

Chapter 3

Source signals

This chapter describes the time-domain cross-correlation used by the relative localisation system as well as the motivation behind the choice of maximum length sequences (MLS) as the source signal. Several classes of signals are considered and their performance is experimentally evaluated against the required characteristics of the source signal to be used in an underwater acoustic localisation system. The main criteria are the cross-correlation peak detection performance and interference robustness when used in noisy, reverberant environments. In addition, the effects of uncorrelated additive white noise, mixing of multiple MLS signals as well as detrimental effects of non-linear transducers on the transmitted MLS signal's spectral properties are also discussed in the subsequent sections. An empirical method for overcoming the frequency filtering introduced by the transducer to improve the cross-correlation performance is presented as well.

3.1 Full-range cross-correlation of time-domain signals

In the relative localisation system presented in this thesis, cross-correlation is used as the main tool for extracting delays between signal channels. The usual approach suggested in the literature involves the generalised cross-correlation (GCC) or its variations first suggested by Knapp and

Carter (1976). Applications of this method including the work done by Carter (1981) and Piersol (1981) is based on performing the cross-correlation in the frequency domain, i.e, using the Fourier transforms of the original time-domain signals to produce a cross-power spectrum. Arguments supporting this approach include lower processing complexity and being able to achieve a narrow peak in the resulting cross-correlogram by applying appropriate pre-filters for the signals. The latter requires *a priori* knowledge about the signal characteristics to properly apply the pre-filters, which otherwise would make the delay estimation system susceptible to external noise.

In this research, cross-correlation of signals is performed in the time-domain. With relatively short duration signals (as explained in the next section), the issue of computational complexity does not arise and the narrowness of the cross-correlogram peak is maintained by the characteristics of the source signal being used and compensating for the transducer effects.

For two finite, discrete sequences $s_1(n)$ and $s_2(n)$ of length N (with $s_1(n) = s_2(n) = 0$ for $n < 0$ and $n \geq N$) the cross-correlation can be expressed as:

$$R_{s_1s_2}(\tau) = \sum_{n=i}^{N-|k|-1} s_1(n)s_2(n-\tau) \quad (3.1)$$

where $i = \tau$, $k = 0$ for $\tau \geq 0$, and $i = 0$, $k = \tau$ for $\tau < 0$. The auto-correlation of the sequence $s_1(n)$ can be derived by setting $s_2(n) = s_1(n)$ in the above formula (Burdic, 1984; Proakis and Manolakis, 1996).

By using the relationship:

$$R_{s_1s_2}(\tau) = R_{s_2s_1}(-\tau) \quad (3.2)$$

negative lags can be measured as well. Considering signed lags instead of an absolute value helps to avoid the *left-right ambiguity* otherwise encountered in TDOA based source localisation. Combining cross-correlations for both positive and negative lags, a full-range cross-correlation function can be constructed, which represents all possible delays between the two sequences. For the localisation system, the sequences consist of discrete samples of the acoustic signals output by the analogue to digital converter. These **sample-domain** signal lags can be easily converted to the **time-domain** by considering the sampling frequency of the converter (f_s) and the **spatial-domain**¹ by considering both the sampling frequency and propagation speed of the signal (v). The quantity τ_0 which maximises $R_{s_1s_2}(\tau)$ corresponds to the delay between the two signals. The resolution of τ_0 depends on the characteristics of the signal waveform being used, as explained in the following sections.

¹. The distance travelled by a wavefront of the signal during the delay period.

3.2 Choice of signal waveform

In acoustic source localisation, it is preferable to avoid pure tone periodic signals since the TDOA can only be unambiguously estimated for frequencies whose wavelengths are at least twice the base distance of the receivers. Therefore, for a base distance of 0.3 m used by the proposed system explained in chapter 4.4.3, the frequency needs to be 2500 Hz or lower. Cross-correlation of such signals results in a wide peak in the cross-correlogram causing the delay estimates to have a lower resolution. In addition, due to frequency dependent fading effects caused by multipath propagation (Sozer et al., 2000; Kilfoyle and Baggeroer, 2000) encountered in underwater acoustics, the signals of choice for localisation applications have been mostly non-periodic and broadband.

For the localisation application presented in this thesis, the primary requirement of the source signal is to provide a unique narrow peak in the cross-correlogram. The secondary requirements are to provide a high signal-to-noise ratio (SNR) when used in noisy environments¹ (*i.e.* robustness against interference) and to maintain a broad frequency spectrum to counter the fading effects of the underwater channel. While a ‘flat’ frequency spectrum is highly desirable, it is extremely difficult to find transducers that faithfully reproduce such a response.

Based on prior research in the fields of acoustical localisation in air (Girod et al., 2006), sonar signal processing (Nielsen, 1991; Waite, 2002) and room acoustics (Bradley, 1996), three broad classes of signal waveforms were identified as chirps, pseudo-noise and linear feedback shift register sequences. Samples from each of those classes were evaluated against the performance criteria mentioned above and compared.

3.2.1 Chirps, pseudo-noise and shift-register sequences

Sine sweeps or ‘chirps’ are signals where the base frequency is changed with time. The rate of change in frequency is referred to as the ‘chirp rate’ where linear (linear frequency modulated - LFM) and logarithmic chirp rates are commonly used to synthesise source signals. These signals are popular in sonar and radar applications where the reflected signal is utilised for localisation of a target rather than localising the sound source itself (Ma and Goh, 2006). Studies have revealed that bats use ‘down sweeps’, among other waveform structures, for precise echolocation of prey (Neuweiler, 2003). Sine sweeps with logarithmically increasing chirp rates are considered to be better suited for moving target detection as they behave more gracefully with the introduced doppler shifts in frequency. In addition, sine sweeps with arbitrary bandwidths can be synthesised to match the operational bandwidth of the transducers used.

¹. Noise in this context includes ambient background noise, stray impulsive signals emitted from nearby sources as well as reflections of the main signal when used in reverberant, cluttered environments.

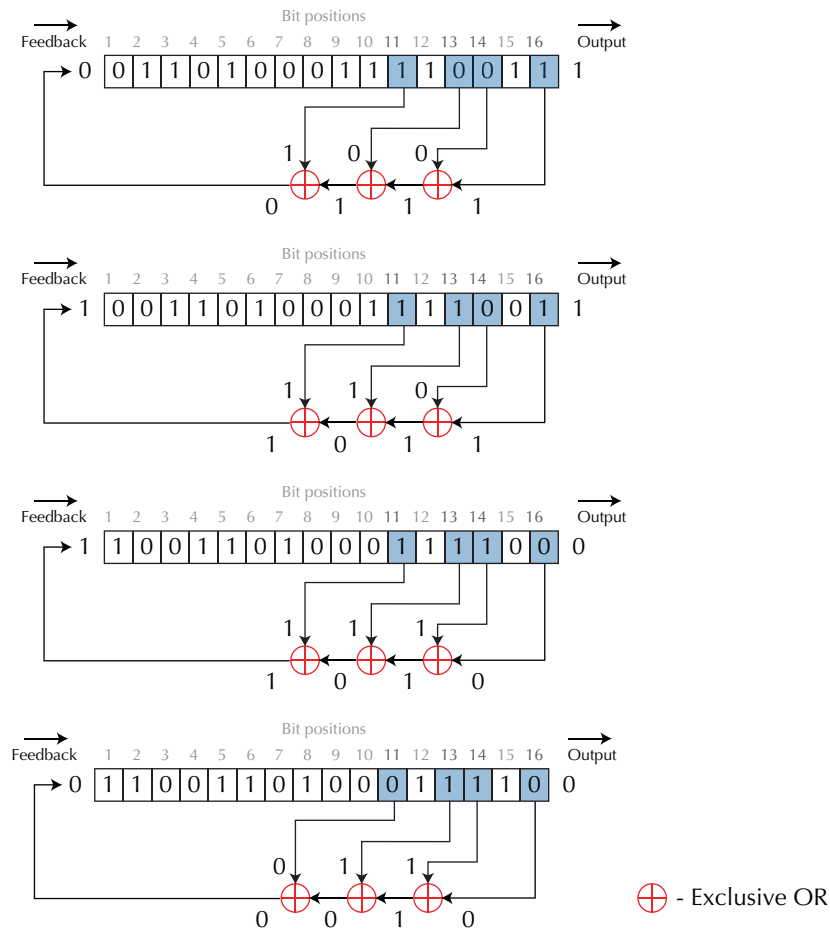


Figure 3.1: Generation of a linear feedback shift register sequence adapted from Aguirre and Kerr (2009) is depicted above (some errors in the third iteration had been corrected). The bit positions of ‘taps’ in the shift register corresponds to the exponents of the primitive polynomial $x^{16} + x^{14} + x^{13} + x^{11} + 1$.

Pseudo-noise consists of a wide variety of broadband signals which have a quasi-flat frequency spectrum however without a deterministic progression of frequency as in chirps. ‘White noise’ referred to in many theoretical analyses, falls under this category as well. These signals are generated based on a random number sequence. The broadband qualities of the signal depends on the quality of the pseudo-random generator used. Pseudo-noise signals can either be uniform, or binary. Normalised uniform signals are generated with a pseudo-random sequence of uniformly distributed real numbers in the range 0 (or -1) to 1 while a normalised binary signal can be generated with the same pseudo-random sequence with an additional function which outputs 0 (or -1) and 1 depending on the random number being less than or greater than 0.5 (or being negative or positive). Unlike chirp signals, pseudo-noise signals do not display gradual phase transitions.

The third class of signals, linear feedback shift register sequences based signals have identical spectral properties to pseudo-noise signals but also have a number of desirable statistical

properties including being deterministic sequences and the possibility to generate multiple unique sequences of the same length having identical spectral properties but with minimal correlation between the sequences (Golomb, 1982; Golomb and Gong, 2004). These sequences and signals are widely used in applications such as cryptography and digital communication systems. A polynomial over the Galois Field $GF(2)$ can be used with a linear feedback shift register as shown in figure 3.1 to generate a binary sequence. When the generating polynomial (such as $x^{16} + x^{14} + x^{13} + x^{11} + 1$ used in the example is figure 3.1) is primitive, the class of generated sequences are known as *maximum length sequences* (MLS) or *m-sequences* and they contain every possible sequence which can be produced by the shift register (Peterson and Weldon, 1972; Cohn and Lempel, 1977). Further description of properties and characteristics of MLS signals is given in section 3.2.2.

Experimental evaluation of different signal waveforms

In order to compare the cross-correlation performance under ‘real’ conditions, sample signals of five different types were transmitted and received using actual transducers (same **projectors** and **hydrophones** used in the experimental evaluation of the relative localisation system) in an enclosed reverberant underwater environment¹. Each of the signals (1.3ms in duration - 127 samples, sampled at 96 000Hz) were repeated 25 times at 5.0Hz. The signal amplitudes of all source signals were normalised such that the transmission power was equivalent for each type. The compared source signals were:

- An MLS signal of length 127.
- A chirp with an up sweep of 750 - 48 000Hz with a logarithmic chirp rate².
- A chirp with an up sweep of 750 - 48 000Hz with a linear chirp rate.
- A pseudo-noise signal based on the MT19937 pseudo-random generator³.
- A uniform pseudo-noise signal based on the MT19937 pseudo-random generator.

MLS	Chirp (Logarithmic)	Chirp (Linear)	Pseudo-noise	Pseudo-noise (uniform)
34.03dB	29.71dB	32.27dB	34.02dB	30.64dB

Table 3.1: Average signal to noise ratios of the different source signals used for cross-correlation as illustrated in figure 3.3.

¹ Cylindrical tank with corrugated metal walls filled with tap water. Diameter 4.2m, depth 1.5m.

² Due to the length of the signal and the sampling rate used, the lowest and highest producible frequencies are 750Hz and 48 000Hz.

³ Implementation of the Mersenne Twister pseudo-random generator are given by Hoe (2002) and Matsumoto (2007) while the theory behind the implementation is given by Matsumoto and Nishimura (1998).

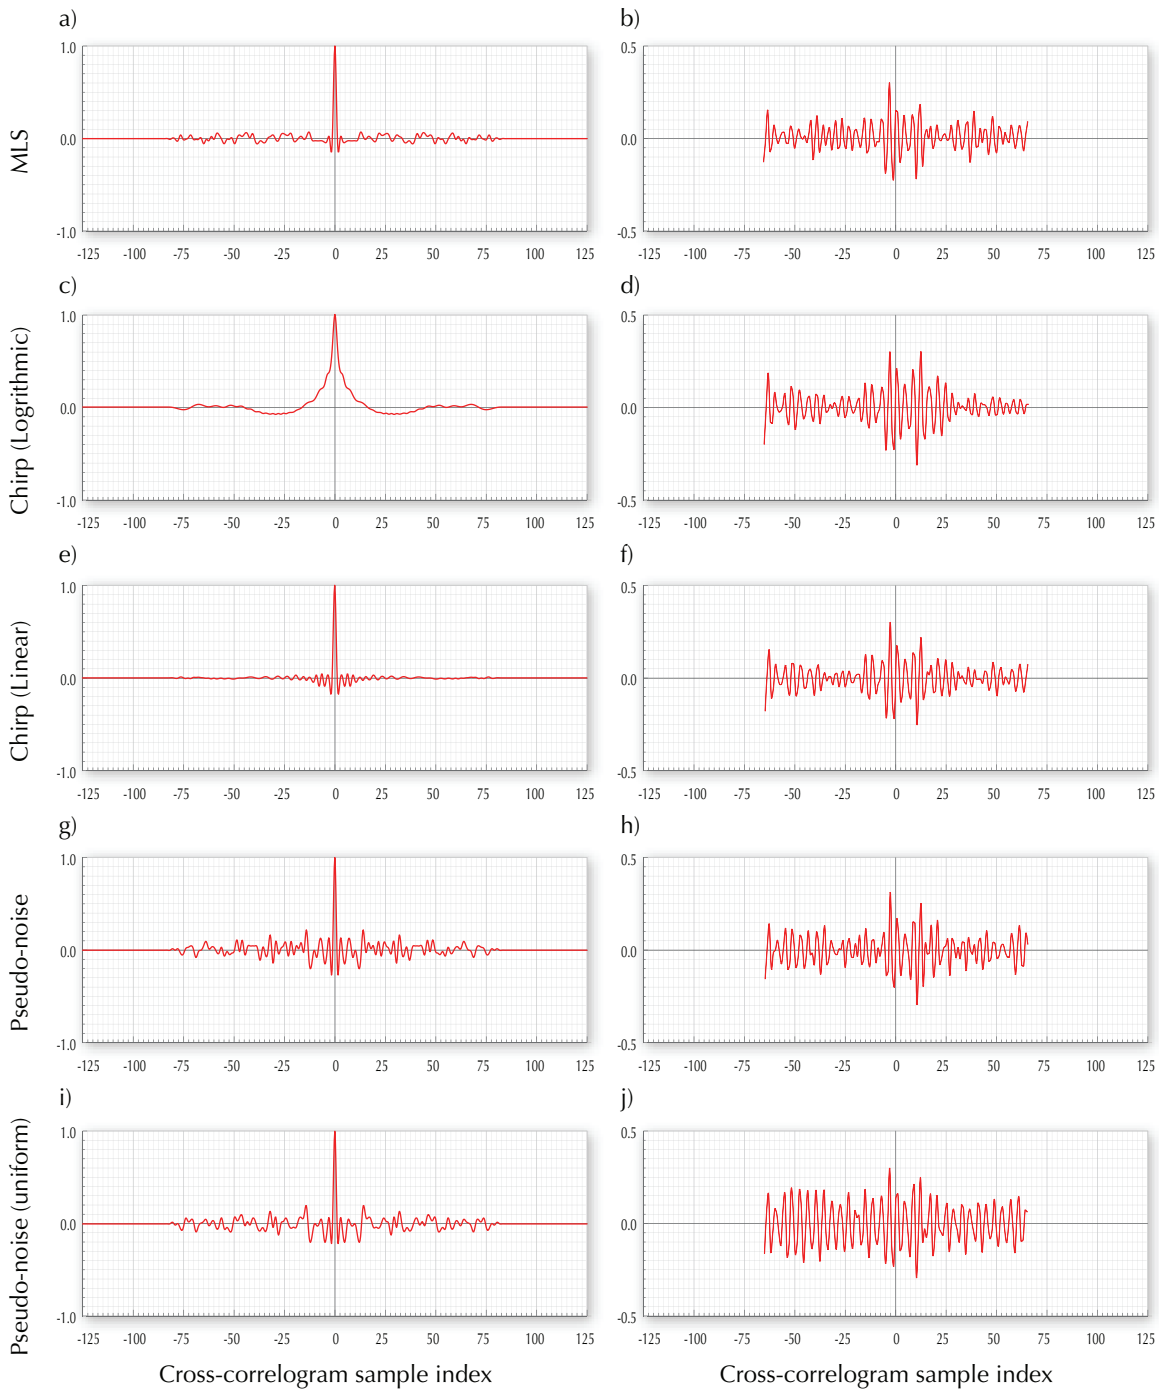


Figure 3.2: The left column shows the auto-correlograms of the five evaluated source signals while the right column shows the cross-correlograms of the experimentally recorded versions of the same signals. The y -axes of these plots represent normalised amplitude.

The signal source was placed 2.0 m away from a pair of receivers and the received signal channels were cross-correlated in the time-domain as described in section 3.1. To compare the ‘real’ and ideal performances, representative cross-correlograms of each of the different experimentally recorded signals are given alongside the auto-correlograms of the source signals in figure 3.2. Due

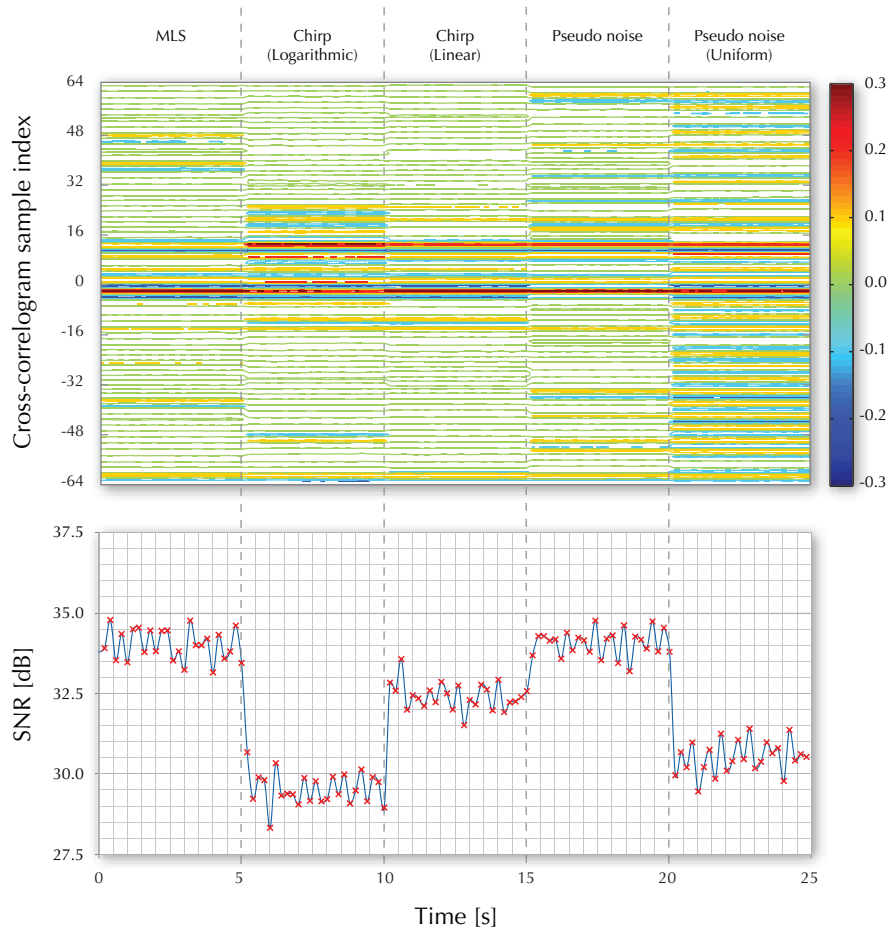


Figure 3.3: Contour plot of the 125 cross-correlograms with colours representing peak heights shown on top with the measured SNR for each of the corresponding signals.

to the placement of the source signal relative to the receivers, cross-correlations of the experimentally recorded signals have a peak in the vicinity of -3 samples. Figure 3.3 shows the contour plots of all 125 cross-correlograms (25 each for the five source signal types) along with the measured signal-to-noise ratio (SNR)¹. The average SNR for each signal type is given in table 3.1. By inspecting the auto-correlation of the different signals shown in figure 3.2, the linear chirp appears to have the most unique peak with side-lobe peak amplitudes at a minimum while the logarithmic chirp has the widest peak. The side-lobe peak amplitudes of the two pseudo-noise signals are comparatively higher than those of the MLS signal. Observing the performance of the experimentally recorded signals, it is apparent that each of them were severely affected by reflected signals which are represented by the many significant side-lobe peaks present in the cross-correlogram. The most prominent secondary peak near the vicinity of $+12$ samples correspond to the bottom reflected signal received on one channel, according to the experimental

¹. The methods used for SNR measurement is explained in section 6.6.

configuration of the transducers within the test tank. When considering the cross-correlation peak detection performance as well as the delivered SNR, MLS signals appeared to perform better than the other compared signals.

3.2.2 Maximum Length Sequences

As mentioned earlier, an MLS is a pseudo-random binary sequence. The statistical properties of these sequences which were first introduced by Golomb (1982) in his first edition published in 1967 have been widely studied since then (Peterson and Weldon, 1972; Cohn and Lempel, 1977; Dunn and Hawksford, 1993; Vanderkooy, 1994). The most attractive of its properties is its auto-correlation function which is essentially a *Dirac delta function* (single sharp peak at zero-shift). Some additional properties of these sequences as given by Aguirre and Kerr (2009) are:

- The number of '0's in the sequence is one less than the number of '1's (balance property)
- For any q which is relatively prime to N , if you choose every q^{th} element in an MLS until the length of the new sequence is N , the resulting sequence is also an MLS.
- Another MLS can also be created by adding a shifted version of the original sequence to the original sequence.

Acoustically transmitted MLS signals (instead of '0's and '1's, these signals consist of -1 and $+1$) are widely used in fields such as room acoustics to measure the impulse response of linear systems without actually using impulsive signal sources (Borish and Angell, 1983; Bradley, 1996). Farina (1998) presents an experimental study of using MLS signals instead of impulsive sources for underwater bottom profiling and concludes that the MLS approach yields a higher SNR and better spatial resolution compared to conventional methods.

While behaving similarly to band-limited white noise in the frequency domain with a *quasi-flat* spectrum up to the Nyquist frequency, MLS signals have the additional advantage of being fully deterministic which makes them accurately reproducible and the ability to have multiple unique signals with the same spectral properties with guaranteed minimal cross-correlation amongst them. These properties combined with the experimental performance mentioned in the previous section makes MLS signals an ideal candidate for the source signal to be used in the relative localisation system presented in this thesis.

There exists many standard algorithms for the generation of MLSs by using generating polynomials with different degrees which result in MLS signal sets of different lengths. The degree n of the generating polynomial, which is also the length of the shift register (also known as the degree of the MLS signal) governs the length l_{MLS} of the sequence as:

$$l_{MLS} = 2^n - 1 \quad (3.3)$$

The length of the sequence and the employed sampling rate of the analogue to digital converter determines the duration of the signal. Even though longer MLS signals give a better resolution of the cross-correlation peak resulting in higher precision of the consequent estimation, a longer duration has its drawbacks. Due to undesirable echoes in cluttered or enclosed environments, higher processing overheads and lower update rate for the overall estimation system (presented in chapter 5) associated with a longer duration signal, a relatively short MLS signal of degree 7 (length 127) is employed as a compromise. The MLS generation algorithm (see appendix B) produces 18 unique sequences of this length which exhibits extremely low correlation between each other.

3.3 Cross-correlation of MLS signals

Plots a) and b) in figure 3.4 depict two unique MLS signals picked out of the different length-127 sequences produced by the generating algorithm. The original pure square-wave signals have been conditioned with sub-sample interpolation and the amplitude have been normalised. The auto-correlation of the first signal is shown in plot c) with a maximum lag equal to the length of the signal. The highly desirable narrowness and height of the correlation peak can be observed from this plot. Plot d) is the result of a cross-correlation between the two different MLS signals shown in plots a) and b). As is evident, different MLS signals show almost no correlation amongst each other.

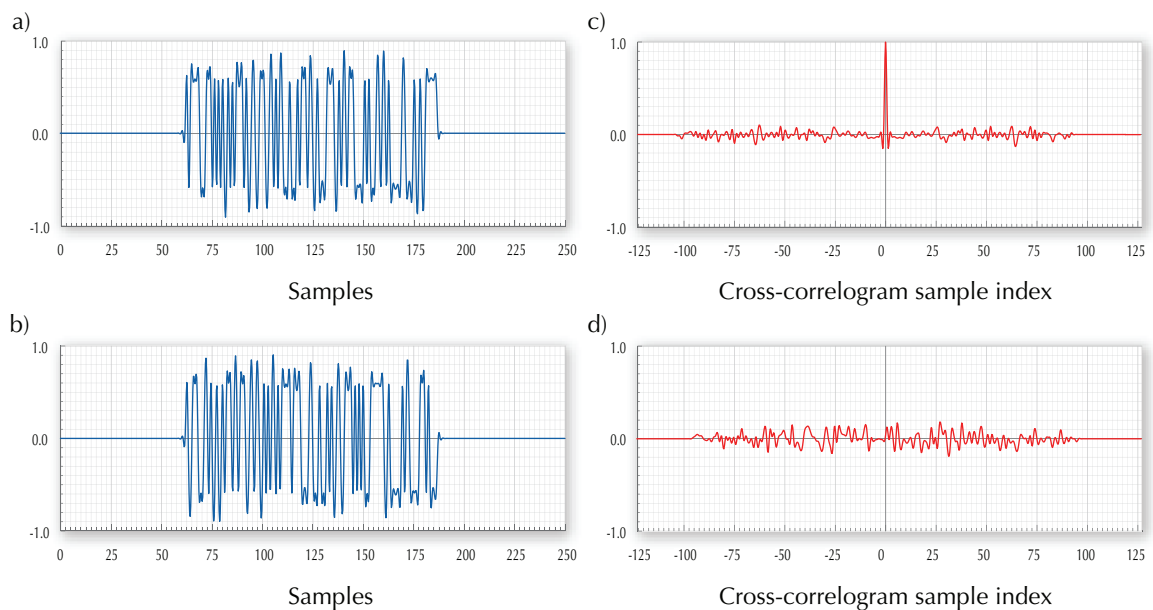


Figure 3.4: Plot c) shows the narrow sharp peak resulting from the auto-correlation of the MLS signal shown in a) while plot d) shows the cross-correlation between the two different MLS signals given in plots a) and b). In each plot, the *y-axis* represents normalised amplitude.

3.3.1 Effect of uncorrelated noise

The robustness of MLS signals against uncorrelated noise, even at relatively short lengths, is demonstrated as follows. Two copies of the MLS signal shown in figure 3.4a are used where one is shifted by 8 samples, such that the first channel (figure 3.5.b) leads. Then the two signals are contaminated with additive white Gaussian noise and cross-correlated. The average signal powers were equal (-3 dBFS) when mixing white noise with the MLS signal representing a SNR of 0dB. The two signals (plots a) and b)) and the resulting cross-correlogram (plot c)) is shown in figure 3.5. The position and width of the peak of the cross-correlation as well as the area surrounding the peak remains unaltered in this case. However, the height of the peak has been reduced to about 70% of its original value compared to the auto-correlation plot in figure 3.4c.

3.3.2 Effect of mixing and shifting

The plots in figure 3.6 show the behaviour of the cross-correlation for shifted signals. Instead of cross-correlating a simple sample-shifted signal, which might suggest a role played by the relative difference in the leading edges of the two signals, a mixed signal was used on one channel. The two different MLS signals shown in plots b) and e) were shifted by 16 samples and added such that the signal in b) leads. This mixed signal was cross-correlated separately with its component signals, each shifted by 8 samples such that the mixed signal leads. If it was merely leading edge detection, in both cases the relative delay between the two channels would have been measured as -8 samples. However, as seen in plots c) and f), the lags are correctly represented in the cross-

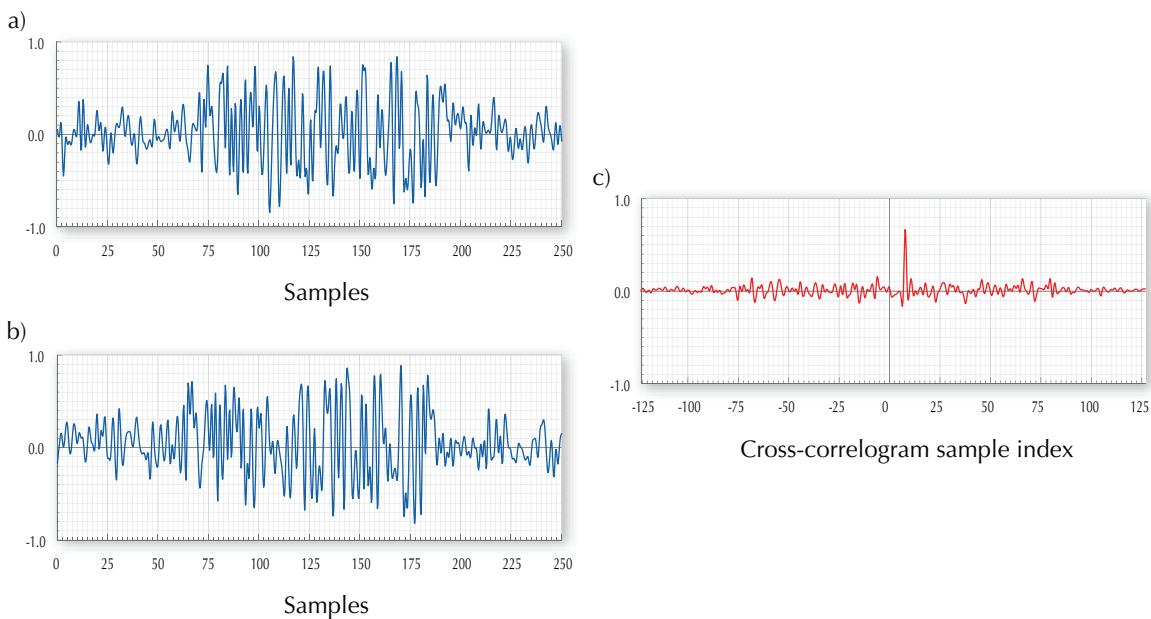


Figure 3.5: Plot shown in c) results from the cross-correlation of the shifted noisy MLS signals shown in a) and b). In each plot, the *y-axis* represents normalised amplitude.

correlogram with peaks at -8 samples and $+8$ samples respectively. In addition, the height reduction of the two cross-correlogram peaks are similar to that observed in the noise contamination earlier. This suggests that the sporadic peaks introduced to the frequency spectrum during the mixing process behaves as uncorrelated noise, hence lowering the effective SNR of the signals.

3.3.3 Effect of non-linear transducers

As it was shown by figure 3.5.c, contamination by additive white noise does not contribute to a noticeable deterioration of the cross-correlation performance of MLS signals, apart from a slightly lower height for the peak. In theory, addition of two 'flat' frequency spectra should again result in a 'flat' frequency spectrum, hence the spectral properties of the MLS signal which provides the narrowness of the peak are preserved. However, this is not necessarily the case when these signals are transmitted and received via transducers with a non-linear frequency response

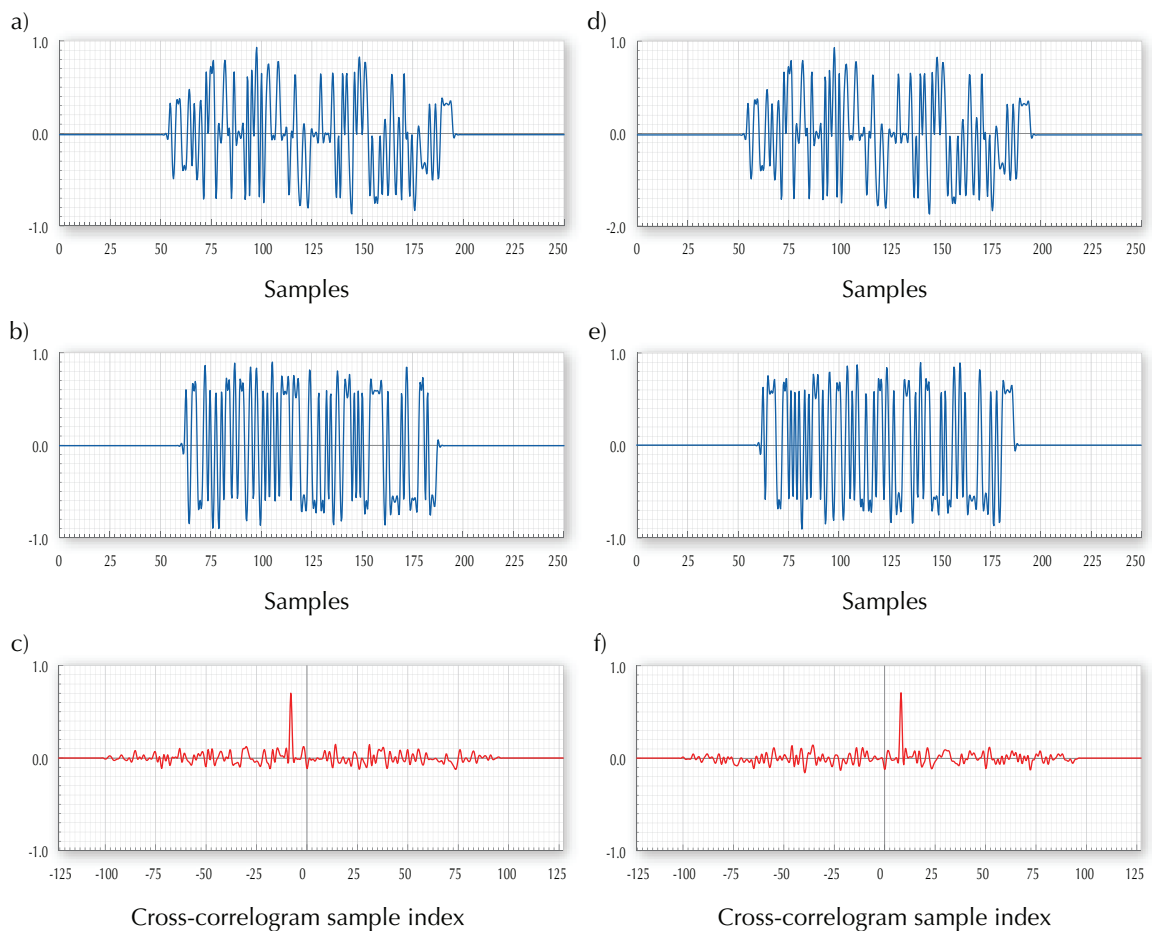


Figure 3.6: Plots in a) and d) represents a signal which is a mixture of the signals in b) and e). This mixed signal is cross-correlated with each of the two component signals and the resulting cross-correlograms are depicted in c) and f). In each plot, the *y-axis* represents normalised amplitude.




Figure 3.7: The frequency response curve of the hydrophone reproduced from the AQ-2000 data sheet (Benthos, 2001)

within the bandwidth of the signal. Due to the sampling rate used ($f_s=96\,000\text{Hz}$), the signals are band limited by $48\,000\text{Hz}$. The frequency response of the Benthos AQ-2000 transducers which are used as transmitting projectors as well as receiving hydrophones in the relative localisation system is shown in figure 3.7. As seen from this logarithmic plot, the transducers have a resonance near $20\,000\text{Hz}$ and an anti-resonance near $25\,000\text{Hz}$ which results in a highly non-linear response within the signal bandwidth.

To test the effect of this potential frequency filtering effect introduced by the transducers, the frequency response shown in figure 3.7 was empirically modelled¹ and implemented as an FFT filter. The filter was then applied to white noise contaminated MLS signals shown earlier in plots a) and b) of figure 3.5 and the two filtered signals were cross-correlated. This setup is illustrated in figure 3.8.a. The effect it has on the cross-correlogram shown in plot a) of figure 3.9 can be compared to the cross-correlogram shown in plot c) of figure 3.5, which was produced by the same source signals but without frequency filtering. Resonance of the transducer near $20\,000\text{Hz}$ appears as the dominant frequency in the resulting cross-correlogram. For comparison, plot c) of figure 3.9 shows the cross-correlogram of two signal channels received via two AQ-2000 transducers (separated by 0.3m) when an MLS (length-127, duration 1.3ms) signal was transmitted via another AQ-2000 transducer. The signal travelled a distance of 2.0m underwater and the cross-correlation reveals a delay of $+8$ samples ($83.33\mu\text{s}$) between the channels. This setup is schematically illustrated in figure 3.8.c. Both the cross-correlograms, one from cross-correlating the ‘real’ and the other from cross-correlating the ‘simulated’ signals, shows the dominance of the $20\,000\text{Hz}$ resonance of the transducer throughout the plots. Even though the position and height of the peak is not affected, the uniqueness of the peak had been lost by being surrounded by an envelope of decaying side-lobes. As revealed in chapter 7, this decreases the accuracy of the localisation system when used in enclosed and cluttered environments due to peaks caused by reflected signals.

¹. The shape of the frequency response curve was replicated as the shape of a frequency response curve of an FFT filter.

In order to address this issue, another filter was empirically modelled which had the inverse frequency response of the one modelled earlier to represent the response of the transducer. The results of applying the new FFT filter¹ to the simulated and the real signals used earlier and cross-correlating the channels are shown in plots b) and d) of figure 3.9 respectively. The source signals, before and after being filtered is shown in appendix A (figure A.1). The setups used in these instances are schematically depicted in figures 3.8.b and 3.8.d. As depicted by the plots the uniqueness, narrowness and the height of the peak is restored making it possible to unambiguously locate it.

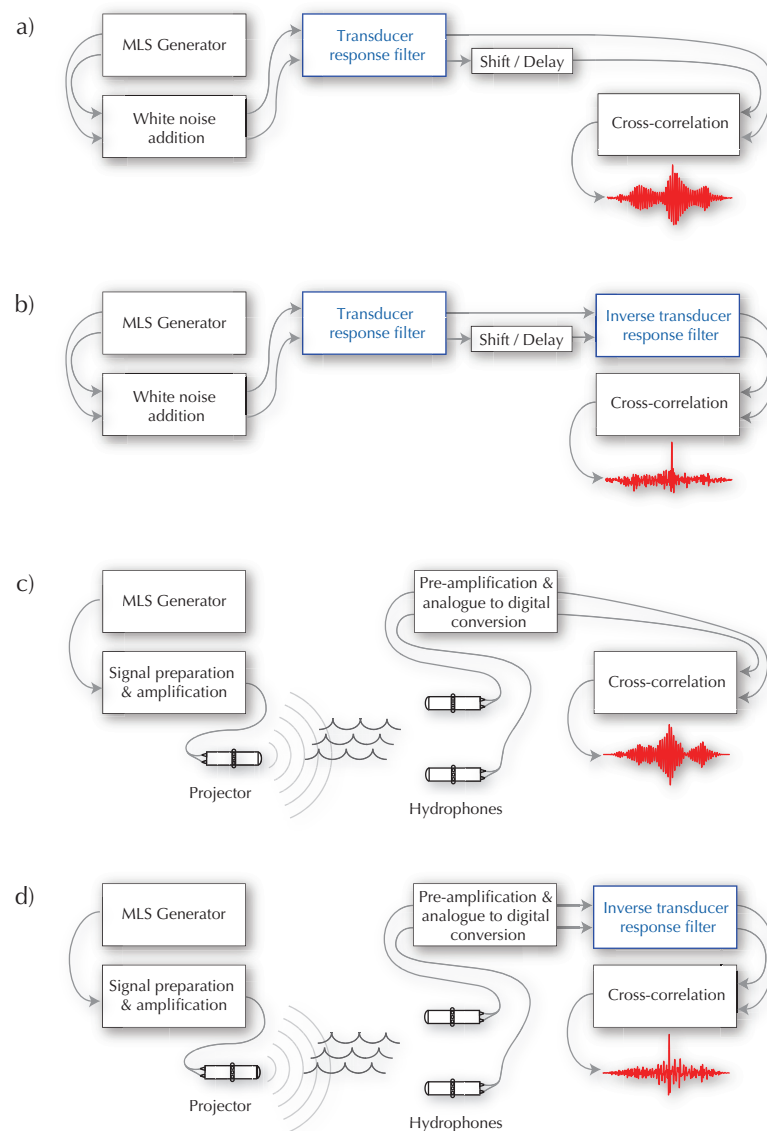


Figure 3.8: The four different setups producing the four different cross-correlograms shown in figure 3.9. The source signals that provide inputs to the cross-correlation are given in section A.1.

¹ The filter was applied to the two signal channels prior to being cross-correlated.

3.4 Discussion

An experimental evaluation of the performance of MLS signals alongside other types of broadband signals with regard to cross-correlation peak detection and measured SNR was presented earlier in this chapter. In a harsh reverberant acoustical environment, the MLS signals performed better with regard to peak detection and delivered SNR, in comparison with chirps and pseudo-noise signals. Further statistical properties of MLS signals, their behaviour characteristics with regard to exposure to noise and the effect mixing and shifting signals are given in the subsequent sections.

Finally, the effects of the non-linear frequency response of transducers on the transmitted MLS signals are discussed and a strategy to address this effect is proposed. It must be emphasised that the filter process described in the earlier section only accounts for the transducer characteristics and not those of the propagation medium. Since channel characteristics of the underwater medium greatly varies with depth, temperature and salinity as well as environmental features such as the composition and texture of the bottom (sediment/sand/vegetation), modelling the transducer characteristics are more practical. In the context of an autonomous underwater vehicle, it is far more convenient to account for the characteristics of on-board sensors than to have access to a model of the channel characteristics of the operating medium. However, the underwater channel does indeed have an effect on the transmitted MLS signal as seen in figure 3.9.d

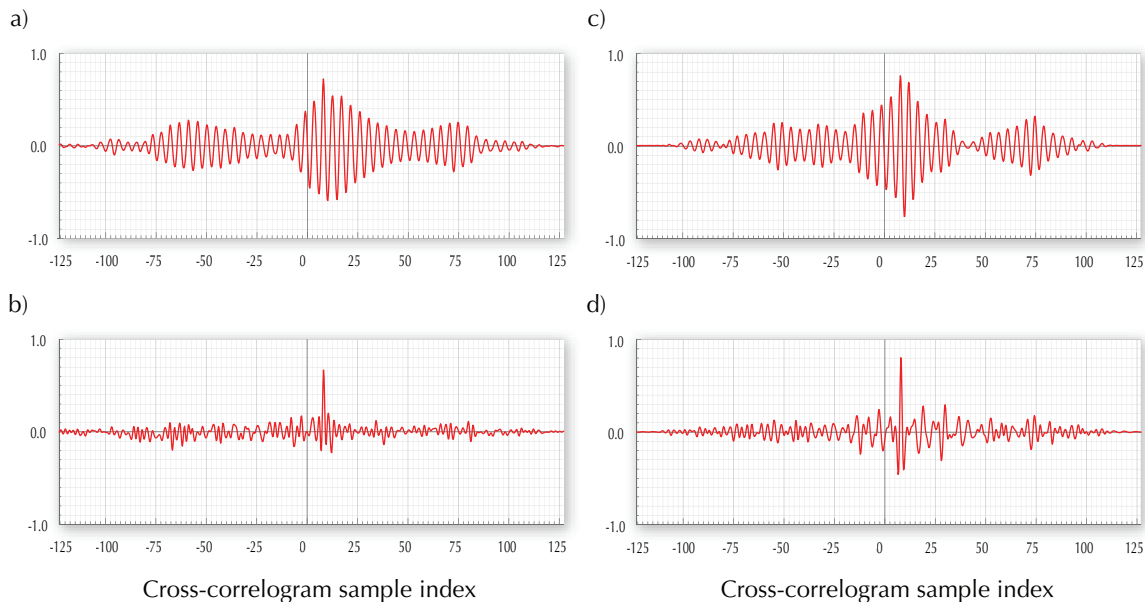


Figure 3.9: Plots in a) and b) represents the cross-correlograms of the shifted MLS signals contaminated with additive white noise, first filtered with the transducer frequency response model, then filtered with the inverse of that filter. Plot c) shows the cross-correlogram resulting from two actual signal channels (with a shift of +8 samples) which were transmitted and received using the transducers. Plot d) shows the resulting cross-correlogram when the inverse transducer filter was applied to the signal channels prior to cross-correlation. In each plot, the *y-axis* represents the normalised amplitude of the signals.

where the filtering process does not completely reconstruct the original cross-correlogram (figure 3.4.c). Furthermore, the power of the received signals are greatly reduced by the inverse filtration process since most of the transmitted signal power is around the resonance frequency of the transducer. This attenuation affects the distance which the MLS signals can be effectively transmitted at a given transmission power. For the application concerned, the achievable operating distances at a low transmission power and the cross-correlation peak detection performance given by the use of MLS signals are amply sufficient and far surpasses other alternative signal waveforms which were evaluated.

The next chapter describes the distance and angle estimations carried out during the process of localising source signal pings. The methodology and basic measurement schemes are also described along with identification and analysis of different classes of errors affecting the estimated quantities.