updates of the affected infrastructure, whereas other new features could conceivably require more disruptive or costly changes.[40])

OPERATIONAL DESIGN DECISION: ENABLING HUMAN INTERVENTION

Perhaps the most important decision that must be made about the operation of a national wireless broadband network for public safety concerns the role of human operators to allocate resources and set priorities. Whatever the decision, it will significantly affect network technology, priority roaming agreements, and governance. Yet, this issue has often been missing from the debate. Some in the policy realm may see it as a technical issue, but on an issue like this there is no way to separate design of technology and design of organizations and how they operate, which demands a sociotechnical perspective.[41] As Bostrom and Heinen said, "a work system is made up of two jointly independent but correlative interacting systems – the social and the technical."[42] This section will discuss the pros and cons of a fully automated system versus a system in which human operators make real-time decisions about priority and resource allocation, and it will show the capabilities of LTE technology.

In an automated priority system, the priority parameters (such as the QCI and ARP values in LTE) are assigned according to policies and decision rules that are based only on factors that the system either (1) knows *a priori* or (2) can detect without any human intervention. These include static factors (e.g., user identity, as stored in a subscriber profile) and dynamic but detectable factors

(e.g., device location as determined through GPS). These factors are then used in a predefined decision rule (e.g., if roaming, then ARP level = 1; else ARP level = 2).

A purely automated priority system can prioritize resource requests appropriately in a great many circumstances, but not all. Factors like the intention of a first responder or their perceived level of danger cannot be detected by a network and therefore cannot be used for resource allocation. For example, an automated system could base priority decisions on the fact that a police officer is making a voice call from his patrol vehicle on a highway outside of his normal jurisdiction; it just cannot tell whether he is pursuing a known fugitive or driving back to the police station. To handle these cases, an intervention-enabled system could be designed such that decisions depend, at least in part, on factors that require human intervention. In the example above, the police officer or a dispatcher could explicitly indicate that this particular session is of elevated importance.

While allowing human intervention may yield better resource allocation in some situations, human intervention also brings some complications for policy and governance that must be addressed. It requires that humans be available with both the authority and expertise to make decisions that affect the network, and do so within appropriate time constraints. Moreover, there can be first responders from many different public agencies simultaneously responding to emergencies, each with its own incident commander, and its own urgent needs. Even without roaming, it must be determined who is allowed to intervene in ways that affect resource allocation, and how.

With the addition of priority roaming, solutions must be found that are effective for all first responders that are roaming onto commercial networks, while also treating commercial traffic appropriately. Moreover, someone must define procedures by which those with situational knowledge (whether they are incident commanders, dispatchers, or individual first responders) communicate that knowledge in real time back to the networks, and ideally those methods would be common across the many public safety agencies and commercial networks. It is also possible that the additional technical and operational complexity will also affect cost.

To assess the trade-offs, decision makers should understand the needs of the public safety community, the functionality provided by automated priority systems, and the additional functionality provided by intervention-enabled systems and then balance these against the additional complexity and challenges human intervention presents before deciding on which method to employ. This section will present a few potential operational arrangements for enabling human intervention as well as a few examples of possible technical implementations.

The good news for policymakers is that LTE makes possible a wide range of options for enabling human intervention in priority decisions, so technology need not be an impediment. To demonstrate this, three possible options will be presented as examples, each of which has its advantages and disadvantages. One option is to give a centralized public safety entity full control to make QoS decisions for public safety users even while those users are roaming (which means these decisions would affect commercial users as well). Another option is to allow commercial operators to hold final control over the QoS decisions that affect their network, while public safety only provides input to the commercial operator. A third option is to leave the decision up to individual public safety users, and allow them to affect the QoS they receive in response to their current situation. To demonstrate technical feasibility, the following are some examples of how to implement these options; other approaches are also possible. (See the appendix for more information on the technology.)

CENTRALIZED CONTROL BY PUBLIC SAFETY

The first option could be supported by implementing the (optional) PCRF function in the public safety network and "home-routing" traffic back through that network, even when public safety users are roaming on commercial networks. The QoS level these users receive would then be controlled by the

PCRF in public safety's home network (H-PCRF). Therefore, if a public safety representative maintains control over the H-PCRF (or multiple H-PCRF's if there are multiple regional public safety networks), then public safety could potentially intervene and update the policies to reflect current situations even for users who are roaming on commercial networks, with no action required by the commercial operator. This level of control is most like what public safety agencies with typical LMR systems are accustomed to. It has the advantage of giving a public safety representative maximal control over resources, although it is not clear who that representative should be, especially in scenarios when multiple public safety agencies simultaneously respond to serious emergencies. However, commercial networks may be reluctant to allow public safety to make decisions that affect the QoS observed by commercial subscribers. As shown previously, their concerns should be lessened tremendously if public safety makes agreements with multiple commercial providers in every region and spreads the roaming load across multiple providers. This may or may not occur in practice.

COMMERCIAL OPERATOR CONTROL

Alternatively, the commercial operator could maintain control over QoS when public safety users are roaming by employing 'local breakout,' and placing control for QoS in the PCRF of the visited network (V-PCRF). The operator may receive public safety input in a variety of ways, but the operator would have responsibility for the network element that controls priority. This arrangement is consistent with how commercial operators tend to view their customers, none of whom would normally be allowed to directly control QoS. It better protects commercial customers, but it gives public safety less ability to put communications resources where they are most needed.

INDIVIDUAL CONTROL

Finally, LTE also supports the use of *terminal-initiated* QoS control. As discussed further in the appendix, with this approach, a terminal can signal the network and request that a dedicated bearer with the desired level of QoS be established.[43] For example, a first responder might press an emergency button on the handset, which would cause the handset to request a different level of QoS as reflected in ARP and QCI parameters. Thus, control is passed on to the first responders in the field, who can make decisions based on their instantaneous situational needs. These first responders are best able to assess their own needs, but increasing their own priority has the effect of reducing the resources available to others, and they cannot easily assess the needs of others in their cell.

CONCLUSIONS

In many ways, firefighters, police, and paramedics are at the front line of homeland security emergencies, and they have long been forced to make due with communications systems that are unnecessarily prone to failure, limited in functionality, and overpriced. Wireless broadband functionality could revolutionize the way public safety responds to emergencies by bringing capabilities to first responders they have never had before. Also providing users of these new public safety systems with roaming access to commercial networks, on a priority basis, can provide far greater aggregate capacity, geographic coverage, and service reliability than would be available from dedicated public safety networks alone. Policies and arrangements should be adopted that take advantage of these benefits.

In the United States, which is in the process of creating the first nationwide public safety network under the auspices of FirstNet, priority roaming should be one of the core elements of an initial roll-out strategy. FirstNet must show results nationwide very quickly and on a limited budget, and they can begin offering services over commercial networks right away even in areas where they have will ultimately offer services over their own infrastructure.[44]

While some may be looking for a single roaming partner per region, the more agreements public safety negotiates with commercial carriers, the more capacity they

will have available during emergencies and the less commercial users will be affected by public safety's traffic. If public safety is able to roam onto just 10 MHz of commercial spectrum, they would have enough roaming capacity (even without the capacity provided by the dedicated public safety infrastructure) to support a hypothetical emergency response to a major chemical and biological terrorist attack on Washington D.C. If public safety is able to roam onto all of the 700MHz band classes (i.e., LTE bands 13, 14, 17), they would have over 60 Mbps of roaming capacity per cell sector (45+ Mbps downlink; 15+ Mbps uplink): enough to support an emergency ten times greater than the DC-based scenario studied. Moreover, concerns about the impact on consumers are vastly overstated. If public safety only has a roaming agreement with one tier-1 commercial carrier, only about 10 percent of that carrier's capacity will be utilized by public safety users during the disaster scenario studied, which is unlikely to cause significant congestion for commercial subscribers. Clearly, more agreements with commercial carriers would mean even less impact on average for each of the commercial subscribers affected.

To make this kind of priority access possible, LTE provides a wide range of priority-related capabilities that are likely to be important to public safety. Thus, even during periods of congestion, an LTE-based system with well-crafted rules could meet the QoS requirements that public safety deems most important, both for users on dedicated public safety networks and those that are roaming on commercial networks. However, challenging issues must be addressed when crafting service level agreements for roaming between public safety and commercial carriers. These include determining which optional elements of the LTE standard should be adopted and how, and defining a common understanding of the quality of service that public safety can expect despite what could be significant divergence in important technical design decisions that are outside the scope of the standard.

To meet these challenges, there should be a single entity responsible for understanding the needs of the many public safety agencies, and creating a consistent set of guidelines for how agreements can be constructed, thereby providing increased commonality while still meeting regional needs by allowing some degree of divergence. In the United States, this might be the National Institute of Standards and Technology (NIST), which was given funding and responsibility to "conduct research and assist with the development of standards, technologies and applications to advance wireless public safety communications" in the Middle Class Tax Relief and Job Creation Act of 2012,[45] but other agencies could also try to play this role. The goal should be to identify the minimum level of specificity that can meet public safety needs, while enabling innovation and evolution. Moreover, to accommodate the inevitable changes in technology, and changes in public safety needs, this same entity should play an active role in relevant standards bodies such as 3GPP to represent the needs of public safety. This entity should have expertise in the technical issues, an understanding of public safety needs, the responsibility to continually solicit feedback from the public safety community, and the authority to act on their behalf. In the United States, there are a number of candidates to perform this function, including NIST, FirstNet, and the Department of Homeland Security; one should be chosen.

Another important issue is the extent to which priorities and resource allocation should be controlled through human intervention. In addition to changing the shape of roaming agreements with commercial carriers, this will significantly affect the technology, policy, and governance structures that public safety needs. We have shown that LTE supports a wide range of arrangements, including placing high levels of control on resource allocation in the hands of a central public safety authority, each commercial operator, individual emergency responders, or some combination thereof. There are complex non-technical factors to consider.

There is clearly more work to do. While this article has shown that much of what public safety needs is present in the standard, this only helps if devices are produced using these capabilities. For public safety to use priority roaming, they need mobile devices that operate in both the public safety

spectrum band and (at least one of) the spectrum bands licensed to commercial carriers. In addition, any of the possible features of LTE described in this article that are found to be essential to public safety must actually be implemented in the devices in question. Further technical analysis may be needed to determine precisely how to do design devices of this kind at minimum cost. Moreover, either public policies or business agreements will be needed that make sure the devices are produced, and at sufficient scale to drive down costs. There is not yet agreement on how best to achieve this.

The most important issues – beyond the scope of this article – may be more organizational than technical. In a large-scale emergency, it is not unusual for first responders from many public safety agencies to respond, and with priority roaming multiple commercial carriers may also be involved. Nevertheless, there will be times when limited resources must be allocated to the most important needs. The existence of technology that allows for intelligent prioritization only increases the need for governance structures that allow for effective and rapid decision-making even when it is not obvious which person or organization is in charge.

ABOUT THE AUTHORS

Ryan Hallahan *recently received his PhD in Engineering & Public Policy from Carnegie Mellon University where he was a Bertucci Graduate Fellow. His dissertation focused on methods of improving public safety wireless communication systems in the United States, including analyzing the cost of a nationwide wireless broadband network, and the advantages of and challenges to leveraging commercial infrastructure and spectrum to supplement dedicated public safety networks. His papers have been presented at the Telecommunications Policy Research Conference (TPRC) and published in Telecommunications Policy. He has a Master of Science in Electrical & Computer Engineering from Carnegie Mellon University and a Bachelor of Science in Engineering Physics from the University of California, Berkeley.*

Jon M. Peha *is a full professor at Carnegie Mellon University in the Department of Engineering & Public Policy and the Department of Electrical & Computer Engineering, and served as associate director of the university's Center for Wireless and Broadband Networking. From 2008 to 2011, Dr. Peha served in the US government, first as chief technologist of the Federal Communications Commission, and then as assistant director of the White House Office of Science & Technology Policy where he focused on communications policy and research policy. He has previously been chief technical officer of three high-tech start-ups, and a member of technical staff at SRI International, AT&T Bell Laboratories, and Microsoft. He has addressed telecom and e-commerce issues on legislative staff in the US House and Senate, and helped launch and lead a US government interagency program to assist developing countries with information infrastructure. Dr. Peha consults for industry and government agencies around the world. He is an IEEE fellow and an AAAS fellow. He holds a PhD in Electrical Engineering from Stanford.*

ACKNOWLEDGEMENTS:

The authors gratefully acknowledge the financial support of the John and Claire Bertucci Fellowship.

APPENDIX: Overview of LTE Concepts AND Mechanisms

This section serves as an introduction to the relevant LTE concepts and mechanisms that are discussed in the main body of this paper. This section will first present an overview of the LTE standard in section, then discuss the fundamental QoS concepts in section, and finally discuss the available QoS policy control mechanisms in section.

OVERVIEW OF LTE

LTE, or Long Term Evolution, refers to the Release 8 iteration of the 3rd Generation Partnership Project's (3GPP) mobile network technology.[46] There are two main components of a LTE network: the radio access network (E-UTRAN) and the packet core network (EPC). Both of these components were designed to ensure that LTE is a packet-switched, all-IP standard in contrast to previous voice-centric, circuit-switched architectures.[47] The E-UTRAN has two main elements: the User Equipment (UE) and the E-UTRAN base station (eNodeB or eNB). The UE is a generic term for the handsets and other devices that subscribers use to communicate with the eNodeB's over the network's allocated spectrum. The eNodeB handles all radio access related functions and each eNodeB communicates with the packet core. Service providers can have their own separate core networks, but share eNodeB's, since each eNodeB can be connected to multiple cores. Each packet core or EPC will typically include the following network elements: a Serving Gateway (S-GW), a PDN Gateway (PDN-GW), a Policy and Charging Rules Function (PCRF), and a Mobile Management Entity (MME), each of which is discussed in greater detail in (Johnson 2010).[48] Figure 4 is a generic representation of an LTE network architecture, which shows the general relationship between the main elements of an LTE network, based on diagrams in.[49]

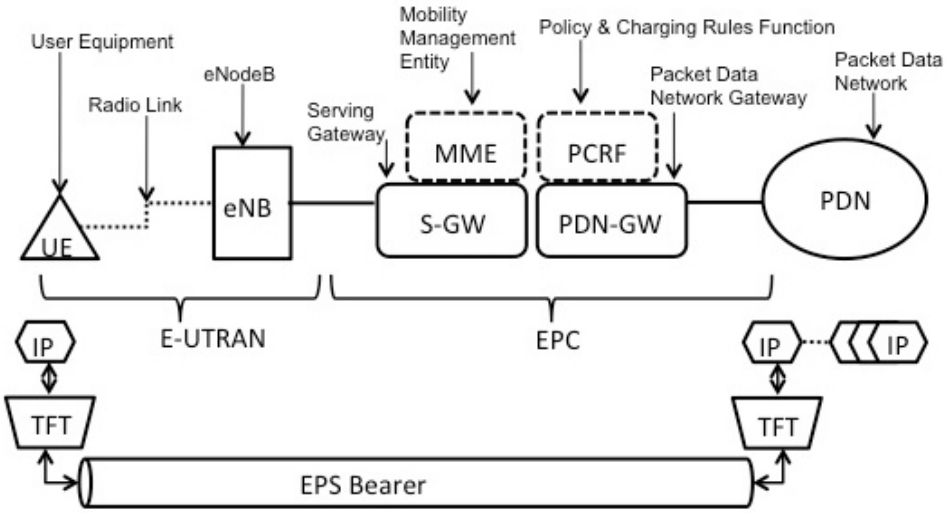

Figure 1. The main elements of an LTE network and their general relationship to each other

FUNDAMENTAL QOS CONCEPTS IN LTE

The LTE standard provides a robust set of technical mechanisms that enable different users and services to receive preferential treatment (both in terms of priority access and preemption capability) across the E-UTRAN and EPC. The fundamental concepts that are central to priority access are discussed in this section. The EPS bearer is a logical channel for data flows that require the same QoS treatment and is discussed in the Overview section; the parameters that differentiate bearers based on their QoS requirements are discussed in the subsequent section.

OVERVIEW OF EPS BEARERS: DEFAULT VS. DEDICATED & GBR VS. NON-GBR

At a high level, a bearer is the term used to describe a 'virtual' channel established between the endpoints of the network (i.e., from the UE to the PDN-GW). A bearer is 'virtual' in the sense that all traffic from the user device is carried across the same physical channel (the radio channel) back to the network, but many virtual channels can be created to distinguish between how different traffic should be treated over the same physical channel. There are two types of EPS bearer in LTE: default and dedicated. Every UE has at least one default bearer that is established when the UE first attaches to the network and remains available for the duration of the connection. A UE can have anywhere from zero to several dedicated bearers established at any given time and each is set up and taken down on an as-needed basis.

Dedicated bearers are used when the QoS requirements for some traffic is different than the QoS provisions provided by the default bearer. Furthermore, all traffic requiring the same QoS-level treatment will be carried on the same bearer (e.g., if a device is making a voice call and streaming video at the same time, and both applications require the same level of QoS, the traffic from both will be mapped to the same bearer). In LTE, packets are filtered into the appropriate bearer using a Traffic Flow Template (TFT). At a high level, the TFT is just a list of source/destination IP addresses and TCP/UDP port combinations that identify which IP packets (based on their header information) should be assigned to which bearer.[50]

Thus, the bearer forms the fundamental unit for discussing the QoS mechanisms available in an LTE network. Furthermore, there are two possible types of dedicated bearer: guaranteed bit rate (GBR) bearers and non-guaranteed bit rate (non-GBR) bearers. For a GBR bearer, the system guarantees a minimum bit rate will be provided to that bearer once it is established. This means that GBR bearers sending at a bit rate less than or equal to their GBR can assume that packet drops as a result of congestion will not occur. For a GBR bearer, a maximum bit rate (MBR) may also be specified which caps the maximum bit rate that bearer will receive.[51]

The network guarantees no minimum bit rate for non-GBR bearers.Thus, there are no guarantees as to the amount of traffic a non-GBR bearer can support at any given time, which could potentially result in packet loss during times of congestion. In addition, non-GBR bearers for the same device may be capped in the aggregate bandwidth they receive by using the aggregate maximum bit rate (AMBR) parameter. The AMBR can be specified at either the APN level (APN-AMBR) or the UE level (UE-AMBR).[52] For example, the UE-AMBR can be used to cap the aggregate bit rate that is allocated to all non-GBR bearers used by a given UE.

The decision of which type of bearer should be used (GBR vs. non-GBR) depends upon the service that is carried by that bearer. As discussed by Olsson et al.,[53] GBR bearers are typically used for services where it is better to block them initially rather than degrade the service after it has already started. For example, it may be desirable to block a real-time voice call before it begins during times of congestion, rather than admit the service and then have the voice be unintelligible since the guaranteed bit rate cannot be maintained. (Note that many real-time services can actually adapt to the available bit rate to some degree, but there is still a minimum bit rate below which they cannot operate properly.) On the other hand, non-GBR bearers are typically used for applications such as web browsing and email, as these applications do not require a guaranteed bit rate. However, simply because an application does not require a GBR does not make it less important than applications that require a GBR; the relative importance of applications can depend on a number of additional factors. As discussed by Olsson et al.,[54] the

choice of which bearer to use for each service is up to the operator and their configuration. Table 2 summarizes the types of bearers available and the bit rate and QoS treatment parameters (discussed in the next section) available to each.[55]

Table 1. The bit rate and QoS treatment parameters available to each of type bearer

		Type of Bearer	
Bit Rate Parameter		GBR	Non-GBR
GBR:	Guaranteed Bit Rate	**X**	
MBR:	Maximum Bit Rate	**X**	
APN-AMBR:	APN Aggregate Maximum Bit Rate		**X**
UE-AMBR:	UE Aggregate Maximum Bit Rate		**X**
QoS Parameter			
QCI:	Quality Class Identifier	**X**	**X**
ARP:	Allocation and Retention Priority	**X**	**X**

OVERVIEW OF BEARER-LEVEL QOS PARAMETERS

To support QoS requirements, the EPS bearer includes several parameters that dictate the preferential treatment a bearer may receive. Each bearer, including both GBR and non-GBR bearers, is associated with the following bearer level QoS parameters: the QoS Class Identifier (QCI) and the Allocation and Retention Priority (ARP). The QCI parameter dictates the packet-level preferential treatment a bearer receives, while the ARP parameter dictates the preferential treatment individual bearer receives when they are being established. These parameters may be specified independently of the other, allowing for many different QCI+ARP combinations for each bearer.

When bearers are being established (or modified) on the network and resources are limited, the network may need to make decisions regarding which bearer requests should be accepted and which should be rejected (this usually occurs when available radio capacity is limited and typically involves GBR bearers). The primary role of the ARP parameter is to facilitate this decision making process.[56] To do so, the ARP parameter contains three components: a single scalar value and two separate flag values. The scalar value contains information about the priority level of a bearer, while the two flags refer to the preemption capability and preemption vulnerability of the bearer.

The ARP priority level is used to ensure that the request of the bearer with the higher priority level is given preference over lower priority bearers. During periods where resources are limited, the network may choose to drop bearers of low priority to free up required resources. The pre-emption capability flag defines whether or not a given bearer is allowed to preempt (i.e., force the system to drop) other bearers of lower priority level. On the other hand, the preemption vulnerability flag defines whether or not a given bearer is susceptible to preemption (i.e., being dropped) even by bearers with a higher ARP priority level.

Once bearers are established using the access control mechanisms provided by the ARP parameter, the nodes in the network still need to know how to treat the packets for each bearer. During times of congestion, bearers (who have been established) will compete for limited resources. This means that at individual nodes (e.g., eNodeB), the limited resources need to be

allocated to individual packets from many different bearers. The QCI parameter tells the nodes how to prioritize those resources among the packets (the ARP value has no effect on this).

The QCI parameter is specified by a simple scalar value. There is one-to-one mapping of standardized QCI values to standardized QoS characteristics. Table 3 summarizes the QCI that have already been standardized including: their priority level, packet delay budget, packet error loss rate, and examples of services that will typically be mapped to each QCI.[57] Thus, the QCI parameter is used by the eNodeB to determine the packet forwarding treatment of each bearer (e.g., scheduling weights and queue management thresholds). This treatment is pre-configured by the operator owning the access node (e.g., eNodeB), such that the QoS requirements associated with a given QCI are met.[58]

Table 2. The Standardized QCI Values and their Standardized QoS Characteristics

Resource Type	QCI	Priority	Packet Delay Budget	Packet Error Loss Rate	Example Services
GBR	1	2	100 ms	10^-2	Conversational Voice
	2	4	150 ms	10^-3	Conversational Video (Live Streaming)
	3	3	50 ms	10^-3	Real Time Gaming
	4	5	300 ms	10^-6	Non-Conversational Video (Buffered Streaming)
Non-GBR	5	1	100 ms	10^-6	IMS Signalling
	6	6	300 ms	10^-6	Video (Buffered Streaming) TCP-based (e.g., www, e-mail, chat, ftp, p2p file sharing, etc.)
	7	7	100 ms	10^-3	Voice, Video (Live Streaming), Interactive Gaming
	8	8	300 ms	10^-6	Video (Buffered Streaming) TCP-based (e.g., www, e-mail, chat, ftp, p2p file sharing, etc.)
	9	9	300 ms	10^-6	QCI typically used for the default bearer of a UE/PDN

OVERVIEW OF POLICY CONTROL AND ROAMING

The LTE standard also provides several features for controlling and initiating the QoS mechanisms discussed previously. This section first discusses the policy and charging control framework in LTE and the network elements required to support this framework, then compares network-initiated QoS control to terminal-initiated QoS control, and finally, discusses how QoS policies can be controlled when users roam on to other networks.

OVERVIEW OF POLICIES AND CHARGING CONTROL (PCC)

Policy and Charging Control (PCC) is the concept in LTE that enables flow-based policy control (e.g., QoS management) and charging control.[59] The main component of this concept is the Policy and Charging Rules Function (PCRF), which is an optional element in the LTE architecture that is responsible for providing policy control decision and charging control functionalities that are enforced by the Policy and Charging Enforcement Function (PCEF). (Where a policy is just a set of rules that determines how a specific IP flow is treated and the QoS it receives.) The Application Function (AF) interacts with application level signaling and extracts session information that it provides to the PCRF, while the Subscription Profile Repository (SPR) contains subscription and policy information for individual users. Figure 5, based on diagrams in Olsson, et al.,[60] illustrates the relationship between these network elements.

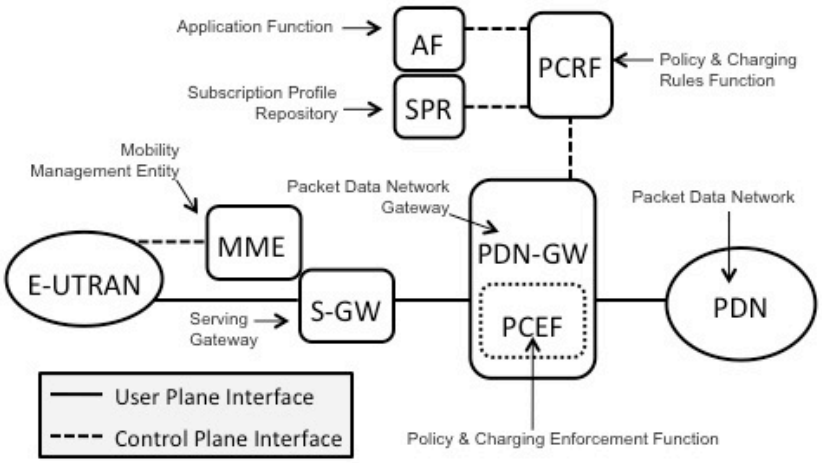

Figure 1. The main elements of the LTE PCC framework and their general relationship to each other

QOS CONTROL – NETWORK-INITIATED VS. TERMINAL-INITIATED

As discussed in greater detail by Olsson et al.,[61] there are two different methods available to establish a dedicated bearer for a given level of QoS in a LTE network: network-initiated QoS control and terminal-initiated QoS control. In network-initiated QoS control, the network signals the UE to establish a dedicated bearer with a given level of QoS. Ultimately, it is the PCRF that makes this decision, although it may consult the AF and/or SPR in the process. The exact details of this process depend on a number of factors and are not central to this article. The key is that with network-initiated QoS control, it is the responsibility of the network to

detect and infer what QoS resources are needed by the user or application, without explicitly being told.

In terminal-initiated QoS control, it is the terminal that signals the network and requests that a dedicated bearer with the desired level of QoS be established.[62] This means that the terminal must be aware of the specifics of how QoS is handled in the access network, which is not the case with network-initiated QoS control (where terminals can be access QoS-agnostic). Additionally, in a terminal-initiated QoS scheme, the terminal must be able to interface with the network to convey the QoS request (e.g., using an Application Programming Interface [API]).[63] However, terminal-initiated QoS control means that a PCRF is not needed to send QoS information to the network (although a PCRF can still be used, if desired, to authorize QoS requests made by terminals).[64]

ROAMING: HOME-ROUTED VS. LOCAL-BREAKOUT

There are two main roaming scenarios that are supported by the PCC framework in LTE: "home-routed" and "local-breakout." In home-routed roaming, the user in the visited network (i.e., the user who is roaming) is connected to the PDN through a PDN-GW that resides in the home network. Thus, all traffic for that user is routed from the visited network (where the roaming user is connected to the visited E-UTRAN) back through the home network before it exits to external packet networks (e.g., the internet). In local-breakout roaming, the user in the visited network is connected to the PDN through a PDN-GW in the visited network. Thus, all traffic for that user is routed through the visited network only and never enters the home network.

The PCC architecture was designed to enable the PCRF in the home network (H-PCRF) to communicate with the PCRF in the visited network (V-PCRF) and, when allowed by the visited network, control and authorize all resources for roaming users in the visited network.[65] The exact QoS control the H-PCRF has over its roaming users depends upon which PCC network elements are connected and how they are configured on both networks. In some cases, the V-PCRF may be allowed to either accept or reject (but not change) policy decisions made by the H-PCRF thereby allowing the visited operator some degree of control over the resource usage in its radio access network (i.e., E-UTRAN).[66]

---

[1] Jon M.Peha, "The Need for Fundamental Reform in Public Safety Spectrum and Communications Policy," *Wireless Future Program: Working Paper #15* (October 2006), http://www.ece.cmu.edu/~peha/safety.html; Peha, "How America's Fragmented Approach to Public Safety Wastes Spectrum and Funding," *Paper presented at the 33rd Telecommunications Policy Research* Conference (2005), http://www.ece.cmu.edu/~peha/safety.html; Peha, "Fundamental Reform in Public Safety Communications Policy," *Federal Communications Law Journal* 59, no. 3 (June 2007): 517-546, http://www.ece.cmu.edu/~peha/safety.html; Peha, "How America's Fragmented Approach to Public Safety Wastes Money and Spectrum," *Telecommunications Policy* 31, no. 10-11 (November 2007): 605-618.

[2] Ryan Hallahan and Jon M. Peha, "Quantifying the Costs of a Nationwide Broadband Public Safety Wireless Network," *Paper presented at the 36th Telecommunications Policy Research Conference* (2008), http://www.ece.cmu.edu/~peha/safety.html; Hallahan and Peha, "The Business Case of a Nationwide Wireless Network that Serves both Public Safety and Commercial Subscribers" *Paper presented at the 37th Telecommunications Policy Research Conference* (2009), http://www.ece.cmu.edu/~peha/safety.html; Hallahan and Peha, "Quantifying the Costs of a Nationwide Public Safety Wireless Network," *Telecommunications Policy* (Elsevier) 34, no. 4 (May 2010): 200-220; Jon M. Peha, "A 'Successful' Policy for Public Safety Communications, Comment before the Federal Communications Commission in the matter of implementing a broadband interoperable public safety network in the 700 MHz band," *PS Docket No. 06-229* (May 26, 2008), http://www.ece.cmu.edu/~peha/safety.html.

[3] 3G Americas, "Global 3G status UMTS / UMTS-HSPA/ HSPA+ / LTE" (July 2010), http://www.3gamericas.org/documents/Global%20Status%20Update%20July%2027%202010.pdf.

[4] Recently in the United States, Congress passed the Middle Class Tax Relief and Job Creation Act of 2012, which created the First Responder Network Authority (FirstNet). FirstNet was granted the license to 20 MHz of 700MHz spectrum and was tasked with establishing a nationwide, interoperable public safety broadband network. 112th Congress, H.R. 3630: Middle Class Tax Relief and Job Creation Act of 2012 (December 9, 2011), http://www.govtrack.us/congress/bills/112/hr3630.

[5] Federal Communications Commission (FCC), *Connecting America: The National Broadband Plan* (Washington, DC: March 16, 2010) http://www.broadband.gov/plan/; FCC, "Final Report of the Technical Advisory Board for First Responder Interoperability: Recommended Minimum Technical Requirements to Ensure Nationwide Interoperability for the Nationwide Public Safety Broadband Network" (May 22, 2012), http://transition.fcc.gov/Daily_Releases/Daily_Business/2012/db0621/FCC-12-68A3.pdf.

[6] National Commission on Terrorist Attacks Upon the United States, *The 9/11 CommissionRreport: Final Report of the National Commission on Terrorist Attacks Upon the United States* (2004), http://www.gpoaccess.gov/911/pdf/fullreport.pdf.

[7] FCC, *Connecting America*.

[8] Motorola Solutions and Verizon Wireless Alliance, "Motorola Solutions and Verizon Wireless Alliance Delivers Industry Leading Public Safety LTE Solutions" (2011), http://www.motorola.com/web/Business/Solutions/Business%20Solutions/Mission%20Critical%20Communications/LTE_for_Government_and_Public_Safety/_Documents/_Static_files/Motorola_Verizon_Alliance.pdf,

[9] Linda K. Moore, *Establishment of the First Responder Network Authority: Business Models and Congressional Oversight,* CRS Memo (Washington, DC: Congressional Research Service, 2012); 112th Congress, H.R. 3630: Middle Class Tax Relief and Job Creation Act of 2012.

[10] Jon M. Peha, "A Public Private Approach to Public Safety Communications," *Issues in Science and Technology*, National Academy Press (July, 2013), http://www.ece.cmu.edu/~peha/safety.html.

[11] Compared to commercial systems, it is extremely difficult to determine what the worst-case load will be on a public safety system in any given year, but, given the mission of public safety agencies, meeting this worst-case load is very important. At the same time, the peak to average ratio of load on a public safety system is much higher than on a commercial system, meaning that if public safety designs its network to meet the absolute worst-case load scenario, much of the capacity on their network will lay idle most of the time. Doing so can substantially increase the cost of the public safety network, but makes inefficient use of that infrastructure (Hallahan and Peha, "Quantifying the Costs of a Nationwide Broadband Public Safety Wireless Network" and "The Business Case of a Network that Serves both Public Safety and Commercial Subscribers"). On the other hand, if public safety can make use of commercial networks in addition to their own in times of extreme emergency, it is more likely they will be able to support the increased load in these instances.

[12] In those rare instances when all infrastructure in a region is down, including systems owned by public safety and all commercial carriers, emergency responders may turn to satellite services or deploy portable infrastructure such as the cell on wheels (COW) or cell on light truck (COLT). The US National Broadband Plan (FCC, *Connecting America*) recommended storing equipment for this purpose.

[13] NPSTC, "Priority and QoS in the Nationwide Public Safety Broadband Network v1.0" (April 20, 2012), http://www.npstc.org/download.jsp?tableId=37&column=217&id=2304&file=PriorityAndQoSDefinition_v1_0_clean.pdf; PSCR, "QoS Information" (December 1, 2010), http://www.pscr.gov/projects/broadband/700mhz_demo_net/stakeholder_mtg_122010/day_1/5.1_qos_priority_preemption-pscr_intro.pdf; Roberson and Associates. "Public Safety Priority Access to Shared Commercial Networks." March 2, 2011. http://fjallfoss.fcc.gov/ecfs/comment/view?id=6016171906.

[14] FCC, "Final Report of the Technical Advisory Board for First Responder Interoperability."

[15] Hallhan and Peha, "The Business Case of a Nationwide Wireless Network." The next generation of LTE technology, LTE Advanced, will offer providers the ability to combine noncontiguous spectrum bands, which could allow a single network that had both public safety and commercial spectrum to provide even higher data rates. See Jeanette Wannstrom, "Carrier Aggregation Explained" (3GPP, May 2012), http://www.3gpp.org/Carrier-Aggregation-explained.

[16] (Johnson 2010) Chris Johnson, *Long Term Evolution IN BULLETS* (Northhampton, UK: CreateSpace, 2010).

[17] Ryan Hallahan and Jon M. Peha, "Compensating Commercial Carriers for Public Safety Use: Pricing Options and the Financial Benefits and Risks," Paper presented at the 39th annual Telecommunications Policy Research Conference (2011), http://www.ece.cmu.edu/~peha/safety.html.

[18] Spectrum Coalition for Public Safety, "Public Safety Spectrum: How Much do We Need for Data?" (2005), http://www.spectrumcoalition.dc.gov/img/PS_Whitepaper_10-25-05.pdf.

[19] Hallahan and Peha, "Compensating Commercial Carriers for Public Safety Use."

[20] Morgan Stanley, *The Mobile Internet Report,* December 15, 2009, http://www.morganstanley.com/institutional/techresearch/pdfs/mobile_internet_report.pdf.

[21] See FCC, *Connecting America*. While the total number of cell sites is similar, this is still a very rough approximation. In particular, the total area covered as well as the distribution of cell sites throughout rural and urban areas isn't necessarily the same on commercial and public safety networks. Exact estimates would require detailed data on the exact locations of all cell sites in the commercial networks, and the design of a public safety network that has to be built. Nevertheless, this approximation can still give useful insight as to the order of magnitude of capacity available and utilization during an emergency scenario.

[22] FCC, "Public Safety Nationwide Interoperable Broadband Network: A New Model for Capacity, Performance and Cost" (FCC white paper, June 2010), http://fcc.gov/pshs/docs/releases/DOC-298799A1.pdf.

[23] FCC, *Connecting America*; FCC, "Mobile Broadband: The Benefits of Additional Spectrum (FCC staff technical paper, October 2010), http://download.broadband.gov/plan/fcc-staff-technical-paper-mobile-broadband-benefits-of-additional-spectrum.pdf.

[24] Morgan Stanley, *The Mobile Internet Report*.

[25] Hallahan and Peha, "Compensating Commercial Carriers for Public Safety Use."

[26] Hallahan and Peha, "Quantifying the Costs of a Nationwide Public Safety Wireless Network."

[27] FCC, "First report and order and third notice of proposed rulemaking in the matter of the development of operational, technical and spectrum requirements for meeting federal, state and local public safety agency communication requirements through the year 2010," *WT Docket No. 96-86,* September 29, 1998, http://wps.ncs.gov/documents/fcc98191.pdf ; FCC, "Second report and order in the matter of the development of operational, technical and spectrum requirements for meeting federal, state and local public safety agency communication requirements through the year 2010" *WT Docket No. 96-86,* July 13, 2000, http://wps.ncs.gov/documents/242.pdf.

[28] Ryan Hallahan and Jon M. Peha, "Policies for public safety use of commercial wireless networks." *Paper presented at the 38th Telecommunications Policy Research Conference.* 2010a. http://www.ece.cmu.edu/~peha/safety.html.

[29] To guarantee the public some access to cellular resources during emergencies, WPS users are limited to priority on 25 percent of a cell site's capacity during periods of congestion. Robert K. Ackerman, "Cellular Priority System Begins Operation," *SIGNAL Magazine* (March 2003).

[30] Ibid.

[31] Roberson and Associates, "Public Safety Priority Access to Shared Commercial Networks."

[32] 3GPP, *Technical Specification 23.401, 'General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access (Release 8)'* (December 2008), http://www.3gpp.org/ftp/Specs/html-info/23401.htm

[33] Motorola. "Comment before the Federal Communications Commission in the matter of Public Safety and Homeland Security Bureau seeks comment on interoperability, out of band emissions, and equipment certification for 700MHz public safety broadband networks," *PS Docket No. 06-229,* July 19, 2010, http://fjallfoss.fcc.gov/ecfs/document/view?id=7020549858.

[34] PSCR, "QoS Information."

[35] (Dahlman, et al. 2008, Sesia, Toufik and Baker 2009) Erik Dahlman, Stefan Parkvall, Johan Sköld, and Per Beming, *3G Evolution: HSPA and LTE for Mobile Broadband,* 2nd ed. (Oxford: Academic Press, 2008); Stefania Sesia, Issam Toufik, and Matthew Baker, *LTE – The UMTS Long Term Evolution: From Theory to Practice* (Chichester, UK: Wiley, 2009).

[36] (Olsson, et al. 2009) Magnus Olsson, Shabnam Sultana, Stefan Rommer, Lars Frid, and Catherine Mulligan, *SAE and the Evolved Packet Core: Driving the Mobile Broadband Revolution* (Oxford: Academic Press, 2009).

[37] Hallahan and Peha, "Quantifying the Costs of a Nationwide Broadband Public Safety Wireless Network;" Hallahan and Peha, "The Business Case of a Nationwide Wireless Network."

[38] 3GPP, "Technical Specification 24.301, 'Non-Access-Stratum (NAS) protocol for Evolved Packet System (EPS); Stage 3 (Release 8)'." December 2008c. http://www.3gpp.org/ftp/Specs/html-info/24301.htm; 3GPP "Technical Specification 29.212, 'Policy and charging control over Gx reference point (Release 8)'." December 2008d. http://www.3gpp.org/ftp/Specs/html-info/29212.htm.

[39] 3GPP, "Technical Specification 24.301."

[40] Hallahan and Peha, "Policies for Public Safety Use of Commercial Wireless Networks."

[41] Hamad Akbari, Hamad, and Frank Land, "Theories Used in Research: Sociotechnical Theory" (n.d.), http://www.istheory.yorku.ca/sociotechnicaltheory.htm.

[42] Robert P. Bostrom and Stephen J. Heinen, "MIS problems and failures: A socio-technical perspective," *MIS Quarterly* 1, no. 3 (1977).

[43] 3GPP, "Technical Specification 24.301."

[44] Peha, "A Public Private Approach to Public Safety Communications."

[45] 112th Congress, H.R. 3630; Moore, *Establishment of the First Responder Network Authority*.

[46] 3GPP, *3GPP – Specifications* (2010), http://www.3gpp.org/specifications

[47] Johnson, *Long Term Evolution IN BULLETS*.

[48] Ibid.

[49] Olsson, et al., *SAE and the Evolved Packet Core*.

[50] Ibid.

[51] Johnson, *Long Term Evolution IN BULLETS*.

[52] Ibid.

[53] Olsson et al., *SAE and the Evolved Packet Core*.

[54] Ibid.

[55] Johnson, *Long Term Evolution IN BULLETS*.

[56] 3GPP, *Technical Specification 23.401*.

[57] 3GPP, "Technical Specification 23.203, 'Policy and charging control architecture (Release 8)'" (December 2008), http://www.3gpp.org/ftp/Specs/html-info/23203.htm

[58] 3GPP, *Technical Specification 23.401*.

[59] Olsson et al., *SAE and the Evolved Packet Core*.

[60] Ibid.

[61] Ibid.

[62] 3GPP, "Technical Specification 24.301."

[63] Olsson et al., *SAE and the Evolved Packet Core*.

[64] Ibid.

[65] Ibid.

[66] The actual enforcement of the QoS policy decisions will occur in the home network for 'home-routed' traffic and in the visited network for 'local-breakout' traffic. This is because the PCEF (the function responsible for policy enforcement) resides in the PDN-GW and 'home-routed' traffic passes through a PDN-GW in the home network, while 'local-breakout' traffic is routed through a PDN-GW in the visited network.