

Robust Optimal and Adaptive Pulse Compression for FM Waveforms

By

Dakota Henke

Submitted to the Department of Electrical Engineering and Computer Science and the
Graduate Faculty of the University of Kansas
in partial fulfillment of the requirements for the degree of
Master of Science

Dr. Shannon Blunt, Chairperson

Committee members

Dr. Christopher Allen

Dr. James Stiles

Date defended:

September 03, 2015

The Thesis Committee for Dakota Henke certifies
that this is the approved version of the following thesis :

Robust Optimal and Adaptive Pulse Compression for FM Waveforms

Dr. Shannon Blunt, Chairperson

Date approved: September 03, 2015

Abstract

The least-squares mismatched filter (LS MMF) is a pulse compression method used to suppress range sidelobes. This work provides a description of the adjustments needed such that the LS MMF, initially derived for codes, can be robustly applied to arbitrary FM waveforms. Additionally, the effects of range straddling and Doppler on the LS MMF are examined. The effects of straddling on mismatch loss is well known, what is less appreciated is the effect straddling has on the range sidelobes. This work outlines methods that alleviate some of the degradation in sidelobe levels due to straddling. Making the LS MMF more robust to Doppler is also investigated. Adaptive Pulse Compression (APC) is another pulse compression algorithm that has recently been modified to be applicable to FM waveforms. This work includes analysis of these modifications via simulation and measured data. The effects of straddling and Doppler on APC are also examined. These robust pulse compression methods are applied to measured data, demonstrating their viability for real FM-based radar systems.

Acknowledgements

I would like to thank my advisor, Dr. Shannon Blunt, and the rest of the University of Kansas Radar Systems Lab for their help in making both this thesis and my research these last two years possible. I would also like to thank the US Naval Research Lab for funding my work and the work of others in the Radar Systems Lab. The University of Kansas Electrical Engineering and Computer Science Department also deserves recognition for providing me with opportunities to be involved in meaningful work for the past six years. Last but not least, I would like to thank my wife, Brittany, and daughter, Everette, for helping me push through the completion of this thesis, as they provided motivation when I needed it most.

Contents

1	Introduction and Background	1
1.1	Radar Signal Processing Basics	2
1.2	Radar Pulse Compression	5
1.3	Radar Waveforms	7
1.4	The Least-Squares Mismatch Filter	12
1.5	Adaptive Pulse Compression	14
2	Least Squares Mismatched Filtering	16
2.1	LS MMF Modification for Application to FM Waveforms	16
2.2	Idealistic Simulations and Results	22
2.3	Straddling Effects	30
2.4	Averaging Methods	33
2.4.1	Filter Averaging	34
2.4.2	Waveform Averaging	37
2.4.3	Combination Averaging	40
2.5	Doppler Effects	43
2.6	Doppler Averaging	45
3	Analysis of Adaptive Pulse Compression	53
3.1	APC Modifications for Application to FM Waveforms	53
3.2	Idealistic Simulations and Results	60

3.3	Straddling Effects on APC	63
4	Pulse Compression of Free-Space Measurements	73
4.1	Test Setup	73
4.2	Pulse Compression Results	75
4.2.1	Least-Squares Mismatch Filter Responses	75
4.2.2	APC Responses to Measured Data	78
4.2.3	Comparison of the Averaged LS MMFs and APC	79
4.2.4	Effects of Super-Resolution on Measured Responses	82
5	Conclusions	88
5.1	Future Work	90
A	Appendix	97

List of Figures

1.1	Visual representation of range response PSL, 3 dB range resolution, and SNR. . . .	4
1.2	Matched filter responses for unmodulated and modulated pulses. Range resolution and PSL improve when modulated pulses are used.	6
1.3	Taken from [1], characteristics of a linear FM waveform. Plot (a) shows the frequency as a function of time, (b) the spectrum of the LFM, and (c) its auto-correlation.	10
1.4	Taken from [1], characteristics of a non-linear FM waveform. Plot (a) shows the frequency as a function of time, (b) the spectrum of the NLFM, and (c) its auto-correlation.	11
2.1	The spectrum of an LFM with three oversampling rates: $K = 1$, corresponding to the 3 dB bandwidth $K=3$, and $K=5$	17
2.2	Matched filter and LS MMF responses for LFM and PD with $K = 1$, $M = 1$, and $BT = 64$. The matched filter outperforms the LS MMF in terms of MML and PSL, while the LS MMF has lower sidelobe levels near the mainlobe, with an increase near the ends of the filter response.	19
2.3	LS MMF responses for LFM and PD with oversampling factors $K = 2$ through $K = 7$, $M = 1$, and $BT = 64$. As K increases, sidelobe levels improve (decrease) monotonically, while MML increases monotonically. Although there are more values in the LS MMF responses as K increases, the pulse width remains constant for all LS MMFs.	21

2.4	LS MMF responses for LFM and PD with $K = 5$, filter length factors of $M = 1$ to $M = 4$, and $BT = 64$. As M increases, the PSLs decrease monotonically and the MML remains relatively constant for $M > 1$	22
2.5	Matched filter and LS MMF responses for LFM for five diagonal loading values. These diagonal loading values (δ), in decreasing order, are 16.63 dB, 6.63 dB, -3.37 dB, -13.37 dB, and 0. The responses on the right are the zoomed-in mainlobes of the responses on the left. As δ increases, MML decreases while mainlobe width and sidelobe levels increase.	24
2.6	Matched filter and LS MMF responses for PD for five diagonal loading values. These diagonal loading values (δ), in decreasing order, are 18.86 dB, 8.86 dB, -1.14 dB, -11.14 dB, and 0. The responses on the right are the zoomed-in mainlobes of the responses on the left. As δ increases, MML decreases while mainlobe width and sidelobe levels increase.	25
2.7	LS MMF responses for LFM with beamspoilage. The number of rows that were zeroed above and below the m^{th} row of $\tilde{\mathbf{A}}$, in decreasing order, are 13, 10, 7, 5, 2, and 0. The responses on the right are the zoomed-in mainlobes of the responses on the left. As more rows are zeroed, the mainlobe width increases and surpasses the matched filter mainlobe width. The minimum MML occurs when 7 rows are zeroed.	28
2.8	LS MMF responses for PD with beamspoilage. The number of rows zeroed above and below the m^{th} row of $\tilde{\mathbf{A}}$, in decreasing order, are 15, 12, 9, 6, 3, and 0. The responses on the right are the zoomed-in mainlobes of the responses on the left. As more rows are zeroed, the mainlobe width increases and surpasses the matched filter mainlobe width. The minimum MML occurs when 9 rows are zeroed.	29
2.9	A pictorial example of straddling, using a sample LFM. The black dots at integer multiples of T_S represent the discretized waveform model for filter construction. The red dots, shifted $0.5 T_S$ in delay, represent the "worst-case" straddled samples.	31

2.10	Matched filter and LS MMF responses for LFM and PD, for both ideal and "worst-case" straddled sampling. The matched filter shows minimal degradation in performance when straddling is present. The LS MMF shows a significant increase in PSL with straddling compared to the ideal sampled case.	32
2.11	LS MMF responses for $L = 1$ to $L = 5$ filter averaging for LFM with worst-case straddling. As L increases, the PSL improves monotonically but with diminishing improvements for successive values of L	35
2.12	LS MMF responses for $L = 1$ to $L = 5$ filter averaging for PD with worst-case straddling. As L increases, the PSL improves monotonically but with diminishing improvements for successive values of L	36
2.13	LS MMF responses with $P = 1$ to $P = 5$ waveform averaging for LFM with worst-case straddling. As P increases, the PSL improves monotonically but with diminishing improvements for successive values of P	38
2.14	LS MMF responses with $P = 1$ to $P = 5$ waveform averaging for PD with worst-case straddling. As P increases, the PSL improves monotonically but with diminishing improvements for successive values of P	39
2.15	LS MMF responses with filter and waveform averaging for LFM with worst-case straddling. The responses on the left were generated with L varied from 1 to 5 and $P = 2$; the responses on the right were generated with L being varied from 1 to 5 and $P = 5$. The bottom plots are the zoomed-in views of the sidelobes near the mainlobe of the LS MMF response.	41
2.16	LS MMF responses with filter and waveform averaging for PD with worst-case straddling. The responses on the left were generated with L varied from 1 to 5 and $P = 2$; the responses on the right were generated with L being varied from 1 to 5 and $P = 5$. The bottom plots are the zoomed-in views of the sidelobes near the mainlobe of the LS MMF response.	42

2.17	Matched filter and LS MMF responses for LFM and PD, for ideal and Doppler shifted received waveforms. The matched filter shows little degradation in performance with Doppler. The LS MMF has an increase in sidelobe levels near the mainlobe with Doppler, compared to the ideally sampled case.	45
2.18	LS MMF responses for LFM and PD for Doppler shifts between $+0.05$ to -0.05 radians. At zero Doppler, the sidelobes for the LS MMF are at their lowest. Sidelobe levels increase monotonically as the Doppler shift increases in magnitude. . .	46
2.19	Doppler averaged LS MMF responses with $J = 2$ and $J = 5$ for LFM for Doppler shifts from $+0.05$ to -0.05 radians. Averaging was linearly done from -0.04 to $+0.04$ radians. The LS MMF has the lowest sidelobes at zero Doppler, and sidelobes increase monotonically moving away from zero. The sidelobe levels of the $J = 2$ and $J = 5$ LS MMFs are worse than those found in the non-averaged $J = 1$ case.	48
2.20	LS MMF responses with Doppler averaging for LFM with three different Doppler-shifted received waveforms (0 , -0.02 , and -0.04 radians) and $J = 1$, $J = 2$, $J = 3$, and $J = 5$. Averaging occurs on the interval of -0.04 to $+0.04$ radians. From this figure, no noticeable change in PSL is observed when using Doppler averaging. . .	49
2.21	Doppler averaged LS MMF responses with $J = 2$ and $J = 5$ for PD for Doppler shifts from $+0.05$ to -0.05 radians. Averaging was linearly done from -0.04 to $+0.04$ radians. The LS MMF has the lowest sidelobes at zero Doppler, and sidelobes increase monotonically away from zero. The sidelobe levels of the $J = 5$ LS MMF are worse than those for the $J = 2$ LS MMF. The sidelobe levels from both the $J = 2$ and $J = 5$ responses are worse than those found in the non-averaged $J = 1$ case.	50

2.22	Doppler averaged LS MMF responses for PD with three Doppler shifted received waveforms (0, -0.02 , and -0.04 radians) and with $J = 1$, $J = 2$, $J = 3$ and $J = 5$. Averaging occurred on the interval of -0.04 to $+0.04$ radians. From this figure, no noticeable change in PSL is observed when using Doppler averaging.	51
2.23	Doppler averaged LS MMF responses for LFM and PD for Doppler shifts from -0.05 to $+0.05$ for $J = 5$. Averaging was linearly done from -0.04 to 0 radians. The best sidelobe performance occurs at -0.02 radians, halfway between the end points of the averaging interval. Moving away from -0.02 radians, sidelobe levels increase. The Doppler averaging was completed assuming that the target is receding relative to the radar system.	52
3.1	Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $K = 1$. Neither the filters nor APC are able to discern the smaller target with confidence.	54
3.2	Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $K = 3$. The LS MMF is able to discern the smaller target, but the matched filter and APC cannot.	55
3.3	Matched filters, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $K = 5$. The LS MMF is able to discern the smaller target, but the matched filter and APC cannot.	55
3.4	LS MMF and APC responses for LFM and PD with two and three adaptive stages and $K = 5$. The targets are more distinguishable with three adaptive stages of APC than with two.	56
3.5	Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $z = 1$ row zeroed, which results in a resolution that is slightly coarser than APC implemented in Figure 3.3. Only the LS MMF is able to discern the smaller target.	57

3.6	Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $z = 3$ rows zeroed, resulting in a coarser resolution than APC responses when $z = 1$. APC is able to somewhat distinguish the target for PD, but is unable to do so for LFM.	58
3.7	Matched filters, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $z = 5$ rows zeroed out on either side of n , which results in a resolution that is coarser than APC implemented in Figure 3.6. The PD-based implementation of APC is able to discern the smaller target with confidence, but the LFM-based implementation is still unable to discern the smaller target.	59
3.8	LS MMF and APC responses for LFM and PD with two and three adaptive stages. Five rows were zeroed out above and below n for APC. There is no noticeable performance gain from using a third adaptive stage of APC; the smaller target is distinguishable after two adaptive stages for both LFM and PD.	60
3.9	Matched filter, LS MMF and APC responses for LFM. The top two plots show the responses of the matched filter and APC, with APC initialized using the matched filter. The bottom two plots show the responses of the LS MMF and APC, with APC initialized using the LS MMF. When the LS MMF is used to initialize APC, only one adaptive stage is needed to discern the smaller target. Two adaptive stages are for target detection when APC is initialized with the matched filter.	61
3.10	Matched filter, LS MMF and APC responses for LFM. The top two plots show the responses of the matched filter and APC, with APC initialized using the matched filter. The bottom two plots show the responses of the LS MMF and APC, with APC initialized with the LS MMF. When the LS MMF or the matched filter are used to initialize APC, only one adaptive stage is needed and the filter can discern the smaller target. This was not seen with LFM because LFM does not have as low of range sidelobes as the PD filter responses.	62

3.11	Matched filter, LS MMF and APC responses for LFM in a target-rich environment. Two adaptive stages are needed to discern the smaller targets when the matched filter is used to initialize APC, but only one adaptive stage is needed when the LS MMF is used as the initialization filter.	63
3.12	Matched filter, LS MMF and APC responses for PD in a target-rich environment. As with the plots generated using LFM, two adaptive stages are needed to discern the smaller targets when the matched filter is used to initialize APC, but only one adaptive stage is needed when APC is initialized with the LS MMF.	64
3.13	Matched filter, LS MMF and APC responses for LFM in a two-target environment with worst-case straddling. The top two plots show the responses of the matched filter and APC initialized with the matched filter. The bottom two plots show the responses of the LS MMF and APC initialized with the LS MMF. Neither the matched filter nor the LS MMF can discern the smaller target. APC initialized with the matched filter requires two adaptive stages to discern the smaller target, while when APC is initialized with the LS MMF, only one adaptive stage is needed.	65
3.14	Matched filter, LS MMF and APC responses for PD in a two-target environment with worst-case straddling. The top two plots show the responses of the matched filter and APC initialized with the matched filter. The bottom two plots show the responses of the LS MMF and APC initialized with the LS MMF. Neither the matched filter nor the LS MMF can discern the smaller target. The smaller target can be found after one adaptive stage for both the matched filter and LS MMF initializations.	66
3.15	Matched filter, LS MMF and APC responses for LFM in a target-rich environment with worst-case straddling. APC initialized with the matched filter requires two adaptive stages to discern the smaller targets, while when APC is initialized with the LS MMF, only one adaptive stage is needed, which corroborates the observations for the two target scenario.	67

3.16	Matched filter, LS MMF and APC responses for PD in a target-rich environment with worst-case straddling. The matched filter initialization of APC is able to discern the smaller targets using two adaptive stages. In this LS MMF initialization of APC, only one adaptive stage is needed to discern the smaller targets.	68
3.17	Taken from [2], this figure visually depicts the eclipsing effect at both the beginning and end of the pulse repetition interval.	69
3.18	Matched filter, LS MMF and APC responses for LFM in a target-rich environment with worst-case straddling and eclipsing. The LS MMF was used to initialize APC. Unlike the matched filter and LS MMF, which experience an increase in sidelobes due to eclipsing, APC is relatively robust to eclipsing effects. The second iteration of APC yields some slight improvement in target detection, even in the eclipsed regions.	70
3.19	Matched filter, LS MMF and APC responses for PD in a target-rich environment with worst-case straddling and eclipsing. The LS MMF was used to initialize APC. Unlike the matched filter and LS MMF, which experience an increase in sidelobes due to eclipsing, APC is relatively robust to eclipsing effects. The second iteration of APC yields some slight improvement in target detection, even in the eclipsed regions.	71
3.20	Matched filter, LS MMF and APC responses for LFM and PD in a two-target environment. The larger target has a Doppler shift of 0.04 radians, while the smaller target has no Doppler shift. The LS MMF was used to initialize APC. For both waveforms, the LS MMF and APC are able to discern the smaller target, and APC further suppresses sidelobes below the levels in the LS MMF response.	72
4.1	Equipment used to record measurements. The radar is quasi-monostatic, with the right antenna used to transmit and the left to receive. The antennas are roughly 15 meters from the ground and are directed at two large buildings on the University of Kansas campus.	74

4.2	Radar field of view. The star near the bottom of the map gives the location of the antennas, and the targets that may produce significant reflections are denoted by triangles. The antennas are aimed at the Dole Institute of Politics, but given the beamwidth of the antennas, returns from the Lied Center are expected.	76
4.3	Pulse compressed responses for both the LFM and PD, using the matched filter, non-averaged LS MMF, and averaged LS MMFs. The LS MMFs have better performance with respect to sidelobe levels than the matched filter.	77
4.4	Pulse compressed response of APC for both waveforms, with matched filter initialization and LS MMF initialization. The LS MMF initialized version of APC does a better job of suppressing sidelobes after one adaptive stage.	79
4.5	Matched filter, LS MMF, and APC pulse compressed responses for both waveforms with 25% and 75% eclipsing. The matched filter and LS MMF responses experience significant degradation in sidelobe levels compared to APC, making them less capable when eclipsing occurs.	80
4.6	Pulse compressed responses with LFM for the matched filter, LS MMFs, and APC out to a range of 500 meters (and -100 meters). There is little difference in the filter responses beyond 125 meters.	81
4.7	Pulse compressed responses with PD for the matched filter, LS MMFs, and APC out to a range of 500 meters (and -100 meters). These pulse compression responses, when compared to those for LFM, have lower sidelobes and deeper nulls between targets.	81
4.8	A zoomed-in version of Figure 4.6, focused on the first 150 meters (and -50 meters) so differences between the pulse compression methods can be observed. The LS MMFs and APC are able to discern more targets than the matched filter.	83

4.9	A zoomed-in version of Figure 4.7, focused on the first 150 meters (and -50 meters) so differences between the pulse compression methods can be observed. The LS MMFs and APC are able to discern more targets than the matched filter. These responses, when compared to the LFM pulse compression responses, demonstrate better sidelobe suppression and target detection.	84
4.10	Pulse compressed responses of the LS MMF with a varying number of zeroed rows. For LFM, 7 zeroed rows results in the optimal LS MMF response with respect to MML. Zeroing out fewer rows results in finer resolution, deeper nulls, and significant (> 1 dB) MML. Zeroing out more than the optimal number of rows results in coarser resolution, shallower nulls, and lower sidelobe levels.	85
4.11	Pulse compressed responses of the LS MMF with a varying number of zeroed rows. For PD, 9 zeroed rows results in the optimal LS MMF response with respect to MML. Zeroing out fewer rows results in finer resolution, deeper nulls, and significant (> 1 dB) MML. Zeroing out more than the optimal number of rows results in coarser resolution, shallower nulls, and lower sidelobe levels.	86
A.1	MMLs and PSLs of the filter averaged LS MMF response for each straddled subshift of the received waveform. Generated for LFM and L varied from 1 to 5. . . .	98
A.2	MMLs and PSLs of the filter averaged LS MMF response for each straddled subshift of the received waveform. Generated for PD and L varied from 1 to 5. . . .	99
A.3	MMLs and PSLs of the waveform averaged LS MMF response for each straddled subshift of the received waveform. Generated for LFM and P varied from 1 to 5. . .	100
A.4	MMLs and PSLs of the waveform averaged LS MMF response for each straddled subshift of the received waveform. Generated for PD and P varied from 1 to 5. . . .	101
A.5	PSLs of the LS MMF response for each straddled subshift of the received waveform. Generated for LFM, $P = 2$ and $P = 3$, and L varied from 1 to 5.	102
A.6	PSLs of the LS MMF response for each straddled subshift of the received waveform. Generated for LFM, $P = 4$ and $P = 5$, and L varied from 1 to 5.	103

A.7	PSLs of the LS MMF response for each straddled subshift of the received waveform. Generated for PD, $P = 2$ and $P = 3$, and L varied from 1 to 5.	104
A.8	PSLs of the LS MMF response for each straddled subshift of the received waveform. Generated for PD, $P = 4$ and $P = 5$, and L varied from 1 to 5.	105
A.9	LS MMF responses with Doppler averaging for LFM and PD with Doppler shifts from $+0.05$ to -0.05 radians. Sidelobes increase as J increases. Generated for $J = 3$ and $J = 4$, and averaging linearly from -0.04 to $+0.04$ radians.	106
A.10	Doppler averaged LS MMF responses for LFM with Doppler shifts of 0 to -0.05 radians. Generated for $J = 2, 3$, and 4 and linearly averaged from 0 to -0.04 radians.	107
A.11	Doppler averaged LS MMF responses for PD with Doppler shifts of 0 to -0.05 radians. Generated for $J = 2, 3$, and 4 and linearly averaged from 0 to -0.04 radians.	108
A.12	Matched filter, LS MMF, and APC responses for LFM and PD with two and three adaptive stages and $K = 1$. APC is able to discern the smaller target while the matched filter and LS MMF are unable to do so.	109
A.13	Matched filter, LS MMF, and APC responses for LFM and PD with two and three adaptive stages and $K = 3$. The LS MMF is able to discern the smaller target, and APC is able to as well, except in the LFM, $K = 3$, two adaptive stages scenario. . .	110
A.14	Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with two and three adaptive stages and $z = 1$ row zeroed. LS MMF is able to discern the smaller target; APC is able to discern the target with three adaptive stages, but with two stages, it is less able to do so.	111
A.15	Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with two and three adaptive stages and $z = 3$ rows zeroed. APC and the LS MMF can discern the target in all scenarios.	112
A.16	Pulse compressed responses of the measured data for LFM and PD with 50% eclipsing.	113

List of Tables

2.1	MML and PSL of the matched filter and LS MMF responses for both waveforms with $K = 1$ (resulting in a coded, not FM, waveform), $M = 1$, and $BT = 64$. The PSLs of the LS MMF responses are worse than those of the matched filter.	19
2.2	MML, sidelobe level, PSL, resolution (relative to the matched filter), and condition number of matrix $\mathbf{A}^H \mathbf{A} + \delta \mathbf{I}$ for the matched filter and diagonally loaded LS MMFs for LFM. As diagonal loading value δ increases, MML and condition number decrease while mainlobe width, sidelobe levels and PSLs increase.	24
2.3	MML, sidelobe level, PSL, resolution (relative to the matched filter), and condition number of matrix $\mathbf{A}^H \mathbf{A} + \delta \mathbf{I}$ for the matched filter and each diagonally loaded LS MMF response for PD. As diagonal loading value δ increases, MML and condition number decrease while mainlobe width and sidelobe levels increase. Unlike what was observed for LFM, PSLs decrease as loading values increase.	25
2.4	MML, sidelobe level, PSL, resolution (relative to the matched filter), and condition number of $\tilde{\mathbf{A}}^H \tilde{\mathbf{A}}$ for the matched filter and LS MMFs for LFM. As the number of zeroed rows in $\tilde{\mathbf{A}}$ increases, sidelobe levels (and PSL) decrease. The resolution degrades and condition number increases with more beamspoiling. For LFM, zeroing 7 rows above and below the m^{th} row of $\tilde{\mathbf{A}}$ yields the LS MMF response with the least amount of MML.	28

2.5	MML, sidelobe level, PSL, resolution (relative to the matched filter), and condition number of $\tilde{\mathbf{A}}^H \tilde{\mathbf{A}}$ for the matched filter and LS MMFs for PD. As the number of zeroed rows in $\tilde{\mathbf{A}}$ increases, sidelobe levels (and PSL) decrease. The resolution degrades and condition number increases with more beamspoiling. For LFM, zeroing 9 rows above and below the m^{th} row of $\tilde{\mathbf{A}}$ yields the LS MMF response with the least amount of MML.	29
2.6	MML and PSL of matched filter and LS MMF responses, with and without straddling. Straddling has a relatively small effect on the matched filter responses, but it causes a significant increase in PSL for the LS MMF responses. Neither filter experiences a significant increase in MML.	33
2.7	Worst-case and average PSLs and MMLs of the $L = 1$ to $L = 5$ filter averaged LS MMF response for LFM and PD. The averaged PSL and MML values correspond to the average PSL and MML across all 40 straddled LS MMF responses.	36
2.8	Worst-case and average PSLs and MMLs of the $P = 1$ to $P = 5$ filter averaged LS MMF response for LFM and PD. The averaged PSL and MML values correspond to the average PSL and MML across all 40 straddled LS MMF responses.	38
2.9	Worst-case PSL with L varied from 1 to 5 and P varied from 1 to 5 for LFM. The numbers on the axis corresponds to values of L and P used to construct the LS MMF which corresponds to the PSL value at the intersection of the L^{th} column and P^{th} row. PSL improves as the product LP increases and the best PSL occurs when L and P equal five.	42
2.10	Worst-case PSL with L varied from 1 to 5 and P varied from 1 to 5 for PD. The numbers on the axis corresponds to values of L and P used to construct the LS MMF which corresponds to the PSL value at the intersection of the L^{th} column and P^{th} row. PSL improves as the product LP increases and the best PSL occurs when L and P equal five.	43

2.11	MML and PSL of the matched filter and LS MMF responses for both waveforms, with and without Doppler. The LFM matched filter and LS MMF responses and the PD matched filter response are relatively unaffected by the Doppler shift. The LS MMF corresponding to PD response exhibits roughly 7 dB in PSL degradation due to Doppler.	44
2.12	The ω_{max} values used to generate the Doppler averaged LS MMFs.	47
2.13	The ω_{max} values used to generate the Doppler averaged LS MMFs when relative direction of travel is known.	50
A.1	Worst-case and averaged PSL and MML of the filter averaged LS MMF responses for both waveforms. The averaged PSL and MML correspond to the average value of the PSL and MML across all 40 straddled LS MMF responses. Generated using $K = 5$, $M = 4$ and $BT = 64$, with $P = 2$ and L varied from 1 to 5.	100
A.2	Worst-case and averaged PSL and MML of the filter averaged LS MMF responses for both waveforms. The averaged PSL and MML correspond to the average value of the PSL and MML across all 40 straddled LS MMF responses. Generated using $K = 5$, $M = 4$ and $BT = 64$, with $P = 3$ and L varied from 1 to 5.	101
A.3	Worst-case and averaged PSL and MML of the filter averaged LS MMF responses for both waveforms. The averaged PSL and MML correspond to the average value of the PSL and MML across all 40 straddled LS MMF responses. Generated using $K = 5$, $M = 4$ and $BT = 64$, with $P = 4$ and L varied from 1 to 5.	102
A.4	Worst-case and averaged PSL and MML of the filter averaged LS MMF responses for both waveforms. The averaged PSL and MML correspond to the average value of the PSL and MML across all 40 straddled LS MMF responses. Generated using $K = 5$, $M = 4$ and $BT = 64$, with $P = 5$ and L varied from 1 to 5.	103

Chapter 1

Introduction and Background

With the advent of new technologies, use of the radio frequency (RF) spectrum space is becoming more and more limited. The RF spectrum is becoming more crowded, and spectrum allocation is shifting towards favoring commercial applications over radar systems [3, 4]. As a result, the need for waveforms that are well contained spectrally and have good time domain (auto-correlation) properties has warranted tremendous interest from the radar community [5, 6]. This work provides the derivation and analysis of pulse compression algorithms that can be applied to spectrally efficient waveforms and that yield significant improvements in the pulse compressed time domain properties.

Before delving into the details of this work, background information is needed to establish a baseline of both basic knowledge relevant to its understanding and a brief synopsis of related work; both are provided in this chapter. In Chapter 2 the application of the least-squares mismatch filter to FM waveforms is derived and simulations corresponding to these results are analyzed. Chapter 3 provides analysis of the adaptive pulse compression algorithm when it is applied to FM waveforms. This chapter also analyzes adjustments that have been made to adaptive pulse compression to improve its performance. Chapter 4 then analyzes the performance of the least-squares mismatch filter and adaptive pulse compression on measured data. In the conclusion, Chapter 5, the success of the least-squares mismatch filter and adaptive pulse compression when applied to

FM waveforms is discussed, and potential sources of future work are presented. This thesis summarizes and extends [7], which was presented at the 2015 IEEE International Radar Conference in Arlington, VA.

1.1 Radar Signal Processing Basics

The two aspects of a radar system that are generally controlled via radar signal processing are the waveform and filters. The radar waveform refers to both the mathematical representation and the electromagnetic wave that is emitted by a radar system. Physically emitted radar waveforms can either be pulsed or continuous wave. For this work, only pulsed radar waveforms are considered. A pulsed waveform is implemented on a radar system by limiting the time in which the transmitter is "on". The radar system then "listens" for the reflections of the waveforms. When the radar waveform is an unmodulated pulse, it has a bandwidth (B) that is inversely proportional to the pulse width (T), but the received signal power is proportional to the pulse duration. Bandwidth is a measure of the frequency spectrum occupied by a signal. It is usually measured at the 3 dB (half power) point in the power spectral density (PSD), which is the spectral content squared. The bandwidth of the signal is therefore inversely proportional to the range resolution of the radar system, which determines the ability of the radar to distinguish targets radially (in range). Signal strength and bandwidth oppose each other, since a longer pulse is needed to get enough energy on the target for signal reception, but shorter pulses are better for target discrimination (Ch. 1 of [8]).

Filtering is needed to better detect returned radar signals. Frequency specific filters, such as low pass or band pass filters, are used to reject unwanted signals and noise from frequency bands outside of the one the pulsed radar waveform occupies. For the pulsed waveform described in the previous paragraph, a matched filter is used to process the returned signals. The matched filter is designed by taking the time-reversed complex conjugated version of the transmitted waveform. It is then convolved with the received signal reflections to generate a range response. In practice, a Fourier transform is used instead of convolution, since it is computationally more efficient to

multiply in the frequency domain and then take the inverse Fourier transform than it is to convolve two signals in the time domain. Additional waveform-specific filters can also be designed for noise and unwanted signal rejection within the frequency range of interest, a topic which is discussed later.

The range response generated using the matched filter can be judged based on several criteria. These criteria, namely signal-to-noise ratio (SNR), range resolution, and peak/integrated sidelobe level (PSL/ISL), are determined by mathematically examining the range response [9]. SNR is the difference in dB between the signal power of a target relative to that of the noise floor. This value is the metric that ultimately determines whether a target can be detected. If the SNR is too small, the target may not be distinguishable from the noise. Range resolution is a measure of how closely together in range two targets can be such that they can still be distinguished from each other in the range response. This metric is measured by determining the width of the auto-correlation (convolution of the waveform with itself) mainlobe. It may also be referred to as the 3 dB resolution, because the mainlobe width is measured 3 dB down from its peak. The last performance metric involves the sidelobes. Sidelobes correspond to the non-target related peaks in an auto-correlation or range response plot. Two metrics, PSL and ISL, are related to the sidelobes. The peak sidelobe level (PSL) is measured by taking the ratio of the largest sidelobe to the mainlobe peak and is expressed mathematically in

$$\text{PSL}(\tau) = \max \left| \frac{\chi(\tau)}{\chi(0)} \right| \text{ for } \tau \in [\tau_m, T], \quad (1.1)$$

where τ_m correspond to the first delay outside of the mainlobe of the filter response, χ is the filter response, T is the pulse width, and 0 refers to the mainlobe peak. The peak sidelobe value is usually given in dB using $20 * \log_{10}(\text{PSL}(\tau))$. The integrated sidelobe level (ISL) is a measure of the entire sidelobe profile; it is the ratio of the integrated total of the filter response outside of the mainlobe to the integrated total of the filter response inside the mainlobe. It is calculated as

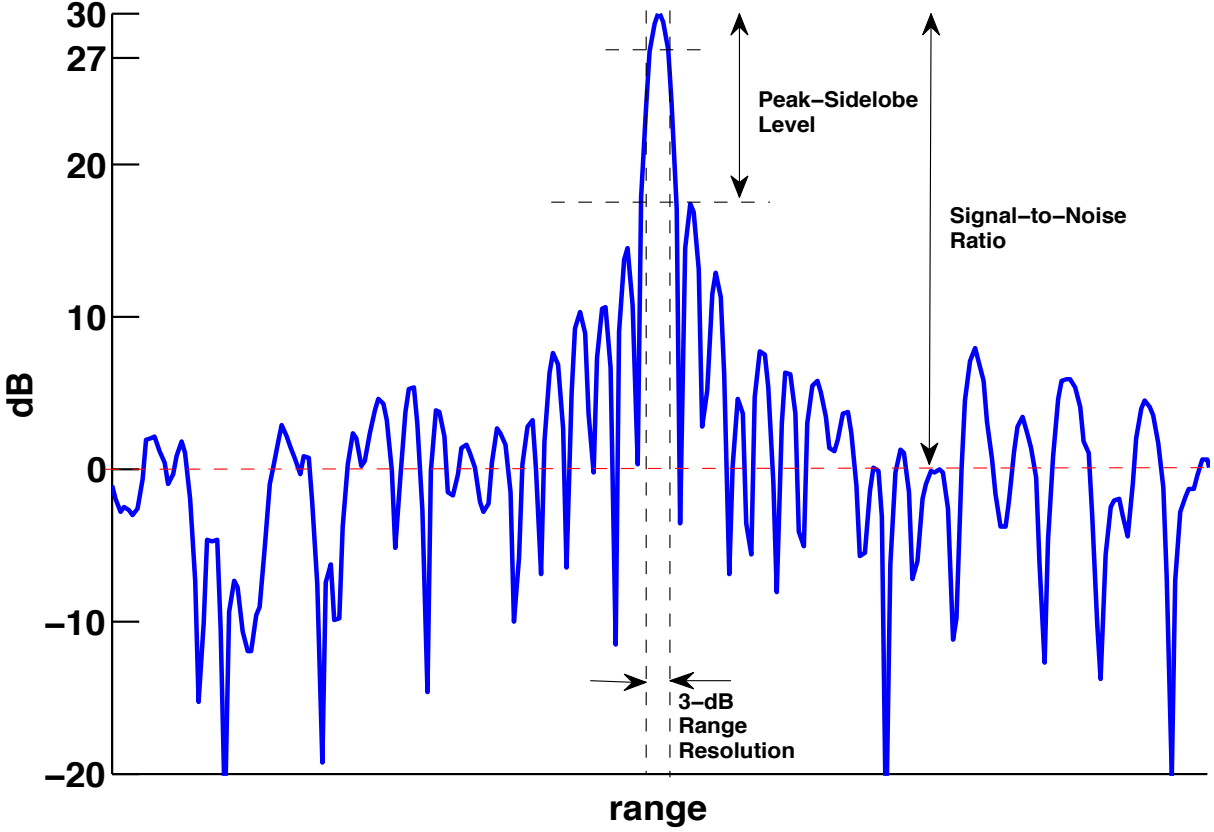


Figure 1.1: Visual representation of range response PSL, 3 dB range resolution, and SNR.

follows:

$$ISL(\tau) = \frac{\int_{\tau_m}^T |\chi(\tau)| d\tau}{\int_0^{\tau_m} |\chi(\tau)| d\tau} \quad (1.2)$$

and is usually expressed in dB as $20 * \log_{10}(ISL(\tau))$. Figure 1.1 provides a visual example of SNR, range resolution, and PSL. ISL is not shown in the figure, but to calculate it, the area outside of the vertical dashed lines would be integrated.

There are several aspects of radar that require attention from a practical or implementation point of view, including fidelity, robustness, and computational cost. Fidelity is a measure of how accurately a physical signal matches its idealistic mathematical model. It is important that signal processing algorithms yield high fidelity signals so that accuracy is maximized. Robustness is another concern because signal processing algorithms that work in a variety of scenarios, and

not just one specific, ideal scenario, are more practical for real-world applications. Robustness is needed so that non-ideal effects do not completely degrade performance of the radar system. The last practical concern is computational cost, that is, the time it takes for signal processing to take place when it is deployed in real time. The amount of time allotted for signal processing varies based on application. For example, compare automotive collision radar to weather radar; collision avoidance radar needs to be able to process signals in a matter of tenths of a second, while weather radar can take tens of seconds. Signal processing techniques that are ideal for one scenario may not be applicable to another situation due to processing time or performance constraints. More processing time often allows for better response characteristics.

1.2 Radar Pulse Compression

Up to this point, only pulsed radars with unmodulated pulses have been discussed. Radar pulse compression modulates the transmitted pulse. The unmodulated pulsed waveform has a pulse width T and bandwidth B , which are inverses of each other such that the time-bandwidth product $BT \approx 1$. By modulating the pulsed waveform, the bandwidth of the signal increases such that $BT \gg 1$, which helps suppress sidelobe levels and increases the capabilities of signal processing (Ch. 20 of [8]). Without pulse compression, in order to get a desired range resolution (proportional to bandwidth), the only way to improve the returned signal strength is to increase transmitted power. With pulse compression, the transmitted power can be set to achieve the desired return signal strength while the pulsed waveform is modulated to get the desired bandwidth. The matched filter, as described previously, is commonly used for pulse compression, as it yields the response with the maximum SNR because the filter is matched to the waveform. Figure 1.2 shows the non-pulse compressed matched filter response for an unmodulated pulse and the pulse compressed matched filter response for a modulated pulse.

The time-bandwidth product (BT) of the pulse compressed waveform is known as the pulse compression ratio or pulse compression gain. This value refers to the improvement in range resolu-

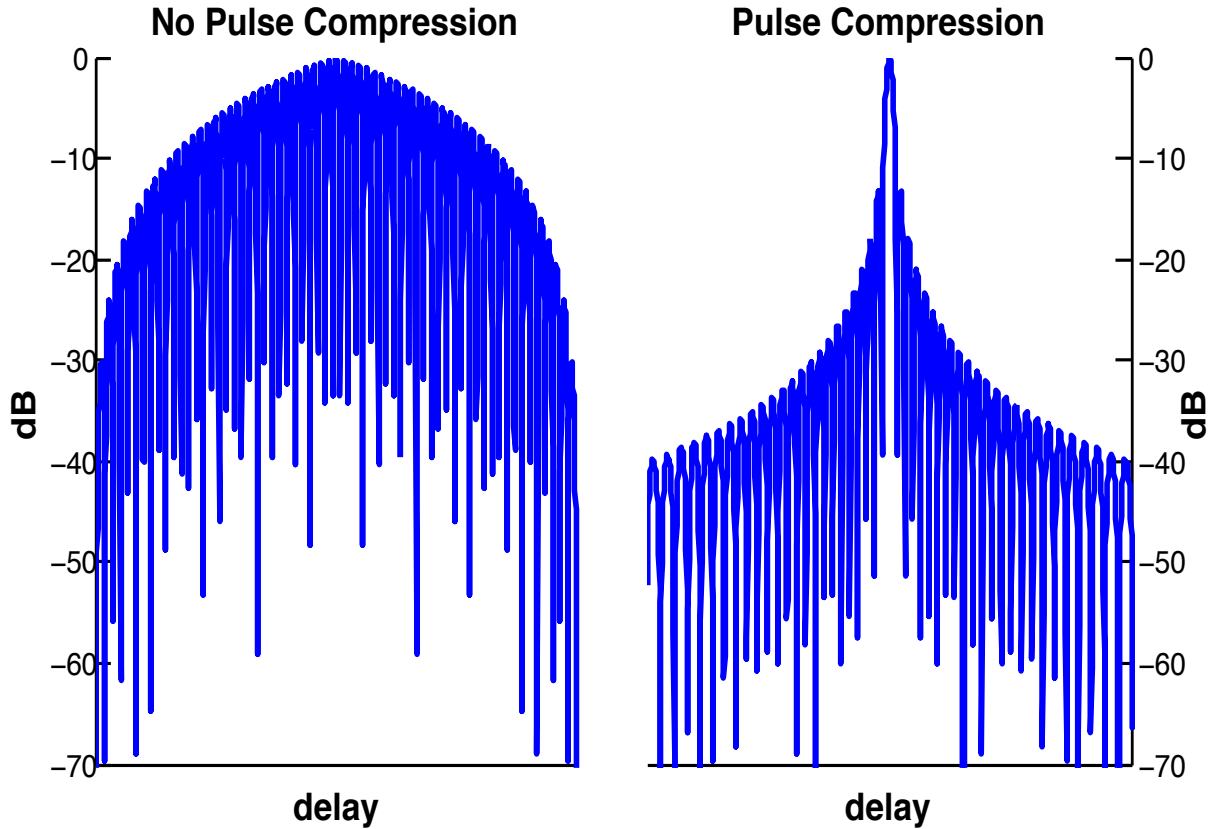


Figure 1.2: Matched filter responses for unmodulated and modulated pulses. Range resolution and PSL improve when modulated pulses are used.

tion that the modulated pulse has relative to that of the unmodulated pulse. The pulse compression gain refers to the improvement in SNR that is realized in the pulse compressed response relative to the non-pulse compressed response (Ch. 20 of [8]).

While the matched filter is easily implemented, results in no SNR loss, and has good range resolution, it has range sidelobes that, depending on the waveform, can be relatively high. High range sidelobes are undesirable for target detection because they can mask other targets with SNRs below the sidelobes. To reduce sidelobe levels, amplitude tapers or windows are commonly used. Amplitude tapering/windowing applies a weighting to all sample values, with a higher weight being applied to the mainlobe and lower weights to the sidelobes [9]. This multiplication in the time domain is equivalent to a convolution of the windowing spectrum and the range response spectrum, which causes spectral spreading. Spectral spreading results in a widening of the mainlobe

(degraded resolution) and SNR loss in the range response.

1.3 Radar Waveforms

The performance of the un-windowed matched filter response or auto-correlation plot is dependent on the waveform that is used, as the matched filter is simply the time-reversed complex conjugated version of the waveform. Thus, pulsed compressed waveforms can be designed to have desirable auto-correlation properties. The first pulsed compressed waveform to be discussed is the linear frequency modulated (LFM) waveform. This waveform is commonly used because it is simple to implement. It was given the name LFM because the frequency of the waveform changes linearly with respect to time. It is known to have auto-correlation peak sidelobes at -13.7 dB, and its auto-correlation plot can be seen in Figure 1.3. It is also commonly used because, when compared to other frequency modulated (FM) waveforms, LFM has the best range resolution. The spectrum of the LFM is, for $BT \gg 1$, essentially rectangular in that it rolls off sharply after its 3 dB bandwidth (Ch. 20 of [8]). This rolloff is desirable because it allows for the waveform to occupy the maximum bandwidth with minimal spectral energy intruding into neighboring frequencies. Other FM waveforms, which do not roll off as quickly, may need to have a bandwidth that is less than the LFM so as not to have an unacceptable amount of energy in neighboring frequency bands.

Another class of radar waveforms is coded waveforms. Unlike FM waveforms, coded waveforms modulate the phase of the transmitted waveform rather than the frequency. The most basic form of a coded waveform is a binary code, which is based off of the binary idea of '1's and '0's. Instead of transmitting a '1' or '0', the phase of the transmission is toggled between 0 and π . Binary codes are generally presented as a sequence of phase values that is N values long, where each value in the sequence is known as a chip. An example of a binary code is a Barker code, which has auto-correlation sidelobes equal to 1 and has a mainlobe peak of N (Ch. 20 of [8]). When expressed in dB, the PSL of a Barker code is $-20 * \log_{10}(N)$ and is the minimum possible PSL for a code of length N . Binary codes' sudden changes in phase, which are discontinuities, result in

significant spectral spreading, making them difficult to implement because spectral containment is a priority for radar waveforms. Additionally, if any filters are used to bandlimit the signal, some of the spectral content may be removed or skewed. The resulting signal would then be different than the signal used to construct the matched filter, known as mismatching. Mismatching can result in SNR loss and degradation in resolution (widening of the mainlobe).

When a chip can assume any phase value (0 to 2π) the resulting waveform is known as a polyphase code. By allowing the phase to assume any value, lower sidelobe levels can be achieved without the same degree of spectral spreading seen with the binary code. Spectral spreading is generally less for polyphase codes because the transition between phase values is usually smaller than π . By designing polyphase codes with small chip-to-chip transitions, spectral spreading resulting from discontinuities can be significantly decreased (Ch. 20 of [8]). Polyphase coded waveforms can also have lower sidelobe levels because there are more degrees of freedom in terms of phase values, which allow the codes to be optimized. Polyphase code optimization is a topic of great interest in the radar signal processing community, as there does not appear to be a floor for sidelobe levels and polyphase codes can be implemented with only a slight amount of spectral spreading if optimized correctly [10–13]. The downside to these optimized polyphase codes is the difficulty involved in designing them, since there are an infinite number of phase values and N can theoretically be infinity. They can also be difficult to implement in a physical system, depending on the degree of precision required for the phase values.

The LFM described previously is in a class of radar waveforms known as frequency modulated (FM) waveforms. These waveforms are continuous and have no discontinuities like the sudden transitions in phase seen in polyphase coded waveforms. Because of the lack of discontinuities, the spectrum of FM waveforms is relatively well contained and can be implemented in a real system with high fidelity. Additionally, the lack of discontinuities means that the transmitter is always transmitting at full power. For coded waveforms, the discontinuities are typically accompanied by a period of time where the transmitted power is reduced. This phenomenon can be understood if the unit circle is considered. A coded waveform has chip values that correspond to points on

the unit circle. Being located on the unit circle corresponds to the full transmitted power. When phase transitions between chips occur, the transition can be represented by drawing a straight line between the two points on the circle. Any point inside the unit circle corresponds to a transmitted power less than the maximum. If the transition line goes through the origin, no power will be transmitted, which could have detrimental effects on the radar system. FM waveforms are also Doppler tolerant, meaning the filter response is not degraded significantly when observed targets are moving relative to the radar. Due to these characteristics, FM waveforms are highly desirable for use in radar systems.

Non-linear FM (NLFM) waveforms are continuous FM waveforms whose frequency changes non-linearly with respect to time. NLFM waveforms, if constructed correctly, can have lower auto-correlation sidelobes than the LFM for the same BT . Because of the improved performance of the NLFM and the growing need for spectrally efficient waveforms, a lot of work is currently being done on physically realizable optimal NLFM waveforms [1, 14–26]. Whereas the spectrum of the LFM has very quick rolloff outside of its 3 dB bandwidth, NLFM waveforms have slightly worse spectral rolloffs for the same BT . However, with this slight degradation in spectral containment, the auto-correlation of NLFM waveforms generally has lower sidelobe levels [1]. An example of NLFM is shown in Figure 1.4. An LFM, with the same BT as the NLFM, is shown in Figure 1.3. The relationships between waveforms' spectrum shapes and time-sidelobe levels can be easily observed, as well as the increased auto-correlation mainlobe width seen for NLFM. The NLFM responses shown in Figure 1.4 demonstrate a spectrum shape and auto-correlation that is similar to a tapered LFM matched filter response, only for the NLFM case, there is no SNR loss in the auto-correlation mainlobe.

Because polyphase codes can be optimized but are difficult to implement in a real system with high fidelity, the ability to convert a polyphase code into an FM waveform would eliminate the spectral spreading and distortion experienced when a polyphase coded waveform is implemented. Such a waveform is known as an optimized FM waveform and can be implemented using the continuous phase modulation (CPM) framework, derived from the communications community [27].




Figure 1.3: Taken from [1], characteristics of a linear FM waveform. Plot (a) shows the frequency as a function of time, (b) the spectrum of the LFM, and (c) its auto-correlation.

CPM is implemented using the transition value α (difference in phase values) between successive chip phases. Since the phases are located on the unit circle, there are two "paths" for each phase transition, clock-wise and counter-clockwise around the circle [27]. The value that should be used for α is the smallest value in magnitude. Thus, α is always between $-\pi$ and π . The output of the CPM framework is an NLFM waveform and as such is continuous and differentiable. Instead of having sharp phase transitions which "cut through" the unit circle, the CPM waveform has a continuous phase transition that stays on the unit circle and results in a transmission that is always transmitting with maximum power output, thereby avoiding spectral spreading observed for the polyphase code. When a polyphase code is implemented via the CPM framework, the resulting NLFM waveform is referred to as a polyphase-coded FM (PCFM) waveform [27].

The optimization of coded waveforms allows for the generation of optimized PCFM wave-




Figure 1.4: Taken from [1], characteristics of a non-linear FM waveform. Plot (a) shows the frequency as a function of time, (b) the spectrum of the NLFM, and (c) its auto-correlation.

forms through the use of the CPM framework [28]. Polyphase codes can be optimized based on several criteria, such as ISL, PSL, or other desired sidelobe parameters or characteristics. The resulting optimized FM waveforms have similar sidelobe behavior as the polyphase codes, but are physically realizable in a real system, and they are spectrally efficient. The number of chips in the polyphase code used to construct the PCFM waveform is roughly equivalent to the BT of the PCFM waveform. This process is discussed in great detail in [27, 28]. For this work, two PCFM waveforms, implemented with CPM, are used. The first of these is the LFM, which is described in detail previously. This waveform is used because its pulse compressed responses can be used as a benchmark for the second waveform, which is an optimized FM waveform denoted as 'performance diversity' and is described in detail in [28]. Both of these waveforms are implemented with a polyphase code with a $BT = 64$.

1.4 The Least-Squares Mismatch Filter

The only pulse compression method examined thus far has been the matched filter. With the matched filter, it has been shown that windowing can be used to suppress range sidelobes, but at the cost of mainlobe SNR loss and degraded range resolution. Another pulse compression method, known as mismatch filtering, can be used to generate a desired range response. (Note: it is not referred to as an auto-correlation because the matched filter is no longer used.) As the name implies, the mismatch filter utilizes mismatching in order to reduce range sidelobe level [29], but, as with tapering, SNR loss, now referred to as mismatch loss (MML), and resolution degradation can occur. MML is expressed mathematically in

$$\text{MML}(\chi) = \left| \frac{\chi_{\text{M}}(0)}{\chi(0)} \right|, \quad (1.3)$$

where $\chi_{\text{M}}(0)$ is the matched filter's mainlobe peak for the waveform being considered, and $\chi(0)$ is the mainlobe peak of the mismatch filter that is designed using the same waveform. This work considers a specific type of mismatch filter known as the least-squares mismatch filter (LS MMF).

The LS MMF, developed by Ackroyd and Ghani in 1973 [29], was initially derived for polyphase radar codes. The idea behind the LS MMF is that a small penalty in the mainlobe signal-to-noise ratio (SNR) can be incurred to allow range sidelobe levels to be further suppressed. Furthermore, the LS MMF, which is two to four times longer than the matched filter, has essentially the same real-time computational cost as the matched filter.

The filter derived in [29] is calculated using a matrix \mathbf{A} , given by

$$\mathbf{A} = \begin{bmatrix} s_1 & 0 & \cdots & 0 \\ \vdots & s_1 & & \vdots \\ s_N & \vdots & \ddots & 0 \\ 0 & s_N & & s_1 \\ \vdots & & \ddots & \vdots \\ 0 & \cdots & 0 & s_N \end{bmatrix}. \quad (1.4)$$

The matrix \mathbf{A} has a Toeplitz structure that is composed of delay shifted versions of the code values. The dimensions of \mathbf{A} are $(M+1)N-1 \times MN$, where MN is the length of the mismatch filter. The scalar N is the code length, and M is a scalar value typically set at 2, 3, or 4. The superscript $(\cdot)^H$ corresponds to the Hermitian operator, or complex conjugate transpose. The filter equation can then be written as

$$\mathbf{h} = (\mathbf{A}^H \mathbf{A} + \delta \mathbf{I})^{-1} \mathbf{A}^H \mathbf{e}_m, \quad (1.5)$$

where \mathbf{A} is the Toeplitz matrix defined in (1.4), the elementary vector \mathbf{e}_m contains all zeros except for the m^{th} (middle) element, which is a '1', δ is a diagonal loading factor, and \mathbf{I} is an $MN \times MN$ identity matrix.

The LS MMF, as discussed [29], has one significant pitfall: mismatch loss (MML), which is the SNR loss seen in the mainlobe of the LS MMF response relative to the auto-correlation of the waveform. MML can be partially mitigated through use of the loading factor δ . By increasing δ , mismatch loss can be reduced, but some of the improvement in sidelobe levels is lost.

Here, the LS MMF is modified to enable application to arbitrary FM waveforms and to enhance robustness to range straddling and Doppler effects. This new formulation is evaluated using simulated and measured data.

1.5 Adaptive Pulse Compression

Adaptive Pulse Compression (APC), as derived in [30–32] (with experimental results in [33]), is an iterative pulse compression algorithm initially derived for polyphase codes. In each iteration, APC implements a minimum mean-square error (MMSE) estimator that behaves similarly to that of adaptive beamforming, but in the range domain. Each of these adaptive filters has a form that is similar to that of the Minimum Variance Distortionless Response (MVDR) algorithm (Ch. 15 of [34]). The equations for the adaptive filter coefficients are given by

$$\mathbf{w}(n) = \frac{(\mathbf{C}(n) + \mathbf{R})^{-1} \mathbf{s}}{\mathbf{s}^H (\mathbf{C}(n) + \mathbf{R})^{-1} \mathbf{s}} \quad (1.6)$$

and

$$\mathbf{C}(n) = \sum_{\tau=-N+1}^{N-1} \rho(n + \tau) \mathbf{s}_\tau \mathbf{s}_\tau^H, \quad (1.7)$$

where $\mathbf{w}(n)$ is the adaptive filter coefficients for range cell n , \mathbf{R} is the noise covariance matrix, \mathbf{s} is the vector of polyphase code values, $\mathbf{C}(n)$ is the structured signal correlation matrix, and ρ is the estimated power in the matched filter response for range cell n . Unlike the matched filter and LS MMF, these filter coefficients, as the name implies, adaptively change to estimate the value of each range cell. The APC algorithm is initialized with the matched filter response, and APC uses this initial baseline to adaptively generate filter coefficients to reduce sidelobe levels.

Due to the ability of APC to significantly reduce sidelobe levels, it is a topic that has been adapted for many different applications, such as multistatic radar [35–42], Doppler compensation [43–45], single pulse imaging [44, 46], and radar pulse compression repair [2, 47–49]. The work on multistatic APC is also important because it allows spectrum sharing, which is another way to combat the problem of a crowded spectrum, as this pulse compression algorithm could coexist with other waveforms operating at the same frequencies.

The APC algorithm has recently been modified to permit application to arbitrary FM waveforms. Here this new version of APC is evaluated using simulated and measure data to ascertain

its robustness to practical pulse compression effects of range straddling and Doppler.

Chapter 2

Least Squares Mismatched Filtering

To the best of our knowledge, the LS MMF has not previously been successfully applied to FM waveforms. We are not claiming we are the first to do so, but, given a sizable literature review, and having published one paper on the topic already, we have yet to learn of anyone who has done it before. Furthermore, the use of the LS MMF is a departure from stretch processing, which is typically used for the LFM waveform (Ch. 20 of [8]). In regards to current digital signal processing endeavors, FM waveforms that are developed in response to the growing need for spectral efficiency and low autocorrelation sidelobes can be used in conjunction with the LS MMF to achieve even greater sidelobe suppression and/or super resolution [50].

2.1 LS MMF Modification for Application to FM Waveforms

In order for the LS MMF to be applied to an FM waveform (which is continuous and differentiable, unlike a code), adjustments to the least squares structure need to be made. Implementing the LS MMF with FM waveforms requires that \mathbf{A} now be composed of delay shifted version of the discretized waveform s instead of the polyphase code values. It is also necessary to "oversample" the waveform, that is, sample the waveform beyond the 3 dB bandwidth. When a polyphase code is used for the LS MMF, the code is sampled at N , the code length. Oversampling is not used when the LS MMF is implemented for a polyphase code because the phase transitions of the coded

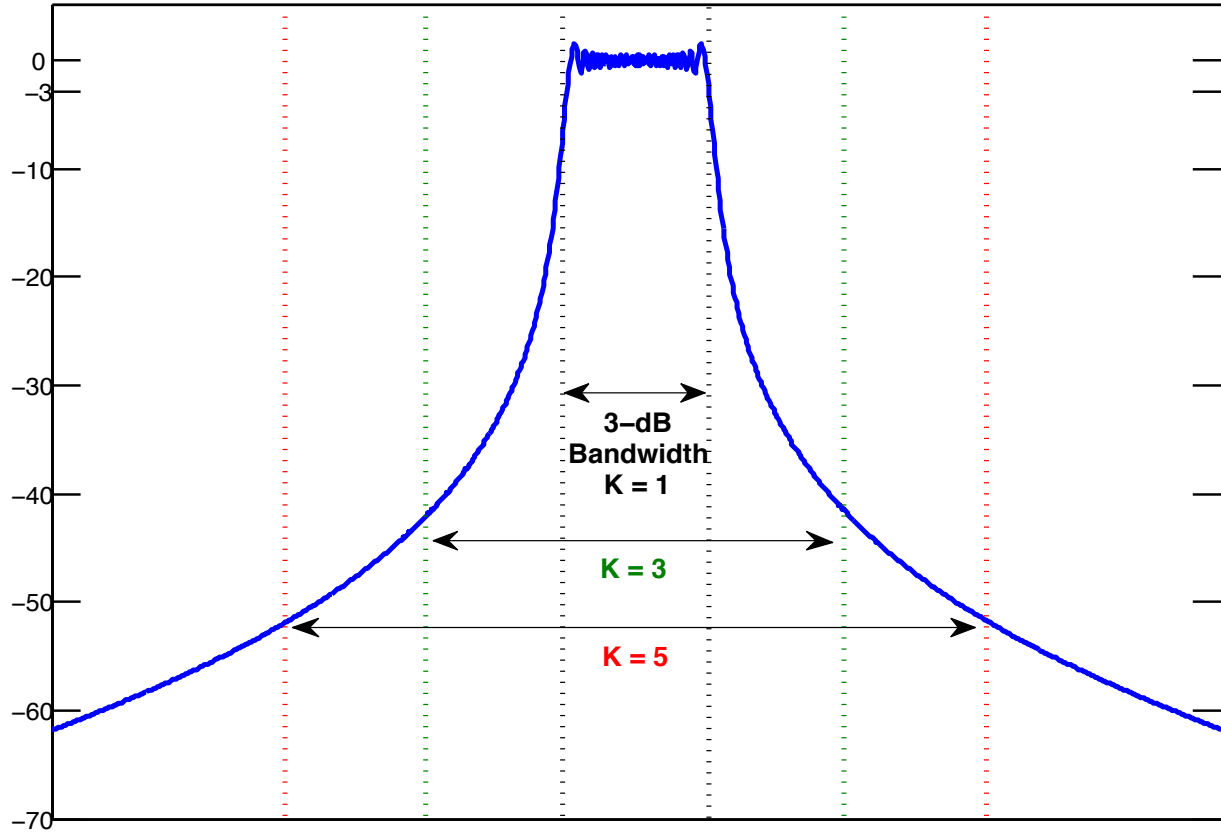


Figure 2.1: The spectrum of an LFM with three oversampling rates: $K = 1$, corresponding to the 3 dB bandwidth $K=3$, and $K=5$.

signal are abrupt and the phase remains constant over the entire chip interval. Therefore, having multiple samples per chip interval would be redundant, as each sample would correspond to the same phase. The frequency of the FM waveform, however, is a continuous function of time. Thus, to ensure a higher fidelity model of the waveform, it needs to be sampled at a rate higher than the 3 dB bandwidth. The oversampling factor is denoted by the scalar K , and $K = 1$ corresponds to the case where no oversampling occurs. When no oversampling occurs, the waveform is sampled at the 3 dB bandwidth, which is approximately equivalent to the time-bandwidth product BT of the waveform. The waveform resulting from $K = 1$ will be a code, not a discretized FM waveform. Figure 2.1 provides a pictorial example of oversampling, showing how more of the spectrum is sampled as K increases.

When using the LS MMF with an oversampled waveform, matrix $\mathbf{A}^H \mathbf{A}$ may be susceptible

to ill-conditioning issues [50]. Ill-conditioning occurs when a matrix that is to be inverted has a large condition number. The condition number is equal to the maximum eigenvalue of the matrix divided by the minimum eigenvalue of the matrix. If the condition number is too large, the resultant inverted matrix may experience undesirable noise enhancement, which in this scenario degrades the performance of the LS MMF response. The effects of ill-conditioning can be alleviated by using a non-zero loading factor in (1.5).

This work will consider two FM waveforms in its analysis: linear FM (LFM) and an optimized FM waveform previously denoted as 'performance diversity' (PD) based on how it was designed in [28]. Both waveforms are implemented as the polyphase coded FM (PCFM), as described in Section 1.3, and have a time-bandwidth (BT) product of approximately $64 (N)$. Figure 2.2 shows the matched filter and LS MMF responses for LFM and PD with an oversampling factor of $K = 1$. With an oversampling factor of 1, the LFM matched filter response is that of a P3 code (Ch. 20 of [8]) rather than the expected LFM response (-13.7 dB peak sidelobe level). In both plots, the matched filter outperforms the LS MMF in terms of mismatch loss (MML) and peak sidelobe level (PSL), as described in (1.3) and (1.1), respectively. The LS MMF has lower sidelobe levels near the mainlobe compared to the matched filter, but near the ends of the filter responses, the LS MMF realizes a significant increase in sidelobe levels, which is why it has a worse PSL than the matched filter. The MML and PSL values from the matched filter and LS MMF responses for both waveforms can be found in Table 2.1. The LS MMF is used because it is known to have significant improvement in sensitivity (sidelobe levels) than the matched filter [29]. This improvement is not realized in these results, however, which demonstrates the need for adjustments to the LS MMF structure so that it can be applied to FM waveforms.

The more oversampled the waveform, the lower the sidelobes of the LS MMF response, since oversampling results in higher fidelity. Pulsed waveforms, which are used in this analysis, have a theoretically infinite bandwidth, since they are time limited. Though still undersampled relative to Nyquist sampling, increasing the oversampling factor reduces error resulting from aliasing, which causes degradation due to mismatch [27]. However, these lower sidelobes come at a cost

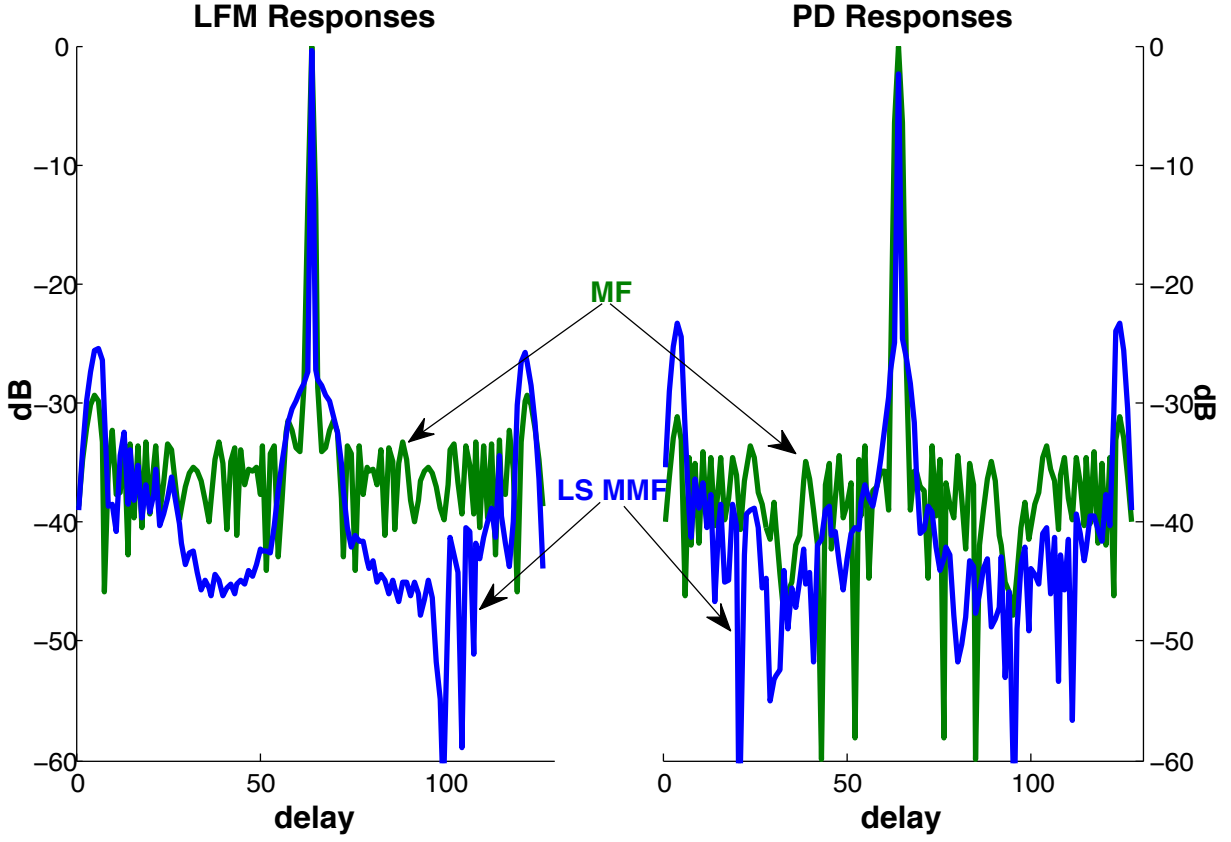


Figure 2.2: Matched filter and LS MMF responses for LFM and PD with $K = 1$, $M = 1$, and $BT = 64$. The matched filter outperforms the LS MMF in terms of MML and PSL, while the LS MMF has lower sidelobe levels near the mainlobe, with an increase near the ends of the filter response.

Table 2.1: MML and PSL of the matched filter and LS MMF responses for both waveforms with $K = 1$ (resulting in a coded, not FM, waveform), $M = 1$, and $BT = 64$. The PSLs of the LS MMF responses are worse than those of the matched filter.

	<i>LFM</i>		<i>PD</i>	
	MML (dB)	PSL (dB)	MML (dB)	PSL (dB)
Matched	0	-29.27	0	-31.13
LS MMF	0.47	-25.43	2.36	-23.24

of increased mainlobe SNR loss (mismatch loss). This loss occurs because the LS MMF tries to suppress all delays not corresponding to the match point. The more oversampled the waveform is, the more similar the delay shifted columns in matrix \mathbf{A} become. Thus, the LS MMF suppresses sidelobes closer to the mainlobe, eventually leading to "self suppression", which causes mismatch loss and super-resolution [50].

This pattern is clearly demonstrated in Figure 2.3, where the oversampling factor is varied from 2 to 7. The peak sidelobes of these responses occur near the ends of the filter responses, as seen in Figure 2.2. The improvement in sidelobe level diminishes as K increases. Since little improvement was observed beyond the value of $K = 5$, this value will be the largest oversampling factor used in the following simulations. The total number of samples in the discretized version of the waveform is then equal to KN ($\approx KBT$). For LFM at $K = 5$, the MML of the LS MMF response is 12.2 dB and has a PSL of -34.2 dB. The LS MMF response for PD yields an MML of 14.6 dB and a PSL of -36.97 dB.

The filter length factor, M , is varied from 1 to 4 in Figure 2.4 to show that the improvements in sidelobe level can be realized by making the filter longer (increasing the degrees of freedom). As with increasing the oversampling factor, increasing the filter length leads to a higher MML from $M = 1$ to $M = 2$. However, unlike what was observed when K was varied, the MML remains relatively constant as M increases from 2 to 4. The "floor" of the sidelobes appears to be lower for even values of M than it is for odd values (see $M = 2$ compared to $M = 3$). The spurious peaks present in these four LS MMF responses are undesirable, as they could lead to false target detection. The improvement in sidelobe level diminishes as M gets larger, as was observed with K . Since a value of $M = 4$ yields the best sidelobe levels, this value will be used from this point forward.

While the PSLs improve monotonically as both K and M increase (on the ranges that were simulated), the MML also increases monotonically. Fortunately, as described in the following section, there are ways to reduce significantly the amount of MML present in the LS MMF response, such that in some cases it is nearly zero. In the case of $K = 5$ and $M = 4$, which will be the case for all

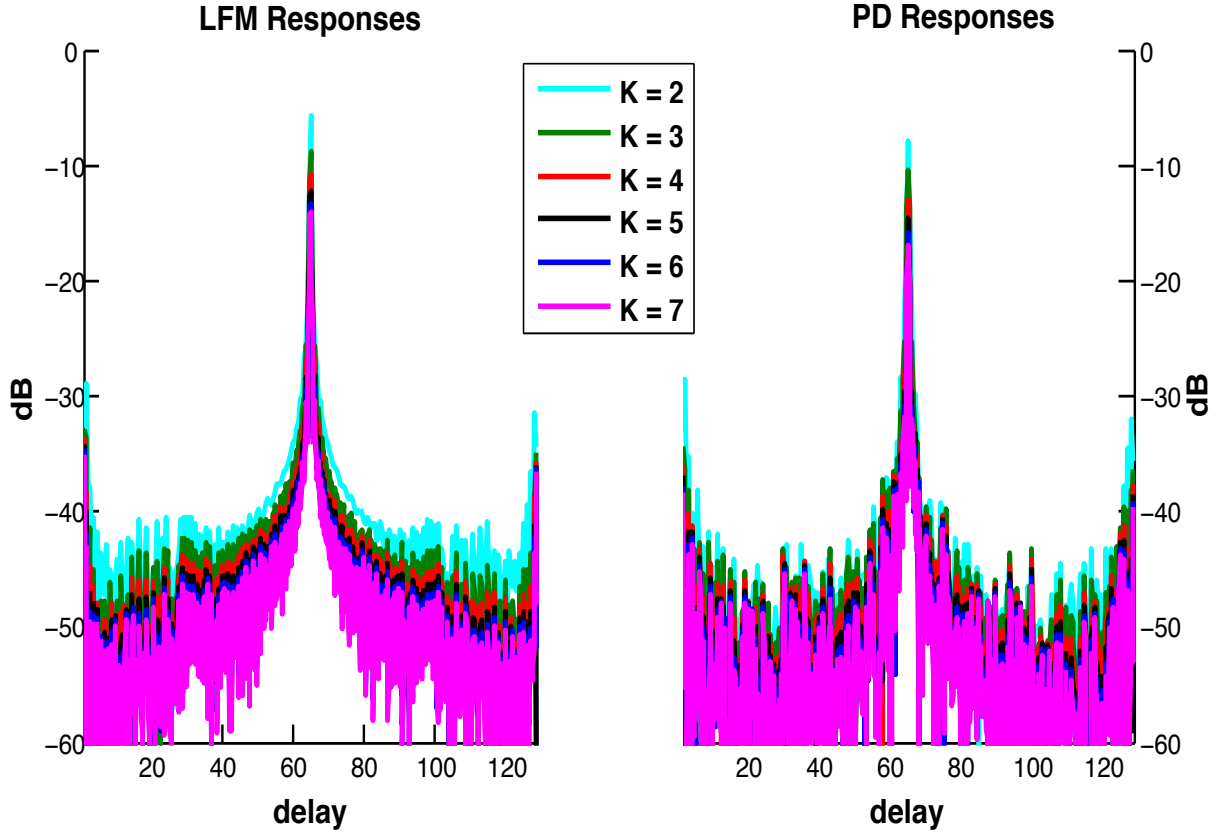


Figure 2.3: LS MMF responses for LFM and PD with oversampling factors $K = 2$ through $K = 7$, $M = 1$, and $BT = 64$. As K increases, sidelobe levels improve (decrease) monotonically, while MML increases monotonically. Although there are more values in the LS MMF responses as K increases, the pulse width remains constant for all LS MMFs.

following simulations (unless otherwise stated), the LS MMF corresponding to the LFM response has an MML of 24.65 dB and a PSL of -19.3 dB. The LS MMF corresponding to the PD response has an MML of 24.67 dB and a PSL of -19.4 dB. These MML and PSL values are unacceptable if our goal is to have enhanced sensitivity, although these LS MMF responses do have resolution that is roughly 10% that of their respective matched filters. For most applications, an acceptable MML is a fraction of a dB and desirable PSLs are less than -40 dB. In order to achieve acceptable MML and PSL values, some, if not all, of the improvement in resolution will need to be sacrificed.




Figure 2.4: LS MMF responses for LFM and PD with $K = 5$, filter length factors of $M = 1$ to $M = 4$, and $BT = 64$. As M increases, the PSLs decrease monotonically and the MML remains relatively constant for $M > 1$.

2.2 Idealistic Simulations and Results

In the LS MMF responses shown in Figure 2.4, the mainlobe resolution becomes finer while the mismatch loss increases. The initial ($K = 5$, $M = 4$) LS MMF response has a resolution that is roughly 10% that of the matched filter resolution. The LS MMF can also experience ill-conditioning issues, especially when oversampling is used. These undesirable results in the LS MMF are why δ is present in (1.5). The matrix $\delta \mathbf{I}$ is needed for diagonal loading of the matrix $\mathbf{A}^H \mathbf{A}$ in an attempt to mitigate these unwanted results. Diagonal loading adds the same value to every element on the diagonal of matrix $\mathbf{A}^H \mathbf{A}$. The value of the diagonal loading term δ varies based on application. Here it was varied from 0 to the largest eigenvalue λ_{MAX} of the matrix $\mathbf{A}^H \mathbf{A}$. The results of the diagonally loaded LS MMF responses are given in Figures 2.5 and 2.6 and in

Tables 2.2 and 2.3. The loading factor δ was set at 0, λ_{MAX} , and various scalings thereof including $\lambda_{\text{MAX}} - 30$ dB, $\lambda_{\text{MAX}} - 20$ dB, and $\lambda_{\text{MAX}} - 10$ dB. The maximum eigenvalue of $\mathbf{A}^H \mathbf{A}$ for LFM was 16.63 dB, and for PD λ_{MAX} was 18.86 dB.

With diagonal loading, ill-conditioning issues are eliminated because the condition number decreases. MML decreases as the diagonal loading value increases and the mainlobe width increases, causing coarser resolution. In Figure 2.5, which uses LFM, as the loading value increases, the sidelobe levels and PSLs also increase. However, in Figure 2.6, which uses PD, the MML improves at a quicker rate than the sidelobe level, which means that the PSLs actually improve as δ increases. As the loading value increases, the mainlobe width approaches that of the matched filter, and the overall shape of the LS MMF response approaches that of the matched filter. This result effectively makes using the LS MMF unnecessary when significant amounts of diagonal loading are used because the purpose of the LS MMF is to realize a filter response with either finer resolution or improved sensitivity (sidelobe levels) than that of the matched filter.

Diagonal loading, which reduces MML but degrades sidelobe suppression, is not sufficient for use as the primary means of reducing MML in the LS MMF response. That is not to say that some degree of diagonal loading should not be used, as it is an effective method of combating ill-conditioning, but for our purposes and for simplicity, we will henceforth assume that $\delta = 0$. The ineffectiveness of diagonal loading is the impetus for the second alteration to the LS MMF structure that is explored: beamspoiling.

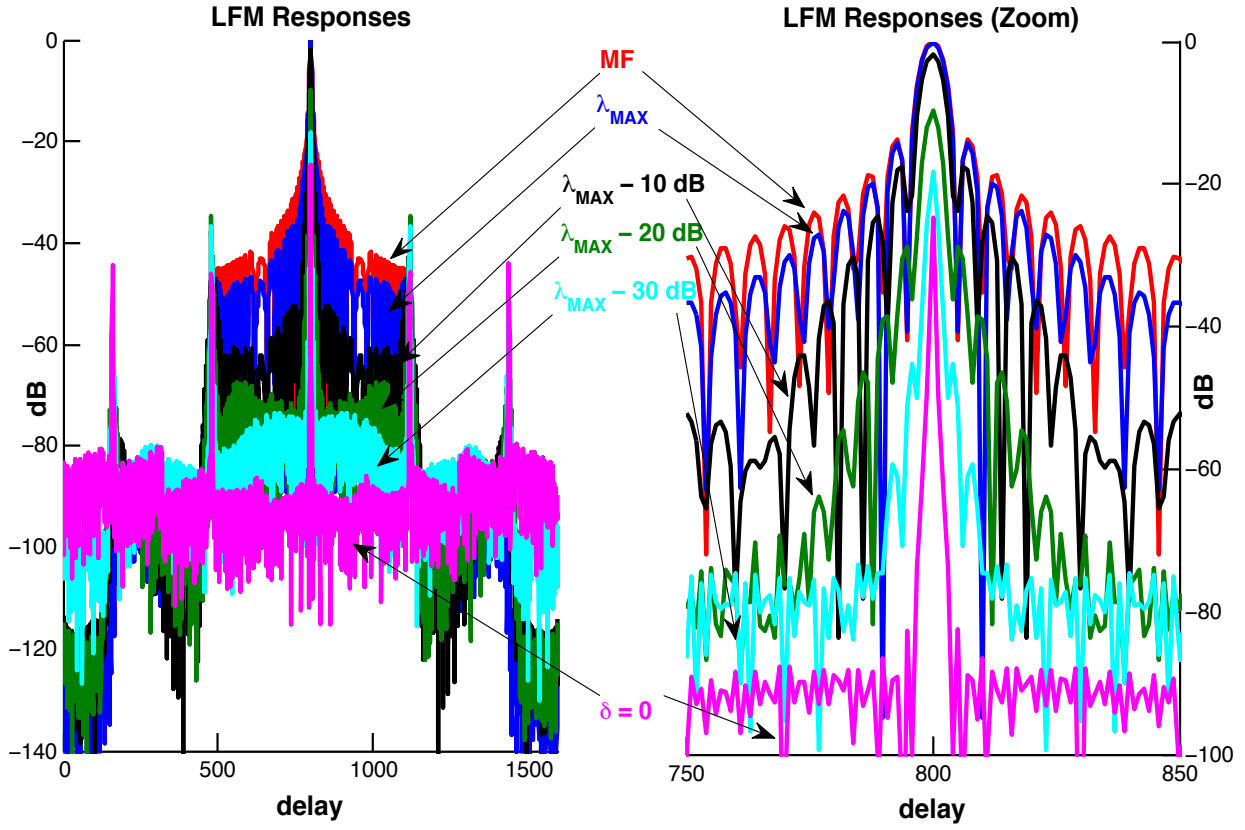


Figure 2.5: Matched filter and LS MMF responses for LFM for five diagonal loading values. These diagonal loading values (δ), in decreasing order, are 16.63 dB, 6.63 dB, -3.37 dB, -13.37 dB, and 0. The responses on the right are the zoomed-in mainlobes of the responses on the left. As δ increases, MML decreases while mainlobe width and sidelobe levels increase.

Table 2.2: MML, sidelobe level, PSL, resolution (relative to the matched filter), and condition number of matrix $\mathbf{A}^H \mathbf{A} + \delta \mathbf{I}$ for the matched filter and diagonally loaded LS MMFs for LFM. As diagonal loading value δ increases, MML and condition number decrease while mainlobe width, sidelobe levels and PSLs increase.

δ	MML (dB)	SL (dB)	PSL (dB)	Resolution	Condition Number
Matched	0	-13.77	-13.77	100%	N/A
λ_{MAX}	0.07	-14.29	-14.22	94.5%	2.00
$\lambda_{MAX} - 10$ dB	1.77	-17.60	-15.83	80.8%	11.0
$\lambda_{MAX} - 20$ dB	9.74	-28.87	-19.13	52.5%	101
$\lambda_{MAX} - 30$ dB	18.37	-36.53	-18.16	26.9%	956
0	24.65	-43.94	-19.29	12.3%	2.11×10^4

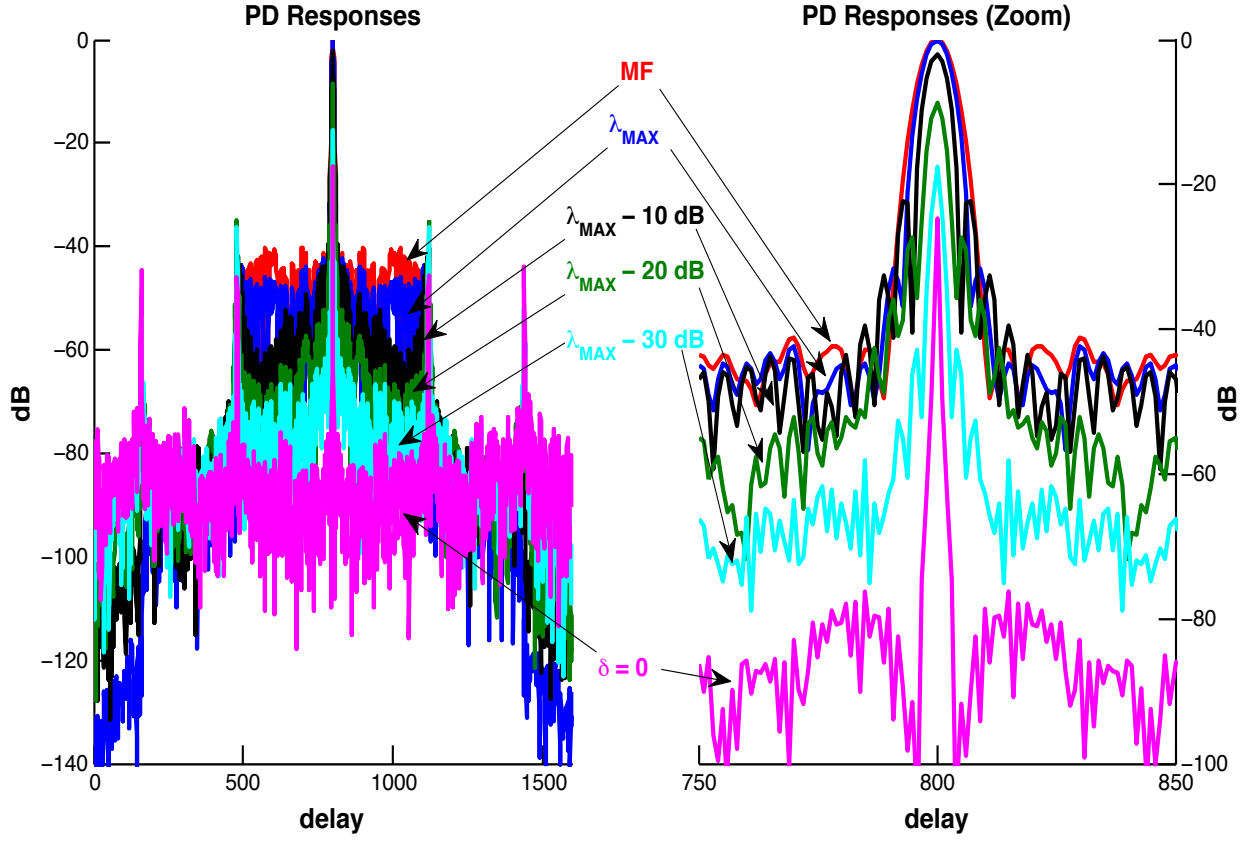


Figure 2.6: Matched filter and LS MMF responses for PD for five diagonal loading values. These diagonal loading values (δ), in decreasing order, are 18.86 dB, 8.86 dB, -1.14 dB, -11.14 dB, and 0. The responses on the right are the zoomed-in mainlobes of the responses on the left. As δ increases, MML decreases while mainlobe width and sidelobe levels increase.

Table 2.3: MML, sidelobe level, PSL, resolution (relative to the matched filter), and condition number of matrix $\mathbf{A}^H \mathbf{A} + \delta \mathbf{I}$ for the matched filter and each diagonally loaded LS MMF response for PD. As diagonal loading value δ increases, MML and condition number decrease while mainlobe width and sidelobe levels increase. Unlike what was observed for LFM, PSLs decrease as loading values increase.

δ	MML (dB)	SL (dB)	PSL (dB)	Resolution	Condition Number
Matched	0	-40.22	-40.22	100%	N/A
λ_{MAX}	0.15	-42.32	-42.17	88.6%	2.00
$\lambda_{MAX} - 10$ dB	2.01	-40.17	-38.16	67.3%	11.0
$\lambda_{MAX} - 20$ dB	8.62	-35.04	-26.42	43.8%	101
$\lambda_{MAX} - 30$ dB	17.56	-36.12	-18.56	21.7%	969
0	24.67	-44.02	-19.35	9.6%	3.05×10^4

Beam spoiling is a method commonly used in beam forming of phased arrays to suppress the energy in the sidelobes by allowing the array's mainlobe width to increase. This idea can be applied here because we want to degrade the resolution of the LS MMF response (analogous to the antenna mainlobe width), in hopes of suppressing sidelobes and countering the deleterious effects of super-resolution [27]. This process is implemented by zeroing out rows above and below the m^{th} row in matrix \mathbf{A} , yielding

$$\tilde{\mathbf{A}} = \begin{bmatrix} s_1 & 0 & \cdots & \cdots & 0 & 0 \\ \vdots & s_1 & \ddots & \ddots & \vdots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & \cdots & 0 & 0 \\ 0 & 0 & \cdots & \cdots & 0 & 0 \\ \hline s_n & \vdots & \ddots & \ddots & s_1 & 0 \\ \hline 0 & 0 & \cdots & \cdots & 0 & 0 \\ 0 & 0 & \cdots & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & s_N & \vdots \\ 0 & 0 & \cdots & \cdots & 0 & s_N \end{bmatrix}, \quad (2.1)$$

where the horizontal lines delineate the m^{th} row. The LS MMF equation (1.5) can now be written as

$$\mathbf{h} = (\tilde{\mathbf{A}}^H \tilde{\mathbf{A}} + \delta \mathbf{I})^{-1} \tilde{\mathbf{A}}^H \mathbf{e}_m. \quad (2.2)$$

where $\tilde{\mathbf{A}}$ has replaced \mathbf{A} .

As the number of zeroed rows increases, the MML and sidelobe levels of the LS MMF decrease, and the mainlobe widens. This trend holds true up to a certain number of zeroed rows, which varies based on the LS MMF parameters and the particular waveform. After this point, the mismatch loss starts to degrade, though sidelobe levels still decrease and resolution continues to become coarser. Since MML tends to be a high priority as it translates directly into the SNR of

a received signal, the LS MMF is optimized such that it minimizes MML. The optimal LS MMF is determined by zeroing out different numbers of rows on both sides of the m^{th} row of $\tilde{\mathbf{A}}$, and finding which number of zeroed rows corresponds to the smallest amount of MML. Though, as shown in Figures 2.7 and 2.8 and in Tables 2.4 and 2.5, the PSL improves as more and more rows are zeroed out, MML increases and the condition number becomes larger, which may make the $\tilde{\mathbf{A}}^H \tilde{\mathbf{A}}$ matrix ill-conditioned. Note that regardless of waveform, the number of zeroed rows that minimizes MML results in a mainlobe width that is just slightly wider than that of the respective waveform's matched filter response.

The LS MMF responses for LFM with beamspoilage are shown in Figure 2.7, and the responses for PD are shown in Figure 2.8. For both waveforms, as the number of zeroed rows increases, sidelobe levels (and PSL) decrease. Additionally, the mainlobe width increases (causing coarser resolution) and eventually surpasses the mainlobe width of the matched filter response (see Tables 2.4 and 2.5). MML is the one parameter that does not exhibit monotonic behavior. Initially, as more rows are zeroed, MML decreases. After a certain number of rows are zeroed, the MML then starts to increase. The number of rows that need to be zeroed to minimize MML varies based on the waveform that is used and the oversampling factor K . All of the LS MMF responses exhibit spurious peaks, which could lead to false target detection. As more beamspoilage is implemented, the height of the spurious peaks relative to the average sidelobe levels decreases.

For LFM, the minimum MML occurs when 7 rows are zeroed above and below the m^{th} row of matrix $\tilde{\mathbf{A}}$. Table 2.4 shows that with 7 zeroed rows, the 3-dB resolution is 14% wider than that of the matched filter, but the LS MMF has a PSL that is nearly 53 dB better than that of the matched filter. For PD, the minimum MML occurs when 9 rows are zeroed. Table 2.5 shows that with 9 zeroed rows, the 3-dB resolution is 6.4% wider than that of the matched filter, but the LS MMF has a PSL that is nearly 37 dB better than that of the matched filter.

It is also interesting to consider the combination of diagonal loading and beamspoilage. Beamspoilage, as shown previously, enables tremendous improvement in PSL at the cost of increased MML and possible ill-conditioning. By implementing a small degree of diagonal loading along




Figure 2.7: LS MMF responses for LFM with beamspoiling. The number of rows that were zeroed above and below the m^{th} row of $\tilde{\mathbf{A}}$, in decreasing order, are 13, 10, 7, 5, 2, and 0. The responses on the right are the zoomed-in mainlobes of the responses on the left. As more rows are zeroed, the mainlobe width increases and surpasses the matched filter mainlobe width. The minimum MML occurs when 7 rows are zeroed.

Table 2.4: MML, sidelobe level, PSL, resolution (relative to the matched filter), and condition number of $\tilde{\mathbf{A}}^H \tilde{\mathbf{A}}$ for the matched filter and LS MMFs for LFM. As the number of zeroed rows in $\tilde{\mathbf{A}}$ increases, sidelobe levels (and PSL) decrease. The resolution degrades and condition number increases with more beamspoiling. For LFM, zeroing 7 rows above and below the m^{th} row of $\tilde{\mathbf{A}}$ yields the LS MMF response with the least amount of MML.

Zeroed Rows	MML (dB)	SL (dB)	PSL (dB)	Resolution	Condition Number
Matched	0	-13.77	-13.77	100%	N/A
0	24.65	-43.94	-19.29	12.3%	2.11×10^4
2	14.03	-50.58	-36.55	75.3%	4.16×10^4
5	2.49	-56.82	-54.33	94.5%	9.72×10^4
7	0.73	-67.18	-66.45	114.6%	2.24×10^5
10	1.24	-84.35	-83.11	142%	2.24×10^6
13	2.03	-96.89	-94.86	172.1%	2.89×10^7

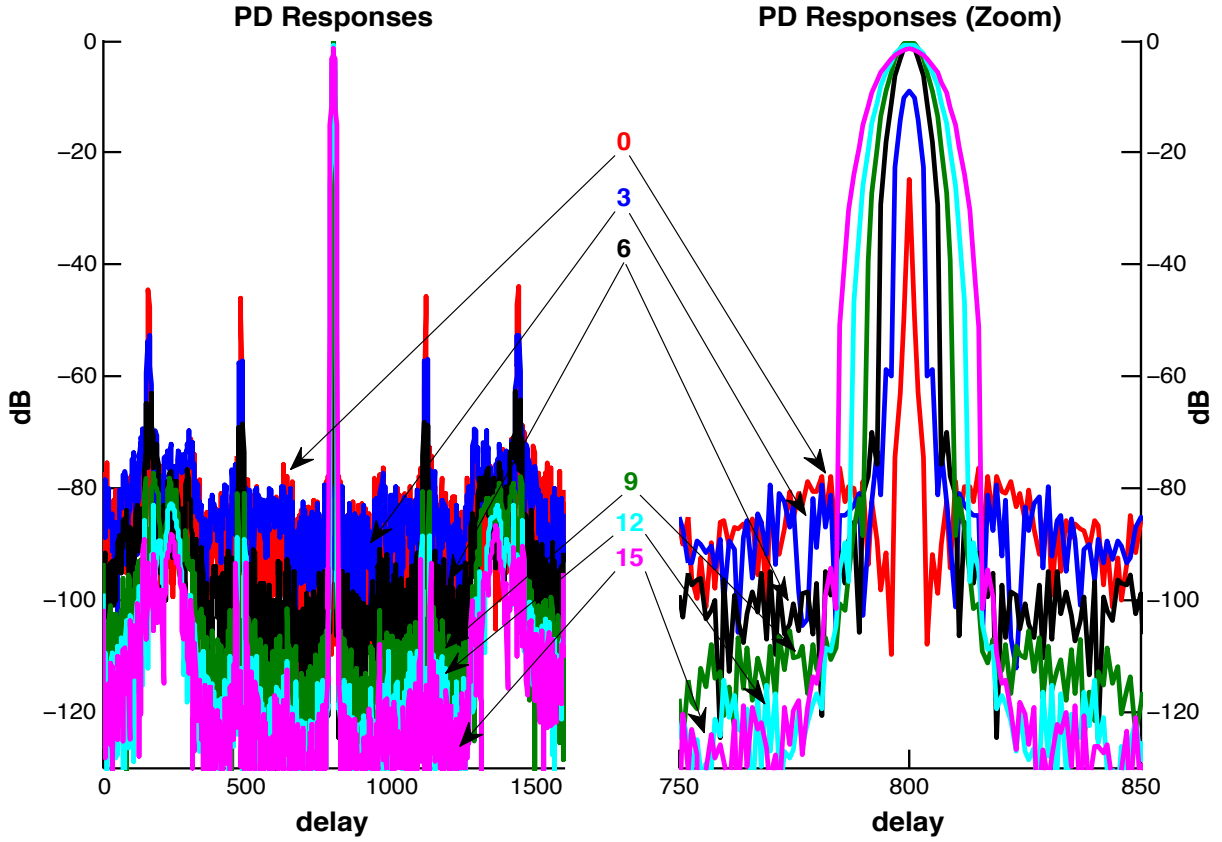


Figure 2.8: LS MMF responses for PD with beamspoiling. The number of rows zeroed above and below the m^{th} row of $\tilde{\mathbf{A}}$, in decreasing order, are 15, 12, 9, 6, 3, and 0. The responses on the right are the zoomed-in mainlobes of the responses on the left. As more rows are zeroed, the mainlobe width increases and surpasses the matched filter mainlobe width. The minimum MML occurs when 9 rows are zeroed.

Table 2.5: MML, sidelobe level, PSL, resolution (relative to the matched filter), and condition number of $\tilde{\mathbf{A}}^H \tilde{\mathbf{A}}$ for the matched filter and LS MMFs for PD. As the number of zeroed rows in $\tilde{\mathbf{A}}$ increases, sidelobe levels (and PSL) decrease. The resolution degrades and condition number increases with more beamspoiling. For LFM, zeroing 9 rows above and below the m^{th} row of $\tilde{\mathbf{A}}$ yields the LS MMF response with the least amount of MML.

Zeroed Rows	MML (dB)	SL (dB)	PSL (dB)	Resolution	Condition Number
Matched	0	-40.22	-40.22	100%	N/A
0	24.67	-44.02	-19.35	9.6%	3.05×10^4
3	9.07	-52.57	-43.50	47.3%	6.06×10^4
6	0.93	-62.63	-61.70	80.1%	1.89×10^5
9	0.09	-77.35	-77.27	106.4%	1.02×10^6
12	0.67	-82.63	-81.96	142.0%	3.22×10^6
15	1.49	-86.81	-85.32	184.7%	7.50×10^6

with beamspoiling, the degradation in MML and ill-conditioning may be slightly mitigated without suffering too much of an increase in PSL. Fortunately, in our simulations, the time-bandwidth products, oversampling rates, filter length factors, and number of zeroed rows did not result in $\tilde{\mathbf{A}}^H \tilde{\mathbf{A}}$ being ill-conditioned. The number of zeroed rows that minimizes MML (7 for LFM, 9 for PD) will be used for all forthcoming simulations involving the LS MMF, unless otherwise noted. Figures 2.7 and 2.8 display the "ideal" filter responses (no straddling or Doppler effects) for the LS MMF and matched filter. Now non-ideal effects that result from implementing these filters and waveforms will be examined.

2.3 Straddling Effects

The first non-ideal effect to be examined is known as range straddling, which occurs when the sampled version of the waveform obtained at the receiver (either in simulation or in a real system) possesses a delay offset relative to the sampled model of the waveform that is used in filter construction (Ch. 18 of [8]). This effect is always present in a real system because it is impossible to perfectly sample the received signal commensurate to the model used for filter construction. Range straddling occurs as a delay between 0 and T_S , where T_S is the sampling period, which corresponds to the chip interval for a coded waveform ($K = 1$). In order to design a filter with some degree of robustness, the "worst-case" scenario needs to be improved. For straddling effects, it is assumed that the worst-case occurs when a delay of $0.5 T_S$ occurs in the received signal. This assumption is made because samples that differ by half of the sampling interval should be the most different from the samples used in filter construction, furthest away from the desired samples with respect to time.

Figure 2.9 provides a pictorial example of worst-case straddling. In this figure, black dots represent the discretized model of the waveform used for filter construction; in ideal sampling, these same samples (black dots) would be used to model the received waveform. Red dots represent the worst-case straddled samples. Sampling at this straddling offset would theoretically lead to




Figure 2.9: A pictorial example of straddling, using a sample LFM. The black dots at integer multiples of T_S represent the discretized waveform model for filter construction. The red dots, shifted $0.5 T_S$ in delay, represent the "worst-case" straddled samples.

the most model mismatch, resulting in the largest degradation in performance in the LS MMF response. While such a small difference may seem negligible, especially when K is sufficiently large, it should be noted that achieving very low sidelobe levels (less than -60 dB PSL) requires rather extreme fidelity.

Implementing straddling in simulation, the worst-case LS MMF responses are generated for the two previously used waveforms (LFM and PD). The worst-case matched filter responses are also generated. These responses are shown in Figure 2.10; the significant increase in sidelobe levels are immediately observed for the LS MMF response. This increase in sidelobe levels takes place because with relatively low initial sidelobe levels, the LS MMF is a highly sensitive filter. Due to this higher degree of sensitivity, any model mismatch results in a significant penalty in sensitivity/sidelobe levels. The matched filter, on the other hand, is relatively insensitive, and realizes little performance degradation due to straddling, though some MML is observed.

Specifically, as shown in Table 2.6, the LS MMF response for LFM experiences an increase of




Figure 2.10: Matched filter and LS MMF responses for LFM and PD, for both ideal and "worst-case" straddled sampling. The matched filter shows minimal degradation in performance when straddling is present. The LS MMF shows a significant increase in PSL with straddling compared to the ideal sampled case.

nearly 14 dB in PSL, while the LS MMF response for PD, which had lower unstraddled sidelobe levels than LFM, suffers roughly 22 dB in PSL degradation. The increase in sidelobes resulting from straddling does not occur as spurious peaks, as seen in the idealistic case. Rather, the straddled LFM sidelobes look like rounded peaks, while the PD sidelobes are plateau-like. For both LS MMF responses, the increase in sidelobes occurs near the mainlobe. As a result, the spurious peaks seen in the idealistic LS MMF responses have been masked by the sidelobes from straddling.

Table 2.6: MML and PSL of matched filter and LS MMF responses, with and without straddling. Straddling has a relatively small effect on the matched filter responses, but it causes a significant increase in PSL for the LS MMF responses. Neither filter experiences a significant increase in MML.

	<i>LFM</i>		<i>PD</i>	
	MML (dB)	PSL (dB)	MML (dB)	PSL (dB)
Matched	0	−13.77	0	−40.22
Straddled Matched	0.15	−13.65	0.09	−38.94
LS MMF	0.73	−67.18	0.09	−77.35
Straddled LS MMF	0.85	−53.21	0.20	−55.00

2.4 Averaging Methods

Since straddling effects are a common occurrence in received radar signals, the mitigation of their effects is something of great interest and usefulness. Publications exist addressing the mitigation of mismatch loss resulting from straddling (e.g., [51]), but, as far as we have found, successfully mitigating the effects on sidelobe levels has not yet been explored. The approach presented here uses an averaging method that, in principle, enables more of the radar waveform to be represented by the LS MMF without any additional implementation cost, as all filter construction can be done offline.

The first of these averaging methods is denoted as "filter averaging". This method involves constructing L LS MMFs and averaging them together to get one filter, which has the same length as the initial LS MMF. The second of these averaging methods is denoted as "waveform averaging". This method further oversamples the waveform beyond K that is used for filter construction, and then averages P adjacent discrete samples to get an "averaged waveform". This waveform is then used to construct an LS MMF as described in (2.2). Lastly, these two averaging methods are combined to examine the effects of implementing waveform averaging along with filter averaging.

2.4.1 Filter Averaging

The filter averaging method requires that the discrete model of the waveform be further oversampled by a factor L such that the total sampling rate is equal to $LKBT$. The new variable L refers to the number of filters to be averaged together, while K and BT refer to the oversampling factor and time-bandwidth product, respectively. The new discretized waveform is then polyphase decomposed (Ch. 11 of [52]) into L waveforms, which correspond to every L th sample in the initial oversampled waveform. For example, if $L = 2$, the first, third, fifth, etc., discrete samples of the highest oversampled waveform would be used to construct one of the discretized waveforms, and then the second, fourth, sixth, etc., would correspond to the second. Once these waveforms are obtained, the LS MMF can be designed using (2.2). The resulting L LS MMFs are then averaged together to form one final "averaged filter" as

$$\bar{\mathbf{h}} = \frac{1}{L} \sum_{\ell=0}^{L-1} \mathbf{h}_{\ell}. \quad (2.3)$$

For the initial ($L = 1$) LS MMF, it was assumed that the worst-case straddling response occurred at $0.5 T_S$. However, now that sampling is done at a higher rate (for $L > 1$), it can no longer be assumed that the worst-case samples occur at $0.5 T_S$ because this delay may be used to construct the averaged filter in (2.3). To find the new worst-case straddling delay, 40 subshifts of the waveform are used to represent the received waveform, such that "sampled" versions correspond to straddling delays $n/40 T_S$ where n goes from 1 to 40. Each of these discretized waveforms is then convolved with the LS MMF and the peak sidelobe level (PSL) is calculated. The response with the highest PSL is deemed the worst-case.

Presented in Figure 2.11 and 2.12 and in Table 2.7 are the results for the LS MMF using the averaged filter method for $L = 1$ to $L = 5$ for both LFM and PD. Figure 2.11 corresponds to the results for LFM. Notice how the PSL improves as L increases, but with diminishing improvements. The mismatch loss (MML) is relatively constant (less than a 0.01 dB increase) as L increases, while an improvement of nearly 4 dB is seen between the worst-case PSL for $L = 1$ and $L = 5$.




Figure 2.11: LS MMF responses for $L = 1$ to $L = 5$ filter averaging for LFM with worst-case straddling. As L increases, the PSL improves monotonically but with diminishing improvements for successive values of L .

In Figure 2.12 the results for the LS MMF responses for PD are presented. These responses are similar to those observed in the LS MMF response for LFM, but with slightly better performance due to the use of a better waveform. Table 2.7 contains the worst-case and average PSLs and MMLs for the LS MMFs generated using LFM and PD, when filter averaging is used ($L = 1$ to $L = 5$). This table shows that worst-case and average PSL improve while worst-case and average MML worsen as L increases. Shown in Figures A.1 and A.2 in Appendix A are plots displaying the PSL and MML values of the LS MMFs corresponding to each of the 40 straddling shifted versions of the received waveform. As seen in Figure 2.10, the sidelobes that result from filter averaging for LFM are rounded peaks, while the sidelobes for PD are relatively flat.




Figure 2.12: LS MMF responses for $L = 1$ to $L = 5$ filter averaging for PD with worst-case straddling. As L increases, the PSL improves monotonically but with diminishing improvements for successive values of L .

Table 2.7: Worst-case and average PSLs and MMLs of the $L = 1$ to $L = 5$ filter averaged LS MMF response for LFM and PD. The averaged PSL and MML values correspond to the average PSL and MML across all 40 straddled LS MMF responses.

L	<i>LFM</i>				<i>PD</i>			
	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)
1	-47.6	0.86	-56.1	0.79	-50.1	0.18	-59.0	0.13
2	-49.9	0.86	-58.1	0.79	-52.4	0.18	-61.2	0.13
3	-50.7	0.86	-58.3	0.79	-53.2	0.18	-61.4	0.13
4	-51.2	0.86	-58.3	0.79	-53.6	0.18	-61.5	0.13
5	-51.5	0.86	-58.4	0.79	-53.9	0.18	-61.5	0.13

2.4.2 Waveform Averaging

The second averaging method that was explored is called waveform averaging. Similar to the filter averaging method, waveform averaging requires the initial waveform to be further oversampled, this time by a factor of P . Thus, the total sampling rate is now $PKBT$. This factor corresponds to the number of discrete samples that are averaged together to construct the averaged waveform. In the nominal, non-averaged case, $P = 1$. When P is larger than 1, P adjacent samples are averaged together. For example, if $P = 3$, the 1st, 2nd, and 3rd samples of the highest sampled waveform are averaged to generate the 1st sample of the averaged waveform. Then, the 4th, 5th, and 6th samples are averaged to generate the 2nd sample of the averaged waveform, and so on. This averaging process is described by

$$\bar{s}_n = \frac{1}{P} \sum_{p=0}^{P-1} s_{Pn+p}, \quad (2.4)$$

where n is the sample index of the final averaged waveform. This averaged waveform is then used to construct an LS MMF via (2.2).

The worst-case LS MMFs are determined using the method described previously in Section 2.4.1. These worst-case responses are presented in Figures 2.13 and 2.14 and in Table 2.8 using the waveform averaging method for $P = 1$ to $P = 5$, for LFM and PD. Again, as was seen for filter averaging, PSL improves as P increases, but with diminishing improvements. As before, the MML is relatively constant (less than a 0.01 dB increase) as P increases, while an improvement of nearly 4 dB is seen between the worst-case PSL for $P = 1$ and $P = 5$.

Table 2.8 contains the worst-case and average PSLs and MMLs for the LS MMFs generated using LFM and PD, when waveform averaging is used ($L = 1$ to $L = 5$). This table shows that the worst-case PSL and average PSL improve while worst-case and average MML worsen as P increases. Shown in Figures A.3 and A.4 in Appendix A are plots displaying the PSLs and MMLs of the LS MMFs corresponding to each of the 40 straddling shifted versions of the received waveform. The sidelobe shapes for LFM and PD are the same as was observed previously in Figures 2.11 and 2.12.

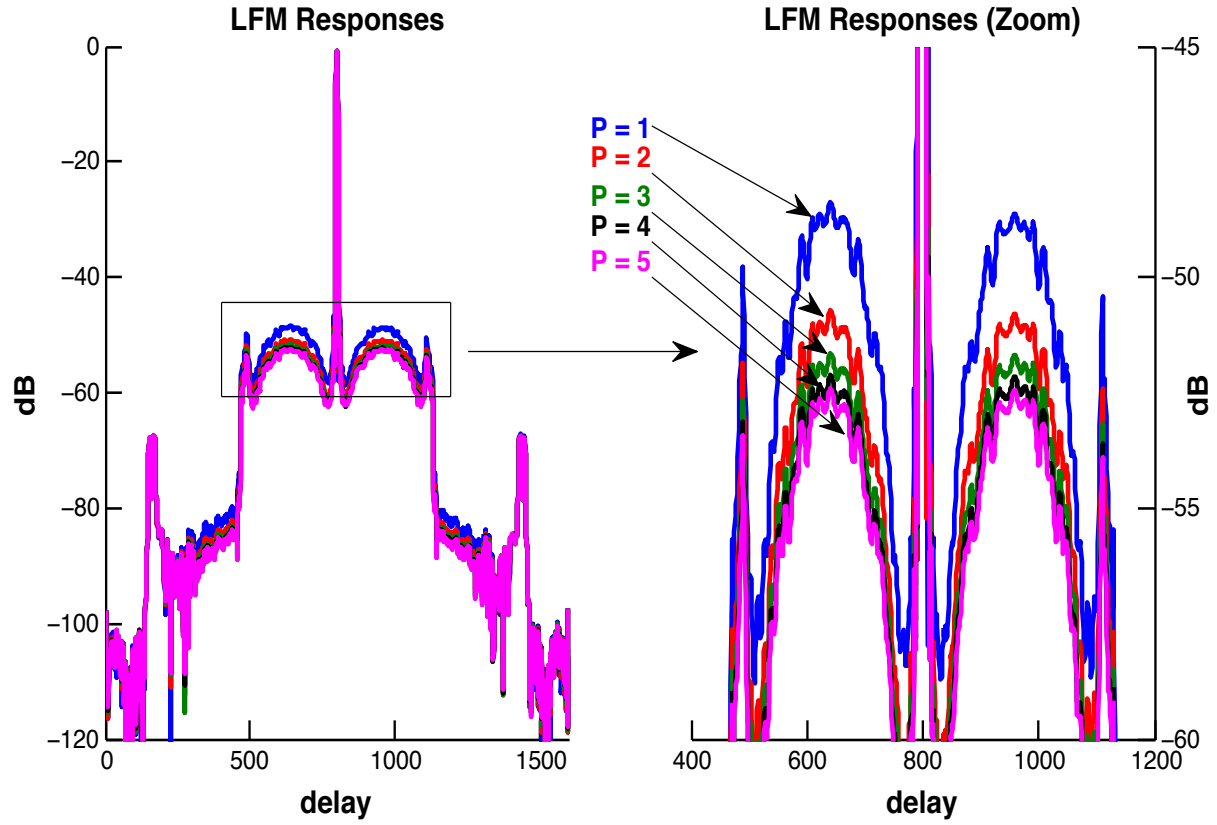


Figure 2.13: LS MMF responses with $P = 1$ to $P = 5$ waveform averaging for LFM with worst-case straddling. As P increases, the PSL improves monotonically but with diminishing improvements for successive values of P .

Table 2.8: Worst-case and average PSLs and MMLs of the $P = 1$ to $P = 5$ filter averaged LS MMF response for LFM and PD. The averaged PSL and MML values correspond to the average PSL and MML across all 40 straddled LS MMF responses.

P	<i>LFM</i>				<i>PD</i>			
	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)
1	-47.6	0.86	-56.1	0.79	-50.1	0.18	-59.0	0.13
2	-49.9	0.86	-58.7	0.78	-52.3	0.18	-61.8	0.13
3	-50.8	0.86	-59.1	0.78	-53.2	0.18	-62.2	0.13
4	-51.3	0.86	-59.3	0.78	-53.6	0.18	-62.3	0.13
5	-51.6	0.86	-59.3	0.78	-53.9	0.18	-62.4	0.13

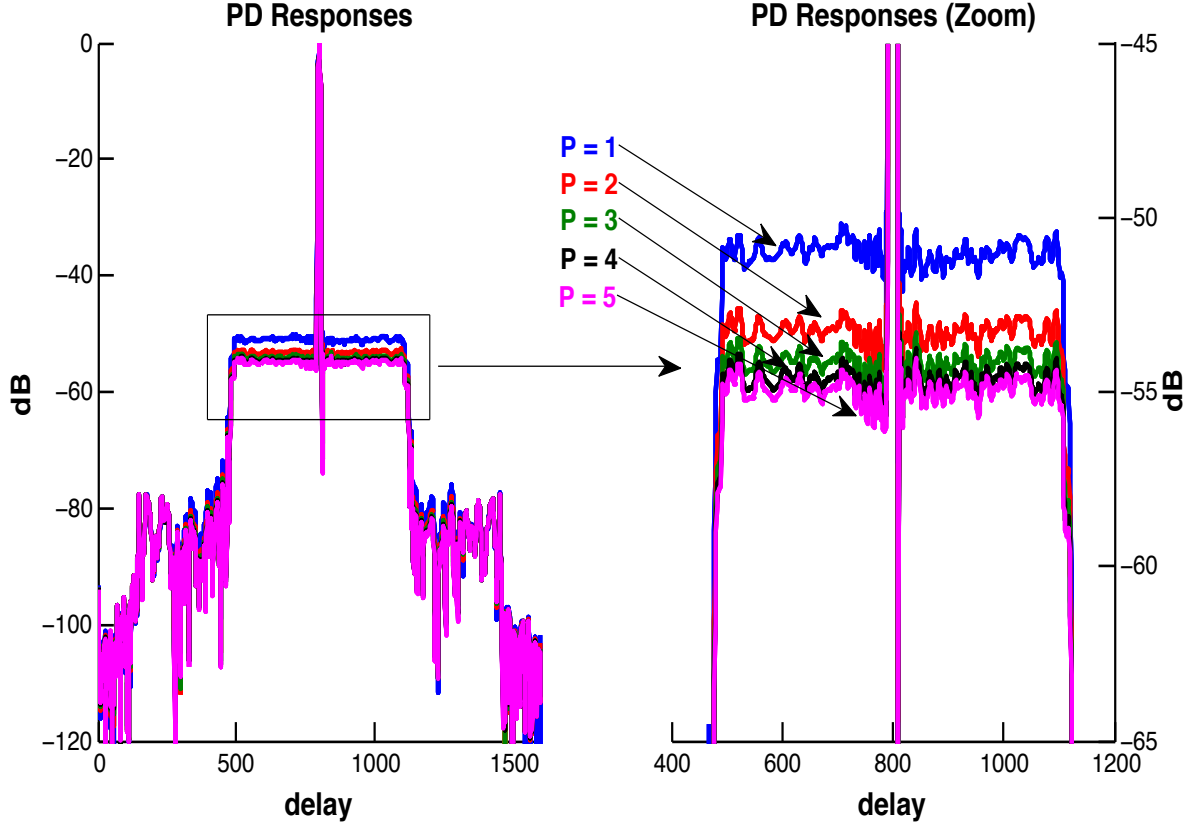


Figure 2.14: LS MMF responses with $P = 1$ to $P = 5$ waveform averaging for PD with worst-case straddling. As P increases, the PSL improves monotonically but with diminishing improvements for successive values of P .

When comparing filter averaging (see Table 2.7) and waveform averaging (see Table 2.8), no significant difference in performance metrics is observed. The PSLs are within 1 dB of each other for $P = L$. This observation can possibly be explained by the fact that, for the same degree of averaging ($P = L$), the initial discretized versions of the waveform are sampled at the same rate for the two averaging methods ($PKBT = LKBT$). Furthermore, filter averaging and waveform averaging use two different "operations": filter construction and averaging (filter or waveform). Filter averaging generates the LS MMFs and then averages them together, while waveform averaging implements averaging first, then generates the LS MMF. Thus, the similarities between the two averaging methods can be explained at a relatively simple level because they are sampled at the same rate and involve the same operations, though the order is switched.

2.4.3 Combination Averaging

Filter and waveform averaging can be combined to generate what is called the waveform-filter or combination averaging method. Thus, the initial discretized version of the waveform is further oversampled such that the new sampling rate is $PLKBT$, where P and L are the number of samples averaged for the waveform and the number of filters that are averaged together, respectively. If P and L are varied between 1 and 5, as was done previously, there are now 25 LS MMFs to examine. For example, for the $P = 3, L = 2$ LS MMF, the initial waveform needs to be sampled at $6KBT$. Waveform averaging is performed before filter averaging. Thus, the 1st, 2nd, and 3rd samples are averaged together, and then the 4th, 5th, and 6th, and so on. The resulting model of the waveform is then polyphase decomposed into two discrete waveforms (using every other sample) so that two LS MMFs can be constructed. Looking at it from the initial, highest sampled waveform, samples 1, 2, 3, 7, 8, 9... are included in the waveform used to construct the first LS MMF, while samples 4, 5, 6, 10, 11, 12... are used for the second LS MMF. These two LS MMFs are then averaged together. This process is completed for all values of L and P from 1 to 5.

Figures 2.15 and 2.16 and Tables 2.9 and 2.10 display the results from the waveform-filter averaging method. The $P = 1$ responses correspond to the results seen in the filter averaging section (L varying from 1 to 5), and likewise, the $L = 1$ responses correspond to the results in the waveform averaging section (P varying from 1 to 5). Again, these results were generated for LFM and PD.

Figure 2.15 gives the LS MMF responses for LFM with worst-case straddling, and Figure 2.16 gives the LS MMF responses for PD with worst-case straddling. Again, notice that as L increases, PSLs improve but with diminishing improvements as L gets larger. When comparing the $P = 2$ plots to the $P = 5$ plots, the same diminishing improvement is observed. Furthermore, the range of worst-case sidelobe levels for the $P = 5$ responses is smaller than the sidelobe range in the $P = 2$ responses. As observed when only one averaging method is used, MML is relatively constant (with less than a 0.01 dB increase) as P and L increase. The improvement in PSL, as L increases from 1 to 5, is greater for the $P = 2$ case than when $P = 5$, though the PSL values are better in the $P = 5$

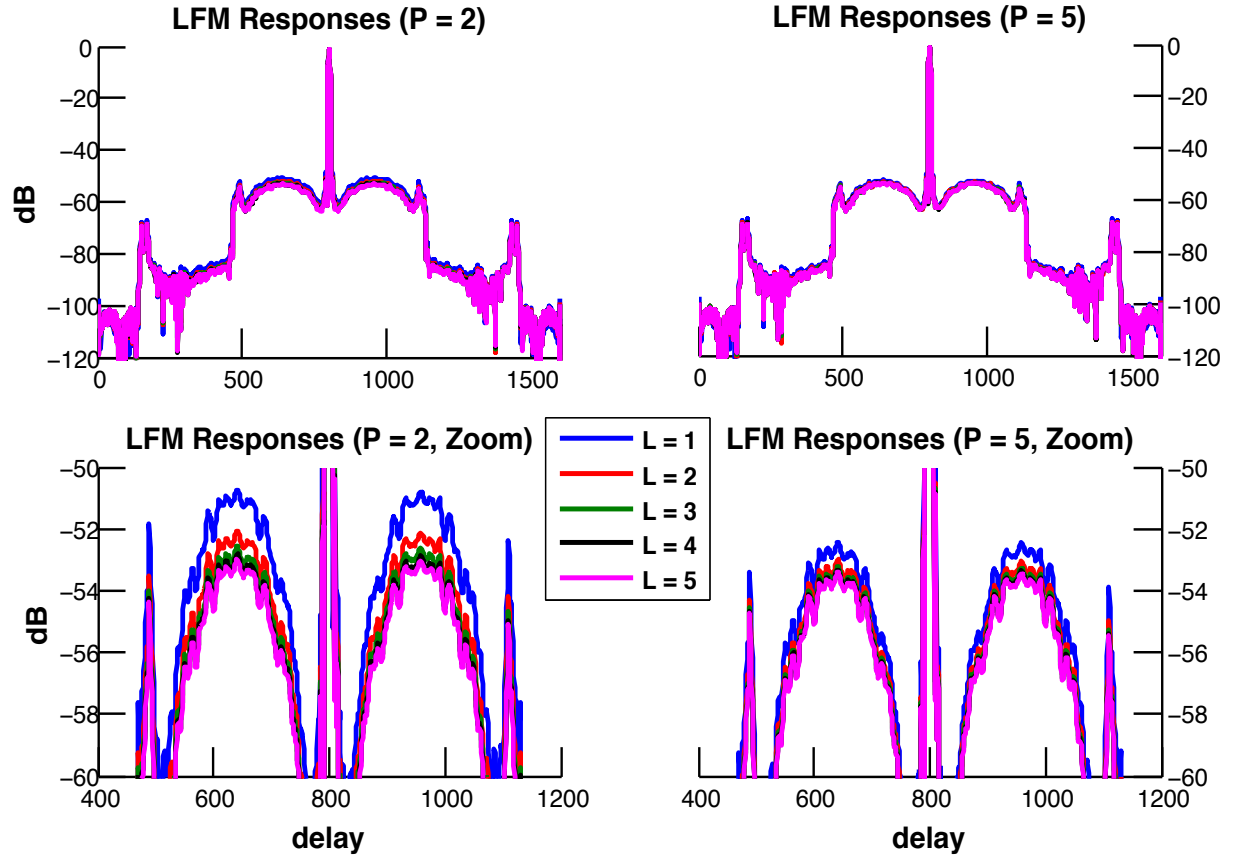


Figure 2.15: LS MMF responses with filter and waveform averaging for LFM with worst-case straddling. The responses on the left were generated with L varied from 1 to 5 and $P = 2$; the responses on the right were generated with L being varied from 1 to 5 and $P = 5$. The bottom plots are the zoomed-in views of the sidelobes near the mainlobe of the LS MMF response.

case.

As seen in Tables 2.9 and 2.10, the best PSL occurs when L and P equal five. Notice that the PSLs are roughly symmetric about the main diagonal of the table. Provided in Appendix A are figures and tables that were used to find the values in Tables 2.9 and 2.10. Figures A.5 through A.8 display the PSLs for all combination averaged (P and L varied from 2 to 5) LS MMFs for each of the 40 straddling shifted versions of the received waveform. Tables A.1 through A.4 give the worst-case and average PSLs and MMLs corresponding to those figures.




Figure 2.16: LS MMF responses with filter and waveform averaging for PD with worst-case straddling. The responses on the left were generated with L varied from 1 to 5 and $P = 2$; the responses on the right were generated with L being varied from 1 to 5 and $P = 5$. The bottom plots are the zoomed-in views of the sidelobes near the mainlobe of the LS MMF response.

Table 2.9: Worst-case PSL with L varied from 1 to 5 and P varied from 1 to 5 for LFM. The numbers on the axis corresponds to values of L and P used to construct the LS MMF which corresponds to the PSL value at the intersection of the L^{th} column and P^{th} row. PSL improves as the product LP increases and the best PSL occurs when L and P equal five.

$P \backslash L$	1	2	3	4	5
1	-47.6	-49.9	-50.7	-51.2	-51.5
2	-49.9	-51.2	-51.7	-52.0	-52.1
3	-50.8	-51.7	-52.1	-52.2	-52.4
4	-51.3	-52.0	-52.2	-52.4	-52.5
5	-51.6	-52.1	-52.4	-52.5	-52.5

Table 2.10: Worst-case PSL with L varied from 1 to 5 and P varied from 1 to 5 for PD. The numbers on the axis corresponds to values of L and P used to construct the LS MMF which corresponds to the PSL value at the intersection of the L^{th} column and P^{th} row. PSL improves as the product LP increases and the best PSL occurs when L and P equal five.

$P \backslash L$	1	2	3	4	5
1	-50.1	-52.4	-53.2	-53.6	-53.9
2	-52.3	-53.6	-54.1	-54.4	-54.5
3	-53.2	-54.1	-54.5	-54.7	-54.8
4	-53.6	-54.4	-54.6	-54.8	-54.9
5	-53.9	-54.5	-54.8	-54.9	-54.9

2.5 Doppler Effects

Straddling effects, and the averaging methods used to mitigate them, occur in the range/time domain. Next, consider the effects of fast-time Doppler shifts on filter responses and determine if, through averaging techniques, it is possible to mitigate the effects of Doppler. Like that of slow-time Doppler, fast-time Doppler results in phase modulation of the received waveform relative to the transmitted waveform. Unlike slow-time Doppler, which considers phase changes on a pulse-to-pulse basis, the phase change from fast-time Doppler occurs within a single pulse and results in slight waveform distortion. Thus, the observed Doppler values are generally smaller than those seen in slow-time Doppler, since the observation interval is significantly shorter. For this analysis, consider the following scenario: the system is operating at 10 GHz, with a pulse width of 10 μ s, observing a target that is moving at 60 m/s. These parameters, using the formula $2vf/c \times T$ (Ch. 8 in [8]), result in a Doppler shift of 0.04 radians over the pulse width. FM waveforms are relatively Doppler tolerant, and as a result, the degradation in PSL that is expected should be noticeably less than the degradation observed from straddling.

To implement a Doppler shift on the received waveform, the original model of the received waveform is multiplied via the Hadamard product (element by element multiplication) by a vector of the same length of the waveform, composed of phase shifts (Ch. 17 in [8]). These phase shifts start at 0 and end at the desired Doppler shift, with every element in between linearly spaced

Table 2.11: MML and PSL of the matched filter and LS MMF responses for both waveforms, with and without Doppler. The LFM matched filter and LS MMF responses and the PD matched filter response are relatively unaffected by the Doppler shift. The LS MMF corresponding to PD response exhibits roughly 7 dB in PSL degradation due to Doppler.

	<i>LFM</i>		<i>PD</i>	
	MML (dB)	PSL (dB)	MML (dB)	PSL (dB)
Matched	0	−13.77	0	−40.22
Doppler Matched	0.00	−13.75	0.00	−40.14
LS MMF	0.73	−67.18	0.09	−77.35
Doppler LS MMF	0.73	−67.18	0.09	−77.00

between 0 and the desired Doppler shift. This multiplication is mathematically described as

$$\mathbf{s}_{Dop} = \mathbf{s} \odot \exp(-j \times \boldsymbol{\omega}_{vec}), \quad (2.5)$$

where $\boldsymbol{\omega}_{vec} = [0, \frac{\omega_{max}}{N}, \dots, \frac{\omega_{max}(N-1)}{N}, \omega_{max}]$, and \odot is the Hadamard product.

In Figure 2.17 are the idealistic LS MMF and matched filter responses for LFM and PD. Plotted with these responses are the LS MMF responses and matched filter responses of the same filter, but with the received model of the waveform subjected to the maximum Doppler shift, 0.04 radians. The matched filter, with and without Doppler, shows minimal degradation in performance when Doppler is present. The LS MMF, however, has an increase in sidelobe levels near the mainlobe. For LFM, this increase does not affect the PSL, since the peak sidelobe occurs further away from the mainlobe in the response. The PSL of the LS MMF response for PD, however, increases by roughly 7 dB because PD is not as robust to Doppler as LFM. It is evident in Table 2.11 that the matched filter and LS MMF responses corresponding to LFM are relatively unaffected by the Doppler shift. This result is attributed to the range-Doppler relationship that is a well-known characteristic of LFM (Ch. 20 of [8]). Again, as was noticed with straddling, Doppler effects have a relatively small effect on the matched filter responses while the effects of Doppler on the LS MMFs are more pronounced, since the LS MMF has relatively low sidelobe levels in the idealistic case.

Figure 2.18 shows the LS MMF responses for LFM and PD for Doppler shifts between ± 0.05

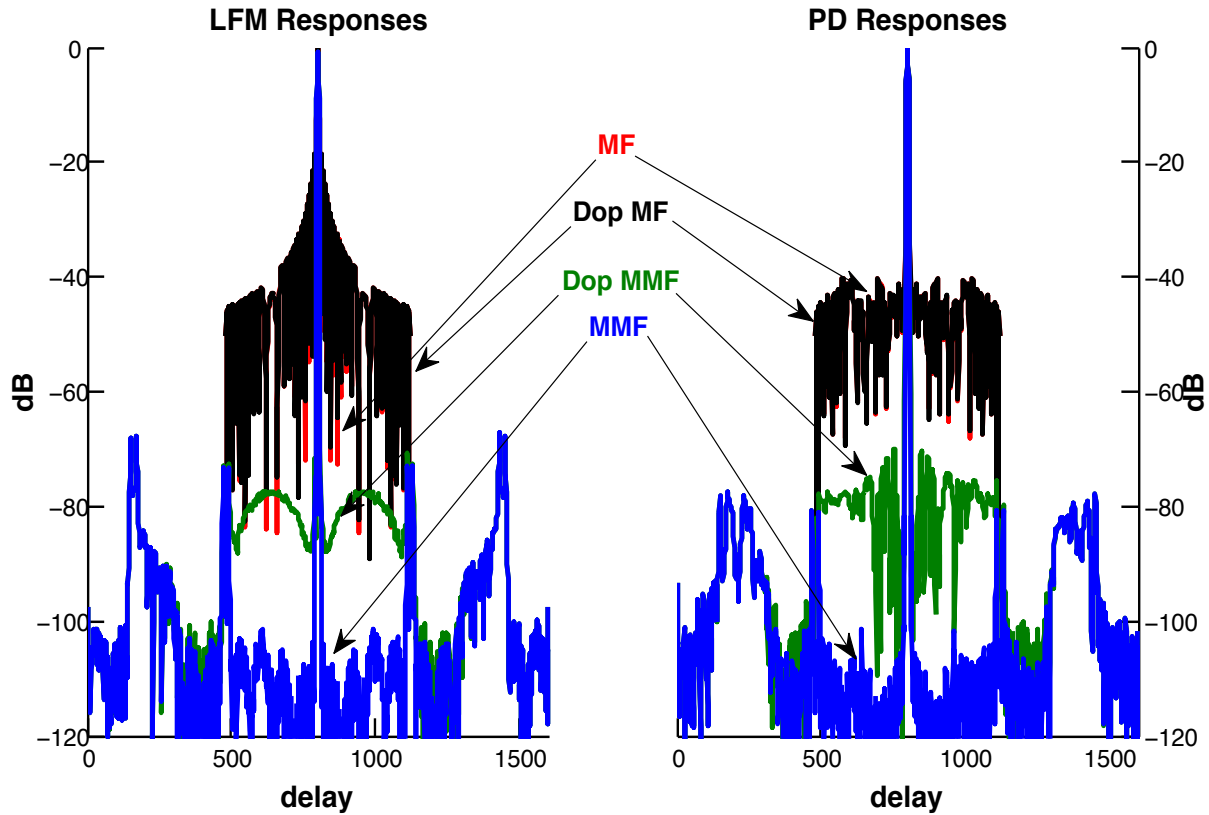


Figure 2.17: Matched filter and LS MMF responses for LFM and PD, for ideal and Doppler shifted received waveforms. The matched filter shows little degradation in performance with Doppler. The LS MMF has an increase in sidelobe levels near the mainlobe with Doppler, compared to the ideally sampled case.

to -0.05 radians. This range-Doppler plot shows the sidelobe levels, in color, as the Doppler shift varies. At zero Doppler, the LS MMF has the lowest sidelobes. Moving away from zero, the sidelobes increase monotonically. This pattern is mirrored across the zero Doppler cut in the range-Doppler plot.

2.6 Doppler Averaging

Now averaging methods similar to those used for straddling will be implemented in an attempt to compensate for Doppler. Doppler averaging is not a new idea (e.g., [53]), but it has not yet been explored with respect to the LS MMF or using the averaging method that is implemented

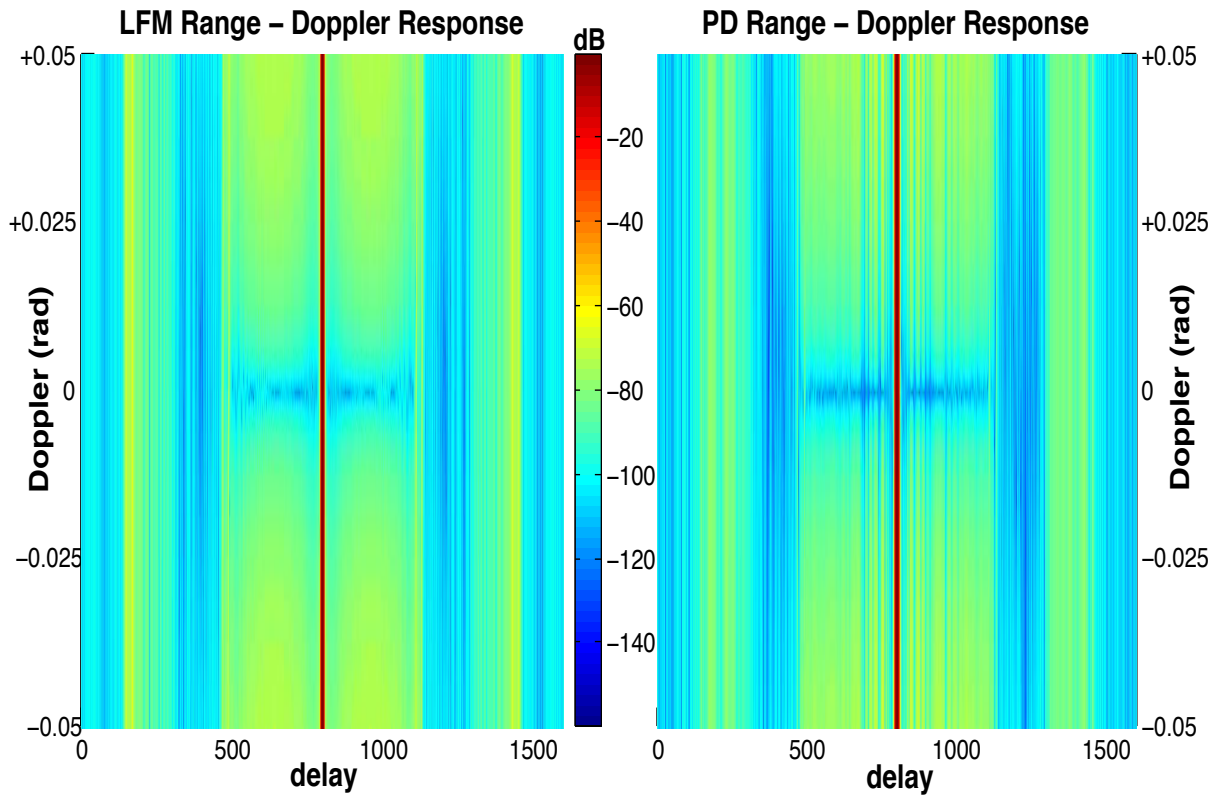


Figure 2.18: LS MMF responses for LFM and PD for Doppler shifts between $+0.05$ to -0.05 radians. At zero Doppler, the sidelobes for the LS MMF are at their lowest. Sidelobe levels increase monotonically as the Doppler shift increases in magnitude.

here. Since only Doppler up to 0.04 radians is used, 0.04 is the "end point" of the averaging. To average across Doppler, J LS MMFs are averaged together. These LS MMFs differ from one another by the phase shift applied to the model of the discretized version of the waveform used in filter construction. The maximum Doppler shift for each of these discretized waveforms varies linearly between -0.04 and $+0.04$ radians. For example, if three filters ($J = 3$) are averaged, one waveform model will have a Doppler shift of -0.04 radians, the second a Doppler shift of 0 radians, and the last, $+0.04$ radians. The Doppler shifts used to generate the Doppler averaged LS MMFs are provided in Table 2.12 for $J = 1$ to $J = 5$.

Averaging is done across both positive and negative Doppler shifts because the target's relative direction (approaching or receding) is unknown. This averaging process is completed using (2.5)

with ω_{max} values from Table 2.12 to generate the discretized versions of the waveform. These discretized waveforms are then used to generate J LS MMFs using (2.2) and averaged according to (2.3), replacing L with J .

To test this averaging method, straddling is ignored. 128 different Doppler-shifted versions of the waveform are used to represent the possible models of the received waveform. These models of the waveform are linearly spaced in Doppler between -0.05 and $+0.05$ radians. The filter responses for the non-Doppler averaged and the Doppler averaged LS MMFs for all 128 Doppler-shifted waveforms are used to generate the range-Doppler plots.

Figure 2.19 shows the range-Doppler plot of the LS MMF corresponding to LFM, with Doppler averaging, for Doppler shifts between $+0.05$ and -0.05 radians. As in the previous, non-Doppler averaged case, at zero Doppler, the LS MMF has the lowest sidelobes. As Doppler increases in magnitude, the sidelobes increase monotonically. Doppler averaging does not appear to produce any improvement in the LS MMF response relative to the non-Doppler averaged scenario. Also, when compared to Figure 2.18, the sidelobe levels of the $J = 2$ and $J = 5$ LS MMFs are worse than those found in the non-averaged, $J = 1$ case.

The LS MMF responses for three different Doppler shifted received waveforms (-0.04 , -0.02 , and 0 radians) and four different values of J (1,2,3, and 5) are provided in Figure 2.20. For each Doppler shift, the LS MMF response does not noticeably change as P varies. There does not appear to be any significant improvement or degradation in performance when Doppler averaging is used.

The same simulations were run for PD, shown in Figure 2.21 and Figure 2.22, and behaviors similar to those seen with LFM were observed. Though Doppler averaging had no significant effect

Table 2.12: The ω_{max} values used to generate the Doppler averaged LS MMFs.

J	$\omega_{max}/0.04$
1	0
2	$-1, 1$
3	$-1, 0, 1$
4	$-1, -1/3, 1/3, 1$
5	$-1, -1/2, 0, 1/2, 1$

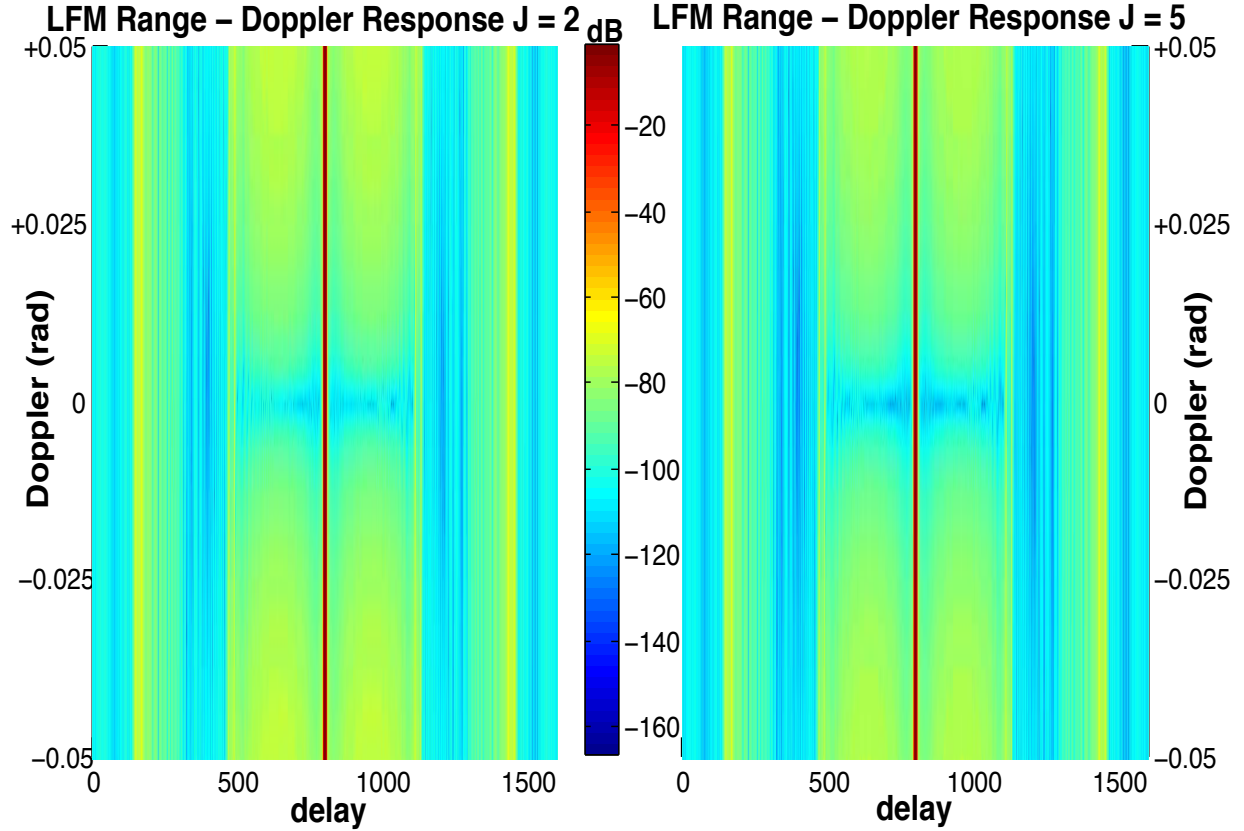


Figure 2.19: Doppler averaged LS MMF responses with $J = 2$ and $J = 5$ for LFM for Doppler shifts from $+0.05$ to -0.05 radians. Averaging was linearly done from -0.04 to $+0.04$ radians. The LS MMF has the lowest sidelobes at zero Doppler, and sidelobes increase monotonically moving away from zero. The sidelobe levels of the $J = 2$ and $J = 5$ LS MMFs are worse than those found in the non-averaged $J = 1$ case.

on the PSL for LFM, it does appear to have a slightly more pronounced effect for PD. This effect is most noticeable in the range-Doppler plot in Figure 2.21. For PD, the sidelobe levels in the $J = 5$ case appear to be slightly higher than those in the $J = 2$ case, both of which are higher than those seen when $J = 1$. Thus, instead of mitigating the effects of Doppler, Doppler averaging appears to further increase the sidelobe levels.

Given that this first Doppler averaging method proved fruitless, the averaging method was tweaked under one assumption: whether the target is approaching or receding from the radar system is known. If this assumption is made, half of the averaging range can be eliminated, and it is no longer necessary to consider both positive and negative Doppler shifts. By assuming the relative

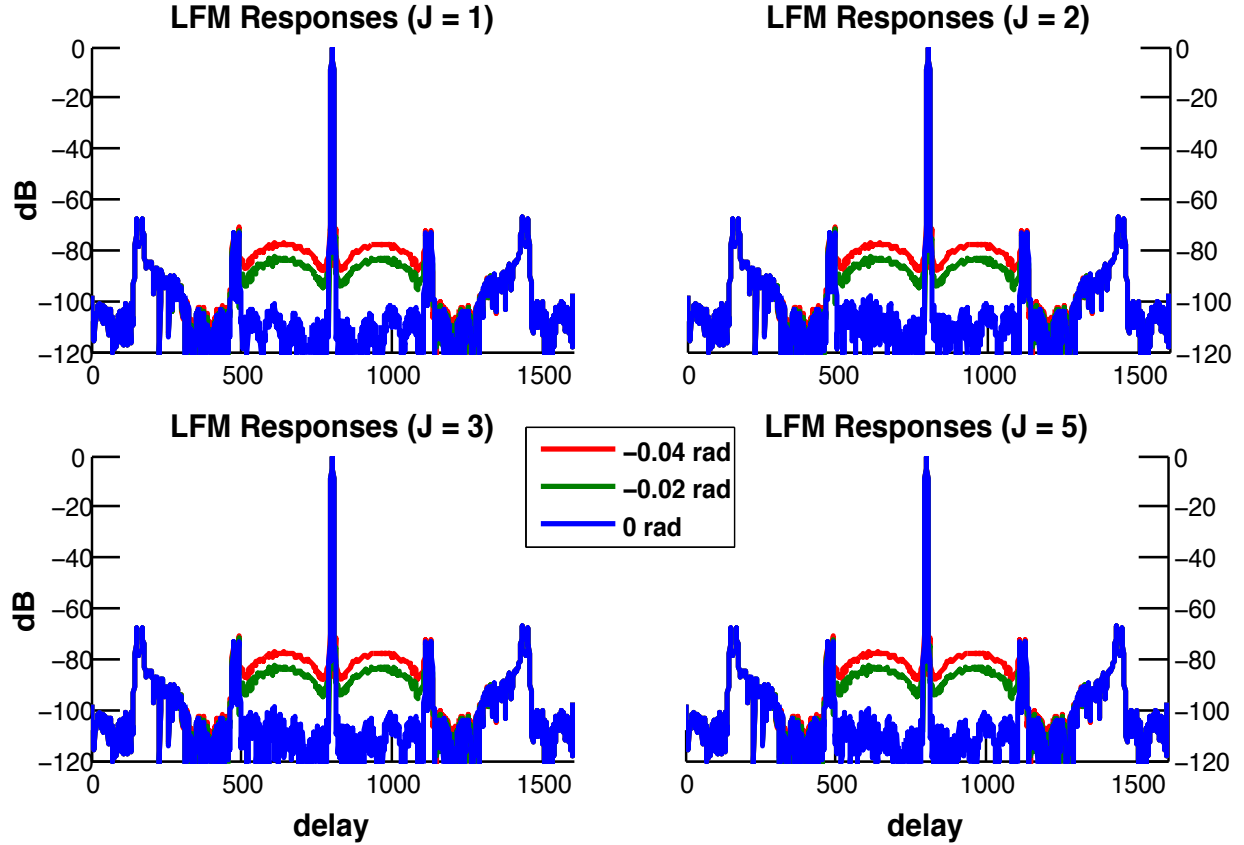


Figure 2.20: LS MMF responses with Doppler averaging for LFM with three different Doppler-shifted received waveforms (0, -0.02 , and -0.04 radians) and $J = 1$, $J = 2$, $J = 3$, and $J = 5$. Averaging occurs on the interval of -0.04 to $+0.04$ radians. From this figure, no noticeable change in PSL is observed when using Doppler averaging.

direction of the target, averaging is performed linearly between the maximum Doppler shift and zero. This change is implemented by altering the averaging process described in the beginning of the current section and using the ω_{max} values from Table 2.13. This new Doppler averaged LS MMF has averaging that is linearly spaced between ± 0.04 and 0 radians, assuming the target is approaching/receding.

A range-Doppler plot is generated for the new Doppler averaged LS MMF, with the received waveform model having a Doppler shift between -0.05 and 0 radians, shown in Figure 2.23. The lowest sidelobe levels now occur at -0.02 radians, which is halfway between the end points of the averaging interval. Moving away from -0.02 radians, sidelobe levels increase. This range-Doppler plot illustrates the fact that sidelobe levels can be improved through averaging, if the

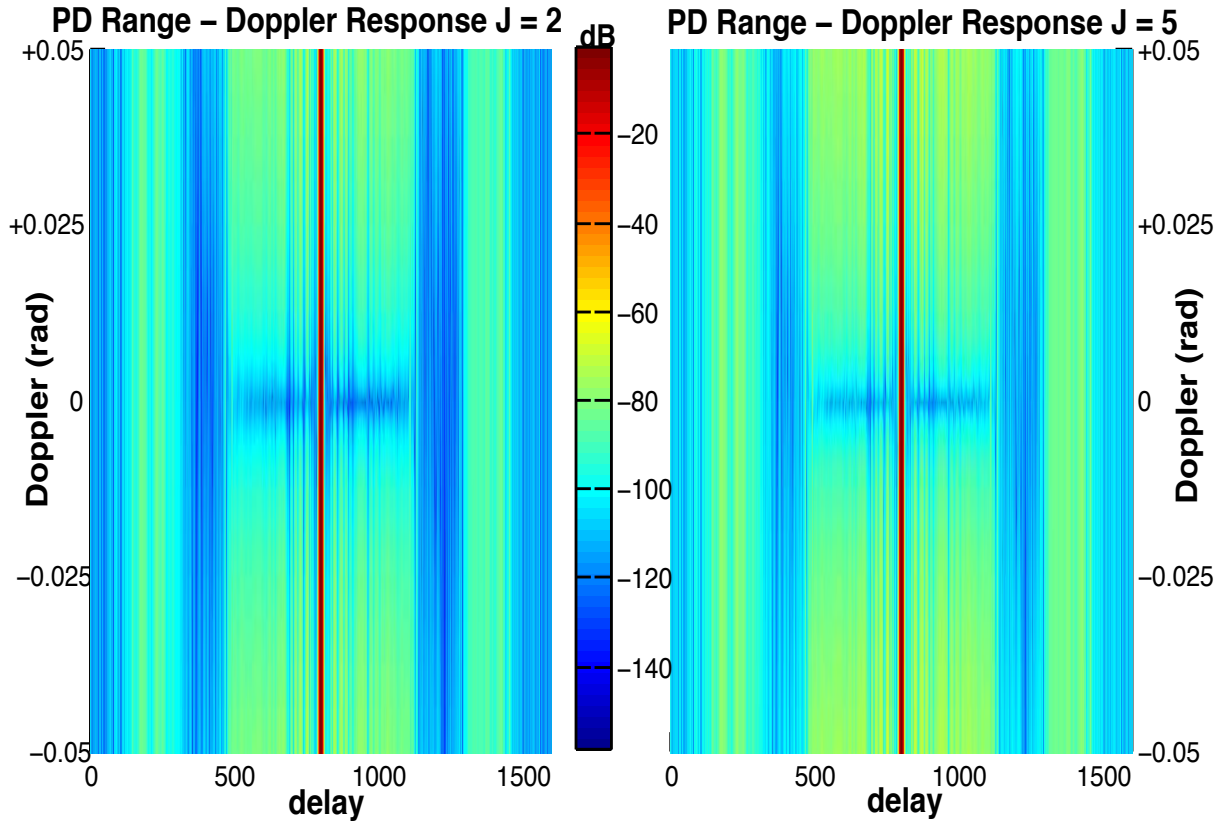


Figure 2.21: Doppler averaged LS MMF responses with $J = 2$ and $J = 5$ for PD for Doppler shifts from $+0.05$ to -0.05 radians. Averaging was linearly done from -0.04 to $+0.04$ radians. The LS MMF has the lowest sidelobes at zero Doppler, and sidelobes increase monotonically away from zero. The sidelobe levels of the $J = 5$ LS MMF are worse than those for the $J = 2$ LS MMF. The sidelobe levels from both the $J = 2$ and $J = 5$ responses are worse than those found in the non-averaged $J = 1$ case.

Table 2.13: The ω_{max} values used to generate the Doppler averaged LS MMFs when relative direction of travel is known.

J	$\pm\omega_{max}/0.04$
1	$1/2$
2	$0, 1$
3	$0, 1/2, 1$
4	$0, 1/3, 2/3, 1$
5	$0, 1/4, 1/2, 3/4, 1$

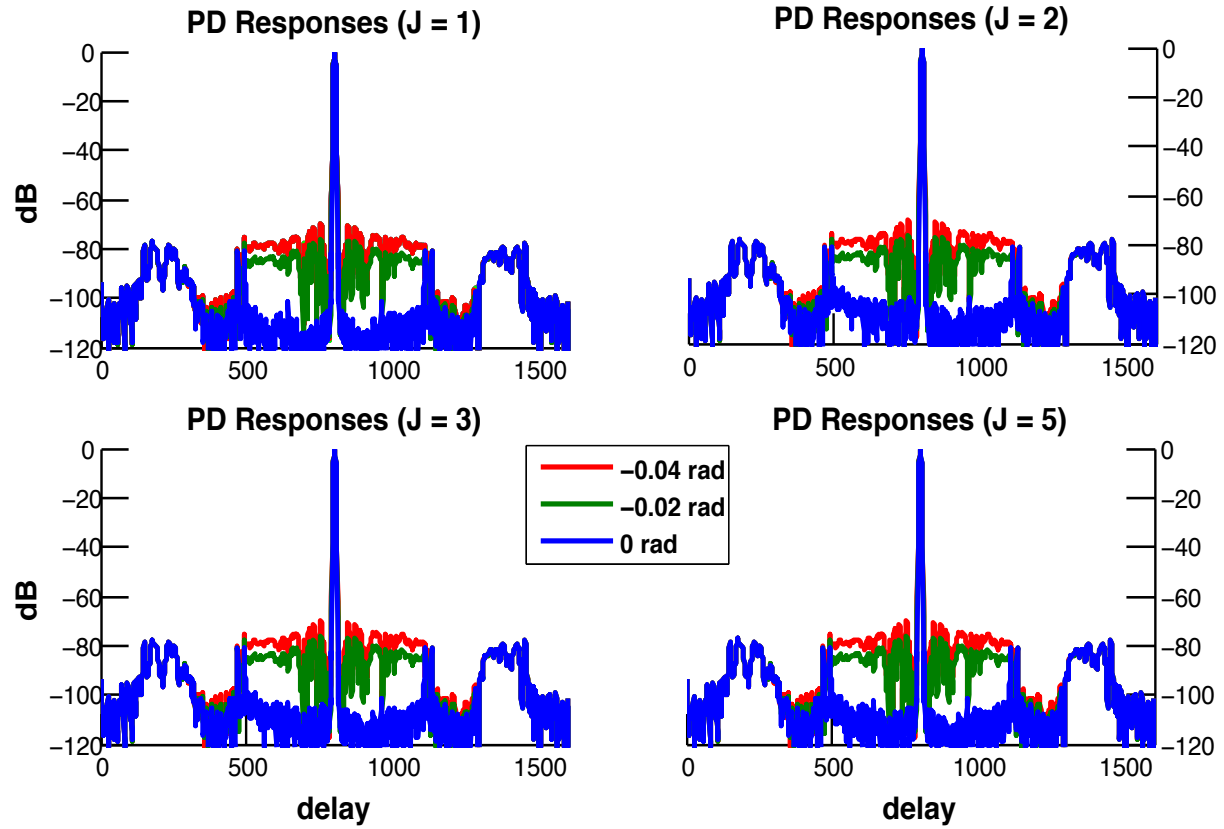


Figure 2.22: Doppler averaged LS MMF responses for PD with three Doppler shifted received waveforms (0, -0.02 , and -0.04 radians) and with $J = 1$, $J = 2$, $J = 3$ and $J = 5$. Averaging occurred on the interval of -0.04 to $+0.04$ radians. From this figure, no noticeable change in PSL is observed when using Doppler averaging.

relative direction (approaching or receding) of the target is known and if there is a limit on the fast time Doppler shift that can be expected. This improvement can be seen when comparing the non-Doppler averaged LS MMF range-Doppler plot in Figure 2.18 to Figure 2.23. By extension, a more robust bank of Doppler-tuned filters may be formed using this approach.




Figure 2.23: Doppler averaged LS MMF responses for LFM and PD for Doppler shifts from -0.05 to $+0.05$ for $J = 5$. Averaging was linearly done from -0.04 to 0 radians. The best sidelobe performance occurs at -0.02 radians, halfway between the end points of the averaging interval. Moving away from -0.02 radians, sidelobe levels increase. The Doppler averaging was completed assuming that the target is receding relative to the radar system.

Chapter 3

Analysis of Adaptive Pulse Compression

Next, the alterations of the Adaptive Pulse Compression (APC) algorithm such that it is applicable to FM waveforms are examined. The advancements made in APC include decreasing the computational complexity/cost via polyphase decomposition (FAPC) [54, 55], application of a gain constraint to account for SNR losses incurred in FAPC [56], repairing the APC algorithm so that it can estimate into the eclipsed regions [2], and application of APC to polarimetric waveforms [20]. Process gains from each of these advancements are used in this work, as they have led to both improved performance and practicality of the APC algorithm. The modifications of APC such that it is applicable to FM waveforms was not part of this work. Rather, this work focuses on the analysis of APC when FM waveforms are used, and both simulation and measured results are presented.

3.1 APC Modifications for Application to FM Waveforms

As with the LS MMF, APC was initially derived for polyphase coded waveforms. Structurally, several changes are needed for APC. When oversampling occurs, a polyphase decomposition version denoted as fast APC (FAPC) has been shown to perform as well as APC with decreased computational cost [54]. For polyphase decomposition, the decimation factor that is used is equal to the oversampling factor K . With FAPC, the adaptive filter coefficients \mathbf{w} and the discretized version of the waveform \mathbf{s} from (1.6) and (1.7) are now represented via polyphase decomposition and shown




Figure 3.1: Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $K = 1$. Neither the filters nor APC are able to discern the smaller target with confidence.

in

$$\mathbf{w}_k(n) = \frac{(\mathbf{C}_k(n) + \mathbf{R}_k)^{-1} \mathbf{s}_k}{\sum_{i=0}^{K-1} \mathbf{s}_i^H (\mathbf{C}_i(n) + \mathbf{R}_i)^{-1} \mathbf{s}_i} \quad (3.1)$$

and

$$\mathbf{C}_k(n) = \sum_{\tau=-N+1}^{N-1} \rho(n + \tau) \mathbf{s}_{k,\tau} \mathbf{s}_{k,\tau}^H. \quad (3.2)$$

Matrices \mathbf{C}_k and \mathbf{R}_k are also decimated by K due to polyphase decomposition used in FAPC.

Oversampling relative to the 3 dB bandwidth is the first alteration needed such that APC can be applied to FM waveforms. As described in Section 2.1, FM waveforms need to be sampled at a higher rate to provide adequate fidelity. Three iterations of APC are used and oversampling factors of 1, 3, and 5 are examined. Figure 3.1 shows the matched filter, LS MMF, and APC responses for both the linear frequency modulated (LFM) and optimal FM 'performance diversity' (PD) waveforms when no oversampling ($K = 1$) is used. The target scenario has two targets, one with an SNR of 80 db (at a range of 100) and the other with an SNR of 15 dB (at a range of 95).




Figure 3.2: Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $K = 3$. The LS MMF is able to discern the smaller target, but the matched filter and APC cannot.

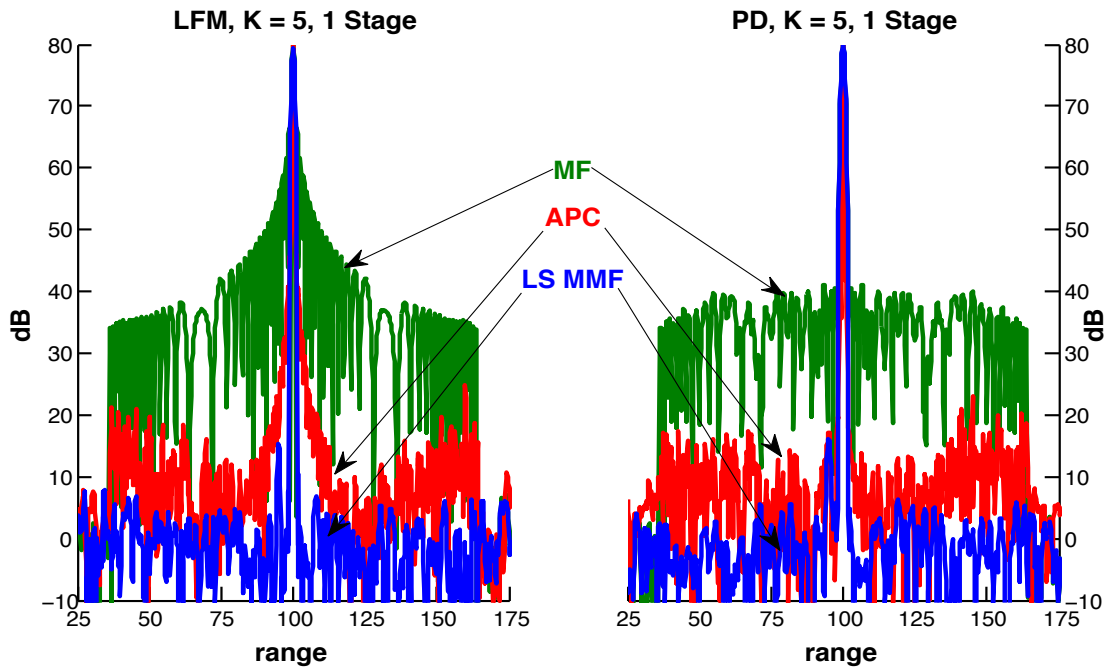


Figure 3.3: Matched filters, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $K = 5$. The LS MMF is able to discern the smaller target, but the matched filter and APC cannot.

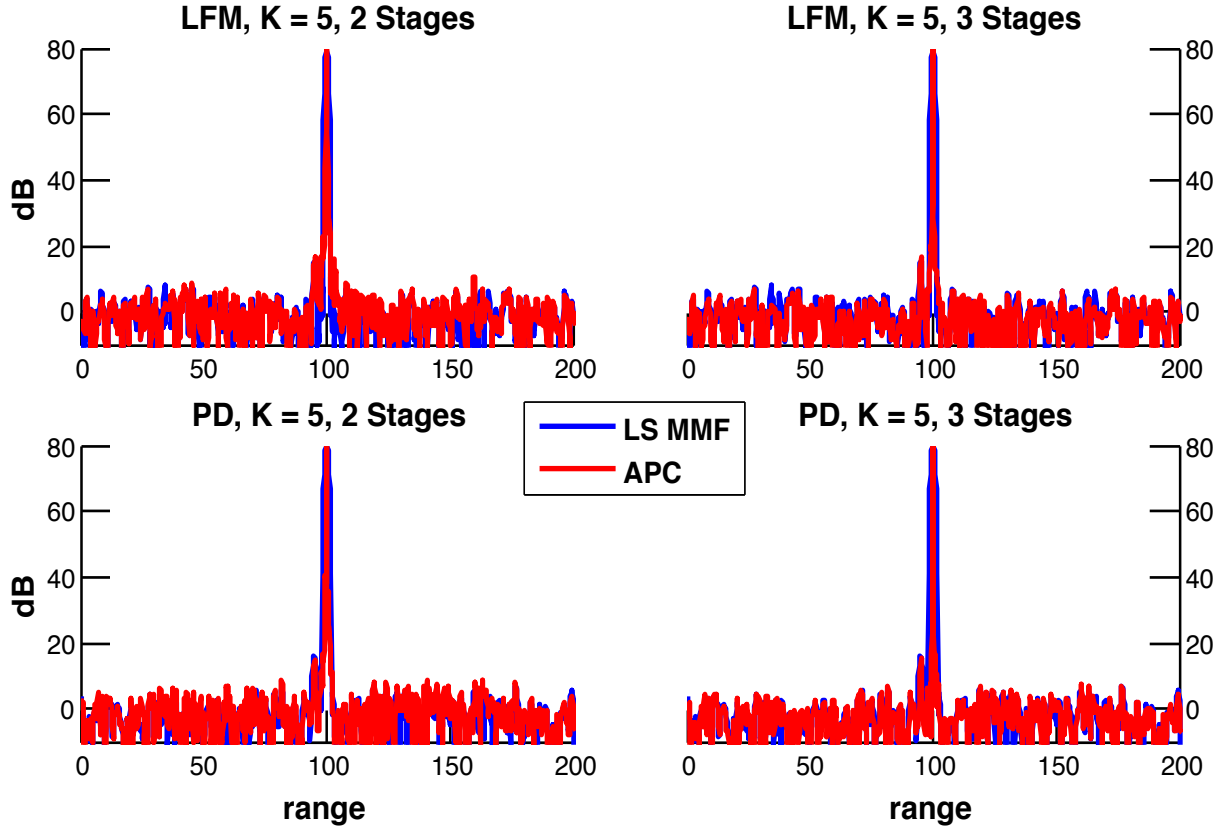


Figure 3.4: LS MMF and APC responses for LFM and PD with two and three adaptive stages and $K = 5$. The targets are more distinguishable with three adaptive stages of APC than with two.

This target scenario will be used for all figures, until noted otherwise. The matched filter and LS MMF are unable to discern the smaller target. With one adaptive stage, APC appears to be able to find the smaller target, but also presents some false targets. With an oversampling factor of $K = 3$, as shown in Figure 3.2, the matched filter is unable to discern the smaller target, but the LS MMF is able to do so, to some degree. With one stage of APC, neither the LFM-based nor PD-based implementations of APC can discern the smaller target. The same is true for an oversampling factor of $K = 5$, represented in Figure 3.3.

When APC is implemented with two or three adaptive stages, as in Figure 3.4, in which $K = 5$, the APC response is able to discern the smaller target. With the third adaptive stage, the null between the large and small target becomes larger, making the targets more distinguishable. Because it was used for the LS MMF and produces desirable results, from this point forward, $K = 5$.




Figure 3.5: Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $z = 1$ row zeroed, which results in a resolution that is slightly coarser than APC implemented in Figure 3.3. Only the LS MMF is able to discern the smaller target.

The second alteration that is needed, again taking from the LS MMF adaptations in section 2.1, is the idea of beamspooling. For APC, beamspooling is done on $\mathbf{C}_k(n)$, where $\mathbf{C}_k(n)$ is the decimated structured signal correlation matrix, by zeroing out rows above and below the current range index n . The number of rows that are zeroed is denoted by z . Oversampling results in super-resolution and significant mismatch loss (MML), as defined in (1.3). By zeroing rows, MML is reduced and sidelobe levels are suppressed, but degradation in resolution occurs.

Figure 3.5 gives the matched filter, LS MMF, and APC responses for LFM and PD with one adaptive stage of APC and $z = 1$ row zeroed above and below n in \mathbf{C}_k . The matched filter is unable to discern the smaller target, while the LS MMF can. This will hold true for all simulations regarding zeroing rows in \mathbf{C}_k since the zeroing process does not alter the discretized waveform, matched




Figure 3.6: Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $z = 3$ rows zeroed, resulting in a coarser resolution than APC responses when $z = 1$. APC is able to somewhat distinguish the target for PD, but is unable to do so for LFM.

filter, or LS MMF. With one adaptive stage, neither the LFM- nor PD-based APC implementations can find the smaller target. When $z = 3$ rows are zeroed out and one adaptive stage is used, as in Figure 3.6, the PD-based APC implementation can now find the smaller target.

When $z = 5$ rows are zeroed on either side of the current range index n in \mathbf{C}_k and one adaptive stage is used, the LFM-based APC implementation is still unable to find the smaller target, as shown in Figure 3.7. The PD-based APC does a better job at discerning the target than was seen in previous APC results. As more rows are zeroed in \mathbf{C}_k , super-resolution diminishes and APC is better able to suppress sidelobes. This trend was also observed for the LS MMF when rows in \mathbf{A} were zeroed. Due to the desirable performance of APC when $z = 5$ rows were zeroed above and below the current range cell n in \mathbf{C}_k , this number of zeroed rows will be used from this point




Figure 3.7: Matched filters, LS MMF, and APC responses for LFM and PD. APC was implemented with one adaptive stage and $z = 5$ rows zeroed out on either side of n , which results in a resolution that is coarser than APC implemented in Figure 3.6. The PD-based implementation of APC is able to discern the smaller target with confidence, but the LFM-based implementation is still unable to discern the smaller target.

forward. It should be noted that the desired number of zeroed rows appears to correspond with the oversampling factor K . This reduces mismatch loss because it prevents the adaptive filter from causing a super-resolution condition, leading to self-suppression as described previously.

When APC is implemented with two or three adaptive stages, as in Figure 3.8, with $z = 5$ rows zeroed out on either side of n in \mathbf{C}_k , the APC responses for LFM and PD are able to discern the smaller target. This time, there is no noticeable performance gain from using a third adaptive stage of APC. Since $K = 5$ will be used throughout the rest of this work, a maximum of two adaptive stages of APC will be implemented.




Figure 3.8: LS MMF and APC responses for LFM and PD with two and three adaptive stages. Five rows were zeroed out above and below n for APC. There is no noticeable performance gain from using a third adaptive stage of APC; the smaller target is distinguishable after two adaptive stages for both LFM and PD.

3.2 Idealistic Simulations and Results

With this FM-capable structure, simulations for APC using LFM and PD can be performed. Because APC is an iterative process, sidelobes can be suppressed beyond that which was seen for the LS MMF. In the first simulation, two targets are placed 5 range cells apart with the larger target having an SNR of 80 dB and the smaller target having an SNR of 15 dB. APC is simulated using the matched filter to initialize the APC algorithm. The next simulations that were completed used the LS MMF, which has the same implementation cost as the matched filter, for initialization. Since the LS MMF has lower sidelobe levels than the matched filter, some improvement in the convergence time of APC is expected.




Figure 3.9: Matched filter, LS MMF and APC responses for LFM. The top two plots show the responses of the matched filter and APC, with APC initialized using the matched filter. The bottom two plots show the responses of the LS MMF and APC, with APC initialized using the LS MMF. When the LS MMF is used to initialize APC, only one adaptive stage is needed to discern the smaller target. Two adaptive stages are for target detection when APC is initialized with the matched filter.

Figure 3.9 shows the responses of the matched filter and APC for LFM when APC is initialized with the matched filter as well as the responses for the LS MMF and APC when APC is initialized with the LS MMF. When the matched filter is used to initialize APC, two adaptive stages are needed to discern the smaller target. Only one adaptive stage is needed when the LS MMF is used to initialize APC. The simulations corresponding to PD are presented in Figure 3.10. In these simulations, when the matched filter is used as the initialization of APC, the first adaptive stage can discern the smaller target, depending on SNR requirements. Only one adaptive stage is needed to discern the smaller target when the LS MMF is used to initialize APC.

Figures 3.11 and 3.12 provide simulations that show the capability of APC in a randomly




Figure 3.10: Matched filter, LS MMF and APC responses for LFM. The top two plots show the responses of the matched filter and APC, with APC initialized using the matched filter. The bottom two plots show the responses of the LS MMF and APC, with APC initialized with the LS MMF. When the LS MMF or the matched filter are used to initialize APC, only one adaptive stage is needed and the filter can discern the smaller target. This was not seen with LFM because LFM does not have as low of range sidelobes as the PD filter responses.

generated, target-rich environment. This environment was generated with 15 randomly placed targets with SNR values between 80 dB and 15 dB. These simulations are generated for both LFM and PD using the matched filter for initialization. With matched filter initialization, two adaptive stages are needed to discern the smaller targets, as observed in the two target scenario. Only one adaptive stage is necessary to discern the targets when the LS MMF is used to initialize APC. When only considering the matched filter and LS MMF responses, the matched filter is unable to uncover all of the targets, while the LS MMF is able to do so.

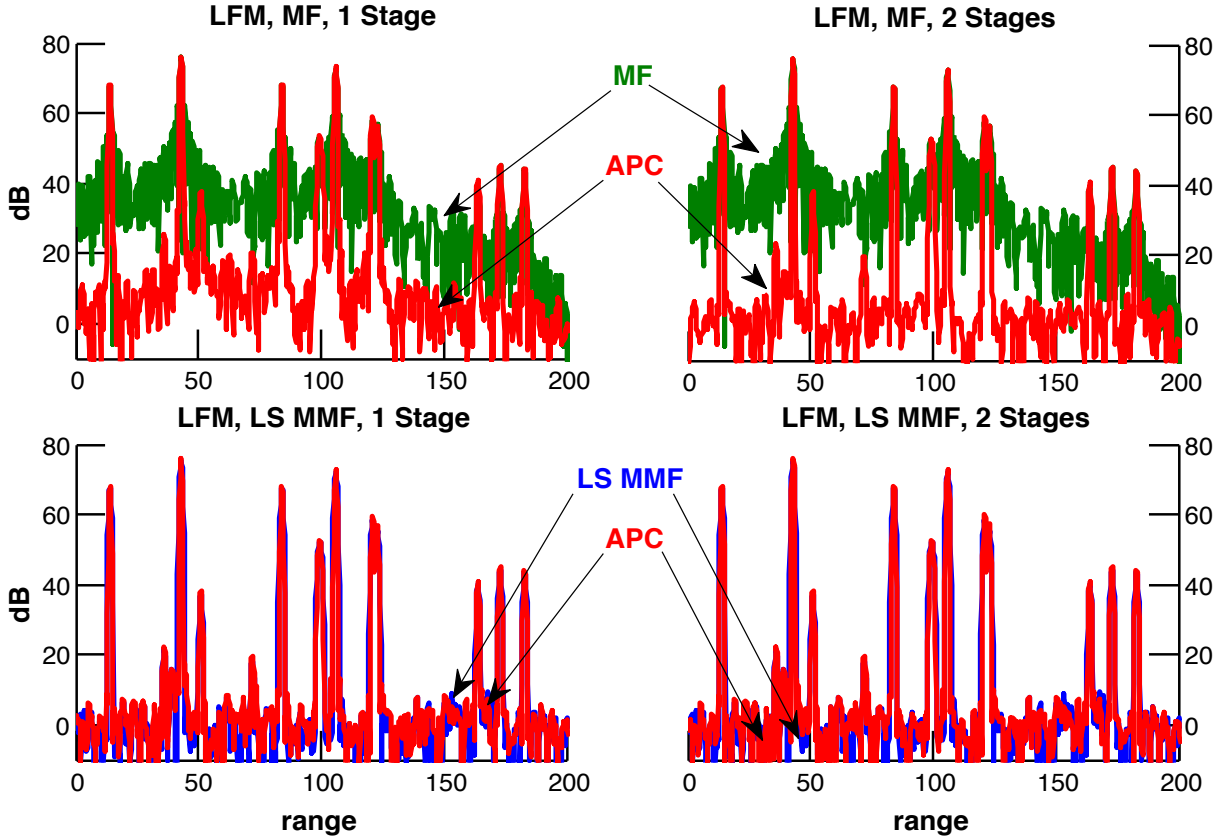


Figure 3.11: Matched filter, LS MMF and APC responses for LFM in a target-rich environment. Two adaptive stages are needed to discern the smaller targets when the matched filter is used to initialize APC, but only one adaptive stage is needed when the LS MMF is used as the initialization filter.

3.3 Straddling Effects on APC

Up to this point, the APC simulations were generated using ideal sampling; that is, the model used for the received waveform is an exact match to the model used in filter construction. Now, the effects of worst-case straddling ($0.5 T_S$) on APC will be considered. The two scenarios previously depicted in Figures 3.9 and 3.10 are now simulated using the straddling delay of $0.5 T_S$. In Figure 3.13, APC is unable to discern the smaller target when straddling is present for LFM with one adaptive stage, but is able to after the second adaptive stage. The matched filter alone is unable to discern the smaller target, which was observed in the non-straddled results. When the LS MMF is used to initialize APC, the smaller target is discernible with one adaptive stage. The LS MMF




Figure 3.12: Matched filter, LS MMF and APC responses for PD in a target-rich environment. As with the plots generated using LFM, two adaptive stages are needed to discern the smaller targets when the matched filter is used to initialize APC, but only one adaptive stage is needed when APC is initialized with the LS MMF.

alone, however, can no longer discern the smaller target due to the increase in sidelobe levels, resulting from straddling effects.

The plots in Figure 3.14 were generated using PD and the two target scenario. When the matched filter is used to initialize APC, the first adaptive stage can discern the smaller target, depending on SNR requirements. The target is clearly distinguishable after two adaptive stages. Only one adaptive stage is needed when the LS MMF is used to initialize APC. Neither the matched filter nor the LS MMF can discern the smaller target due to the increased sidelobe levels from straddling.

In Figure 3.15, the filter responses are shown for the same target-rich scenario used in Section 3.2. When the matched filter is used as the initialization filter of APC, two adaptive stages are




Figure 3.13: Matched filter, LS MMF and APC responses for LFM in a two-target environment with worst-case straddling. The top two plots show the responses of the matched filter and APC initialized with the matched filter. The bottom two plots show the responses of the LS MMF and APC initialized with the LS MMF. Neither the matched filter nor the LS MMF can discern the smaller target. APC initialized with the matched filter requires two adaptive stages to discern the smaller target, while when APC is initialized with the LS MMF, only one adaptive stage is needed.

needed to discern the smaller targets, as was needed in the two target scenario. Only one adaptive stage is needed for target recognition when the LS MMF is used to initialize APC, though two adaptive stages does result in additional sidelobe suppression. Neither the matched filter nor the LS MMF alone can discern the smaller targets, though the LS MMF can distinguish more targets than the matched filter.

The same target-rich environment is used in Figure 3.16, now using PD. When the matched filter is used to initialize APC, two adaptive stages are needed to discern the smaller targets, as was needed in the two target scenario. Only one adaptive stage is needed to find all targets when the LS MMF is used to initialize APC, though two adaptive stages does result in additional sidelobe

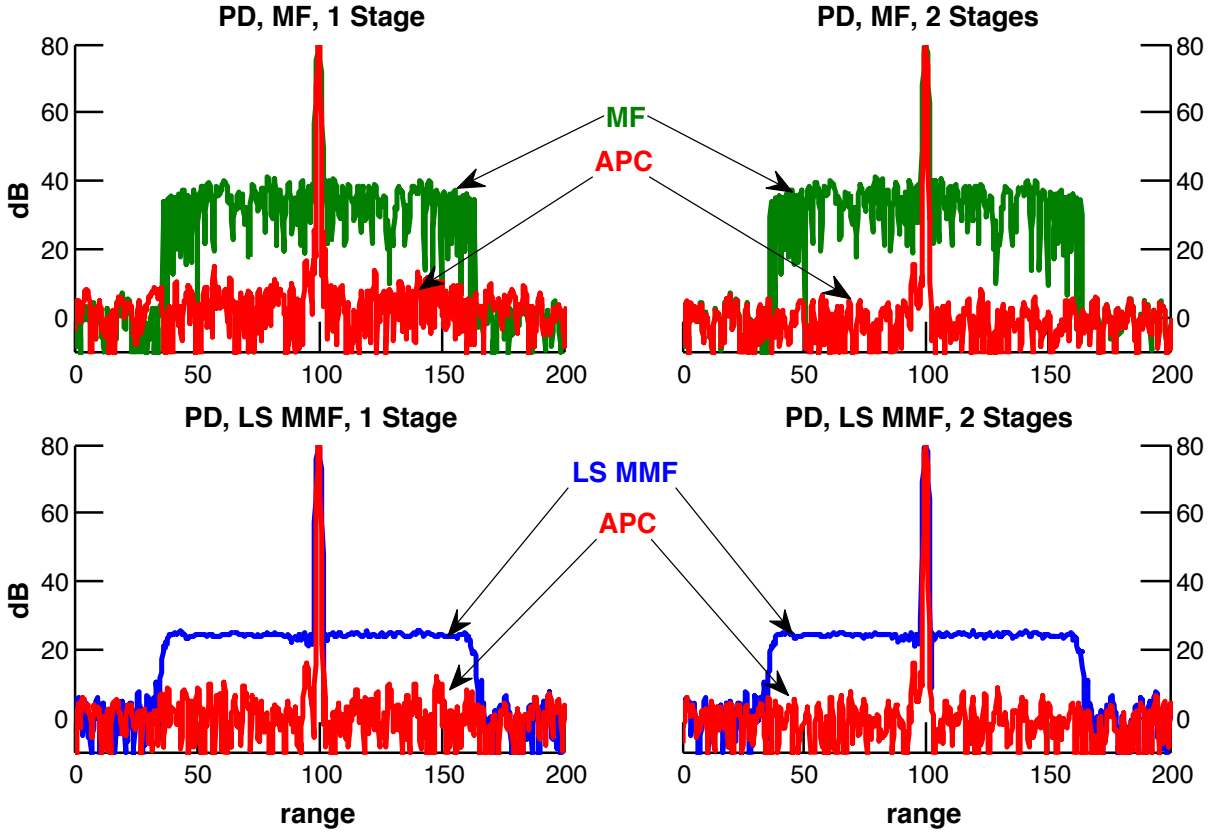


Figure 3.14: Matched filter, LS MMF and APC responses for PD in a two-target environment with worst-case straddling. The top two plots show the responses of the matched filter and APC initialized with the matched filter. The bottom two plots show the responses of the LS MMF and APC initialized with the LS MMF. Neither the matched filter nor the LS MMF can discern the smaller target. The smaller target can be found after one adaptive stage for both the matched filter and LS MMF initializations.

suppression. Neither the matched filter nor the LS MMF can discern the smaller targets, though the LS MMF distinguishes more targets than the matched filter.

After comparing the performance of APC using two different initialization filters (matched filter and LS MMF), it is apparent that the lower sidelobe levels of the LS MMF allow APC to converge quicker than it does when the matched filter is used in initialization. Thus, from this point forward, APC will only be implemented using the LS MMF as the initialization filter, unless otherwise noted. Referring to the performance of APC with regards to straddling, unlike the LS MMF, which realizes an overall increase in sidelobe levels due to range straddling, the adaptive filter coefficients, for a given range cell estimate, impose nulls at the relative delay offsets of

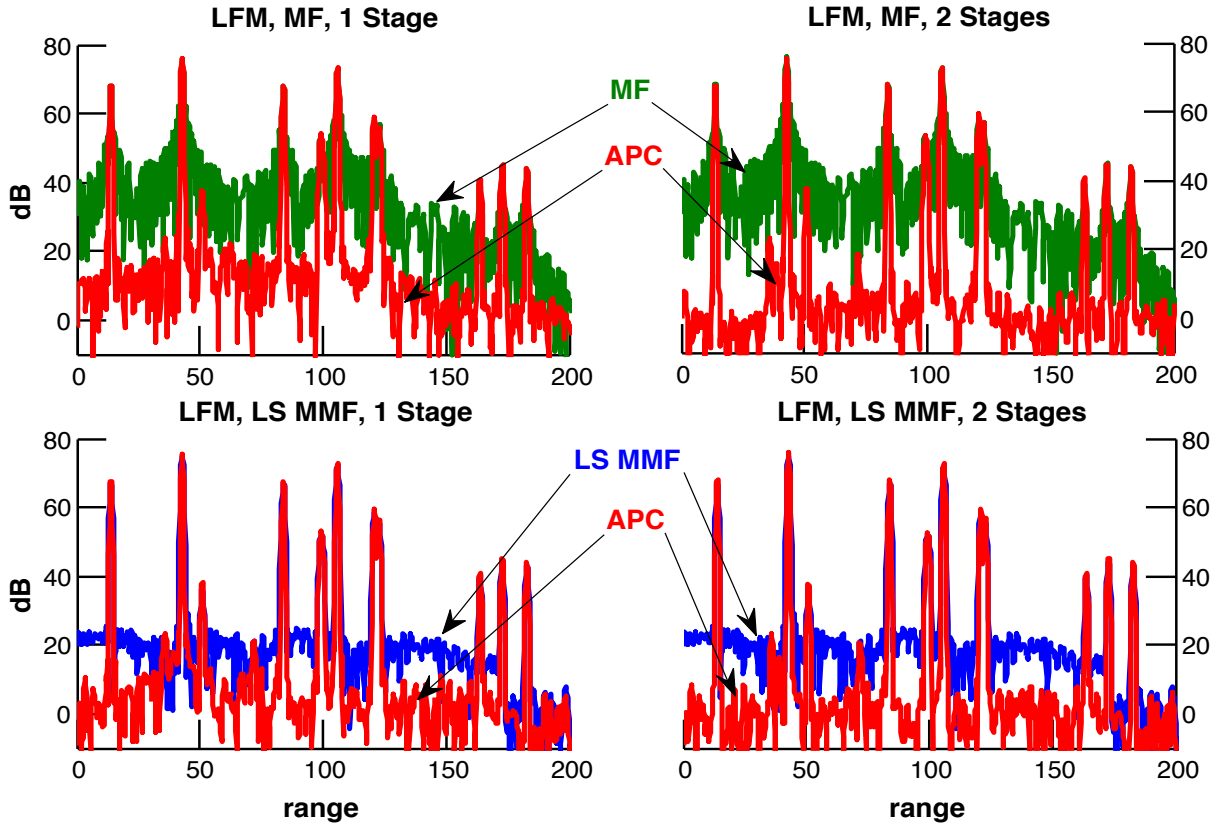


Figure 3.15: Matched filter, LS MMF and APC responses for LFM in a target-rich environment with worst-case straddling. APC initialized with the matched filter requires two adaptive stages to discern the smaller targets, while when APC is initialized with the LS MMF, only one adaptive stage is needed, which corroborates the observations for the two target scenario.

large scatterers. Such nulls are less affected by straddling mismatch because the null width is generally sufficient to still null the interfering scatterer despite the small mismatch [7]. Therefore, as observed when straddling effects were imposed on the APC responses, APC has robustness to the effects of straddling that the LS MMF does not.

The first derivation of APC was unable to estimate into the eclipsed regions, where some portion of the pulse width of the received waveform is unavailable. In practice, this circumstance occurs when reflections from targets are first received or at the end of the signal's reception, which is shown visually in Figure 3.17. For eclipsing repair, each eclipsed range cell is estimated using only a portion of the pulse and applying a gain constraint. When eclipsing repair is not used, the range window will collapse for successive adaptive stages. The equations needed for this eclipsing

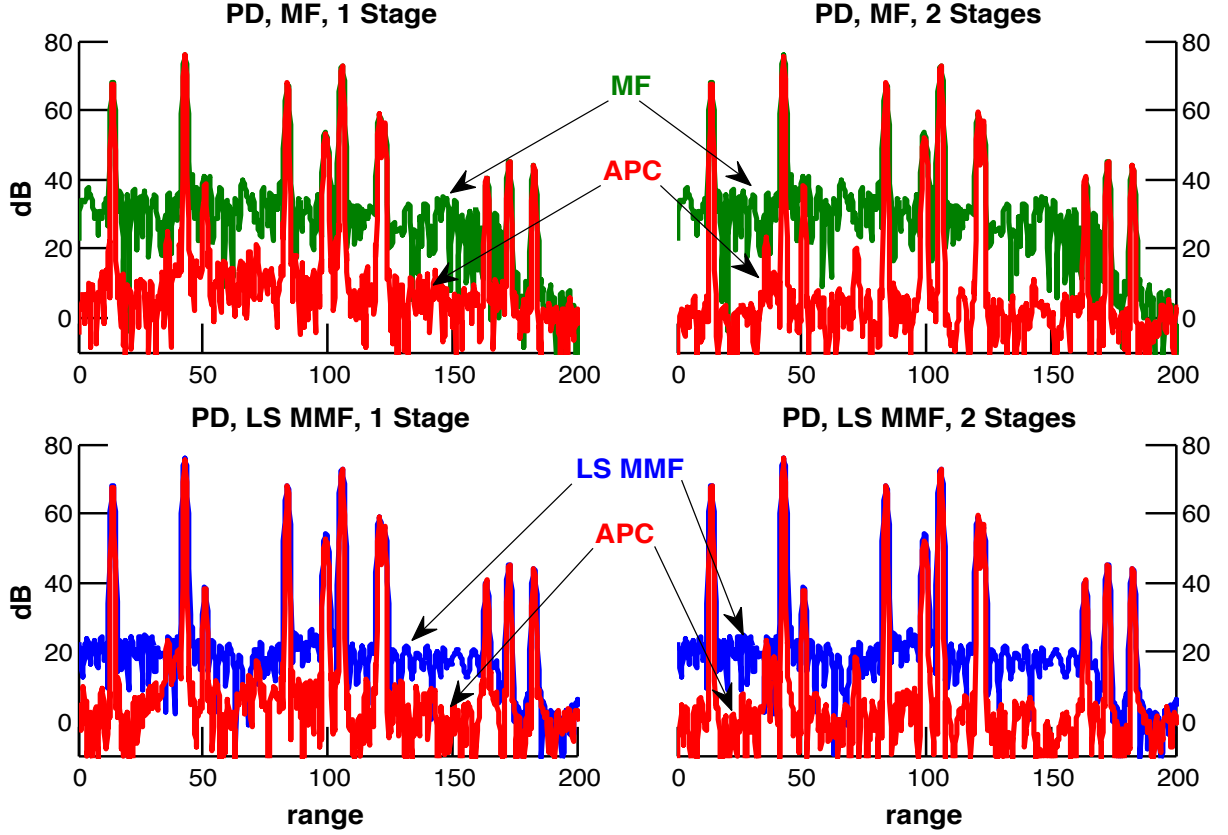


Figure 3.16: Matched filter, LS MMF and APC responses for PD in a target-rich environment with worst-case straddling. The matched filter initialization of APC is able to discern the smaller targets using two adaptive stages. In this LS MMF initialization of APC, only one adaptive stage is needed to discern the smaller targets.

repair are

$$\mathbf{w}_k(n = -q) = \frac{\|\mathbf{s}_{N,-q}\|}{\|\mathbf{s}_N\|} \frac{(\mathbf{C}_k^{(n)}(0) + \mathbf{R}_k)^{-1} \mathbf{s}_{k,-q}}{\sum_{i=0}^{K-1} \mathbf{s}_{i,-q}^H (\mathbf{C}_i^{(n)}(0) + \mathbf{R}_i)^{-1} \mathbf{s}_{i,-q}} \quad (3.3)$$

and

$$\mathbf{w}_k(n = Q - 1 + q) = \frac{\|\mathbf{s}_{N,q}\|}{\|\mathbf{s}_N\|} \frac{(\mathbf{C}_k^{(n)}(Q - 1) + \mathbf{R}_k)^{-1} \mathbf{s}_{k,q}}{\sum_{i=0}^{K-1} \mathbf{s}_{i,q}^H (\mathbf{C}_i^{(n)}(Q - 1) + \mathbf{R}_i)^{-1} \mathbf{s}_{i,q}}, \quad (3.4)$$

the development of which were described in [2, 6]. The adaptive filter coefficients in (3.3) and (3.4) are the same as (3.1), but for range cells in the eclipsed region. The differences in these equations are the indices Q , which is the total number of non-eclipsed range cells and q , which denotes the




Figure 3.17: Taken from [2], this figure visually depicts the eclipsing effect at both the beginning and end of the pulse repetition interval.

early or late eclipsed samples, where q goes from 1 to $N-1$ and N is the number of samples in a pulse width. Negative values of q refer to the "early" eclipsed samples and positive values of q refer to the "late" eclipsed samples. The ratio $\frac{\|s_{N,\pm q}\|}{\|s_N\|}$ is needed to prevent over-compensation from the unity gain constraint [56].

The previous worst-case straddling target rich scenario, seen in Figures 3.15 and 3.16, are now replicated with an additional five targets, which are placed in the eclipsed region. The performance of APC in these regions can be compared to that of the matched filter and LS MMF, as shown in Figures 3.18 and 3.19. In these figures, there is slight improvement in target detection when a second adaptive stage is used. In the eclipsed regions (-50 to 0 and 200 to 250), the sidelobe performance of the matched filter and LS MMF is significantly degraded, while APC does not appear to be significantly degraded due to eclipsing.




Figure 3.18: Matched filter, LS MMF and APC responses for LFM in a target-rich environment with worst-case straddling and eclipsing. The LS MMF was used to initialize APC. Unlike the matched filter and LS MMF, which experience an increase in sidelobes due to eclipsing, APC is relatively robust to eclipsing effects. The second iteration of APC yields some slight improvement in target detection, even in the eclipsed regions.

The last phenomenon examined with respect to APC is the effect of Doppler. For comparison, refer to the idealistic APC simulations in Figures 3.9 and 3.10. In the first two-target scenario, in Figure 3.20, the Doppler shift for the small target is zero, while a Doppler shift of 0.04 radians is applied to the larger target. If any increase in sidelobe levels occurs in APC, the smaller target may be masked. For both waveforms, the LS MMF and APC responses are able to distinguish the smaller target, while the matched filter cannot. The LS MMF response corresponding to PD does have slightly higher sidelobes than seen for the non-Doppler LS MMF response. APC is still able to suppress sidelobes below the levels observed in the LS MMF response. This capability of APC is covered in more detail in [43].




Figure 3.19: Matched filter, LS MMF and APC responses for PD in a target-rich environment with worst-case straddling and eclipsing. The LS MMF was used to initialize APC. Unlike the matched filter and LS MMF, which experience an increase in sidelobes due to eclipsing, APC is relatively robust to eclipsing effects. The second iteration of APC yields some slight improvement in target detection, even in the eclipsed regions.




Figure 3.20: Matched filter, LS MMF and APC responses for LFM and PD in a two-target environment. The larger target has a Doppler shift of 0.04 radians, while the smaller target has no Doppler shift. The LS MMF was used to initialize APC. For both waveforms, the LS MMF and APC are able to discern the smaller target, and APC further suppresses sidelobes below the levels in the LS MMF response.

Chapter 4

Pulse Compression of Free-Space Measurements

To verify the simulations seen in Chapters 2 and 3, which showed that the Least-Squares Mismatch Filter (LS MMF) and Adaptive Pulse Compression (APC) can be applied to FM waveforms, free-space measurements were gathered. These measurements were taken on the roof of Nichols Hall on the University of Kansas campus; the radar system in use was aimed towards two large buildings on campus. Since the signal processing methods discussed in this paper are all post-processing techniques, only one data set was needed for each waveform. The pulse compressed range responses can be generated using the LS MMFs and APC and the performance of these methods compared against that of the matched filter.

4.1 Test Setup

The data was gathered using a quasi-monostatic radar configuration. The term "quasi"-monostatic is used because there were two antennas (one transmit, one receive) located within 2 meters of each other. A Tektronix DPO70k oscilloscope was attached to the vertically polarized port of a dual-pole Vivaldi antenna (3-dB beamwidth of $\sim 45^\circ$), which was used to measure received waveforms. A Tektronix AWG70k arbitrary waveform generator was used to generate the desired waveforms,



Figure 4.1: Equipment used to record measurements. The radar is quasi-monostatic, with the right antenna used to transmit and the left to receive. The antennas are roughly 15 meters from the ground and are directed at two large buildings on the University of Kansas campus.

which were transmitted through an identical dual-pole Vivaldi antenna with vertical polarization. The test set up and field of view, from the radar system point of view, can be seen in Figure 4.1. The antenna on the right, used to transmit, and the antenna on the left, used to receive, are both shown clearly in the picture. Because separate antennas were used for transmit and receive, and because of their proximity to each other, the direct path is the largest received signal. Though both the horizontally and vertically polarized ports are shown in this picture, only the vertically polarized ports were used to collect data.

Data corresponding to the linear FM (LFM) waveform and the optimized FM waveform previously denoted as 'performance diversity' (PD) were gathered using this system. The transmitted signal had a center frequency of 2.3 GHz, a pulse width of 768 ns, a bandwidth of 83 MHz, and a

transmit power of 24 dBm. Data was captured for 1 second and a pulse repetition frequency (PRF) of 10 kHz was used. With a pulse width of 768 ns and a bandwidth of 83 MHz, a time-bandwidth product of approximately 64 ($768 \text{ ns} \times 83 \text{ MHz} = 63.77$) was achieved, which matched the time-bandwidth product used in simulation. With a PRF of 10 kHz and a total transmission time of 1 second, 10,000 pulses were transmitted (100 microseconds of data per pulse). The received data corresponding to these 10,000 pulses were coherently integrated before being processed using the LS MMFs and APC. No moving targets were within the field of view during data collection. Re-sampling of the data was necessary so that the sampling rate of the data would match that of the respective waveform used in filter construction.

The radar system was placed on the roof of Nichols Hall on the University of Kansas campus, roughly 15 meters above the ground. The antennas were aimed at two other buildings on campus, the Lied Center and the Dole Institute of Politics. Closer to the antennas in the field of view are several trees, signs, and light poles (shown in Figure 4.2) that are good reflectors; this is evident in the pulse compressed responses. Though the antennas are aimed towards the large buildings on campus, the return signals from the signs and trees are of greater interest, as these returns occur within the pulse width of the return from the direct path response and thus reside within the sidelobes of the matched filter response.

4.2 Pulse Compression Results

4.2.1 Least-Squares Mismatch Filter Responses

Figure 4.3 shows the pulse compressed responses of the matched filter, non-averaged least-squares mismatched filter (MMF label), filter averaged ($L = 5$) LS MMF, waveform averaged ($P = 5$) LS MMF, and combination averaged ($L = P = 5$) LS MMF. The LS MMFs are designed in the manner deemed optimal in Chapter 2 ($K = 5$, $M = 4$, 7/9 rows zeroed for LFM/PD).

The direct path response was set at a range of 0 meters. After the direct path, it is somewhat difficult to tell differences between the filtering methods, though it is clear that all of the MMFs



Figure 4.2: Radar field of view. The star near the bottom of the map gives the location of the antennas, and the targets that may produce significant reflections are denoted by triangles. The antennas are aimed at the Dole Institute of Politics, but given the beamwidth of the antennas, returns from the Lied Center are expected.




Figure 4.3: Pulse compressed responses for both the LFM and PD, using the matched filter, non-averaged LS MMF, and averaged LS MMFs. The LS MMFs have better performance with respect to sidelobe levels than the matched filter.

have better sidelobe suppression than the matched filter, and as the amount of averaging increases, sidelobe levels are further suppressed. Looking at the "negative range", that is, the time in which the receiver had been turned on but had not yet received the entire pulse from the transmitter via the direct path, performance differences between the MMFs are readily observed. As shown in simulation, the averaged MMFs (L and/or $P > 1$) have lower sidelobe levels than the non-averaged MMF. Additionally, as seen in simulation, the $L = P = 5$ MMF has lower sidelobes than both the $L = 5$ MMF and $P = 5$ MMF. These trends are clear for both the LFM and PD waveforms.

4.2.2 APC Responses to Measured Data

Using the LS MMF responses as a benchmark, the APC responses were generated for the same range data. With this measured free-space data, behaviors noticed in simulation are expected to be reciprocated after processing real data. From Chapter 3, recall that a second iterative stage of APC yields slightly better results than when just one stage is used. However, on this data, it appears that the APC algorithm converges after just one iteration, so multiple adaptive stages were not further explored.

Like what was done in simulation, the APC responses generated using both initialization filters were compared. As seen in Chapter 3, initializing with the LS MMF results in faster convergence (when multiple adaptive stages are needed) and lower sidelobes in the first adaptive stage than was observed with the matched filter initialization of APC. Provided in Figure 4.4 are the results for both LFM and PD with matched filter and LS MMF initialization. For both waveforms, the LS MMF initialized version of APC has slightly lower sidelobes than the matched filter initialized version, verifying what was seen in simulation. Because of this observation, the remaining APC responses in this chapter will use the LS MMF for initialization.

The last phenomenon that APC was able to account for in simulation was eclipsing. Since the direct path is the largest received signal, and as such causes the largest sidelobes, it is easier to observe eclipsing if some of the data at the beginning of the pulse repetition interval (PRI, which is the inverse of the PRF) is ignored rather than data at the end of the PRI. This situation is synonymous to the receive antenna switching on too late and missing some of the pulse arriving from the direct path. To replicate this scenario, data located within the first pulse width of the direct path return is omitted.

The performance of APC was examined when the first pulse is eclipsed 25%, 50% and 75% of the total duration. Shown in Figure 4.5 are the matched filter, LS MMF, and APC responses when 25% and 75% of the pulse is eclipsed. For 50% eclipsing, see Figure A.16 in Appendix A. When the direct path is eclipsed by 25%, all 3 pulse compression methods experience some degree of sidelobe degradation. The matched filter and LS MMF undergo a significant increase in




Figure 4.4: Pulse compressed response of APC for both waveforms, with matched filter initialization and LS MMF initialization. The LS MMF initialized version of APC does a better job of suppressing sidelobes after one adaptive stage.

sidelobe levels, and thus are less capable for target detection. The APC response, however, sees little degradation in performance when eclipsing occurs. When eclipsing is increased to 75%, the sidelobes of the matched filter and mismatched filter increase even further. The APC response still does not experience a significant degradation in sidelobe levels and has only slightly coarser resolution than when no eclipsing occurs. However, some performance degradation is observed, as APC can no longer discern the target that is located within the first 10 meters of the direct path.

4.2.3 Comparison of the Averaged LS MMFs and APC

Now that the feasibility of the LS MMF and APC on real data gathered using FM waveforms has been shown, these processing methods are compared in Figures 4.6 and 4.7. Figure 4.6 shows the

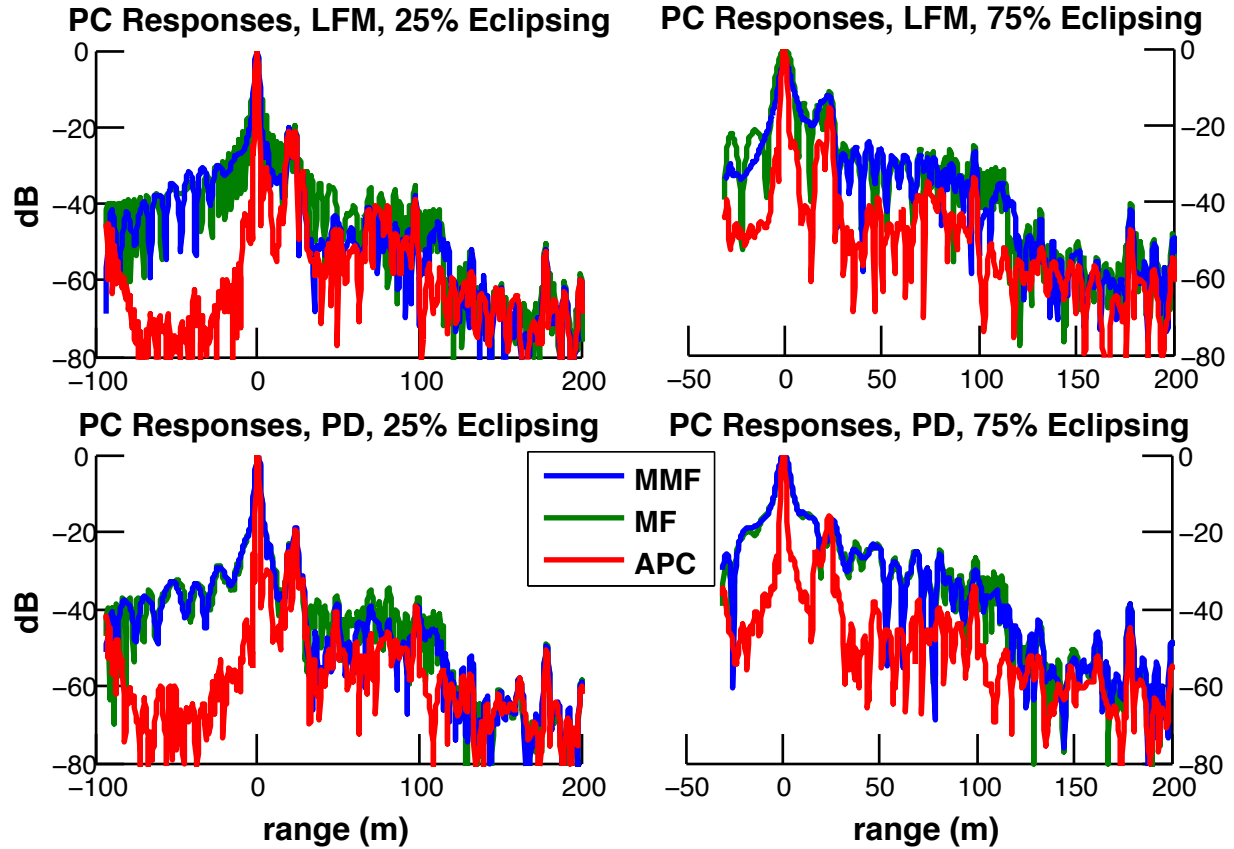


Figure 4.5: Matched filter, LS MMF, and APC pulse compressed responses for both waveforms with 25% and 75% eclipsing. The matched filter and LS MMF responses experience significant degradation in sidelobe levels compared to APC, making them less capable when eclipsing occurs.

range response when LFM is used, while Figure 4.7 shows the range response for PD. Each of these figures shows the processed results of the matched filter, non-averaged LS MMF, filter averaged ($L = 5$) LS MMF, waveform averaged ($P = 5$) LS MMF, combination averaged ($L = P = 5$) LS MMF, and APC out to a range of 500 meters.

The pulse compressed responses for APC were generated using one adaptive stage, with the LS MMF used in initialization. For both waveforms, all filtering methods have essentially the same performance beyond 125 meters. At 200 meters, the building facade (left side of the field of view) can be seen. Beyond the building, in range, the return from the hill can be seen (dead center of the field of view). The Lied Center (~ 275 m) and Dole Center (~ 350 m) are visible in the responses of each pulse compression method. However, within the first 125 meters and in negative range,




Figure 4.6: Pulse compressed responses with LFM for the matched filter, LS MMFs, and APC out to a range of 500 meters (and -100 meters). There is little difference in the filter responses beyond 125 meters.




Figure 4.7: Pulse compressed responses with PD for the matched filter, LS MMFs, and APC out to a range of 500 meters (and -100 meters). These pulse compression responses, when compared to those for LFM, have lower sidelobes and deeper nulls between targets.

differences in performance between the six pulse compression methods are observed. The pulse compression responses in Figure 4.7, when compared to those in Figure 4.6, have lower sidelobes and deeper nulls between targets. As a result, there appear to be more distinct targets present in this response than in the LFM responses.

Figures 4.8 and 4.9 present the same data as in Figures 4.6 and 4.7, but now the range -50 m to 150 m is examined, rather than from -100 m to 500 m . This change allows the differences between filters to be more visible. For both the LFM and PD responses, all LS MMFs and APC can discern more targets than the matched filter. The non-averaged LS MMF is less capable at target discrimination than the averaged ones. The combination averaged LS MMF has better sidelobe suppression than the filter or waveform averaged LS MMF. When comparing the combination averaged LS MMF to APC, it is difficult to tell which pulse compression method is the best for sidelobe suppression. In the negative range region, APC has slightly higher sidelobe levels than the combination averaged LS MMF. APC does not suppress nulls as well as the LS MMFs, noticeable at the null at ~ 10 meters. However, when looking out slightly past 100 meters, APC appears to do better with peak discrimination. This may be because the LS MMFs have spurious peaks that occur at the ends of the pulse width, which APC is able to suppress.

Comparing the simulations from Chapters 2 and 3 to the measured results in this section, it appears that APC is more capable than the LS MMFs when there are multiple targets with large SNRs and/or significant dynamic range. APC appears to be able to suppress nulls between closely spaced targets better than the non-adaptive methods that were analyzed. Compared to the LFM pulse compression responses, the PD responses in Figure 4.9 have lower sidelobes, so they are able to uncover/reveal targets with more certainty, especially in the "signs" region.

4.2.4 Effects of Super-Resolution on Measured Responses

The last pulse compressed responses examined were those for the LSMMF when the number of zeroed rows in matrix $\tilde{\mathbf{A}}$ is varied. (See Section 2.2, for an explanation of this process). By varying the number of zeroed rows, the range resolution of the LS MMF response can be controlled. More




Figure 4.8: A zoomed-in version of Figure 4.6, focused on the first 150 meters (and -50 meters) so differences between the pulse compression methods can be observed. The LS MMFs and APC are able to discern more targets than the matched filter.

zeroed rows yields coarser resolution, while fewer rows results in finer resolution. The LFM responses in Figure 4.10 are generated when matrix $\tilde{\mathbf{A}}$ has 3, 5, 7, and 9 rows zeroed above and below the m^{th} row, with 7 zeroed rows corresponding to the optimal LS MMF response, which has the least mismatch loss (MML), defined in (1.3). Zeroing fewer than 7 rows results in an LS MMF response with finer resolution. Likewise, zeroing out more rows results in a response with coarser resolution.

When 3 rows are zeroed, the LS MMF response has the finest resolution of the cases that were examined. This situation, however, results in MML that is roughly 8 dB worse than the optimal LS MMF response. The LS MMF responses in Figure 4.10 are normalized with respect to the peak of the LS MMF response with 7 zeroed rows. With resolution this fine, it appears an additional target




Figure 4.9: A zoomed-in version of Figure 4.7, focused on the first 150 meters (and -50 meters) so differences between the pulse compression methods can be observed. The LS MMFs and APC are able to discern more targets than the matched filter. These responses, when compared to the LFM pulse compression responses, demonstrate better sidelobe suppression and target detection.

can be seen within 5 meters of the direct path, as well as several others in target rich areas. Nulls between targets are deeper as well, which results from the finer resolution. When the number of zeroed rows is increased to 5, the resolution is coarser than it was with 3 rows, but MML is reduced by nearly 6 dB. If more than the optimal number of zeroed rows (7) is used, resolution becomes coarser and the nulls between targets become shallower. Zeroing out rows beyond the optimal amount, however, does slightly suppress sidelobes, an observation that is visible in Figure 4.10 beyond 100 meters.

The PD responses in Figure 4.11 are generated for 3, 6, 9, and 12 zeroed rows. Zeroing 9 rows yields the optimal LS MMF response with respect to MML and the LS MMF responses in this figure are normalized with respect to the peak of the optimal LS MMF response. When 3 rows are




Figure 4.10: Pulse compressed responses of the LS MMF with a varying number of zeroed rows. For LFM, 7 zeroed rows results in the optimal LS MMF response with respect to MML. Zeroing out fewer rows results in finer resolution, deeper nulls, and significant (> 1 dB) MML. Zeroing out more than the optimal number of rows results in coarser resolution, shallower nulls, and lower sidelobe levels.

zeroed, the resolution is the finest, but it results in MML that is roughly 8 dB worse than the optimal LS MMF response. With resolution this fine, additional targets become apparent (e.g., within the first 5 meters), and nulls between targets are deeper than in the optimal LS MMF response. With 6 rows zeroed, the resolution is coarser than with 3 rows, but MML is reduced by nearly 7 dB. This response has slightly better target recognition than the optimal LS MMF response, but it also has MML that is approximately 1 dB worse. Going beyond the optimal number of zeroed rows, resolution gets coarser and the nulls between targets get smaller. This trend is most evident in the targets located around 20 meters. Zeroing out rows beyond the optimal amount, however, does slightly suppress sidelobes, an observation visible in this figure beyond 100 meters.




Figure 4.11: Pulse compressed responses of the LS MMF with a varying number of zeroed rows. For PD, 9 zeroed rows results in the optimal LS MMF response with respect to MML. Zeroing out fewer rows results in finer resolution, deeper nulls, and significant (> 1 dB) MML. Zeroing out more than the optimal number of rows results in coarser resolution, shallower nulls, and lower sidelobe levels.

The following discussion applies to both the responses with the LFM and PD, but can be best observed with the PD responses, especially at 20 meters in range. With only 3 zeroed rows, a noticeable improvement in resolution can be seen, as targets that are not apparent in the other LS MMF responses are now visible. The finer resolution also increases the nulls between targets, such as those at 5 meters and 20 meters. However, this improvement in resolution results in a sizable amount of MML, nearly 8 dB, in both waveform responses for the direct path peak. This loss in SNR is evident in the entire LS MMF response. When the number of zeroed rows is increased to five for LFM and six for PD, the resulting MML is less than 2 dB. Though the resolution is not as narrow as with 3 zeroed rows, there is some improvement in resolution, again referring to the LS

MMF response for PD for the targets at 20 meters. When the number of zeroed rows is increased beyond that which minimizes MML, the resolution is coarser. These LS MMF responses (9 zeroed rows for LFM, 12 for PD) have a "smoother" response, i.e., they have relatively shallow nulls and are thus not as able to discern close targets.

Chapter 5

Conclusions

This work has shown the feasibility of both the least-squares mismatch filter (LS MMF) and adaptive pulse compression (APC) for applications involving frequency-modulated (FM) waveforms. FM waveforms are examined because they are spectrally efficient and power efficient, can be Doppler tolerant, and are physically realizable. The initial derivations of LS MMF and APC [29, 30] only consider polyphase coded waveforms for use in these pulse compression methods. This work has shown through simulation and measured data that the LS MMF and APC methods are compatible with FM waveforms, provided some adjustments to account for the difference between polyphase codes and FM waveforms. For the LS MMF, oversampling relative to the 3 dB bandwidth and beamspoiling in the delayed signal matrix are required to yield desirable results with the LS MMF responses. Likewise, for APC, oversampling relative to the 3 dB bandwidth and beamspoiling in the signal correlation matrix are needed.

In addition to being applied to FM waveforms, the LS MMF was also designed such that it is robust to straddling effects, improves sidelobe levels relative to the matched filter, and yields acceptable mismatch loss (MML) on the order of tenths of a dB. An averaging process was developed to improve the performance of the LS MMF when straddling effects and, to a lesser degree, Doppler effects are present. Furthermore, the super-resolution capabilities of LS MMF were examined and shown to provide noticeable improvement in range resolution in exchange for modest

mismatch loss. With respect to straddling effects, the averaging process used with the LS MMF was found to improve peak sidelobe levels (PSLs) by 3-5 dB relative to that of the non-averaged LS MMF. Furthermore, with the averaged LS MMF response, spurious peaks are essentially eliminated. If the target's direction of movement relative to the radar is assumed, an improvement in PSL of the Doppler-averaged LS MMF is realized as well. These improvements for the LS MMF responses are particularly significant because the only additional computational cost relative to the matched filter is offline design time. For use in a real-time system, the LS MMF essentially realizes no increased computational cost relative to the matched filter, which means the new LS MMF can be deployed on existing systems with minimal adjustments.

This work also analyzed, via simulation and measured data, recent modifications that have been made to adaptive pulse compression (APC) that made it applicable to FM waveforms. One main concern with iterative pulse compression algorithms such as APC is computational cost. While APC has been improved by [54, 55], more than one iterative stage is still usually required for desired sidelobe suppression. By initializing APC with the LS MMF in lieu of the matched filter, it was shown in this paper that one iterative stage can provide significant sidelobe suppression in a target-rich environment (at least from those seen in our simulations and measured data). Additionally, adjustments have previously been made for APC that enable it to estimate into the eclipsed regions [2], which here were shown to be valid for FM waveforms via simulations and measured results. In the eclipsed regions, APC suffers little performance degradation relative to the LS MMF or matched filter.

With the measured results, the practical application of LS MMF and APC using FM waveforms was shown to be viable. This result is significant because optimized FM waveforms (e.g. designed through the continuous phase modulation (CPM) framework) from [27, 28] can be implemented in a high-power radar system, and using the LS MMF and APC pulse compression methods significant sidelobe reduction can be realized. These optimized FM waveforms are of significant interest in the radar community because they are well-contained spectrally and have low auto-correlation sidelobes. Spectral containment is becoming more and more necessary for radar waveforms, as the

spectrum is becoming more crowded, especially due to growing consumer market needs [3].

5.1 Future Work

There are several avenues of future work stemming from this research. One item of interest is adjustments to the LS MMF structure to eliminate spurious peaks in the idealistic LS MMF responses. Spurious peaks, as seen in Figures 2.7 and 2.8, can sometimes lead to false target detection. Though not fully explored, there has been promising preliminary work regarding the elimination of these undesired peaks. Another idea that has arisen from this work is the possibility of "fusing" multiple LS MMFs such that the super-resolution LS MMF responses can be used in target-rich environments (where target discrimination is needed). LS MMF responses with great sidelobe level suppression can then be used when target scenarios with high dynamic range are examined. Finally, the work here was empirical, that is, the improvements realized were obtained through simulation and observations. What could next be explored is a parametrized mathematical model that yields an optimizable cost function that would enable an optimizable result to be calculated based on desired LS MMF response criteria.

References

- [1] A. W. Doerry, “Generating nonlinear FM chirp waveforms for radar,” Sandia National Laboratories, Albuquerque, NM, Report SAND2006-5856, September 2006.
- [2] S. Blunt, K. Gerlach, and E. Mokole, “Pulse compression eclipsing-repair,” in *IEEE Radar Conference*, May 2008, pp. 1–5.
- [3] H. Griffiths, L. Cohen, S. Watts, E. Mokole, C. Baker, M. Wicks, and S. Blunt, “Radar spectrum engineering and management: technical and regulatory issues,” *Proceedings of the IEEE*, vol. 103, no. 1, pp. 85–102, Jan 2015.
- [4] H. Griffiths, S. Blunt, L. Cohen, and L. Savy, “Challenge problems in spectrum engineering and waveform diversity,” in *IEEE Radar Conference*, April 2013, pp. 1–5.
- [5] M. Wicks, E. Mokole, S. Blunt, R. Schneible, and V. Amuso, *Principles of Waveform Diversity & Design*. SciTech Publishing, 2010.
- [6] F. Gini, A. de Maio, and L. Patton, *Waveform Design and Diversity for Advanced Radar Systems*. IET, 2012.
- [7] D. Henke, P. McCormick, S. Blunt, and T. Higgins, “Practical aspects of optimal mismatch filtering and adaptive pulse compression for FM waveforms,” in *IEEE International Radar Conference*, May 2015, pp. 1149–1155.
- [8] M. Richards, J. Scheer, and W. Holm, *Principles of Modern Radar: Basic Principles*. SciTech Publishing, 2010.

- [9] N. Levanon and E. Mozeson, *Radar Signals*. Wiley – IEEE Press, 2004.
- [10] C. Nunn and G. Coxson, “Polyphase pulse compression codes with optimal peak and integrated sidelobes,” *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 45, no. 2, pp. 775–781, April 2009.
- [11] J. Jenshak and J. Stiles, “A fast method for designing optimal transmit codes for radar,” in *IEEE Radar Conference*, May 2008, pp. 1–6.
- [12] G. Coxson and J. Russo, “Efficient exhaustive search for optimal-peak-sidelobe binary codes,” in *IEEE Radar Conference*, April 2004, pp. 438–443.
- [13] B. Lewis and F. Kretschmer, “Linear frequency modulation derived polyphase pulse compression codes,” *IEEE Trans. Aerospace and Electronic Systems*, vol. AES-18, no. 5, pp. 637–641, Sept 1982.
- [14] P. Tan, J. Jakobosky, J. Stiles, and S. Blunt, “On higher-order representations of polyphase-coded FM radar waveforms,” in *IEEE International Radar Conference*, May 2015, pp. 467–472.
- [15] J. Jakobosky, S. Blunt, and T. Higgins, “Ultra-low sidelobe waveform design via spectral shaping and LINC transmit architecture,” in *IEEE Radar Conference*, May 2015, pp. 1021–1026.
- [16] J. Jakobosky, S. Blunt, and B. Himed, “Optimization of "over-coded"; radar waveforms,” in *Radar Conference, 2014 IEEE*, May 2014, pp. 1460–1465.
- [17] J. Jakobosky, S. Blunt, M. Cook, J. Stiles, and S. Seguin, “Transmitter-in-the-loop optimization of physical radar emissions,” in *IEEE Radar Conference*, May 2012, pp. 0874–0879.
- [18] J. Jakobosky, P. Anglin, M. Cook, S. Blunt, and J. Stiles, “Non-linear FM waveform design using marginal Fisher’s information within the CPM framework,” in *IEEE Radar Conference*, May 2011, pp. 513–518.

- [19] S. Blunt, M. Cook, E. Perrins, and J. de Graaf, "Cpm-based radar waveforms for efficiently bandlimiting a transmitted spectrum," in *IEEE Radar Conference*, May 2009, pp. 1–6.
- [20] P. McCormick, J. Jakabosky, S. Blunt, C. Allen, and B. Himed, "Joint polarization/waveform design and adaptive receive processing," in *IEEE International Radar Conference*, May 2015, pp. 1382–1387.
- [21] L. Ryan, J. Jakabosky, S. Blunt, C. Allen, and L. Cohen, "Optimizing polyphase-coded FM waveforms within a LINC transmit architecture," in *IEEE Radar Conference*, May 2014, pp. 0835–0839.
- [22] J. Kurdzo, B. Cheong, R. Palmer, and G. Zhang, "Optimized NLFM pulse compression waveforms for high-sensitivity radar observations," in *IEEE International Radar Conference*, Oct 2014, pp. 1–6.
- [23] T. Collins and P. Atkins, "Nonlinear frequency modulation chirps for active sonar," *IET Radar, Sonar, and Navigation*, vol. 146, no. 6, pp. 312–316, Dec 1999.
- [24] C. Cook, "A class of nonlinear FM pulse compression signals," *Proceedings of the IEEE*, vol. 52, no. 11, pp. 1369–1371, Nov 1964.
- [25] J. Johnston and A. Fairhead, "Waveform design and doppler sensitivity analysis for nonlinear fm chirp pulses," *Communications, Radar and Signal Processing, IEE Proceedings F*, vol. 133, no. 2, pp. 163–175, April 1986.
- [26] I. Gladkova, "Design of frequency modulated waveforms via the zak transform," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 40, no. 1, pp. 355–359, Jan 2004.
- [27] S. Blunt, M. Cook, J. Jakabosky, J. De Graaf, and E. Perrins, "Polyphase-coded FM (PCFM) radar waveforms, part I: implementation," *IEEE Trans. Aerospace and Electronic Systems*, vol. 50, no. 3, pp. 2218–2229, July 2014.

- [28] S. Blunt, J. Jakabosky, M. Cook, J. Stiles, S. Seguin, and E. Mokole, “Polyphase-coded FM (PCFM) radar waveforms, part II: optimization,” *IEEE Trans. Aerospace and Electronic Systems*, vol. 50, no. 3, pp. 2230–2241, July 2014.
- [29] M. H. Ackroyd and F. Ghani, “Optimum mismatched filters for sidelobe suppression,” *IEEE Trans. Aerospace and Electronic Systems*, vol. AES-9, no. 2, pp. 214–218, March 1973.
- [30] S. Blunt and K. Gerlach, “Adaptive pulse compression,” in *IEEE Radar Conference*, 2004.
- [31] S. Blunt and K. Gerlach, “Adaptive pulse compression via MMSE estimation,” *IEEE Trans. Aerospace and Electronic Systems*, vol. 42, no. 2, pp. 572–584, April 2006.
- [32] S. D. Blunt and K. Gerlach, “A novel pulse compression scheme based on minimum mean-square error reiteration,” in *IEEE International Radar Conference*, 2003.
- [33] A. Shackelford, J. de Graaf, S. Talapatra, S. Blunt, and K. Gerlach, “Adaptive pulse compression: preliminary experimental measurements,” in *IEEE Radar Conference*, April 2007, pp. 234–237.
- [34] S. Haykin, *Adaptive Filter Theory*. Pearson Education, Inc., 2014.
- [35] S. Blunt and K. Gerlach, “Aspects of multistatic adaptive pulse compression,” in *IEEE International Radar Conference*, May 2005, pp. 104–108.
- [36] S. Blunt and K. Gerlach, “A generalized formulation for adaptive pulse compression of multistatic radar,” in *IEEE Workshop on Sensor Array and Multichannel Processing*, July 2006, pp. 349–353.
- [37] S. Blunt and K. Gerlach, “Multistatic adaptive pulse compression,” *IEEE Trans. Aerospace and Electronic Systems*, vol. 42, no. 3, pp. 891–903, July 2006.
- [38] S. Blunt and K. Gerlach, “Joint adaptive pulse compression to enable multi-static radar,” in *1st International Waveform Diversity and Design Conference*, 2004.

- [39] K. Gerlach, A. Shackelford, and S. Blunt, "Combined multistatic adaptive pulse compression and adaptive beamforming for shared-spectrum radar," *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 1, pp. 137–146, June 2007.
- [40] S. Blunt, W. Dower, and K. Gerlach, "Hybrid adaptive receive processing for multistatic radar," in *IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, Dec 2007, pp. 5–8.
- [41] S. Blunt, W. Dower, and K. Gerlach, "Hybrid interference suppression for multistatic radar," *IET Radar, Sonar, and Navigation*, vol. 2, no. 5, pp. 232–333, October 2008.
- [42] K. Gerlach, A. Shackelford, and S. Blunt, "A novel approach to shared-spectrum multistatic radar," in *IEEE Radar Conference*, April 2006.
- [43] S. Blunt, K. Smith, and K. Gerlach, "Doppler-compensated adaptive pulse compression," in *IEEE Radar Conference*, April 2006.
- [44] S. Blunt, A. Shackelford, K. Gerlach, and K. Smith, "Doppler compensation & single pulse imaging using adaptive pulse compression," *IEEE Trans. Aerospace and Electronic Systems*, vol. 45, no. 2, pp. 647–659, April 2009.
- [45] T. Cuprak and K. Wage, "Doppler compensated adaptive pulse compression through covariance matrix tapers," 2015, submitted to *IEEE Transactions on Aerospace and Electronic Systems*.
- [46] S. Blunt, A. Shackelford, and K. Gerlach, "Single pulse imaging," in *International Waveform Diversity and Design Conference*, 2006.
- [47] S. Blunt and K. Gerlach, "Adaptive repair of pulse-compressed radar waveforms: seeing the trees despite the forest," in *International Waveform Diversity and Design Conf.*, 2004.
- [48] S. Blunt and K. Gerlach, "Adaptive pulse compression repair processing," in *IEEE International Radar Conference*, May 2005, pp. 519–523.

- [49] K. Gerlach and S. Blunt, “Radar pulse compression repair,” *IEEE Trans. Aerospace and Electronic Systems*, vol. 43, no. 3, pp. 1188–1195, July 2007.
- [50] S. Blunt, K. Gerlach, and T. Higgins, “Aspects of radar range super-resolution,” in *IEEE Radar Conference*, April 2007, pp. 683–687.
- [51] A. Klein and M. Fujita, “Detection performance of hard-limited phase-coded signals,” *IEEE Trans. Aerospace and Electronic Systems*, vol. AES-15, no. 6, pp. 795–802, Nov 1979.
- [52] J. Proakis and D. Manolakis, *Digital Signal Processing*, ser. Prentice Hall international editions. Pearson Prentice Hall, 2007.
- [53] Y. Abramovich and G. San Antonio, “Range sidelobe control for LFM CW waveforms in OTHR with direct digital receivers,” in *IEEE International Radar Conference*, May 2015, pp. 1088–1092.
- [54] S. Blunt and T. Higgins, “Achieving real-time efficiency for adaptive radar pulse compression,” in *Radar Conference, 2007 IEEE*, April 2007, pp. 116–121.
- [55] S. Blunt and T. Higgins, “Dimensionality reduction techniques for efficient adaptive pulse compression,” *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 46, no. 1, pp. 349–362, Jan 2010.
- [56] T. Higgins, S. Blunt, and K. Gerlach, “Gain-constrained adaptive pulse compression via an MVDR framework,” in *Radar Conference, 2009 IEEE*, May 2009, pp. 1–6.

Appendix A

Appendix

Figures A.1 through A.4 provide the mismatch loss (MML) and peak sidelobe level (PSL) for the LS MMF response for each of the 40 subshifted versions of the received waveform used to observe straddling effects. The largest PSL for each value of L or P is deemed the worst-case PSL. The LS MMFs were constructed with an oversampling factor $K = 5$, filter length factor $M = 4$, and the optimal number of rows in $\tilde{\mathbf{A}}$ (7 for LFM and 9 for PD). These MML and PSL plots are provided to supplement the discussion in Sections 2.4.1 and 2.4.2 and to show where the values used in Tables 2.7 and 2.8 originate.

In these figures, notice how the worst-case MML and PSL values now occur for straddled waveforms at the ends of the straddling interval and the best-case MML and PSL values occur near $0.5 T_S$, which is the opposite of what was observed for the non-averaged LS MMFs. The worst-case and best-case values occur at these straddling intervals because the averaged LS MMFs are constructed by using waveforms that are linearly spaced through the straddling interval, which is centered at $0.5 T_S$. As such, the best-case is no longer at the beginning of the straddling interval, but it occurs in the middle. As L or P increases, the worst-case sidelobe level improves, but the best-case sidelobe level gets worse. This result of averaging is acceptable because a robust LS MMF was desired, meaning that the improvement of the worst-case results was given priority over improving the best-case results.




Figure A.1: MMLs and PSLs of the filter averaged LS MMF response for each straddled subshift of the received waveform. Generated for LFM and L varied from 1 to 5.

Figures A.5 through A.8 provide the PSLs for the LS MMF response for each of the 40 subshifted versions of the waveform used to observe straddling effects. The LF MMFs used to generate these responses were constructed using the filter-waveform or combination averaged method. Again, as was noticed in the PSL plots in Figures A.1 through A.4, the worst-case now corresponds to straddling shifts near the ends of the straddling interval. The largest PSL for each combination of L and P is deemed the worst-case PSL. These values are provided in Tables 2.9 and 2.10 for all examined combinations of L and P . These are provided to supplement the discussion in Section 2.4.3.

The behavior of the PSL plots generated from combination averaging is similar to those observed when filter and waveform averaging were implemented independently of each other.

Tables A.1 through A.4 provide the worst-case and averaged MML and PSL for Figures A.5




Figure A.2: MMLs and PSLs of the filter averaged LS MMF response for each straddled subshift of the received waveform. Generated for PD and L varied from 1 to 5.

through A.8. These MML and PSL values are generated from the LS MMFs designed with L varied from 1 to 5 and with P varied from 2 to 5 for both the LFM and PD waveforms. The averaged MML and PSL values are generated by averaging the MML and PSL from the 40 straddled LS MMF responses.

In these four tables, the worst-case PSL improves monotonically for both waveforms as L and P increase, though with diminishing improvements as L and P get larger. Worst-case and averaged MML remain relatively constant for both waveforms as L and P increase. The average PSL degrades as the degree of averaging increases, contrasting the results seen for the worst-case PSL. The reason that the average PSL does not improve, even though worst-case improves, is that the degradation in the "better" PSLs (near the best-case straddled LS MMF) is greater than the improvement realized in the "worse" PSLs (near the worst-case straddled LS MMF).




Figure A.3: MMLs and PSLs of the waveform averaged LS MMF response for each straddled subshift of the received waveform. Generated for LFM and P varied from 1 to 5.

Table A.1: Worst-case and averaged PSL and MML of the filter averaged LS MMF responses for both waveforms. The averaged PSL and MML correspond to the average value of the PSL and MML across all 40 straddled LS MMF responses. Generated using $K = 5$, $M = 4$ and $BT = 64$, with $P = 2$ and L varied from 1 to 5.

$P = 2$	<i>LFM</i>				<i>PD</i>			
L	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)
1	-49.9	0.86	-58.7	0.78	-52.3	0.18	-61.8	0.13
2	-51.2	0.86	-58.6	0.79	-53.6	0.18	-61.7	0.13
3	-51.7	0.86	-58.5	0.79	-54.1	0.19	-61.6	0.13
4	-52.0	0.86	-58.5	0.79	-54.4	0.18	-61.6	0.14
5	-52.1	0.86	-58.4	0.79	-54.5	0.19	-61.5	0.13




Figure A.4: MMLs and PSLs of the waveform averaged LS MMF response for each straddled subshift of the received waveform. Generated for PD and P varied from 1 to 5.

Table A.2: Worst-case and averaged PSL and MML of the filter averaged LS MMF responses for both waveforms. The averaged PSL and MML correspond to the average value of the PSL and MML across all 40 straddled LS MMF responses. Generated using $K = 5$, $M = 4$ and $BT = 64$, with $P = 3$ and L varied from 1 to 5.

$P = 3$	<i>LFM</i>				<i>PD</i>			
L	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)
1	-50.8	0.86	-59.1	0.78	-53.2	0.18	-62.2	0.13
2	-51.7	0.86	-58.7	0.79	-54.1	0.19	-61.8	0.13
3	-52.1	0.86	-58.5	0.79	-54.5	0.19	-61.6	0.13
4	-52.2	0.86	-58.5	0.79	-54.7	0.18	-61.6	0.13
5	-52.4	0.86	-58.5	0.79	-54.8	0.18	-61.5	0.13




Figure A.5: PSLs of the LS MMF response for each straddled subshift of the received waveform. Generated for LFM, $P = 2$ and $P = 3$, and L varied from 1 to 5.

Table A.3: Worst-case and averaged PSL and MML of the filter averaged LS MMF responses for both waveforms. The averaged PSL and MML correspond to the average value of the PSL and MML across all 40 straddled LS MMF responses. Generated using $K = 5$, $M = 4$ and $BT = 64$, with $P = 4$ and L varied from 1 to 5.

$P = 4$	LFM				PD			
L	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)
1	-51.3	0.86	-59.3	0.78	-53.6	0.18	-62.3	0.13
2	-52.0	0.86	-58.7	0.79	-54.4	0.18	-61.8	0.13
3	-52.2	0.86	-58.5	0.79	-54.6	0.18	-61.6	0.13
4	-52.4	0.86	-58.5	0.79	-54.8	0.18	-61.6	0.13
5	-52.5	0.86	-58.5	0.79	-54.9	0.19	-61.5	0.13




Figure A.6: PSLs of the LS MMF response for each straddled subshift of the received waveform. Generated for LFM, $P = 4$ and $P = 5$, and L varied from 1 to 5.

Table A.4: Worst-case and averaged PSL and MML of the filter averaged LS MMF responses for both waveforms. The averaged PSL and MML correspond to the average value of the PSL and MML across all 40 straddled LS MMF responses. Generated using $K = 5$, $M = 4$ and $BT = 64$, with $P = 5$ and L varied from 1 to 5.

$P = 5$	LFM				PD			
L	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)	Worst PSL (dB)	Worst MML (dB)	Avg. PSL (dB)	Avg MML (dB)
1	-51.6	0.86	-59.3	0.78	-53.9	0.18	-62.4	0.13
2	-52.4	0.86	-58.7	0.79	-54.5	0.19	-61.9	0.13
3	-52.4	0.86	-58.6	0.79	-54.8	0.18	-61.6	0.13
4	-52.5	0.86	-58.5	0.79	-54.9	0.19	-61.6	0.13
5	-52.5	0.86	-58.5	0.79	-54.9	0.19	-61.5	0.13




Figure A.7: PSLs of the LS MMF response for each straddled subshift of the received waveform. Generated for PD, $P = 2$ and $P = 3$, and L varied from 1 to 5.

The range-Doppler plots in Figure A.9 are used to supplement the discussion in Section 2.6 for the first Doppler averaging method (when no assumptions are made about the target's direction of travel). This figure displays the range-Doppler plots corresponding to the Doppler averaging LS MMF when averaging was done linearly between -0.04 and $+0.04$ with $J = 3$ and $J = 4$.

When more LS MMFs are averaged together, going from $J = 3$ to $J = 4$, the sidelobe levels actually get worse. This is most noticeable near the zero Doppler cut of the plot, as the width of the blue section, denoting lower sidelobes, gets narrower. Sidelobe levels further away from zero Doppler get larger as J is increased from 3 to 4.

The range-Doppler plots in Figures A.10 and A.11 are used to supplement the discussion in Section 2.6 for the second Doppler averaging method (the target's direction of travel is assumed to be known). These figures display the range-Doppler plots corresponding to the Doppler averaging




Figure A.8: PSLs of the LS MMF response for each straddled subshift of the received waveform. Generated for PD, $P = 4$ and $P = 5$, and L varied from 1 to 5.

LS MMF when averaging was done linearly between -0.04 and 0 with $J = 2$, $J = 3$, and $J = 4$ and for both LFM and PD.

In the one-sided Doppler averaged range-Doppler plots, for both LFM and PD, it was assumed that the target was receding relative to the radar system. A receding target corresponds to a negative Doppler shift, and as a result, averaging was done linearly between 0 and -0.04 . When compared to the base range-Doppler plot in Figure 2.18, the Doppler averaged LS MMF has better sidelobe performance for negative Doppler shifts. Furthermore, it appears that sidelobe levels, for negative Doppler values, do improve as J increases.

Figures A.12 through A.15 provide the pulse compressed responses for APC with the adjustments required so that APC could be applied to FM waveforms. Figures A.12 and A.13 show the responses for two oversampling factors ($K = 1$ and $K = 3$) and using 2 and 3 adaptive stages. There




Figure A.9: LS MMF responses with Doppler averaging for LFM and PD with Doppler shifts from $+0.05$ to -0.05 radians. Sidelobes increase as J increases. Generated for $J = 3$ and $J = 4$, and averaging linearly from -0.04 to $+0.04$ radians.

are two targets in the range profile, one at 100 with an SNR of 80 dB and another at a range of 95 with an SNR of 15. The oversampling factor, with respect to APC, is discussed in Section 3.1.

Comparing the APC responses for two and three adaptive stages, there appears to be no benefit for using a third adaptive stage. After two adaptive stages, APC is able to discern the smaller target while the matched filter and LS MMF cannot. With $K = 1$, the filter responses of the waveforms are that of a polyphase code and *not* an FM waveform. When $K = 3$, the filter responses are that of an FM waveform and as such, the LS MMF is now able to discern the smaller target. The LFM-based APC response, however, now appears to require 3 adaptive stages to discern the smaller target.

Figures A.14 and A.15 show the responses for two amounts of beamspooling ($z = 1$ and $z = 3$) and using 2 and 3 adaptive stages. Beamspooling, as it is used for APC, is discussed in Section 3.1.




Figure A.10: Doppler averaged LS MMF responses for LFM with Doppler shifts of 0 to -0.05 radians. Generated for $J = 2, 3$, and 4 and linearly averaged from 0 to -0.04 radians.

As z increases, APC does not appear to improve in terms of target detection. However, by increasing the value of z , the amount of MML in the APC response is reduced, but some degradation in resolution is realized. With 2 and 3 adaptive stages of APC, the improvement in sidelobe levels that results from increasing the value of z is not observed. This improvement can be seen in Section 3.1, which shows the APC responses with only one adaptive stage and varying the amount of beamspoiling z .

Figure A.16 provides the pulse compressed responses for the measured data when the direct path is eclipsed by 50%. The pulse compression methods used are the matched filter, optimal non-averaged LS MMF, and APC for both waveforms. This supplements the discussion in Section 4.2.2.

The performance of the matched filter and LS MMF when eclipsing occurs is degraded as the




Figure A.11: Doppler averaged LS MMF responses for PD with Doppler shifts of 0 to -0.05 radians. Generated for $J = 2, 3$, and 4 and linearly averaged from 0 to -0.04 radians.

sidelobes increase significantly. This decreased performance inhibits the ability of the matched filter and LS MMF to detect targets near the direct path, which APC is still able to identify. The performance of APC realizes negligible degradation with 50% eclipsing, causing slightly coarser resolution.

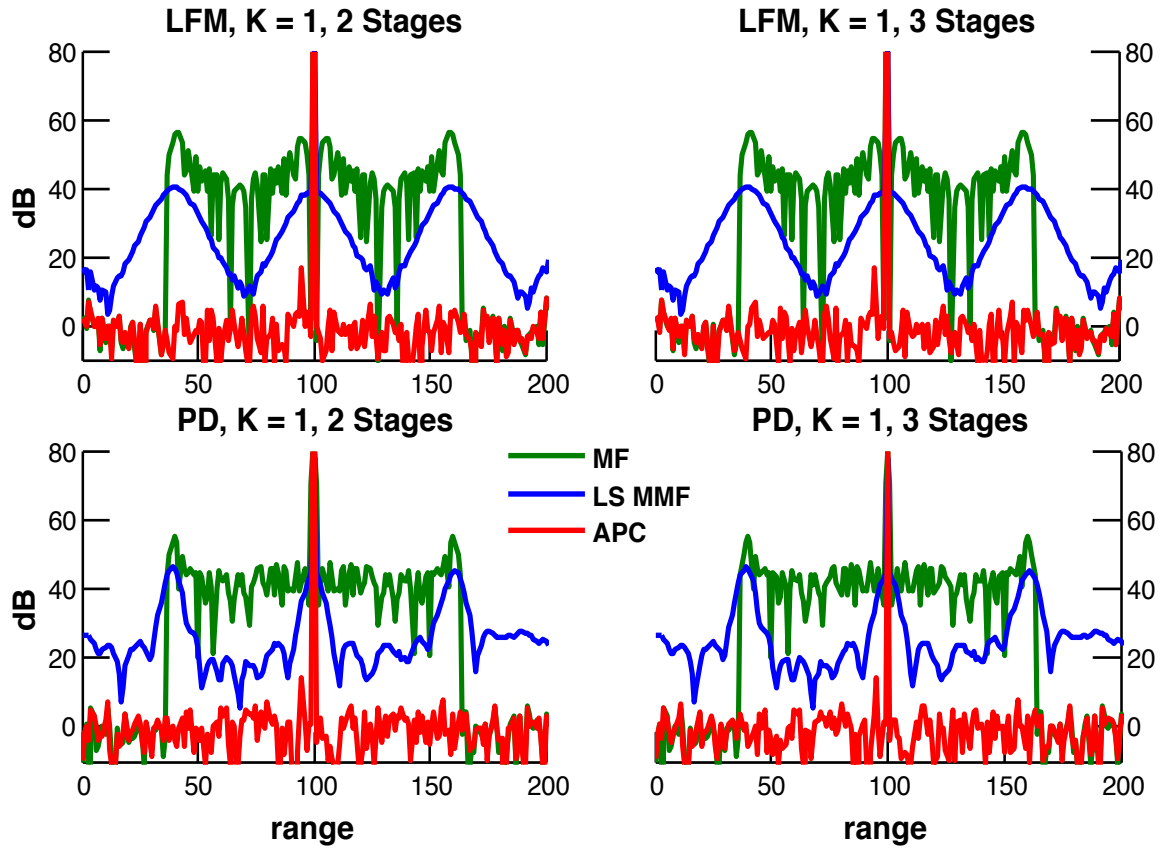


Figure A.12: Matched filter, LS MMF, and APC responses for LFM and PD with two and three adaptive stages and $K = 1$. APC is able to discern the smaller target while the matched filter and LS MMF are unable to do so.




Figure A.13: Matched filter, LS MMF, and APC responses for LFM and PD with two and three adaptive stages and $K = 3$. The LS MMF is able to discern the smaller target, and APC is able to as well, except in the LFM, $K = 3$, two adaptive stages scenario.




Figure A.14: Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with two and three adaptive stages and $z = 1$ row zeroed. LS MMF is able to discern the smaller target; APC is able to discern the target with three adaptive stages, but with two stages, it is less able to do so.




Figure A.15: Matched filter, LS MMF, and APC responses for LFM and PD. APC was implemented with two and three adaptive stages and $z = 3$ rows zeroed. APC and the LS MMF can discern the target in all scenarios.




Figure A.16: Pulse compressed responses of the measured data for LFM and PD with 50% eclipsing.