VirginiaTech
College of Engineering

Wireless @ Virginia Tech
432 Durham Hall (0350)
Blacksburg, Virginia 24061
540/231-2958 Fax: 540/231-2968
www.wireless.vt.edu

November 8, 2012

Lawrence Strickling
Assistant Secretary for Communications and Information
Department of Commerce
1401 Constitution Avenue NW
HCHB Room 7324
Washington DC 20230

Att: FirstNet Conceptual Network NOI

Enclosed we provide a brief response to the FirstNet NOI regarding the conceptual network architecture. The focus on our comments is on the information assurance aspects of LTE and contains a summary of some of our preliminary analysis. This work is still in progress and we would be pleased to share details of our current and future findings on this issue.

Sincerely,

Dr. Jeffrey H. Reed
Director, Wireless@ Virginia Tech

BEFORE THE
DEPARTMENT OF COMMERCE

In the Matter of	)	
	)	
National Telecommunications and	)	
Information Administration	)	Docket No. 120928505–2505–01
	)	RIN 0660–XC002
Development of the Nationwide	)	
Interoperable Public Safety Broadband	)	
Network	)	

COMMENTS OF WIRELESS @ VIRGINIA TECH

Dr. Jeffrey H. Reed
Willis G. Worcester Professor
Director of Wireless @ Virginia Tech
432 Durham Hall, MC0350
1145 Perry Street
Blacksburg, VA 24061
reedjh@vt.edu
(540) 231 2972

Marc Lichtman
Graduate Research Assistant
Virginia Tech
marcll@vt.edu

INTRODUCTION AND EXECUTIVE SUMMARY

The Wireless @ Virginia Tech research group appreciates the opportunity to respond to the National Telecommunications and Information Administration (NTIA) request for comments on the Development of the Nationwide Interoperable Public Safety Broadband Network. This comment is regarding the vulnerability of LTE to intentional and sophisticated jamming attacks.

If LTE technology is to be used for the air interface of the public safety network, then we should consider the types of jamming attacks that could occur five or ten years from now. It is very possible for radio jamming to accompany a terrorist attack, for the purpose of preventing communications and increasing destruction. Likewise it is possible for criminal organizations to

create mayhem among public safety personnel by jamming. In addition, it is possible for a jammer to increase its effectiveness by employing a sophisticated strategy. This is especially likely when every technical aspect of the target signal is known. An example strategy would be to target specific control or synchronization signals, in order to increase the geographic range of the jammer and better avoid detection. The availability of low-cost and easy to use software-defined radios makes this threat even more realistic.

Preliminary research has been performed to show the extent to which LTE is vulnerable to jamming. It was shown that extremely effective attacks can be realized, using fairly low complexity. It would be in the interest of public safety to put forth an effort to find solutions to the described problem, and ultimately improve the wireless interface of Public Safety LTE.

Preliminary Research

In order to show the vulnerability of the LTE wireless interface to jamming, we present a series of efficient attacks that are designed to cause denial of service to one or more LTE cells. Each attack targets one or more LTE subsystems, either in the downlink or uplink signal. The selected attacks described below represent a balance between effectiveness and complexity.

Synchronization Signal Jamming (SSJ): When a UE wants to connect to an eNodeB, it has to first go through a series of synchronization steps. First, it detects the Primary Synchronization Signal (PSS) which allows the UE to synchronize to each slot and gives it the cell ID. Next, it detects the Secondary Synchronization Signal (SSS) which tells the UE the cell ID group, which method of duplexing is used, and the cyclic prefix length. The SSS also allows the UE to detect when each radio frame starts. Both the PSS and SSS are mapped to the central 62 subcarriers

(not including the DC subcarrier). After synchronizing with the PSS and SSS, the UE receives more information about the cell by decoding the Master Information Block (MIB). The MIB contains information essential for initial access to a cell [1]. It consists of 14 bits that contain the downlink system bandwidth, the Physical Control Format Indicator Channel (PHICH) size, and information allowing frame synchronization. It is mapped to the central 72 subcarriers, and appears in slot 1 of each frame. The three signals are not present in all ten subframes, but they are always mapped to the same subcarriers.

The Synchronization Signal Jamming (SSJ) attack is designed to deny the UE access to the PSS, SSS, and MIB. The jamming waveform used for the SSJ attack is noise that spans the center 73 subcarriers. The DC subcarrier is included for the sake of complexity, even though it does not contain information. The SSJ attack does not involve jamming specific symbols (it uses a 100% duty cycle), so the jammer does not have to be synchronized to the eNodeB. The SSJ attack is simply a brute force method of denying the UE three different mechanisms that it needs to access a cell. The act of only jamming certain subcarriers allows the SSJ attack to have roughly a 3 dB gain over traditional barrage jamming, which can be thought of as an increase in jamming radius for a jammer that is power constrained.

Primary Synchronization Signal Jamming: Detecting the PSS is the first step a UE takes in accessing a cell. The PSS uses a sequence length of 63, and the center element is nulled because the downlink DC subcarrier is never used for transmission. There are three PSS sequences used in LTE, and each one corresponds to one of the three sectors. The UE must detect the PSS without any knowledge of the channel, so it finds the timing offset that corresponds to the

maximum cross-correlation for each of the three sequences, and uses it to synchronize in the time domain. For FDD, the PSS only occurs in slots 0 and 10 (there are 20 total slots per frame).

The SSJ attack discussed previously injects noise into the subcarriers that contain the PSS. An attack that only targets the PSS can be realized by only jamming the symbols that contain the PSS. However, the jammer would have to cause a fairly high jammer-to-signal ratio, because the PSS is designed to be detected at high interference levels, so that the UE can also detect neighboring cells.

A more effective method of causing a PSS attack would be to simply transmit one of the three PSS sequences, and thus create a bogus PSS. If the jammers received power at the UE is greater than the eNodeB's, then the UE is most likely going to synchronize to the bogus PSS. This is because a cross-correlation process is used to detect the PSS non-coherently. A jammer using this method would not need receiving capability, because it would simply start the bogus PSS transmission at a random time, leading to uniformly distributed timing relative to the correct PSS signal. If the UE synchronized to the bogus PSS, then is not synchronized in time to the eNodeB, and it will not know when each OFDM symbol starts, and hence will not be able to detect the SSS or decode the MIB.

This attack appears to work, until considering the cell reselection procedure. If a cell does not provide a certain level of quality, then the UE begins the cell reselection procedure, where it tries to access the cell with the next strongest signal. The solution is to spoof all three PSS sequences. A jammer transmitting three bogus PSSs only has to transmit six symbols in every frame, on 62 subcarriers. A downside to PSS jamming is that it will not immediately cause Denial of Service (DOS). It will prevent new UEs from accessing the cell(s), and cause UEs in idle mode to reselect a bogus cell. Therefore PSS jamming is not effective for an attack

intended on causing immediate DOS. However, it is sufficient for an attack that will last a long period of time. Because the jammer barely has to transmit anything, the PSS jamming attack offers roughly 20 dB of gain relative to the barrage jamming attack. This results in an extremely efficient jammer.

Fortunately, this type of attack can be prevented by employing a cell reselection implementation that is able to blacklist "bogus synchronization signals", by keeping track of the time-delay in the cross-correlation. Although this is not required by the LTE specifications, adding this type of mitigation may be worthwhile in Public Safety UE.

Physical Uplink Control Channel Jamming: The Physical Uplink Control Channel (PUCCH) is used to send the eNodeB a variety of control information, including scheduling requests, Hybrid Automatic Repeat Request (HARQ) acknowledgements, and channel quality indicators. The PUCCH is mapped to the resource blocks on the edges of the system bandwidth. These resource blocks are evenly split between the two edges of the system bandwidth, and the UE rapidly alternates between the two sets of resource blocks, for the purpose of frequency diversity. When a UE is transmitting on the PUCCH it is not transmitting anything else [2]. This allows PUCCH jamming to be feasible, even with SC-FDMA in use. PUCCH jamming is possible when the only a priori knowledge is the system bandwidth and location in the uplink signal in the frequency domain.

The signal sent on the PUCCH by the UE depends on the type of information it wants to send. For scheduling requests, all the UE has to do is transmit energy in its assigned slot. This causes the eNodeB to assign the specific UE additional uplink resources. This means PUCCH jamming will cause the eNodeB to assign every active UE additional uplink resources (which

they probably do not need), and cause degradation of service. The signal transmitted by the UE for ACK and NACK responses is not as straightforward; it involves modulating the ACK or NACK indicator onto a predefined sequence which is then cyclically shifted and scrambled. This system is meant to allow multiple PUCCH transmissions to exist in the same time and frequency slot. Successful PUCCH jamming will cause ACKs to not reach the eNodeB, resulting in retransmissions and further degradation of service. The last type of control information sent on the PUCCH is channel state information, which is used by the UEs to send the eNodeB information about the channel quality. The eNodeB uses this information to assign subcarriers to users that experience better channel conditions on the corresponding frequencies, as well as choose which modulation scheme to use. As in the ACK indicators, the information is sent to the eNodeB through modulation with a UE-specific sequence, which is then scrambled and cyclically shifted. The corruption of channel quality information is not as detrimental to the LTE service as missed ACKs, but it is likely to help accelerate the process of causing DOS. PUCCH jamming offers roughly 5 dB of gain compared to barrage jamming, because the jammer can focus its energy into the control channel subcarriers.

Conclusion

These comments describe extremely effective attacks can be realized, using fairly low complexity. It would be in the interest of public safety to take measures to reduce the vulnerability of Public Safety LTE, and lower the likelihood of an effective jamming attack. Certainly there are important cost advantage of keeping the public safety LTE system compatible with commercial devices and systems. Seeking solutions that achieve this compatibility while

providing protection are desirable. We thank you for considering our views, and are eager to address any subsequent questions.

References

[1] Matthew Baker and Tim Moulsley. Downlink physical data and control channels. In Stefania Sesia, Issam Toufik, and Matthew Baker, editors, *LTE, The UMTS Long Term Evolution: From Theory to Practice*, chapter 9. John Wiley & Sons Ltd, Chichester, West Sussex, United Kingdom, second edition, 2011.

[2] 3GPP Technical Report 36.211, "Physical Channels and Modulation", www.3gpp.org.

Dr. Jeffrey H. Reed
Willis G. Worcester Professor
Director of Wireless @ Virginia Tech
Virginia Tech

Marc Lichtman
Graduate Research Assistant
Virginia Tech
