Resilient Peer-to-Peer Ranging using Narrowband High-Performance Software-Defined Radios

Vincent Beech, Paul Groves, Paul Wright

Vincent Beech is a research student at University College London (UCL) working on robust ranging and timing using a narrowband ad-hoc radio system provided by Terrafix. He holds an MSci in Mathematical Science from Queen Mary University of London and is currently working towards a Ph.D in Engineering.
Dr Paul Groves is an Associate Professor at UCL, where he leads a program of research into robust positioning and navigation. He is interested in all aspects of navigation and positioning, including multi-sensor integrated navigation, improving GNSS performance under challenging reception conditions, and novel positioning techniques. He is author of the book Principles of GNSS, Inertial and Multi-Sensor Integrated Navigation Systems. He holds a BA/MA and a DPhil in Physics from the University of Oxford.
Dr Paul Wright is a Technical Development Engineer who provides support for the mobile data communications and navigation systems provided commercially by Terrafix. He has supported Bob Mason's (Terrafix Chief Scientific Officer) collaborative projects with external organisations in the development of navigation technologies and systems.

Abstract: In emergency and military scenarios, standard positioning and communications infrastructure cannot be relied upon. Ad-hoc mobile networks offer a decentralized alternative, so this paper addresses the challenge of incorporating a ranging capacity to an existing ad-hoc digital communications system. However, signal protocols and signal processing chains that are optimised for efficient communications do not easily support the precise timing that is required for accurate range measurements. For this system, there is the further challenge that it uses narrow bandwidth signals in the $10-20$kHz range. For perspective, each data symbol is 2-3 orders of magnitude larger than a chip of a Global Navigation Satellite System (GNSS) signal. Therefore, signal timing measurements must be accurate to a tiny fraction of a symbol if a useful positioning capability is to be achieved.

By applying innovative synchronization techniques to a High-Performance Software-Defined Radio (HPSDR) network, ranging data collected under benign conditions exhibits a standard deviation of only 2.61m despite using these very narrowband signals. This paper will describe how such a level of performance was achieved and verified experimentally.

This paper begins with an overview of the radio network used in this project and how the conclusion that two-way ranging would be the best choice of position fixing for this network rather than other options available. Then the paper continues with a summary of some preliminary work where range estimates have been computed using the radio's existing synchronization protocols. These primitive ranging protocols form the foundation upon which more advanced ranging protocols will be implemented. What follows is the body of the paper where the three key components of signal processing that led to the precise ranging achieved are discussed: timing recovery, non-coherent integration and clock drift correction. Finally, a breakdown of the best results achieved by implementing these key components are discussed followed by a brief discussion of further work that could be done.

1 Background

This project is focused on a Peer-to-Peer (P2P) Private Mobile Radio (PMR) network provided by Terrafix Limited. They're the leading supplier of mobile data and communications systems to the UK ambulance service, security services, police and other emergency based organisations [1]. The PMR network is made up of HPSDR units that move independently and create ad-hoc wireless communication links with those within their signal's range. They are capable of establishing, organizing, and maintaining themselves without a fixed centralized infrastructure as they collectively carry out vital network functions, such as management, packet forwarding and routing [2]. The PMR network is designed to provide short- to medium-range communications for mission-critical military applications, but could also be used for emergency search and rescue operations too. The purpose of this project is to take this robust communications network and develop a method of positioning that is also robust so that it functions continuously, even in harsh environments.

Certain constraints had to be set to ensure the positioning system was not only feasible, but could provide the necessary precision without undermining the security or high-performance of the communications network. No sig-

nificant hardware alterations could be made to the radios - this would likely exceed the allocated budget and is also beyond the scope of this project. No prior knowledge of the deployment region can be required nor a dependency on any infrastructure that's not part of the PMR network. The system would need to be able to position the radios to within at most a few hundred metres to be even potentially useful to Terrafix's clients, but within one hundred metres would be more preferable. Less than this would be even better, but this is not absolutely necessary as the positioning system developed should integrate with radio's existing Global Positioning System (GPS) receiver.

Nowadays, many positioning needs are met using signals received from at least one Global Navigation Satellite System (GNSS), as a receiver is all that's needed for a device to compute its position. The infrastructure, such as the satellites and control system, are already established. This includes the Terrafix radios, which are already fitted with GPS receivers, but this is not sufficient to meet the PMR network's positioning needs. Upon reception, a GNSS satellite transmission can be several orders of magnitude weaker than the ambient noise, so they are particularly vulnerable to interference and jamming. A one-watt signal can challenge a commercial GNSS receiver from as far as 100km away and even a military receiver cannot tolerate much interference during initial acquisition. The Terrafix radios may also need to operate indoors, which are difficult environments for all GNSS receivers as satellite transmissions are usually completely blocked by the building [3]. Filtering techniques are used to protect against typically intermittent narrowband signals that also operate on the same radio channels as GNSSs. Thus incidental interference is easily dealt with by using dual or multi-frequency receivers, but this would require significant hardware changes to the Terrafix radios. Furthermore, there's the possibility that others will inadvertently transmit on GNSS channels, if there is a flaw in their transmitter design, or deliberately do so if they wish to jam the channel. Even signals transmitted correctly on adjacent frequencies can "overspill" weak signals onto GNSS bands causing small but not negligible interference [4]. The Terrafix radios need an alternative means of positioning that can see them through periods of GNSS signal outage, hence the need for this research.

Standard alternatives to GNSS signals are not used as the PMR network is not designed to transmit or receive on other frequencies as much of the networks security revolves around their proprietary channels. Therefore, this would require huge hardware changes that would be costly and beyond the scope of this project. Furthermore, many signals such as a bluetooth and Wireless Local Access Network (WLAN) signals are short range, usually less than 100m, so they would not work without a very dense network [5]. Many of these alternatives are readily available in dense urban environments and these Signals Of OPportuinity (SOOP) can be utilized for positioning purposes [6]. However, the PMR network will not always be deployed in urban centres and, even if they were, the local infrastructure will likely be damaged or destroyed completely in a battlefield setting.

1.1 Position Fixing Techniques

Some of the common positioning fixing techniques in use today are angulation, pattern matching, proximity and ranging, but most them can't be easily applied to the PMR network. The Terrafix radios are fitted with an omnidirectional antenna, so to perform angulation would require the radio to be fitted with directional antenna which are normally much larger and expensive. Pattern matching requires a database of information collected for every point of a grid laid over the region of operation. Such information may include terrain height, landmarks, signals received (or not received) and at what signal power. First of all, the PMR will rarely have access to such information, but, in a military or disaster situation, this information regarding the region will not stay accurate for long. Furthermore, the signal strength of signals emitted by the Terrafix radios themselves, possibly the only signals available, will constantly change as they move about. Certain proximity methods could work reliably, but they're generally only accurate for short-range signals, especially in a mobile network. A Terrafix radio's transmission can still be detectable several kilometres away, so it would be impossible position to a desirable precision. To conclude, many of the positioning fixing techniques available would be too expensive, impractical, or incapable of positioning with sufficient precision. Some could potentially compute positions to an acceptable precision, but, in doing so, will undermine the purpose the network [7].

Positioning via ranging proved most appropriate for this network as it can provide an acceptable precision with little to no hardware alterations. Ranging can also integrate well with the radio's existing communications system and GPS receiver. One-way ranging was considered as it requires minimal transmissions, thereby reducing channel traffic and consumption of limited battery power. However, each radio of the PMR network possesses a local crystal oscillator clock that is not synchronized with those of other radios. Without a common time origin, the clocks of different radios will exhibit a relative clock offset. This clock offset between two radios is the difference in time as measured by their separate clocks at the same point in "absolute time". This refers to the Newtonian idea of absolute, true or mathematical time that progresses uniformly without regard to external influence, as all relativistic effects here are negligible. Clock synchronization parameters would need to be computed as part of the position solution of a one-way ranging system, but this would require a larger network with better signal geometry that can't be guaranteed. A two-way ranging protocol is implemented despite the added burden to the channel and batteries because the clock

synchronization parameters required can be computed with only two radios. Furthermore, most of the error due to the relative offset between different clocks is cancelled simply in the process of computing the range estimate (see Section 1.3).

1.2 Preliminary Work

Ranging experiments are conducted using two Terrafix radios of the PMR network that engage in two-way ranging. They are connected via a 1.5m length of coaxial cable due to licensing restrictions and fitted with two 40dB attenuators to prevent overloading the antennas. Range estimates are made once per second and printed to a terminal window on a Windows 7 operating system via a wired connection made with the radios. Modifications to the radio's signal processing are made by editing the source code of the Digital Signal Processor (DSP) installing it onto the radio. Then range estimates are made and collected by running the radios in real time rather than the more common approach of post-processing collected signal samples.

The PMR network is already capable of secure reliable communications, which was developed prior to this project by Terrafix, and forms the foundation upon which the ranging system is built. As the Terrafix radios have only a single core Central Processing Unit (CPU), they use a multi-threading computer architecture to achieve a high-performance regarding processing speed, efficient use of resources and responsiveness to user-initiated commands. A multi-threaded CPU possesses the ability to divide a process into multiple threads, that likely share common resources of the CPU, such that they're capable of being executed concurrently. For example, the process of receiving a message is divided into threads for interrupt handling, demodulation, message formation and display to the User Interface (UI). Although, this introduces certain difficulties when computing and extracting precise time information as each thread can be interrupted or delayed to execute a thread of another process. This can create huge variability in the execution time of a complete process and, as many threads and processes may have access to shared memory blocks, data can be easily lost or overwritten. To prevent this, a First-In-First-Out (FIFO) data buffer is implemented to hold timing data, maintain order, and ensure data is not overwritten.

Timestamps are made upon each frame transfer from the radio's Analog-to-Digital Converter (ADC) to the DSP, which marks the beginning of the receive signal processing chain. Each frame contains four data symbols that are timestamped by estimating where the midpoint of each is within the frame during symbol timing recovery, then adjusting the frame timestamp accordingly (see Section 2). These timestamps are placed in a FIFO data buffer to be retrieved in the near future should they form part of a successful message decode. The Time Of Flight (TOF) is computed by using the timestamps of self-received signals in place of computing transmission timestamps as illustrated in Figure 1 which shows the experiment schematic. This greatly simplifies the timestamping procedure and maintains similar processing time errors across all timestamps, which then cancel better when calculating the TOF (see Section 1.3). Timestamps made at Radios A and B are denoted by t_A and t_B respectively with an additional subscript of T to denote a transmission time and R to denote a reception time.

Using the existing protocols for symbol timing recovery, TOF measurements have been computed accurate to just a few microseconds, bar the occasional outlier. The jitter associated with the corresponding range is quantified by the Standard Deviation (SD) of ten thousand range estimates. Fixed biases are quantified by the Root-Mean-Square (RMS) of the ranges estimates, which are easily compensated for by calibration. The jitter, however, can only be compensated for in real-time by modifying the Terrafix radio's DSP source code to remove or avoid the jitter. These protocols, however, were never designed for TOF measurement - they only have accuracy and precision enough to demodulate symbols. Initial ranging estimates exhibited a SD in the several thousands of kilometres as they were calculated from timestamps made when messages received the inbox of the UI. This is because there are many threads executed between the transfer of message symbols from the ADC to the DSP to the inbox, each of which will contribute to the jitter. By taking timestamps earlier in the receive process chain and exploiting timing information extracted from the existing synchronization protocols, range estimates exhibiting a SD of 640m has been achieved. By implementing new synchronization techniques for precise timing recovery and clock synchronization, the SD can be lowered to a much more desirable level (see Section 2).

1.3 Clock Synchronization

The radio transceivers of the PMR network have half duplex connections with one another which enables them all to perform two-way ranging. Radios belonging to a half duplex communications system are capable of both transmitting and receiving signals, just not at the same time. If two radios A and B measured the times t_A and t_B respectively in accordance to their local clocks at an instant in absolute time, then the clock offset, δt, is defined as:

$$\delta t \triangleq t_B - t_A \tag{1}$$

This two-way time transfer scheme is illustrated in Figure 2 where the additional subscript of T or R to each time, t, denotes a transmission or reception time respectively.

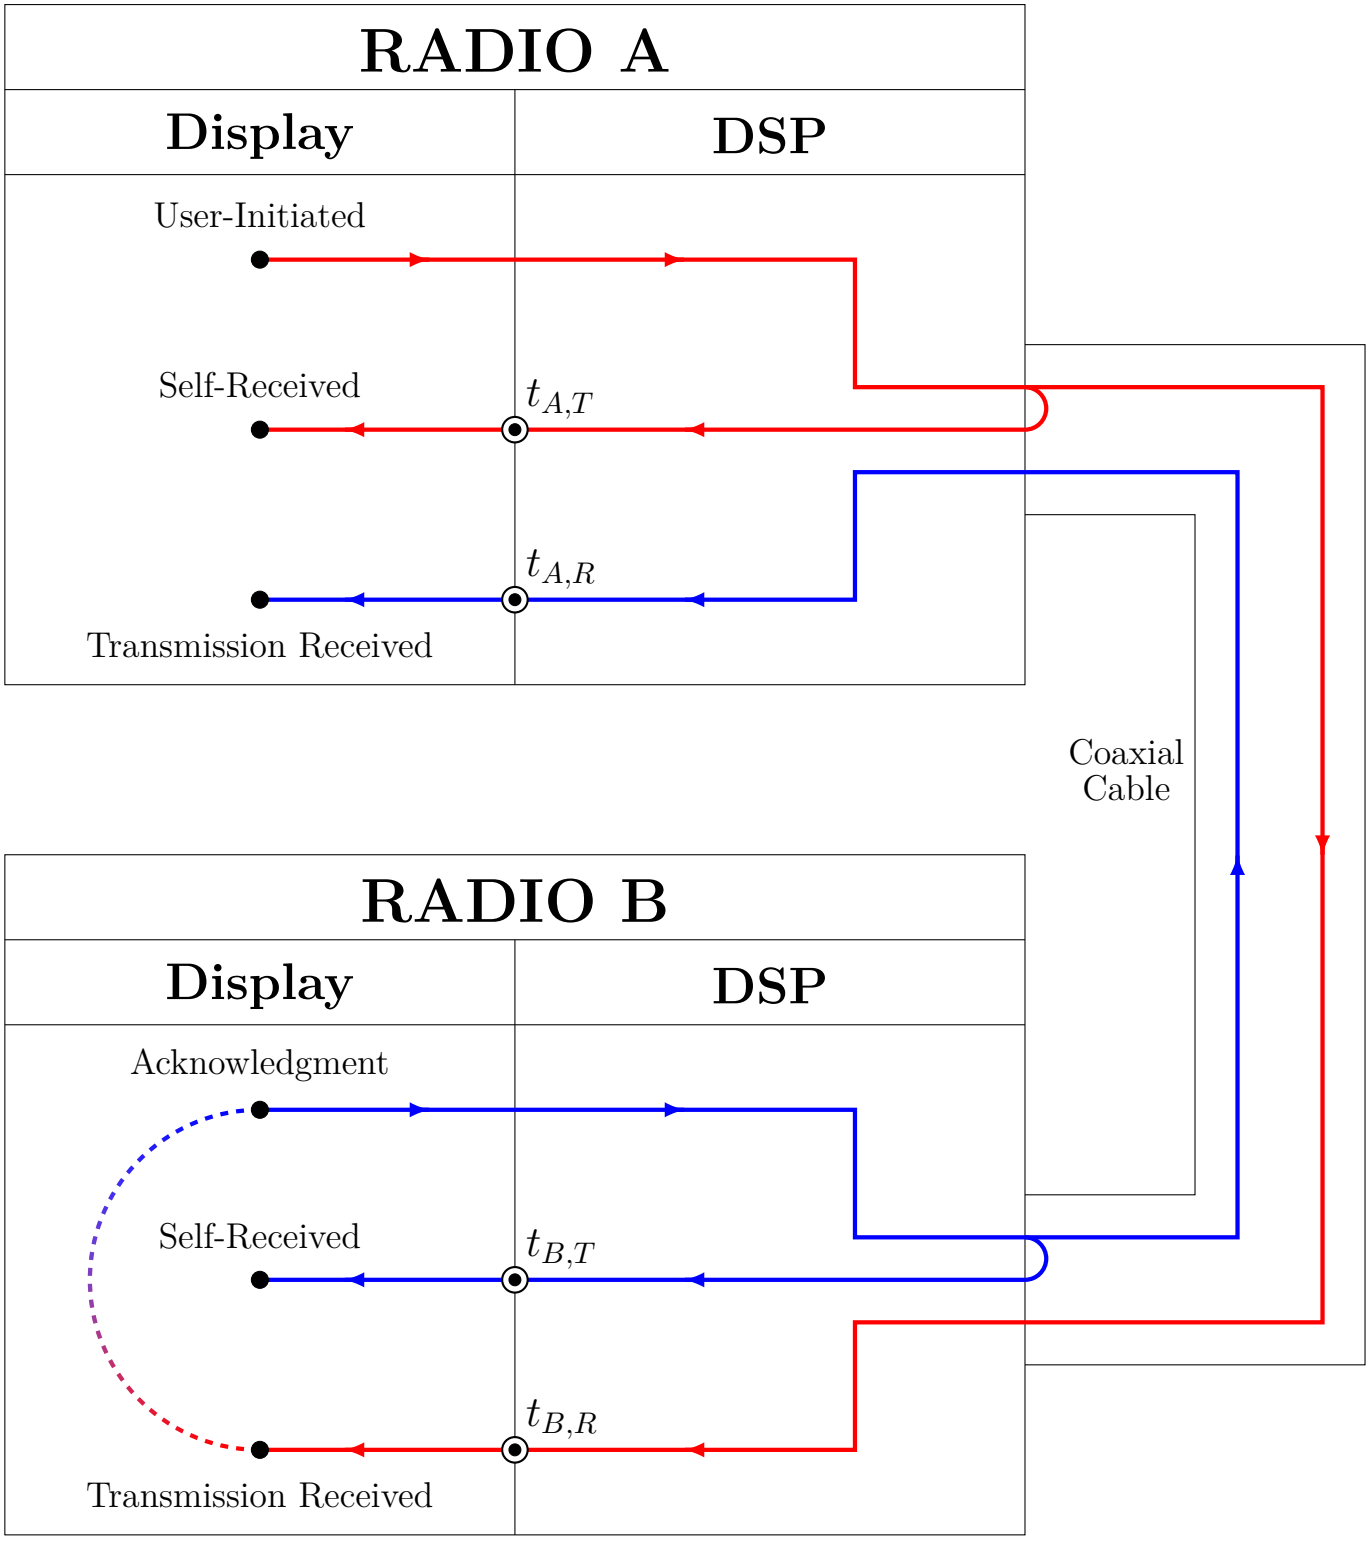

Figure 1: A schematic of the signal exchange made between two radios of the PMR network to estimate the range between them. A user-initiated signal (red) is sent from Radio A to Radio B where an acknowledgement signal (blue) is passively transmitted back to Radio A.

The propagation time, τ, cannot be calculated as simply the difference between transmission and reception timestamps at opposite sides of the communications link. This is because the values of t_A are not known with respect to the clock at B, nor the values of t_B with respect to the clock at A. By adding the clock offset, δt, to the timestamps

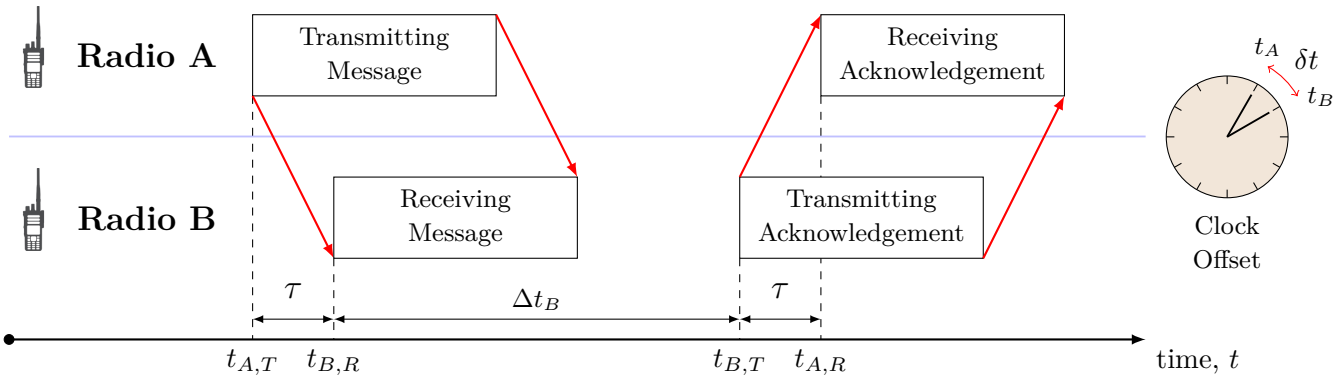

Figure 2: A graphical representation of the two-way time transfer scheme implemented into the Terrafix radios. Signals are transmitted between two radios with a non-zero clock offset of δt [8] p.96. Radio B takes a time of Δt_B to respond to the message sent by Radio A, which is called the response latency.

made at A, their timing is then aligned with the timing at B and the TOF, τ, may be calculated:

$$\tau = t_{B,R} - (t_{A,T} + \delta t) \tag{2}$$
$$\tau = (t_{A,R} + \delta t) - t_{B,T} \tag{3}$$

The TOF of each signal are treated as equal despite the fact the radios are mobile because they are handheld devices, so the change in range is negligible over the response latency. For a typical response latency of $\sim$ 100ms, the change in distance between the radios would only be a few tens of centimetres at most. Viewing Equations 2 and 3 as a pair of simultaneous equations in terms of δt and τ, they can be solved to find:

$$\delta t = \tfrac{1}{2}\left[(t_{B,R} - t_{A,T}) - (t_{A,R} - t_{B,T})\right] \tag{4}$$
$$\tau = \tfrac{1}{2}\left[(t_{B,R} - t_{A,T}) + (t_{A,R} - t_{B,T})\right] \tag{5}$$

Environmental conditions such as humidity and temperature affect the frequency of a crystal oscillator which can cause radios to become unsynchronized over time, even if they initially had no clock offset. It's possible to control a crystal oscillators condition to stabilise its frequency, but the stable frequency of each oscillator will still differ from those of others, thereby causing them to drift. Measurements of clock drift between the local clocks of the Terrafix radios, as computed in Section 4, have shown the drift to be constant over several hours, give or take about 75ns. The clock offset at any point in time is the initial offset plus the integral of the clock drift since initialization. Conversely, the clock drift at any point in time is the derivative of the clock offset. The vast majority of the TOF error due to relative clock drift at this point is incurred during the response latency period, as this is far greater than the TOF itself.

Thus far in the formulae, the clock offset has been treated as a constant, but typically this is not the case. Therefore, each of the four timestamps made, $t_{A,T}$, $t_{B,R}$, $t_{B,T}$ and $t_{A,R}$, will have a unique respective clock offset at that time: $\delta t_{A,T}$, $\delta t_{B,R}$, $\delta t_{B,T}$ and $\delta t_{A,R}$. However, the range error due to the relative clock drift over the TOF of the signal between radios A and B would be only $\sim$ 1mm per kilometre of separation. Therefore, the problem can be simplified by only considering the clock offsets associated with the timestamps made at Radio A as $\delta t_{B,R} \approx \delta t_{A,T}$ and $\delta t_{B,T} \approx \delta t_{A,R}$. Furthermore, as $\tau \ll \Delta t_B$, one can treat the Round-Trip-Time (RTT), Δt, and the response latency as equal because:

$$\Delta t \triangleq t_{A,R} - t_{A,T} \approx t_{B,T} - t_{B,R} = \Delta t_B \tag{6}$$

Recall from Equation 1 that $t_B = t_A + \delta t$, therefore, by making the substitution $t_A \mapsto t_A + \delta t$, the TOF corrected for relative clock drift, $\tilde{\tau}$, can be found:

$$\begin{aligned}\tilde{\tau} =& \frac{1}{2}\left[(t_{B,R} - t_{A,T}) + (t_{A,R} - t_{B,T})\right] \\ &+ \frac{1}{2}(\delta t_{A,R} - \delta t_{A,T}) \\ =& \tau + \delta^2 t\end{aligned} \tag{7}$$

where

$$\delta^2 t = \frac{1}{2}(\delta t_{A,R} - \delta t_{A,T}) \tag{8}$$

is the clock drift correction. Under closer inspection, one can see that the clock drift correction deduced in Equation 8 is simply half the clock drift incurred over the RTT, Δt. Therefore, the clock drift correction, $\delta^2 t$, may be reformulated as:

$$\delta^2 t = \frac{1}{2}(\delta t_{A,R} - \delta t_{A,T}) \equiv \frac{1}{2}\frac{d(\delta t)}{dt}\Delta t \tag{9}$$

If the derivative in Equation 9 can be estimated, that is the clock drift per unit time, then the range error due to clock drift over the response latency period can also be found (see Section 4).

2 Symbol Timing Recovery

When a signal is received, information regarding the timing that was used to create it is needed to successfully demodulate the signal. This procedure is known as timing recovery, which can be implemented in a variety of ways, but the Terrafix radios use oversampling. This is when the signal is sampled at a rate higher than the symbol rate, then symbol timing recovery is achieved by selecting the sample that best represents each symbol. In doing so, the sampled signal has then been downsampled back to the symbol rate.

The existing symbol timing recovery protocols of the Terrafix radios are not conducive to precise ranging for several reasons. Firstly, timestamps can only be made to the nearest sample and the oversampling rate is only eight, so it's not possible to obtain the timing of each symbol to an acceptable precision, especially considering the duration of each symbol. Timing is not recovered for each symbol individually, instead a downsampling index, $j \in [0, ..., 7]$, is selected based on the absolute amplitude of the samples. Then every eighth sample from j is selected to downsample the signal, but this can also be utilized to timestamp all symbols in a block. The size of this block is carefully chosen so that it's long enough to produce an accurate downsampling index, but not so long that significant error due to clock drift would be incurred in the time to receive it. This method is sufficient to demodulate the sampled signal for communications, but not nearly sufficient for useful ranging.

2.1 The Zero-crossing Method

The symbol zero-crossings refer to points in the symbol spectrum where the amplitude of demodulated samples go from positive to negative, or vice versa. Generally, a zero-crossing will occur between two samples, so this new method of timing recovery will allow timestamps to be made at what could be seen as a "sub-sample" level. If the signal's Signal-to-Noise Ratio (SNR) is low, then the demodulated samples may appear on the wrong side of the zero line, hence create an erroneous zero-crossing. A sample of the symbol spectrum collected under laboratory conditions is plotted in Figure 3 and, though it may sometimes appear close to the zero line, it's still very high considering the scale of the amplitude axis. Attenuation experiments have demonstrated that the zero-crossing method is robust down to a Received Signal Strength Indicator (RSSI) value approaching the radio's signal detection limit. However, in practice, signals can experience more than just attenuation, such as interference, so the new timing protocols may need to consider several samples together before accepting a zero-crossing has occurred. From the difference between one zero-crossing and the next, one can estimate the timing of the symbols between them by making the reasonable assumptions that the symbols are spaced evenly and are between 4 to 12 sample periods in duration.

When a change in sign has been detected between the k^{th} demodulated sample, d_k, and the previous, d_{k-1}, at times t_k and t_{k-1} respectively, the time of the zero-crossing, t_z, can be estimated as:

$$t_z = t_k - d_k\left(\frac{t_k - t_{k-1}}{d_k - d_{k-1}}\right) = t_k + \nabla t_z \tag{10}$$

where

$$\nabla t_z = -d_k\left(\frac{t_k - t_{k-1}}{d_k - d_{k-1}}\right) \tag{11}$$

is the adjustment required to obtain the timestamp for the zero-crossing, t_z, from the timestamp of the current sample, t_k. Deltas, Δ, are used to denote intervals whereas Nablas, ∇, are used to denote adjustments. This is simply the zero-crossing of the straight line between points (t_{k-1}, d_{k-1}) and (t_k, d_k) in the time-amplitude plane. The difference between this crossing and the previous can then used to determine the number of samples between the two crossings. This difference will be close to a multiple of eight as eight signal samples are measured by the ADC for each symbol. Therefore, dividing the difference by eight and rounding to the nearest integer will estimate the number of symbols between the two zero-crossings. By making the reasonable assumption that the symbols are evenly spaced between the two zero-crossings, one can deduce timing estimations for the midpoint of each symbol. These estimations are illustrated in Figure 3 by red dashed lines for a portion of a sampled and demodulated signal.

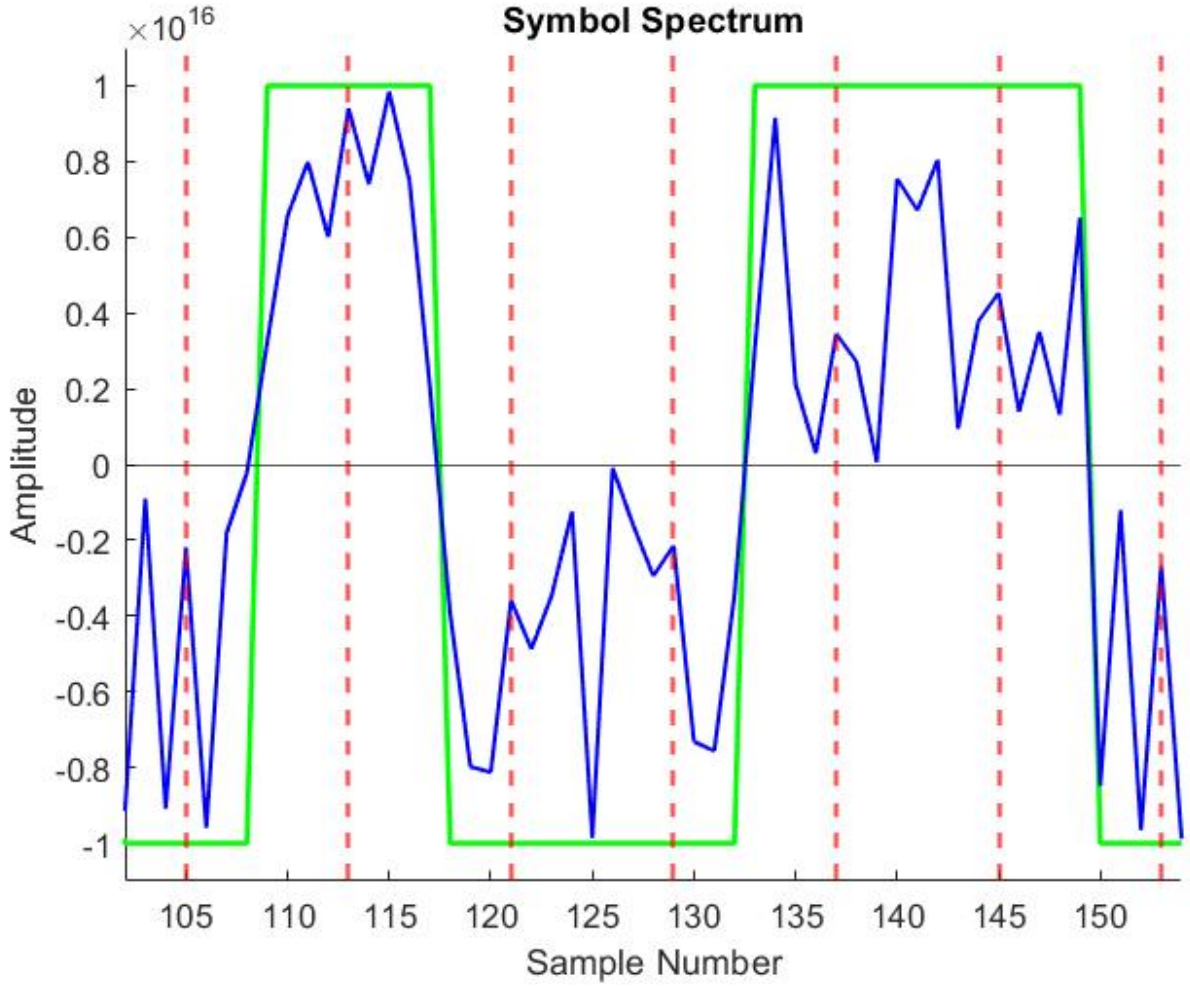

Figure 3: The symbol spectrum of a portion of a real demodulated and sampled signal (blue) which produced the demodulated chips (green) by selecting sample numbers closest to the red dashed lines.

3 Non-coherent Integration

In the prior work, only the leading symbol timestamp of a packet is used despite the fact that timing information needs to be collected for every symbol; this is very inefficient in both memory and potential timing precision. By applying the same procedure to every symbol of a packet and averaging the timestamps produced, a far greater precision can be achieved as much of the jitter cancels. This is done by integrating symbol timestamps from the leading symbol timestamp until the final symbol timestamps of a data packet and then dividing by the number of symbols in a packet. The PMR network uses packet switching to transmit data, which is a method of formatting data into packets consisting of a header block used to direct the packet to its destination and then the actual data payload.

Coherent integration requires synchronous signal reception, which is usually accomplished using carrier recovery methods. When reception is synchronous, each symbol of the signal is sampled at the same points determined by the oversampling rate. The radios of the PMR network differentially encode data to eliminate the adverse effects of received signal phase offset relative to their local oscillator clocks. This allows for asynchronous reception, but this means that the timing of a received signal's sampling is independent of the signal's Time-Of-Arrival (TOA) [9]. Synchronous and asynchronous sampling are illustrated in Figure 4 for three sampled symbols superimposed over one another. The Gaussian Minimum-Shift Keying (GMSK) demodulation process performed by the radios of the Terrafix PMR network is asynchronous, so the sample points of each symbol will drift from the previous. This modulation commonly used for non-linear, severely band-limited channels as the sinusoidal pulse shaping keeps side-lobes low while maintaining an efficient detection performance [10]. Asynchronicity introduces a jitter to the timestamps made, but this can be lessened by averaging over many symbol timestamps. Baseband data is differentially encoded and modulated onto both in-phase and quadrature components of a carrier to produce a signal that is resistant to amplitude loss when the radio clocks are out-of-phase [11]. Differential encoding is a technique whereby data symbols are decoded not

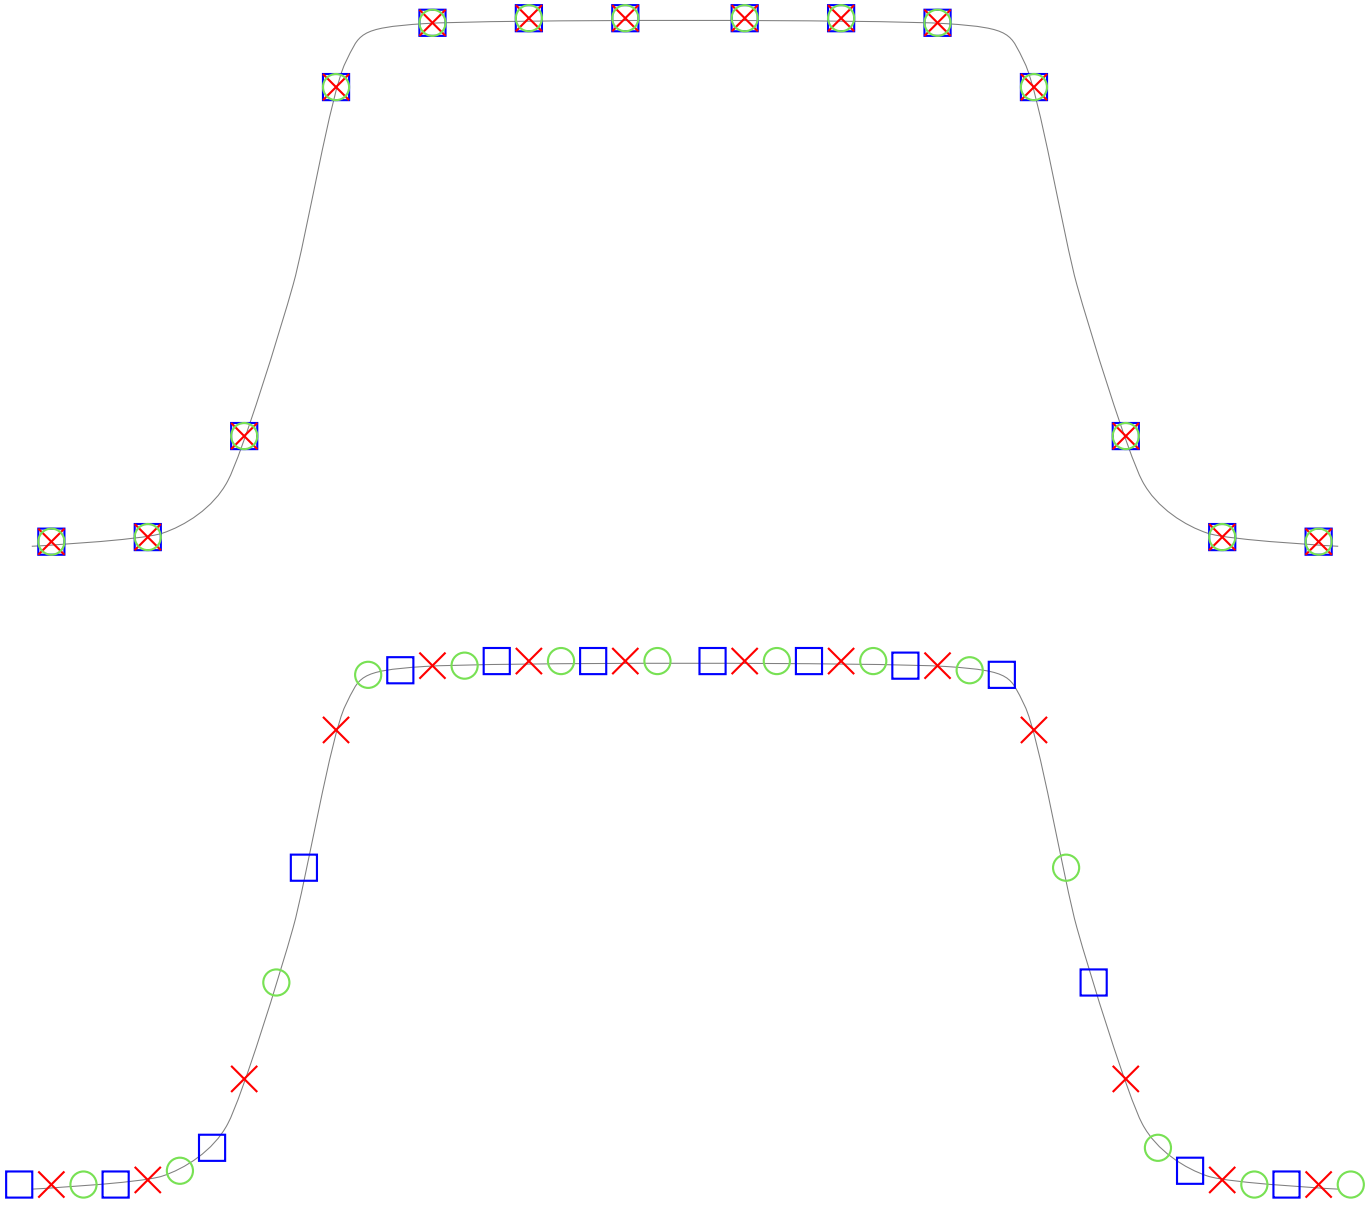

Figure 4: The difference between synchronous (top) and asynchronous (bottom) sampling when demodulating a band-limited signal. Three sampled symbols are superimposed: the first represented as boxes, the second as crosses and the third as circles. When sampling is synchronous, the symbol samples are taken at the same points, but when sampling asynchronous, symbol samples will drift.

from the signals state, but from a change in the signals state. This avoids the need for a carrier recovery protocol to synchronize the radios for demodulating a signal.

The efficacy of this averaging is limited by the relative phase offset between the received signal and local clock, which grows linearly over time for the Terrafix radios due to a constant clock drift. A longer integration allows more time for the two radio clocks to drift while timestamps are being computed, so, beyond a certain point, the error due to clock drift will exceed the benefit gained from lengthening the integration. Although, the integration length is limited by the processing time itself which would eventually surpass a practical duration.

4 Clock Drift Correction

The effect that relative clock drift has on ranging is twofold: there's drift over the response latency at Radio B for each range estimate and there's drift over the integration period for each timestamp. By making the clock drift corrections from one symbol to the next during integration would be most effective as this minimizes the time the clocks can drift unchecked. However, this would be an unnecessary use of resources as the range accuracy error due

to clock drift between symbols is about 4cm and, therefore, negligible. As timestamps need to be shared between two radios in order to correct the error due to the clock drift between them in the TOF estimates, then it's not practical to make corrections more frequently than once per packet. The range accuracy error due to clock drift over one packet is about 23m, which will incur a significant jitter, so it's appropriate to make a clock drift corrections every packet.

The response latency at Radio B is typically $\sim$ 100ms and the nominal uncertainty of the oscillator clocks used by the Terrafix radios is 1ppm. Therefore, the drift over the response latency can be as much as $\sim \pm$100ns. However, notice the formula for the TOF expressed in Equation 5 can be rewritten as:

$$\begin{aligned}\tau &= \frac{1}{2}\left[(t_{A,R} - t_{A,T}) - (t_{B,T} - t_{B,R})\right] \\ &= \frac{1}{2}\Delta t - \frac{1}{2}\Delta t_B\end{aligned} \tag{12}$$

So if the error in Δt_B is no more than $\sim \pm$100ns, the actual error in τ due to clock drift over the latency period should be within $\sim \pm$50ns, which corresponds to a range error of $\sim \pm$15m. The range estimates appear to follow a Pearson type VII distribution (see Section 5), and, like a normal distribution, the vast majority of data points lie within at least three standard deviations from the mean. Hence, by making a clock drift correction, the improvement in jitter, as measured by the standard deviation of range estimates, should be approximately 5m.

A formula for clock drift per unit time is deduced here differentially by calculating the difference in clock offset between consecutive range estimates (see [8] pp.173-176). Recall from Equation 9 that an estimate of the clock drift per unit time is needed in order to cancel the clock offset effectively. Referring to Figure 2, one can see that the timestamps $t_{A,T}$ and $t_{B,R}$ are separated by τ in absolute time, so the difference between $t_{A,T} + \tau$ and $t_{B,R}$ will be the clock offset $\delta t_{A,T}$. By applying the same logic with timestamps $t_{B,T}$ and $t_{A,R}$, the clock offsets for the n^{th} range estimate when timestamps are made at A can be defined by:

$$\delta t_{n,A,T} = t_{n,B,R} - (t_{n,A,T} + \tau_n) \tag{13}$$

$$\delta t_{n,A,R} = t_{n,B,T} - (t_{n,A,R} - \tau_n) \tag{14}$$

The Terrafix radios exhibit a constant clock drift, so by averaging Equations 13 and 14, an average relative clock offset, $\overline{\delta t}_n$, can be calculated over the n^{th} range estimate:

$$\begin{aligned}\overline{\delta t}_n &= \frac{1}{2}(\delta t_{n,A,T} + \delta t_{n,A,R}) \\ &= \frac{1}{2}\left[(t_{n,B,R} - t_{n,A,T}) - (t_{n,A,R} - t_{n,B,T})\right]\end{aligned} \tag{15}$$

Notice the resemblance of Equation 15 with Equation 4, as one might expect. This average clock offset always represents the clock offset at the midpoint of the RTT because the clock drift is constant. As Equation 15 depends only on the measured timestamps, it can be used to estimate the derivative of δt_n with respect to n as:

$$\frac{d(\delta t_n)}{dn} \approx \frac{\overline{\delta t}_n - \overline{\delta t}_{n-1}}{n - (n-1)} = \overline{\delta t}_n - \overline{\delta t}_{n-1} \tag{16}$$

Similarly, the derivative of time, t, with respect to the range estimate number n can be estimated as:

$$\frac{dt}{dn} \approx t_{n,A,T} - t_{n-1,A,T} \sim 1\text{s} \tag{17}$$

Range estimates are made at a rate of one per second, but this is governed by the radio display which is not nearly as regular as the DSP. This is not the maximum rate, any rate below about 2 per second would work, but 1 per second is chosen to minimize the data collection time while avoiding packet collisions. The true rate can differ from the defined rate by as much as several tens of milliseconds, so the true rate is calculated for each range estimate to account for this fluctuation. Finally, the clock drift with respect to time can be deduced from Equations 16 and 17 as the following:

$$\frac{d(\delta t_n)}{dt} \equiv \frac{d(\delta t_n)}{dn}\frac{dn}{dt} \approx \frac{\overline{\delta t}_n - \overline{\delta t}_{n-1}}{t_{n,A,T} - t_{n-1,A,T}} \tag{18}$$

5 Ranging Results

Continuing from the preliminary work, the jitter associated with range estimates is measured by calculating the Standard Deviation (SD) over ten thousand range estimates. The SD dropped from approximately 640m to 496m once

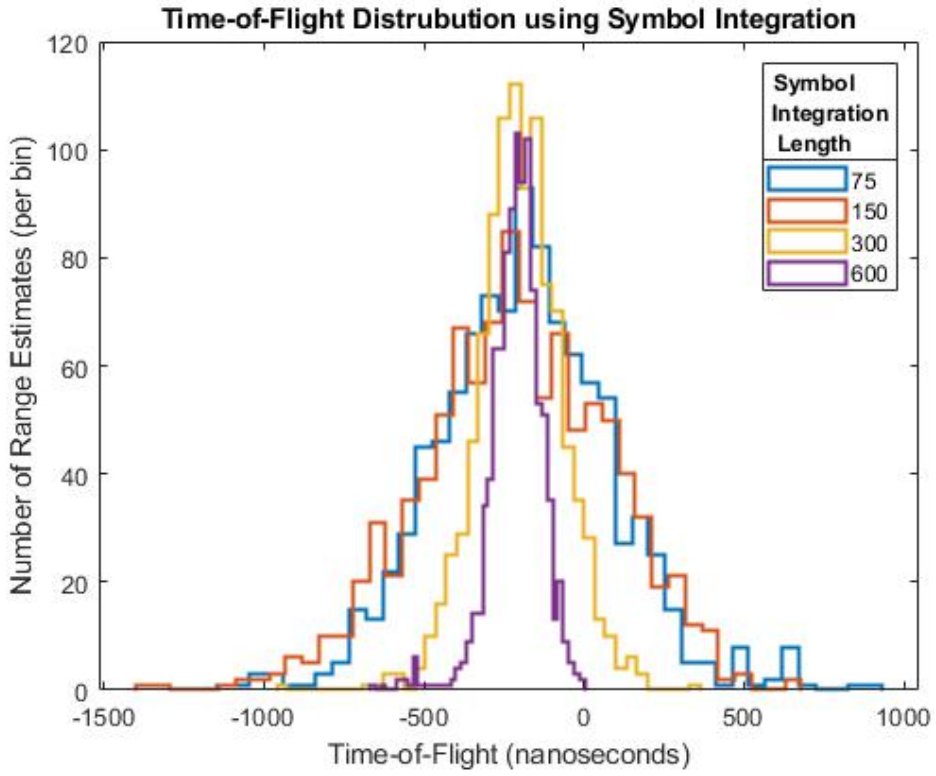

Figure 5: The distribution of range estimates made for different integration sizes: 75, 150, 300 and 600 symbols. These are illustrated by stair histograms of 40 bins.

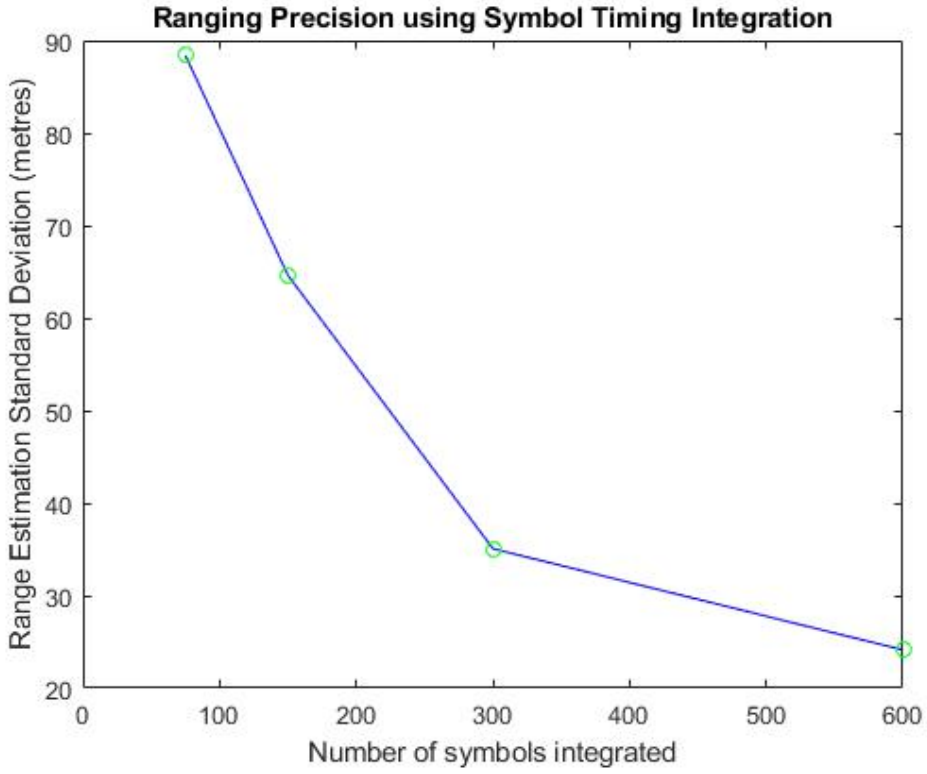

Figure 6: The standard deviation of range estimates made for different integration sizes: 75, 150, 300 and 600 symbols.

the zero-crossing method of symbol timing recovery was implemented. This method produces a superior performance to the original protocols, but not by a huge margin.

The benefit of the zero-crossing method becomes apparent when the symbol timing is averaged over several symbols to produce a single timestamp for the message as whole of a higher precision. Averaging over a message of 600 symbols using the zero-crossing method causes the SD of the resulting range estimates to plummet to only 24.13m - a drop of over 95% from the 496m SD achieved with a single symbol. For comparison, averaging over 600 symbols using the original timing recovery protocols produces a SD of 478m - barely below the zero-crossing method with no averaging whatsoever. This is because the original timing recovery protocols create a skew in the error distribution of the timestamps. Therefore, the precision gained from symbol timing integration is hardly worth the additional processing time and energy consumption. This skew is caused by how the signal is downsampled, as the timing of each symbol is not computed individually but collectively across an entire block. Hence, the error associated with the timestamp of each symbol is not independent of others, so they don't cancel as well when the timestamps are averaged. Computing the symbol timing using the zero-crossing method produces a far more symmetric error distribution, thereby causing a major drop in SD once averaged. The range estimation distributions and decline in SD are displayed in Figures 5 and 6 respectively for integrations over 75, 150, 300 and 600 symbols using the zero-crossing method. Integrating also greatly dilutes outliers which can be seen in the reduction of the minimum and maximum range estimates from $\sim \pm 3,700$m to $\sim \pm 135$m. There is a steep drop in SD just between 1 and 75 symbols integrated, however, this is not shown on the graph in order to maintain focus on results below 100m, which are considered useful.

The clock drift correction implemented in Section 4 reduced the SD from 24.13m to 19.53m - in close accordance with the hypothesised improvement of 5m. The clock drift correction is made here as this is the first instance that the SD has been low enough for any improvement to be detectable. The previous result of 24.13m was stable to about the nearest 1m, whereas the SD seen using the zero-crossing method and no averaging could fluctuate by over 10m when the experiment was repeated. The 19.53m SD has also proven to be stable to about the nearest 1m when the experiment is repeated several times.

A ranging message can be made up of more than 600 symbols, so the integration length is extended to see when does the decline in SD plateau. The range estimation distributions and decline in SD are displayed in Figures 7 and 8 respectively including messages of up to 76,800 symbols in length. Following the pattern in Figures 8, it can be inferred from the symbol rate that to double the message length again to 153,600 symbols would increase the transmission time to over 9 seconds. The SD of the range estimates, however, will likely only reduce by about 50cm. As the Terrafix radios are mobile units, they have a limited power supply that would drain very quickly should the integration length be extended any further.

Of the ranging experiments conducted, a best SD of 2.61m was achieved - a further reduction of approximately 87% from the previous best of 19.53m. All ten thousand range estimates fall within a statistical range of approximately only 16m. However, as can be seen in Figure 8, range estimates with a SD as small as 6m can be achieved with an integration length of only 10,000 symbols. Beyond this many symbols, we begin to see diminishing returns as the signal transmission time becomes very long, hence increasing power consumption and the odds of packet collisions, but the SD drops by only a few centimetres.

6 Conclusions

In this paper, an innovative means of extracting useful time information from a narrowband transceiver has been presented. The jitter of range estimates, quantified by the Standard Deviation (SD) of ten thousand estimates, has been measured as 2.61m, despite the PMR system never being designed for positioning. Even outlier range estimates are kept within approximately 8m of the mean.

In this narrowband system, the duration of a single symbol is $\sim 60\mu$s, so, based on the TOF estimations made, timing recovery to approximately one thousandth of a symbol has been achieved. Notwithstanding the multi-threaded system architecture and very narrow bandwidth, range estimates of significance have been produced.

7 Future Work

The next step for this research is to measure how the ranging precision is affected by multipath interference. As the Terrafix radios generally only transmit over a radius of a few kilometres, the maximum signal path delay is less than that of GPS signals. However, due to the narrow bandwidth of the PMR network, Non-Line-Of-Sight (NLOS) signals with a large path delay could still interfere with more direct signals. Therefore, range errors due to multipath interference may be greater for radios of the PMR network than their GPS receivers. A multipath mitigation protocol such as the MUltiple SIgnal Classification (MUSIC) algorithm could be implemented, though this will likely need to be in post-processing as it's computationally intensive [12]. The Multipath Estimating Delay Locked Loop (MEDLL) technique used by GNSS could also be applied to the PMR network, which is not as computationally intense as MUSIC.

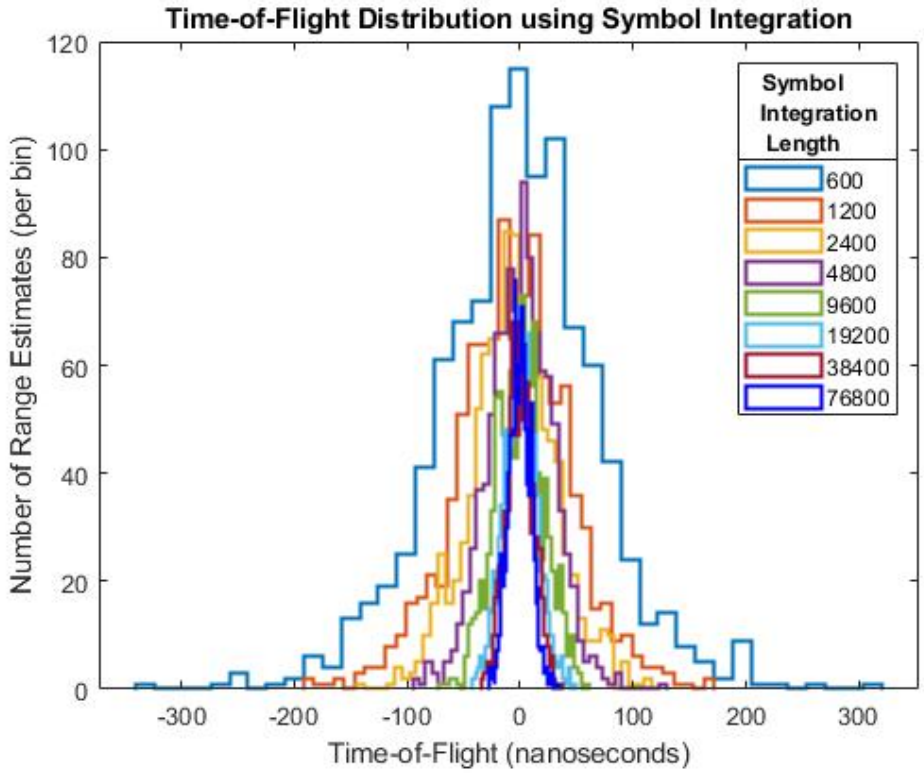

Figure 7: The distribution of range estimates across different integration sizes from 600 to 76800 symbols in length. These are illustrated by stair histograms of 40 bins.

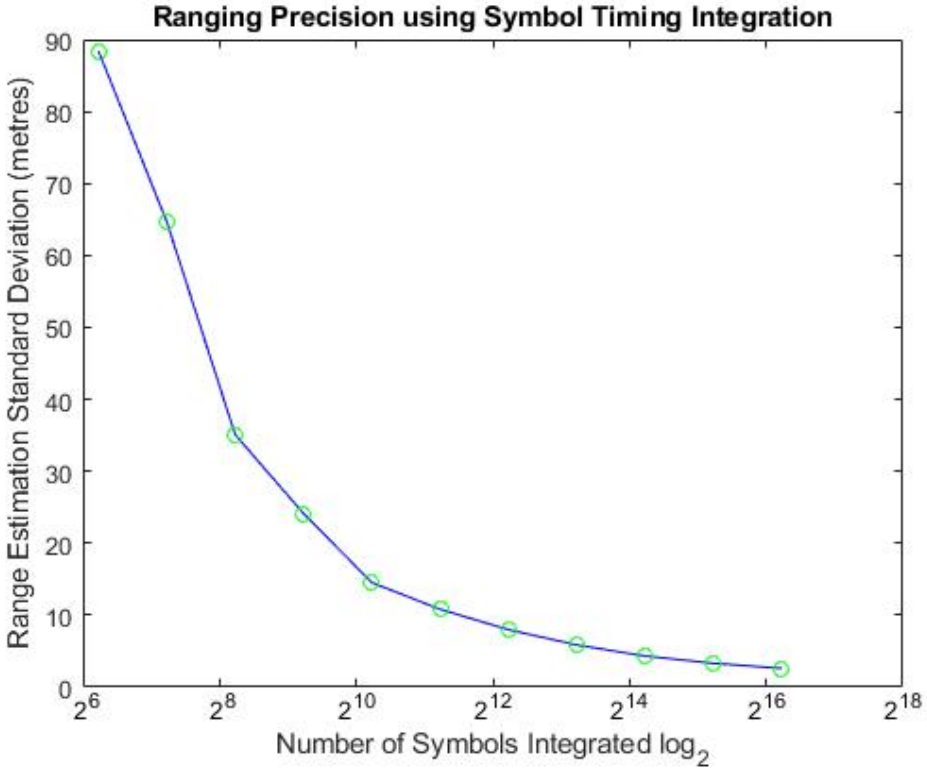

Figure 8: The standard deviation of range estimates across all integration sizes, ranging from 75 to 76,800 symbols. The number of symbols integrated are displayed logarithmically with a base of 2 to more clearly show the rapid decline in SD between experiments when the number of symbols is doubled.

If the chosen method's efficacy is sufficient for the PMR network, then hardware alterations to increase processing speeds may be made to accommodate real-time multipath mitigation.

Existing positioning algorithms [13] [14] [15] may be applied to ad-hoc systems such as the Terrafix radios to enable network-wide positioning. Such algorithms will make use of communication links between radios in the PMR network to improve upon the positions made from individual ranging data alone. This method is most effective when the positions of several radios are computed together, rather than independently, as all information is available.

8 Acknowledgements

I (Vincent Beech) would like to acknowledge the contributions of Dr Bob Mason, Technical Director for Terrafix before he retired part-way through this project, who provided us with valuable knowledge and wisdom. In particular, for his idea to timestamp self-received signals in place of transmissions to simplify the timestamping procedures and reduce errors in the range estimates (see Section 1.2).

I would also like to acknowledge the contributions of Dr David Selviah, who has been secondary supervisor to this project and provided guidance regarding writing and equipment.

Finally, I would like to acknowledge the contributions of all those at Terrafix who have provided advice and assistance over years regarding this project. I particularly appreciate the time Joe Woodward has taken to help me understand various protocols of the radios and for writing clear orderly source code that was simple to build upon.

References

[1] Terrafix: Total solution providers; 2017. Accessed: 11/01/18. Available from: `http://www.terrafix.co.uk/about.html`.

[2] Raja ML, Baboo CDSS. An Overview of MANET: Applications, Attacks and Challenges. International Journal of Computer Science and Mobile Computing (IJCSMC). 2014;3:408–417.

[3] Sklar JR. Interference mitigation approaches for the global positioning system. Lincoln laboratory journal. 2003;14(2):167–179.

[4] White paper on GNSS Interference; 2012. Available from: `https://www.septentrio.com`.

[5] Palmer D. Position estimation using the Digital Audio Broadcast (DAB) signal. University of Nottingham; 2011.

[6] Webb TA. Differential Positioning using Signals of Opportunity. UCL (University College London); 2013.

[7] Groves PD. Principles of GNSS, inertial, and multisensor integrated navigation systems. 2nd ed. Boston/London: Artech house; 2013.

[8] Bensky A. Wireless positioning technologies and applications. Boston/London: Artech House; 2008.

[9] Baringbing J; Signal Processing Graz University of Technology; Speech Communications Laboratory. Noncoherent Detection and Differential Detection. 2006;.

[10] Gronemeyer S, McBride A. MSK and offset QPSK modulation. IEEE Transactions on Communications. 1976;24(8):809–820.

[11] Practical Guide to Radio-Frequency Analysis and Design; 2020. Accessed: 04/06/20. Available from: `https://www.allaboutcircuits.com/textbook/radio-frequency-analysis-design/`.

[12] Lin W, Liu J, Zhou Y, Huang J. Estimation of TOA based MUSIC algorithm and cross correlation algorithm of appropriate interval. In: AIP Conference Proceedings. vol. 1820. AIP Publishing; 2017. p. 080011.

[13] Huang B, Yao Z, Cui X, Lu M. Distributed GNSS Collaborative Positioning Algorithms and Performance Analysis. In: China Satellite Navigation Conference (CSNC) 2015 Proceedings: Volume III. Springer; 2015. p. 427–437.

[14] Čapkun S, Hamdi M, Hubaux JP. GPS-free positioning in mobile ad hoc networks. Cluster Computing. 2002;5(2):157–167.

[15] Capkun S, Hubaux JP. Secure positioning in wireless networks. IEEE Journal on Selected Areas in Communications. 2006;24(2):221–232.