

DESIGN OF A LOW POWER CARGO SECURITY DEVICE USING A
MICROPOWER ULTRA-WIDEBAND IMPULSE RADAR

A Thesis
presented to
the Faculty of California Polytechnic State University,
San Luis Obispo

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in Electrical Engineering

by
Brian Matthew Wihl
December 2013

© 2013
Brian Matthew Wihl
ALL RIGHTS RESERVED

COMMITTEE MEMBERSHIP

TITLE: Design of a Low Power Cargo Security Device Using a
Micropower Ultra-Wideband Impulse Radar

AUTHOR: Brian Matthew Wihl

DATE SUBMITTED: December 2013

COMMITTEE CHAIR: John Oliver, PhD
Associate Professor of Electrical Engineering

COMMITTEE MEMBER: Bridget Benson, PhD
Assistant Professor of Electrical Engineering

COMMITTEE MEMBER: Lynne Slivovsky, PhD
Professor of Electrical Engineering

ABSTRACT

Design of a Low Power Cargo Security Device Using a Micropower Ultra-Wideband Impulse Radar

Brian Matthew Wihl

Each year, thousands of cargo containers are broken into during shipping, costing billions of dollars in lost and damaged goods. In addition to removing its contents, intruders can also add unwanted and dangerous materials to a container, posing a threat to National Security. The possibilities of cargo container break-ins require that the containers go through check points at which they are physically searched. These searches often require the opening of the container, unloading and inspecting all cargo, and then loading the container and resealing it. This is a long and costly process.

Because of the high costs of break-ins and inspections, many security devices have been developed to ensure the safety and detect the tampering of cargo containers. Most of these mechanisms involve more intricate door locks and electronic seals that are able to add a degree of security to the containers. Other “smart” cargo security devices exist, which employ a variety of sensors to detect intrusion, however, none of the current solutions are reliable and practical enough to eliminate the necessity for frequent inspection of cargo containers. The shipping industry is in need of a reliable, unobtrusive, low-cost, low-effort cargo security device.

Over the last two decades, Lawrence Livermore National Laboratory (LLNL) has been developing a micropower impulse radar capable of detecting objects and motion within a short to medium range. Due to its past uses for intrusion and motion detection, the LLNL micropower impulse radar is a top prospect for a sensor technology used in a cargo security device.

This paper describes the design of a low-power, low-cost cargo security device which uses the LLNL micropower impulse radar for the detection of shipping container intrusions. With the evaluation of the impulse radar as well as various other sensors, a device was created which successfully detected intrusions over 98% of the time with the capability of lasting 5 to 6 months when powered by two AA batteries.

TABLE OF CONTENTS

LIST OF TABLES	vi
LIST OF FIGURES	vii
1. Introduction.....	1
1.1. System Problem Statement	2
1.2. Overview	3
2. Existing Technologies.....	5
2.1. Passive Infrared	6
2.2. Ultrasonic.....	7
2.3. Microwave	8
2.4. Tomographic Sensor Networks	9
3. LLNL Impulse Radar.....	11
4. Design	15
4.1. Selection of Parts	16
4.2. Passive Infrared Sensor.....	17
4.2.1. Vibration, Shock, and Movement Sensor	18
4.2.2. Processor	19
4.2.3. Memory.....	20
4.2.4. Temperature/Humidity.....	21
4.3. Bluetooth	21
4.3.1. Power Consumption.....	23
4.3.2. Communicating with BLE and the nRF8001	28
4.4. Event Driven Firmware	34
5. Intrusion Detection.....	39
6. Results.....	46
6.1. Power Consumption	46
6.2. Intrusion Detection	48
6.2.1. Test Development	51
6.2.2. Door Event Tests.....	53
6.2.3. Container Motion Tests.....	57
Conclusion	60
Bibliography	62

LIST OF TABLES

Table 1: Results for door opening with empty container test case. 54

Table 2: Results for door opening and entry with empty container test case..... 55

Table 3: Results for door opening with loaded container test case..... 56

Table 4: Results for door opening and entry with loaded container test case..... 57

Table 5: Results for movement test case..... 58

Table 6: Results for vibration test case. 59

LIST OF FIGURES

Figure 1: PIR sensors detection pattern [10].....	7
Figure 2: Ultrasonic sensor operation.	8
Figure 3: Microwave sensor [10].....	9
Figure 4: Tomographic sensor network [15].....	10
Figure 5: LLNL micropower impulse radar block diagram [13].	13
Figure 6: Operation of LLNL radar as a motion detector.	14
Figure 7: System block diagram for CSD.....	16
Figure 8: Assembled micropower impulse radar.	17
Figure 9: Current profile of BLE device in connection state [14].	24
Figure 10: Connection between application controller and nRF8001 [14].	26
Figure 11: Current consumption of nRF8001 versus connection and advertisement interval settings [14].....	27
Figure 12: BLE stack [14].....	29
Figure 13: Illustration of service pipes used by the nRF8001 [14].	30
Figure 14: nRF8001 setup process [14].	31
Figure 15: Typical communication between application controller and nRF8001 [14].	33
Figure 16: Interrupt routines for sensors and communication.	35
Figure 17: Log file from CSD.....	42
Figure 18: Windowed intrusion algorithm.	43
Figure 19: Intrusion detection algorithm beginning with a radar event.....	45
Figure 20: Current profile for CSD.....	47
Figure 21: Current consumption of major system components.	48
Figure 22: Loaded cargo container used for test LLNL radar operation.	50
Figure 23: Illustration of multipath.....	51

1. Introduction

Transportation of cargo by sea is the most popular shipping method used throughout the world, accounting for approximately 90% of the world's international trade [8]. This translates to 200 million cargo container shipments per year, with over 10 million of these containers being shipped to the United States [9]. During a shipment, the container is moved between ships, trains, and trucks, passing through many different handlers and shipping organizations. Cargo containers themselves are simple, unsecure devices; they are made of relatively thin corrugated steel that can be easily cut and have doors that can be unhinged and removed.

Each year, thousands of cargo containers are broken into during shipping, costing billions of dollars in lost and damaged goods. In addition to removing its contents, intruders can also add unwanted and dangerous materials to a container, posing a threat to National Security. The possibilities of cargo container break-ins require that the containers go through check points at which they are physically searched. These searches often require the opening of the container, unloading and inspecting all cargo, and then loading the container and resealing it. The searches are done purely through manual labor, making the process extremely costly and impractical [9].

Because of the high costs of break-ins and inspections, many security devices have been developed to ensure the safety and detect the tampering of cargo containers. Many of the mechanisms involve more intricate door locks and electronic seals and are able to add a degree of security to the containers. None of the present solutions, however, are reliable and practical enough to eliminate the necessity for frequent inspection of cargo containers. While the door is a major weak point in the security of a container and

the addition of a tamper detecting seal and special lock may help protect it, many intrusions involve the removal of the door entirely or the cutting away of a wall. Most available security devices fail to detect such break-ins. Some “smart” cargo security devices exist, employing a variety of sensors to detect intrusion, however, many of these systems use technology that is unreliable in detection of intrusions or are impractical due to power consumption, complexity, and/or size. The shipping industry is in need of a reliable, unobtrusive, low-cost, low-effort cargo security device (CSD).

Over the last two decades, Lawrence Livermore National Laboratory (LLNL) has been developing an ultra low-power impulse radar capable of detecting objects and motion within a short to medium range. The radar has been used in many different applications and is a proven reliable technology. Due to its past uses for intrusion and motion detection, the LLNL micropower impulse radar is a top prospect for a sensor technology used in a cargo security device. This paper discusses a CSD developed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory for a sponsor consisting of various government agencies concerned with the secure shipment of cargo in, out, and within the U.S. Because of the sensitivity of the project to national security, some of the details have been omitted.

1.1. System Problem Statement

The overall goal of the cargo security sensor stated by the sponsor is to monitor and report the tamper state of a cargo container reliably and with minimal impact on commerce. This goal is to be reached by the automated and inexpensive detection and reporting of the change in the tamper state of a cargo container that is initially known to be secure, using a cargo security device (CSD) consisting of a battery powered impulse

radar that monitors the cargo container volume, logs intrusions, and reports to a remote reader. The critical requirements of the sensor are that it enables rapid and inexpensive approval for border crossings and has high reliability, specifically greater than 95% probability of detection per door event and less than 4% probability of combined false alarm and critical failure per trip, including not alarming due to container motion or door racking. Important requirements of the CSD are that it has a compact form factor, does not affect the loading, unloading, or sealing of the container, withstands environmental hazards (temp, humidity, vibration, shock, ice, fungus, radio communications, lighting, static discharge, etc.), is able to send the tamper status of the container automatically to a portable reader through easy to use software application, is easily maintainable, can be remotely monitored within a range of 3 meters, and is economically viable for private sector stakeholders. Other desirable goals of the CSD are that its cost is minimized, is battery powered with sufficient battery life, and improves border security.

1.2. Overview

This paper describes the use of the LLNL micropower for the detection of shipping container intrusions as well as the design of a low-power, low-cost cargo security device. The project was to design a mother card consisting of support hardware for the radar which could be mated with the existing LLNL radar daughter card to meet the afore mentioned requirements. In Section 2, currently available systems and technologies for cargo security will be discussed. A more in depth description of the LLNL micropower impulse radar and how it operates is presented in Section 3. The design of the CSD including the selection of parts, technologies, and the firmware is covered in Section 4. Section 5 describes the intrusion detection algorithm which utilizes

the various sensors in order to accomplish the sponsor's desired detection and error rates. Finally, Section 6 provides the results of the project by reporting the power consumption of the device and its intrusion detection abilities through the results of many different test cases.

2. Existing Technologies

The most common currently available cargo security devices are tamper-evident seals. These seals are placed on the container door's latch and must be removed in order to open the door. Once the seals are locked, they must be destroyed in an unfixable way in order to undo the seal and open the door. Therefore a broken seal indicates the opening of the container's door and can be used to know if the door has been opened before it should have. Most tamper seals have visible and non visible signs that the seal has been broken. While these devices are capable of exposing an unwanted breach through the opening of a container's door, it is not able to detect any other type of intrusion. The intruder could cut into the container without breaking the seal, and if the evidence of cutting the container goes unnoticed then the intrusion will not be detected. Furthermore, these seals are only capable of exposing an intrusion that has already happened. They are not able to notify the shippers as the break-in is occurring and in most cases cannot tell the inspectors when the intrusion occurred.

In more recent years, a select number of "smart" cargo security systems have come to the market, employing several different sensors to detect intrusions. The goal of these systems is to provide better and more complete detection of intrusions not only involving the container door, but also the walls, roof, and floor. Many of the smart devices can use wireless communication to notify shippers when an intrusion is occurring as well as allow inspectors to check intrusion data without opening the container. The sensors used across different security devices include passive infrared detectors, ultrasonic sensors, microwave based sensors, and tomographic motion detectors.

2.1. Passive Infrared

Passive infrared detectors sense a change in heat radiated by object in an environment. The sensor settles to the environment and detects changes in heat which are usually caused by the introduction of a person's, or animals, body. Although this technology has been used in motion detectors for years and is fairly reliable, it requires that there is an unobstructed view of the area that needs to be guarded. Shipping containers are typically full of cargo, blocking the sight of a passive infrared sensor and therefore making it ineffective. The sensor by itself has a narrow field of view, therefore fresnel lenses or segmented parabolic mirrors are used to focus infrared energy in different directions. Figure 1 shows a typical detection pattern that is the result of the lens/mirror; note the gaps in the pattern. These gaps are areas that the PIR sensor is unable to detect motion in because of the way in which the lens focuses the IR energy. As the distance from the sensor increases, the gaps become wider, and the PIR is more likely to fail to detect motion. Passive infrared sensors are also sensitive to heat and lighting changes in the environment. While lighting changes may not be an issue inside cargo containers, relatively swift changes in environmental temperature can be. PIR sensors are low power devices and therefore able to be used in a battery powered system.




Figure 1: PIR sensors detection pattern [10].

2.2. Ultrasonic

Ultrasonic sensors measure the distance to an object by emitting acoustic waves at high frequencies and measuring the wave that returns to the sensor after being reflected of the object. If the object is in motion, then the frequency of the return wave will be different than that of the emitted wave, known as the Doppler shift. The operation of an ultrasonic device is depicted in Figure 2. These sensors have the benefits of being low cost and dependable, making the overall system inexpensive and reliable [4]. Although ultrasonic sensors do not have coverage gaps, the sensors are prone to the same issues as PIR devices, requiring line of sight to detect objects as well as being affected by noise in the environment. Furthermore, ultrasonic sensors have a narrow detection area, often making it necessary to have multiple sensors at different locations and viewing angles. However, operating multiple ultrasonic sensors near each other introduces the issue of interference between the sensors [6]. Research of the use of ultrasonic sensors for

occupancy detection done by Kianpisheh, Mustaffa, Limtrairut and Keikhosrokiani and discusses the negative effects of temperature changes on the reliability of the sensors measurement; changes in temperature can cause the sensor to falsely detect objects [5].



Figure 2: Ultrasonic sensor operation.

2.3. Microwave

Microwave sensors are similar to ultrasonic; the sensor emits microwaves and listens for a return wave, looking for a Doppler shift in the return frequency due to object motion. Figure 3 illustrates the operation of a microwave sensor, showing a coverage area where the emitted beam will reflect off objects and return to the sensor. Using electromagnetic waves in the microwave frequencies has the benefit of being able to penetrate many materials that are not conductive, no longer requiring line of sight to detect object, and being less susceptible to noise from the environment. Microwave sensors are, however, susceptible to interference from other electronic devices operating

at or near the same frequency [7]. Like ultrasonic sensors, microwave sensors have a narrow field of view and are typically used for detecting motion in long, thin regions [10]. Historically, microwave sensors tend to consume a large amount of power, which isn't ideal for a battery powered device.




Figure 3: Microwave sensor [10].

2.4. Tomographic Sensor Networks

Tomographic motion detecting systems use a mesh network of radio transceivers, as seen in Figure 4, which detect the changes in the radio waves between each transceiver. Like microwave sensors, tomographic motion detectors have the ability to “see” through many different types of material and cover a wide area even with obstacles. Also like microwave sensors, this technology has high power consumption. Each node in a tomographic sensor network must have a sensor for detecting the motion of nearby objects, a processor for acquiring the sensor data and initiating communication, and communication hardware to allow for the sharing of sensor data with the rest of the network. The communication hardware can either be wired (PCI, USB, RS232, etc.), or wireless (Bluetooth, WiFi, ZigBee, etc). Requiring communication lines to be routed to

every sensor in the network makes this technology invasive and more difficult to set-up. As mentioned by Goh et al. in their research on wireless sensor based tomography systems, using batteries to power each wireless sensor node is illogical due to the limited lifetime and would require frequent replacing [1]. Ali Elkateeb points out that the use of a general processor for acquiring sensor data and handling communication does not align with the system's needs and therefore is inefficient [2]. Instead, he proposes the use of soft-core processors designed specifically for the sensor nodes to be implemented on FPGAs. While this reduces power consumption, it increases design complexity and cost. Gungor and Hancke propose the use of energy harvesting techniques to power the sensor nodes, but this concept once again adds complexity to the system and is most likely ineffective inside a cargo container's environment [3]. Overall, tomographic motion detecting systems require a large amount of hardware, increasing system power consumption. The amount of hardware, in addition to the need for a mesh network to be setup in every container, also eliminates this technology as a low-effort solution.




Figure 4: Tomographic sensor network [15].

3. LLNL Impulse Radar

The CSD's functionality is built around the ability of LLNL's micropower impulse radars. This technology is what separates this CSD from other options that are already available. This study's main goal is to implement the micropower radar into a security device and compare its detection ability to that of current sensor devices currently being used. The micropower impulse radar is an ultra wide-band sensor which transmits and receives narrow pulses of energy across many frequencies, the width of the band being determined by the duration of the pulse. With a wide band of frequencies being reflected off objects and received by the radar, detailed information about the environment can be obtained. Also, since various materials reflect and absorb different frequencies of radiation, sending a broad band of frequencies gives the radar excellent signal penetration. This enables the radar to "see" through most non-conductive materials and not require line-of-sight to intruding objects, a problem that PIR and ultrasonic sensor have. A side effect, which has become one of the greatest features of the micropower impulse radars, is that the power consumption of the radar is very low due to the short time duration of the transmitted pulses. Typically, ultra-wideband radars emit pulses with durations between .2 and 1 nanosecond [12]. These properties of the impulse radars are in contrast to traditional radars which transmit energy in a single frequency or multiple frequencies over time. This leads to larger pulses and therefore higher power consumption and less information about the environment per unit time. In comparison to the PIR, ultrasonic, microwave and tomographic network technologies discussed in section 2, the impulse radar's power consumption is significantly lower.

The impulse radar uses the range gating concept for detecting objects at specific distances. The range gating method sets an interval that occurs some amount of time after a pulse is emitted for which the return signal is captured. Since the distance of the object is related to the time the return wave is detected, shifting the time interval, called a range gate, makes the radar look for objects at different distances. Figure 5 shows the major components of the micropower impulse radar as well as the general operating principle of generating and emitting short pulses, receiving the return echoes, and inspecting the return at specified times according to the set range gate. The radar is able to open a single range gate per pulse, allowing it to greatly reduce the reception of unwanted noisy signals. To further improve the performance of the radar in a noisy environment, many samples of the return signals are averaged together. Also, the period of the emitted pulses is not constant, but is randomly changed by a randomization signal. The random periods of the pulses along with any outside noise effectively averages out all noise. The random period has the added benefit of forming the emitted radar signal as to appear as random noise to other radio devices, making the operation of the radar difficult to detect [13]. In addition, many impulse radars can coexist in the same environment since each individual radar will have a unique signal to determine when to emit and receive pulses.

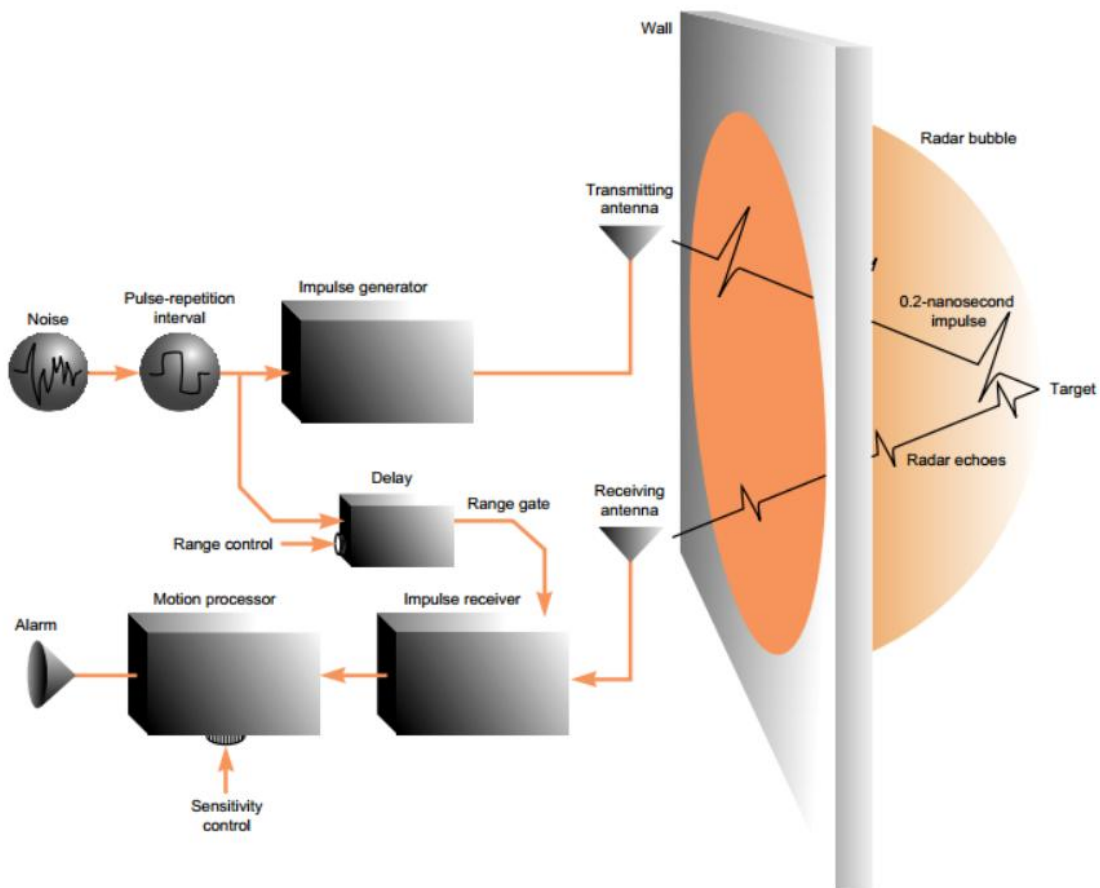


Figure 5: LLNL micropower impulse radar block diagram [13].

By looking at the difference in frequency of the signal at a specific range gate over time and applying a threshold to that difference, the radar can act as a motion detector, sensing moving objects at the distance set by the range gate. Sweeping the range gate, which is accomplished by incrementally increasing the delay of the start time of the range gate for each pulse (controlled by the delay block in Figure 5), will make the radar sense any motion within a maximum distance. When paired with an omni-directional antenna, the radar creates a “sensing bubble” (orange sphere in Figure 5) in which any motion within a certain distance triggers an alarm. Because of the wavelengths of the signals emitted, the radar can detect objects as small as 15 centimeters located 15

centimeters to up to 50 meters away [13]. Range and gain settings allow for the sensitivity and the diameter of the sensing bubble to be changed. Figure 6 illustrates how moving objects at different distances from the radar affect the return signal with reference to time, revealing how the range gating method of thresholding the signal at specific times is used to set the detection distance. The radar outputs an analog signal of the pulse return as well as a digital alarm signal used for detecting motion if motion is present. It also has an I²C interface for setting the range and gain values.

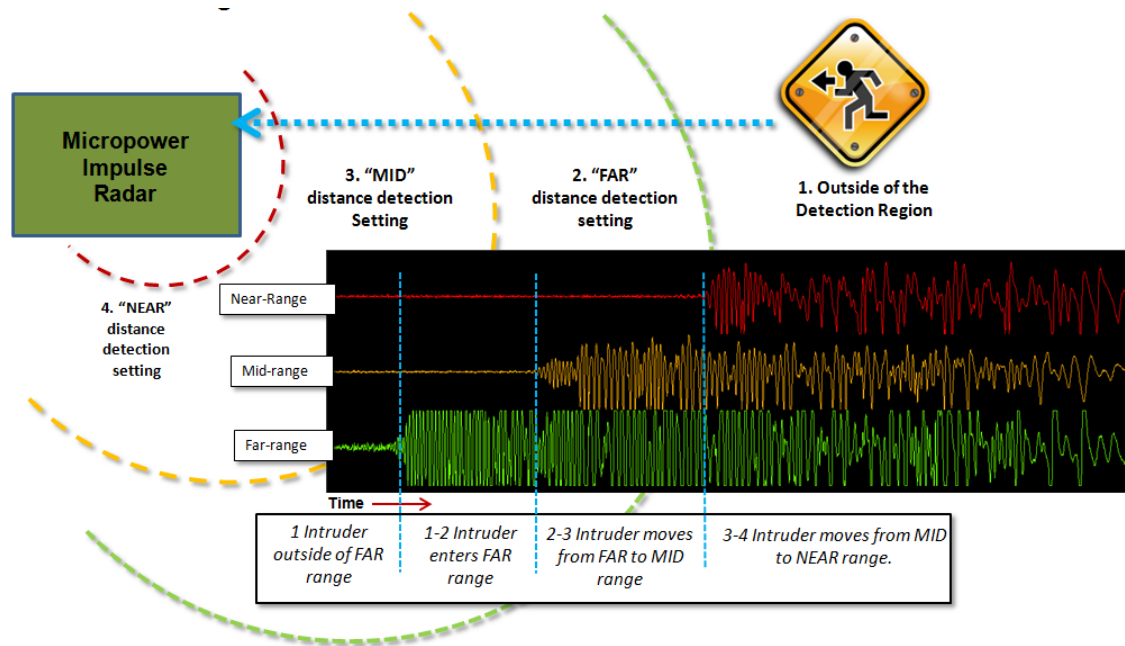


Figure 6: Operation of LLNL radar as a motion detector.

4. Design

Based of the requirements for the CSD, a general parts list and initial hardware system plan was created consisting of the components shown in Figure 7: a microprocessor for analyzing and manipulating sensory data, an LLNL impulse radar for primary detection of intrusions, a passive infrared sensor for increasing reliability of intrusion detection through redundancy of motion detection, a vibration, shock, and movement sensor for monitoring the motion of the container in order to rule out false alarms, a flash memory for storing intrusion data as well as any important system information, and a wireless communication device for remote communication with the CSD by an external reading device. Figure 8 shows the radar components assembled on a daughter card measuring approximately 1 inch by 1 inch, making it one of the smallest radars. Between its motion detecting capability, low power consumption, low cost (approximately \$10) and small size, the LLNL micropower impulse radar is a great fit as a motion sensor for the CSD. This daughter card is a previous product created by LLNL for use in other projects. The daughter card connects to the motherboard, which contains the all the other hardware described above and shown in Figure 7.

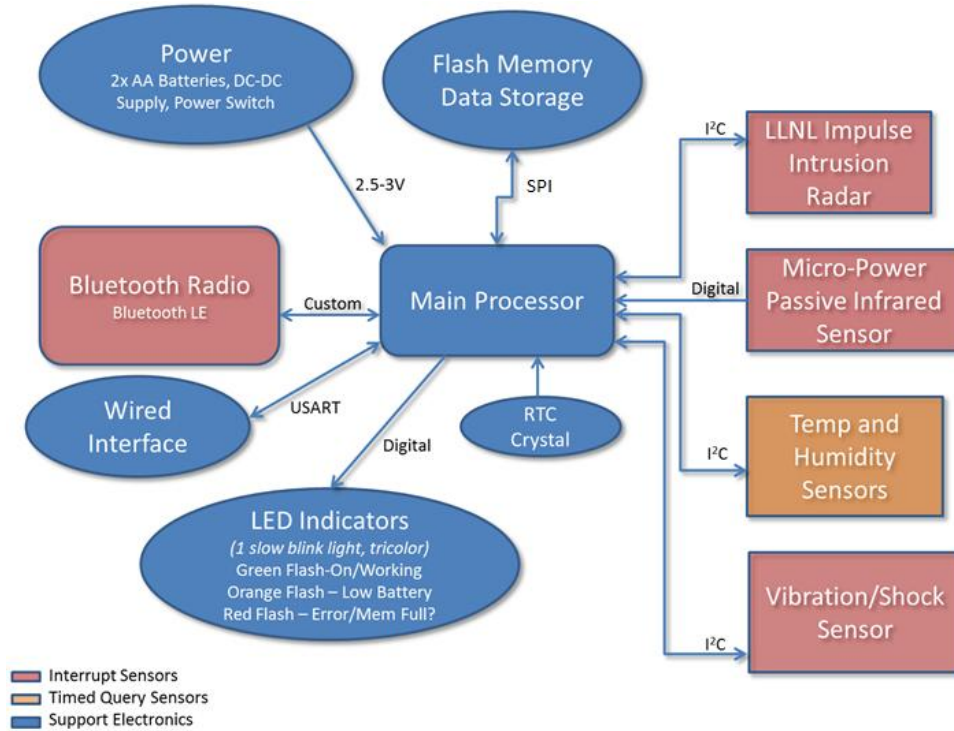


Figure 7: System block diagram for CSD.

4.1. Selection of Parts

The initial hardware plan covered the critical needs as well as a few of the important goals of the CSD. Guided by the remaining important and desired requirements of the CSD, more specific parts were selected: the Atmel ATxmega34A4U for the microprocessor, the Panasonic EKMB1303112K for the passive infrared sensor, the Freescale MMA8452Q accelerometer for a vibration/shock sensor, the Atmel AT45DB161E for the flash memory, the Honeywell HIH-6131 for a temperature/humidity sensor, and the Nordic nRF8001 Bluetooth LE for wireless communications. The guiding forces behind the selection of mentioned components, beside that they were able to meet the critical and important goals, were that they be ultra low-power, low-cost, small in physical size, and readily available. Due to the typical use

case of the CSD, the sleep mode power properties of the components were considered, as opposed to the active or standby power modes, since the device will be spending most of the time in an idle state waiting for an alarm or to communicate with a reading device.




Figure 8: Assembled micropower impulse radar.

4.2. Passive Infrared Sensor

A passive infrared (PIR) sensor will be included in the CSD design for redundancy of motion detection as well as in aiding in ruling out false alarms that the LLNL radar is prone to. Although most PIR sensors are relatively identical in all areas, the new Panasonic EKMB line of sensor not only offer long distance detection, but also are the lowest power devices on the market. The PIR's proprietary internal circuit allows the sensor to go into a sleep mode until motion is detected. In this sleep mode, the sensor has a current draw as low as 1 μA . The EKMB line was designed specifically for occupancy detecting systems that are designed to be powered by photovoltaic cells or small battery cells. Due to the length of the shipping containers, the CSD requires the

longer distance sensor, specifically the EKMB1303112K, with a range of 12 meters. This particular sensor has the same sleep current of 1 uA, but a standby current between 6-12 uA. The sensor has a digital output that takes approximately five seconds to settle after detecting motion.

4.2.1. Vibration, Shock, and Movement Sensor

In order to rule out false alarms due to the movement of the container, the CSD must have a means of detecting when the container is in motion. The Freescale MMA8452Q is a low-power digital 3-axis accelerometer. In its typical mode where the sensor is sampling the acceleration in each of the three axes, the sensor draws a maximum of 165 uA at its highest sampling rate. The feature that makes this particular accelerometer ideal is its ability to generate interrupts when certain types of motion occur. The MMA8452Q has embedded DSP functions in order to detect freefall/motion, transient (shock/vibration), orientation, shake, and single/double tap. Two of the five detectable types of motion can be configured to generate interrupts on two of the sensors pins. Control registers for each type of motion can be modified to change the sensors sensitivity, for example, a transient threshold register controls how violent of a shock causes an interrupt to trigger. By configuring the motion and transient detection functions to generate interrupts, the CSD will be able to detect when the container is in motion and/or is vibrating. To further reduce the power consumption of the accelerometer, it can be put into an auto-wake/sleep mode in which it lowers the sampling rate until a motion event is detected. This mode fits the CSD's use case since it will be in a low power mode until an event occurs, at which point it will wake and gather information from all its

sensors. The sensor communicates using I²C, allowing it to share the serial bus with other sensors in the system.

4.2.2. Processor

There are many microprocessors that support the features needed to meet the requirements of the CSD. However, due its flexibility, low-power options in various sleep and power-down states, and low-cost, the Atmel ATxmega34A4U microprocessor was selected. This Atmel processor can be operated at a voltage as low as 1.6 V and can be put in a deep power-down mode in which the processor will only draw 1 uA. In its deep power-down mode, the processor disables its system clock and will only be woken up by external interrupts or an interrupt from its watchdog timer. Since the system clock is disabled in the deep power-down mode, it is important that the processor supports many asynchronous interrupts, allowing the processor to enter and stay in the power-down mode for most of the time, and only be awoken when a sensor or external component requires data to be serviced.

The processor also includes a real-time clock module, which reduced the need for a separate component to keep track of the time that will be used for time stamping data (mainly alarms) in memory. The removal of an extra component for the system reduces overall cost, power consumption, and size. Since the system clock is disabled in the deep power-down mode, an external 32.768 kHz crystal is needed to keep time during the idle state. This does not, however, add hardware to the design since a 32.768 kHz rail-to rail oscillator is needed for the nRF8001 Bluetooth chip. The ATxmega34A4U allows for the digital output of the external clock on one of its GPIO pins, allowing the processor and the Bluetooth chip to “share” the oscillator by attaching the crystal to the ATxmega and

then connecting the nRF8001's clock input to the GPIO of the ATxmega with the digital clock output.

Encryption of the data saved on and transmitted from the CSD may be required to ensure security; the ATxmega34A4U contains an AES/DES encryption/decryption hardware module. Using hardware to encrypt/decrypt data eliminates the need for the implementation of a costly (in processing time and therefore power) encryption/decryption algorithm in software.

The flexibility of the processor is seen through its many general purpose I/Os, support for common communication protocols, and diverse clocking options. This flexibility allows for an open design which supports future changes, modifications, and additions to the system without having to change the processor and consequently the firmware.

4.2.3. Memory

The important features to use for selecting the flash memory were size and power consumption. The Atmel AT45DB161E is a 16Mbit flash memory, having more than enough memory space to store alarm logs, system information (such values for identification and communication), and allowing future features that may require the use of non-volatile memory. More importantly, however, is that this flash memory can operate at 2.3 V, one of the lowest voltage flash chips available, and draws only 3uA while in sleep mode. Due to the typically higher operating voltages of flash memories, this component was the deciding factor in the voltage level at which the system would require. The system could now operate at 2.5 V (the impulse radar's operating voltage), however a system voltage of 2.7 V was used (regulated from the voltage of two AA

batteries in parallel, so approximately 3 V). This memory communicates using SPI, which is a highly supported serial protocol and allows for multiple slaves to be connected to the same wire bus (as long as each slave has its own select signal). Other components in the design also use SPI, allowing for these components to share lines, reducing system complexity and allowing for flexibility.

4.2.4. Temperature/Humidity

A sensor to measure and record the temperature and humidity of the interior of the cargo container was added to meet the CSD's desired requirement of sensing the environment of the container. Implementing this sensor into the design also proves that environmental sensors using common communication protocols can be easily added to the CSD in the future, allowing for a flexible system which can be modified to meet the user's specific needs. The temperature and humidity sensor selected was the Honeywell HIH-6131 due to its low operating voltage and sleep-mode current; 2.3 V and .6 uA respectively. Most humidity sensors have a long settling time, some taking several days, after extreme conditions, such as those during soldering or extremely moist environments, during which the readings are highly inaccurate; the Honeywell sensor has a settling time of only 5 hours. The HIH-6131 uses I²C for communication, allowing other components in the system which use I²C to share a bus. The sensor is also available in an ultra small form factor, is low-cost, and is industry leading in accuracy.

4.3. Bluetooth

Bluetooth was selected as the wireless technology for communication between the CSD and an external reader. The motivation for choosing Bluetooth was that it is a well supported by mobile devices including laptops, tablets, and phones, making the later

development of a CSD reading device flexible. Bluetooth also supports the range of operation desired for the CSD and reader and is one of the lower power wireless standards. A concern for the CSD was that due to the fact that it would be housed inside a metal container, wireless communication to devices outside of the container would not be effective, however, initial tests of Bluetooth devices operating across the container walls proved that it was able to consistently communicate within 5 meters of the container, especially when the power was set to 0dBm.

Although Bluetooth is one of the lowest power options for wireless communication, there were still concerns for its power consumption since the CSD is desired to last at least 3 months on a single “charge” (set of batteries) and will not have a large battery due to the goal of making the device physically small. In order to reduce power consumption of the wireless communications, the Bluetooth hardware could be put into a low-power sleep mode until communication is needed. This would require a way of waking up the hardware once communication is desired, which is signified by a reader device trying to bond with the CSD Bluetooth hardware. A solution for discovering the reader would be to have a “sniffing” circuit which would detect power in the Bluetooth band. Once the power detected in the band is above a certain threshold, the circuit would signal the system to wake up the Bluetooth hardware and bond with the reader. Several possible circuits to accomplish this were simulated, however none were consistent in noisy environments (which would be the case due to the ultra-wide band impulse radar) and the circuits would add unwanted complexity to the CSD.

A new solution to the wireless power consumption issue was found in the recent development of Bluetooth Low Energy (BLE), which is part of the Bluetooth 4.0

specification. Although BLE was introduced to the core specification in 2010, the first device to implement it did not arrive until late 2011 with many more supporting devices arriving by early 2012. More importantly, BLE development hardware was not widely and readily available until mid 2012. Although as of 2012 BLE was not as supported as other wireless protocols, its power saving abilities and high flexibility in connection and advertisement settings made it a top prospect. Although it was not heavily supported, BLE was incorporated into the core specification and therefore was likely to become well supported over time.

4.3.1. Power Consumption

Bluetooth LE achieves its low power consumption by greatly reducing the duty cycle of the communication and therefore the amount of time the radio is on compared to other Bluetooth specifications. Figure 9 depicts a BLE device's typical behavior when connected to a device; the device is mostly in an idle state, only becoming active periodically to transfer data. When a BLE device is advertising, its behavior is very similar to when it is connected; only turning on periodically to send an advertisement packet. During the advertisement, BLE uses 3 channels with the radio being on for only .6 to 1.2ms per total advertisement period. Furthermore, BLE allows for the advertisement interval to be set between 20 milliseconds to 10 seconds, enabling a choice between lowering power consumption and increasing connection response time. In comparison, other Bluetooth specifications typically advertise on 32 channels, requiring the radio to be on for 22.5 milliseconds during the total advertisement period. Bluetooth does not allow for the setting of the advertisement interval. BLE also enables the connection interval to be set between 1.25 milliseconds to 4 seconds when a connection

is established with a peer device. This allows for the reduction of power consumption during the connected state in sacrifice of data rate and is meant for use cases when a device will be connected for long periods of time. Since BLE devices will spend the vast majority of their time in a sleep mode, the sleep current will become the main factor in overall power consumption. As the advertisement and connection intervals are lengthened and practically become infinitely large when compared to the radio's active time, the peak power consumption (during transmit and receive) no longer increase the devices average power consumption.




Figure 9: Current profile of BLE device in connection state [14].

Due to the flexibility of the advertisement interval and the effect that increasing the interval has on reducing the power consumption, in addition to BLE device's already low consumption, the BLE device can be left on and constantly advertising without significantly reducing the battery life of the CSD. This would allow an external reading device the ability to discover and connect to the CSD at any time. Since a connection

delay time of 10 seconds is allowable, the advertisement interval can be set to its maximum possible time and the power consumption will be minimized. The connection interval can be set at its lowest setting, allowing for the fastest data rate and will not significantly increase the CSD's overall power consumption since it will spend most of its time in the advertising state waiting for an external reader to try to connect.

Because of the periodic nature of the advertisement and connection states, BLE can have a longer connection setup time and a lower data rate. With the advertisement interval set to its largest value, it is possible that a connection will not be established for up to 10 seconds while the discovering device waits for the next advertisement packet. Due to the small packet sizes and lower packet frequency because of connectional intervals, application throughput using BLE is around 270 kbps. The original purpose for BLE was to allow devices that will be transferring small amounts of data (around 20 bytes) at a low frequency.

Specifically, the Nordic nRF8001 was selected for the BLE hardware. The nRF8001 includes the radio, link layer and host stack all in one integrated chip. The chip is low cost, low power (includes DC/DC converter to increase efficiency by 20%), and is physically small (5 millimeters by 5 millimeters). The nRF8001 is a slave device and uses a serial interface, named the Application Controller Interface by Nordic, which is almost identical to SPI except for an added control line that the slave device uses to indicate when it has data to send to the master controller. The typical use of the nRF8001 can be seen in Figure 10, where the MOSI, MISO, and SCK lines are analogous to those in SPI and REQN is the same as the SS line. The RDYN is an added line that indicates data is ready for transfer to the application controller. The nRF8001 offers a wide range

of advertisement and connection interval settings, allowing fine tuning of its power consumption versus responsiveness.

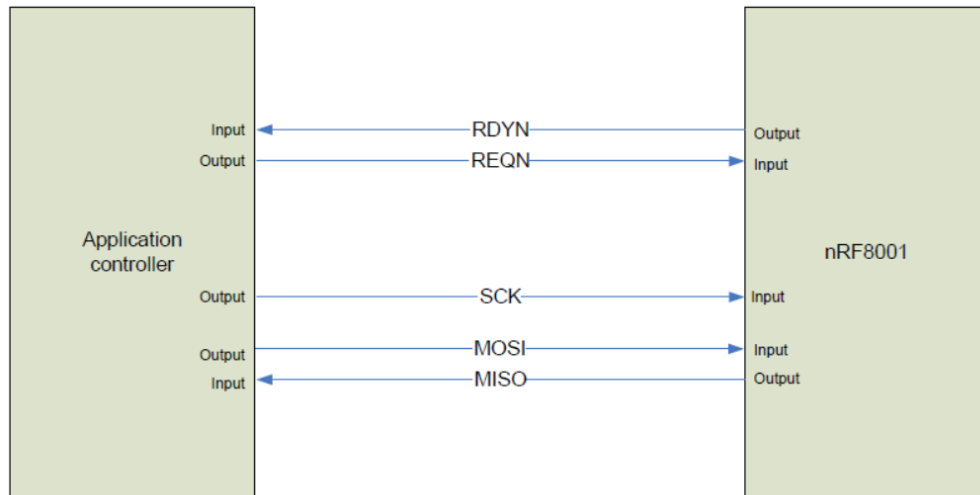


Figure 10: Connection between application controller and nRF8001 [14].

During the advertisement interval the nRF8001 broadcasts its identity to any listening devices that are trying to connect in advertisement packets containing, among other information, the specific nRF8001's MAC address, device name, and connection settings. The advertisement interval determines the amount of time between each transmission of an advertisement packet. When a device attempts to connect to the nRF8001, it will send a connect packet that contains the desired connection interval, which sets the period between transmissions of data packets. The graphs in Figure 11 show the relationship between the advertisement and connection interval length and the nRF8001's current consumption, revealing diminishing returns on power savings once the intervals are set approximately above half a second. As seen on the plots, the average current draw for an advertisement interval of about .5 seconds is roughly 30 uA, further

supporting the idea of leaving the nRF8001 on and advertising during the system's idle state without harsh effects on the CSD's power consumption.




Figure 11: Current consumption of nRF8001 versus connection and advertisement interval settings [14].

4.3.2. Communicating with BLE and the nRF8001

Communication using BLE is based on the GATT (Generic Attribute) protocol. In the GATT protocol, devices take on one of two roles: client and server. The client is the device that will be requesting and receiving data from the server while the server handles the client's requests and responds with the desired data. Data is transferred between the server and client by setting and reading characteristic values. Characteristics are created to represent the data that will be transferred and are stored on the GATT server. For example if a heart rate sensor will be sending beats per minute, a characteristic will be created that stores the heart rate in a single byte on the GATT server. The client will request an update on the heart rate and the sensor will set the characteristic value. Multiple characteristics can be grouped together to create a service. Each characteristic can have a descriptor which describes the use and context of the characteristic. To communicate, the client requests the server to update the characteristic value and, once the server has done so, then reads the characteristic value from the server. A service was created for the CSD that allowed the client (a reader device) to send requests in the form of a single byte integer which the server (the CSD) would receive and respond to accordingly.

The application controller is concerned with the GATT layer of the BLE stack, which is seen on the top of the BLE stack in Figure 12. The layers seen in Figure 12 show that the BLE stack is vastly different from those of other Bluetooth standards. For this reason, a BLE device is not compatible with devices using other versions of Bluetooth unless it is a dual-mode chip, meaning it also contains the Bluetooth stack necessary for communicating with other versions. For the nRF8001, the ACI allows the application

controller to communicate with the GATT layer; however this is done in an indirect way using what Nordic has coined as service pipes. Instead of setting the values of GATT characteristics, the application controller writes to service pipes which are assigned to specific characteristics, like the pipes shown in Figure 18. In order to read and write to the characteristics, a pipe must be opened for each desired operation on each characteristic value. For example, Figure 13 represents a GATT service which is defined by two characteristics. The first characteristic will be read from in two different ways: a read that acknowledges the reception of the data, and a read that does not send an acknowledgment. Since these are two different operations, two separate pipes will be opened. The second characteristic only supports being written to and therefore only requires a single pipe.




Figure 12: BLE stack [14].

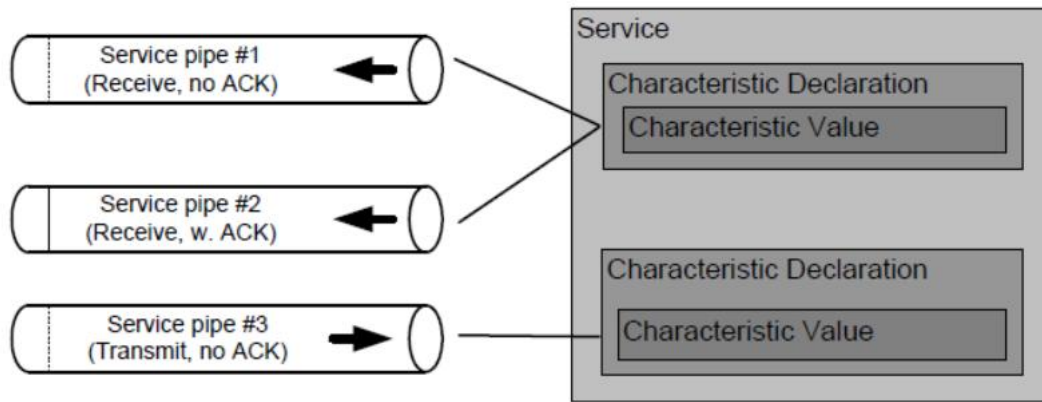


Figure 13: Illustration of service pipes used by the nRF8001 [14].

Before the nRF8001 can be used for communication, it must go through a setup process. Setup messages are created using Nordic’s nRFG Studio, an application which is used to configure all the wireless radios which Nordic sells. It is in nRFG Studio that new characteristics are created and service pipes are set for communication. Device attributes are also set using the application, such as the setting of the device name, advertisement interval and connection interval. nRFG Studio outputs a text file containing the byte values that will form the setup messages for the nRF8001. Each message created by the application contains a setting for configuring the radio and supporting hardware. The messages also contain the information needed to create the characteristics and services used for communication between the client and server, as well as the service pipes that will be associated with the characteristics. The messages are sent one-by-one from the application controller to the nRF8001 as shown in Figure 14, with the reception of each setup message acknowledged by a `CommandResponseEvent` from the BLE device, indicating that it is ready for the next message (the responses from the device in Figure

14). Once all the setup messages have been received, the nRF8001 issues a response indicating that the transaction is complete and then sends a DeviceStartedEvent.

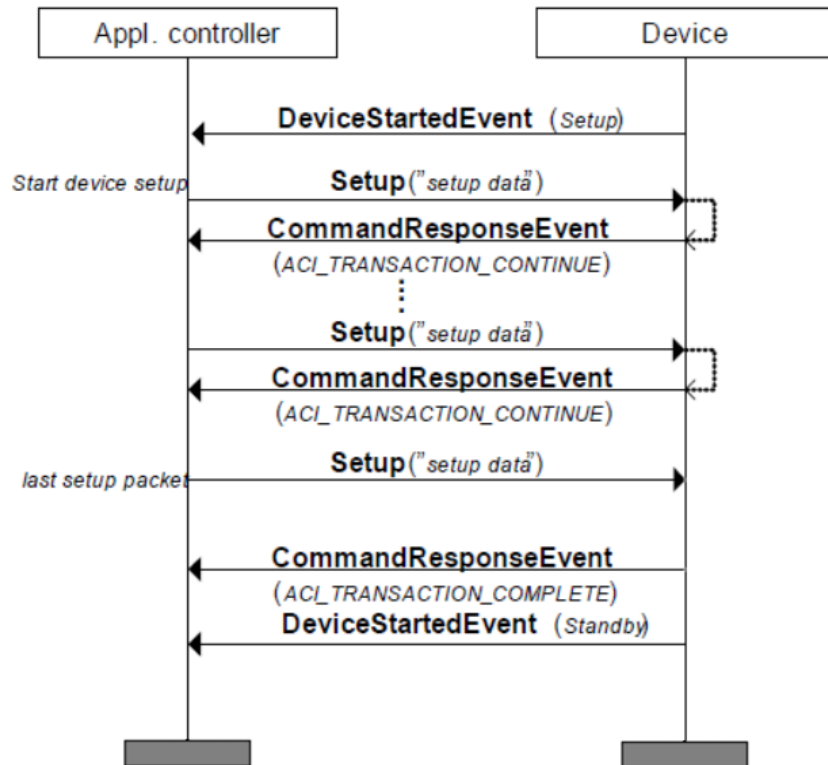


Figure 14: nRF8001 setup process [14].

After the nRF8001 has completed its setup sequence, it is controlled by sending system and data commands and responds with system and data events. Commands and events are represented as data packets containing a header for identification, payload, and a checksum. The main events used are the:

- DeviceStartedEvent
- CommandResponseEvent
- ConnectedEvent
- PipeStatusEvent
- DataCreditEvent
- DataReceivedEvent

- PipeErrorEvent
- DisconnectedEvent

The DeviceStartedEvent tells the application controller whether the nRF8001 has successfully powered on, what operation mode it is in, the available space in the data buffers, or an error code describing the failure that occurred during startup. The CommandResponseEvent is the reply from the nRF8001 after the application controller has issued a command. This packet contains confirmation or error codes depending on the outcome of the command. A ConnectedEvent tells the application controller that a peer device has connected, the information of the peer device (such as MAC address and clock accuracy), and the connection settings (connection interval, latency, and timeout). After a connection is setup, a PipeStatusEvent occurs which describes which pipes are open for transferring data on as well as the direction (read/write) of the open pipes. After data is successfully sent through a pipe, a DataCreditEvent will be issued, indicating that more room is open on the data buffers. When data is received by the BLE device and is ready for transfer to the application processor, a DataReceivedEvent will be issued containing the data and the pipe it was received on. If the data is unsuccessfully sent, then a PipeErrorEvent will occur, containing error codes indicating the type of failure. When the connection is ended for whatever reason, a DisconnectedEvent will be sent to the application controller describing the reason for the disconnection. The application processor typically uses the Connect and SendData commands to control the nRF8001. The Connect command tells the Nordic chip to go into its advertising state with a set timeout and advertisement interval. The SendData command is used to transfer data on a specific pipe, and therefore set the corresponding characteristic value. A typical scenario between the application controller and the nRF8001 can be seen in Figure 15 where the

nRF8001 issues a DeviceStartedEvent and the application processor establishes a connection then begins transferring data. Because the nRF8001 operates by sending asynchronous events to the application processor, the best way to handle communication is by creating an interrupt subroutine which will be called when the nRF8001 has an event ready. The processor will receive the event, determine the correct response, then sends a command accordingly.

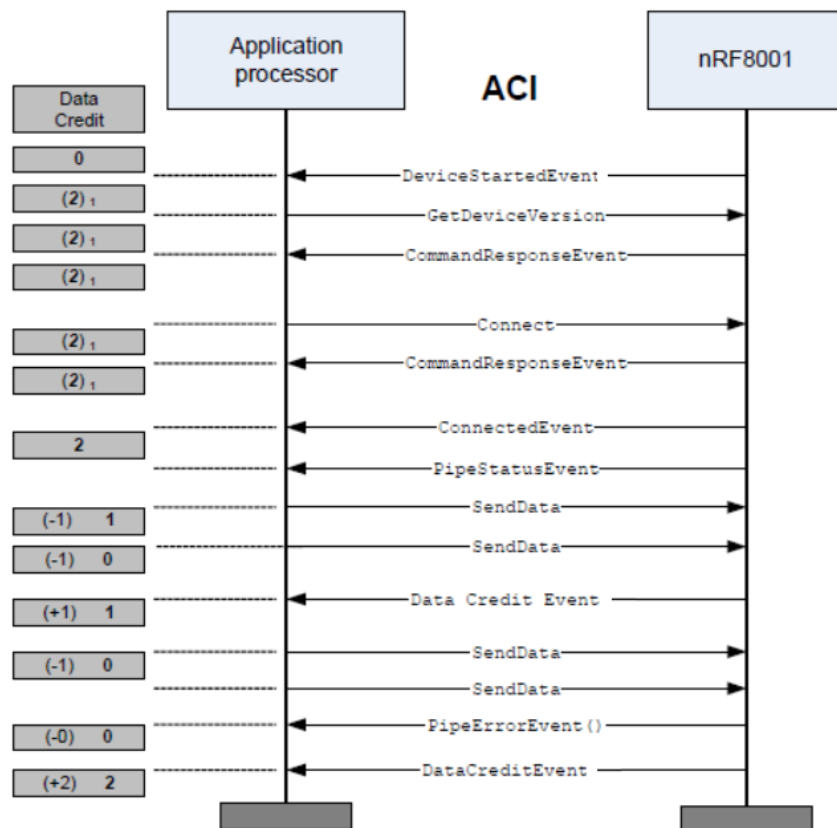


Figure 15: Typical communication between application controller and nRF8001 [14].

The nRF8001 allows for a maximum of 20 bytes of data to be sent in each packet. While 20 bytes is well suited for the transfer of remote sensing data (the original purpose for BLE), it is not ideal for transferring large amounts of data in a small period of time, as

is the case for the CSD's communication with a reading device. In order to send the intrusion log and system information that is stored in the flash memory, the data must be split up into 20 byte blocks and sent packet-by-packet. With a minimum connection interval 1.25 ms (meaning a packet will be sent every 1.25 ms) the maximum data rate of log data is 128 kbps (this is lower than the 270 kbps quoted by Nordic since only 20 bytes out of a packets 41 bytes consist of application data). Although this is a low data rate compared to many other wireless technologies, it is perfectly fine for the CSD since it will only be transferring a small amount of data at a time (a log entry consists of 16 bytes, meaning that the CSD can transfer 1000 intrusion logs per second).

4.4. Event Driven Firmware

To reduce the amount of power consumption for the CSD, each component must be operating in a sleep mode whenever possible. Every component except for the nRF8001 and micropower impulse radar can be put in a power down mode while it is not sensing. The accelerometer will sleep until either a strong enough shock or movement is detected, the temperature and humidity sensor will be sleeping until a reading is made and logged in memory, the memory will be in low power mode until there is data to store, the PIR sensor will sleep until motion is detected, and the processor will basically be off until one of the sensors is triggered or a devices tries to connect to the nRF8001. The radar is always on since it does not have a sleep mode and consumes very little power as does the nRF8001 during its advertisement state.

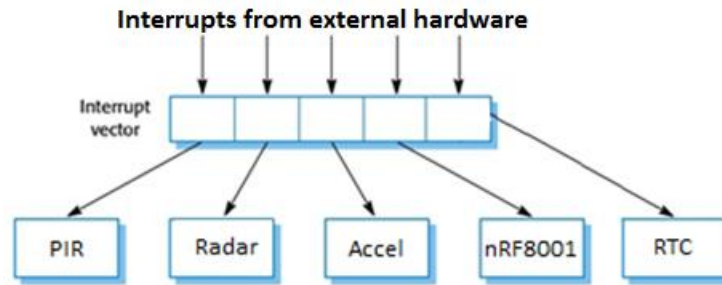


Figure 16: Interrupt routines for sensors and communication.

The firmware is designed to maximize power saving by placing all devices in their low power modes and configuring interrupts and interrupt service routines to handle sensor information when it is ready. As seen in Figure 16, all the external sensors as well as the nRF8001 and real time clock module generate interrupts which wake the processor from its deep sleep state. On startup, the firmware first configures all its IO including serial buses, enables the interrupts for the accelerometer, real time clock, Bluetooth, radar, and PIR sensor, enables the processors low power sleep mode, and configures the accelerometer to generate interrupts on shock and movement. The ATxmega's external interrupts that were chosen for the sensors are all asynchronous, meaning that they do not depend on the clock of the processor so that when the system clock is disabled in the low power mode, the processor will still wake up and service the interrupts. The firmware then reads the system settings and memory status from the flash memory. System settings include values for device identification, Bluetooth radio settings, and the frequency to make temperature and humidity readings. The memory status informs the processor of the amount of memory that is currently being used through the number of intrusion entries. After the system settings and memory status are loaded, the firmware configures the Bluetooth radio. The setup messages for the nRF8001 that are generated by the

nRFGo Studio are read from memory and then sent to the Bluetooth device. After the last message has been sent, the firmware waits for a DeviceStartedEvent from the nRF8001 indicating a successful startup. If the DeviceStartedEvent holds an error message, then the setup is retried until success or after a set number of tries. If the setup is not completed successfully after the set number of tries, then the firmware discontinues the startup sequence and enters an error state denoted by a red flashing LED. Once the nRF8001 has started up successfully, the firmware issues a Connect command to the nRF8001, placing it in its advertisement state so that an outside reader device may connect. Once the nRF8001 is advertising, the processor is placed in its low power mode, waiting for an interrupt to wake it up.

The intrusion detection algorithm is implemented in the interrupt subroutines for the PIR sensor, radar, and accelerometer. Each interrupt routine time stamps when the event occurred and then checks for other events that have happened within the time windows mentioned in the description of the intrusion algorithm. The accelerometer interrupt routine will also add an entry to the log in flash memory if it senses a shock greater than its set threshold, allowing the cargo owner to know if the container was subject to any harsh movements. If any intrusions are detected as a result of the algorithm, then an alarm instance is created in the processors local memory and stamped with the time when the intrusion was detected. Since more events can occur while the time window is open for the intrusion, it is not immediately written to memory until the time window closes. This enables more data to be collected in order to present better and more accurate information about the intrusion. The RTC interrupt routine keeps track of system time and is responsible for carrying out any constant time operations such as

reading the temperature and humidity sensor and the battery voltage level, then writing the values to the log in flash memory. The RTC routine also keeps track of the time windows used for the intrusion detection algorithm. When the time window closes, the current alarm entry held in the processor's local memory has the alarm duration added to it and is written to flash memory and added to the log.

The Bluetooth interrupt routine handles any events that are issued from the nRF8001. The routine first reads in the event packet's data and then decodes which type of event it is. The routine handles three events: `ConnectedEvent`, `DisconnectedEvent`, and `DataReceivedEvent`. If a `ConnectedEvent` is received, the firmware then waits for the passkey to be sent in order to authorize the connecting device for communication. If the passkey is entered correctly, then the connecting device is allowed to stay connected and transfer data. If the passkey is incorrect, a `Disconnect` command is sent to the nRF8001 and the connecting device is denied a connection to the CSD.

Once a device is connected, it may send commands to the CSD asking for data or changing settings. Data sent from the connected device is seen as `DataReceivedEvents` to the firmware. The data held in the packet is decoded and the command is carried out. The connected device may set the CSD's system time, request the transfer of the intrusion log, arm the CSD for detecting intrusions, unarm the CSD, clear the log file, set the radar range and gain, request the device status, or change the passkey. To change the system time, the global variables used by the RTC interrupt routine are modified to match the time sent in the time change packet.

When the connected device requests the intrusion log, the handling of sending the log is deferred from the interrupt routine since it may take some time to complete,

allowing any interrupts, other than those from the nRF8001, to still be handled. When transferring the intrusion log, first the total number of entries is sent followed by each entry, one at a time, along with a sequence number indicating which entry in the log is currently being sent. The sequence number allows for additional error checking during the transmission as well as indication of when the transfer is complete.

If the connected device issues an arm command, the radar, accelerometer, and PIR sensor's interrupts are enabled and any intrusion detection variables, such as current time windows, are zeroed. If the command is to disarm the device, the previously mentioned interrupts are disabled and any pending alarms that have not yet been written to memory are added to the log. The clear entry command causes the number of saved entries variable to be set to zero, meaning that new entries will overwrite the old entries. When the connected device requests the CSD's system status, a data packet is sent which holds flags indicating if an intrusion alarm has occurred, the accelerometer has triggered, and if the battery is low. These values are used to quickly check the CSD and know if an intrusion has occurred. If an intrusion hasn't occurred, then the container can remain sealed and won't require an inspection. If the intrusion alarm flag indicates a break in, then the log file can be downloaded and inspected to further aid in deciding whether to inspect the container or not. The system status also includes the current temperature, humidity, battery voltage, and system time. The change passkey command replaces the current passkey saved in memory with the new passkey sent in the packet.

5. Intrusion Detection

Intrusions are detected using the micro-impulse radar sensor as well as the passive infrared sensor. Each sensor is subject to noise in various environments. In areas which contain medium to heavy electromagnetic interference, the micro impulse radar can trigger falsely (a false trigger being when the radar alarm is triggered for any reason other than the motion of a nearby object). In testing the radar sensor, when put near a wireless router with a significant amount of radio traffic the radar could be triggered. Because a PIR sensor is detecting the infrared radiation emitted from the environment (triggering when a large enough change in the radiated energy occurs) and infrared radiation of objects are closely related to the objects' temperature, environments with temperatures that vary over relatively short periods of time, especially when the temperatures are around those of a human's body, can cause the PIR sensor to trigger falsely. An example would be a container that has been warmed by the sun being occasionally struck by medium to high breezes of cool air. The wind can cause parts of the container to cool relatively quickly and therefore cause the PIR sensor to trigger.

Since the radar and PIR sensor operate in extremely different areas of the electromagnetic spectrum, these devices can be used together to rule out the situations which cause false positives for each sensor individually. In addition, with a better understanding of how the sensor operates and, more importantly, how the sensor reacts to the environments it is in, an effective and accurate algorithm can be developed for intrusion detection. Although the radar alarm is sensitive to electromagnetic interference, its behavior when triggering due to the interference is quite different than when it triggers due to motion of nearby objects. When subject to interference, the radar triggers with a

small and regular period, as opposed to the longer irregular period when triggering due to the detection of motion. The PIR sensor also reacts differently when triggering due to changes in temperature of static environment versus triggering due to a non-static environment and in opposite manner to that of the radar. The PIR sensor will trigger regularly with a small period when detecting motion, and less regularly with a larger period when detecting the change in temperature of a static environment. To understand this difference in triggering behavior, the example of the warming/cooling container mentioned earlier can be explored deeper; the warm container is cooled relatively quickly by a cool breeze, causing the PIR sensor to trigger falsely. The container then warms up slowly due to the sun (since the heating is slow, it does not cause the PIR to trigger again). Eventually the container is cooled again and the PIR will trigger, but the process takes a long time when compared to the rate of triggering caused by the motion of an object. Since the wind is irregular and the warming of the container takes time, the trigger interval is irregular. On the other hand, when the PIR is detecting motion (and more importantly the motion of a human being), the movement is constant and quick, causing the PIR to trigger more frequently and at a more consistent interval.

As mentioned earlier, a false positive for each sensor individually is defined as the sensor being triggered for any reason other than the motion of a nearby object. However, a false positive for the entire cargo security device is anytime an intrusion is detected for any other reason than an object entering into the container. Although the two definitions may sound nearly identical, there are key differences which are important to explore. After the contents of a cargo container are loaded and the doors are shut and sealed, the container is then subject to a harsh environment. The container may be loaded onto a

truck, ship, and/or train and depending on the loading method may experience swaying motions, short violent motions, vibrations, and/or tilting/rotating. While the container is in transit it is also subject to many different types of motion. These movements can cause the contents in the container to move around and shift as well as cause the walls and other components of the container to vibrate, and, if the movements are large enough, will cause the radar and/or PIR to trigger. Since scenarios which can make the container engage in the previously mentioned behavior is by no means uncommon during shipping, having the means to detect and rule out such scenarios is a requirement for an accurate cargo security sensor. To detect the motion of the container, a three axis accelerometer was added to the design. The accelerometer can detect when the container is in motion, its orientation (such as when it is tipping), as well as when it has experienced abrupt shock-type motion (in the event that the container is dropped or struck by another object) and rule out alarms that are generated because of the motion.

To develop the intrusion detection algorithm, the CSD was placed inside a container and reported any sensor events along with the time at which they occurred. Different types of interactions with the container were administered, including opening and closing the doors, vibrating the container, and shifting the cargo while the CSD was in motion. Each administered interaction along with its time was noted and compared to the CSD's log file in order to find a relationship between the type of interaction and the way in which the sensors react. An example of the CSD's log file can be seen in Figure 17.


```

Entry 4 : 15:50, 12/17/2012, Type: Alarm, Duration: 0
Entry 5 : 15:00, 12/17/2012, Type: Temp/Hum/Batt, Temp(deg C) = 22, Hum(%) = 44, Batt(mV) = 3018
Entry 6 : 14:24, 12/17/2012, Type: Alarm, Duration: 0
Entry 7 : 14:24, 12/17/2012, Type: Accel: X: 2.812000 Y: 2.687000 Z: 3.968000
Entry 8 : 14:00, 12/17/2012, Type: Temp/Hum/Batt, Temp(deg C) = 22, Hum(%) = 43, Batt(mV) = 3016
Entry 9 : 13:00, 12/17/2012, Type: Temp/Hum/Batt, Temp(deg C) = 22, Hum(%) = 43, Batt(mV) = 3029
Entry 10 : 12:00, 12/17/2012, Type: Temp/Hum/Batt, Temp(deg C) = 22, Hum(%) = 43, Batt(mV) = 3022
Entry 11 : 11:00, 12/17/2012, Type: Temp/Hum/Batt, Temp(deg C) = 22, Hum(%) = 43, Batt(mV) = 3011
Entry 12 : 10:58, 12/17/2012, Type: Alarm, Duration: 1

```

Figure 17: Log file from CSD.

The intrusion detection algorithm uses a windowed approach to correlate radar, PIR, and accelerometer events in order to decide if a true intrusion has occurred. When an event occurs, a window of time is set wherein the device considers subsequent events related to the prior events which initially opened the window. Events are instances when one of the sensors are triggered, causing an interrupt for the processor to handle. Due to the different behavior of the sensors mentioned earlier, each window is specific to its corresponding sensor. An example of a time window can be seen in Figure 18: when a PIR event occurs, a five second window is opened for detecting a radar event. If a radar event occurs within the window, an alarm will be saved. Figure 18 also contains the window for two radar events. If an initial radar event occur (meaning that the radar is triggered while there is no previously opened window), then a window is opened one second after the initial radar event and lasts for five seconds. The delay in opening the window reduces the amount of false alarms caused noise.

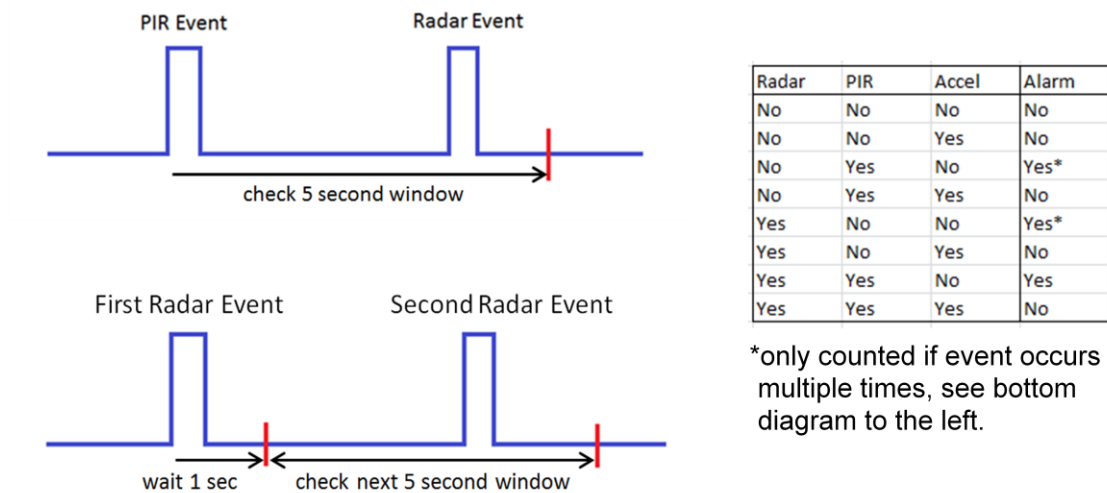


Figure 18: Windowed intrusion algorithm.

The detection of an intrusion will begin with one of two possible initial events: a radar or PIR event. If an initial radar event occurs, as is shown in Figure 19, the algorithm first checks the accelerometer; if the accelerometer has triggered recently due to vibration, shock, a considerable amount of motion, or a significant change in orientation, then the radar event will be disregarded since it is most likely due to the motion of the container. If the accelerometer has not triggered recently, the next window to check is the intrusion detection window, or alarm window. Since an intrusion can last anywhere from a couple seconds to hours, multiple intrusions would be logged for the single actual intrusion if there was not a delay after each logged alarm in which the algorithm attributes detected motion to the previous intrusion. If the alarm window has not expired, then the radar event is recorded as the most recent radar occurrence. If the alarm window has expired, then a new intrusion may be detected depending on the state of the radar window. Since the radar is subject to noise which typically causes it to trigger regularly and frequently (a period of less than one second), the radar's window

begins one second after it "initially" triggers (an initial trigger is when the radar event occurs after not occurring for six seconds) and lasts for five seconds. If the radar window has expired then the radar event is recorded as the "initial" radar event, unless the PIR has triggered recently (within the last five seconds). If a PIR event has recently occurred, then an intrusion will be logged and the radar, PIR, and alarm windows will be reset. If the PIR has not triggered recently, then the radar event will be recorded as the latest occurrence that will be used to detect intrusions when subsequent events occur. If the radar window has not expired and the radar has not triggered within the last second, then an intrusion is logged and the radar and alarm windows are reset. If the radar has triggered within the last second, the radar event is saved and marked as the latest occurrence which will be used later for determining future intrusions.

Since radar and PIR events that occur within a certain time of an accelerometer event are disregarded, it is possible that an intrusion can take place but not be detected if the container is moving or vibrating during the break-in. However, the CSD logs accelerometer events which can be inspected to discover if the container was subject to an extensive amount of motion. Accelerometer events occurring periodically for some time interval or when the container should not be in motion (identified by comparing the accelerometer event's time to the shipping schedule) are an indication of a possible intrusion. Unauthorized accelerations which do not correlate with the shipping schedule can be considered an intrusion either due to the possibility that the whole container was "stolen" (moved without permission), or the intruders attempted to neutralize the sensor by making it attribute any motion detected to the movement of the whole container (for

example, one intruder continuously strikes the outside of the container while the other intruder breaks-in).

The intrusion detection algorithm is similar when a PIR event occurs first, the only differences being the window sizes. Since the PIR has a several second settle time after it has detected motion, the window for checking for another PIR event does not open till several seconds after the first event. Figure 18 includes a truth table that summarizes the different series of event which cause an alarm. A ‘yes’ in the table signifies that the type of event occurred within the time window. Cases where only one type of event occurs only result in an alarm when the type of event happens multiple times and within the time window of the initial event.

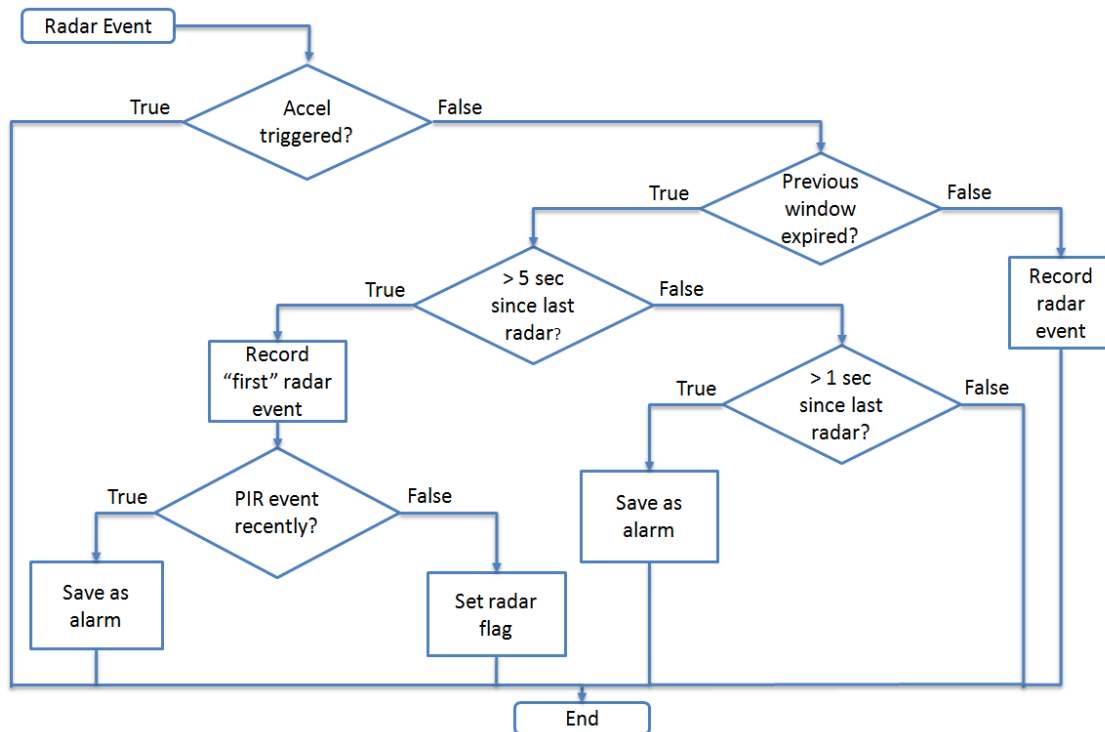


Figure 19: Intrusion detection algorithm beginning with a radar event.

6. Results

This section discusses the results of the project. From the requirements of the sponsor two overall goals are seen: create a CSD that requires little attention and detects container intrusions reliably. For the CSD to require little attention, it must be able to operate inside a container for long periods of time without needing maintenance. The major type of maintenance needed for remote sensors is the replacements or recharging of power sources. The first results section discusses the power consumption of the designed CSD. The subsequent sections present the test cases developed and the CSD's performance for detecting container intrusions.

6.1. Power Consumption

As a result of the careful selection of low power components as well as the event driven design of the firmware, the CSD accomplishes extremely low power consumption, especially when being compared to other smart container security sensors. In its typical use, the CSD spends most of its time in an idle state within which nearly every component is in a low power sleep mode. The processor is nearly completely off with its system clock disabled and interrupts enabled to wake it up on the event of an real time clock (RTC) tick (happening once per second), an accelerometer, PIR, or radar event, when it is time to take a temperature/humidity reading, or a reading device attempts to communicate. The accelerometer is in a sleep mode waiting for the container to move or vibrate significantly enough to cause an interrupt. The PIR sensor, temperature/humidity sensor, and flash memory are in their low power sleep modes. The radar is active but only consuming micro amps of current. The only device that is active is the Nordic BLE chip, which is in its advertisement state. The nRF8001's current draw spikes to about 12 mA

when it transmits an advertisement packet, but due to its advertisement scheme discussed previously its average current consumption is approximately 10 μA . In its idle state with the LED blinking at 1 Hz, the CSD's average current is .4 mA. A typical current consumption profile of the CSD in its idle state can be seen in Figure 20, showing the system mostly sleeping except when handling an interrupt event and advertising over Bluetooth.




Figure 20: Current profile for CSD.

Figure 21 breaks down the current consumption of the system during its idle state by each major component. The accelerometer, radar, and processor account for most of power consumption (the 'other' category is composed of the smaller system components, with its current consumption being dominated by the blinking status LEDs). During testing, when the device wakes up from its idle state, the highest current draw observed was 9 mA, making the overall average current consumption approximately .5 mA. Due to its low power consumption, the CSD can be powered by two AA size batteries and still have an estimated life of 5 to 6 months without needing the batteries to be replaced.

Being able to be powered by AAs increases the ease of use of the device, removing the need to have a specific charger or worry about replacing a rechargeable battery. AA batteries are readily available around the world, nearly guaranteeing the ability to power the CSD at all times.




Figure 21: Current consumption of major system components.

6.2. Intrusion Detection

After selecting the impulse radar as a possible sensor for the CSD, tests were conducted to inspect its usefulness inside a container. The radar was set inside the container powered by a coin cell battery with an LED connected to the digital alarm output. The radar was monitored by a stationary viewer watching for the LED to turn on when the container was breached. 100 test runs were administered with each run consisting of an intruder opening the container door, walking 6 feet into the container, and then exiting and closing the door. The test runs were done at various speeds, with the shortest occurrence happening in about 10 seconds and the longest for approximately one

minute. The LED was expected to light up when the door was opened, while the intruder moved in the container, and then go off once the door was closed. The initial results of the tests were promising, showing that every entry into the container was detected by the radar. Because of the penetrating abilities of the radar, cargo inside the container did not affect the detection of motion; people moving on the other side of cargo were still detected. Figure 22 shows a test case in which the radar sensor was placed in the back of a loaded container and an intruder moved around the entrance. There was approximately 10 feet of miscellaneous packaging, mainly consisting of paper products, blocking the line of sight between the intruder and the radar. The radar was still able to detect any significant motion of the intruder including door openings.

In addition to the radar's penetrating ability, the multipath generated from the reflection of the emitted signals off the cargo container's walls allowed the radar to "see around" and obstructions which it could not penetrate. This allows a single impulse radar to be used to detect motion anywhere within the container, as opposed to needing several sensor in different locations as was the case for the sensor types discussed in section 2. Figure 23 conveys the concept of multipath and how the emitted signals from a transmitter can reach a receiver in multiple ways due to reflections off the container wall. This led to the radar being able to detect not only any motion inside the container, but also the motion of the container walls. Detecting the motion of all the walls allows intrusions from any direction, not just the container doors, to be detected, something that all other cargo sensors lacked. However, an unwanted side effect is that it also detects the vibration of the walls, an event that happens quite often when the container is being moved or it is dropped. Being able to detect all motion inside the container also has

another unwanted side effect of alarming due to shifting cargo, which also occurs when the container is in motion. Since these false alarms are caused by the motion of the container and not the movement of an object within, the CSD will need to be able to detect container movement and rule out alarms triggered as a result. The detection of movement will be addressed by a sensor capable of measuring any vibration, shock, and/or motion of the container as discussed later.



Figure 22: Loaded cargo container used for test LLNL radar operation.




Figure 23: Illustration of multipath.

Tests for how well the CSD was capable of detecting intrusions were developed to support the sponsor's requirements. The sponsor explicitly stated that the CSD must be capable of a 95% detection rate per door event and a 4% error rate. A door event is defined as any unauthorized opening of the door while the container is supposed to be sealed. A door event may or may not include an intruder's entrance into the container. The sponsor defines the error rate as the percentage of false negatives and false positives caused by the system failing and the percentage of false positives due to container motion causing the CSD to alarm.

6.2.1. Test Development

To test the detection rate for door opening events, two cases were used: one case where the door was opened and closed, and another where the door was opened, an intruder entered the container, and then exited and closed the door. For the first case, the door was opened no less than six inches for at least two seconds before it was closed. The opening width and duration for the first test case reflects the average width and time needed to quickly place a small object inside the container. For the second case the

intruder moved about the container for at least five seconds, never going farther than half the container length. Again, these test parameters represent the average case of an intrusion as found by the sponsor and LLNL. To test the CSD's ability to detect break-ins from any wall, the sensor was placed in seven different locations within the container with 50 intrusion events occurring while the sensor was at each location. The CSD was firmly attached to the walls so that the sensors motion would match the containers. In addition, each case was administered with the container empty and with the container approximately half full. An empty container allowed the CSD to have line-of-sight with all of cargo area, while the partially loaded container tested the sensor's ability to detect intrusions with obstacles blocking its direct view. A successful detection of an intrusion is defined as the CSD logging at least one alarm entry per intrusion. Since each test case occurs for a duration of time, the sensor may detect multiple instances of motion per intrusion.

Testing the detection error rate due to container motion is a more difficult task. While a container is being shipped, it will be loaded and travel on ships, trains and trucks, causing the container to rock and tilt. In addition, the containers are often dropped or struck by an object. There are two main types of motion that the container can experience: lower frequency movements (surge, sway, and yaw) and high frequency movements (shock and vibration). To replicate container motion during shipping, ideally the container would be sealed and then lifted by a crane and moved around. However, this is an extensive test setting which would be hard to repeat. Since the CSD's motion is coupled to the container's (due to it being tightly attached to the walls), the ability of the sensor to rule out alarms due to container movement can be tested by moving the sensor

and checking if it still logs alarms. To test if the CSD could rule out alarms during periods of lower frequency motion, the sensor was attached to a rotating arm that was manually operated. The arm would be rotated at different speeds with the CSD attached in different orientations. Since the motion of a container on a ship, train, or truck varies greatly, the sensitivity of the motion detection can be tuned. For the test case, the threshold was set to .5 g, the minimum value stated by the U.S. Department of Transportation for cargo to withstand during transportation [11]. If an alarm was logged while the accelerometer reported a reading of .5 g or above, then the alarm was considered an error. Creating a test case for vibration and shock sensitivity was simpler; the CSD was mounted to the wall and the container was struck so that the walls would vibrate. During the test, there were no intrusions into the container; therefore an alarm logged at anytime during the test was considered an error. Furthermore, if an accelerometer event was not logged when the vibration occurred, it was considered an error.

6.2.2. Door Event Tests

The test results for the door opening test cases, as seen in Tables 1, 2, 3, and 4, show that the CSD accomplishes the sponsor's requirements for the detection rate per door event with the overall detection rate being 98.7%. Furthermore, the test cases reveal that certain mounting locations inside the container are more ideal for the reliable detection of intrusions; the sensor being mounted on the container's roof near the door produced the lowest error rates. Conversely, the detection rate was lowest when the sensor was mounted on the right wall near the rear of the container. When comparing the

test results for an empty container to the partially loaded container it can be seen that cargo has a slight adverse effect on the detection rate, lowering it from 99.5% to 98%.

Table 1 reports the results for the door opening event tests in an open container with no intruder entry. The overall detection rate for this test case was 98.9%, with false negatives occurring when the sensor was placed on the roof or the right or left wall near the back of the container. Since the sensor was able to detect all intrusions when located on any wall near the front as well as when it was on the back wall facing the front, it appears that a combination of the sensors orientation (not facing toward the door) and its location (near the back of the container) caused the detection rate to drop.

Table 1: Results for door opening with empty container test case.

Sensor Location	True Positives	False Negative	False Positive	Total
Left Wall Near Door	50	0	0	50
Left Wall Near Rear	49	1	0	50
Right Wall Near Door	50	0	0	50
Right Wall Near Rear	48	2	0	50
Rear Wall	50	0	0	50
Roof Near Door	50	0	0	50
Roof Near Rear	49	1	0	50
Totals	346	4	0	350
Detection Rate	98.9			

Table 2 contains the detection data for the test case of a door event as well as intruder entry into an empty container. The CSD was able to detect every intrusion that was administered. When comparing Table 1 to Table 2, it can be seen that the sensor is more sensitive to intruder motion within the container than it is to the motion of the door.

If the CSD did not detect the opening or closing of the door, it was able to detect the presence of the intruder. One reason for the higher detection rate is most likely attributed to the PIR sensor's effectiveness at detecting motion when it has an unobstructed line-of-sight with the intruder. Although this was not a test case directly required by the sponsor, who was more concerned with strictly door event and no entry, the results are important since intrusions typically involve the introduction of new objects (whether it be the intruders or unwanted cargo) or the removal of existing cargo. The motion of such objects inside the container increases the detection ability of the CSD.

Table 2: Results for door opening and entry with empty container test case.

Sensor Location	True Positives	False Negative	False Positive	Total
Left Wall Near Door	50	0	0	50
Left Wall Near Rear	50	0	0	50
Right Wall Near Door	50	0	0	50
Right Wall Near Rear	50	0	0	50
Rear Wall	50	0	0	50
Roof Near Door	50	0	0	50
Roof Near Rear	50	0	0	50
Totals	350	0	0	350
Detection Rate	100			

The same test summarized in Table 1 was repeated but with the container partially loaded with cargo. It is important to note that with the container loaded, the sensors line-of-sight with the door was blocked when it was placed on the left wall near the rear, on the right wall near the rear, and on the rear wall. Table 3 shows the results of the test and mostly coincides with the results seen in Table 1. These results, however, depict the

effects of the cargo on the detection rate. When the sensors line-of-sight with the intrusion was blocked, the detection rate dropped. This is specifically seen with the increase in false negatives when the sensor was placed near the rear of the container.

Table 3: Results for door opening with loaded container test case.

Sensor Location	True Positives	False Negative	False Positive	Total
Left Wall Near Door	50	0	0	50
Left Wall Near Rear	48	2	0	50
Right Wall Near Door	50	0	0	50
Right Wall Near Rear	47	3	0	50
Rear Wall	45	5	0	50
Roof Near Door	50	0	0	50
Roof Near Rear	50	0	0	50
Totals	340	10	0	350
Detection Rate	97.1			

Table 4 is the results of the test case where an intruder entered a partially loaded container and then exited. The same trend can be seen when comparing Tables 4 and 3 as when comparing Tables 2 and 1: intruder entry increases the detection rate of the CSD. However, when comparing Tables 4 and 2, the effects of the partially loaded container are seen through the lower detection rate reported in Table 4. When the container was empty the CSD could rely on the PIR sensor to detect motion, but with a loaded container and an obstructed view, the radar became the primary sensor. As with the other test results, with the exclusion of the sensor located on the rear wall in Table 3, the CSD had the lowest detection rate when located on the right wall near the rear. The poor behavior

of the sensor located on the rear wall in Table 3 is likely due to the large amount of cargo that was present between the sensor and the intrusion. The right wall near the end of the container appears to be the least ideal location for the CSD.

Table 4: Results for door opening and entry with loaded container test case.

Sensor Location	True Positives	False Negative	False Positive	Total
Left Wall Near Door	50	0	0	50
Left Wall Near Rear	49	1	0	50
Right Wall Near Door	50	0	0	50
Right Wall Near Rear	48	2	0	50
Rear Wall	49	1	0	50
Roof Near Door	50	0	0	50
Roof Near Rear	50	0	0	50
Totals	346	4	0	350
Percent Error	98.9			

6.2.3. Container Motion Tests

The results for the movement tests, listed in Table 5, find that the CSD logged an alarm due to its motion for 3% of the test samples. Most of the errors (five out of six) occurred when the sensor was either mounted facing up or down. A possible explanation for this behavior is that the accelerometer filters the acceleration measured on its Z-axis, assuming that the sensor will be oriented so that this is the axis aligned with Earth's gravity. The accelerometer was mounted on the motherboard in that manner. When the sensor is oriented in either of those two positions, the movement administered during the test will be measured on the sensor's Z-axis, which will be filtered. The filtering may have reduced the amount of motion perceived by the accelerometer, bringing it under the

set threshold that triggers a motion event and therefore causing the CSD to not report container motion.

Table 5: Results for movement test case.

Sensor Orientation	True Positives	False Negative	False Positive	Total
Facing Up	47	0	3	50
Facing Down	48	0	2	50
On Left Side	50	0	0	50
On Right Side	49	0	1	50
Totals	194	0	6	200
Error Rate	3.00			

Table 6 reports the results for the vibration tests, showing that the CSD alarmed due to container vibration 3.4% of the time. The CSD had the most errors when located on walls that are near the stronger structural components of the container. A possible explanation is that the locations farther away from the container's structure (mainly the roof) have a stronger response to container vibration and movement. These locations experience greater, or longer, accelerations that the CSD is more capable of detecting. Combining the results from Tables 5 and 6 confirm that the CSD's detection error rate due to container movement is under the 4% sponsor requirement, with the sensor falsely alarming 3.3% of the time.

Table 6: Results for vibration test case.

Sensor Location	True Positives	False Negative	False Positive	Total
Left Wall Near Door	46	0	4	50
Left Wall Near Rear	47	0	3	50
Right Wall Near Door	48	0	2	50
Right Wall Near Rear	49	0	1	50
Rear Wall	48	0	2	50
Roof Near Door	50	0	0	50
Roof Near Rear	50	0	0	50
Totals	338	0	12	350
Error Rate	3.43			

Conclusion

This study has found that the LLNL micropower ultra-wideband appears to be a suitable solution for the detection of intrusions in cargo shipping containers. Specifically, through the testing of the standalone radar as well as the CSD using the radar, this study has shown that the technology is highly reliable for the detection of container intrusions. The radar's penetrating ability, low susceptibility to noise and behavior due to multipath are features that make it an effective break-in sensor inside a container. A CSD employing the radar technology along with supporting sensors and hardware (PIR, accelerometer, processor, etc.) is capable of detecting container break-ins with low error rates. By selecting low-power options for the support hardware and the low-power operation of the impulse radar, the CSD can be powered by common, easily accessible batteries (such as AAs) and not require replacements for up to six months. The project also investigated the new Bluetooth Low Energy standard and used it in a unique way. It was found that BLE devices provide ultra low-power solutions for wireless communication, and, although not originally designed for bulk data transfer, can be used for the transmission of relatively large amounts of data in short periods of time (approximately 200 kbps). Through the use of the latest Bluetooth Low Energy devices, cargo container inspections will no longer require the opening and examination of the container. Instead, a reading device can wirelessly connect to the CSD inside the container and download the tamper status, allowing inspectors to know whether further investigation is needed. The end result is a reliable and inexpensive CSD which requires minimal effort, has a low impact on the loading, unloading, and sealing of the container,

and enables the rapid approval of shipping containers during inspection. The design proposed in this study can be used when developing future CSDs.

Bibliography

- [1]. Chiew Loon Goh, Nor Muzakir Nor Ayob, Ruzairi Abdul Rahim, Herlina Abdul Rahim, Muhammad Jaysuman Pusppanathan, Mohd. Hafiz Fazalul Rahiman, Leow Pei Ling, Zulkarnay Zakaria, “Study on Wireless Sensor Based Industrial Tomography Systems”, *Sensors & Transducers*, Vol. 154, Issue 7, July 2013, pp. 71-81
- [2]. Ali Elkateeb, “Soft-Core Processor Design for Sensor Networks”, *International Journal of Computer Networks & Communications (IJCNC)*, Vol. 3, No. 4, 2011.
- [3]. Vehbi C. Gungor, and Gerhard P. Hancke, “Industrial Wireless Sensor Networks: Challenges, Design Principles and Technical Approaches”, *Industrial Electronics*, Vol. 56, No. 10, 2009, pp. 4258-4265
- [4]. S.W. Eun Kim, “Performance Comparison of Loop/Piezo and Ultrasonic Sensor-Based Traffic Detection Systems for Collecting Individual Vehicle Information”, *Proceedings of the 5th World Congress Intelligent Transport Systems*, 1998
- [5]. Amin Kianpisheh, Norlia Mustaffa, Pakapan Limtrairut, and Pantea Keikhosrokiani, “Smart Parking System (SPS) Architecture Using Ultrasonic Detector”, *International Journal of Software Engineering & Its Applications*, Vol. 6, Issue 3, July 2012, p 51
- [6]. Bernhard Wirnitzer, Wolfgang M. Grimm, Hauke Schmidt, Robert Bosch GmbH, “Interference Cancelling in Ultrasonic Sensor Arrays by Stochastic Coding and Adaptive Filtering”, *International Conference on Intelligent Vehicles*, 1998.

- [7]. F.P. Martinez, F.C. Galeano, “New Microwave Sensors for Intrusion Detection Systems”, *33rd Annual International Carnahan Conference on Security Technology*, 1999, pp. 49-53
- [8]. International Maritime Organization, “International Shipping Facts and Figures – Information Resources on Trade, Safety, Security, Environment”, March 2012
- [9]. Su Jin Kim, Guofeng Deng, Sandeep K.S. Gupta, Mary Murphy-Hoye, “Enhancing Cargo Container Security during Transportation: A Mesh Networking Based Approach”, *Conference on Technologies for Homeland Security*, May 2008, pp. 90-95
- [10]. United States Nuclear Regulatory Commission, “Intrusion Detection Systems and Subsystems”, March 2011
- [11]. United States Department of Transportation, “Federal Motor Carrier Safety Administration’s Cargo Securement Rules”, December 2003
- [12]. Igor Y. Immoreev, Sergey V. Samkov, The-Ho Tao, “Short-Distance Ultra-Wideband Radars. Theory and Designing”, *Aerospace and Electronic Systems Magazine*, Vol. 20, Issue 6, June 2005, pp. 9-14
- [13]. Stephen Azevedo, Thomas E. McEwan, “Micropower Impulse Radar”, *Potentials*, Vol. 16, Issue 2, May 1997, pp. 15-20
- [14]. Nordic Semiconductor, “Single-chip Bluetooth low energy solution”, nRF8001 datasheet, January 2012 [Revised August 2013]
- [15]. Xandem, “Tomographic Motion Detection (TMD) User Guide”, TMD datasheet, September 2013