# Jamming and Interference Induced Denial of Service Attacks on IEEE 802.15.4 Based Wireless Networks

Jacob Brodsky, PE, Anthony McConnell, PE
Washington Suburban Sanitary Commission
Laurel, Maryland
jbrodsk@wsscwater.com, amcconn@wsscwater.com

**Abstract:** This paper documents a denial of service attack using only primitive concepts and inexpensive hardware against IEEE 802.15.4 low power networks. The efficacy of swept frequency attacks, dead-carrier attacks, and Bipolar Phase Shift Keying with square waves is documented. This paper will also document the ease and inexpensive nature of the components needed to conduct these attacks, and the technical proficiency required.

We express significant concerns of the ramifications of the Clear Channel Assessment feature of the IEEE 802.15.4 specification in a jamming context.

Designs can be hardened against such attacks by using directional antennas, by ensuring paths have a physical line of sight, and monitoring signal strength. Recommendations for mesh networking include keeping track of hop counts, as well as noise floor and eye closure monitoring. We recommend obtaining direction finding equipment and practicing with it to find stray signal sources.

While Part 15 users may not act against other signal sources, they can at least identify where these sources come from and take measures to avoid them.

**Keywords:** Wireless, IEEE802.15.4, Jamming, Denial of Service, Clear Channel Assessment

## 1 Introduction

We have been watching the new ISA 100 [1] standards as they grow. However, as both of us have a significant telecommunications background we naturally had concerns about the actual physical layer of any low power RF standards. The ISA-100 uses the IEEE 802.15.4 [2] as its "physical layer". So we decided to research the IEEE 802.15.4 physical layer with real devices and real hardware.

Before going any further, we need to discuss the ground rules and the ethics behind our tests.

### 1.1 Legal and Ethical Issues

First, we limited our experiments to the guidelines of 47 CFR 15.23 (Home Built Devices) [3]. Good engineering practice also dictated that we conduct the experiments

where they would not cause interference to others. Much of the distance testing was done in Mr. Brodsky's hay field. This served us well because it is wide open area; with no uncontrolled wireless users closer than 1000' in any direction. No usable 802.11b or 802.11g signals were found (We did notice two signals, however they were so weak that they weren't reliably detectable). This was deemed sufficient for our experiments.

Second, we kept power outputs low. Our devices operated at similar magnitudes of effective radiated power levels as the 802.15.4 devices we were experimenting with. We were primarily interested in the relative efficacy of these methods, not how wide an area we could jam.

Third, we deliberately used unmodified off-the-shelf devices.

Fourth, philosophically, this is a test of a technology. The products named here are meant to be representative of the industry as a whole. We're trying to show what simple jamming methods can do and how one might defend against them.

Fifth, we focused on simple jamming techniques. Although we have the background and the test equipment to pursue more complex methods, most attackers do not. We're interested in the likely attack scenarios, not the exotic ones. We believe that those with the resources to perform difficult jamming efforts will find that it is counterproductive to expend too much money and effort against these targets when less expensive methods can be more effective.

## 2 Regulations

Most evangelists of wireless technology neglect to mention 47 CFR 15.5(b). But we will quote it here:

> "Operation of an intentional, unintentional, or incidental radiator is subject to the conditions that no harmful interference is caused and that interference must be accepted that may be caused by the operation of an authorized radio station, by another intentional or unintentional radiator, by industrial, scientific and medical (ISM) equipment, or by an incidental radiator."

Basically, it says that if something causes a wireless network to malfunction, the users have no legal recourse. And to make matters worse, if the wireless network interferes with a licensed activity, the unlicensed wireless activity must shut down. This regulation is not unique to the United States. Most other countries have adopted very similar regulations.

Second, there are the regulations concerning the limits of unlicensed operation. The relevant regulation here is 47 CFR 15.247. Without quoting everything, most of the issues that concern us are as follows:

- Up to 1 Watt (+30 dBm) in to 6 dBi antenna
- For every 3 dB of gain above 6 dBi antenna, reduce power by 1 dB.

In other words, you can use antenna gain, or you can use up to one watt of power, but you can't use both at the same time.

## 2.1 Specifications and Vulnerability Discussions

### 2.1.1 Power Availability and Consumption

These devices are limited by very low power consumption requirements. In theory, this should make the receiver vulnerable to strong signals. The CC2480 receiver in our sample appears to be much better than we expected, though there was no specification for blocking signal performance or some form of Inter-Modulation Distortion (IMD).

Often in low power, low noise amplifiers, these performance issues become quite significant. For example, too much signal in the first stage of a receiver can actually be rectified by the transistor or diodes, and this DC energy can change bias of the amplifier or mixer, effectively cutting off all signals to the following stages. The lower the power of the first stage amplifier, the less energy it takes to do this. This is called "Blocking." It is a common test done with almost all receivers.

A second hazard is IMD. IMD products result from non-linearities of the active devices in a receiver. This isn't as much of a problem as it used to be with earlier semiconductors, but it is still quite significant. One can measure this in many ways, and it is worth noting that people write books about this subject. Our curiosity, however, is how much the noise floor of this receiver goes up when a strong signal presents itself and is de-spread. As we found out later, this issue is irrelevant because other aspects of the specification get in the way.

### 2.1.2 Clear Channel Assessment (CCA)

One of the interesting features of IEEE 802.15.4 is that it attempts to ensure that it doesn't stumble on any existing activity on the ISM band. Basically, it is supposed to conduct "energy detection" on the channel and back off or change channels if it finds activity. Activity is defined as being no more than 10 dB higher than the weakest signal the receiver can reliably detect (Sections 6.5.3.3, 6.9.7, and 6.9.9 of the specification). Since the standard mandates sensitivity less than -85 dBm, this corresponds with no more than -75 dBm of energy measured at the antenna terminals.

There are two possible reasons for doing this: One is that these ISM devices should be good RF citizens by not stepping on other legitimate signals. Two, radiating on a busy channel may not be able to get the remote node to reliably demodulate the packet without errors. As these are meant to be very low power battery operated devices, it could waste precious battery power. If it isn't likely the signal will get through, then it shouldn't waste battery power by trying to transmit anyway.

Note, however, that the latter assumption isn't always valid. The far ends of each signal path may not be able to hear what the transmitting station hears, so there is a significant chance that there won't be a collision on the air after all.

There are two basic modes of CCA. The first is just plain energy detection. The second is signal detection.

In the first mode, the receiver simply measures how much energy is present in the passband of the receiver. It does not pay much attention to what the signal might be.

The second mode of CCA is signal detection. If the receiver detects another possible IEEE 802.15.4 signal, no matter what strength, it is supposed to inhibit transmission.

These two methods could be used together according to the specification(in other words Mode 1 AND Mode 2 or Mode 1 OR Mode 2). The specification calls this Mode 3.

If an 802.15.4 device is configured to do so, it is supposed to conduct a CCA check and then change channels to a clearer frequency, if the channel has activity.

### 2.1.3 IEEE 802.15.4, Annex E

The IEEE 802.15.4 standard isn't the only such standard on the 2.4 GHz ISM spectrum. There is a veritable zoo of other ISM wireless standards out there, and some of the most ubiquitous standards are the IEEE 802.11b and 802.11g standards. This issue is addressed in Annex E of the 802.15.4 standard.

Figure 1 below from Annex E seems to indicate that while an 802.11b signal is transmitting, that the 802.15.4 equipment needs to be far away to have any hope of communicating.

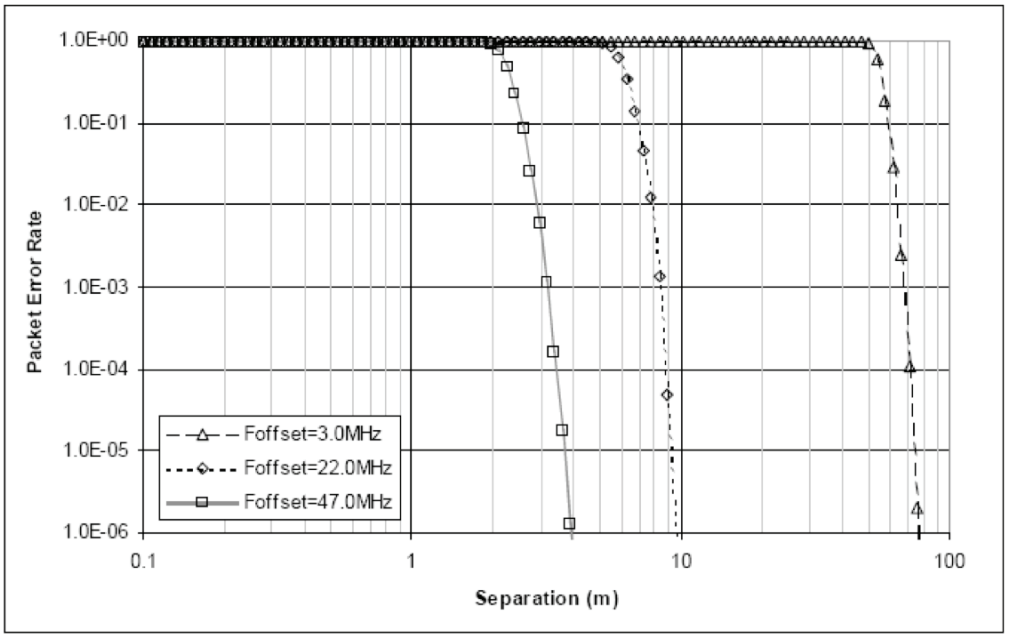


**Figure E.3—IEEE 802.15.4 receiver, IEEE 802.11b interferer**

*Figure 1 – IEEE 802.15.4 Annex E Figure E.3*

Note that for an IEEE 802.15.4 signal source anywhere inside the passband of an IEEE 802.11b channel; one must station the 802.15.4 transceivers at least 50 meters away to communicate. This comes from the Foffest = 3.0 MHz curve and a packet error rate of 1.0E-01, a maximum error rate for most applications.

Also note that while both the 802.11b channels and 802.15.4 use 5 MHz steps (except for channel 14 of 802.11b) that the 802.11b signal is 22 MHz wide, while the 802.15.4 signal is only 2 MHz wide. In practice this means that it may only take four 802.11b/g signals to cover the 2.4 GHz ISM band completely.

Interestingly, most of the study in Annex E was concerned with ensuring that other services don't get interfered with by the 802.15.4 specification. This was one of the few graphs that indicated there might be a problem with other services interfering with 802.15.4.

The actual tests turned out to not to be so alarming. This is one area where things turned out to be better than expected.

## 2.2 Performance of Test Units

Our test units came from a CC2480 engineering test kit from Texas Instruments. The test units from TI were meant to be 1 mW devices with nearly isotropic antennas. They were capable of communications over a distance of almost precisely 100 feet in an open field. However, for most of our tests, we kept these units about 4 feet apart. This is approximately 10 wavelengths, which represents a reasonable far-field antenna approximation.

Texas Instruments is to be commended for the RF portion of the design of these chips. These units significantly exceed most of the RF specifications in IEEE 802.15.4. The device sensitivity is -92 dBm instead of -85 as specified, the adjacent channel selectivity was specified at 0 db and these devices have 40 dB of rejection, the stability and frequency tolerance is substantially better as well.

Because we were experimenting with several configurations, we used well packaged devices with connectors. The cost of our packaged VCO was under $70. The packaged mixer was under $40.

Our VCO and mixer modules were ordered from Mini-Circuits online. We used the ZX-95-2755+ VCO and a ZX05-30W-S+ diode ring mixer module. The mixer is characterized up to 4 GHz.

If we were interested in production with modest quantities of fewer than 25 units, Mini-Circuits offers broadband amplifier modules suitable for micro-stripline installation for under $2 each. They offer VCO modules for as little as $20, and mixers for $7. The signal sources we used could be built with very inexpensive 7555 timer chips. Potent jamming systems could be built in small quantities for about $40 each.

# 3 Coexistence

## 3.1 Practical Measurements

The graph in Figure 1 looks ominous. We have to remember that it will not be common for both signals to be on the air at the same time. On a lightly loaded network, proximity should not be a major problem. At those traffic levels, it may even be possible for 802.11b/g and 802.15.4 to share the same frequencies.

However, we do not recommend sharing the same frequencies. The problem is that should something that causes lots of traffic show up on the network, the assumption of lightly loaded airwave channels may not be reasonable.

Our experiments showed that without a significant level of traffic it was hard to get the 802.15.4 devices to miss a single poll, even at close range. We simulated our traffic with a very high volume of large packet pings. This was planned to be a relatively symmetric traffic loading.

Nevertheless, the theoretical results in Annex E should not be ignored. We suggest that readers study the bandwidths and channel plans for 802.11b and 802.15.4 and coordinate something reasonable on the plant. Note that under Part 15 your only authority to coordinate any unlicensed activities comes from being a property owner. People can legally fly kites or unmanned aircraft overhead with unlicensed devices.

# 4 Jamming

For the purposes of our discussion, we are going to borrow the term "burn-through distance" from electronic warfare in Radar applications. In this context we're using it to describe the distance a jamming transmitter has to maintain to keep interfering with an RF link. In other words, given an 802.15.4 pair of transceivers approximately 1.2 meters apart, how far must a jammer be, perpendicular to the pair, to keep the two from talking. (Note: the 1.2 meter figure is approximately 10 times the wavelength of 2.4 GHz, at which one can apply the far-field antenna approximations)

## 4.1 Unmodulated Carrier

First, we attempted to validate the assertion that a dead carrier anywhere near the signal wouldn't do much damage. Unfortunately, it did.

An unmodulated carrier, on the active channel caused the test devices to stop communicating. It appears the test units were confined to channel 16 of the 802.15.4 band plan. We speculate that the Clear Channel Assessment was inhibiting activity, as defined in Section 6.7.7 and 6.7.9 of the IEEE 802.15.4 specification, though we lacked the Software Development environment to know if this was indeed the cause of the problem. The burn-through distance was about 30 feet. Keep in mind that the dead carrier was about six times stronger (+8 dBm) than the signal source.

Given the relatively low expense and simplicity of this attack, one could very easily purchase sixteen of these devices and set them up, one per channel. However, there is an easier and less expensive way to generate more effective noise.

## 4.2 Bipolar Phase Modulation

Next, we considered that if we could put this energy on all sixteen channels at once, that it would inhibit all 802.15.4 traffic. We figured that phase modulating the VCO with a 500 kHz square wave would do this pretty well.

We took the VCO output fed it to the LO port on a diode ring mixer. Using a function generator, we inserted a 500 kHz square wave in to the IF port of the mixer. The output on the RF port was radiated from an antenna.

This was also very successful at jamming the 802.15.4 connection. The burn through distance was approximately 15 feet. Keep in mind, however, that the signal we were radiating was at the same level as the other 802.15.4 devices. With reasonable amplification and antenna gain, this could black out a fairly wide area and several channels at once. It would take a grand total of maybe three of these units to cover most of the band with so much noise that it would shut down all sixteen channels very easily.

## 4.3 Chirping

We tried a gross Frequency Sweeping attack against the 802.15.4 signal. We swept the VCO through all the channels rapidly with a 100 kHz triangle wave imposed upon the VCO tuning port. Running the oscillator directly in to the antenna we had at least a five fold advantage against the incoming signal. Even then, our burn-through distance was only 1.2 meters or so –approximately the distance between the units themselves. We thought of trying various other sweep modulation frequencies. We suspect that a properly configured 5 MHz sine wave signal at the appropriate modulation index ought to cover every channel with an FM sideband. However, adjusting it properly is not easily done without a spectrum analyzer, so we discarded this approach as a likely candidate for a simple jammer.

## 4.4 Other Methods

We did not use any amplifiers or unusual antennas. The antenna we used in our jammers could not have more than 3 dBi of gain. However, amplifiers and antennas are inexpensive and readily available. We purchased TI's CC2591 amplifier kits for future testing. The Power Amplifier (PA) section has an output of approximately +22 dBm, with about 17 dB of gain. Amplifiers up to one watt are legally available for the 2.4 GHz band. We purchased a 1 watt amplifier for 802.11b/g on E-Bay for $60. It would work equally well for our jammer.

The methods we've outlined so far can be augmented with home built antennas. There are plenty of designs on the Internet for 2.4 GHz "Coffee Can" antennas. In practice, one should probably use a slightly smaller diameter can.

Future tests will probably entail pulsing the CC2590 amplifier on and off at a rate of around 300 kHz or so. It has very rapid PA keying, going to full power in under one micro-second.

In addition the following other methods are worth mentioning here:

### 4.4.1 Repeaters

This is an insidious method that we are merely mentioning here for completeness. Repeater jamming takes a garbled copy of the signal and puts it back on the air. The garbled signal will cause significant problems for the incoming signal to properly de-spread. There are several ways to build such devices. Repeater attacks will be notoriously

difficult to find as they're not on the air all the time, they can be very low powered, and are only present when the desired signals are on the air. The downside is that while the expense of construction may be low, the technical prowess required to build one is significant.

### 4.4.2 Brute Force: The Microwave Oven

We should not overlook an ordinary signal source: a plain old microwave oven. It is powerful, ranging from 500 watts to 1 kW, and very inexpensive. Find a way to power it, defeat the door interlock and it will saturate the front ends of every single 2.4 GHz receiver for several thousand feet.

Aside of the obvious issues of legality, there are significant safety hazards to the perpetrator. This energy must be carefully directed at a target. The operator must keep away from the unit to avoid suffering life threatening RF burns.

# 5 Mitigation Strategies

We feel such discussion is very dangerous without some suggestion of defense. With that in mind we present the following mitigation strategies. The fundamental thinking behind all this is that one should not remain ignorant of the physical layer of the network.

## 5.1 Antennas

One quick look through a search engine will yield all sorts of interesting recipes for building tin can antennas for 2.4 GHz. There are also many other designs for directional antennas, ranging from printed circuit board yagi antennas, to full size dish designs.

We want to make two points: First, omni-directional antennas are not good choices for rejecting unwanted signals. Directional antennas are useful, not just because they make it possible to use less power, but also because they reject signals from unwanted directions.

Second, avoid pointing antenna radiation patterns at the horizon. The goal in this application is not distance. It is a reliable, secure link. We recommend setting up major nodes at ground level in the middle of a plant campus and pointing all peripheral antennas toward that central ground point. This is exactly the reverse of what most people would think to do. Most would think to put the central antennas up high, where it could cover maximum ground. We are proposing that if one uses omni-directional antennas, that it be kept sheltered at ground level somewhere in the center of the campus, so as to make it less likely that someone from beyond the fence line will have a harder time interfering with it.

Looking back at the regulations in Part 15.247, and assuming a 0 dBm output from the 802.15.4 device, one would need an antenna gain of at least 36 dBi before attenuating the power becomes a legal necessity. While such antennas exist, the size would be unwieldy in this application. For all practical purposes the antenna can be as big as we want it to be.

## 5.2 Link Monitoring

We also notice a very strong tendency in among IT network engineers to compartmentalize and layer technologies. However, this is one instance where it might be smarter to avoid the layering.

The first element in this is characterization of the link itself. Many advocates suggest a "lick-em and stick-em" approach, and never mind the propagation path. We suggest that users should not do this unless they don't care whether they get the data or not. We recommend direct, line-of-sight paths to the destination nodes.

We recommend monitoring link quality at all times. At a minimum, this would include Received Signal Strength Indication (RSSI). In particular, a sudden increase in RSSI should be cause for investigation. However, if at all possible, eye closure data could help too. The latter indicates how much distortion or ISI the incoming signals have.

Neither of these is available through the CC2480 application interface. In fairness to TI, they do sell other chips with RSSI in the application interface, and this is outlined in Texas Instruments Design Note 505 [4].

## 5.3 Mesh Networking

Some advocate massive mesh networks to overcome problems with jamming. While there is some merit to setting up networks this way, we have misgivings, particularly given what we have observed with CCA. Will the mesh network be effective against a deliberate jamming attack.

The first issue with a large mesh network is that each node pretty much has to have an omni-directional antenna. If one is to integrate a large mesh network, there needs to be a hop count for each packet of data. As hop counts from certain sites grow, we would expect someone to take action. Of course, how many would know to do this, and how would it be alarmed in a manner that operators could understand and react to appropriately?

## 5.4 Direction Finding

Finally, no plant would be secure without direction finding gear. Unfortunately, this is one area where the tools simply do not exist. The jamming systems outlined here are all pretty easy to find, given that they radiate most of the time. A more sophisticated jamming system based upon a repeater might not be.

Unfortunately, the very features that make IEEE 802.15.4 devices battery powered also make it very difficult to find. We are working on inexpensive methods to locate and identify these devices.

# 6 Conclusions

The state of the art for wireless security is still immature. It does not take physical security issues in to account at the physical layer. As a bare minimum, we think all devices should have directional antennas made available, and RSSI visible at the application layer. If we were to use Packet Error rates, the damage would already be happening. The goal is to find out before the error rates go up significantly.

Unlike conventional wired media, there is no physical access control to wireless media. Literally anyone or anything can jump on the media for this network. We strongly recommend that users consider the ramifications of all permutations of denials of service. We also strongly recommend the use of directional antennas pointed away from the horizon. Aside of the gain from directing the energy toward where it needs to go, it also avoids picking up signals from other places that might represent attack points.

We also believe that all wireless usage on a plant campus should be coordinated on company property. While there is legally no way to coordinate frequencies outside the plant, the distances to the company fence line may be significant enough that it may not matter.

Furthermore, as ISM band use grows, we caution users that they are not in a future-proof application here. Tomorrows new killer wireless applications may have to be tightly policed from plant premises to prevent interference. Given how pervasive such devices are and how many more are growing every day, this will not be easy.

Despite the potential savings, we do not recommend the use of wireless technology for any except non-critical applications.

---

**About the Authors** - Jacob Brodsky and Anthony McConnell are Registered Control Systems Engineers at the Washington Suburban Sanitary Commission, with a combined experience of 45 years.

Mr. Brodsky began his career there in 1986 as a Tele-Communications Technician while attending evening classes at the Johns Hopkins University. He earned his Bachelors Degree in Electrical Engineering in 1990.

Mr. McConnell has many years of various telecommunications experience ranging from in-car cameras for NASCAR, police and fire radio communications, radio paging networks, and utility communications networks.

Mr. Brodsky and Mr. McConnell, as licensed amateur radio operators, have participated in specially authorized tests by the Federal Communications Commission on Direct Sequence Spread Spectrum technology in the early 1990s, before 802.11 technologies were available.

## Glossary

**Decibel, dB**: A decibel is 10 log (P1/P2). It is ten times the log of the ratio of two power readings. It is also 20 log (V1/V2); assuming the impedance of the voltages is the same. This is merely a different way to measure the same thing.

**dBm**: Decibels compared to one milliWatt. A 0 dBm signal is one milliWatt. A +10 dBm signal is 10 milliWatts. A -10 dBm signal is 0.1 milliWatts.

**dBi**: Antenna performance measured as decibels of gain compared to an Isotropic Radiator. For example, a dipole antenna is theoretically 2.2 dBi.

**ISI**: Inter-Symbol-Interference. This is a form of distortion or interference that results from non-linearities in a Quadrature-Amplitude Modulation receiver or transmitter system.

**ISM**: Industrial Scientific & Medical. This is a piece of electromagnetic spectrum intended for unlicensed RF uses. Common uses include RF heating (such as a microwave oven), Plasma excitation, and unlicensed communications networks.

**Isotropic Radiator**: An antenna that radiates equally in every spherical direction, without loss.

**Process Gain**: When a spread spectrum signal is de-spread, any signals or noise across the spectrum of the original signal will drop by the degree to which the signal is spread. This is called the Process Gain. The greater the ratio of the spreading bandwidth to the data communications rate, the greater the Process Gain can be.

**VCO**: Voltage Controlled Oscillator: An oscillator whose frequency can be controlled using a variable input voltage.

## References

[1] ISA 100 Committee on Wireless Systems for Automation, ISA100 Standards Overview and Status, Road Show Presentation, 2008.

[2] IEEE Computer Society, IEEE Std 802.15.4-2006, Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs), 2006.

[3] Federal Communications Commission, 47 Code of Federal Regulations (CFR) Part 15 Radio Frequency Devices, 2004.

[4] S. Namtvedt, Design Note 505: RSSI Implementation and Timing, Texas Instruments, 2007.