

ELE 302: BUILDING REAL SYSTEMS

Professors: Andrew A. Houck and Antoine Kahn

Independent Project: Ultrasonic Positioning System

Michael Danielczuk, Andrew Kim, Monica Lu, and
Victor Ying

Due Friday, May 15, 2015

Contents

1	Introduction	3
2	Multilateration	3
2.1	Time difference of arrival multilateration	4
2.2	Position calculations	4
2.2.1	A closed form solution?	6
2.2.2	A least squares regression	7
3	Hardware systems	8
3.1	Ultrasonics	8
3.1.1	Ultrasonic transmitter stations	9
3.1.2	Ultrasonic receivers for cars	9
3.2	Serial radio communication	12
4	Signals	12
4.1	Transmitted signals and synchronization	13
4.1.1	Synchronization	13
4.1.2	Sending Pings	14
4.1.3	Timeouts and Error Checking	15
4.2	Digital processing of received signals	16
4.2.1	Digital filtering	16
4.2.2	Timing	18
5	Testing and debugging scaffolds	19
6	Performance and sources of bad data	20
7	Applications	21
7.1	Navigation	21
7.1.1	H-Bridge	21
7.1.2	General autonomous navigation	22
7.1.3	Simpler navigation demonstration	23
7.2	Mapping	23
A	C code compiled in PSoC Creator for car	26
A.1	main.c	26
A.2	position.c	28
A.3	drive.c	35
A.4	speed.c	37
A.5	steer.c	47
A.6	shell.c	54
A.7	usb_uart.c	60

B	Arduino code for transmitter stations	64
B.1	Master transmitter station	64
B.2	Slave transmitter stations	68
C	Arduino code for testing tools	71
C.1	tx_test_2	71
C.2	tx_test_4	73
C.3	rx_test	75
D	MATLAB code for mapping	76

1 Introduction

Our indoor positioning system works by the same principle as a satellite navigation system, with ultrasonic pings substituted for radio waves. Signals are transmitted from stations at known locations, and signals traveling different distances to a given receiver will take different amounts of time to get there. With enough information about distances to known locations, the receiver's position can be calculated. As the speed of sound at room temperature is a mere 1126 feet per second, it is easy for general-purpose microcontrollers to measure the timing of ultrasonic signals to a precision that would enable spatial resolution of better than one foot. Our 25 kHz ultrasound has a wavelength of about half an inch, so if our receive circuitry allows for timing of ultrasonic signals to within several cycles, we might expect to achieve spatial resolution of a few inches. In fact, this is exactly what we observe.

The next three sections of this report document the design of the positioning system at three levels of abstraction. First, at the highest level of abstraction, there is Section 2, which explains the principles and mathematics of positioning used, assuming the ability to measure differences in time of flight. The other two sections concern the hardware and software systems needed to obtain these measurements. Section 3 discusses the physical hardware systems we had to build, while Section 4 discusses the way the various components of the system communicate and the signals and processing that are needed to get accurate measurements of differences in times of flight.

2 Multilateration

A positioning system generally either relies upon the ability to either measure different distances to known points (multilateration), or different angles to known points (multiangulation). Multiangulation could be done by use of cameras or other angle-resolved sensors, but using cameras located on mobile robots such as our cars would be difficult as each car would need to process probably several different video feeds covering a 360 field of view and correctly identify beacons at known locations, and transform the apparent position of beacons in the video frames to the angles in which the beacons are oriented. Additionally, putting a bunch of cameras on every robot would be somewhat expensive. If doing multiangulation, it may make more sense to have a fixed number of cameras at known locations processing the angles it sees cars at. Unfortunately, this requires the cameras to actively track and communicate with each and every robot, which does not scale well to many independent robots.

Multilateration can be done by measuring the varying time of flight of signals between receivers on the robots to be located and transmitter stations at known locations. The benefit of this approach is that each receiver does not have to communicate any information to any transmitter or any other receiver. If the transmitters are transmitting a known signal, then the receiver can simply time the signals it receives to get information about the time of flight of the signals.

2.1 Time difference of arrival multilateration

It might seem easy to measure distances simply by measuring the time from when a transmitter’s microcontroller starts trying to transmit and the time the receiver circuitry sends a detectable signal to a microcontroller on the car. Unfortunately, there are sources of delay that would contribute to this time measurement other than the time it takes for the ultrasound to propagate through the air. The transmitting and receiving ultrasonic transducers act as mechanical band-pass filters, as they are built to physically resonate at 25 kHz. They, together with the more significant band-pass filter in the receiver circuitry, will add some delay to the signal. We could carry out experiments to measure and correct for this delay, but with the receive circuitry in particular, this would require either being careful with tolerances to ensure each receive amplification and filter circuit is similar enough to the next, or individually characterizing and calibrating the correction term for different receivers. It would also mean every receiver would need to be synchronized with every transmitter to measure the time of transmission and time of receiving each signal on a shared clock, and such synchronization generally requires two-way communication between the transmitters and receivers, or one-way communication with known (or very small) latencies, which would impact the ability of the system to scale as more receivers are added.

In the approach used in our system, as well as in satellite navigation systems, there is a small fixed number of transmitter stations that are synchronized with each other and send out signals at fixed known times relative to each other, but the receivers do not communicate with the transmitter stations to synchronize their clocks. Knowing that the transmitters are broadcasting certain signals at certain times, the receiver can simply observe the time difference at which it sees signals from different transmitters to get a difference in the lengths of the signals flights. Since all the transmitters have identical hardware and are sending out the same sort of signal, and these signals are received through the same receiver hardware for each receiver, any difference in the signals’ times of flight must be due to differences in times of propagation of the ultrasound through the air. This has the advantage that the receiver does not need to communicate with the transmitter stations in any form to synchronize its clock with their clocks. In general, calculating position using a time difference of arrival measurements requires one more measurement (i.e., one more transmitter) than using absolute time of flight measurements. However, adding one more transmitter station that is identical to the rest requires no additional hardware design, so it is a small fixed cost to pay for making the positioning system more scalable, as, without the need to synchronize receivers, now there can be any number of receivers which do not need to interact with each other or the transmitters to calculate their individual positions.

2.2 Position calculations

The problem of multilateration based on time differences of arrival was explored and solved in the context of “sound ranging”, a practice used beginning in World War I to locate enemy artillery by comparing the time at which the sound of the enemy guns firing reached sound recording devices placed at known locations along the front. In sound ranging, the enemy

guns at an unknown position were the transmitters of the sound signal. For our case, the transmitter stations are at known locations and the receiver's position is what we wish to know, but the ideas in the math are the same.

We place the transmitters in a rectangle of known width X and length Y . The transmitters are arbitrarily numbered zero through three going counterclockwise, starting in the third quadrant. The origin of our coordinate system is arbitrarily placed below the center of this rectangle along the z -axis, at the height of the car receiver. The receiver is at a point $(x, y, 0)$. Let Z be the height of the transmitter plane above the height of the receiver; this is known because the car travels only on the flat floor. x and y are the unknowns we wish to calculate.




Figure 1: Our coordinate system.

We can write the measurements the receiver takes as being the difference between each distance to transmitters one through three and the distance to transmitter zero:

$$\Delta_1 = D_1 - D_0 = \sqrt{\left(x - \frac{X}{2}\right)^2 + \left(y + \frac{Y}{2}\right)^2 + z^2} - \sqrt{\left(x + \frac{X}{2}\right)^2 + \left(y + \frac{Y}{2}\right)^2 + z^2}$$

$$\Delta_2 = D_2 - D_0 = \sqrt{\left(x - \frac{X}{2}\right)^2 + \left(y - \frac{Y}{2}\right)^2 + z^2} - \sqrt{\left(x + \frac{X}{2}\right)^2 + \left(y + \frac{Y}{2}\right)^2 + z^2}$$

$$\Delta_3 = D_3 - D_0 = \sqrt{\left(x + \frac{X}{2}\right)^2 + \left(y - \frac{Y}{2}\right)^2 + z^2} - \sqrt{\left(x + \frac{X}{2}\right)^2 + \left(y + \frac{Y}{2}\right)^2 + z^2}$$

Given that we can measure Δ_1 , Δ_2 , and Δ_3 , we wish to solve these three equations for the two unknowns x and y . In the next two subsections, we will discuss two approaches to this calculation that we tried.

A Mathematica notebook file showing the derivations of all the formulas discussed below can be found at <https://github.com/YingVictor/ultrasonic-positioning/raw/master/MultilaterationDerivationMathematica.nb>. The exact implementation of these calculations in software can be found in the file `position.c` located in appendix section A.2.

2.2.1 A closed form solution?

The problem is overdetermined, but if we treat Z as an additional unknown, it turns out this problem has a closed-form solution that can easily be found using computer algebra software such as Mathematica:

$$\begin{aligned}
 x &= \frac{\Delta_1(\Delta_2 - \Delta_3)(\Delta_1 - \Delta_2 - \Delta_3)}{2X(\Delta_1 - \Delta_2 + \Delta_3)} \\
 y &= \frac{\Delta_3(\Delta_2 - \Delta_1)(\Delta_3 - \Delta_2 - \Delta_1)}{2Y(\Delta_1 - \Delta_2 + \Delta_3)} \\
 Z &= \pm \frac{1}{2XY(\Delta_1 - \Delta_2 + \Delta_3)} \left[X^2Y^2 (\Delta_1^2 + (\Delta_2 - \Delta_3)^2) (\Delta_3^2 + (\Delta_2 - \Delta_1)^2) \right. \\
 &\quad \left. - X^2 (\Delta_3^2(\Delta_1 - \Delta_2)^2(\Delta_1 + \Delta_2 - \Delta_3)^2 + X^2Y^2(\Delta_1 - \Delta_2 + \Delta_3)^2) \right. \\
 &\quad \left. - Y^2 (\Delta_1^2(\Delta_3 - \Delta_2)^2(\Delta_3 + \Delta_2 - \Delta_1)^2 + X^2Y^2(\Delta_1 - \Delta_2 + \Delta_3)^2) \right]^{\frac{1}{2}}
 \end{aligned}$$

This geometry gives fairly simple closed form solutions for x and y , and we can throw away the calculation for Z . The solution can be calculated directly except for the locations where the denominator tends to zero. This only occurs when $\Delta_1 - \Delta_2 + \Delta_3 \rightarrow 0$. Substituting in the definitions of Δ_1 , Δ_2 , and Δ_3 and simplifying, it can be demonstrated that this only occurs when $x \rightarrow 0$ or $y \rightarrow 0$, which, intuitively, is when the receiver is equidistant from some of the transmitters. When calculating our position we can first check for these two cases and change the calculation appropriately. Specifically, if $\Delta_1 \approx 0$, then we can estimate our position as

$$\begin{aligned}
 x &= 0 \\
 y &= \frac{\Delta_3\Delta_2}{2Y} \\
 Z &= \pm \frac{1}{2Y} \sqrt{(\Delta_3^2 - Y^2) (Y^2 - \Delta_2^2) - X^2Y^2}
 \end{aligned}$$

And if $\Delta_3 \approx 0$, then we can estimate our position as

$$\begin{aligned}
 x &= \frac{\Delta_1\Delta_2}{2X} \\
 y &= 0 \\
 Z &= \pm \frac{1}{2X} \sqrt{(\Delta_1^2 - X^2) (X^2 - \Delta_2^2) - X^2Y^2}
 \end{aligned}$$

This calculation turns out to work decently. Unfortunately, by treating Z as an unknown, we are throwing away information rather than making use of all the information we have to get a better approximation. We could actually calculate the value of Z and compare it against the known value as a sort of sanity check to help throw out bad measurements, but the calculation for Z is sufficiently complex that it seemed not worth the trouble. This particular solution tends to magnify errors in measurement, particularly near the x - and y -axes where certain factors in both the numerators and denominators become very small and measurement errors can change these factors substantially. Using the approach just described where near the x - and y -axes we switch to a different approximation, the positioning calculation becomes discontinuous, so the errors one sees are not always locally consistent as one moves around in a small area. The entire problem of position calculation could be approached in other ways that do not give closed-form solutions but require iterative numerical approximation methods.

2.2.2 A least squares regression

In an overdetermined problem, a common approach to finding an approximate solution is to minimize some metric that is the sum of some squares of errors. Actual GPS receivers do this, and, with some encouragement from Prof. Houck, we ended up doing this as well. In particular, the metric we minimized was

$$f(x, y) = \left(\Delta_1 - \hat{\Delta}_1(x, y)\right)^2 + \left(\Delta_2 - \hat{\Delta}_2(x, y)\right)^2 + \left(\Delta_3 - \hat{\Delta}_3(x, y)\right)^2$$

and so the problem we were solving was

$$\arg \min_{x, y} f(x, y)$$

where Δ_n indicates the actual measured differences described earlier, and $\hat{\Delta}_n(x, y)$ indicates the expected value for Δ_n given particular values of x and y . If you review how we defined the Δ_n , you may notice that all of these deltas are with respect to D_0 , the distance to transmitter zero located at $(-\frac{X}{2}, -\frac{Y}{2}, Z)$. This may seem like we have arbitrarily chosen to give more weight in the errors to this first distance. In fact this choice seems sensible in the context of how the transmitters synchronize, which is potentially a significant source of errors. As discussed in Section 4.1, TX0 is the master that determines when ultrasonic signals should begin being sent out, and the other transmitters follow its lead.

The method of minimization that we use is essentially Newton's method. Given a guess $\mathbf{r} = (x, y)$, we calculate the next guess as

$$\mathbf{r}_{\text{next}} = \mathbf{r} - \alpha \frac{f(\mathbf{r})\nabla f(\mathbf{r})}{\|\nabla f(\mathbf{r})\|^2}$$

where the factor $\alpha = 0.1$ is used to keep the guess from moving too far too quickly, which can lead to overshoot and failure to converge. The values of f and ∇f can be calculated

exactly at each point (x, y) , as f and its partial derivatives have perfectly cromulent closed-form expressions in terms of x and y . The expressions are quite large to write out, and so have been left out here, but many terms and factors, such as the expected distances between the four transmitter stations and the point (x, y) , appear repeatedly, so they can be calculated once and reused in the calculations in software. The exact calculation can be seen in appendix section A.2.

If at any point $\nabla f(\mathbf{r}) = \mathbf{0}$, that is, if both partial derivatives of f are zero and we have reached a stationary point, we stop the iterative minimization method. Otherwise, we continue iterating to get better approximations for x and y until either 100 iterations have been performed, which is generally more than enough, or the value of the metric f drops below 0.01 ft^2 , at which point it is sufficiently small that we don't care to keep performing the iterative minimization.

Minimizing this error metric has the useful benefit that, if we find the value of the metric is very close to zero, we can be quite confident that we are basing our position calculation off of good measurements. Conversely, if the metric is large after the minimization, then we may wish to throw this set of measurements and calculation away, as the measurements themselves may have been suspect. Sources of bad measurements from ultrasonic signals are discussed in Section 6. In our experience, when the receiver is stationary, our position calculations always give a consistent position to within a third of a foot, and the value of the error metric at that minimum is generally less than 0.1 ft^2 . (The value of the metric at the minimum is, in general, likely to be smaller than the square of the actual error because the metric at the minimum is in general appreciably smaller than the metric evaluated at the correct (x, y) values.) In our code, we check that the minimized value of the error metric is less than 0.5 before recording and reporting the new calculated position. If this check fails, we simply throw the current measurement away.

3 Hardware systems

3.1 Ultrasonics

We decided to use 25 kHz ultrasonic signals for this project. This frequency is far enough from the typical 40 kHz used by the ultrasonic rangefinders commonly used by hobbyists for us to be confident that we would avoid picking up the rangefinder pings from other robots in the room using ultrasonic rangefinding, which might otherwise interfere with our positioning. We decided to use a lower ultrasonic frequency rather than a higher one because lower frequencies experience less attenuation traveling through the air. A lower frequency and hence a longer wavelength might sacrifice spatial resolution, but it was felt early on in the design process that maximizing the chances that we would be able to receive ultrasonic signals from across the room was a much greater concern. We decided to buy simple ultrasonic transducers so we could design circuits ourselves to maximize the range of our system. To this end, we purchased Prowave 250ST180 and 250SR180 transducers, which are the versions of the 18 mm diameter ultrasonic transducers offered by Prowave, optimized for transmission and

reception, respectively. These ultrasonic transducers cost less than \$8 each, but were still relatively expensive compared to other common ultrasonic transducers that might cost only a couple dollars. They were chosen for their relatively high advertised sensitivity and sound pressure level.

3.1.1 Ultrasonic transmitter stations




Figure 2: XBee radio module mounted to an Arduino, which connects to an ultrasonic transmitter. Note the use of the book-based precision mounting system.

We used an Arduino Uno as a low-cost microcontroller for our transmitter stations, which need to communicate with each other for synchronization purposes and control the Prowave 250ST180 ultrasonic transducers. The Prowave 250ST180 is simply connected directly to two digital output pins of the Arduino, which are capable of outputting 0 or 5 V. The reason we use two output pins rather than, for example, one output pin and ground, is that we can drive the two output pins as the inverse of each other, effectively doubling our amplitude. This is discussed further in section 4.1.2. We observed this improved our reception range a little bit in testing.

3.1.2 Ultrasonic receivers for cars

When receiving ultrasonic signals from across the room, we typically see the Prowave 250SR180 transducer produce signals on the order of a millivolt, and sometimes even lower. To be able to read this signal using a PSoC and to get rid of unwanted noise requires significant amplification and filtering. The circuit we ended up with for this purpose is shown in Figure 3.

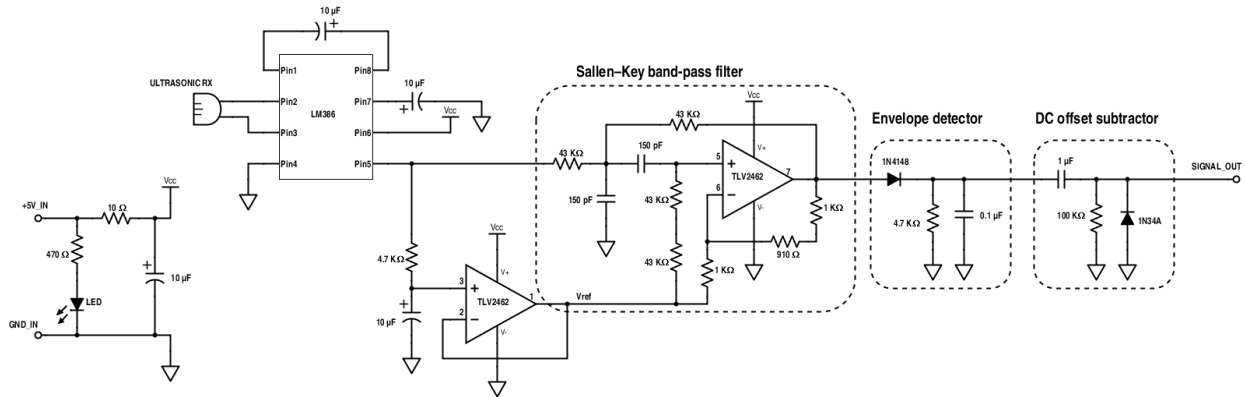


Figure 3: Schematic of ultrasonic receiver amplification and filtering circuitry.




Figure 4: Photograph of ultrasonic receiver amplification and filtering circuitry.

We spent a while trying to build our own differential amplifiers with voltage gain on the order of hundreds, using general purpose op-amps. We repeatedly found that the noise was intolerably bad, it was easy to produce unstable circuits when trying to make the gain so large, and, ultimately, when we got past those problems, we rediscovered the limitations of gain-bandwidth products. After all, to give our 25 kHz signal a gain on the order of a thousand would require gain-bandwidth product on the order of 25 MHz! Finally understanding the amplifier chip specifications we needed to be looking for to provide large gain at ultrasonic frequencies, we went in search of a solution, and found that the LM386 audio amplifier chip was readily available. This chip takes a differential input biased to ground, and produces an output referenced to half the power supply, with a gain that can be set between 20 and 200. As it is designed for audio, it works just fine for our 25 kHz signal, which is just barely above the upper limit for human hearing. It works extremely well providing a large

first stage of amplification with the gain set to 200.

After the initial amplification, we were seeing a decent amount of both 60 Hz noise and high frequency noise in the signal. We wished to put our signal through band-pass filtering to reduce those low and high frequency noises, as well as provide additional amplification to bring the signal up to the level of whole volts. To do this, we decided to build a Sallen–Key band-pass filter, which is a common topology for an active filter that in general provides non-unity gain, which is desirable in this case. This filter theoretically has a center frequency of 24.7 kHz, a Q factor of about 11, voltage gain of about 32 at the center frequency. We found some TLV2462 dual op-amp chips to use that have a gain-bandwidth of about 6 MHz, which is sufficient. In practice, due to poor tolerances in the components, the Q factor and the gain both tend to be lower than the theoretical value, but the gain in all the versions of this circuit that we built was above an order of magnitude at 25 kHz, and it showed a decent ability to reject of 40 kHz signals, which we might pick up from ultrasonic rangefinders used in other projects. Perhaps more importantly, it did quite a good job filtering out 60 Hz and high frequency noise, producing a fairly clean waveform. As the amplification is relative to a reference voltage of half the power supply, the maximum amplitude we can get in the output is half the power supply, or 2.5 V. The two amplification stages together have a voltage gain of several thousand which is usually enough to get clipping of the received signal. Clipping is not a problem in our context, so this is fine.

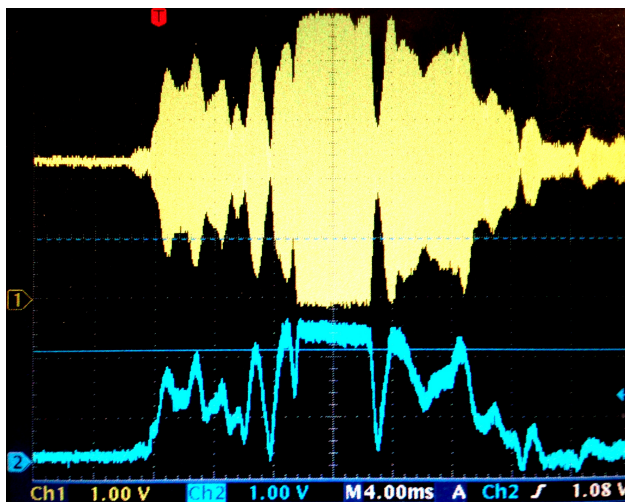


Figure 5: Ch1 (Yellow): Output of the Sallen–Key band-pass filter. Ch2 (Blue): Output of DC offset subtractor, after the envelope detector.

After the amplification and filtering, we have a fairly straightforward envelope detector to demodulate the signal that uses an ordinary silicon diode. After the envelope detector, we subtract off the DC offset in order to get a signal that is referenced to ground. This is essentially achieved using a AC coupling capacitor in the signal line and a pull-down resistor to ground on the output. However, if the DC offset subtraction consisted of only those two components, it would pull the average level of the output to ground, and the output would range from positive to negative voltages, and the voltage of the top and bottom of pulses in

the signal would vary depending on the timing and durations of the pulses. To prevent this, a 1N34A germanium diode was added between the output and ground to limit the output so that it cannot drop far below ground. In particular, it cannot drop below ground by more than the turn-on voltage of the diode. A germanium diode was chosen because of its low turn-on voltage. This has the effect of keeping the parts of the signal that are low close to ground, and the pulses from when ultrasonic pings are received are all to positive voltages. A captured waveform demonstrating the function of the envelope detector and DC offset subtractor is shown in Figure 5.

3.2 Serial radio communication

To synchronize all of the transmitters, we needed a system of communication between the transmitters with either very low or very stable latency. A logical wireless communication choice, due to its relative simplicity was the XBee 802.15.4 (series 1) radio module from Digi. In contrast, we could communicate with a transport protocol such as TCP or UDP over an internet connection using WiFi, but with such a heavy protocol stack, it becomes difficult to reason about the sources of latency and to know whether latency is symmetric, and in all likelihood the latency would be highly variable depending on contention for the WiFi bandwidth available in the room. XBee modules are essentially just modems that translate between RS-232 on a physical wire and the 802.15.4 wireless communication standard. They can be configured for point-to-point communication, but even easier is just to have all the transmitter stations XBee modules on the same channel, so that each character sent by one transmitter station is broadcast and received by all the others. While the configuration tool for XBee modules can be strangely finicky, once the XBees are configured, using them to communicate could not be simpler: they merely need to be provided with 5 V power and ground, and connected with two wires for sending and receiving data. Since we were able to obtain an exclusive channel for use in the lab, and since we are not sending much data between the transmitter stations, there is little contention in our radio communication network, and we do not expect latencies to vary due to data being held in buffers, waiting to be transmitted. As shown in Figure 2, each transmitter station has an Arduino with a shield connecting it to an XBee module. We were able to use the Arduino SoftwareSerial library to communicate between each Arduino via the XBees.

4 Signals

Measuring varying time of flights of a transmitted signal could be done by measuring the phase of a periodic signal, or calculating and finding the maximum of a cross correlation function between the received signal and the signal the transmitter is known to have transmitted, for some arbitrary transmitted signal. GPS works using the latter approach with pseudorandom signals, as this maximizes the time resolution of the time of flight measurements. In our case, we felt it would not be feasible to encode much information in the signals we send out, as there is a fair amount of distortion and echoing as the sound travels through

the air, as well as ringing in the transmitter and receiver that would distort the signal. Therefore, we took a much simpler approach: we measure the time of the first rising edge for each ultrasonic ping signal received. This time should correspond to the time it took for the ultrasonic ping to travel the line-of-sight distance between the transmitter station and the receiver. The exact duration of echos and ringing after the first rising edge then does not matter.

4.1 Transmitted signals and synchronization

To allow the car to know where it was in the room, we developed a system of four transmitters stations, one in each corner of the room, that send ultrasonic pulses at well-defined intervals. Each of the transmitters is pointed towards the opposite side of the room. This is because the transmitter is directional, that is, the strength of the signal depends on the angle of transmission. In this configuration, the receiver will be able to receive the pings if they are directly underneath the transmitter simply because the signal is stronger close to the transmitter, and if they are further away, the angle will help improve reception. If the last three pulses are sent at known times with respect to the first pulse during each series of pulses, then the time differences of arrival of each pulse at the receiver can be used to determine the car's position. Thus, all of the transmitters need to be synchronized such that they always send an ultrasonic pulse at the same time with respect to the first pulse, sent by the master transmitter. All of the code for the transmitter stations is given in Appendix A.4.

4.1.1 Synchronization

While radio waves are fast, XBee communication has latency much higher than the time for the propagation of radio waves. To ensure synchronization, the first step was to determine the latency in serial communication between the master transmitter and each of the slave transmitters. To find the latencies, the master transmitter sends a character (either 'a', 'b', or 'c', depending on which slave transmitter it wishes to communicate with) to the radio channel, then waits to receive the same character in return from the correct slave transmitter. Each slave transmitter is programmed to immediately respond to its own lowercase character (again 'a' for the first slave transmitter, 'b' for the second, and 'c' for the third) by transmitting the same character back to the channel. The master transmitter is then able to calculate the round-trip latency by keeping track of the time to send and receive the character. We found typical round-trip latencies to vary from about 17 ms to 20 ms. We suspect much of this time is due not necessarily to the XBees, but to the software and hardware stack on the Arduino that sends and receives RS-232. The master transmitter tests the round-trip latency for a slave transmitter five times, and divides the total of the round-trip latencies by ten to get an estimate for the one-way latency to each slave transmitter station. It then transmits another character, this time an uppercase letter corresponding to the slave transmitter it wishes to communicate with ('A' for the first slave transmitter, 'B' for the second, and 'C' for the third). Upon sending the uppercase character, the master transmitter then sends a series of four bytes encoding what it measured the one-way latency

for the slave transmitter to be. The slave transmitter then sends back an acknowledgment character (the same uppercase character that it received from the master to start the transaction) to tell the master that it has received all of the bytes. We make the assumption here that the round-trip latency is symmetric — that the time taken to transmit from master to slave is the same as to transmit from slave to master. Since all the transmitter stations have identical hardware, we find this assumption to be quite reasonable, and we did a few quick tests to confirm the assumption. When repeated for each of the three transmitters (five round-trip latency tests and transmission of the average one-way latency time), this process ensures near-perfect synchronization. Although latencies can occasionally change, we found that they usually stayed relatively constant when the transmitter stations are not being disturbed, meaning the average round-trip latency is not skewed and the one-way latencies are fairly accurate.

4.1.2 Sending Pings

With the latency measurement and communication process completed, the master transmitter sends out a lowercase ‘p’ over its XBee to the other transmitters to tell them that it is about to send a pulse. It then sends out its ultrasonic pulse, using two pins connected to the actual ultrasonic transmitter via a twisted pair. To gain more control over the ultrasonic pulse output, we used low-level pin manipulation to turn the pins of the Arduino connected to the ultrasonic transmitter on and off. In this way, we were able to simultaneously turn on one pin while turning off another, which cannot be done using the standard `digitalWrite()` command. The two pins are turned on and off opposite to each other, allowing for twice the voltage range. One pin would be switched on to +5V while the other was switched to ground, and then vice versa. With this manipulation, the differential signal sent to the ultrasonic transducer is a square wave with peak to peak voltage of 10V, double what we would achieve if we tied one pin of the transmitter to ground and only switched the other one. We also have good control over timing using this method, as one very fast register write switches both pins, and can use this control to send the pings at exactly 25 kHz, the optimal rated frequency for the transmitters and receivers that we purchased.

When the master sends out the ‘p’ character followed by its ping, each of the slave transmitters receive the ‘p’ character a few milliseconds later according to their respective latencies. They can then start timing how long it has been since the master transmitter first began transmitting its ping. Once the required time has elapsed (100 ms intervals between each ping), the slave transmitter begins transmitting its ping through the same low-level pin manipulation discussed above. Thus, the first slave transmitter waits 100 ms after receiving the send pulse character from the master minus the known latency time that it has received from the master plus a small amount of processing time due to the fact that the master must first write the ‘p’ character and then send the ping (the actions are not simultaneous on the master end). Each consecutive slave transmitter waits an extra 100 ms more before sending its pulse, again subtracting the known latency time and adding the processing delay. 100 ms was chosen as the delay time between each ping because while the transmission of the ping itself lasts only 5 ms, the echoes created from the signal bouncing around the room

cause the received signal to stretch out much more, sometimes up to about 25 ms, and if the receiver is across the room from one transmitter while very close to another, the signal from the transmitter far away could take as much as 30 ms longer to reach the receiver. We want enough time between pings received such that there is no doubt about the reception of each ping, so 100 ms is a safe period of time between transmissions. The overall result is four 25 kHz ultrasonic pings sent out 100 ms apart.

After one cycle of four pings sent out, the master then waits 300 ms before beginning the next cycle of latency tests and pings. This delay is necessary to allow the receiver to distinguish which ping is the ping sent from the master transmitter and which are the pings sent from each of the slave transmitters: TX0, the master transmitter, always transmits first after a quiet period of at least a few hundred milliseconds. Thus, the total cycle time between when the master sends one of its pings and the next must be greater than around 500 ms assuming the 100 ms between pings so that the receiver can distinguish the master ping based on the longer delay time between the last ping and the master ping as compared to the normal 100 ms time between pings. Since the latency tests also take time, in the final version of our system, the transmitters send out a round of pings every two thirds of a second or so. However, if improvements were to be made to speed up the transmission such as to get more position calculations per second on the receiving end, both the delay between pings and the delay between sets of pings could be reduced. Also important is that the slave transmitters sit idle after transmitting their ping. Thus, the master can begin doing latency testing on the first slave transmitter immediately after it sends its ping but before the second slave transmitter sends its ping, effectively interweaving pinging and latency testing so that latency tests do not impinge on the total timing of the system. In effect, then, the overall frequency of transmission cycles is only limited by the pulse-stretching experienced by the receiver and the time needed to distinguish the master ping, not by latency test times.

4.1.3 Timeouts and Error Checking

As a note, we found that bytes are sometimes lost over the course of communicating between devices on the radio channel. This loss can be an issue if it causes the system to hang while, for example, the master looks for an acknowledgment from one of the slaves or one of the bytes containing the latency time is never received by the slave. To combat this issue, we created a timeout structure; if the master is ever waiting for a response from one of the slaves (either an acknowledgment or as part of a latency test), it will automatically timeout after 50 ms. If the timeout occurs while testing for latency, then that test is thrown out and not counted as part of the averaging. If an acknowledgment is not received, the master will simply move on without it. On the slave side, a timeout can also occur when receiving the bytes corresponding the latency from the master. If one of these bytes is lost, the slave simply uses the last known latency time and ignores the corrupted new data. This error checking and timeout structure helps to make the system more robust so that if a transmitter turns off or bytes are lost, the entire system can recover and continue operating without manual resetting.

4.2 Digital processing of received signals

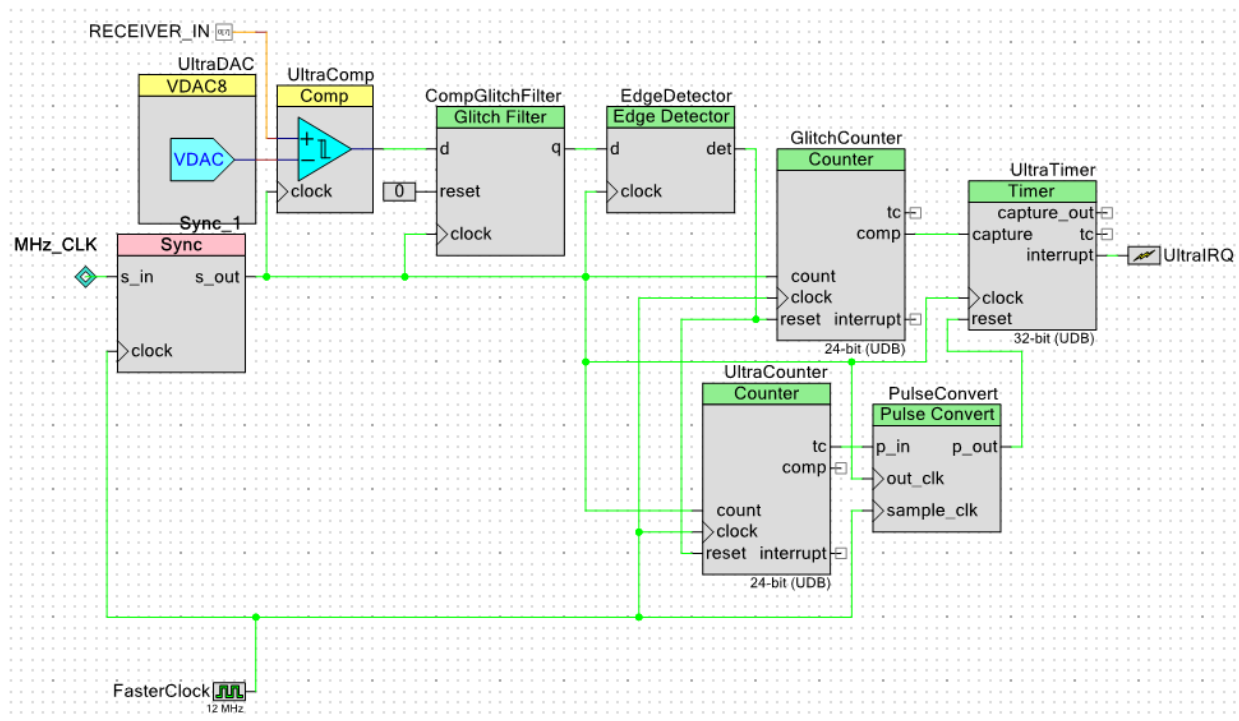


Figure 6: Programmable hardware used to capture timing information for positioning.

4.2.1 Digital filtering

Once the signal has been received, amplified, and filtered using the circuit described in section 3.1.2, we are left with the signal shown as the yellow waveform in Figure 7. To turn this waveform into a series of four clean square pulses so that timing can easily be done requires several pieces of programmable hardware. The schematic detailing the programmable hardware and connections referenced in this section can be found in Figure 6.

The first step was to use a comparator to turn the analog input into a series of digital pulses so that the edge time was rigidly defined and voltage levels could be raised even further from about the peak 2V received to 5V. We used a comparator level of 0.7V, high enough to avoid the comparator staying high for various echoes of the pings that had been sent out, but low enough to give accurate timing, since we want to record when the ping first reaches the receiver. This level of 0.7V was generated with an 8-bit voltage digital-to-analog converter. The output after the comparator step is shown as the blue waveform in Figure 7.

As is evident from the waveform representing the output of the comparator, due to the inconsistent nature of the analog waveform being received, further processing is necessary to generate a single pulse corresponding to the reception of each ping. Thus, we have a




Figure 7: Ch1 (Yellow): Output from the receiver circuitry that is fed into the PSoC. Ch2 (Blue): Output of the comparator.

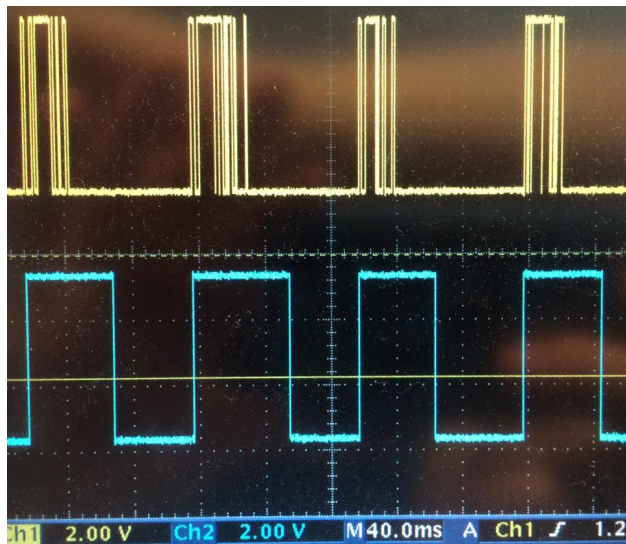


Figure 8: Ch1 (Yellow): Output of the comparator. Ch2 (Blue): Output of the filtering system that is fed into the timer.

system of two glitch filters to further refine the signal so that it is suitable for timing. A very useful, but somewhat unknown component, the glitch filter can be used either to eliminate small dips in a high signal or short spikes in the midst of a low signal. The former can be done by setting the glitch filter to bypass high mode such that its output remains high for a set period after seeing a high value while the latter can be achieved in the normal mode, equivalent to bypass low.

The first glitch filter (CompGlitchFilter in Figure 6), set to normal mode, directly takes

the output of the comparator as input and has a very short period of only 40 μs , which is the period of a 25 kHz signal. This eliminates possible spikes that do not correspond to pings, possibly generated from small amounts of noise. However, it will not eliminate any significant part of the pings, since they always last longer than 40 μs for the initial spike when they are received.

The second glitch filter is made up of two components, an edge detector and a counter (labeled GlitchCounter in Figure 6). This glitch filter removes the occasional times when the signal goes low in the middle of the ping signal, as seen in the comparator waveform (yellow) in Figure 8. We made our own glitch filter out of two components here due to a limitation of the Glitch Filter component available in PSoC Creator: it can only maintain a moving window of length at most 256. Thus, with a clock at 1 MHz, its maximum filtering time would be only 0.256 ms, much shorter than a possible drop in the signal, which we have measured at up to 10 ms in length. To allow for a longer filtering time, the clock can be slowed, but this could cause a loss of precision in our timing measurements, since we ideally want a filtering time of about 25 ms (or about twice as long as possible drops in the signal), which would require a clock speed of about 10 kHz. With a 10 kHz clock speed, timing is only accurate to 0.1 ms, corresponding to about one tenth of a foot. That's not so bad at all, but it seems there's no reason that timing should be even close to a limiting factor in measurement precision. To get around this filtering length/measurement precision tradeoff, we created our own glitch filter using an edge detector and a counter. The edge detector outputs pulses on both edges of its input (essentially the output of the comparator). These are used to reset the counter, so the counter is always counting the number of cycles of the 1 MHz clock that have passed since the last edge. The counter outputs a high pulse if its count value is less than 25000 (corresponding to 25 ms with a 1 MHz clock). Thus, the output of the counter goes high as soon as it gets the initial edge and stays high even as the signal drops low for small periods of time. Since the high parts of the signal never last for more than 25 ms, this combination of components effectively imitates the function of a glitch filter. The output of this two-part filter can be seen in Figure 8, where the yellow waveform represents the input to the filtering (the comparator output) and the blue waveform represents the output of the filtering system: four clean pulses.

4.2.2 Timing

Now all that is left is timing the positive edges of the filtered received signal. To do this, we use a Timer component, named UltraTimer in Figure 6. Very conveniently, it can hold a queue of four values before triggering an interrupt.

In case some bad input gets past the filter, to make sure the queue is empty when the next round of pings starts, we reset the timer during the intervening quiet period between each set of four pings. To identify that quiet period, we have a counter component named UltraCounter that counts the number of cycles of a 1 MHz clock since the last time an edge was seen in the received signal, much like GlitchCounter does. Unlike GlitchCounter, UltraCounter counts up to a maximum value of 200000, or 200 ms. We use the terminal count output of UltraCounter as a reset signal for UltraTimer, to ensure UltraTimer's queue

is emptied at some point during the quiet period, which is more than 200 ms long. 200 ms is longer than the delay between two received pings within a set of four pings should be, so this reset only happens during the reset period before the master transmitter sends out its next ping. Since the UltraCounter is a UDB Counter component, it counts edges of its "count" input, and must be operating at a higher speed clock. To convert the terminal count signal pulse from the faster clock domain to the 1 MHz clock domain used by the timer, a pulse converter component is used.

When UltraTimer sees four positive edges on the filtered received signal and captures those four times, it triggers an interrupt handler that reads those captured values and converts the time differences into differences in distances in feet. Before doing so, some sanity checks on the captured values are performed. First, since a timer reset should occur every cycle of four pings, we know something is wrong if the captured timer value indicates it has been a long time since the last reset of the timer. If any captured timer value indicates it was captured more than a second since the previous measurement, something has gone wrong and we throw out the entire measurement. Furthermore, we throw away the entire measurement if any of the differences in distances are much larger than the dimensions of the rectangle of transmitters. Only if the measurements make it through these sanity checks do we carry out the positioning math described in section 2.2.2.

5 Testing and debugging scaffolds

Before we were doing integration testing of our positioning system, we developed various testing tools to test one part of the system by simulating inputs as if from the rest of the system. Here we will briefly discuss a few such test setups.

First, to test ultrasonic reception and the basic ability to time the difference between two pings, we had `tx'test'2`. In this test, one Arduino Uno controls two transmitters, and alternately sends out pings from them. We put the transmitters on the ends of long twisted wire pairs, so that we could move them close and further away, and see whether we could measure the time difference between when the pings are received changed. Different receiver circuits could then be tested. The code for this test is included in appendix section C.1.

To test the positioning math given relatively clean reception, we had `rx'test`, whose code is included in appendix section C.3. Here, the Arduino sends out a series of four square pulses spaced unevenly, looking a lot like the blue waveform in Figure 8. A wire would be connected from the Arduino output pin directly to the PSoC, and the Arduino is essentially pretending to be the receiver circuitry, so that there were no actual ultrasonic components involved in this test setup. We could calculate what the timings of the received pulses should be at arbitrarily chosen locations, program the Arduino to send pulses with those timings, and then check whether the PSoC was correctly calculating the position we chose.

Finally, we had a version of `rx'test` that actually uses ultrasonics, named `tx'test'4`. The code for `tx'test'4`, given in appendix section C.2, would run on an Arduino connected to a single transmitter. The Arduino would send out a series of four pings with certain timings. The actual ultrasonic receiver circuit would be plugged in to the PSoC, and we would see

whether the PSoC was correctly calculating the position those timings should indicate. Since all the pings are being sent out a single transmitter, they all experience the same time of flight to the receiver, so the time of flight does not introduce any differences to the relative timings of the received pings. This could be tested with the transmitter on the other side of the room, and essentially the one transmitter is pretending to be four synchronized transmitters at different locations.

We did do some initial testing of radio communication by hooking Arduinos with XBees to separate laptops and communicating across the room. The XBees always just seemed to work, and are rated to have a range of at least a few hundred feet, which is more than the length of the room. Having seen this, testing of the synchronization of transmitter stations was done simply by bringing all the transmitter stations to a single lab bench and using an oscilloscope to look at the relative timing of the pings they emitted. This did not require any special code. As practically none of the latency we are concerned with arises from the time for propagation of radio signals, we assumed that if the transmitters can correctly synchronize when they are within a few feet of each other, as assume they can still more or less correctly synchronize when they are spread around the room. The fact that positioning worked fairly well when we had the whole system running indicates this assumption was good enough.

6 Performance and sources of bad data

When the car is sitting stationary and has a clear line of sight to all four transmitter stations, the positions it will calculate from repeated measurements will generally vary by a few inches, and a new measurement occurs after each round of pings, which happens once every two thirds of a second or so. However, there are several ways the positioning system can be foiled and produce poor results.

If something is blocking the direct line of sight between the car and a transmitter station, generally the ultrasonic ping from the transmitter station will still reach the car, because ultrasound has no trouble traveling around obstacles. However, this will delay the first rising edge at the receiver for this ping, leading the car to believe it is further from that transmitter station and ruining the positioning calculation.

Sharp noises with appreciable high-frequency content, especially loud clapping, are readily observed to cause the ultrasonic transducers at the receivers to ring briefly as if receiving an ultrasonic ping. This does not appear to be a problem in our receiver circuitry, but rather has to do with the resonance of the physical transducer itself. Usually, such extra pings will fail the sanity checks or the final error metric check, and the whole measurement will be thrown out, so the problem is not that this introduces error into positioning calculations, but that it delays the next new position information. In our testing, one person clapping continuously a couple yards away from the receiver could indefinitely prevent the car from getting a position fix.

If the car is moving, because the signals from the four transmitter stations are received at different points in time, they will also be received when the car receiver is at different points

in space. All of our positioning calculations described in Section 2 assume all measurements pertain to a single unknown location of the receiver, $(x, y, 0)$. When the car is moving, this assumption is false, and as shown in Figure 11, this can lead to errors that are much larger than when stationary. Intuitively, one might expect errors on the same order as the distance the car travels during a round of pings, which last a few hundred milliseconds. Thus, if the car is traveling a couple feet per second, one can expect errors on the order of a foot or more.

7 Applications

Having built an indoor ultrasonic positioning system, we set out to demonstrate its capabilities.

7.1 Navigation

One application of our positioning system would be to create an autonomous navigation system, whereby a car could travel to defined waypoints, either making stops at desired locations or tracing out a desired path. With radio communication, cars could interact and travel routes around each other, avoiding collisions, or following a path previously traveled by another car.

7.1.1 H-Bridge




Figure 9: The board of the H-bridge circuit.

To enable the car to navigate robustly around the room, the ability to drive the car in reverse would be helpful, so we added an H-bridge circuit (as seen in Figure 9) to our motor control. The H-bridge makes it possible for the PSoC to run the motor in both directions.




Figure 10: Schematic of the H-bridge circuit..

The gate of each transistor corresponds to a control signal from the PSoC. Inputs A and B controlling the NMOS's can come directly from the PSoC. Because the PMOS transistors had to deliver 7.2V to the motor and the PSoC only outputs a maximum of 5V, we need to put some hardware between the PSoC output and the gate of the PMOS's to raise the signal voltage to 7.2V. For one car, we chose the LM311P comparator IC to output 7.2V when the PSoC signal is above a reasonable threshold of 3.6V and output 0V when the PSoC signal is under 3.6V. For the other car, we simply ran the PSoC output to simple single-transistor inverters, constructed out of a general-purpose 2N7000 n-type pull down transistor and a 1 k Ω pull-up resistor.

With such an H-bridge circuit, driving forward can be done by, for example, turning on the PMOS controlled by A' and sending a PWM signal to A. Backwards driving can be done by turning the PMOS controlled by B' and sending a PWM signal to B. Rheostatic braking can also be achieved, for example by simultaneously turning on both NMOS's while leaving the PMOS's off. In the software that controls all this, we ensure any time we're switching between any of these three modes, (forwards, backwards, and braking), all the transistors are first turned off for a millisecond, to guarantee there is no time at which a short circuit path exists. This can be seen in appendix section A.4.

7.1.2 General autonomous navigation

Unfortunately, figuring out how to get the car to navigate to an arbitrary (x, y) position using only the positioning system proved more difficult than initially expected. If we knew the orientation of the car, getting to an arbitrary position seems fairly straightforward. At any given time, take the desired orientation to be the direction between the position of the car and the desired position. If the difference between the current orientation and desired orientation is small, we could simply perform PID or any other form of feedback control to point the car towards the desired destination. If the difference is large, we could code special

turning maneuvers for the car to rotate itself around.

Unfortunately, the only way we had to figure out the orientation of the car in general is to make some movement larger than the effective spatial resolution of our positioning system, and look at the change in position. After identifying an initial orientation, in principle one could keep updating an estimate of orientation by a combination of odometry and continuing to look at the direction in which the most recent positions have moved. Unfortunately, we started investigating this relatively late, and ran out of time before coming up with a working solution. The noise in our positioning is such that merely taking our orientation to be the direction of the change between the two most recent positions is far too noisy. In general, it seems estimating orientation based on a fixed-length history of recent positions might work well for some speeds, but would do poorly if the speed of the car varied, and a special approach would have to be taken when the car comes to a stop. It seems possible in principle to come up with a good algorithm to look at a dynamically adjusting number of previous points and determining the recent trend in change direction, but we have not found this algorithm. As for odometry, we have the Hall effect sensor to tell us about our speed, and we always know something about the current steering angle as we have control over that. Unfortunately, it seems that steering is sufficiently nonlinear and/or the Hall effect sensor doesn't provide enough resolution that orientation could be tracked accurately through turns. Using inertial sensing (i.e., a gyroscope) would probably help a lot.

7.1.3 Simpler navigation demonstration

Of course, to demonstrate the ability to use the positioning system to navigate, it is not truly necessary to have the car autonomously navigate itself to an arbitrary (x, y) position from any arbitrary position and orientation. It is quite easy to start the car in a known orientation in a known area, and have it move forward or backward until it sees its x - or y -coordinate reach a certain threshold, before it takes some other action, such as changing direction. Having the car oscillate between certain y -coordinates, or stop and go at certain locations, turns out to be quite simple.

7.2 Mapping

For the mapping application of our positioning system, we again decided to use the XBee radio modules to communicate between the PSoC board on the car and our computer. The goal was to have the positioning information collected by the car sent to the computer so that it could be plotted in real time. To accomplish this goal, we used an XBee module on the car with a shield similar to those used for the Arduino transmitters to connect to the PSoC and a XBIB-U-DEV board to connect the XBee module to the computer via USB. These XBees were configured such that they were using a different channel to communicate than the XBees used in the transmitter stations to avoid possible interference that could derail both systems of communication.

On the car side, all communication was coordinated by the PSoC through the provided UART module. Using this module, we can set TX and RX pins to send and receive informa-

tion via the XBee connected to those designated pins. Whenever data is generated from the receiver, it is written to the UART module and sent over the XBee to the XBee attached to the computer. The position string sent to the computer is always the same: 'X' then current x-coordinate, then 'Y' then current y-coordinate. An example string, if the car were located at (2.41, -5.34), would be "X2.41Y-5.34".




Figure 11: A partial map plotted in Matlab using coordinates received from the car as it drove around the racetrack. At about $(-5, 0)$ you can see where the car veered off the track, and was manually returned to the track, before it continued along the track. On the other side of the curve, at around $(-7, 7)$, there is one noticeable bad data point, which appeared to be caused by a human blocking line of sight between the car and one of the transmitter stations.

On the computer side, we used Matlab to communicate via serial with the XBee due to its relatively simple serial port I/O commands and ability to easily turn points received into a plot. To create a real-time connection between the car, we first connect to the correct COM port and open it using the correct Baud Rate. Then we set up the plot according to the parameters of our grid such that the plot will be to the same scale of the room. We used

the `animatedline` object in Matlab to integrate new information into the plot as we received it from the car. The `animatedline` is very efficient for this application as it does not require the user to store each individual point in an array and update it each time a new point is received. Instead, the `animatedline` object has an `addPoint()` function that allows the user to add new points to the line that are then stored in the object dynamically, requiring only a refresh of the plot and no vector separately storing points. Then, as each string containing the data points is received via the attached XBee, we can parse the string according to the known formatting and extract the x and y coordinates and feed them into the `animatedline` for a dynamic plot of real-time position. The Matlab script used to generate these plots can be found in Appendix D.

Just as the Arduino code for the transmitters required timeouts and other checks to ensure a robust system, due to occasional data loss, this system also needs similar checks to be sure that it does not quit or produce undesirable results when it receives corrupted data. Thus, we have implemented various checks on the input, such as whether the string contains all four elements in the correct order (character, float, character, float), whether the string contains any information, and a timeout to determine if the connection to the car has been lost. Using the `fgetl()` function, we can take one line of input at a time to process, since we look for a newline character to terminate each string sent from the car. However, one problem that arose with this strategy was that sometime either in the transmission or receiving process, extra characters could be appended onto the front of the string that were either random or whitespace, throwing off our parsing. To avoid these characters being included in the parsing process, we do an initial paring of the string down to just the necessary information before the string format check; specifically, we iterate through the string until the 'X' character is found, signifying the beginning of the relevant coordinate string. With these error checks, the script runs smoothly, plotting the points received effectively and exiting gracefully by closing the serial port if an error occurs.

A C code compiled in PSoC Creator for car

Cypress's APIs use capitalized letters at the start of words. Therefore, whenever we define our own functions, we use only lowercase letters, which helps distinguish functions we write from the provided APIs, and reduces the chance of a name collision.

A.1 main.c

```
1 /* =====
2  * main.c
3  * Victor A. Ying and Monica Lu
4  *
5  * This is the main program. It contains very little code.
6  * =====
7  */
8
9 #include <project.h>
10 #include <stdio.h>
11
12 #include "usb_uart.h"
13 #include "shell.h"
14 #include "drive.h"
15 #include "speed.h"
16 #include "steer.h"
17 #include "position.h"
18
19 /*
20  * MAIN PROGRAM
21  */
22 int main(void) {
23     // Turn on interrupts
24     CyGlobalIntEnable;
25
26     // Initialize USB serial connection for I/O
27     // usb_uart_init(USBUART_5V_OPERATION); //!!NOTE!! Make sure this matches your
        board voltage!
28
29     // Initialize LCD display
30     LCD_Start();
31
32     // Initialize radio communication
33     UART_Start();
34
35     // Begin positioning
```

```

36     position_init();
37
38     // Initialize navigation stuff
39     drive_init();
40
41     // Main loop performs user interface/communication actions.
42     // Other actions are performed in interrupt handlers.
43     for (;;) {
44         //     shell_handle_received_chars();
45         //     speed_display_info();
46
47         // Display position to LCD
48         if (position_data_available()) {
49             char buf[32], radiobuf[32];
50             float x, y;
51             uint16 counter = 0u;
52             uint8 status = CyEnterCriticalSection();
53             x = position_x();
54             y = position_y();
55             sprintf(radiobuf, "X%.2fY%.2f\n", x, y);
56             UART_PutString(radiobuf);
57             /*
58             LCD_Position(0,0);
59             LCD_PrintNumber(counter++);
60             sprintf(buf, "Error:%.2f ", error());
61             LCD_Position(0,4);
62             LCD_PrintString(buf);
63             sprintf(buf, "X:%.2f Y:%.2f ", x, y);
64             LCD_Position(1,0);
65             LCD_PrintString(buf);
66             */
67             CyExitCriticalSection(status);
68         }
69     }
70 }
71
72 /* [] END OF FILE */

```

A.2 position.c

```
1 /* =====
2 *
3 * position.c
4 * Michael Danielczuk, Andrew Kim, Monica Lu, and Victor Ying
5 *
6 * Provides positioning using time difference of arrival
7 * multilateration with four transmitters arranged in a
8 * rectangle. Assumes the transmitters send out pings in turn
9 * with a certain amount of spacing time between when pings are
10 * sent, and that these pings are sent in a counterclockwise
11 * order.
12 *
13 * =====
14 */
15
16 #include <project.h>
17 #include <math.h>
18 #include <limits.h>
19 #include <stdio.h>
20
21 #include "position.h"
22
23
24 /*
25 * CONSTANTS
26 */
27
28 #define X 23.5 // distance between first and second transmitters in feet
29 #define Y 33.75 // distance between second and third transmitters in feet
30 #define Z 7.583
31 #define CLOCK_FREQ 1000000 // Hz
32 #define WAVE_SPEED 1135.0 // ft/s
33 #define TX_SPACING 100 // ms
34 #define EPSILON 0.5 // ft
35 #define DEL_FACTOR 0.1 // ft
36 #define MAX_ERROR 0.5 // ft^2
37 #define ERROR_THRESHOLD 0.01 // ft^2
38 #define MAX_ITERATIONS 100
39
40 //#define SHOW_GARBAGE // Uncomment this to check if sanity checks are failing
41 #define PRINT_CONVERGENCE // Comment this out to make positioning silent
42
43 /*
```

```

44 * STATIC FUNCTION PROTOTYPES
45 */
46
47 static CY_ISR_PROTO(positioningHandler) ;
48
49
50 /*
51 * GLOBAL VARIABLES
52 */
53
54 static float x = 0.0, y = 0.0; // the current position
55 static float fxy = 0.0; // the current error
56 static uint8 new_data = 0u; // Boolean indicating whether new data available
57
58
59 /*
60 * FUNCTIONS
61 */
62
63 /*
64 * position_init:
65 * Start positioning.
66 */
67 void position_init(void) {
68     UltraCounter_Start();
69     GlitchCounter_Start();
70     UltraTimer_Start();
71     UltraComp_Start();
72     UltraDAC_Start();
73     UltraIRQ_Start();
74     UltraIRQ_SetVector(positioningHandler);
75 }
76
77 /*
78 * position_data_available:
79 * returns nonzero if new data since the last time this function was called.
80 */
81 uint8 position_data_available(void) {
82     uint8 status = CyEnterCriticalSection();
83     uint8 ret = new_data;
84     new_data = 0u;
85     CyExitCriticalSection(status);
86     return ret;
87 }
88

```

```

89 /*
90 * position_?:
91 * Getter functions for position in units of feet, with the origin at the
92 * center of the rectangle formed by the transmitters.
93 */
94 float position_x(void) {
95     return x;
96 }
97 float position_y(void) {
98     return y;
99 }
100
101 float error(void) {
102     return fxy;
103 }
104
105 float fabsf(float num) {
106     if (num >= 0.0)
107         return num;
108     else
109         return -num;
110 }
111
112 /*
113 * positioningHandler:
114 * Interrupt handler run after sequence of four pings, calculating position.
115 */
116 static CY_ISR(positioningHandler) {
117     uint32 time[4];
118     float diff[4];
119     int i, iters;
120     float new_x, new_y, new_fxy;
121
122     // Get the times of arrival
123     for (i = 0; i < 4; i++) {
124         time[i] = UltraTimer_ReadCapture();
125
126         // If more than a second since the last reset, then throw away this
127         // set of measurements
128         if (time[i] == 0u || time[i] < ULONG_MAX - CLOCK_FREQ) {
129 #ifdef SHOW_GARBAGE
130             x = (float)i;
131             y = (float)time[i];
132             new_data = 1u;
133 #endif

```

```

134     return;
135 }
136 }
137
138 // Calculate differences in distances in feet
139 for (i = 1; i < 4; i++) {
140     diff[i] = (float)((int32)(time[0] - time[i])
141                    - i*(CLOCK_FREQ/1000*TX_SPACING))
142            * (WAVE_SPEED/CLOCK_FREQ);
143
144     // If difference is much larger than the size of the rectangle of
145     // transmitter stations, the data is probably bad, so throw it away
146     if (fabsf(diff[i]) > X + Y) {
147
148 #ifdef SHOW_GARBAGE
149         x = (float)i;
150         y = diff[i];
151         new_data = 1u;
152 #endif
153         return;
154     }
155 }
156
157 // Positioning using Newton's method
158 new_x = x;
159 new_y = y;
160 iters = 0;
161 do {
162     float dfx, dfy, gradient_magnitude_squared;
163     float dist[4], error[4];
164     char buf[32];
165     uint8 status;
166
167     // Calculate what the distances should be based on our most recent (x,y)
168     dist[0] = sqrt((new_x+X/2)*(new_x+X/2) + (new_y+Y/2)*(new_y+Y/2) + Z*Z);
169     dist[1] = sqrt((new_x-X/2)*(new_x-X/2) + (new_y+Y/2)*(new_y+Y/2) + Z*Z);
170     dist[2] = sqrt((new_x-X/2)*(new_x-X/2) + (new_y-Y/2)*(new_y-Y/2) + Z*Z);
171     dist[3] = sqrt((new_x+X/2)*(new_x+X/2) + (new_y-Y/2)*(new_y-Y/2) + Z*Z);
172
173     // Calculate disagreement between hypothetical distances and measurements
174     for (i = 1; i < 4; i++)
175         error[i] = (dist[i]-dist[0]) - diff[i];
176
177     // Calculate our metric as the sum of the squares of the errors
178     new_fxy = 0.0;

```



```

179     for (i = 1; i < 4; i++)
180         new_fxy += error[i]*error[i];
181
182     // Calculate the partial derivatives of the metric
183     dfx = 2*error[2] * (((new_x - X/2)/dist[2]) - (new_x + X/2)/dist[0]);
184     dfx += 2*error[3] * (((new_x + X/2)/dist[3]) - (new_x + X/2)/dist[0]);
185     dfx += 2*error[1] * (((new_x - X/2)/dist[1]) - (new_x + X/2)/dist[0]);
186
187     dfy = 2*error[2] * (((new_y - Y/2)/dist[2]) - (new_y + Y/2)/dist[0]);
188     dfy += 2*error[3] * (((new_y - Y/2)/dist[3]) - (new_y + Y/2)/dist[0]);
189     dfy += 2*error[1] * (((new_y + Y/2)/dist[1]) - (new_y + Y/2)/dist[0]);
190
191     // Quit now if we're already at a stationary point
192     gradient_magnitude_squared = dfx*dfx + dfy*dfy;
193     if (gradient_magnitude_squared == 0.0)
194         break;
195
196     // Otherwise, update according to a version of Newton's method
197     new_x -= DEL_FACTOR * new_fxy * dfx / gradient_magnitude_squared;
198     new_y -= DEL_FACTOR * new_fxy * dfy / gradient_magnitude_squared;
199
200 #ifndef PRINT_CONVERGENCE
201     // Show convergence happening on the LCD
202     status = CyEnterCriticalSection();
203     sprintf(buf, "dX:%.1f dY:%.1f %d ", dfx, dfy, iters);
204     LCD_Position(1,0);
205     LCD_PrintString(buf);
206     sprintf(buf, "X:%.1f Y:%.1f ", new_x, new_y);
207     LCD_Position(0,0);
208     LCD_PrintString(buf);
209     sprintf(buf, " %.1f ", new_fxy);
210     LCD_Position(0,13);
211     LCD_PrintString(buf);
212     CyExitCriticalSection(status);
213 #endif
214
215     iters++;
216 } while ((fabsf(new_fxy) > ERROR_THRESHOLD) && (iters < MAX_ITERATIONS));
217
218 if (fabsf(new_fxy) < MAX_ERROR) {
219     x = new_x;
220     y = new_y;
221     fxy = new_fxy;
222     new_data = 1u;
223 }

```

```

224
225     // Clear interrupt
226     UltraTimer_ReadStatusRegister();
227 }
228
229
230 /* [] END OF FILE */

```

position.h

```

1 /* =====
2 *
3 * position.h
4 * Michael Danielczuk, Andrew Kim, Monica Lu, and Victor Ying
5 *
6 * Provides positioning using time difference of arrival
7 * multilateration with four transmitters arranged in a
8 * rectangle.
9 *
10 * =====
11 */
12
13 #ifndef POSITION_H
14 #define POSITION_H
15
16 #include <project.h>
17
18 /*
19 * position_init:
20 * Start positioning.
21 */
22 void position_init(void) ;
23
24 /*
25 * position_data_available:
26 * returns nonzero if new data since the last time this function was called.
27 */
28 uint8 position_data_available(void) ;
29
30 /*
31 * position_?:
32 * Getter functions for position in units of feet, with the origin at the
33 * center of the rectangle formed by the transmitters.
34 */

```

```
35 float position_x(void) ;
36 float position_y(void) ;
37
38 /*
39  * error:
40  * Getter function for error in units of feet squared.
41  */
42 float error(void) ;
43
44 #endif
45
46 /* [] END OF FILE */
```

A.3 drive.c

```
1 /* =====
2  * drive.c
3  * Monica Lu and Victor Ying
4  *
5  * High-level movement program.
6  * =====
7  */
8
9 #include <project.h>
10
11 #include "drive.h"
12 #include "speed.h"
13 #include "steer.h"
14
15
16 /*
17  * drive_init:
18  * Begins driving the car.
19  */
20 void drive_init(void) {
21     steer_init();
22     speed_init();
23
24     steer_pid_start();
25     speed_forward();
26     speed_pid_start("0.5");
27 }
28
29 /*
30  * magnet_callback:
31  * Called on every hall sensor tick with a running total distance traveled.
32  */
33 void magnet_callback(void) {
34     // Do about half a lap
35     if (distance_traveled >= 50.0)
36         speed_brake();
37
38     /*
39     // Do two laps and maximum speed
40     if (distance_traveled < 1.0)
41         ;
42     else if ((distance_traveled > 3.0 && distance_traveled < 20.0) ||
43             (distance_traveled > 100.0 && distance_traveled < 118.0))
```

```
44     speed_set(9.0);
45     else if (distance_traveled < 194.0)
46         speed_set(5.7);
47     else
48         speed_brake(); */
49 }
50
51
52 //[] END OF FILE
```

drive.h

```
1 /* =====
2  * drive.h
3  * Monica Lu and Victor Ying
4  *
5  * High-level movement program.
6  * =====
7  */
8
9 #ifndef DRIVE_H
10 #define DRIVE_H
11
12
13 /*
14  * drive_init:
15  * Begins driving the car.
16  */
17 void drive_init(void) ;
18
19 /*
20  * magnet_callback:
21  * Called on every hall sensor tick with a running total distance traveled.
22  */
23 void magnet_callback(void) ;
24
25 #endif
26
27 //[] END OF FILE
```

A.4 speed.c

```
1 /* =====
2 * speed.c
3 * Monica Lu and Victor Ying
4 *
5 * Speed measurement and control.
6 * =====
7 */
8
9 #include <project.h>
10 #include <stdio.h>
11 #include <stdlib.h>
12 #include <limits.h>
13
14 #include "speed.h"
15 #include "steer.h"
16 #include "drive.h"
17 #include "usb_uart.h"
18
19 #define DISTANCE_PER_TICK 0.1285 // in feet
20 #define DERIV_CONTROL_AVERAGING 3
21 #define TIMER_COUNTS_PER_SECOND 1000000.0 // timers are fed by 1 MHz clock
22 #define PID_INTERVALS_PER_SECOND 100.0 // PID control is reevaluated every 10 ms
23
24 enum state {
25     COAST = 0u,
26     FORWARD = 1u,
27     BACKWARD = 2u,
28     BRAKE = 3u,
29 };
30
31 // Because the speed sensor is on the right side of the car, it underestimates
32 // speed during right turns and overestimates during left turns. This is for a
33 // compensating correction term.
34 #define SPEED_ADJUSTMENT_COEFFICIENT (-0.12)
35
36
37 static CY_ISR_PROTO(hall_handler) ;
38 static CY_ISR_PROTO(speed_pid_handler) ;
39 static void speed_pid_control(void) ;
40
41 static enum state current_state;
42
43 float speed = 0.0;
```

```

44 float distance_traveled = 0.0;
45 float power_output = 0.0;
46
47 static uint16 count_since_magnet = 0u;
48 static uint8 speed_pid_enabled = 0u; // Boolean value indicating whether or not
    to do PID speed control.
49 static float speed_setpoint = 0.0; // in feet per second
50 static float kp = 0.1; // in normalized control output adjustment per (feet per
    second) error.
51 static float ki = 0.2; // in normalized control output adjustment per (feet per
    second) error per second.
52 static float kd = 0.0; // in normalized control output adjustment per (change in
    feet per second per second).
53
54
55 static void set_state(enum state desired) {
56     current_state = desired;
57     Drive_Control_Reg_Wakeup();
58     Drive_Control_Reg_Write(desired);
59 }
60
61
62 /*
63 * speed_init:
64 * Initializes timers, pwm, etc. for controlling speed. Initializes
65 * to coasting state.
66 */
67 void speed_init(void) {
68     uint8 status = CyEnterCriticalSection();
69
70     speed_coast();
71
72     Hall_Timer_Start();
73     Hall_IRQ_Start();
74     Hall_IRQ_SetVector(hall_handler);
75
76     Speed_PID_Timer_Start();
77     Speed_PID_IRQ_Start();
78     Speed_PID_IRQ_SetVector(speed_pid_handler);
79
80     Drive_PWM_Start();
81
82     CyExitCriticalSection(status);
83 }
84

```

```

85 /*
86  * speed_display_info:
87  * Print current speed and normalized controller output to the LCD.
88  */
89 void speed_display_info(void) {
90     char strbuf[17];
91
92     uint8 status = CyEnterCriticalSection();
93
94     // Print out info
95     sprintf(strbuf, "Speed: %.6f", speed);
96     LCD_Position(0,0);
97     LCD_PrintString(strbuf);
98     if (speed_pid_enabled) {
99         sprintf(strbuf, "Control: %.6f", power_output);
100        LCD_Position(1,0);
101        LCD_PrintString(strbuf);
102    }
103
104    CyExitCriticalSection(status);
105
106    CyDelay(50u);
107 }
108
109 /*
110  * speed_pid_start:
111  * Sets the speed to the floating point value represented by desired_speed,
112  * and enables PID speed control
113  */
114 void speed_pid_start(const char* desired_speed) {
115     uint8 status = CyEnterCriticalSection();
116
117     if (*desired_speed != '\0')
118         speed_setpoint = atof(desired_speed);
119     speed_pid_enabled = 1;
120
121     CyExitCriticalSection(status);
122 }
123
124 /*
125  * speed_set:
126  * Sets the speed_setpoint.
127  */
128 void speed_set(float desired) {
129     uint8 status = CyEnterCriticalSection();

```



```

130
131     speed_setpoint = desired;
132
133     CyExitCriticalSection(status);
134 }
135
136 /*
137 * speed_coast:
138 * Disables PID speed control, and sets main drive motor PWM compare value
139 * to 0, and sets state to coasting.
140 */
141 void speed_coast(void) {
142     uint8 status = CyEnterCriticalSection();
143
144     speed_pid_enabled = 0;
145     Drive_PWM_WriteCompare(0);
146     set_state(COAST);
147
148     CyExitCriticalSection(status);
149 }
150
151 /*
152 * speed_brake:
153 * Disables PID speed control, and sets main drive motor PWM compare value
154 * to 0, and sets state to braking.
155 */
156 void speed_brake(void) {
157     uint8 status = CyEnterCriticalSection();
158
159     speed_pid_enabled = 0;
160     Drive_PWM_WriteCompare(0);
161     if (current_state == FORWARD || current_state == BACKWARD) {
162         set_state(COAST);
163         CyDelay(1);
164     }
165     set_state(BRAKE);
166
167     CyExitCriticalSection(status);
168 }
169
170 /*
171 * speed_forward:
172 * Changes the state to forward motion.
173 */
174 void speed_forward(void) {

```

```

175     uint8 status = CyEnterCriticalSection();
176
177     if (current_state == BRAKE || current_state == BACKWARD) {
178         set_state(COAST);
179         CyDelay(1);
180     }
181     set_state(FORWARD);
182
183     CyExitCriticalSection(status);
184 }
185
186 /*
187 * speed_backward:
188 * Changes the state to forward motion.
189 */
190 void speed_backward(void) {
191     uint8 status = CyEnterCriticalSection();
192
193     if (current_state == BRAKE || current_state == FORWARD) {
194         set_state(COAST);
195         CyDelay(1);
196     }
197     set_state(BACKWARD);
198
199     CyExitCriticalSection(status);
200 }
201
202 /*
203 * speed_set_power:
204 * Sets the main drive motor PWM compare value.
205 */
206 void speed_set_power(const char *valstr) {
207     uint8 status = CyEnterCriticalSection();
208
209     if (speed_pid_enabled) {
210         usb_uart_putline(
211             "Cannot manually set power while PID control is running!");
212     }
213     else {
214         uint16 cmp;
215
216         power_output = atof(valstr);
217         if (power_output > 1.0)
218             power_output = 1.0;
219         else if (power_output < 0.0)

```

```

220         power_output = 0.0;
221         cmp = (uint16)(power_output * USHRT_MAX);
222         Drive_PWM_WriteCompare(cmp);
223     }
224
225     CyExitCriticalSection(status);
226 }
227
228 /*
229  * speed_set_k*:
230  * Set PID control constants.
231  */
232 void speed_set_kp(const char *val) {
233     uint8 status = CyEnterCriticalSection();
234     kp = atof(val);
235     CyExitCriticalSection(status);
236 }
237 void speed_set_ki(const char *val) {
238     uint8 status = CyEnterCriticalSection();
239     ki = atof(val);
240     CyExitCriticalSection(status);
241 }
242 void speed_set_kd(const char *val) {
243     uint8 status = CyEnterCriticalSection();
244     kd = atof(val);
245     CyExitCriticalSection(status);
246 }
247
248
249 /*
250  * hall_handler:
251  * Should run every time a magnet is seen.
252  * Recalculates the current speed estimate.
253  */
254 static CY_ISR(hall_handler) {
255     static uint32 prev;
256     uint32 val = Hall_Timer_ReadCapture();
257     uint32 diff = prev - val;
258     float time;
259     uint8 status;
260
261     prev = val;
262
263     time = (float)diff / TIMER_COUNTS_PER_SECOND;
264     speed = DISTANCE_PER_TICK / time; // in units of feet/second

```

```

265 speed /= 1.0 + steer_output * SPEED_ADJUSTMENT_COEFFICIENT;
266
267 count_since_magnet = 0;
268
269 distance_traveled += DISTANCE_PER_TICK /
270     (1.0 + steer_output * SPEED_ADJUSTMENT_COEFFICIENT);
271
272 status = CyEnterCriticalSection();
273 magnet_callback();
274 CyExitCriticalSection(status);
275
276 // Clear interrupt
277 Hall_Timer_ReadStatusRegister();
278 }
279
280 /*
281 * speed_pid_handler:
282 * Should run once every 10 ms interval. Does speed control.
283 */
284 static CY_ISR(speed_pid_handler) {
285     uint8 saved_interrupt_status;
286
287     // If a magnet is not seen for a long time, decrease the speed estimate
288     if (count_since_magnet > 0) {
289         float time_since_magnet = count_since_magnet / PID_INTERVALS_PER_SECOND;
290         float speed_estimate = DISTANCE_PER_TICK / time_since_magnet;
291         speed_estimate /= 1.0 + steer_output * SPEED_ADJUSTMENT_COEFFICIENT;
292         if (speed_estimate < speed)
293             speed = speed_estimate;
294     }
295     count_since_magnet++;
296
297     saved_interrupt_status = CyEnterCriticalSection();
298     if (speed_pid_enabled &&
299         (current_state == FORWARD || current_state == BACKWARD))
300         speed_pid_control();
301     CyExitCriticalSection(saved_interrupt_status);
302 }
303
304 static void speed_pid_control(void) {
305     float error, deriv, next_riemann_sum,
306         prop_control, integ_control, deriv_control;
307     static float riemann_sum = 0.0;
308     static float prev_errors[DERIV_CONTROL_AVERAGING];
309     static uint8 prev_errors_index = 0;

```

```

310 float change;
311 uint16 pwm_cmp;
312 uint8 status;
313
314 // Begin control calculations
315 error = speed_setpoint - speed;
316
317 // Proportional control
318 prop_control = kp * error;
319
320 // Integral control
321 next_riemann_sum = riemann_sum + error / PID_INTERVALS_PER_SECOND;
322 integ_control = ki * next_riemann_sum;
323
324 // Derivative control
325 change = error - prev_errors[prev_errors_index];
326 deriv = change * PID_INTERVALS_PER_SECOND / (float)DERIV_CONTROL_AVERAGING;
327 prev_errors[prev_errors_index] = error;
328 prev_errors_index = (prev_errors_index + 1) % DERIV_CONTROL_AVERAGING;
329 deriv_control = kd * deriv;
330
331 // Add it all together, with limiting to valid duty cycle values
332 power_output = 0.15 + prop_control + integ_control + deriv_control;
333 if (power_output < 0.0) {
334     power_output = 0.0;
335 }
336 else if (power_output > 1.0) {
337     power_output = 1.0;
338 }
339 else {
340     // Anti-windup: only allow integrator to build up if not saturating
341     riemann_sum = next_riemann_sum;
342 }
343
344 // Scale control value to a PWM compare value
345 pwm_cmp = (uint16)(USHRT_MAX * power_output);
346 status = CyEnterCriticalSection();
347 Drive_PWM_WriteCompare(pwm_cmp);
348 CyExitCriticalSection(status);
349 }
350
351
352 //[] END OF FILE

```

speed.h

```
1 /* =====
2 * speed.h
3 * Monica Lu and Victor Ying
4 *
5 * Speed measurement and control.
6 * =====
7 */
8
9 #ifndef SPEED_H
10 #define SPEED_H
11
12
13 extern float speed;
14 extern float distance_traveled;
15 extern float power_output;
16
17
18 /*
19 * speed_init:
20 * Initializes timers, interrupts, pwm, etc. for speed control. Initializes
21 * to coasting state.
22 */
23 void speed_init(void) ;
24
25 /*
26 * speed_display_info:
27 * Print current speed and normalized controller output to the LCD.
28 */
29 void speed_display_info(void) ;
30
31 /*
32 * speed_pid_start:
33 * Sets the speed to the floating point value represented by desired_speed,
34 * and enables PID speed control
35 */
36 void speed_pid_start(const char* desired_speed) ;
37
38 /*
39 * speed_set:
40 * Sets the speed_setpoint.
41 */
42 void speed_set(float speed_setpoint) ;
43
```

```

44 /*
45  * speed_brake:
46  * Disables PID speed control, and sets main drive motor PWM compare value
47  * to 0, and sets state to braking.
48  */
49 void speed_brake(void) ;
50
51 /*
52  * speed_coast:
53  * Disables PID speed control, and sets main drive motor PWM compare value
54  * to 0, and sets state to coasting.
55  */
56 void speed_coast(void) ;
57
58 /*
59  * speed_forward:
60  * Changes the state to forward motion.
61  */
62 void speed_forward(void) ;
63
64 /*
65  * speed_backward:
66  * Changes the state to forward motion.
67  */
68 void speed_backward(void) ;
69
70 /*
71  * speed_set_power:
72  * Sets the main drive motor PWM compare value.
73  */
74 void speed_set_power(const char *cmpstr) ;
75
76 /*
77  * speed_set_k*:
78  * Set PID control constants.
79  */
80 void speed_set_kp(const char *val) ;
81 void speed_set_ki(const char *val) ;
82 void speed_set_kd(const char *val) ;
83
84
85 #endif
86
87 //[] END OF FILE

```

A.5 steer.c

```
1 /* =====
2  * steer.c
3  * Monica Lu and Victor Ying
4  *
5  * Line detection and steering control.
6  * =====
7  */
8
9 #include <project.h>
10 #include <stdio.h>
11 #include <stdlib.h>
12
13 #include "steer.h"
14 #include "usb_uart.h"
15
16
17 #define FASTEST_CLK_FREQ 48000000.0 // 48 MHz
18 #define EXPECTED_ROW_LEN (0.000045 * FASTEST_CLK_FREQ) // in clock cycles
19 #define ACCEPTABLE_ROW_MARGIN 0.2
20 #define MIN_ACCEPTABLE_ROW_LEN ((uint16)(EXPECTED_ROW_LEN * (1.0 -
    ACCEPTABLE_ROW_MARGIN)))
21 #define MAX_ACCEPTABLE_ROW_LEN ((uint16)(EXPECTED_ROW_LEN * (1.0 +
    ACCEPTABLE_ROW_MARGIN)))
22 #define PID_INTERVALS_PER_SECOND 60.0 // Camera is 30 fps interlaced
23 #define DERIV_CONTROL_AVERAGING 4
24 #define STEERING_CENTER 1500 // 1.5 ms pulse = steer straight ahead
25
26
27 static CY_ISR_PROTO(camera_handler) ;
28 static void steer_pid_control(void) ;
29
30
31 float steer_output = 0.0;
32
33 static float measurement = 0.0;
34 static uint8 steer_pid_enabled = 0; // Boolean value indicating whether or not
    to do PID steering control.
35 static float kp = 0.75;
36 static float ki = 0.0;
37 static float kd = 0.03;
38
39
40 /*
```



```

41 * steer_init:
42 * Initializes timers, interrupts, pwm, etc. for steering the car.
43 */
44 void steer_init(void) {
45     steer_stop();
46
47     Camera_Comp_Start();
48     Camera_Counter_Start();
49     Camera_Timer_Start();
50     Camera_IRQ_Start();
51     Camera_IRQ_SetVector(camera_handler);
52
53     Steering_PWM_Start();
54 }
55
56 /*
57 * steer_display_info:
58 * Print current measurements and normalized controller output to the LCD.
59 */
60 void steer_display_info(void) {
61     char strbuf[17];
62
63     // Print out info
64     sprintf(strbuf, "Meas.: %f ", measurement);
65     LCD_Position(0,0);
66     LCD_PrintString(strbuf);
67     sprintf(strbuf, "Steer: %f ", steer_output);
68     LCD_Position(1,0);
69     LCD_PrintString(strbuf);
70     CyDelay(50u);
71 }
72
73 /*
74 * steer_pid_start:
75 * Enables PID steering control
76 */
77 void steer_pid_start(void) {
78     steer_pid_enabled = 1u;
79 }
80
81 /*
82 * steer_stop:
83 * Disables PID steering control.
84 */
85 void steer_stop(void) {

```

```

86     steer_pid_enabled = 0u;
87 }
88
89 /*
90 * steer_set:
91 * Sets the steering angle.
92 */
93 void steer_set(const char *valstr) {
94     uint8 status = CyEnterCriticalSection();
95
96     if (steer_pid_enabled) {
97         usb_uart_putline(
98             "Cannot manually set steering while PID control is running!");
99     }
100    else {
101        uint16 cmp;
102        steer_output = atof(valstr);
103        if (steer_output > 1.0)
104            steer_output = 1.0;
105        else if (steer_output < -1.0)
106            steer_output = -1.0;
107        cmp = (int16)(500.0 * steer_output) + STEERING_CENTER;
108        Steering_PWM_WriteCompare(cmp);
109    }
110
111    CyExitCriticalSection(status);
112 }
113
114 /*
115 * steer_set_k*:
116 * Set PID control constants.
117 */
118 void steer_set_kp(const char *val) {
119     kp = atof(val);
120 }
121 void steer_set_ki(const char *val) {
122     ki = atof(val);
123 }
124 void steer_set_kd(const char *val) {
125     kd = atof(val);
126 }
127
128 static CY_ISR(camera_handler) {
129     uint16 row_start_time = Camera_Timer_ReadCapture();
130     uint16 black_start_time = Camera_Timer_ReadCapture();

```

```

131  uint16 black_end_time = Camera_Timer_ReadCapture();
132  uint16 row_end_time = Camera_Timer_ReadCapture();
133  uint16 black_mid_time, row_mid_time;
134  uint8 saved_interrupt_status;
135
136  // Only update our measurement of where the black strip is if the row looks
137  // like a full row containing a single black strip. (This effectively tosses
138  // out data at intersections.)
139  uint16 row_length = row_start_time - row_end_time;
140  if (row_length >= MIN_ACCEPTABLE_ROW_LEN &&
141      row_length <= MAX_ACCEPTABLE_ROW_LEN) {
142      black_mid_time = black_end_time + (black_start_time - black_end_time)/2u;
143      row_mid_time = row_end_time + row_length/2u;
144
145      measurement = (float)(int16)(black_mid_time - row_mid_time) * 2 /
146                    (float)row_length;
147  }
148
149  // Do PID steering control
150  saved_interrupt_status = CyEnterCriticalSection();
151  if (steer_pid_enabled)
152      steer_pid_control();
153  CyExitCriticalSection(saved_interrupt_status);
154
155  // Clear interrupt
156  Camera_Timer_ReadStatusRegister();
157  }
158  /*
159  * steer_pid_control:
160  * Should run once after every new video frame. Does steering control.
161  */
162  static void steer_pid_control(void) {
163      float error, deriv, next_riemann_sum,
164            prop_control, integ_control, deriv_control;
165      static float riemann_sum = 0.0;
166      static float prev_errors[DERIV_CONTROL_AVERAGING] = {0};
167      static uint8 prev_errors_index = 0;
168      float change;
169      uint16 pwm_cmp;
170      uint8 status;
171
172      // Begin control calculations
173      error = - measurement;
174

```

```

175 // Proportional control
176 prop_control = kp * error;
177
178 // Integral control
179 next_riemann_sum = riemann_sum + error/PID_INTERVALS_PER_SECOND;
180 integ_control = ki * next_riemann_sum;
181
182 // Derivative control
183 change = error - prev_errors[prev_errors_index];
184 deriv = change * PID_INTERVALS_PER_SECOND / (float)DERIV_CONTROL_AVERAGING;
185 prev_errors[prev_errors_index] = error;
186 prev_errors_index = (prev_errors_index + 1) % DERIV_CONTROL_AVERAGING;
187 deriv_control = kd * deriv;
188
189 // Add it all together, with limiting to valid duty cycle values
190 steer_output = prop_control + integ_control + deriv_control;
191 if (steer_output < -1.0) {
192     steer_output = -1.0;
193 }
194 else if (steer_output > 1.0) {
195     steer_output = 1.0;
196 }
197 else {
198     // Anti-windup: only allow integrator to build up if not saturating
199     riemann_sum = next_riemann_sum;
200 }
201
202 // Scale control value to a PWM compare value
203 pwm_cmp = (int16)(500.0 * steer_output) + STEERING_CENTER;
204 status = CyEnterCriticalSection();
205 Steering_PWM_WriteCompare(pwm_cmp);
206 CyExitCriticalSection(status);
207 }
208
209 //[] END OF FILE

```

steer.h

```

1 /* =====
2 * steer.h
3 * Monica Lu and Victor Ying
4 *
5 * Line detection and steering control.
6 * =====

```

```

7  */
8
9  #ifndef STEER_H
10 #define STEER_H
11
12
13 extern float steer_output; // Ranges from -1.0 to 1.0.
14
15
16 /*
17  * steer_init:
18  * Initializes timers, interrupts, pwm, etc. for steering the car.
19  */
20 void steer_init(void) ;
21
22 /*
23  * steer_display_info:
24  * Print current steering info to the LCD.
25  */
26 void steer_display_info(void) ;
27
28 /*
29  * steer_pid_start:
30  * Enables PID steering control
31  */
32 void steer_pid_start(void) ;
33
34 /*
35  * steer_stop:
36  * Disables PID steering control.
37  */
38 void steer_stop(void) ;
39
40 /*
41  * steer_set:
42  * Sets the steering angle.
43  */
44 void steer_set(const char *valstr) ;
45
46 /*
47  * steer_set_steer_k*:
48  * Set PID control constants.
49  */
50 void steer_set_kp(const char *val) ;
51 void steer_set_ki(const char *val) ;

```

```
52 void steer_set_kd(const char *val) ;  
53  
54  
55 #endif  
56  
57 //[] END OF FILE
```

A.6 shell.c

```
1 /* =====
2  * shell.c
3  * Victor A. Ying and Monica Lu
4  *
5  * a very minimal shell.
6  * =====
7  */
8
9 #include <project.h>
10 #include <stdio.h>
11 #include <string.h>
12 #include <stdlib.h>
13
14 #include "speed.h"
15 #include "steer.h"
16 #include "usb_uart.h"
17 #include "drive.h"
18
19 /*
20  * vshell_do_command()
21  * A function for parsing and dispatching command lines.
22  * Must be passed a null-terminated string!
23  */
24 void shell_do_command(const char* line) CYREENTRANT {
25     char cmd[SHELL_MAX_COMMAND_LENGTH + 1];
26     uint8 i = 0u;
27
28     // Ignore leading whitespace
29     while (isspace(*line))
30         line++;
31
32     // Identify command as the longest prefix not containing a space.
33     while (*line != '\0' && !isspace(*line)) {
34         cmd[i] = *line;
35         i++;
36         line++;
37     }
38     cmd[i] = '\0';
39
40     // Throw away any whitespace between the command and the arguments
41     while (isspace(*line))
42         line++;
43
```

```

44 // Now we can execute the identified command
45
46 // Do nothing in case of empty command
47 if (strcmp(cmd, "") == 0)
48     ;
49
50 // Shell built-in commands
51 else if (strcmp(cmd, "repeat") == 0) {
52     char numloops[SHELL_MAX_COMMAND_LENGTH];
53     i = 0;
54
55     while (*line != '\0' && !isspace(*line)) {
56         numloops[i] = *line;
57         i++;
58         line++;
59     }
60     numloops[i] = '\0';
61     while (isspace(*line))
62         line++;
63
64     for (i = 0; i < atoi(numloops); i++) {
65         shell_do_command(line);
66     }
67 }
68
69 // Other commands
70 else if (strcmp(cmd, "serial") == 0) {
71     usb_uart_print_info();
72 }
73 else if (strcmp(cmd, "lcd") == 0) {
74     LCD_ClearDisplay();
75     LCD_PrintString(line);
76 }
77 else if (strcmp(cmd, "pid") == 0) {
78     speed_pid_start(line);
79 }
80 else if (strcmp(cmd, "brake") == 0) {
81     speed_brake();
82 }
83 else if (strcmp(cmd, "coast") == 0) {
84     speed_coast();
85 }
86 else if (strcmp(cmd, "fw") == 0) {
87     speed_forward();
88 }

```



```

89     else if (strcmp(cmd, "bw") == 0) {
90         speed_backward();
91     }
92     else if (strcmp(cmd, "power") == 0) {
93         speed_set_power(line);
94     }
95     else if (strcmp(cmd, "speedkp") == 0) {
96         speed_set_kp(line);
97     }
98     else if (strcmp(cmd, "speedki") == 0) {
99         speed_set_ki(line);
100    }
101    else if (strcmp(cmd, "speedkd") == 0) {
102        speed_set_kd(line);
103    }
104    else if (strcmp(cmd, "steerpid") == 0) {
105        steer_pid_start();
106    }
107    else if (strcmp(cmd, "steerstop") == 0) {
108        steer_stop();
109    }
110    else if (strcmp(cmd, "steerset") == 0) {
111        steer_set(line);
112    }
113    else if (strcmp(cmd, "steerkp") == 0) {
114        steer_set_kp(line);
115    }
116    else if (strcmp(cmd, "steerki") == 0) {
117        steer_set_ki(line);
118    }
119    else if (strcmp(cmd, "steerkd") == 0) {
120        steer_set_kd(line);
121    }
122    // If command was not any of the above...
123    else {
124        char8 strbuf[128];
125
126        sprintf(strbuf, "Command \"%s\" not recognized", cmd);
127        usb_uart_putline(strbuf);
128    }
129 }
130
131 /*
132 * shell_handle_char()
133 * A function for handling keypresses from a

```

```

134 * terminal, and building up command lines.
135 */
136 void shell_handle_char(char c) {
137     static char linebuf[SHELL_MAX_COMMAND_LENGTH + 1];
138     static uint8 line_index = 0;
139
140     if (c < 127 && c >= 32) { // If the user pressed a character
141         #ifdef SHELL_LCD_DEBUG
142             // For debugging purposes, display the character to the LCD
143             LCD_ClearDisplay();
144             LCD_PutChar(c);
145         #endif
146
147         if (line_index < SHELL_MAX_COMMAND_LENGTH) {
148             // Add character to current line
149             linebuf[line_index] = c;
150             line_index++;
151
152             // Echo character to terminal.
153             usb_uart_putchar(c);
154         }
155     }
156     else if (c == '\b' || c == '\177') { // If the user pressed backspace
157         #ifdef SHELL_LCD_DEBUG
158             // For debugging purposes, display to the LCD
159             LCD_ClearDisplay();
160             LCD_PrintString("Backspace");
161         #endif
162
163         if (line_index > 0) {
164             // Clear one character on the terminal
165             usb_uart_putstr("\b \b");
166
167             // Discard one character from the command buffer
168             line_index--;
169         }
170     }
171     else if (c == '\r') { // If the user has hit the Return/Enter key
172         #ifdef VSHELL_LCD_DEBUG
173             // For debugging purposes, display to the LCD
174             LCD_ClearDisplay();
175             LCD_PrintString("Return");
176         #endif
177
178         // Echo newline to terminal

```

```

179     usb_uart_putCRLF();
180
181     // Do command
182     linebuf[line_index] = '\0';
183     shell_do_command(linebuf);
184
185     // Print a prompt to indicate ready for next command
186     usb_uart_putstring("$ ");
187     line_index = 0;
188 }
189 else { // For any other characters/keypress
190     #ifdef SHELL_LCD_DEBUG
191     char8 strbuf[128];
192
193     // For aid in debugging, print out ASCII codes in hex
194     sprintf(strbuf, " %X", (int)c);
195     LCD_PrintString(strbuf);
196     #endif
197 }
198 }
199
200 /*
201 * shell_handle_recieved_chars()
202 * For every character that has been received, calls handler with
203 * that character. Does nothing if no data has been received since
204 * the previous time this function was called.
205 */
206 void shell_handle_received_chars(void) {
207     while(USBUART_DataIsReady()) { // Check for input data from PC
208         uint8 i, count, buffer[128];
209
210         count = USBUART_GetData(buffer, sizeof(buffer)); // Get any data from PC
211         for (i = 0; i < count; i++) { // For each character recieved
212             shell_handle_char(buffer[i]); // Pass the character to the handler
213         }
214     }
215 }
216
217 /* [] END OF FILE */

```

shell.h

```

1 /* =====
2 * shell.h

```

```

3 * Victor A. Ying and Monica Lu
4 *
5 * a very minimal shell.
6 * =====
7 */
8
9 #ifndef SHELL_H
10 #define SHELL_H
11
12 #define SHELL_MAX_COMMAND_LENGTH 63
13
14 /*
15 #define SHELL_LCD_DEBUG // Uncomment this line to print keypresses to the LCD
16 */
17
18 /*
19 * shell_do_command()
20 * A function for parsing and dispatching command lines.
21 * Must be passed a null-terminated string!
22 */
23 void shell_do_command(const char* line) CYREENTRANT ;
24
25 /*
26 * shell_handle_char()
27 * For every character that has been received, calls handler with
28 * that character. Does nothing if no data has been received since
29 * the previous time this function was called.
30 */
31 void shell_handle_char(char c) ;
32
33 /*
34 * shell_handle_recieved_chars()
35 * For every character that has been received, calls handler with that character.
36 * Does nothing if no data has been received since the previous time this
37 * function
38 * was called.
39 */
40 void shell_handle_received_chars(void) ;
41
42 #endif
43 //[] END OF FILE

```

A.7 usb_uart.c

```
1 /* =====
2 * usb_uart.c
3 * Victor A. Ying and Monica Lu
4 *
5 * Manage the serial connection over USB
6 * =====
7 */
8
9 #include <project.h>
10 #include <stdio.h>
11
12 #include "usb_uart.h"
13
14
15 /*
16 * usb_uart_init()
17 * Initialize USBUART for I/O. Should be called after CYGlobalIntEnable;.
18 * NOTE: Make sure mode matches your board voltage!
19 */
20 void usb_uart_init(uint8 mode) {
21     USBUART_Start(0, mode);
22     while(!USBUART_GetConfiguration())
23         ;
24     USBUART_CDC_Init();
25 }
26
27 /*
28 * usb_uart_print_info()
29 * Sends information about the USBUART connection to the USBUART connection.
30 */
31 void usb_uart_print_info(void) {
32     char8 strbuf[128];
33
34     // USBUART info
35     uint32 rate = USBUART_GetDTERate();
36     uint8 stop_bits = USBUART_GetCharFormat();
37     uint8 parity = USBUART_GetParityType();
38     uint8 data_bits = USBUART_GetDataBits();
39     sprintf(strbuf,
40         "UART over USB data rate is %lu bits per second", rate);
41     usb_uart_putline(strbuf);
42     sprintf(strbuf,
43         "UART over USB characters consist of %i data bits",
```

```

44     (int)data_bits);
45     usb_uart_putline(strbuf);
46     sprintf(strbuf, "UART over USB parity mode; %s",
47             (parity == USBUART_PARITY_NONE) ? "none" :
48             (parity == USBUART_PARITY_ODD) ? "odd" :
49             (parity == USBUART_PARITY_EVEN) ? "even" :
50             (parity == USBUART_PARITY_MARK) ? "mark" :
51             (parity == USBUART_PARITY_SPACE) ? "space" :
52             "?ERROR?");
53     usb_uart_putline(strbuf);
54     sprintf(strbuf, "UART over USB stop bits: %s",
55             (stop_bits == USBUART_1_STOPBIT) ? "1" :
56             (stop_bits == USBUART_1_5_STOPBITS) ? "1.5" :
57             (stop_bits == USBUART_2_STOPBITS) ? "2" :
58             "?ERROR?");
59     usb_uart_putline(strbuf);
60 }
61
62 /*
63  * usb_uart_put*()
64  * Wrappers around Cypress's provided USBUART_Put*() functions
65  * that wait for previous Tx to finish before continuing.
66  */
67 void usb_uart_putdata(const uint8* pData, uint16 length) {
68     while(!USBUART_CDCIsReady())
69         ;
70     USBUART_PutData((uint8 *)pData, length);
71 }
72 void usb_uart_putstring(const char8* str) {
73     while(!USBUART_CDCIsReady())
74         ;
75     USBUART_PutString((char8 *)str);
76 }
77 void usb_uart_putchar(char8 txDataByte) {
78     while(!USBUART_CDCIsReady())
79         ;
80     USBUART_PutChar(txDataByte);
81 }
82 void usb_uart_putCRLF(void) {
83     while(!USBUART_CDCIsReady())
84         ;
85     USBUART_PutCRLF();
86 }
87
88 /*

```

```

89 * usb_uart_putline()
90 * An additional helpful function for printing a string followed by CRLF.
91 */
92 void usb_uart_putline(const char8* str) {
93     usb_uart_putstring(str);
94     usb_uart_putCRLF();
95 }
96
97 /* [] END OF FILE */

```

usb_uart.h

```

1 /* =====
2 * usbuart_helpers.h
3 * Victor A. Ying and Monica Lu
4 *
5 * Functions to manage the serial
6 * connection over USB.
7 * =====
8 */
9
10 #ifndef USB_UART_H
11 #define USB_UART_H
12
13 #include "USBUART.h"
14 #include "shell.h"
15
16 /*
17 * usb_uart_init()
18 * Initialize USBUART (and possibly LCD) for I/O
19 * NOTE: Make sure mode matches your board voltage!
20 */
21 void usb_uart_init(uint8 mode);
22
23 /*
24 * usb_uart_print_info()
25 * Sends information about the USBUART connection to the USBUART connection.
26 */
27 void usb_uart_print_info(void);
28
29 /*
30 * usb_uart_put*()
31 * Wrappers around Cypress's provided USBUART_Put*() functions
32 * that wait for previous Tx to finish before continuing.

```

```
33 */
34 void usb_uart_putdata(const uint8* pData, uint16 length);
35 void usb_uart_putstr(const char8* string);
36 void usb_uart_putchar(char8 txDataByte);
37 void usb_uart_putCRLF(void);
38
39 /*
40 * usb_uart_putline()
41 * An additional helpful function for printing a string followed by CRLF.
42 */
43 void usb_uart_putline(const char8* string);
44
45
46 #endif
47
48 //[] END OF FILE
```

B Arduino code for transmitter stations

B.1 Master transmitter station

```
1 /*****
2 master_transmitter.ino
3
4 Communicate with slaves to establish radio communication latency,
5 then start off the sequence of ultrasonic pings.
6
7 Hardware Hookup:
8   The XBee Shield makes all of the connections you'll need
9   between Arduino and XBee. If you have the shield make
10  sure the SWITCH IS IN THE "DLINE" POSITION. That will connect
11  the XBee's DOUT and DIN pins to Arduino pins 2 and 3.
12 *****/
13 // We'll use SoftwareSerial to communicate with the XBee:
14 #include <SoftwareSerial.h>
15
16 #define NUM_TRANSMITTERS 3
17 #define NUM_TESTS 5
18 #define TX_PORT PORTB
19
20 enum {
21   TX_PIN_1 = 12,
22   TX_PIN_2 = 13,
23   TX_PIN_1_MASK = 1u << (TX_PIN_1 % 8u),
24   TX_PIN_2_MASK = 1u << (TX_PIN_2 % 8u),
25   FREQ = 25000u, // Hz
26   PERIOD = 1000000u / FREQ, // s
27   HALF_PERIOD = PERIOD / 2u, // s
28   TIMEOUT = 50000u, // s
29   DURATION = 5000u, // s
30   BAUD_RATE = 9600u, // bps
31 };
32
33
34 // XBee's DOUT (TX) is connected to pin 2 (Arduino's Software RX)
35 // XBee's DIN (RX) is connected to pin 3 (Arduino's Software TX)
36 SoftwareSerial XBee(2, 3); // RX, TX
37
38 void setup()
39 {
40   pinMode(TX_PIN_1, OUTPUT);
41   pinMode(TX_PIN_2, OUTPUT);
```

```

42
43 // Set up both ports at 9600 baud. This value is most important
44 // for the XBee. Make sure the baud rate matches the config
45 // setting of your XBee.
46 XBee.begin(BAUD_RATE);
47 Serial.begin(BAUD_RATE);
48 }
49
50 void loop()
51 {
52   int i, j;
53   unsigned long startTime, totalLatTime, averageLatTime, beginning;
54   byte b[4];
55   unsigned long startTotTime = micros();
56
57   for (i = 0; i < NUM_TRANSMITTERS; i++) {
58     totalLatTime = 0u;
59     int successCount = 0;
60     for (j = 0; j < NUM_TESTS; j++) {
61       startTime = micros();
62       XBee.write((char)('a' + i));
63       while (!XBee.available() && (micros() - startTime < TIMEOUT));
64       if (XBee.available()) {
65         XBee.read();
66         totalLatTime += (micros()-startTime);
67         successCount++;
68       }
69       //Serial.println(totalLatTime);
70     }
71     if (successCount > 0) {
72       averageLatTime = totalLatTime/(2*successCount);
73       Serial.print("Average Time: ");
74       Serial.println(averageLatTime);
75       XBee.write((char)('A' + i));
76       LongToBytes(averageLatTime, b);
77       for (j = 0; j < 4; j++) {
78         XBee.write(b[j]);
79       }
80       beginning = micros();
81       while (!XBee.available() && (micros()-beginning < TIMEOUT));
82       Serial.print("Acknowledgement: ");
83       Serial.write(XBee.read());
84       Serial.println();
85     }
86   }

```

```

87  XBee.write('p');
88  sendPing();
89  Serial.print("Total Time: ");
90  Serial.println(micros()-startTotTime);
91  delay(300);
92  }
93
94  void LongToBytes(unsigned long val, byte b[4]) {
95    b[0] = (byte)((val >> 24) & 0xff);
96    b[1] = (byte)((val >> 16) & 0xff);
97    b[2] = (byte)((val >> 8) & 0xff);
98    b[3] = (byte)(val & 0xff);
99  }
100
101  /*
102  unsigned long BytesToLong(byte b[4]) {
103    return ((unsigned long)(b[0] & 0xFF) << 24) + ((unsigned long)(b[1] & 0xFF) <<
104            16)
105            + ((unsigned long)(b[2] & 0xFF) << 8) + (unsigned long)(b[3] & 0xFF);
106  }
107  */
108  void sendPing() {
109
110    //Serial.println("Sending ping");
111
112    // Send out ping from transmitter
113    unsigned long beginning = micros();
114    unsigned long time = micros() - beginning;
115    unsigned long next = time;
116    while (time < DURATION) {
117      // Turn off TX_PIN_2 and turn on TX_PIN_1
118      TX_PORT = TX_PORT & (~TX_PIN_2_MASK) | TX_PIN_1_MASK;
119      next += HALF_PERIOD;
120      while (micros() - beginning < next)
121        ;
122
123      // Turn off TX_PIN_1 and turn on TX_PIN_2
124      TX_PORT = TX_PORT & (~TX_PIN_1_MASK) | TX_PIN_2_MASK;
125      next += HALF_PERIOD;
126      while ((time = micros() - beginning) < next)
127        ;
128    }
129    Serial.println("Ping sent");
130  }

```

```
131 /*
132 // ASCIItoInt
133 // Helper function to turn an ASCII hex value into a 0-15 byte val
134 int ASCIItoInt(char c)
135 {
136     if ((c >= '0') && (c <= '9'))
137         return c - 0x30; // Minus 0x30
138     else if ((c >= 'A') && (c <= 'F'))
139         return c - 0x37; // Minus 0x41 plus 0x0A
140     else if ((c >= 'a') && (c <= 'f'))
141         return c - 0x57; // Minus 0x61 plus 0x0A
142     else
143         return -1;
144 }
145 */
```

B.2 Slave transmitter stations

```
1 /*****
2 slave_transmitter.ino
3
4 Responds to messages from the master and sends out ultrasonic pings
5 at the appropriate time.
6
7 Hardware Hookup:
8 The XBee Shield makes all of the connections you'll need
9 between Arduino and XBee. If you have the shield make
10 sure the SWITCH IS IN THE "DLINE" POSITION. That will connect
11 the XBee's DOUT and DIN pins to Arduino pins 2 and 3.
12 *****/
13 // We'll use SoftwareSerial to communicate with the XBee:
14 #include <SoftwareSerial.h>
15 #define TX_PORT PORTB
16 #define TRANSMITTER_NUMBER 1
17
18 enum {
19     TX_PIN_1 = 12,
20     TX_PIN_2 = 13,
21     TX_PIN_1_MASK = 1u << (TX_PIN_1 % 8u),
22     TX_PIN_2_MASK = 1u << (TX_PIN_2 % 8u),
23     FREQ = 25000u, // Hz
24     PERIOD = 1000000u / FREQ, // s
25     HALF_PERIOD = PERIOD / 2u, // s
26     DURATION = 5000u, // s
27     TIMEOUT = 50000u, // s
28     BAUD_RATE = 9600u, // bps
29     MAX_LAT_TIME = 20000u, // s
30     SOFTWARE_SERIAL_DELAY = 2000u, // s
31 };
32
33 // XBee's DOUT (TX) is connected to pin 2 (Arduino's Software RX)
34 // XBee's DIN (RX) is connected to pin 3 (Arduino's Software TX)
35 SoftwareSerial XBee(2, 3); // RX, TX
36 unsigned long latTime;
37
38 void setup()
39 {
40     // Set up both ports at 9600 baud. This value is most important
41     // for the XBee. Make sure the baud rate matches the config
42     // setting of your XBee.
43     XBee.begin(BAUD_RATE);
```

```

44 //Serial.begin(BAUD_RATE); // Serial for debugging
45
46 pinMode(TX_PIN_1, OUTPUT);
47 pinMode(TX_PIN_2, OUTPUT);
48 }
49
50 void loop()
51 {
52   int i;
53   char c;
54   byte b[4];
55
56   if (XBee.available())
57   { // If data comes in from XBee, send it out to serial monitor
58     c = XBee.read();
59     if (c == (char)('a' + TRANSMITTER_NUMBER - 1)) {
60       XBee.write(c);
61     }
62     else if (c == (char)('A' + TRANSMITTER_NUMBER - 1)) {
63       for (i = 0; i < 4; i++) {
64         unsigned long beginning = micros();
65         while (!XBee.available() && (micros()-beginning < TIMEOUT));
66         if (XBee.available()) {
67           b[i] = XBee.read();
68         }
69         else {
70           break;
71         }
72       }
73       if (i == 4) {
74         unsigned long temp = BytesToLong(b);
75         if (temp < MAX_LAT_TIME) {
76           latTime = temp;
77         }
78       }
79       XBee.write(c);
80     }
81     else if (c == 'p') {
82       delayMicroseconds(SOFTWARE_SERIAL_DELAY);
83       sendPing();
84     }
85   }
86 }
87
88 unsigned long BytesToLong(byte b[4]) {

```

```

89  unsigned long ret = 0u;
90  for (int i = 0; i < 4; i++) {
91      ret += (unsigned long)(b[i] & 0xFF) << 8*(3-i);
92  }
93  return ret;
94  }
95
96  void sendPing() {
97      unsigned long beginning = micros();
98      //Serial.println("Sending ping");
99      while(micros() - beginning < (100000u*TRANSMITTER_NUMBER)-latTime)
100          ;
101
102      // Send out ping from transmitter
103      beginning = micros();
104      unsigned long time = micros() - beginning;
105      unsigned long next = time;
106      while (time < DURATION) {
107          // Turn off TX_PIN_2 and turn on TX_PIN_1
108          TX_PORT = TX_PORT & (~TX_PIN_2_MASK) | TX_PIN_1_MASK;
109          next += HALF_PERIOD;
110          while (micros() - beginning < next)
111              ;
112
113          // Turn off TX_PIN_1 and turn on TX_PIN_2
114          TX_PORT = TX_PORT & (~TX_PIN_1_MASK) | TX_PIN_2_MASK;
115          next += HALF_PERIOD;
116          while ((time = micros() - beginning) < next)
117              ;
118      }
119
120      //Serial.println("Ping sent");
121  }

```

C Arduino code for testing tools

C.1 tx_test_2

```
1 // One Arduino controls two transmitters, alternately sending
2 // out pings from the two of them, simulating two synchronized
3 // transmitter stations sending out pings in turn.
4
5 #define TX_1_PORT PORTD
6 #define TX_2_PORT PORTB
7
8 enum {
9     TX_1_PIN_1 = 2,
10    TX_1_PIN_2 = 3,
11    TX_2_PIN_1 = 10,
12    TX_2_PIN_2 = 11,
13    TX_1_PIN_1_MASK = 1u << (TX_1_PIN_1 % 8u),
14    TX_1_PIN_2_MASK = 1u << (TX_1_PIN_2 % 8u),
15    TX_2_PIN_1_MASK = 1u << (TX_2_PIN_1 % 8u),
16    TX_2_PIN_2_MASK = 1u << (TX_2_PIN_2 % 8u),
17    FREQ = 25000u, // Hz
18    PERIOD = 1000000u / FREQ, // s
19    HALF_PERIOD = PERIOD / 2u, // s
20    DURATION = 5000u, // s
21    SPACING = 100000u, // s
22    CYCLE = 1000000u, // s
23 };
24
25 void setup() {
26     pinMode(TX_1_PIN_1, OUTPUT);
27     pinMode(TX_1_PIN_2, OUTPUT);
28     pinMode(TX_2_PIN_1, OUTPUT);
29     pinMode(TX_2_PIN_2, OUTPUT);
30 }
31
32 void loop() {
33     unsigned long beginning = micros();
34
35     // Send out ping from transmitter 1
36     unsigned long time = micros() - beginning;
37     unsigned long next = time;
38     while (time < DURATION) {
39         // Turn off TX_1_PIN_2 and turn on TX_1_PIN_1
40         TX_1_PORT = TX_1_PORT & (~TX_1_PIN_2_MASK) | TX_1_PIN_1_MASK;
41         next += HALF_PERIOD;
```



```

42  while (micros() - beginning < next)
43      ;
44
45  // Turn off TX_1_PIN_1 and turn on TX_1_PIN_2
46  TX_1_PORT = TX_1_PORT & (~TX_1_PIN_1_MASK) | TX_1_PIN_2_MASK;
47  next += HALF_PERIOD;
48  while ((time = micros() - beginning) < next)
49      ;
50  }
51
52  // Send out a ping from transmitter 2
53  while ((time = micros() - beginning) < SPACING)
54      ;
55  next = time;
56  while (time < SPACING + DURATION) {
57      // Turn off TX_2_PIN_2 and turn on TX_2_PIN_1
58      TX_2_PORT = TX_2_PORT & (~TX_2_PIN_2_MASK) | TX_2_PIN_1_MASK;
59      next += HALF_PERIOD;
60      while (micros() - beginning < next)
61          ;
62
63      // Turn off TX_2_PIN_1 and turn on TX_2_PIN_2
64      TX_2_PORT = TX_2_PORT & (~TX_2_PIN_1_MASK) | TX_2_PIN_2_MASK;
65      next += HALF_PERIOD;
66      while ((time = micros() - beginning) < next)
67          ;
68  }
69
70  // Wait for end of cycle
71  while (micros() - beginning < CYCLE)
72      ;
73  }

```

C.2 tx_test_4

```
1 // One Arduino controlling one transmitters, simulates four transmitters
2
3 #define TX_PORT PORTD
4
5 enum {
6     TX_PIN_1 = 2,
7     TX_PIN_2 = 3,
8     TX_PIN_1_MASK = 1u << (TX_PIN_1 % 8u),
9     TX_PIN_2_MASK = 1u << (TX_PIN_2 % 8u),
10    FREQ = 25000u, // Hz
11    PERIOD = 1000000u / FREQ, // s
12    HALF_PERIOD = PERIOD / 2u, // s
13    NUM_TRANSMITTERS = 4u,
14    DURATION = 5000u, // s
15 };
16
17 unsigned const long SPACINGS[NUM_TRANSMITTERS] =
18     {117766u, 95420u, 77814u, 709001u};
19
20 void setup() {
21     pinMode(TX_PIN_1, OUTPUT);
22     pinMode(TX_PIN_2, OUTPUT);
23 }
24
25 void loop() {
26     static unsigned i = 0;
27     unsigned long beginning = micros();
28
29     // Send out ping from transmitter
30     unsigned long time = micros() - beginning;
31     unsigned long next = time;
32     while (time < DURATION) {
33         // Turn off TX_PIN_2 and turn on TX_PIN_1
34         TX_PORT = TX_PORT & (~TX_PIN_2_MASK) | TX_PIN_1_MASK;
35         next += HALF_PERIOD;
36         while (micros() - beginning < next)
37             ;
38
39         // Turn off TX_PIN_1 and turn on TX_PIN_2
40         TX_PORT = TX_PORT & (~TX_PIN_1_MASK) | TX_PIN_2_MASK;
41         next += HALF_PERIOD;
42         while ((time = micros() - beginning) < next)
43             ;
```

```
44 }
45
46 // wait until time for next ping
47 unsigned long spacing = SPACINGS[i];
48 i = (i + 1) % NUM_TRANSMITTERS;
49 while (micros() - beginning < spacing)
50     ;
51 }
```

C.3 rx_test

```
1 // Mimics the receiver board, sending a repeating series of four pings
2 // as if from four transmitter stations to the PSoC, for the PSoC to
3 // attempt position calculations.
4
5 enum {
6     RX_PIN = 13,
7     NUM_TRANSMITTERS = 4u,
8     DURATION = 5000u, // s
9 };
10
11 unsigned const long SPACINGS[NUM_TRANSMITTERS] =
12     {117766u, 95420u, 77814u, 709001u};
13
14 void setup() {
15     pinMode(RX_PIN, OUTPUT);
16 }
17
18 void loop() {
19     static unsigned i = 0;
20     digitalWrite(RX_PIN, HIGH);
21     unsigned long beginning = micros();
22     while (micros() - beginning < DURATION)
23         ;
24     digitalWrite(RX_PIN, LOW);
25     unsigned long spacing = SPACINGS[i];
26     i = (i + 1) % NUM_TRANSMITTERS;
27     while (micros() - beginning < spacing)
28         ;
29 }
```

D MATLAB code for mapping

```
1 if (~isempty(instrfind))
2     fclose(instrfind);
3     delete(instrfind);
4     clear s;
5 end
6
7 s = serial('COM14');
8 s.BaudRate = 9600;
9 s.Timeout = 20;
10 s.Terminator = 'LF';
11 fopen(s);
12
13 figure(1);
14 xlim([-12 12]);
15 ylim([-17 17]);
16 pbaspect([1 1 1]);
17 clear h;
18 h = animatedline('Color','red','Marker','o');
19
20 while 1
21     try
22         [str, count, msg] = fgetl(s);
23         if (~isempty(msg))
24             error(msg);
25         end
26     catch err
27         err
28         disp('A timeout occurred or the user quit the program!')
29         break;
30     end
31
32     if str < 0
33         break;
34     end
35
36     for i=1:size(str)
37         if str(i) ~= 'X'
38             else
39                 str = str(i:end);
40                 break;
41             end
42         end
43     if strcmp(str, '')
```

```
44     continue;
45 end
46 [c, num] = sscanf(str, '%c%f%c%f', 4);
47 if num == 4
48     x = c(2);
49     y = c(4);
50     if (x ~= 0 || y ~= 0)
51         addpoints(h, x, y);
52         drawnow
53     end
54 end
55 end
56
57 fclose(s);
58 delete(s)
59 clear s
```
