



Riverine Imaging

Contract: W911NF-10-2-0069
U.S. Department of Defense, ARL
Final Report

PI: Dr. Nikola Subotic
Co-Director, MTRI
Michigan Technological University
nsubotic@mtu.edu

Contract Manager: Lisa Jukkala
Research and Sponsored Programs
Michigan Technological University
lajukkal@mtu.edu

December 16, 2011

Distribution Statement A:
Approved for public release; distribution is unlimited.

2011221022

Riverine Imaging

**Contract: W911NF-10-2-0069
U.S. Department of Defense, ARL
Final Report**

Authors:

MTRI

*Dr. Nikola Subotic, Dr. Dr. Mark Stuff, Joseph Garbarino, James Ebling,
Christopher Roussi, Brian Thelen, Dr. John Valenzuela, Joel LeBlanc*

Applied Signals Intelligence, Inc.
John McCorkle

December 16, 2011

Distribution Statement A:
Approved for public release; distribution is unlimited.

Michigan Tech Research Institute (MTRI)
3600 Green Court, Suite 100 • Ann Arbor, Michigan 48105

Table of Contents

Three Dimensional Imaging Studies and Demonstration.....	Tab - 1
Information-Theoretic Analysis & Performance Bounds for Super- Resolution (SR) Video Imagery Reconstruction	Tab - 2
Appendix: Riverine Imaging	Tab - 3



Title: Three Dimensional Imaging Studies and Demonstration

Quarterly Report

Volume I & II

Date: January 26, 2011

Volume 1

1.	Introduction	1
2.	Overview of Processing Approach	3
3.	Simulations	6
4.	Local Collections.....	13
	Experiments	15
	Example Results	16
	Conclusions.....	18
5.	Data Analysis.....	19
6.	Description of the McCorkle Antenna.....	39
7.	Summary	51

Volume II

1	Introduction	1
2	Progress in the Fourth Quarter of 2010	2
3	Mathematical Framework	7
	3.1 Forward Model	7
	3.2 General Bayesian Cram'er-Rao Bound	7
	3.2.1 Translations only.....	10
	3.2.2 Performance bounds for the translation only case	12
4	Summary and Plans	12
	Appendices	14
	A Mathematics: Flat-earth approximation to projective transform	14
	B Parameterizing Quadratically Penalized Spaces	15



Title: Three Dimensional Imaging Studies and Demonstration

Quarterly Report

Volume I

Date: January 26, 2011

Three Dimensional Imaging Studies and Demonstration

***N. Subotic, M. Stuff, J. Garbarino, J. Ebling, C. Roussi, B.
Thelen, J. Valenzuela, J. LeBlanc, J. McCorkle***

January 26, 2011

Table of Contents

1.	Introduction	1
2.	Overview of Processing Approach	3
3.	Simulations	6
4.	Local Collections.....	13
	Experiments	15
	Example Results	16
	Conclusions.....	18
5.	Data Analysis.....	19
6.	Description of the McCorkle Antenna.....	39
7.	Summary	51
8.	References.....	Error! Bookmark not defined.

List of Figures

Figure 1: The algorithm development is in support of a 3D penetrative imaging system which monitors boat traffic.....	1
Figure 2: The placement of the RADAR along a river constrains the amount of angular diversity that is created due to the boat motion.	2
Figure 3: Penetrative RADAR and 2D reconstruction produces signatures that can be interpreted as revealing boat infrastructure and cargo.....	3
Figure 4 High level flow diagram for a possible processing system.....	4
Figure 5 Canal Boat Model.....	6
Figure 6 Canal Boat with Water Plane.....	6
Figure 7 Canal Boat 7 Point Model Locations.....	7
Figure 8 River Bend Simulation Motion.....	8
Figure 9 Ranges to Points in 7 Point Canal Boat Model.....	9
Figure 10 Simulated Range compressed signals of 7 Point Canal Boat Model.....	9
Figure 11 Closeup of Simulated Range compressed signals of 7 Point Canal Boat Model.....	10
Figure 12 Simulated Range compressed signals of Xpatch Canal Boat Model.....	10
Figure 13 Closeup of Simulated Range compressed signals of Xpatch Canal Boat Model.....	11
Figure 14 Point Scatterers for 6 Point Foam Cart Model.....	11
Figure 15 Ranges to Points in 6 Point Foam Cart Model.....	12
Figure 16 Simulated Range compressed signals of 6 Point Foam Cart Model.....	12
Figure 17 Closeup of Simulated Range compressed signals of 6 Point Foam Cart Model.....	13
Figure 18: Specifications of the Akela RADAR system used for collections.....	14
Figure 19: Local collection geometry.....	14
Figure 21: RPG used as target. RPG was in a wooden shipping crate during collections.	15
Figure 20: Target layout used on the cart for local collections.	15
Figure 22: Range compressed signature of two trihedral configuration as the sled moves by in time.....	16
Figure 23: Range profile of sled with full target complement.....	17
Figure 24: Example 2D ISAR image of two trihedral configuration.	17
Figure 25: ISAR image of the RPG showing its flange and front.....	18
Figure 26 Range Compressed signals of Motion Compensated 7 Point Model.....	19
Figure 27 Range Doppler Map of Reduced 7 Point Model.....	20
Figure 28 Frequency vs Doppler and Frequency vs Velocity Maps of Reduced 7 Point Model.....	21
Figure 29 CTRV Time, Range, And Velocity Views.....	21
Figure 30 Range Compressed signals of 6 Point Foam Cart Model.....	22
Figure 31 Range Compressed Signals from Motion Compensated 6 Point Foam Cart Model.....	23
Figure 32 Range Compressed Signals and Range Doppler Map, Reduced 6 Point Foam Cart Model.....	23
Figure 33 CTRV Time, Range, And Velocity Views of 6 Point Foam Cart Model.....	24
Figure 34 Paths Through Time-Range-Velocity Cuboid for 6 Point Foam Cart Model, Overlaid on Range Compressed Signals.....	25
Figure 35 True and estimated range histories.....	26
Figure 36 True and estimated velocities.....	26
Figure 37: True scattering locations (blue) and estimated 2D(green) and 3D (red) locations based on the invariant algorithm.....	27

Figure 38 True and estimated reflector configuration from far field geometry analysis. Truth in blue, 2d estimate in green, 3d estimate in red.	28
Figure 39 Angular motion estimates	29
Figure 40 Three dimensional back-projection image with estimated scattering center locations	30
Figure 41 Polar formatted two-dimensional image formed by using motion estimates from far field invariant theory	30
Figure 42 Akela Foam Cart Data Before and After Background Subtraction.....	31
Figure 43 Akela Foam Cart Data Before and After Target Extraction.....	32
Figure 44 Akela Foam Cart Data, Motion Compensation and Spatial Reduction.....	32
Figure 45 Akela Foam Cart Data Time Range Velocity, No Whitening or Extrapolation.....	33
Figure 46 Akela Foam Cart Data Time Range Velocity After Whitening or Extrapolation.....	33
Figure 47 Akela Foam Cart Data Scatterer Paths Shown on Range Compressed signals.....	34
Figure 48: Antenna pair of cardioid type produced by John McCorkle.....	40
Figure 49: Frequency dependent gain plot of McCorkle antenna pair.	40
Figure 50: Frequency dependent SWR of McCorkle antenna pair.....	41
Figure 51: Range Profile of -17.8 dBsm Sphere with (blue curve) and without (red curve) chamber background subtracted	45
Figure 52: Predicted SNR vs. Range Performance for the Akela Radar with 500 mW Power Amp Based on Noise Power Calculated from Sphere to Chamber Noise Level	47
Figure 53: Compressed Range Profiles with Average Range Profile Coherently Subtracted for Akela Radar with All Ports Terminated	48
Figure 54: Predicted SNR vs. Range Performance for the Akela Radar with 500 mW Power Amp Based on Noise Power Calculated from Sphere to Terminated Radar Noise Level.....	50

1. Introduction

Numerous sensing modalities can produce 3D data sets. Two very important surveillance modalities are video (2D space plus time) and 3D Inverse Synthetic Aperture RADAR (ISAR). It is the goal of this task to examine these specific modalities and 1) for video – determine fundamental bounds on sub-sampling, resolution and reconstruction such that large video data volumes can be moved through communication channels; and 2) formulate and demonstrate a 3D ISAR image reconstruction process to be used for penetrative reconstruction of river boat internals as they float by a small, low frequency RADAR. The envisioned system could provide remote, unattended river intelligence, surveillance and reconnaissance (ISR) for several months in a concealed package located near tropical rivers, mountain passes, and locations of special interest.

The kernel of this project is the development of reconstruction algorithms. Specifically the development and implement of novel algorithms for generating 3D representations of targets and their contents using a single penetrative RADAR system. The application area for this technology is river boat surveillance. An overview figure of this application area is shown in Figure 1.



Figure 1: The algorithm development is in support of a 3D penetrative imaging system which monitors boat traffic.

The concept of operations for this sensing modality places restrictions on the type and amount of diversity that can be used in 3D reconstruction. Because the sensor is surveilling a specific part of a river and is fixed in a tree, the boat must provide the angular diversity necessary for 3D reconstruction. We will not consider an interferometric system due to packaging constraints. Figure 2 shows where a potential system will be placed. The implications of these placements is that the boat will either be moving straight through the RADAR beam as it moves down a straight segment of the river or will scribe a curve as it move around a river bend. We are exploring the amount of 3D information inherent in that type of motion.

The ideal types of motion include roll/pitch/and yaw which fills a 3D measurement space with data. However, this isn't compliant with the eventual deployment of such a system and would motivate installing the RADAR to perform surveillance only on areas of the river that have rapids which is where boats don't wish to go. The straight and curved trajectories give less 3D fill. Specifically, as the boat moves through the large RADAR beam, there is a diversity in azimuth angle and in depression angle which provides some inherent 3D data fill. We will be concentrating on that scenario in our development and testing.



Figure 2: The placement of the RADAR along a river constrains the amount of angular diversity that is created due to the boat motion.


Because the overall sensing processing goal is penetrative imaging of river boats to determine their internal structure and whether they are carrying contraband, the sensing modality is a low frequency, high bandwidth RADAR. We are operating in the 500MHz to > 2GHz frequency bands to allow for penetration with modest RADAR powers along with sufficient bandwidth to isolate specific scattering centers in range such that they can be track over time. This quarterly report outlines the progress that has been made in the algorithm development and also describes the supporting data collections and simulations used in the algorithm development.

We have performed signature simulations and data collections to provide input to be used by the algorithm development. Figure 3 shows a 2D SAR example based on a boat simulation. The ISAR reconstruction algorithm is a phase estimation/compensation algorithm on specific scattering center tracks using a rigid body assumption of motion that ties the various tracks together spatially. Invariant geometrical measurements, resolution of three-dimensional shape via comparison of multiple views, fusion with optical, more literal interpretation, simple relation to the user's point of view, greater target discrimination capabilities.

This report is organized as follows: Section 2 provides an overview of the processing approach and algorithm developments attained during this reporting period. Section 3 describes the simulations performed and provides example of the simulations that were used in the signature analysis and algorithm development. Section 4 discusses the local collections that were made with the MTRI RADAR assets to provide empirical measurements for analysis and algorithm development. Section 5 discusses the data and provides results of the data analysis. Section 6 provides information in regards to the antenna that was designed for this task and discusses its performance. A summary is then provided in Section 7.

Xpatch Simulation with Weapons Cache
Radar image formed using wavenumber migration

*However Need to Determine
Inside and Outside*



22

Figure 3: Penetrative RADAR and 2D reconstruction produces signatures that can be interpreted as revealing boat infrastructure and cargo.

2. Overview of Processing Approach

The flow diagram shown in Figure 4 presents a very high level concept for a processing chain for the radar signals. It has ten major components. For each component of the diagram, the envisioned system will contain multiple subcomponents, and there are many possible variations in the implementation details. Later diagrams will illustrate some of those possible components. But first we will describe the purposes of the major components.




Figure 4 High level flow diagram for a possible processing system

The first component comprises an initial data processing stage in which many of the specific radar and environment dependent characteristics of the collected signals may be addressed. These issues include conversion of the data from whatever format is most convenient for data collection into a format which is convenient for processing. Other issues include mitigation of any deterministic signals in the data which have been introduced by electrical leakage in the collection hardware. This appears to be a pervasive problem of small compact systems. This processing stage may also need to mitigate undesirable features in the impulse response of the collection system, to mitigate interference signals which were injected into the data by nearby sources of radio frequency signals, mitigate effects of gaps in the set of collected radio frequencies due to transmit restrictions, and handle irregularities in the collection times of the radar sweeps.

The second component of the envisioned processing system achieves at least three major functions. First, the target signals contained in the data are modified by the processing so that reductions are achieved in the apparent range extent of the target signals, the maximal speed (opening or closing rate) of the target signals, and the accelerations of the target signals. Second, these reductions enable a compression of the target signals into few digital samples, which makes enables more efficient subsequent processing (and/or signal exfiltration). Third, the signal to clutter plus noise ratio is enhanced in the digital remaining samples.

The third component of the processing begins the automated analysis of the target signals. Its purpose is to separate signals from distinct parts of the target, and to estimate the changing ranges to each of those parts of the target. This stage is needed because, encoded within the extracted pulse by pulse descriptions of the ranges to the multiple distinct parts of the target, is all the information which the radar signals contain, concerning the size, shape, and relative motions (angular motions) of the observed parts of the target. This processing stage produces range histories, for a set of distinct physical components of the target. Each range history is a function of time, sampled at the radar's pulse (or sweep) rate, which give the distance from the radar to a particular component of the target. This stage of the processing extracts the range histories for as many target components as can be accurately followed, but it does not need to extract range histories for all of the target components. If the extracted range history estimates are sufficiently precise, four distinct histories will suffice for the subsequent processing stages, but more will generally be extracted to provide redundancy which subsequent components may use to mitigate residual errors in the estimates. At this stage of the processing the system does not require any knowledge of which part of the target a particular range history estimate is associated with. The association of the range histories with specific physical components of the target will be determined later, by subsequent stages of the signal processing.

The fourth component of the processing system examines the range history estimates jointly. This is a robustification stage. This stage is not needed if all of the range estimates from the previous stage are good estimates of the ranges to distinct components of the target. But we include it the system design, because we

have developed methods to use geometric constraints on the joint behavior of the true range histories to rank the quality of a set of estimated range histories, and to eliminate estimates which may contain large errors. (There are various possible ways in which an erroneous range history may have been obtained. Imagine, for example, that the system has extracted the range history from a reflection from a log, which happened to float through the scene while it was being traversed by the target.)

The fifth component uses the range histories to extract the three-dimensional size and shape of the configuration of the target's reflective components. It extracts a complete three-dimensional description of the geometry of the configuration of the set of components that reflected the radar signals (from which the range histories were extracted). All characteristics, such as distances, angles, areas, and volumes, of the configuration become known, and each of the range histories becomes associated with a specific location on the target. The processing strategy for this part is based on geometric theorems which describe how the measured ranges to distinct points on the target are constrained by the size and shape of the target. These constraints lead to systems of equations which determine a fundamental set of geometric invariants, which in turn determine the size and shape. At this stage of the processing, some of the measurements of target and cargo characteristics, which might be directly utilized by an intelligence analyst, begin to emerge from the processing. Yet, there is more processing to be done, so as to extract additional information from the data, and to put that information into useful forms for the end user.

The sixth component uses the detailed size and shape information, from the previous processing stage, along with the associated range histories, to determine all of the observable motions of the target. Because the previous stage has determined the shape of the configuration of the reflectors, and the range histories keep track of the pulse by pulse distances from the reflectors to the radar, constructing the pulse by pulse motions becomes just a matter of solving a small set of equations at each pulse time. The motion solution yields the specific direction from which the target was illuminated at each pulse time. It also refines the initial range to target estimates, and enables the calculation of the pulse by pulse range to any point (reflective or not) on the target. The seventh component uses the motion information to form various image products. These products are based on all the information provided by the sensor signals. The motion estimates from the previous stage, provide all of the auxiliary information which is needed by conventional SAR/ISAR image formation processing algorithms. The image formation processing applies the motion dependant transformations to the clutter reduced radar signals, which were provided by the second stage of the processing. Conventional two-dimensional images can be formed from subsets of the radar signals, using conventional methods (polar formatting, migration, chirp scaling, two-dimensional back-projection, ...). Overlapping subsets can be used to produce movies ("dynamic images") showing the motions of the target (and possible indications of motions of objects or persons on the target). Also all the collected pulses can be simultaneously, to produce images with three spatial dimensions. Conventionally, this is done with back-projection, but more advanced methods can also be used to produce better images by exploiting the natural sparseness of illuminated reflective surfaces in three-dimensional space.

The remaining three components are not yet as well defined by our current work, but are included in the description because of their importance for expanding the audience for the products of the system.

The eighth component is planned to enhance the utility of the shape measurement and image products by corresponding the products with stored additional sources of target and cargo knowledge, such as models and images of known boat types, and previously collected data sets.

The ninth component will use knowledge of boat design, cargo characteristics, boat behavior patterns and regional knowledge, together with the model and image correspondences from the previous component, to enhance the interpretability of the sensor products.


The tenth envisioned component is conceived as primarily an information formatting and display component. Having done everything that can be done efficiently by an automated system to prepare, the last component must be designed to facilitate the exploitation of the combined information by the user.

3. Simulations

Various data simulations were constructed during the course of the program. These simulations supported algorithm development both by providing test data in the period before actual collected data were available, and to simplify and separate phenomenology for closer study.

Two types of simulations were performed, one consisting of simple point scatterer responses, the other utilizing the Xpatch simulation tool. The simple point scatterer simulations run quickly, providing simulations quickly. The Xpatch simulations give some insight into the interactions due to penetration through various dielectric materials, as well as multi-bounce effects. Both are valuable tools during algorithm design and development. To support the simulations, we obtained a faceted model of a river boat. We assigned materials to the boat to match the project con-ops: mainly wood construction, with various smaller metal parts. The boat was also given a motor. Figure 5 shows the canal boat model used in the simulations.

**Metal smokestack,
rail, and motor**



**Wooden roof, sides,
and hull**

Figure 5 Canal Boat Model

The boat was placed in a plane to simulate being on the surface of water.




Figure 6 Canal Boat with Water Plane

The various simulations were designed to emulate the parameters of the expected collected datasets. Simulations were done in the near field at low frequencies (300-2000 MHz) to simulate penetration of RADAR signals through the dielectric materials.

Canal Boat 7 Point Scatterer Simulations

To help us understand issues related to both near field effects and scatterer interaction during typical boat motion, we created a 7 point model based on the canal boat. We chose the 7 points to span the 3 dimensions of the boat, to support simulation of typical scatterer responses and interactions, and provide an unambiguous signature. Figure 7 shows the locations of the 7 points on the model (note: The model is being shown partially transparent to illustrate scatterer locations).




Figure 7 Canal Boat 7 Point Model Locations

The points in the 7 point model were given the motions expected of a canal boat going around a bend in a river. Figure 8 shows some points around a river bend, and a possible sensor location, that were used to generate a path for the points.




Figure 8 River Bend Simulation Motion

This produced the types of ranges histories we expected to see in the collected data. Figure 9 shows the ranges to the points in the 7 point model, Figure 10 shows the range compressed signals created from the simulated signal history created from those ranges, and Figure 11 shows a closeup of the range compressed signals during the period when the boat crosses broadside to the RADAR, causing the scatterers to cross.




Figure 9 Ranges to Points in 7 Point Canal Boat Model




Figure 10 Simulated Range compressed signals of 7 Point Canal Boat Model




Figure 11 Closeup of Simulated Range compressed signals of 7 Point Canal Boat Model

Xpatch Canal Boat Simulation

An Xpatch simulations was run to provide insight into the effects of dielectric materials and multi-bounce effects on the range compressed signals. The canal boat with water plane was used as the target, and it was given the same motion as the 7 point model. We also used the same RADAR parameter. The range compressed signals for the canal boat Xpatch model are shown in Figure 12 and Figure 13.




Figure 12 Simulated Range compressed signals of Xpatch Canal Boat Model




Figure 13 Closeup of Simulated Range compressed signals of Xpatch Canal Boat Model

Foam Cart Point 6 Point Scatterer Simulation

In preparation for our data collection, we performed some local collections, which are discussed later in this document. One of these collections was a cart covered in foam, with 6 reflectors mounted on the top. To help us understand the phenomenology, we used the point-to-point distances between scatterers to create a simulation of this cart, and gave it a straight line motion in front of the sensor. The locations of the scatterers are shown in Figure 14.




Figure 14 Point Scatterers for 6 Point Foam Cart Model

We gave the points the motion shown in Figure 15, and these produced the range histories shown in Figure 15. In Figure 16 and Figure 17 the range compressed signals produced from the simulated signal history are shown.




Figure 15 Ranges to Points in 6 Point Foam Cart Model

Range Histories From Simulated Signal History of 6 Point Model




Figure 16 Simulated Range compressed signals of 6 Point Foam Cart Model

Range Histories From Simulated Signal History of 6 Point Model




Figure 17 Closeup of Simulated Range compressed signals of 6 Point Foam Cart Model

4. Local Collections

Prior to field collections, it was prudent to characterize the proposed system configuration under controlled conditions, and in an environment in which rapid changes and corrections could be made as needed. To this end, we mounted the RADAR system on the roof of our building with a view of a clear, grass-covered area that had a geometry similar to proposed field-collections.

The system itself consisted of an Akela radar, the McCorkle antennas, a 1W power amplifier, and 20ft low-loss cables. The radar specification are depicted in Figure 18 below:

Akeia Radar Specifications

380MHz – 3000MHz, FMCW
 Sample rates from 3KHz – 90KHz
 Built-in, selectable, low-noise amplifier
 Built-in range-gating
 TX pulse length 24ns - 1556ns
 RX delay 18ns - 1556ns
 RX gate 18ns - 1556ns
 Interpulse delay 18ns - 1556ns
 Output power
 17dBm typical, 20dBm maximum
 Input signal limits
 +13dBm (LNA off)
 -7dBm (LNA on)
 4 – channel
 Operates in RX/TX pairs
 2 position encoder inputs





Figure 18: Specifications of the Akela RADAR system used for collections.

The geometry of the collection is depicted in Figure 19 below:

Measurement Setup



Layout



Figure 19: Local collection geometry.

The overhead view above does not accurately reflect the vehicles that were present during the collection. The range gate of the radar was adjusted to reject the vehicles present.

Targets were then built and mounted on a large piece of construction foam. The foam itself is transparent at the frequencies of interest, but served as a rigid platform on which the targets could be arranged as needed. The targets consisted of trihedrals, spheres, top-hats, and an RPG in its wooden shipping crate. Some target arrangements are shown in Figure 20 below. The RPG used in the collection is shown in Figure 21.




Figure 20: Target layout used on the cart for local collections.

The cart itself had 30m ropes attached at the corners to allow it to be dragged by staff that were outside the range-gate extent. The cart was started to the right of the radar beam, and dragged at a relatively constant rate (estimated to be 0.75m/s) through the beam, and well beyond, to provide a target-free background estimate in every data set collected.



Figure 21: RPG used as target. RPG was in a wooden shipping crate during collections.

Experiments

Data were collected under these conditions:

1. 500MHz-3000MHz
2. 501 frequency samples evenly-spaced in frequency over this range
3. 700mW estimated effective radiated power
4. Various target arrangements, ranging from an empty cart (with mounting hardware, shown in the figure above), to a crowded arrangement of up to 6 targets. The experiments performed were:

1. Empty cart
2. 1 trihedral
3. 2 trihedrals
4. all targets (placed as indicated in the figure above, left)
5. RPG on foam, angled away from radar
6. RPG on foam, angled toward the radar
7. RPG elevated above foam, angle away from radar
8. RPG elevated above foam, angled toward the radar
9. RPG in wooden box under the variations indicated above

The reason for varying the RPG elevation was to be able to assess the effects of having targets very close to a ground-plane. It was felt that this might mimic targets being in a boat on the bottom of the hull (close to the water), as well as a target on a deck above the water.

Example Results

Shown below are some of the quick look 2D results obtained during these tests. The data analysis section will describe the 3D processing of these data.

The first example, shown in Figure 22, is a range-compressed image of the foam cart with 2 trihedrals present. All the range-compressed images shown are formed by taking the Fourier Transform of each frequency sweep, and displaying the time course of the sweeps. Increasing time is to the right, and increasing range from the radar is towards the top of these images.




Figure 22: Range compressed signature of two trihedral configuration as the sled moves by in time.

As the targets enter the radar beam, they become increasingly bright, their ranges decrease until they are broadside to the radar, and then they move through, and away from the beam. This image was processed with background clutter suppression, and the target-to-clutter ratio is seen to be over 30dB. This would be considered ideal data for the analysis we wished to perform.

Another example showing the range-compressed image for the sled configuration using all targets present as was shown in Figure 20 is shown in Figure 23.




Figure 23: Range profile of sled with full target complement.

Again, these data appear to be of high quality; there are few system artifacts present after clutter suppression, and the dynamic range of the targets is over 30dB.

Although we did not have motion information recorded for these data, it was observed that the staff

Two Targets




Figure 24: Example 2D ISAR image of two trihedral configuration.

were able to drag the cart at a very regular pace, allowing fully 2-D ISAR images to be formed without the need for adaptive phase estimation. Because the target was in the near-field of the radar, and wavefront curvature was significant, we used a range-migration algorithm to form images.

An example of such an image is shown in Figure 24 for the case of 2 targets mounted on the foam:

In these images, increasing range from the radar is to the right, but due to a sign error in the processing code, the image is the mirror of the actual geometry. However, it can be seen that the target impulse responses are quite compact and symmetric, indicating that the uniform linear path that was estimated is a good approximation to the actual motion. There is at least 40dB of dynamic range in the data, and systematic effects are minimal.

Figure 25 shows the ISAR image of the RPG:

RPG, in box




Figure 25: ISAR image of the RPG showing its flange and front.

Again, this is the mirror image of the actual arrangement, but the ends of the RPG are visible in the image, as is part of the body, even through the wooden box, as would be expected at these frequencies.

Conclusions

The system as configured performs well. Calculation show that there is sufficient power to provide a good signal level at the receive amplifier for targets of this RCS at the ranges expected in the field collection. The antennas and cables appear to be correctly matched over the band of interest, and the range-gate settings can be adjusted to receive targets in the ranges of interest, and reject targets outside that region. The system phase is stable, systematic errors are stationary over the times that we

are likely to be collecting, indicating that clutter suppression via background subtraction is a viable approach. The RMA code has been demonstrated to be substantially correct, and is being incorporated into a “quick look” capability, allowing rapid assessment of the data quality during field collections. The data are of a high enough quality that they can be used in the testing and development of the 3-D processing code.

5. Data Analysis

In this section we provide some example processing results for the simulated and locally collected data sets described previously. These results show the ability of the processing algorithm to determine and track scattering mechanisms in 3D such that a 3D spatial reconstruction can occur.

7 Point Canal Boat

The 7 point canal boat simulation dataset was used to study performance of existing signal separation techniques with these datasets. The data were first motion compensated (Figure 26) using the range to the centroid of the 7 points at each sweep. The goal of the compensation is to slow the range rate of the signals. By compensating the signals from a moving target to a target fixed point, the translational motions of the point are removed, and the only remaining motions are due to rotations of the target about the compensation point. This step reduces the number samples needed to contain the responses from the target. There are two reductions enabled. One reduction is in the number of radio frequency samples that need to be preserved. This is possible because the range compression of the signal occupies fewer range bins after the motion compensation. The other reduction is in the number of times samples. This reduction is possible because the Doppler compression of the target signals occupy fewer Doppler bins after the compensation.




Figure 26 Range Compressed signals of Motion Compensated 7 Point Model

The boat signatures were then extracted from the 2D spatial domain, to help reduced the effects of clutter and the computational burden. Figure 27 shows the result of this extraction.

Range Doppler Map of Reduced 7 Point Canal Boat Model




Figure 27 Range Doppler Map of Reduced 7 Point Model

The goal of this earlier stage of analysis was to produce range history paths for the geometry analysis stage. To do a good job of extracting these range paths, the path extraction was done in a 3 dimensional cuboid with dimensions time, range, and velocity. The 2 dimensional time frequency dataset used to produce the range Doppler map shown above was converted into a 3 dimensional dataset of range velocity slices by taking overlapping subsets of the frequency sweeps. These subsets of sweeps were converted to frequency Doppler, then interpolated along Doppler to produce frequency velocity. This interpolation produces lines of constant velocity in the output. Figure 28 shows the frequency Doppler and frequency velocity maps for the reduced 7 point model.




Figure 28 Frequency vs Doppler and Frequency vs Velocity Maps of Reduced 7 Point Model

Shown in Figure 29 are three visualizations of the time-range-velocity cuboid. This is the cuboid that will be used in the path finding process. These views clearly show the paths of the scatterers, and illustrate that when paths converge in range, they diverge in velocity, providing the leverage needed to separate the scatterers.




Figure 29 CTRV Time, Range, And Velocity Views

6 Point Foam Cart

As described earlier, we performed a local collection with a cart covered in foam carrying 6 reflectors mounted on the top. To help us understand the phenomenology, we used the point-to-point distances between scattering centers to create a simulation of this cart, and moved it in a straight line motion in front of the sensor. Figure 30 displays the range compressed signals associated with the 6 point foam cart model.




Figure 30 Range Compressed signals of 6 Point Foam Cart Model

Motion compensation was performed to allow us to crop the dataset to help reduced the effects of clutter and lessen the computational burden. The motion compensation was performed using a shear averaging algorithm. The range compressed signals formed from the motion compensated signal history are shown in Figure 31, while Figure 32 shows both the range compressed signals and range Doppler map resulting from the spatial cropping.

Motion Compensated Range Compressed Signals From Simulated Signal History of 6 Point Model




Figure 31 Range Compressed Signals from Motion Compensated 6 Point Foam Cart Model




Figure 32 Range Compressed Signals and Range Doppler Map, Reduced 6 Point Foam Cart Model

Using this reduced signal history, the time-range-velocity cuboid was created, as shown in Figure 33.




Figure 33 CTRV Time, Range, And Velocity Views of 6 Point Foam Cart Model

Following construction of the time-range-velocity cuboid, the paths of the scatterers were extracted. This extraction was done in a way that respects physical constraints imposed upon changes in range by the current velocity. We also disallow any velocities that are outside velocities of interest for the problem at hand. These constraints do not allow paths to make transitions that are non-physical.

The paths extracted from the time-range-velocity cuboid for the 6 point foam cart model are shown in Figure 34 as black lines overlaid on top of the range compressed signals.

Range




Figure 34 Paths Through Time-Range-Velocity Cuboid for 6 Point Foam Cart Model, Overlaid on Range Compressed Signals

The maximal difference between the estimated range histories and the true range histories was 1.52 centimeters and the root mean square distance was 7.52 millimeters. Thus the differences are difficult to see in Figure 35 which shows the true and estimated values together.




Figure 35 True and estimated range histories

Differences can be seen more easily in Figure 36, which shows the true and estimated velocities of the synthetic reflectors.




Figure 36 True and estimated velocities

Geometric analysis was applied to both the true and the estimated range histories. The purpose of the geometric analysis is to reconstruct the observable components of the shape and motion of the scattering center configuration. Four versions of the analysis were examined. These are the three and two dimensional versions of the near field, and the three and two dimensional versions of the far field theory. The three-dimensional near field theory is the most general theory. It makes no restriction on the motion of the target, and correctly accounts for spherical spreading of the radio waves. But, despite the fact that it derives from the most correct model for the typical physical situation, current algorithms based on this theory can perform poorly in certain constrained motion situations. These problems arise when the illumination directions are not diverse enough to keep the input data well removed from the singularity conditions associated with certain restricted types of motions.

The errors in the estimated ranges are expected to introduce some error in the geometric reconstructions. Additional errors may be expected in the estimates from the far field estimation process because it ignores wave-front curvature. Further, because the true underlying motion was linear, the three-dimensional reconstruction problem is singular for this case. This should make the three-dimensional solution extremely sensitive to noise in the data.

The straight boat path is an example of a severely degenerate motion case, and this readily degrades the estimates formed with the near field algorithm. This is illustrated by the estimates of the scattering center configuration which were obtained from the near field algorithm. These estimates are shown in Figure 37.




Figure 37: True scattering locations (blue) and estimated 2D(green) and 3D (red) locations based on the invariant algorithm.

For both the two-dimensional (green) and the three-dimensional (red) near field estimates, the coordinates are not near to those of the true coordinates shown in blue. The estimated shape is too narrow, too close to being collinear, to match the shape of the true configuration. The poor estimation may be attributable to the severity of the over-parameterization of the near field invariant equations in the degenerate cases.

The far field estimates are much better. A view from a three-dimensional display, shown in, Figure 38, exhibits the true configuration (in blue), the three-dimensional far field estimate (in red), and the two-dimensional far field estimate (in green).




Figure 38 True and estimated reflector configuration from far field geometry analysis. Truth in blue, 2d estimate in green, 3d estimate in red.

This particular view is nearly normal to the plane of the two-dimensional estimates, so the planarity of the two-dimensional estimates (shown in green) is not apparent (as it would be from other points of view). In this view, the estimates appear to match the truth reasonably well, despite all the reasons why they might not. The most obvious visible deviation is the significant separation of the three-dimensional estimate (shown in red) near the right side of the image, from the true location which appears near the far right edge of the image.

The estimation of the coordinates enables the estimation of the illumination directions, where were controlled by the motion of the target. These illumination direction estimates can then be used to control image formation processes, which convert the original signal history into an image. Given a three-dimensional motion solution, both three-dimensional and two-dimensional image formation process can be used. The qualities of the image products depends not only upon the image formation method, and the accuracy of the motion estimates, but also on the geometric properties of the set of illumination directions. For good two-dimensional images, finding large subsets of the illumination directions which are nearly coplanar are best. For good three-dimensional images, one should use all the data, not a subset, and hope for illumination directions that densely sample a large solid angle.

For this example the estimated illumination directions are shown in Figure 39. The blue curve shows the angle of rotation about the principle axis of rotation as a function of the down-sampled pulse numbers. The green curve gives an indication of the wobbling of the principle axis of rotation.




Figure 39 Angular motion estimates

From the blue curve, we see that there is more than 80 degrees of rotation about the principle axis of rotation. This suggest that the data should support fine resolution imaging in the plane that is orthogonal to the principle axis of rotation. The green curve shows that there is also some wobble of the axis of rotation. Hence the estimate indicates that there is some three-dimensional information in the data. If the green curve were constantly zero, the direction of the estimated axis of rotation would be fixed. In that situation the estimate would imply that a two-dimensional image could be formed using all of the data, and that the resulting two-dimensional image would contain no height dependent defocusing. In that zero-wobble case, there would be no information in the data, concerning variations in the direction of the rotation axis. Because the variation in the green curve is not nearly as large as the blue curve, we expect less resolution in the direction of the principle axis of rotation.

Using the three-dimensional motion estimates, we formed a three-dimensional back-projection images from the radar data. One view of this three-dimensional image is shown in the Figure 40. The image also show markers for the estimated three-dimensional locations of the reflectors which were track though the signal data. These markers were inserted after the three-dimensional back-projection image was formed.




Figure 40 Three dimensional back-projection image with estimated scattering center locations

As expected there are two directions of fine resolution, due to the bandwidth of the radar data, and the large angle of rotation about the principle axis of rotation. But in the a direction of the rotation axis, there is poor resolution due to the limited angular diversity. because there is only modest variation in the estimated axis of the rotation during this collection.

Because of the slow variation in the axis of rotation, reasonably well focused two dimensional images can also be made from subsets of the data. Such and image is shown in Figure 41.




Figure 41 Polar formatted two-dimensional image formed by using motion estimates from far field invariant theory

Foam Cart Data From Akela RADAR

As described earlier, we performed a local collection with a cart covered in foam carrying 6 reflectors mounted on the top. This section describes the processing of that dataset.

To help reduce internal system noise and stationary clutter elements, a background subtraction must be performed on the Akela data. During this step, a principal components analysis was performed on the signal history data to determine which parts of the signal are stationary. These stationary components are then subtracted away. The results of this process on the Akela foam cart data are shown in Figure 42.




Figure 42 Akela Foam Cart Data Before and After Background Subtraction

For the purposes of range history path extraction, we are only interested in those sweeps that contains signals from the target of interest. Those sweeps that contain no target signal were eliminated. The result of this extraction is shown in Figure 43.




Figure 43 Akela Foam Cart Data Before and After Target Extraction

Motion compensation was then performed using a shear averaging technique to slow down the signals and allow spatial cropping. The motion compensated range compressed signals and range compressed signals and range Doppler map of the cropped dataset are show in




Figure 44 Akela Foam Cart Data, Motion Compensation and Spatial Reduction

Due to antenna patterns imposed by the antenna used in the foam cart collection, it was necessary to adjust the amplitudes of the values in the signal history domain to correct a frequency dependent falloff. Additionally, to achieve higher resolution, we also used extrapolation techniques to extend the frequency space, instead of simple zero padding. These two techniques were applied as part of the conversion from the 2 dimensional sweep-frequency data to the 3 dimensional time-frequency-velocity data. Figure 45 shows the time-range-velocity data before extrapolation and whitening, while Figure 46 shows the data after.




Figure 45 Akela Foam Cart Data Time Range Velocity, No Whitening or Extrapolation




Figure 46 Akela Foam Cart Data Time Range Velocity After Whitening or Extrapolation

During a collection, individual scatterers on the boat will come and go as they are shadowed/exposed, and their returns wax and wane due to angular response. The algorithm used to find scatterer paths must account for these effects. We chose to break the time on target into 11 overlapping regions. Each region overlapped half of

the previous region. Path segments were found within each region, and were combined and ranked based on the minimum value along the length of the path (higher minimum value implying a better segment). Figure 47 shows the scatterer path segments that were extracted on top of a projection of the range compressed view of the time-range-velocity cuboid.





Figure 47 Akela Foam Cart Data Scatterer Paths Shown on Range Compressed signals

Theoretical Considerations

Baseline scenario

Consider a sensor at coordinates $(0,0,h)$. Consider a three-dimensional target whose center travels along a straight path defined by $(x + vt, y, z)$, wherein t denotes time and v denotes velocity. This situation is illustrated, from various points of view, in the following figure, for $y=10$ and $z=10$.




Various views of the sensor, target, target path, and rays from sensor to the target path.

The rays from the sensor to the target path all lie in a fixed plane, the plane determined by the sensor location and the line along which the center of the target proceeds. In a far-field scenario (whereby we mean that wavefront curvature is negligible on the scale of the target dimensions), this plane constitutes the slant plane, and only one specific two-dimensional image can be produced. That is the standard range, cross-range SAR (or ISAR, as the distinction depends only on ones frame of reference) image, in which the direct reflections from all scattering centers are orthogonally projected into the slant plane. Under the far field approximation (planar wave-fronts, on the scale of the target, from each pulse), all points on the target that project onto the same point in the slant plane, have the same far field range histories throughout the aperture; hence they contribute indistinguishable signals to the radar data.

“Apparent roll”

In the near field, both the aspect and elevation angles of the rays from the target center to the sensor (and, isomorphically, the aspect and depression angles of the rays from the sensor to the target center) change significantly as the target moves along its path. For the given scenario, these angles are plotted in the following figure:



Aspect and elevation angles to the sensor as seen from a target fixed coordinate frame.


The change in elevation angle gives rise to the beguiling notion of “apparent roll” of the boat relative to the sensor. This “apparent roll” exists even though there is no roll of the boat relative to the planar surface of the water. Hence, there will exist in the radar data, some information about the height (vertical distance from the center of the target) of any scattering center on the target.

Ambiguity in three dimensions


The “apparent roll” observation seems to suggest that there is more three-dimensional information in the near field radar data, concerning a linearly moving target, than there is in the far field approximation of the measurements, wherein only two-dimensional measurements (in the slant plane) can be made (given only the information in the data). But this conclusion is troublesome from a topological perspective.

The transformation from the far field case to the near field case involves a smooth deformation of the metric on the space, as long as the target stays away from the only singularity, which is at the sensor. So as long as the target does not collide with the sensor, topological invariants, such as the dimension of the slant plane, and that of its complement (the span of its normal vector) should be the same as in the far field case. So the dimensionality of the resolvable target shape information should be conserved.

So where does the ambiguity (represented in the far field by the projection along the normal to the slant plane) go? The answer is, it bends around a line through the sensor location, which is parallel to the target trajectory. To illustrate this, we can examine the ranges (exact near field ranges, no approximations) to some distinct points along the circular arc shown in red in the following figure. Each of the seven marked points along the red curve is 1 meter from its nearest neighbor.



A rigid arc of ambiguous possible scattering center locations attached to the rigid boat model. The ranges to seven points along this arc are plotted in the following figure.




Range histories from 7 distinct points along an arc of ambiguity within the target model space. All 7 curves exactly coincide throughout the data collection.

There is only one curve apparent in the figure, because all seven curves are exactly the same. Hence, the radar data will not contain information with which to distinguish the locations of the seven scattering centers.

More generally, the complete set of arcs, through the target space, all of which are parallel to the illustrated arc, forms a foliation of the target space, in which leaf of the foliation (any individual arc) contains a continuum of possible scattering center locations, any of which will contribute exactly the same signal to the radar data. Hence, the complete set of available information about the size and

shape target is still two-dimensional, not three-dimensional. In fact it can still be represented as an image in a two-dimensional slant plane. The only difference from the far field approximation case is that, in the near field, three-dimensional scattering center locations “project” into the slant plane by sliding along the arcs of the foliation rather than along the normal vectors to the slant plane. Even in the near field, to be able to put unambiguous three-dimensional target size and shape information into monostatic radar data, the target must deviate from a straight line path.

The orthographic projections in the following image illustrate an example of the ambiguity. The set of blue scattering center locations are all on the surfaces of the target model. But reflectors at the set of red scattering center locations, traveling as a rigid body along the same trajectory, would generate exactly the same radar signal history. The red set was chosen by choosing a point at random from each leaf of the foliation which contains a blue point.



Orthographic views of scattering center locations attached to the target (blue) and an alternative set of rigid scattering center locations (red), which would reflect the same radar data set from a linearly moving target.

For an algebraic proof, recall that h is the height of the sensor. Recall that each scattering center travels from some starting position, (x, y, z) , along the path, $(x(t), y(t), z(t)) = (x + vt, y, z)$. For any particular scattering center starting position, (x_1, y_1, z_1) , let $R_1 = \sqrt{y_1^2 + (z_1 - h)^2}$. Then let $(x_2, y_2, z_2) = (x_1, R_1, \cos(\theta), h + R_1, \sin(\theta))$, for any arbitrary value, θ . Under the prescribed rigid target motion, the first point will traverse the linear trajectory, $(x_1 + vt, y_1, z_1)$, and the second point will traverse the linear trajectory, $(x_2 + vt, y_2, z_2)$. At any time, t , the range from the sensor, at location, $(0,0,h)$, to the first scattering center location will be

$$\begin{aligned} r_1(t) &= \sqrt{x_1(t)^2 + y_1(t)^2 + (z_1(t) - h)^2} \\ &= \sqrt{(x_1 + vt)^2 + y_1^2 + (z_1 - h)^2} \end{aligned}$$

and the range from the sensor to the second scattering center location will be

$$\begin{aligned}
r_2(t) &= \sqrt{x_2(t)^2 + y_2(t)^2 + (z_2(t) - h)^2} \\
&= \sqrt{(x_2 + vt)^2 + y_2^2 + (z_2 - h)^2} \\
&= \sqrt{(x_1 + vt)^2 + (R_1 \cos(\theta))^2 + ((h + R_1 \sin(\theta)) - h)^2} \\
&= \sqrt{(x_1 + vt)^2 + (R_1 \cos(\theta))^2 + ((R_1 \sin(\theta)))^2} \\
&= \sqrt{(x_1 + vt)^2 + R_1^2(\cos(\theta)^2 + \sin(\theta)^2)} \\
&= \sqrt{(x_1 + vt)^2 + R_1^2} \\
&= \sqrt{(x_1 + vt)^2 + (\sqrt{y_1^2 + (z_1 - h)^2})^2} \\
&= \sqrt{(x_1 + vt)^2 + y_1^2 + (z_1 - h)^2} \\
&= \sqrt{x_1(t)^2 + y_1(t)^2 + z_1(t) - h)^2} \\
&= r_1(t)
\end{aligned}$$

So for any value, θ , and any time, t , the signal from the second scattering center position will be the same as the signal from the first scattering center position.

The fact that the linear motion scenario generates a signal data set that does not support three-dimensional imaging really should not seem surprising. How else, given data generated from such a situation, could we have used the two-dimensional migration program to form a perfectly focused two-dimensional image of the simulated data (which was made from a fully three-dimensional target model)? If there were genuine three-dimensional scattering center location information in the simulated data, the two-dimensional migration algorithm would be unable to fully focus all the signal energy.

The target's perspective

It can be easier for to analyze some aspects of the situation from the inertial reference frame of the target. Imagine an observer inside the target, who does not see out, but it who is able to continuously measure the direction and distance from which the radiation is arriving. This observer can only quantify these directions of radiation arrival with respect to some reference frame which is defined by objects internal to the target. From the observers world view, he and his reference frame are stationary, and changes in the direction of arrival of the radiation must be attributed to movement of the external source of the radiation.

For our scenario the observer will find that the aspect, elevation, and distance from which the radiation comes all change over time. But despite those changes in aspect and elevation, all the illumination directions will fit in a common plane, containing the origin defined by the internal observer.

6. Description of the McCorkle Antenna

The antenna that was designed on this program was intended to demonstrate that designs exists that could fit in very small form-factors, be conformal, and still have the gain and other characteristics needed for this type of application. A printed antenna of a cardioid type was custom-designed by John McCorkle for this program. This

type of antenna can be both broadband, and have reasonable flat gain over that band. The physical antenna produced is shown in Figure 48 below.




Figure 48: Antenna pair of cardioid type produced by John McCorkle.

This shows two such antennae mounted on a backplane. The overall package measures 11"x7"x3". The antenna frequency dependent gain and voltage standing wave ratio (VSWR) performance is shown in Figure 49 and Figure 50 below.




Figure 49: Frequency dependent gain plot of McCorkle antenna pair.

It can be seen that the gain of this antenna in the band above 800MHz is 6dB. The VSWR ranges from 1.25 to 2.5 over this same band. This means that the efficiency of this antenna is relatively low, but is generally acceptable. Because of the poor match at some frequencies, it is important to minimize cable lengths to avoid having reflected power appear in ranges of interest. This type of antenna can be tuned to be very efficient in bands




Figure 50: Frequency dependent SWR of McCorkle antenna pair.

7. Analysis of Akela RADAR for Collection Suitability

The purpose of this section is to document the procedure used to establish the signal-to-noise ratio (SNR) and noise power for the Akela2 radar unit. The lack of information regarding the noise figure, or knowledge of the internal structure of the receiver means a straight forward derivation from the radar equation is not available. For this reason, measurements were made of a target with known radar cross section (RCS), at a known range, and the system noise power was estimated using the radar range equation. Using this established noise power, SNR curves vs. range were generated for targets of various cross section in order to ensure that adequate SNR could be realized in upcoming data collection experiments given the required ranges and RCS.

Signal-to-Noise Ratio

To make reliable radar measurements, the received signal power must be significantly above the receiver noise. This is usually quantified in terms of a signal-to-noise ratio. The standard equation used for monostatic radar measurements is:

$$SNR = \frac{P_T \tau / PRI \cdot G_T \cdot \sigma \cdot G_R \cdot \lambda^2}{4\pi R_T^2 L_T \cdot 4\pi R_R^2 \cdot 4\pi k_B T_{sys} B_n F L_S}$$

where

P_T	average transmit power
τ	transmit pulse width
PRI	pulse repetition interval
R_T	range to target
G_T	transmit antenna gain
G_R	receiver antenna gain
σ	radar cross section

L_T	transmitter loss
F	receiver noise figure
L_S	system losses
λ	wavelength
G_{sp}	processing gain
k_B	boltzman's constant
T_{sys}	system temperature
B_n	noise bandwidth

This can be generalized into the form:

$$SNR = \frac{P_s}{4\pi R_T^4 P_n}$$

where,

P_s	received signal power
P_n	receiver noise power

and,

$$P_s = P_T \frac{\sigma}{4\pi R^4} \frac{G_T G_R \sigma \lambda^2 G_{sp}}{L_S}$$

$$P_n = k_B T_{sys} B_n F$$

By simply substituting the appropriate parameters into the above equation, the Signal-to-Noise ratio (SNR) can be predicted to an accuracy commensurate with the combined accuracy of the parameters entered. Unfortunately, the internal details of the Akela2 radar hardware are unknown, so the receiver noise figure and noise bandwidth are not known.

However, by making a simple, monostatic measurement of a target of known radar cross section, at a known range, with antennas of known gain, and using a known processing gain, the system noise power (P_n), consisting of the $k_B T_{sys} B_n F L_S$ term in the equation, can be calculated from the measurement results.

To begin with, it is important to make a worst case estimate in order to establish bounds on the operational range for targets of various expected radar cross sections. Further, it is important, in the near term, to be able to estimate this performance for the same conditions under which the radar will operate, meaning that the same parameters, such as transmit and receive gate timing, should be used in the measurement to establish the noise power, and hence included in the parameters of the model. To do this, the generalized equation above can be simplified somewhat.

First, the angle dependent antenna gain terms can be replaced by constant values that reflect the worst case expected gain for both the transmit and the receive antennas. In the noise power measurement the target will be located on the boresight of the antennas. Further, the antenna beams to be used in the data collections have extremely broad beams, and will thus have little variability in their gain patterns within the scene. For upcoming experiments, the Large McCorkle antennas will be employed, and have a peak gain of 6 dBi. In order to consider the worst case conditions in generating the predicted SNR performance, the -1 dB point of the

pattern (5 dBi) will be used as the intended collection scenarios will likely limit scene coverage to this area. Both antennas will be the same, so this value will be used for both transmit and receive antenna gain.

The transmitter loss term, L_T , will be combined into the average transmit power term, P_T , with the understanding that any losses due to cabling, connectors, etc. between the radar output and the transmit antenna must be reflected in the P_T term in the equation. The Akela2 is capable of putting out 50mW CW power (nominally). To account for the cables used in the experiments, 1 dB of loss will be included. Also, in order to account for some variability over the band, 1 dB of loss will be included so that the P_T term will be 31.5 mW (15 dBm).

Also, since the measurements are intended to be made using the transmit and receive gating, the power will be reduced by the duty factor. For these experiments, the transmit "on" time is 30 nsec, while the time between pulses, or pulse repetition interval (PRI) is 171 nsec.

The wavelength will be estimated from the middle of the frequency band. Since for the expected data collections the band is expected to cover 500 to 3000 MHz, the wavelength used here will be 17.1 cm.


The processing gain in the equation consists of the usual time bandwidth product realized from the range compression process, which is the number of FFT points processed (also the number of frequency steps in the waveform). The desired transmitted bandwidth as stated above is 2500 MHz, and for the calibration measurements in the chamber, the ambiguous range is set to 30 meters. In the field experiments, the ambiguous range will be set to 100 meters. This means that frequency step size in the stepped linear FM radar signal for the chamber measurements is 5 MHz, resulting in 501 steps/FFT points. For the field experiment, the frequency step size will be 3 MHz and there will 834 steps, hence 834 FFT points for the range compression.

While a target in the chamber can be well controlled as to its range placement, in the field experiments the targets may be located proximate to a range resolution cell boundary, and may therefore spread the target energy over 2 range resolution cells. To compensate for these effects, a factor of 2 processing gain reduction is introduced to account for the spreading across range resolution cells in the SNR prediction for the field experiments, or 417.

System receiver losses are estimated to be restricted to cable losses, and at these low frequencies, are expected to be limited to 1 dB.

To estimate the noise power for the Akela2 radar, a measurement was taken of a sphere with RCS = -17.8 dBsm at a range of 6.1 meters in the MTRI antenna chamber. Clutter in the chamber environment due to the imperfect absorber at the low frequency of the operation, and primarily due to returns from the metal ceiling of the chamber which cannot be covered with absorber due to fire code restrictions, as well as the internal noise generated from feed-through of the transmitted signal internally and at the antennas, requires a background subtraction process in order to enhance the target return. Another measurement of the chamber with sphere removed was taken, and coherently subtracted from the sphere measurement. The resulting range response (averaged over all the sweeps taken) is shown in

Ave Data Rng Profile - Ave BG Rng Profile from Rotary Platform with_sphere_lower_forbg4_data



Figure

51.

Measuring from the peak of the target at 6.1 meters to the nearby surrounding “noise” resulted in an estimated SNR for this target of 14 dB. The measurement was done by eye, and could possibly be improved by integrating the noise over several range bins on either side of the target return (this will be considered for future work). This number was plugged into the rearranged range equation to estimate the system noise power. The parameters used for this estimate are summarized below:

$P_T = 31.5 \text{ mW}$	average transmit power
$\tau = 30 \text{ nsec}$	transmit pulse width
$PRI = 171 \text{ nsec}$	pulse repetition interval
$R_T = 6.1 \text{ m}$	range to target
$G_T = 5 \text{ dBi}$	transmit antenna gain
$G_R = 5 \text{ dBi}$	receiver antenna gain
$\sigma = -17.8 \text{ dBsm}$	radar cross section
$L_S = 1 \text{ dB}$	system losses
$\lambda = 17.1 \text{ cm}$	wavelength
$G_{sp} = 510$	processing gain

And the resulting “measured” noise power for the system is calculated to be:

$$P_n = 3.0968 \times 10^{-6} \text{ mW.}$$

Ave Data Rng Profile - Ave BG Rng Profile from Rotary Platform with_sphere_lower_forbg4_data




Figure 51: Range Profile of -17.8 dBsm Sphere with (blue curve) and without (red curve) chamber background subtracted

Using the noise power above and the parameters for the expected field collections as specified below (discussed above), a set of SNR vs. range curves for several different RCS targets were generated and are shown in




Figure 52. Note that transmit power reflects the use of a power amplifier that boosts the output power to a peak of 500mW that will be employed in the field measurements.

$P_T = 315. \text{ mW}$	average transmit power
$\tau = 30 \text{ nsec}$	transmit pulse width
$PRI = 171 \text{ nsec}$	pulse repetition interval
$R_T = 6.1 \text{ m}$	range to target
$G_T = 5 \text{ dBi}$	transmit antenna gain
$G_R = 5 \text{ dBi}$	receiver antenna gain
$\sigma = -17.8 \text{ dBsm}$	radar cross section
$L_S = 1 \text{ dB}$	system losses
$\lambda = 17.1 \text{ cm}$	wavelength
$G_{sp} = 834$	processing gain




Figure 52: Predicted SNR vs. Range Performance for the Akela Radar with 500 mW Power Amp Based on Noise Power Calculated from Sphere to Chamber Noise Level

Note that the “measured” SNR that resulted in the noise power measure above is actually based on a signal-to-clutter level, and may not accurately depict expected performance under field conditions where the clutter is far more benign than that of the chamber. A different, potentially more accurate, measure of SNR was sought. The radar was operated with the same parameters as for the calibrated sphere measurement above, with the transmit and receive ports terminated by 50Ω loads. The resulting data were averaged over each frequency bin, and this average was coherently subtracted from the data to remove the feed-through signal, yielding a better estimate of the noise level. Since the radar parameters were the same, the resulting range-compressed outputs are on the same scale, and the SNR of the sphere can be estimated by measuring the noise level from this curve (




Figure 53).




Figure 53: Compressed Range Profiles with Average Range Profile Coherently Subtracted for Aklea Radar with All Ports Terminated

From

Residual Range Profile from Antenna Chamber int_noise_port_1T_2R_10_28_2010_data
Range Profile - Average Range Profile




Figure 53, the SNR of the calibrated sphere target is estimated to be 28 dB. Using this to calculate the noise power yields $P_n = 1.2329 \times 10^{-7}$ mW, and with the same parameters as above, the SNR vs. range curves are re-calculated and shown in Figure 54 below.




Figure 54: Predicted SNR vs. Range Performance for the Akela Radar with 500 mW Power Amp Based on Noise Power Calculated from Sphere to Terminated Radar Noise Level

Recently, it was found that there is a non-linear relationship between the Akela radar's noise figure and the receiver gain value. This gain is in an amplifier downstream from the LNA. The following statement is from Allan Hunt from Akela: "For the case where you are using the pulse modulator with the LNA, the SNR improves as you move from a gain setting of 0 to 6 with the improvement being about 10 dB at gain 6. Above that, the SNR starts to decrease."

Unfortunately, the above model and measurements assumed a more well behaved linear relationship in the receiver with any gain being equally applied to the signal and noise. This is clearly erroneous, and based on the above statement, and that the measurements here were taken with gain values of 8 (= 24 dB), it is estimated that the noise power calculated above will be reduced by 5 to 6 dB, and the SNR curves would be improved by the same amount. It is clear that more measurements will be required to get a more accurate model for the Akela SNR. These will be undertaken as cost/schedule restrictions permit.

8. Summary

This report outlines the progress made through December 2010 of the program. During that reporting period, simulations and local collections occurred. Both the simulations and local data collections provided useful data for the algorithm development tasks. In addition, a new antenna system was designed, tested, and fabricated for use in the field collections which will occur in January, 2011.

Progress was made on the signal processing algorithm. 3D reconstructions of simulated and locally collected data were created. In the instance of far field scenarios, the 3D scattering center estimates were quite good.

Our plan is to collect true river boat data in the coming period and perform 3D reconstruction on that data. The data collection will occur in January 2011 on the Pearl river in southern Mississippi



Information-Theoretic Analysis & Performance Bounds for Super- Resolution (SR) Video Imagery Reconstruction

Quarterly Report

Volume II

Date: January 26, 2011

Contents

1	Introduction	1
2	Progress in the Fourth Quarter of 2010	2
3	Mathematical Framework	7
3.1	Forward Model	7
3.2	General Bayesian Cramèr-Rao Bound	7
3.2.1	Translations only	10
3.2.2	Performance bounds for the translation only case	12
4	Summary and Plans	12
	Appendices	14
A	Mathematics: Flat-earth approximation to projective transform	14
B	Parameterizing Quadratically Penalized Spaces	15

Vol. II: Information-Theoretic Analysis & Performance Bounds for Super- Resolution (SR) Video Imagery Reconstruction

1 Introduction

Video and wide-area surveillance imaging systems acquire massive amounts of data owing to demands for high-frame rates and large fields-of-view. Collected data must be transferred to analysts through channels with finite transfer rate capacities. The large mismatch between the volume of the data collected and the rate at which it can be transferred results in long delays between data acquisition and data analysis. We are proposing a joint collection-processing method where the volume of collected data can be reduced while maintaining large fields-of-view (FOV) and high video frame rates by acquiring fewer pixels, i.e., subsampling. Subsampling by a factor of 2 or 4 (equivalent to reducing resolution by a factor of 2 or 4), and maintaining the same field-of-view, results in an overall reduction in data volume by factors of 4 and 16, respectively. Super-resolution image reconstruction algorithms provide a means of recovering fine detail from low resolution imagery, particularly video imagery. That is, in principle, data volumes can be reduced without sacrificing resolution. Alternatively, SR image reconstruction also provides a means for increasing imaging system FOVs while maintaining resolution and current data sizes. The limits of fine detail recovery from lower-resolution (aliased) imagery depend on various factors:

- Scene information content
- Sub-sampling factor (relative to Nyquist)
- Optical and motion-induced blur (known and unknown)
- Signal-to-noise ratio
- Camera frame rate

In the project, we are developing an information-theoretic framework for super-resolution that will provide quantitative limits of achievable super-resolution image reconstruction under general imaging conditions (including "warping" effects of a changing camera perspective) as a function of key variables.

Our approach for development of this formal theory centers on deriving appropriate models for (i) estimation-theoretic lower bounds (e.g., Cramer-Rao bounds) and (ii) Shannon mutual-information between fine-resolution scenes and collected data, with the preliminary emphasis in the early part of the project being the former. Both of these provide a rigorous framework for quantitatively assessing performance of super-resolution (SR) video imagery reconstruction methods in recovering the desired fine-resolution detail. Using these models, we will be deriving bounds (for either estimation accuracy or mutual information) as a function of key scene and acquisition parameters. The bounds will incorporate spatial frequency and temporal domain weightings that emphasize the particular goal, e.g., visual fidelity, or registration, or subtle change detection. The main focus

of the present project is for video monochromatic and/or color data but the framework is flexible enough to encompass hyper-spectral and polarimetric imagery as well.

The organization of this is as follows: the next Section 2 provides a description of the progress for the just-completed quarter of the project. Section 3 provides a more detailed description of the mathematical framework and the associated information-theoretic Cramer-Rao bounds. Finally Section 4 provides a summary of accomplishments and plans for work in the upcoming quarter.

2 Progress in the Fourth Quarter of 2010

A high-level list accomplishments for the quarter just completed is given in Table 1.

Information content metrics for video imagery:
1. Developed general theory for Bayesian Cramer-Rao bounds
2. Incorporation of optical/detector blur, aliasing/sampling, and noise models
3. Matlab implementation complete
4. Translation only case – “arbitrary” image sizes
5. General case – smaller images
In simulation compared theoretical metrics with restoration accuracy:
1. Translation-only case/larger images
2. Full perspective case/smaller images
3. Various levels of sampling and SNR
4. Results look quite promising

Table 1: Accomplishments of Video Super-Resolution Project - Quarter

We have carried out mathematical development for the Bayesian Cramer-Rao bounds (BCRB) for video imagery – the full mathematics of much of this framework is included in Section 3 and the appendices. In particular, we have developed a mathematical framework for deriving estimation-theoretic bounds on image restoration performance as a function of number of frames, aliasing, signal levels, and camera movements during the collection, and thereby provide our main metric for image restoration quality/information. The mathematical analysis for the single channel is complete and documented within a memo/paper form, and a significant portion of this is included in the Section 3. A significant feature of our model is that it is totally general with a forward imaging model that includes optical blur, detector blur, aliasing, Poisson shot noise, and detector noise. We have implemented two forms of the single-channel BCRBs in MATLAB. One is an image-domain method and allows for a general camera perspective – the limitation of this approach is that it is only applicable to smaller 100x100 chip sizes which may be useful for smaller regions of interest applicability. We have also implemented a frequency-domain method – it is able to operate on full scenes, with the scene size only limitation being that of memory, but it is currently limited to only translation methods. Current research is focused on deriving extension of these frameworks to the more general case of large scene sizes and realistic evolution of camera perspectives that

are more general than translations. The main challenges here are computational and that is the reason for the significant research effort in this regards. We have derived a full Bayesian maximum a-posteriori (MAP) image dealiasing and restoration algorithms for multi-frame video – the output results from this appear to be quite good. The current version of this algorithm relies on knowledge of the camera perspectives, and has the nice feature that the estimators have a closed form and do not require an iterative optimization. It is important to note that these models could be appropriate in very well-controlled scenarios where the approximate camera perspectives are well-known. As mentioned above, our mathematical models/framework is general in that it allows for known or unknown camera perspectives, but we have not implemented the general unknown-perspective case in code and tested it out. Figure 1 shows an an example of a de-aliased restoration – on the left are the original images and on the right is the restoration.




Figure 1: Example restoration using the frequency domain method. Left: Four video frames sampled at 0.25 the Nyquist rate. Right: Restoration using the derived closed form estimator.

Another accomplishment on this code is that we extended the regularization components of the dealiasing image restoration code to allow for a general regularization matrix (Bayesian Prior) on the image – this allows us to regularize to the content of the imagery and/or to the goals of the restoration, i.e., focused more on high-frequency information (e.g., edges) as opposed to lower frequency information (e.g., texture). We also derived a computationally efficient method for finding the appropriate magnitude of the regularization coefficient that corresponds to generating minimum mean-squared error restorations. This is useful, as it is in some sense the “optimal” choice with respect to squared error despite being a function of the unknown. We have also begun exploration of a blind regularization selection approach, and demonstrated its efficacy as compared to other current techniques (e.g., generalized cross-validation – GCV) on some simpler problems and are working on extending to our more general situation. This development is very important since the current data-adaptive techniques for selecting the regularization in single/multi-frame image restoration are not a reliable, and this is especially true in the more challenging scenarios of doing significant de-aliasing. To have a reliable technique in this area would be a big enabler for doing accurate and informative image restoration.

The Bayesian Cramer-Rao bound analysis is designed around having a known point-spread-function (PSF) model for the imaging sensor. This is a reasonable working assumption only when relatively accurate estimates of the PSF can be obtained. Along these lines, we generated an algorithm/approach for doing estimation of the system blur utilizing a known target that is relatively optimized for this application. In many systems one is able to collect against a known target before deploying a sensor and this algorithm aids in the design of that target to ensure accurate PSF estimates. Our procedure/framework utilizes a class of targets (2D extensions of the famous Cantor sets in Mathematics) and these were chosen so as to provide significant robustness to mismatches in the assumed sensor model (uncertainty in focal-length, detector pitch, etc). We have generated a MATLAB code that automatically accounts for target resolution (and includes printer capabilities) and suggests a target (both its content, and its scale) to ensure successful experiments. Once you have collected data off of this target, our estimation model jointly estimates the unknown PSF subject along with a perspective mismatch that will always be there and needs to be accounted for. This latter feature allows for a user to not require that the target be perfectly aligned to the sensor's line-of-sight, and allows for reasonable errors in the assumed sensor parameters. The approach was tested and verified on a Sony Power-Shot. Figure 2 shows a version of an original, somewhat classical target typically utilized (on the left) and on the right is the new target which we propose. Via simulations, we have shown our target is a significant improvement over the classical for PSF estimation and it is also much more robust in operations, due to the multi-resolution properties of the target.




Figure 2: (Left) Classical resolution target – These targets tend to be sensitive to collection range due to their concentration of spatial-frequency energy . (Right) Suggested target – This target is designed to avoid aliasing in the printing process, but has asymptotically uniform spatial-frequency content

Our generalization of the image reconstruction algorithms and information metrics to imagery taken at an arbitrary sampling rate requires an “accurate” generalization of the aliasing operator (sampling of collected image data at a rate that is sub-Nyquist) and the “adjoint” of that aliasing operator, where this latter operator is similar to the adjoint/transpose in matrix theory. An accurate model

of both of these are needed, and as the work progressed, we recognized that our current models and implementations were not sufficiently accurate enough for the demand of future applications, especially as we move to the full-video dealiasing implementation. So, we focused on developing, implementing, and verifying aliasing and aliasing-adjoint operators that can generate images acquired at any sampling rate Q from an image acquired at sampling rate P , with the only constraint being $2 \geq P > Q$. As mentioned earlier, the details in developing this operator are quite subtle, and previous implementations had errors, some significant and some not so significant. The physical mechanism that is being simulated can be described in the following way. Given an image acquired with a detector array consisting of $N \times M$ pixels, each of size $R \times S$ microns, with sampling rate $P1$; generate the image that would have been acquired by a detector of sampling rate $P2$ constrained to the same physical array dimensions of $NR \times MS$ microns. We accomplished this by analyzing the aliasing process in the Fourier domain. In the Fourier domain, the equivalent spectrum of the image sampled at rate $P2$ is obtained by appropriately shifting and summing the frequency components of the original image so as to maintain the maximally smooth interpolated signal of the object. The operator was generalized to include (i) resampling with and without aliasing and (ii) energy preservation under both the $L1$ and $L2$ norms. In addition to this “forward” aliasing operator, we also derived and implemented the full adjoint. The adjoint operator is quite complicated, but it is required for computation of the information metrics as well as gradient-based optimization routines for image reconstruction. The subtle, but important changes in the operators are essential estimating “into the tail” of an objects spectral content, which we need to do good adaptive generation of information quality metrics from actual collected data. In Figure 3 is an example of the data as generated with this new code at 5 different sampling factors, with $Q=2$ being Nyquist and the rest corresponding to different undersampling factors.



Figure 3: Equivalent images sampled at (from left to right) $Q=2$, 1.2, 0.85, 0.47, and 0.27

We also developed a mathematical framework/approach for the adaptation of the current approach to information content metrics for color/multispectral image data. The main focus here is on the transformed space of luminance and chrominance. We do not believe that this transformation/framework is the final answer in the complex and dynamic world of color imagery, since there are competing demands for “visual” fidelity as seen by a user and “quantitative” fidelity as demanded by an exploitation algorithm. But it is a very good start, and it is expected that the “prior” information as would be used in the information metrics and in the restoration algorithms will be tailored to these transformed channels. Again, the flexibility of our framework (arbitrary weightings in either the spectral or spatial domain) allows us the ability to emphasize either visual fidelity or exploitation-based fidelity within the video color image restoration. This is important

since there are such a wide and varied set of criteria in this regards. We started the mathematical modeling of the information for the full with two preliminary concepts having been identified – the first is based on the use of sparse matrix representation of the warp operator and the second is based on doing a stitching together (essentially aggregation) of subimage information metrics. We are in the process of fully writing out the mathematical framework of both, and starting the implementation should begin next month. Also, we improved the accuracy of our models and simulations by the incorporation of broadband PSF into forward and inverse model. Real gray-scale imagery is composed of light from across the visible spectrum, so we added more accuracy (i.e., realism) to our forward and inverse models by developing tools that can automatically generate broadband point-spread functions consisting of an arbitrary weighting of spectral components. This tool can now be adapted to match any single-channel image (including component channels of a multispectral image). Figure 4 below shows a version of two point-spread-functions (PSFs), with the one on the left corresponding to a broadband PSF over 500-700 nm range and the one on the right corresponding to a narrowband PSF of 600 nm.




Figure 4: Left - PSF created using 100 spectral components spanning 500-700 nm. Right – Monochromatic PSF at 600nm. Both PSFs were created with 0.2 waves of RMS wavefront error at 600nm.

We have carried out two different models, one in the so-called Fourier domain and one in the image (spatial) domain, and did a formal and rigorous verification that these two methods are equivalent for the case rigid translations for a wide variety of targets and scenarios. The verification of the approximate equivalence of the two methods was an important consistency check that needed to be done and it was found that the information metrics and image reconstruction results were $\sim 3\%$ relative difference. We did not expect the two methods to yield identical answers due to numerical round-off differences between the methods and this level of error is quite acceptable for this stage of the research, and moving forward with our algorithm development efforts.

3 Mathematical Framework

In this section, we provide some of the background mathematics that are critical to modeling the forward propagation model of video data, and the inverse estimation framework that we have developed based on this.

3.1 Forward Model

To begin the section, we first define a number of symbols utilized throughout this section.

Definition of symbols

$\mathbf{f}$ = Nyquist-sampled object ($N_f \times M_f$)	$\mathbf{d}_j$ = jth video frame ($N_d \times M_d$)
$\mathbf{W}$ = geometric warp operator	$\mathbf{v}_j$ = geometric warp parameters for jth frame
$\mathbf{B}$ = space-invariant blur operator	$\mathbf{G}_f$ = guardband ¹ on a grid of size ($N_f \times M_f$)
$\mathbf{A}$ = aliasing operator	$\mathbf{G}_a$ = guardband ² on a grid of size ($N_d \times M_d$)
$\boldsymbol{\varepsilon}_j$ = noise vector for the jth video frame	

For simplicity, we adopt an iid Gaussian noise model, the jth frame in the video sequence is written

$$\mathbf{d}_j = \mathbf{G}_a \mathbf{A} \mathbf{G}_f \mathbf{B} \mathbf{W}(\mathbf{v}_j) \mathbf{f} + \boldsymbol{\varepsilon}_j. \quad (3.1)$$

For a sequence of J video frames we have

$$\mathbf{d} = \text{diag} \left(\{ \mathbf{G}_a \mathbf{A} \mathbf{G}_f \mathbf{B} \mathbf{W}(\mathbf{v}_j) \}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f}) \mathbf{f} + \boldsymbol{\varepsilon}, \quad (3.2)$$

where $\mathbf{d} = (\mathbf{d}_1, \dots, \mathbf{d}_J)$, $\text{diag} \{ \cdot \}_{j=1}^J$ is a block diagonal matrix with J blocks, $\mathbf{v}_j = (v_1, \dots, v_q)$, $\mathbf{1}_J = \text{ones}(J, 1)$, $n_f = N_f M_f$, $n_d = N_d M_d$, $\mathbf{I}_{n_f}$ = the $(n_f \times n_f)$ identity matrix, and $\boldsymbol{\varepsilon} = (\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_J)$.

3.2 General Bayesian Cramèr-Rao Bound

To simplify notation we define $\mathbf{H}_v \triangleq \text{diag} \left(\{ \mathbf{G}_a \mathbf{A} \mathbf{G}_f \mathbf{B} \mathbf{W}(\mathbf{v}_j) \}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f})$ where $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_J)$, so our model becomes

$$\mathbf{d} = \mathbf{H}_v \mathbf{f} + \boldsymbol{\varepsilon}. \quad (3.3)$$

For simplicity, we assume $\mathbf{f} \sim N(\boldsymbol{\mu}_f, \boldsymbol{\Sigma}_f)$ where $\boldsymbol{\Sigma}_f$ is positive definite, and $\mathbf{v} \sim N(\boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v)$ where $\boldsymbol{\Sigma}_v$ is positive definite. Under this assumption a negative log-likelihood function is

$$L(\mathbf{f}, \mathbf{v}) = \frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{H}_v \mathbf{f}\|^2 + \frac{1}{2} (\mathbf{f} - \boldsymbol{\mu}_f)' \boldsymbol{\Sigma}_f^{-1} (\mathbf{f} - \boldsymbol{\mu}_f) + \frac{1}{2} (\mathbf{v} - \boldsymbol{\mu}_v)' \boldsymbol{\Sigma}_v^{-1} (\mathbf{v} - \boldsymbol{\mu}_v) \quad (3.4)$$

where $'$ indicates conjugate transpose (or *Hermitian transpose*), and $\|\cdot\|$ is the Euclidean norm on $\mathbb{C}^{n_d}$.

The information matrix corresponding to (3.4) is a sum of information matrices corresponding to the data-fit term and the prior information term, $\mathbf{S}_T = \mathbf{S}_D + \mathbf{S}_P$, where $\mathbf{S}_T$ is the total information matrix, $\mathbf{S}_D$ is the information matrix corresponding to the data-fit term, and $\mathbf{S}_P$ is the information matrix corresponding to the prior term.

The information matrix corresponding to the prior term is straightforward to compute:

$$\mathbf{S}_P = E \left[\nabla^2 \left(\frac{1}{2} \mathbf{f}' \boldsymbol{\Sigma}_f^{-1} \mathbf{f} + \frac{1}{2} \mathbf{v}' \boldsymbol{\Sigma}_v^{-1} \mathbf{v} \right) \right] = \begin{bmatrix} \boldsymbol{\Sigma}_f^{-1} & \mathbf{0}_{(n_f \times Jq)} \\ \mathbf{0}_{(Jq \times n_f)} & \boldsymbol{\Sigma}_v^{-1} \end{bmatrix}. \quad (3.5)$$

The information matrix corresponding to the data-fit term is:

$$\mathbf{S}_D = E \left[\nabla^2 \frac{1}{2\sigma^2} \|\mathbf{d} - \mathbf{H}_v \mathbf{f}\|^2 \right] \triangleq \begin{bmatrix} \mathbf{S}_{ff} & \mathbf{S}_{fv} \\ \mathbf{S}_{vf} & \mathbf{S}_{vv} \end{bmatrix}. \quad (3.6)$$

In both (3.5) and (3.6) the ∇^2 operator has the block structure shown in figure 5.




Figure 5: Block structure of ∇^2 .

We are ultimately interested in the upper-left block of the inverse of the total information matrix, $\mathbf{S}_T = \mathbf{S}_D + \mathbf{S}_P$. By applying a form of the *matrix inversion lemma* we have

$$[\mathbf{S}_T^{-1}]_{ff} = [(\mathbf{S}_{ff} + \boldsymbol{\Sigma}_f^{-1}) - \mathbf{S}_{fv}(\mathbf{S}_{vv} + \boldsymbol{\Sigma}_v^{-1})^{-1} \mathbf{S}_{vf}]^{-1}. \quad (3.7)$$

Next we apply the *Woodbury* formula to obtain:

$$\begin{aligned} [\mathbf{S}_T^{-1}]_{ff} &= \underbrace{[\mathbf{S}_{ff} + \boldsymbol{\Sigma}_f^{-1}]^{-1}}_{\text{object independent}} \\ &+ \underbrace{[\mathbf{S}_{ff} + \boldsymbol{\Sigma}_f^{-1}]^{-1} \mathbf{S}_{fv} \left([\mathbf{S}_{vv} + \boldsymbol{\Sigma}_v^{-1}] - \mathbf{S}_{vf} [\mathbf{S}_{ff} + \boldsymbol{\Sigma}_f^{-1}]^{-1} \mathbf{S}_{fv} \right)^{-1} \mathbf{S}_{vf} [\mathbf{S}_{ff} + \boldsymbol{\Sigma}_f^{-1}]^{-1}}_{\text{object dependent}}. \end{aligned} \quad (3.8)$$

We next give expression for S_{ff} , S_{fv} , S_{vf} , and S_{vv} .

$$\begin{aligned}
S_{ff} &= E \left[\nabla_{ff}^2 \frac{1}{2\sigma^2} \|\mathbf{d} - \mathbf{H}_v \mathbf{f}\|^2 \right] \\
&= \frac{1}{\sigma^2} \mathbf{H}_v' \mathbf{H}_v \\
&= \frac{1}{\sigma^2} (\mathbf{1}_J' \otimes \mathbf{I}_{n_f}) \text{diag} \left(\left\{ \mathbf{W}'(\mathbf{v}_j) \mathbf{B}' \mathbf{G}_f' \mathbf{A}' \mathbf{G}_a' \mathbf{G}_a \mathbf{A} \mathbf{G}_f \mathbf{B} \mathbf{W}(\mathbf{v}_j) \right\}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f}) \quad (3.9)
\end{aligned}$$

Next we compute the columns of S_{fv} . Each column is formed by first computing the column gradient of L with respect to $\mathbf{f}$, and then performing partial differentiation with respect to the warp parameters. The columns are arranged according to figure 6. So the $\mathbf{v}_{jk}$ column in the above is

$$S_{fv} = \begin{bmatrix} \vdots & \vdots & \vdots & \vdots \\ \partial_{\mathbf{v}_{1_1}}(\nabla_{\mathbf{f}} L) & \partial_{\mathbf{v}_{1_2}}(\nabla_{\mathbf{f}} L) & \dots & \partial_{\mathbf{v}_{J_q}}(\nabla_{\mathbf{f}} L) \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} \quad \begin{matrix} \leftarrow Jq \\ \uparrow n_p \end{matrix}$$

Figure 6: Format of S_{fv} – note that we are assuming $n_f \gg Jq$.

given by

$$\begin{aligned}
[S_{fv}]_{\mathbf{v}_{jk}} &= E \left[\frac{\partial}{\partial \mathbf{v}_{jk}} \nabla_{\mathbf{f}} \frac{1}{2\sigma^2} \|\mathbf{d} - \mathbf{H}_v \mathbf{f}\|^2 \right] \\
&\text{this is the column vector corresponding to the } k\text{th warp parameter of the } j\text{th frame} \quad (3.10) \\
&= E \left[\frac{\partial}{\partial \mathbf{v}_{jk}} \left(-\frac{1}{\sigma^2} \mathbf{H}_v' [\mathbf{d} - \mathbf{H}_v \mathbf{f}] \right) \right] \\
&= -\frac{1}{\sigma^2} \left(\frac{\partial}{\partial \mathbf{v}_{jk}} \mathbf{H}_v' E[\mathbf{d}] - \frac{\partial}{\partial \mathbf{v}_{jk}} \mathbf{H}_v' \mathbf{H}_v \mathbf{f} - \mathbf{H}_v' \frac{\partial}{\partial \mathbf{v}_{jk}} \mathbf{H}_v \mathbf{f} \right) \\
&= \frac{1}{\sigma^2} \mathbf{H}_v' \left(\frac{\partial}{\partial \mathbf{v}_{jk}} \mathbf{H}_v \right) \mathbf{f}. \quad (3.11)
\end{aligned}$$

Now,

$$\frac{\partial}{\partial \mathbf{v}_{j_k}} \mathbf{H}_v = \frac{\partial}{\partial \mathbf{v}_{j_k}} \text{diag} \left(\{ \mathbf{G}_a \mathbf{A} \mathbf{G}_f \mathbf{B} \mathbf{W}(\mathbf{v}_j) \}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f}) \quad (3.12)$$

$$= \text{diag} \left(\left\{ \mathbf{G}_a \mathbf{A} \mathbf{G}_f \mathbf{B} \frac{\partial}{\partial \mathbf{v}_{j_k}} \mathbf{W}(\mathbf{v}_j) \right\}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f}) \quad (3.13)$$

$$= \begin{bmatrix} 0 & \dots & 0 \\ & \ddots & \\ \vdots & \mathbf{G}_a \mathbf{A} \mathbf{G}_f \mathbf{B} \frac{\partial}{\partial \mathbf{v}_{j_k}} \mathbf{W}(\mathbf{v}_j) & \vdots \\ & \ddots & \\ 0 & \dots & 0 \end{bmatrix} (\mathbf{1}_J \otimes \mathbf{I}_{n_f}). \quad (3.14)$$

So we have

$$[\mathbf{S}_{fv}]_{\mathbf{v}_{j_k}} = \left[\frac{1}{\sigma^2} \mathbf{W}'(\mathbf{v}_j) \mathbf{B}' \mathbf{G}_f' \mathbf{A}' \mathbf{G}_a' \mathbf{G}_a \mathbf{A} \mathbf{G}_f \mathbf{B} \frac{\partial}{\partial \mathbf{v}_{j_k}} \mathbf{W}(\mathbf{v}_j) \mathbf{f} \right]. \quad (3.15)$$

Lastly, noting that $\mathbf{S}_{vv}$ is block diagonal, we give the expression for the j th block of $\mathbf{S}_{vv}$:

$$[\mathbf{S}_{v_j v_j}]_{(n,m)} = \frac{1}{2\sigma^2} E \left[\frac{\partial^2}{\partial \mathbf{v}_{j_n} \partial \mathbf{v}_{j_m}} \langle \mathbf{d} - \mathbf{H}_v \mathbf{f}, \mathbf{d} - \mathbf{H}_v \mathbf{f} \rangle \right] \quad (3.16)$$

$$= \frac{1}{\sigma^2} \left\langle \mathbf{G}_a \mathbf{A} \mathbf{G}_f \mathbf{B} \frac{\partial}{\partial \mathbf{v}_{j_n}} \mathbf{W}(\mathbf{v}_j) \mathbf{f}, \mathbf{G}_a \mathbf{A} \mathbf{G}_f \mathbf{B} \frac{\partial}{\partial \mathbf{v}_{j_m}} \mathbf{W}(\mathbf{v}_j) \mathbf{f} \right\rangle. \quad (3.17)$$

where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product on $\mathbb{C}^{n_d}$. In principle, the inverse of the total information matrix can now be computed. While the diagonal elements of the inverse total information matrix provide mean-square error bounds for individual parameters, a natural global scalar performance measure is the mean of the diagonal elements.

3.2.1 Translations only

In the case where there is only translation, we can drop the guardband operators and so we have that

$$\mathbf{S}_{ff} = \frac{1}{\sigma^2} (\mathbf{1}_J' \otimes \mathbf{I}_{n_f}) \text{diag} \left(\{ \mathbf{W}'(\mathbf{v}_j) \mathbf{B}' \mathbf{A}' \mathbf{A} \mathbf{B} \mathbf{W}(\mathbf{v}_j) \}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f}), \quad (3.18)$$

In the case where there is only translational motion between the video frames, our performance bounds can be computed in the Fourier domain. This can be done via the following. Let $\mathcal{F}_a$ denote the unitary version of the 2-D FFT operator on the aliased grid and let $\mathcal{F}_f$ denote the unitary version of the 2-D FFT operator on the fine grid. Note that $\mathcal{F}_a'$ and $\mathcal{F}_f'$ are the normalized inverse 2-D FFT

operators. For this case we use the following notation for the Fourier counterparts of S_{ff} , $\mathbf{W}$, $\mathbf{B}$, and $\mathbf{A}$:

$$\tilde{S}_{ff} = \mathcal{F}'_f S_{ff} \mathcal{F}_f \quad (3.19)$$

$$\tilde{\mathbf{W}} = \mathcal{F}'_f \mathbf{W} \mathcal{F}_f \quad (3.20)$$

$$\tilde{\mathbf{B}} = \mathcal{F}'_f \mathbf{B} \mathcal{F}_f \quad (3.21)$$

$$\tilde{\mathbf{A}} = \mathcal{F}'_a \mathbf{A} \mathcal{F}_a \quad (3.22)$$

Noting that since $\mathbf{B}$ is block-circulant with circulant blocks, the resulting matrix $\tilde{\mathbf{B}}$ is diagonal and the definitions of our aliasing operator and translation operator are essentially given in the Fourier domain, so we can automatically specify the exact form of $\tilde{\mathbf{W}}$ and $\tilde{\mathbf{A}}$. Putting this with the obvious relation of

$$\begin{aligned} \mathcal{F}_f \mathcal{F}'_f &= \mathcal{F}'_f \mathcal{F}_f = \mathbf{I}_f \\ \mathcal{F}_a \mathcal{F}'_a &= \mathcal{F}'_a \mathcal{F}_a = \mathbf{I}_a \end{aligned}$$

where we note that the $'$ denotes the conjugate transpose. Now for $J = 1$, easy to see that

$$\begin{aligned} \tilde{S}_{ff} &= \mathcal{F}'_f S_{ff} \mathcal{F}_f \\ &= \frac{1}{\sigma^2} \mathcal{F}'_f \mathbf{W}'(\mathbf{v}_1) \mathcal{F}_f \mathcal{F}'_f \mathbf{B}' \mathcal{F}_f \mathcal{F}'_f \mathbf{A}' \mathcal{F}_a \mathcal{F}'_a \mathbf{A} \mathcal{F}_f \mathcal{F}'_f \mathbf{B} \mathcal{F}_f \mathcal{F}'_f \mathbf{W}(\mathbf{v}_1) \mathcal{F}_f \\ &= \frac{1}{\sigma^2} (\mathcal{F}'_f \mathbf{W}(\mathbf{v}_1) \mathcal{F}_f)' (\mathcal{F}'_f \mathbf{B} \mathcal{F}_f)' (\mathcal{F}'_a \mathbf{A} \mathcal{F}_f)' (\mathcal{F}'_a \mathbf{A} \mathcal{F}_f) (\mathcal{F}'_f \mathbf{B} \mathcal{F}_f) (\mathcal{F}'_f \mathbf{W}(\mathbf{v}_1) \mathcal{F}_f) \\ &= \frac{1}{\sigma^2} \tilde{\mathbf{W}}'(\mathbf{v}_1) \tilde{\mathbf{B}}' \tilde{\mathbf{A}}' \tilde{\mathbf{A}} \tilde{\mathbf{B}} \tilde{\mathbf{W}}(\mathbf{v}_1) \end{aligned} \quad (3.23)$$

Note the warp parameters in this case are translation parameters and there are two per video frame. It is easy to generalize the above to the general case of $J \geq 2$, and so the result in this case are that the elements of the data-fit information matrix become (omitting guardband operations):

$$\tilde{S}_{ff} = \frac{1}{\sigma^2} (\mathbf{1}'_J \otimes \mathbf{I}_{n_f}) \text{diag} \left(\left\{ \tilde{\mathbf{W}}'(\mathbf{v}_j) \tilde{\mathbf{B}}' \tilde{\mathbf{A}}' \tilde{\mathbf{A}} \tilde{\mathbf{B}} \tilde{\mathbf{W}}(\mathbf{v}_j) \right\}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f}), \quad (3.24)$$

where $\tilde{\cdot}$ denotes the Fourier representation of the quantities and operators³. We note the operators $\{\tilde{\mathbf{W}}, \tilde{\mathbf{B}}\}$ are diagonal and that the product $\tilde{\mathbf{A}}' \tilde{\mathbf{A}}$ can be permuted into block-diagonal form.

Also,

$$[\tilde{S}_{fv}]_{\mathbf{v}_{jx}} = \frac{1}{\sigma^2} \tilde{\mathbf{W}}'(\mathbf{v}_j) \tilde{\mathbf{B}}' \tilde{\mathbf{A}}' \tilde{\mathbf{A}} \tilde{\mathbf{B}} \frac{\partial}{\partial \mathbf{v}_{jk}} \tilde{\mathbf{W}}(\mathbf{v}_j) \tilde{\mathbf{f}} \quad (3.25)$$

$$= \frac{1}{\sigma^2} \tilde{\mathbf{W}}'(\mathbf{v}_j) \tilde{\mathbf{B}}' \tilde{\mathbf{A}}' \tilde{\mathbf{A}} \tilde{\mathbf{B}} (-2\pi l \text{ramp}_x) \tilde{\mathbf{W}}(\mathbf{v}_j) \tilde{\mathbf{f}} \quad (3.26)$$

$$[\tilde{S}_{fv}]_{\mathbf{v}_{jy}} = \frac{1}{\sigma^2} \tilde{\mathbf{W}}'(\mathbf{v}_j) \tilde{\mathbf{B}}' \tilde{\mathbf{A}}' \tilde{\mathbf{A}} \tilde{\mathbf{B}} (-2\pi l \text{ramp}_y) \tilde{\mathbf{W}}(\mathbf{v}_j) \tilde{\mathbf{f}} \quad (3.27)$$

$$[\tilde{S}_{vv}]_{(\mathbf{v}_{ix}, \mathbf{v}_{jy})} = \frac{1}{\sigma^2} \left\langle \tilde{\mathbf{A}} \tilde{\mathbf{B}} (-2\pi l \text{ramp}_x) \tilde{\mathbf{W}}(\mathbf{v}_j) \tilde{\mathbf{f}}, \tilde{\mathbf{A}} \tilde{\mathbf{B}} (-2\pi l \text{ramp}_y) \tilde{\mathbf{W}}(\mathbf{v}_j) \tilde{\mathbf{f}} \right\rangle \delta_{ij}. \quad (3.28)$$

³Fourier representations of operators are obtained by pre and post multiplication by a forward 2D DFT matrix and an inverse 2D DFT matrix respectively.

Note that $[\tilde{S}_{vv}]$ is block diagonal with (2×2) blocks.

3.2.2 Performance bounds for the translation only case

Recall, the total inverse information matrix consists of two parts: an object independent (or system) part, and a part that depends explicitly on the object:

$$[S_T]_{ff}^{-1} = \underbrace{[S_{ff} + \Sigma_f^{-1}]^{-1}}_{\text{object independent}} + \underbrace{[S_{ff} + \Sigma_f^{-1}]^{-1} S_{fv} \left([S_{vv} + \Sigma_v^{-1}] - S_{vf} [S_{ff} + \Sigma_f^{-1}]^{-1} S_{fv} \right)^{-1} S_{vf} [S_{ff} + \Sigma_f^{-1}]^{-1}}_{\text{object dependent}}.$$

We note that the scale of the object is incorporated into the object-independent part through the prior. All of the above operators have Fourier domain representations, and there is an overall scale factor of $\frac{1}{n_f}$ from the discrete version of Parseval's theorem.

Let

$$\Lambda \triangleq [S_{ff} + \Sigma_f^{-1}]^{-1}$$

$$\xi \triangleq S_{fv}$$

$$\eta \triangleq [S_{vv} + \Sigma_v^{-1}],$$

then

$$\text{Tr} [[S_T]_{ff}^{-1}] = \text{Tr} [\Lambda] + \text{Tr} [\Lambda \xi (\eta - \xi' \Lambda \xi)^{-1} \xi' \Lambda]. \quad (3.29)$$

The second term can be simplified by applying the invariance property of the trace under cyclic permutation:

$$\text{Tr} [\Lambda \xi (\eta - \xi' \Lambda \xi)^{-1} \xi' \Lambda] = \text{Tr} \left[\underbrace{(\eta - \xi' \Lambda \xi)^{-1}}_{Jq \times Jq} \underbrace{\xi' \Lambda \xi}_{Jq \times Jq} \right] \quad (3.30)$$

$$= \text{Tr} \left[\underbrace{(\eta - \xi' \Lambda \xi)^{-1}}_{Jq \times Jq} \underbrace{(\Lambda \xi)' \Lambda \xi}_{Jq \times Jq} \right] \quad \text{by the Hermitian symmetry of } \Lambda. \quad (3.31)$$

4 Summary and Plans

In this report, we have provided a brief discussion and presentation of technical details of the work completed in the previous quarter for the Video Super-Resolution Project. General plans for the

next quarter are the following:

- (a) Continue development of two mathematical/numerical approaches for quantifying information for general case of arbitrary set of camera perspectives – one based on sparse matrices and one based on aggregation of subimage information metrics.
- (b) Carry out software implementation of the “aggregation” information-theoretic approach (for computing the estimation bounds in the general case) in (a) and do preliminary testing.
- (c) Implement and test software implementation of the methodology for adaptively estimating object content for the information metric based on image restoration.
- (d) Significant testing/validation of the current approaches for video restoration and generation of information content metrics for color/multispectral image data, based on the luminance/chrominance decomposition approach as developed in the last quarter.
- (e) Review MPEG noise models and make sure that some realistic version is included as part of the overall information-theoretic framework.
- (f) Complete write-up of full-perspective approaches. Do a staged-implementation of the full large-image restoration capability with true joint estimation of the warp and imagery, and in particular the joint of blur and perspective implementation – carry out initial implementation of this joint object/warp estimation and do testing.

Appendices

A Mathematics: Flat-earth approximation to projective transform

For discrete pairs of images with significant camera rotation, a full eight-parameter projective transformation is required to align the images. However, for closely related views a quadratic transformation is a good approximation to the projective transform and is more stable to compute[Kumar]. Let $\vec{d}(x, y, t_1, t_2)$ be the displacement field describing the motion between images $I(x, y, t_1)$ and $I(x, y, t_2)$. Then, the image at time t_1 warped into alignment with the image at time t_2 is

$$\hat{I}(x, y, t_1, t_2) = I\left((x, y) - \left(\vec{d}(x, y, t_1, t_2)\right), t_1\right).$$

The quadratic approximation to the displacement field is given by

$$d_x = a_1 + a_2x + a_3y + a_7x^2 + a_8xy$$

$$d_y = a_4 + a_5x + a_6y + a_7xy + a_8y^2.$$

The implementation of the quadratic transform involves estimating the parameters $\{a_1, \dots, a_8\}$, creating the grid defined by (d_x, d_y) , then interpolating the image at time t_1 onto that grid. For iterative image restoration methods it will likely be necessary to calculate the partial derivatives of the quadratic transform. Since the quadratic transform is an interpolation it can be written as an operator

$$\hat{I}(x, y, t_1, t_2) = Q(\tau) I(x, y, t_1)$$

where $\tau = \{a_1, \dots, a_8\}$.

$$\begin{aligned}
\frac{\partial}{\partial a_i} Q(\tau) I &= \frac{\partial}{\partial a_i} \sum_n \sum_m I[n, m] h(x(\tau) - n\Delta_x, y(\tau) - m\Delta_y) \\
&= \sum_n \sum_m I[n, m] \frac{\partial}{\partial a_i} h(x(\tau) - n\Delta_x, y(\tau) - m\Delta_y) \\
&= \sum_n \sum_m I[n, m] \left(\frac{\partial h}{\partial x} \frac{\partial x}{\partial a_i} + \frac{\partial h}{\partial y} \frac{\partial y}{\partial a_i} \right) \\
&= \sum_n \sum_m I[n, m] \left(\frac{\partial h}{\partial x} \frac{\partial x}{\partial a_i} \right) + \sum_n \sum_m I[n, m] \left(\frac{\partial h}{\partial y} \frac{\partial y}{\partial a_i} \right) \\
&= \sum_n \sum_m \left(I[n, m] \frac{\partial h}{\partial x} \right) \frac{\partial x}{\partial a_i} + \sum_n \sum_m \left(I[n, m] \frac{\partial h}{\partial y} \right) \frac{\partial y}{\partial a_i}.
\end{aligned}$$

So the partial derivative is the sum of two images: (1) the element-by-element product of the original image interpolated with the derivative of the interpolation kernel and the derivative of the new x -grid with respect to the motion parameter, (2) similar to (1) but replace x with y .

$$\frac{\partial}{\partial a_i} Q(\tau) I = \sum_n \sum_m \left(I[n, m] \frac{\partial h}{\partial x} (x(\tau) - n\Delta_x) \right) \frac{\partial x(\tau)}{\partial a_i} + \sum_n \sum_m \left(I[n, m] \frac{\partial h}{\partial y} (y(\tau) - m\Delta_y) \right) \frac{\partial y(\tau)}{\partial a_i}.$$

B Parameterizing Quadratically Penalized Spaces

This section develops a natural parameterization for quadratically penalized weighted least-squares objective functions, and uses this representation to develop an efficient means of searching over the regularization parameter. Consider the objective function

$$\Lambda(y, \lambda) = (y - Bx)^T W (y - Bx) + \lambda x^T P x \quad (B.1)$$

with the solution given by

$$\hat{x} = \arg \min_x \Lambda(y, \lambda) \quad (B.2)$$

$$= (B^T W B + \lambda P)^{-1} B^T W y \quad (B.3)$$

From the structure of the solution one can see that λ is related to the minimizer $\hat{x}$ through the matrix pencil $\lambda (B^T W B, P)$. The weighted inner-product can be eliminated by taking its Cholesky decomposition and defining the problem over a rotated basis

$$\left. \begin{aligned} W_h : W &= W_h^T W_h \\ B_w &= W_h B \end{aligned} \right\} \implies \hat{x} = (B_h^T B_h + \lambda P)^{-1} B_h^T W_h y \quad (B.4)$$

As a result, the remainder of this appendix will assume $\mathbf{W} = \mathbf{I}$ without loss of generality. There exists a unique rotation of the space such that $(\mathbf{B}^T \mathbf{B} + \lambda \mathbf{P})^{-1}$ is congruent to a diagonal matrix which is a function of λ . To see this, consider the generalized eigen-decomposition of $\mathbf{B}^T \mathbf{B}$ and $\mathbf{P}$

$$\exists \mathbf{V}, \mathbf{D} : \mathbf{P}\mathbf{V} = \mathbf{B}^T \mathbf{B} \mathbf{V} \mathbf{D}, (\mathbf{B}\mathbf{V})^T \mathbf{B}\mathbf{V} = \mathbf{I} \quad (\text{B.5})$$

$$\mathbf{V}^T (\mathbf{B}^T \mathbf{B} + \lambda \mathbf{P}) \mathbf{V} = \mathbf{I} + \lambda \mathbf{V}^T \mathbf{P} \mathbf{V} \quad (\text{B.6})$$

$$= \mathbf{I} + \lambda \mathbf{D} \quad (\text{B.7})$$

$$\therefore (\mathbf{B}^T \mathbf{B} + \lambda \mathbf{P})^{-1} = \mathbf{V} \mathbf{D}_\lambda \mathbf{V}^T \quad (\text{B.8})$$

$$\mathbf{D}_\lambda = \text{diag}([1 + \lambda \mathbf{D}_{11}]^{-1}, \dots, [1 + \lambda \mathbf{D}_{pp}]^{-1}) \quad (\text{B.9})$$

$$= \text{diag}(\mathbf{d}_\lambda) \quad (\text{B.10})$$

from which the estimate may be restated as

$$\hat{\mathbf{x}} = \mathbf{V} \mathbf{D}_\lambda (\mathbf{B}\mathbf{V})^T \mathbf{y} \quad (\text{B.11})$$

$$= \mathbf{V} \mathbf{D}_\lambda \mathbf{a} \quad (\text{B.12})$$

$$= \mathbf{V}(\mathbf{d}_\lambda \odot \mathbf{a}) \quad \text{s.t.} \quad \mathbf{a} = ((\mathbf{y}^T \mathbf{B}) \mathbf{V})^T \quad (\text{B.13})$$

Now consider calculating the MSE with respect to this parameterization

$$MSE(\mathbf{x}, \lambda) = \|\hat{\mathbf{x}} - \mathbf{x}\|^2 \quad (\text{B.14})$$

$$\frac{\partial \mathbf{D}_\lambda}{\partial \lambda} = \text{diag}([- \mathbf{D}_{11} (\lambda \mathbf{D}_{11} + 1)^{-2}, \dots, - \mathbf{D}_{pp} (\lambda \mathbf{D}_{pp} + 1)^{-2}]) \quad (\text{B.15})$$

$$= \text{diag}(\dot{\mathbf{d}}_\lambda) \quad (\text{B.16})$$

$$\frac{\partial}{\partial \lambda} [\mathbf{D}_\lambda \mathbf{a}] = \dot{\mathbf{d}}_\lambda \odot \mathbf{a} \quad (\text{B.17})$$

$$\frac{\partial \hat{\mathbf{x}}}{\partial \lambda} = \frac{\partial}{\partial \lambda} [\mathbf{V} \mathbf{D}_\lambda \mathbf{a}] \quad (\text{B.18})$$

$$= \mathbf{V} \frac{\partial}{\partial \lambda} [\mathbf{D}_\lambda \mathbf{a}] \quad (\text{B.19})$$

$$= \mathbf{V} (\dot{\mathbf{d}}_\lambda \odot \mathbf{a}) \quad (\text{B.20})$$

$$\hat{\mathbf{x}}^T \hat{\mathbf{x}} = (\mathbf{D}_\lambda \mathbf{a})^T \mathbf{V}^T \mathbf{V} (\mathbf{D}_\lambda \mathbf{a}) \quad (\text{B.21})$$

$$\frac{\partial}{\partial \mathbf{D}_\lambda} [\hat{\mathbf{x}}^T \hat{\mathbf{x}}] = 2 \mathbf{V}^T \mathbf{V} (\mathbf{D}_\lambda \mathbf{a}) \quad (\text{B.22})$$

$$\frac{\partial}{\partial \lambda} [\hat{\mathbf{x}}^T \hat{\mathbf{x}}] = (\mathbf{D}_\lambda \mathbf{a})^T \mathbf{V}^T \mathbf{V} (\dot{\mathbf{d}}_\lambda \odot \mathbf{a}) + (\dot{\mathbf{d}}_\lambda \odot \mathbf{a})^T \mathbf{V}^T \mathbf{V} (\mathbf{D}_\lambda \mathbf{a}) \quad (\text{B.23})$$

$$= 2 (\mathbf{D}_\lambda \mathbf{a})^T \mathbf{V}^T \mathbf{V} (\dot{\mathbf{d}}_\lambda \odot \mathbf{a}) \quad (\text{B.24})$$

$$\frac{\partial}{\partial \lambda} [\hat{\mathbf{x}}^T \mathbf{x}] = \left(\frac{\partial \hat{\mathbf{x}}}{\partial \lambda} \right)^T \mathbf{x} \quad (\text{B.25})$$

$$= (\dot{\mathbf{d}}_\lambda \odot \mathbf{a})^T \mathbf{V} \mathbf{x} \quad (\text{B.26})$$

Putting this all together, we get an efficient means of computing both the MSE and its derivative with respect to λ in terms of the results of a single $O(n^3)$ calculation to find the canonical space (a generalized eigen-value decomposition).

$$MSE = \|\mathbf{V}(\mathbf{d}_\lambda \odot \mathbf{a}) - \mathbf{x}\|^2 \quad (\text{B.27})$$

$$\frac{\partial}{\partial \lambda} [MSE] = -2 (\mathbf{V}^T \mathbf{V} (\mathbf{d}_\lambda \odot \mathbf{a}) + \mathbf{V} \mathbf{x})^T (\dot{\mathbf{d}}_\lambda \odot \mathbf{a}) \quad (\text{B.28})$$

Riverine Imaging

by

John McCorkle
Applied Signals Intelligence, Inc.
11501 Sunset Hills Rd., Suite 300
Reston, VA 20190
571-313-0681
corporate@asigint.com

For

Kristin Beck
Assistant Contracts Analyst
Sponsored Programs Office
Michigan Technology University
1400 Townsend Drive
Houghton, MI 49931-1295
kgbeck@mtu.edu

September 30, 2011

Contract No. W911NF-10-2-0069
Subagreement #100541Z1

The views, opinions, and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.

Copyright © 2011 by Applied Signals Intelligence, Inc. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, scanning or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of Applied Signals Intelligence, Inc.

Applied Signals Intelligence, Inc.

11501 Sunset Hills Rd., Suite 300

Reston, VA 20190

571-313-0681

corporate@asigint.com

Table of Contents

Table of Contents	3
Table of Figures	4
1 Executive Summary / Conclusion	6
2 Phenomenology	7
2.1 Background -- The Model	7
2.2 Phenomenology Summary	7
2.3 Phenomenology Details	7
3 System Modeling	13
3.1 Overview	13
3.2 Background on chirp and stretch processing	13
3.3 Geometry and sampling (PRF and number of points) calculation	15
3.4 Transmitter Power Calculations	16
3.5 Case-1 - Tx & Rx in separate time slots	19
3.6 Case-2 0 Tx & Rx simultaneously	19
4 RF Hardware Strawman and Power Evaluation	21
4.1 Exciter Strawman Hardware	21
4.2 Receiver Strawman Hardware	22
4.3 RF Hardware Power	23
5 Processor Strawman Design and Power Evaluation	25
5.1 Overview	25
5.2 Digital System Hardware Design	27
5.3 Main Power Consuming DSP Blocks	28
5.4 CFAR	30
5.5 Baseline Radar to Processor Interface	30
5.6 FPGA processing	31
5.6.1 FPGA-1 1-D pulse compression processing	32
5.6.2 FPGA-2 2-D αk Image Former	34
5.7 Power Consumption Analysis	37
5.8 Digital Hardware Power Summary	38
6 Antenna Design	40

Table of Figures

Figure 1: Multipath Gain/Loss Over Direct Path.....	8
Figure 2: Multipath Gain/Loss Over Direct Path.....	9
Figure 3: Correlation between gain, resolution, and sidelobes versus range	9
Figure 4: Received Waveform (Volts).....	10
Figure 5: Resolution & Sidelobes in dB.....	11
Figure 6: Direct & Multipath Received Waveforms at Peak at R = 134 m	11
Figure 7: Bandwidth Gives Range Resolution	14
Figure 8: De-Ramp ("Stretch") Implementation	15
Figure 9: Scenario Geometry.....	16
Figure 10: Energy Calculation.....	16
Figure 11: Characteristics of various window taper functions	17
Figure 12: System losses	17
Figure 13: Equations used to compute transmit power requirements.....	18
Figure 14: Parameter values used for transmit power calculation.....	18
Figure 15: Peak transmit power calc, Tx & Rx in separate time slots.....	19
Figure 16: VCO linearized & locked in a fractional-N (DDS) configuration.....	22
Figure 17: VCO Linearized & Locked.....	22
Figure 18: Receiver.....	23
Figure 19: Battery Power During Imaging.....	24
Figure 20: DSP versus FPGA drivers	26
Figure 21: Block Diagram.....	27
Figure 22: ARM Cortex M3 "Mom" processor.....	28
Figure 23: Memory	28
Figure 24: FPGAs	28
Figure 25: Reason For Preference For Block Floating Point FFT.....	29
Figure 26: Standard Cell-Average Constant False Alarm Rate (CA-CFAR) Detector	30
Figure 27: Baseline Scenario	31
Figure 28: Data flow pipeline through the 6 memories	32
Figure 29: FPGA-1 Pulse Compression Flow	32
Figure 30: FPGA-1 Pulse Compression Operation Count	33
Figure 31: Spotlight formatting	34
Figure 32: FPGA-1 Pulse Compression Operation Count	34
Figure 33: Wavenumber Migration Processing Partitioned Around Corner-Turns	35
Figure 34: Example of Timing on FPGA-2 During Image Formation	36
Figure 35: Metrics Negligible Processing Load	37
Figure 36: FPGA-1 Power Summary Calculation	37
Figure 37: FPGA-2 Power Summary Calculation	38
Figure 38: Memory Power Summary Calculation.....	38
Figure 39: ARM Power Summary Calculation	38
Figure 40: Total Power Summary Calculation	39
Figure 41: Wide bandwidth antennas are often dispersive.....	40
Figure 42: Setup to measure angle dependent dispersion	41
Figure 43: UWB antennas developed since the FCC issued UWB rules.....	42

Figure 44: Folded Slot.....	43
Figure 45: Face-to-Face Two-Antenna S21 (nose-to-nose S21, NOT gain).....	43
Figure 46: Face-to-Face Two-Antenna S21	44
Figure 47: Tight Impulse Response Versus Orientation	44
Figure 48: Measurement Orientation Nomenclature	45
Figure 49: Co-pol E-plane and H-plane Frequency vs Angle energy plots	45
Figure 50: Cross-pol E-plane and H-plane Frequency vs Angle energy plots	46
Figure 51: Beam Patterns Sealed to 700 MHz, 877 MHz, & 1.4 GHz.....	46
Figure 51: First Antenna Delivered To MTRI.....	47
Figure 51: Second Antenna Showing Cavity Backing.....	48
Figure 52: Trimmed ASI390-T2.5-TG1.8.....	49

1 Executive Summary / Conclusion

ASI pursued 5 areas of study.

Phenomenology – Modeling and measurements were analyzed and showed that HH polarization was preferred and gave an SNR suitable to perform the desired ISAR image formation.

System modeling – System modeling was done to establish the radar requirements.

Low Power Radar RF design – A complete strawman transmitter and receiver were designed using the latest generation SiGe, GaAs, PHEMPT technology to demonstrate that an extremely small radar could be built that required less than 2 watts to operate yet meet all performance requirements.

Low Power Radar Digital design – A complete strawman radar processor was designed using current generation FPGA technology, cell-phone RAM, and mobile processors to show that an extremely small processor could be built that operated on less the 200mW.

Antenna design, construction, and testing - An antenna small enough to fit the intended application was fabricated and proved to have all the performance metrics (bandwidth, low dispersion, beamwidth, and small size) needed to perform the high resolution UWB ISAR function in the required small form factor system.

2 Phenomenology

The purpose of this “quick look” analysis was to understand pro’s and con’s of vertical versus horizontal polarization, and why one might be better than the other.

2.1 Background -- The Model

ASI uses an internally developed a fast modeling code with four key features. It models:

- Multipath, including 1, 2, and 3 bounce terms
- Arbitrary bandwidth/resolution waveforms (such as UWB - Ultra Wide Bandwidth)
- Both SAR (moving radar) and ISAR (moving target) cases, and
- Arbitrary ground/water parameters (ϵ , μ , σ).

Regarding the multipath, it includes the direct (one bounce) path (radar to target and back), 2-bounce paths (radar to target to ground and back plus radar to ground to target and back), and 3-bounce path (radar to ground to target to ground and back). Analysis of UWB radar data shows that these paths are by far the most dominant and accurately model UWB radar data. Regarding the SAR and ISAR capability, it models motion along an arbitrary path.

2.2 Phenomenology Summary

After getting model results, field tests were made that confirmed the predictions made by compressed UWB multipath modeling—there was complete agreement. To summarize:

- The RCS of an excessively small target (small empty plastic canoe, no motor, two people) was ~ 0 dBsm broadside and ~ -15 dBsm at end of a SAR aperture.
- The radar impulse response and resolution changes significantly over the range swath due to the significant geometry changes in the one, two, and three-bounce multipath terms.
- HH is superior to VV due to combination of:
- Clutter from water ripple is suppressed in HH polarization due to multipath
 - ~ 25 dB @ 5cm height and >50 dB @ 1 cm height at 100m slant range
 - ~ 20 dB @ 5cm and 45dB @ 1 cm at 50m slant range
- Target (both boat and contents of interest) are enhanced from multipath in HH pol.
- The raw target RCS is larger in HH – Both boat and contents are oriented to have higher RCS at HH
- HH experiences consistent enhancement over VV for targets above the water
 - With 1m target height
 - HH had gain except from 58-78m slant range
 - Worst case loss was 6 dB
 - With 2m target height
 - HH gain/loss curve stretches right from 1m case
 - HH had gain except from 108-132m slant range
 - Worst case loss was 6 dB

2.3 Phenomenology Details

Low height results, representative of water ripple, are shown in Figure 1, while 1m height and 2m height results are shown in Figure 2. Figure 1 shows there is a strong preference for HH polarization to reduce the clutter caused by water ripple. Figure 2 shows three important features. First, that the typical deep nulls that would happen with narrowband signals is mitigated by the fact that the UWB waveform resolves multipath at the shortest ranges - the worst case null depth is reduced to only 6 dB. Second, that there is signal enhancement over much of the range-swath of interest. Third, that this outcome is caused by that fact that at the shortest ranges, the multipath is resolved by the UWB waveform. At other ranges the multipath is in-phase so the various bounces mostly or completely add together.




Figure 1: Multipath Gain/Loss Over Direct Path




Figure 2: Multipath Gain/Loss Over Direct Path

Figure 3 plots three different metrics as a function of range and for each polarization: (1) the gain over the direct path, (2) the received pulse-width (or radar resolution) at the -10dB points, and (3) the peak sidelobe in the time-domain (i.e. down-range) impulse response. These plots highlight and provide further insight into Figure 2. Four key points, labeled A, B, C, and D, are selected for closer analysis. Note that at range B, there is a convergence of the best sidelobes, a peak in resolution, and a peak in the signal enhancement (or gain over the direct path).




Figure 3: Correlation between gain, resolution, and sidelobes versus range

Figure 4 shows the waveform received and its envelope at range A, B, and C. Figure 5 shows the received resolution “pulse width” and sidelobe levels at those same ranges. At range A, the 1-bounce (closest range), 2-bounce (middle range), and 3-bounce (farthest range) terms are all resolved. Note that the 2-bounce return is twice the amplitude of the other terms. This doubling is because there are two paths that are identical length (i.e. in-phase at all frequencies) that add together—radar-target-ground-radar, and radar-ground-target-radar. At range B, there is an in-phase convergence that gives a high resolution, high amplitude, and low sidelobe return. At range C the return is just the reverse—low amplitude, low resolution, and high sidelobes. Figure 6 shows the waveform received and its envelope at range D. This is the range where the different multipath terms add in-phase. Beyond this point, the three isolated returns seen at point-A become more and more aligned, and the phase of the center term becomes 180-degrees out of phase with the other terms, causing an ultra wide bandwidth null. To summarize, the metrics at these points are:

- A – 40 dB sidelobes and multipath resolved (3 lobes)
- B – 70 dB sidelobes and 9 dB gain enhancement due to multipath
- C – 25 dB sidelobes and 5 dB loss due to multipath
- D – Almost 12 dB of gain – (best possible with 4 paths is 12 dB). Note that it is 12 dB as opposed to 6 dB due to the two-way radar path as opposed to a one-way communications path.




Figure 4: Received Waveform (Volts)




Figure 5: Resolution & Sidelobes in dB




Figure 6: Direct & Multipath Received Waveforms at Peak at R = 134 m

Measurements carried out by MTRI showed exactly what the model predicted. (1) There was significantly more clutter (water ripple) “noise” in VV data relative to HH data. The same preference for HH holds with trees – i.e. tree clutter is reduced because the tree RCS is small for HH and large for VV. (2) HH has enhanced target (hidden object) RCS because target objects are generally oriented along hull - not oriented vertically. In other words, they have long horizontal features and almost no vertical features. The measurements showed that presence of metal in a canoe gave higher return in HH data, and made no difference in VV data. (3) HH has enhanced boat RCS. The SNR on the boat was significantly higher in HH. The boat is primarily horizontal, with much smaller vertical features. Furthermore, the water/hull boundary, both far-side and near-side, is (a) a long horizontal feature, and (b) a sharp corner at the water line where the hull breaks the surface of the water, forming a

depression that behaves as a dihedral "corner" reflector, with an especially large return from the far side of the boat. The water/hull boundary is much shorter vertically. Note that this water/hull boundary is an important loud return since it is very predictable and can be used as part of the reference system for an ISAR. The HH data was always about 10dB louder than VV data. The worst case (small empty plastic canoe, no motor, two people) RCS was ~ 0 dBsm broadside and ~ -15 dBsm at end of a SAR aperture.

3 System Modeling

Objectives of this effort were to (1) derive the system requirements that drive the radar design such as minimum and maximum PRF, transmit power, noise figure, duty cycles, etc., and (2) estimating the power budget and understand radar impact on battery life by (3) doing a strawman design to support the power and size budget estimates.

3.1 Overview

Impulse, poly-phase coded, step-frequency, and chirp are a few of many possible waveforms. The system model considered the time/power versus time budget for three architectures:

- Impulse - Peak power & volts required for Impulse - PRF limited by ambiguities
- Pulsed Chirp – Chirp length limited by short range, PRF limited by ambiguities
- Continuous Chirp –simultaneous transmit and receive, PRF defined by sample rate required by ISAR and target motion and dynamics

At the highest level, power is conserved by putting blocks to sleep (power off). Idle modes (clock off) are also helpful, but even when clocks are off, static power will kill battery life. In other words, the integration of a small power over a long time still comes out to be a big number. Integrating zero comes out zero regardless of the integration time. Therefore, the key to low power operation hinges on optimizing sleep time. Because there is usually a significant time delay between a power-off state to a fully operational state, the system can't use many short sleep time-slots, but must use a few long sleep time-slots.

3.2 Background on chirp and stretch processing

Figure 7 illustrates pulse compression theory and practice. The top row depicts how resolution is proportional to the inverse of the system bandwidth. In other words, the wider the bandwidth the better the resolution.

The second row in Figure 7 shows how pulse compression theory works. We start with an impulse. The impulse is passed through an all pass filter $H(\omega)$ that can produce a chirp, or some phase coded waveform. This waveform gets transmitted and received. The received waveform is passed through an inverse filter, which returns the signal back to the impulse we started with. The purpose for doing this is that the transmitted waveform can be a long duration low power signal with low peak-to-peak volts. After compression, it can be equivalent to an incredibly large peak-power impulse that is hard to build due to the high peak voltages and currents required.

The third row in Figure 7 illustrates the case where $H(\omega)$ produces a chirp, which gets compressed to an impulse. In the illustration, the chirp has a 1 GHz bandwidth, lasts 1 ms, and gets compressed to 1 ns. So a 1 mW chirp is compressed to the equivalent of a 1 kW peak power transmitter.




Figure 7: Bandwidth Gives Range Resolution

A key feature of chirp is its ability to use what is called “stretch” or “de-ramp” processing. Figure 8 illustrates stretch processing. The top half of Figure 8 is a plot showing the frequency versus time of the transmitted chirp and returns from targets at two different ranges. The key is that the frequency offset between the transmitted chirp and the returns, is constant and at a different frequency for each range. The lower part of Figure 8 shows a receiver structure where a mixer is used to create that constant frequency that is a function of range. All the offset frequencies are captured by an analog to digital converter. A digital signal processing (DSP) is used to separate the various frequencies to resolve all the different target ranges. In other words, the DSP compresses the signal back to what would have been received had an impulse been transmitted and received instead of the chirp. The DSP requires a Fourier transform and a de-skew operation to make the output identical to an impulse based system.

Why is stretch important? Because, when the compression ratio is high, the ADC bandwidth does not need to cover the entire bandwidth of the transmitted signal. Instead, the ADC bandwidth is proportional to the “range swath” or the distance between the shortest range and farthest range of interest. More precisely:

$$\text{ADC Sample Rate (Bandwidth)} \propto (\text{Max_Range} - \text{Min_Range}) * \text{Ramp_Rate (MHz}/\mu\text{s}) \text{ or}$$

$$\text{ADC Sample Rate (Bandwidth)} \propto (\text{Max_Range} - \text{Min_Range}) * \text{Bandwidth} / (\text{Chirp Duration})$$

Since ADC power goes up with sampling frequency, maximizing the chirp duration minimizes the ADC power. Continuous chirp, i.e. transmitting and receiving simultaneously, maximizes the chirp duration. It minimizes the average power drawn while the radar is active. Where does stretch break down? It breaks down if the compression ratio is too small. Indeed, it can even make the ADC bandwidth larger than the signal bandwidth. Due to end effects, it is most efficient when the compression ratio is high – or in other words, where the chirp is very long relative to the compressed pulse.

Pulsed chirp allows the receiver to operate while the transmitter is off. A pulsed chirp system generally has a huge advantage in dynamic range since the receiver is not trying to operate while the transmitter is potentially blinding the receiver. In long range applications, it is the only way to go. But it drastically raises the battery drain in short range applications. In this case, because of the short range, the chirp cannot be long. Therefore, the efficiencies of high compression ratio chirp disappear. Even if chirp portions were transmitted so that an entire chirp is sent at about the same rate, the chirp rates and ADC clock go up by more than a factor of 2. The air between the radar and nearest range bin must be filled with chirp energy, with enough time for the transmitter to be turned off, the receiver turned on, and all the switching transients to settle, so that the receiver can be receiving the energy from the nearest range. The receiver must then stay active long enough to receive the energy to the farthest bin. This leads to the $>2X$ factor. Beyond this core factor, the hardware required to support generating tightly controlled phase on the chirp sections, with fast settling times, adds complexity and draws all the more power.




Figure 8: De-Ramp ("Stretch") Implementation

The combination of (1) the reduced dynamic range requirement of short range, (2) the short range swath, and (3) the high dynamic range obtainable with today's semiconductor technology, i.e. mixers, amplifiers, and ADC, make the use of continuous chirp a viable and lower power, option.

3.3 Geometry and sampling (PRF and number of points) calculation

Figure 9 shows a side view of the radar/target geometry. Figure 10 shows a top view of the geometry along with a spreadsheet that was used to look at the PRF requirements, the duration of an aperture, and the number of pulses required. The final highlighted columns show the number of samples in an aperture assuming a 250 Hz PRF, which gave some margin over the 208 Hz required for a worst case wide integration angle at close range. It reduces smearing caused by the range sampling widow occurring over a 4ms time span.




Figure 9: Scenario Geometry

Fo Hz	A m	Imaging Angle deg	Speed m/sec	Sampling over Nyquist e.g. 1=A:2 2=A:4 4=A:8	PRF for Request Nyquist spacing direct inbound	Height m	Boat Length m	GHD Range m	Slant Range to Ap center m	Doppler Aperture Length m	Slant Range to Doppler Ap End m	Ap Length m	Slant Range to Ap End m	Distance along boat track to change radar range by $\lambda/4$ at Doppler-ap end points	Time in Ap sec	Min PRF	Min # of Pulses	# of pulses in Ap
1300000000	0.23	60	10	3	280	10												
		104					20	10	14	36	23	16	16	0.048734	1.6	205	335	250 408
		68					20	100	100	136	121	116	116	0.068593	11.6	146	1692	250 2901
		86					10	10	14	26	19	16	16	0.056474	1.6	177	289	250 408
		64					10	100	100	126	119	116	116	0.072374	11.6	138	1603	250 2901
		60					0	10	14	16	16	16	16	0.077144	1.6	130	212	250 408
		60					0	100	100	116	116	116	116	0.076908	11.6	130	1509	250 2901




Figure 10: Energy Calculation

3.4 Transmitter Power Calculations

Range compression gain, ISAR aperture gain, and range and cross-range resolution are a function of the window taper used on each. Figure 11 lists the loss factors for various window-tapers assuming Nyquist sampling – no up-sampling.

Window	1st Sidelobe (dB)	Asymptotic Sidelobe Rolloff (dB/octave)	Coherent Gain	Equiv Noise BW (bins)	6dB BW (bins)	Worst Case Processing Loss (dB)
Rectangular	-13	-6	1	1	1.21	3.92
Triangular	-27	-12	0.5	1.33	1.78	3.07
Hanning	-32	-18	0.5	1.5	2	3.18
Hamming	-43	-6	0.54	1.36	1.81	3.1
Poisson	-24	-6	0.32	1.65	2.08	3.64
Poisson	-31	-6	0.25	2.08	2.58	4.21
Cauchy	-35	-6	0.33	1.76	2.2	3.83
Cauchy	-30	-6	0.28	2.06	2.53	4.28
Gaussian	-55	-6	0.43	1.64	2.18	3.4
Kaiser	-69	-6	0.4	1.8	2.39	3.56
Kaiser	-82	-6	0.37	1.93	2.57	3.74

Figure 11: Characteristics of various window taper functions

Figure 12 shows a calculation of system losses. The loss for detections uses only the first two rows since detection does not involve imaging.

Loss (dB)	Description
1.4	Range Compression Kaiser Window $\beta=5.44$ (~same as Hamming)
1.7	Scalloping loss (reduces to .4/.1 dB with 2X/4X zero-pad oversampling)
1.4	Kaiser for aperture (imaging)
1.7	Scalloping loss for Imaging
6.2	Total
3.1	For Detection

Figure 12: System losses

Figure 13 defines the parameters and lists the equations used for the transmitter power calculation.

L_m	Loss Factor for Multipath
L_p	Loss Factor for 2-way penetration of boat
σ	radar cross section in m^2
$L_R = \frac{G^2 \lambda^2 \sigma L_m L_p}{(4\pi)^3 R^4}$	Range Loss Factor from Tx Antenna port to Rx Antenna port
T_0	290°K
F	$10^{(NF/10)}$ where NF= Noise Figure in dB
$T_e = T_0(F - 1)$	System Effective Noise Temperature
β	Bandwidth Hz,
k	Boltzman's constant = $1.3806488 \times 10^{-23}$ J/K
$\left(\frac{S}{N}\right)_{in} = \frac{P_t L_R}{k T_e \beta}$	S/N at input to receiver for each pulse
L_s	system loss factor
$\left(\frac{S}{N}\right)_{out} = T \beta L_s \left(\frac{S}{N}\right)_{in} = \frac{P_t L_R L_s}{k T_e} = \frac{E_t L_R L_s}{k T_e}$	Average S/N at peak of envelope of matched filter output
$E_t = \frac{k T_e}{L_R L_s n} \left(\frac{S}{N}\right)_{out}$	Energy Per Pulse Required to get SNR_{out} after integration of n pulses

Figure 13: Equations used to compute transmit power requirements

To conserve power, the system runs in two modes. A detection mode with a high sleep time duty cycle. In this mode, the system wakes up, transmits a single pulse, and determines if a target is present that should be imaged. If not, the system goes back to sleep. If a target is detected, the radar goes into the second mode, where it sends enough pulses to image the target.

Figure 14 lists, for imaging an interior object, the parameter values that went into the equations in Figure 13.

$R = 100$ m	Range
$F = 4.5$ (6.5dB NF)	Noise Factor
$G = 2$ (3 dB gain)	Antenna Gain
$L_m = 1$ (no multipath loss)	Multipath Loss
$(S/N)_{out} = 31.6$ (15 dB)	Desired output SNR
$\sigma = -20$ dBsm	cross section of interior target (10 cm metal cylinder with 5 cm radius at 1 GHz)
$n = 2900$	number of pulses integrated to form image
$L_p = -10$ dB	loss penetrating boat
$L_s = .25$	6 dB system loss

Figure 14: Parameter values used for transmit power calculation

Using these parameters, the transmitter requirements are:

$$E_t = 0.34 \text{ mJ @ } 1 \text{ GHz}$$

$$= 0.68 \text{ mJ @ } 1.5 \text{ GHz}$$

Assuming the same power is used for single pulse detection in the detection mode, i.e. $n = 1$ (single pulse detection), $L_p = 1$ (no penetration), and $L_s = .5$ (3 dB system loss) we find $\sigma = 1.6$ dBsm

3.5 Case-1 - Tx & Rx in separate time slots

Figure 15 lists the calculations used to look at a short chirp and impulse based transmitter. Note that the impulse system power is approximately the same voltage as used on the Army Research Lab (ARL) BoomSAR. At this power level, the pulse generator is large and inefficient. While a higher PRF might be used to reduce the voltage requirements, doing so has two drawbacks. First, a higher PRF also requires a higher duty cycle on the ADC and the processing that must digest the increased data flow. Second, a higher PRF ultimately runs into an ambiguity issue. The ADC, for both the 50ns and 2 ns pulse needs to cover the entire transmitted bandwidth. Currently, 8-bit 1.5 Gsps ADCs burn 6 Watts and an 8-bit 2.2 Gsps ADCs burn 8 Watts. In contrast, the 16-bit 20 Msps ADC selected for the continuous system burns only 60 mW.

250 Hz	Minimum PRF (4ms rep rate)
50ns	T/R settling time required prior to receiver being active
15m	Minimum slant range (10m gnd range & h=10m radar)
	CASE-1 – short chirp
~ 50 ns	Maximum pulse duration
13.6 kW	=0.68mJ/50ns = Avg Pwr during pulse
2.3 kV	p-p volts into 50 Ω @ 250 Hz
580V	p-p into 3 Ω , with 16:1 transformer to get to 50 ohms
	CASE-2 - Impulse
2ns	Minimum Length Pulse
340kW	=0.68mJ/2ns = Avg Pwr during pulse
~12kV	p-p into 50 ohms @ 250 Hz

Figure 15: Peak transmit power calc, Tx & Rx in separate time slots

3.6 Case-2 0 Tx & Rx simultaneously

In this case, we minimum the peak-power and the ADC bandwidth by operating at a 100% duty cycle. Given the 250 Hz PRF, we can use a 4ms chirp – (125 Hz Triangle Wave, 4ms up-chirp, 4ms down-chirp). The transmit amplifier only needs to make 0.68mJ/4ms = 170mW or +22dBm (7.9 V p-p into 50 ohms). This low power level has two major impacts. First, small geometry wide bandwidth and efficient transistors in GaAs, InP, GaN, and SiGe can easily withstand this voltage level. Efficiency at higher power levels typically depends on narrowband tuning. A wideband system does not have this luxury, and is therefore usually much less efficient. The viability of using small geometry transistors will raise the efficiency. Second, the receiver has enough dynamic range to manage this relatively low transmitter power level – allowing the receiver to operate while transmitting.

Generation of this 4ms long chirp is significantly easier (smaller, lower power, less expensive) than a 50ns chirp. Furthermore, it is easier to control so as to extend the battery life. For example, the power can be reduced depending on either (a) range – i.e. whether the target is detected at the far side, or close side of the river, or (b) the detected SNR, in combination with the detected speed of the target. The 4 ms time is long enough that the

pulse length could be reduced. For example, the pulse could be reduced to 2 ms and the system could sleep for 2ms. This would save the battery yet not change the sample rate (PRF) that the ISAR needs. If the target is moving slower, then the radar can reduce the power by reducing the PRF and sleeping between pulses. With this approach, the amplifiers remain biased at their most efficient operating point. Another power reduction option is to bypass and switch off the final TX power amplifier.

4 RF Hardware Strawman and Power Evaluation

4.1 Exciter Strawman Hardware

The exciter needs to make an extremely linear chirp in order to achieve a reliable compression ratio of 2×10^6 . Common approaches are:

- Linearize a voltage controlled oscillator (VCO) by driving it with an ADC via a look-up table (LUT).

This open-loop approach is not viable as it cannot achieve the linearity required to achieve the needed compression ratio.

- Directly make the chirp with a DDS (direct digital synthesizer) – a digitally produced waveform driving an ADC.

This approach has the advantage of nearly instant wake-up/sleep and no VCO drift, warm up, settling time, or loop stability issues. It is not viable, however, due to power required to get the required clock rate – on the order of a few Watts.

- Make low bandwidth chirp with a DDS and lock a VCO using a standard PLL to multiply the low-BW chirp.

Figure 17 illustrates a strawman design using the latest generation low power parts to implement this approach. Advantages are that the power is only ~ 0.63 Watts and the PLL loop bandwidth is wide so tracking at endpoints is fast and VCO linearization is finely spaced. Disadvantage is that the phase noise of the DDS is multiplied. As an example, with $M=16$ the DDS might generate a 50MHz to 100 MHz chirp, and the VCO would generate an 800 MHz to 1600 MHz chirp.

- Make low bandwidth chirp with a DDS and lock a VCO using a fractional-N PLL connection.

Figure 16 illustrates a strawman design using the latest generation low power parts to implement this approach. Advantages are that the power is only ~ 0.63 Watts, the DDS output is filtered around the reference frequency providing some noise reduction, the DDS runs slower and those slower running portions of the chirp will draw less power. The disadvantage is that the settling time may not be fast enough.

- Make low bandwidth chirp with a DDS and directly multiply it up to the required frequency range. This has the advantage of allowing the fastest power cycling and extremely fastest settling times. The multiplication factor could be low so that a single amplifier could be used to boost the desired harmonic. Power used in the VCO and amplifiers in Figures 16 and 17 would be used for the RF amplifier. The disadvantage is that high fundamental rejection (i.e. good balance) is needed in the multiplier stage.

The good news is that the last three approaches, and in particular, the diagrams in Figures 16 and 17, show that the required low power level is achievable.




Figure 16: VCO linearized & locked in a fractional-N (DDS) configuration




Figure 17: VCO Linearized & Locked

4.2 Receiver Strawman Hardware

Figure 18, shows a strawman design for the receiver. It uses a classic direct conversion approach. This approach eliminates the second chirp generator that is often used to offset the range swath and center the baseband signal to be at DC in the center of the swath. Doing so saves all the power of another Figure 16 or Figure 17 block except for the crystal oscillator. High dynamic range is achieved by using the latest generation PHEMPT based RF amplifiers, triple balanced mixers, SiGe baseband amplifiers, and 16-bit ADCs. Very high resolution (dynamic range) and high speed ADC's are available at low cost and low power, as shown in Figure 18. The relatively high 20 MHz sample rate of the converter shown, can be used to reduce the duty cycle. Any final design will need to consider the power tradeoffs between the ADC speed/power and the power savings possible via a reduction in the duty cycle.

Figure 16, Figure 17, and Figure 18 take advantage of the latest generation parts to achieve low power. For example, the ultra low power crystal oscillator developed was developed for cell phone & high volume portable applications. The high power output amplifier is configured with GaAs switches to allow bypassing and allow adaptive changes in the duty cycle to enhance battery life. We believe these diagrams that use real parts give strong support for battery life numbers and radar performance numbers.




Figure 18: Receiver

4.3 RF Hardware Power

At 100% duty cycle, the transmitter/exciter draws 930mW and the receiver draws 630 mW, for a total of 1.56W. Because of bias currents often dominating power drain, typically adjusting output power often barely affects battery life. To provide a path to extend the battery life, we will adjust the pulse durations via duty cycle to reduce both transmit and battery power simultaneously. With a sample rate of 20 MHz, and assuming 1.5X Nyquist sampling, the duty cycle can go from 100% (4ms) to 1.5% (0.06ms). This is based on the frequency that the ADC must capture is $2 K R_{max} / C$, where K is the Hz/Sec ramp rate of the chirp, and C is the speed of light. Assuming 1.5x Nyquist Sampling, the ADC sample rate $F_s = 3 F_{off} = 6 K R_{max} / C$, and the maximum ramp rate $K_{max} = F_s C / (6 R_{max})$. So at 20 Msps, the chirp rate can be as high as 600MHz/0.06ms.

Assuming R^4 power scaling, targets at 34m or less can use 1.5% duty cycle. The duty cycle $= (R/100m)^4$. Figure 19 plots the battery watts consumed versus target range, during imaging. It assumes a minimum RCS target. A higher RCS target could use less battery – the entire curve moves down. The knee in the curve is due to 20 Msps sample rate driving the minimum duty cycle and the high PRF needed to capture the higher ISAR angular spin-rate at short ranges.




Figure 19: Battery Power During Imaging

5 Processor Strawman Design and Power Evaluation

5.1 Overview

The purpose of this task was to derive accurate power budget so there was a better understanding on how the complete radar, both RF and Digital sections, impacts battery life. A key part of this work was the derivation of a set of modes (like sleep, blink, run, detect, image) and a system time line or operational concept that sequenced between these modes so as to minimize power. The derived power numbers are supported by mapping the required processing onto real hardware – in this case, a current generation Altera FPGA, cell-phone RAM, and a cell-phone processor, etc. The processing captured the key signal processing steps:

- Range compression
- CFAR detector
- ωk (range migration) processing
- Feature-creation (length, width, RCS, etc.)

Ignored processing included scaling/packing/compressing of the data for transmission, and trajectory estimation. While the trajectory estimation algorithm was ignored, the processing load will likely be smaller after its insertion. This reduction is anticipated because the estimated trajectory will allow decimation of data set prior to FFT processing.

After a survey of DSP and FPGA options, we chose an FPGA approach for baseline, and used an Altera Cyclone III. Why the shift away from CPU's and toward FPGAs for DSP? Until recently, IC scaling has been on both density *and* speed. But today, as the feature sizes shrink further, the IC performance scales primarily in density, and NOT speed. This is a major shift/driver on semiconductor industry. An example outcome is the emergence of multi-core processors as the "standard" processor.

DSP chips do have their advantages. The "C" programmability of DSP is a huge advantage for DSP chips. It makes the development cycle is also much more manageable since there is a large base of engineers who know "C", it is much more easily documented and maintained, and it allows fast compiles so that test-fix-test iteration loops close faster. There are no "routing and timing closure" issues. For all these reasons, DSP chips are advantageous when control is a major part of its job. FPGAs, however, take better advantage of the density scaling and can do more processing in a smaller size and with lower power. Figure 20 lists why this outcome is the case.

DSP	FPGA
Pipeline -- Fixed 1-pipeline fits all Requires blazing speed on short fixed pipeline Can't change balance between different functions (c.g. DSP chip has predefined number of adders and multipliers)	Pipeline -- Custom pipelining on everything Takes advantage of today's scaling of density Capacity to have many parallel functional blocks Can choose more adders and fewer multipliers or vice-versa
Highly serial, each functional block is optimized to the transistor & blazing fast	Highly parallel, functional blocks can be slower because there are so many more of them
Multiplexed Bus Centric Bus shared by multiple blocks Lots of connections-lots of stray C Multiplexing causes toggling that wastes power	Point-to-Point Bus Centric Most "busses" are point-to-point Low number of connections = low stray C No multiplexing, if data does not change, no power is wasted charging/discharging C
1 clock fits all No granular control of what is getting a clock no granular control of what runs fast and what runs slow	Multiple Clock domains easy Fine control over what is clocked & when Takes fine advantage of duty cycle on every block Fine control over what runs fast and slow Examples: Memory controller rate governed by memory Function blocks run at rates limited by their own logic Interface logic runs at rate governed by the interface
General purpose nature results in powered and clocked blocks that are not used for our design	Less "stray C" because unused logic is not clocked/connected If a section cannot be powered down, it burns static power but is not clocked
Already hardened – very optimized chip layout	Power & Performance can be enhanced by hardening the design – c.g. "HardCopy™" where FPGA code is translated into an optimized chip (no longer programmable) – like a custom pipelined DSP.

Figure 20: DSP versus FPGA drivers

After a survey of memory options, we chose cell phone memory for the baseline. The combination of low power and high speed for this memory is a major breakthrough for enabling this application.

After a survey of ultra low power processors for housekeeping functions and the low power detect-mode processing, we chose a "mobile" ARM, again designed for low-power portable consumer electronics products.

We designed a pipelined FPGA architecture to implement the required processing. We implemented with Altera's design tool, the most power consuming (many gates) functions (e.g. an FFT block, etc.) to get reasonably accurate size and power.

The system has two modes - detection and imaging. Each has separate computation path.

Detection runs at a 1 Hz rate to achieve ultra low power. Once per second, the ARM processor wakes up, powers the radar up long enough to collect one or a few pulses, processes the full 100m range of data and either (a) updates the background if no target is detected, or (b) wakes up the FPGA and begins making an ISAR image. This mode is designed to minimize power by maximizing the sleep time.

Imaging runs at 250 Hz radar pulse rate. To save power, this mode only runs when a target is present. The background is not updated since targets are in view. This mode is designed to minimize the powered-up time of the FPGAs. The processing is split between two FPGAs so that when back-to-back images are needed one can still have a sleep cycle, resulting in lower overall power. The output product can take several forms suitable for enhanced processing. These include a range compressed data set, a simple focused data set that is invertible, an auto focused data set that is not invertible, and features like length, width, RCS, and relative locations of largest peaks and their RCS.

5.2 Digital System Hardware Design

Figure 21 shows a block diagram of the digital hardware that includes the processing system. Figure 22 lists the features and functions of the selected "mom" processor. Figure 23 lists the features and functions of the selected memory. Figure 24 lists the features and functions of the selected FPGAs.




Figure 21: Block Diagram

ARM Cortex M3 – (e.g. STM32F103)	Functions
Mobile Oriented	Mom functions – good general cpu
35ma@2.7V (100mw)	Entire Detection Mode
70ma with all peripherals	— Processing (Stretch, CFAR)
5uA deep sleep	— Background memory processing
watch-dog timer (1sec wake up & detect)	Flash Memory management
Single cycle 32-bit multiply	Power management
2 cycle MAC (multiply & accumulate)	turning on/off FPGA, radar, maybe comm's.
64kB on-board memory	

Figure 22: ARM Cortex M3 “Mom” processor

Cellular RAM - MT45W8MW16 = 8Mx16	Functions
Also called PSRAM (Pseudo Static RAM)	Memory to collect full 2D aperture data
<10 uA power down	Memory for 2D FFT during processing
200uA standby	Combined Write & Read on 8k by 1k 256ms
35ma@1.8V = full speed	Scratch pad and Background Record
10/35 ns (in-page/other page)	

Figure 23: Memory

FPGA – Altera CycloneIII LS	Functions
EP3CLS70	FPGA-1 – Imaging mode only (off in detect)
Has 200 18x18 multipliers	1D processing up to first exfiltration point
8k FFT in 8k clocks	FPGA-2 – Imaging mode only (off in detect)
~384 kB memory	2D processing beyond first exfiltration point i.e. past first “corner turn”
60ms power up & config time	Powered off between output of 2D target image chips
150mW static power	Up to 8 targets (e.g. different ranges) can be processed by reducing sleep time

Figure 24: FPGAs

While the two FPGAs could be combined into 1 Larger FPGA, the split allows FPGA-2 to be shut off between images. It is fast enough to finish an image before a new aperture has been collected. This fact allows two benefits. First static power is conserved when there are back-to-back targets. Second, the “excess capacity” can be used to form image chips of multiple targets, such as two targets coming down the river together. They might share the same aperture, but be at different ranges.

5.3 Main Power Consuming DSP Blocks

The FFT's used block floating point to achieve a high dynamic range fixed point implementation. There are three classes of implementations for fixed point add/multiply.

- (1) Shift input bits so that it is impossible to overflow. For example, use only 9 bits out of 16 for a 128 point FFT. Clearly this choice results in a huge loss of dynamic range.
- (2) Fixed Sealing - Perform a 2-bit shift after each butterfly stage.
- (3) Block Floating Point - Perform a bit shift only if needed after each butterfly stage. This method requires output of an exponent to know how data was sealed.

Figure 25 illustrates the difference in the noise floor between fixed scaling and block floating point. Block floating point is >20dB better.




Figure 25: Reason For Preference For Block Floating Point FFT

The power budget is dominated by large blocks – i.e. FFT blocks. The term LE is Altera's term for "Logic Element" which represents a block that can do any 4-input combinatorial logic with or without a register. One large block is an 8k x 1k 2D FFT. This block has the following features:

- Allows variable size (max 8k)
- Uses ~9000 LEs
- Uses 96 18-by-18 Multipliers
- Uses 64 block RAMs
- Pipelined so that there is 1 complex-sample in & out per clock
- ~300mW @ 100MHz
- 2 * 8k * 1k cycles = 160ms with fast RAM
- It takes ~256ms to complete because cell phone RAM access time dominates
Higher speed memory is not viable due to the power dissipation.

Another large block is the 4k FFT that must be done 1600 times. This block has the following features:

- Uses ~6000 LEs
- Uses 48 18-by-18 Multipliers
- Uses 76 block RAMs
- Pipelined so that there is 1 complex-sample in & out per clock
- ~200mW @ 100MHz

Miscellaneous other functions are small enough to neglect relative to the huge 6000 LE and 9000 LE FFT blocks. These include:

- Square root: ~400 LEs, pipelined - 1 answer per clock – only used on 1 vector per image

- Interpolate: pipelined, similar to N-point FIR 1 in & 1 out per clock
- Where N is the number of input samples covered by resampling filter
- This is the largest block behind the FFTs
- Zero pad: 1 clock persample (zero padding is always to internal memory)
- Vector multiply: pipelined 1 in & 1 out per clock regardless of real (1 mult) or complex (4 mult)
- 2D Hanning Window multiply – 1 clock per point
- vectorized so it is always running in 1D in internal memory
- CFAR: pipelined, 1 in & 1 out per clock

5.4 CFAR

Figure 26 shows a standard cell-average constant false alarm rate (CA-CFAR) detector. This structure can be pipelined. An example FPGA pipelined design is in “Design and Implementation of a CFAR Processor for Target Detection” by César Torres-Huitzil, Rene Cumplido-Parra, Santos López-Estrada, Computer Science Department, INAOE, Apdo. Postal 51 & 216, Tonantzintla, Puebla, México. The summary listed below show how the block is tiny relative to the FFTs.

- Process is pipelined to 1 in & 1 output per clock
- Bit width depends on using approximate Mag function or Mag^2 function
- Assume $n=32$, $g=16$ (data is 2X oversampled)
- Vectorize the threshold multiply – i.e. Create Vector $Y = 1366$ multiplies
- Vectorize computation of S_L and of S_R
- $2 * [16 \text{ adds to get started, } 2 \text{ adds (subtracts) per bin after that}] \approx 4 * 1400 \text{ adds}$
- Zero with $n+g$ zeros at ends for startup conditions
- Vector compare between Y and S_L and Y and S_R
- Vectorize the logic on the two resulting 1-bit wide vectors




Figure 26: Standard Cell-Average Constant False Alarm Rate (CA-CFAR) Detector

5.5 Baseline Radar to Processor Interface

Figure 27 shows a data flow representing a baseline radar to processor interface. Listed below are the main parameters

- ADC's run at 3 Msps
- "Pulse" is 4 ms – 12k record
- Detection Mode
- 1 Hz PRF
- ADC feeds ARM processor
- Imaging Mode
- 250 Hz PRF (= 4 ms PRI)
- Transmit, receive, and processing is continuous
- ADC feeds FPGA-1

Worst case analysis list is captured below:

- Target is at the maximum range
- There are a series of targets back-to-back – requires continuous imaging
- 1600 radar records - 6.4 seconds per image
- 100 meter range-swath in the records
- 250 Hz PRF (4ms PRI)
- ARM and FPGA have memory management to handle continuous imaging on multiple targets (multiple ranges with a caravan of back-to-back targets).




Figure 27: Baseline Scenario

5.6 FPGA processing


To have zero-overhead corner turns and maximized the FPGA logic utilization, there are 6 memories, each with its own bus, mapped to the pipelined processing stages. Each memory is required to hold the entire data set. In order to do corner turns at the maximum rate of the pipelined memory, we write rows and read columns 4 vectors at a time. We describe the processing and flow of data through the 6 memories using the colored chart in Figure 28. Each color represents a pipelined process between buffers (i.e. corner turns). The goal is to maximize FPGA non-clocked and power-down durations. FPGA-2 (Purple, Brown, and Blue is powered down while FPGA-1 (Green) is processing.

1D processing To Memory A/B (FPGA-1)	Write-A Rows Ap-1	Write-B Rows Ap-2	Write-A Rows Ap-3	Write-B Rows Ap-4	Write-A Rows Ap-5
1D processing From 1 st to 2 nd corner turn (FPGA-2)		Read-A Columns Ap-1	Read-B Columns Ap-2	Read-A Columns Ap-3	Read-B Columns Ap-4
		Write-C Rows Ap-1	Write-D Rows Ap-1	Write-C Rows Ap-3	Write-E Rows Ap-4
1D Processing after 2 nd Corner turn (FPGA-2)			Read-C Columns Ap-1	Read-D Columns Ap-1	Read-C Columns Ap-3
			Write-E Columns Ap-1	Write-F Columns Ap-1	Write-D Columns Ap-3
Extract sub-image & push to (ARM)				Read-E Write FLASH Ap-1	Read-F Write FLASH Ap-2

Figure 28: Data flow pipeline through the 6 memories

5.6.1 FPGA-1 1-D pulse compression processing

Recall how de-ramp data from stretch processing is generated, as illustrated in Figure 8. The pulse compression processing required to convert the data into the equivalent of what would have been received from an impulse is illustrated in Figure 29. This figure shows background subtraction, filtering and decimation of the ADC samples, the Fourier transform including required zero-padding for upsampling, and the de-skew operation. Note that all these operations are done “at the same time”, that is, in a pipeline that, once full, provides output data in real-time with the input data. Figure 30 shows operation counts for the various blocks. The Filter/Decimate block dominates the processing load and sets the timeline.



- 7.5cm is ~2x above nyquist sample spacing
-- reduces scalloping loss by about 2.7 dB (i.e. from 3.1 dB to about .4 dB)

Figure 29: FPGA-1 Pulse Compression Flow

Processing Block	Clocks	Time	~ Real Multiplies
Filter & Decimate (I & Q)	4k	4m	(61 taps) $4000 * 2 = 488,000$
8k fft	8k		$4(8192)\log_2(8192)=425,984$
Complex multiply over 100 m range	1.3k		$4(1366 \text{ pts}) = 5,464$
Total for 100 m range swath			919,448
Complex multiply over 200 m range	2.7k		$4(2731 \text{ pts}) = 10,924$
Total for 200 m range swath			924,908

Figure 30: FPGA-1 Pulse Compression Operation Count

If/when a trajectory estimator is added, the output data can be put in a "spotlight" mode, centered on the target. Figure 31 illustrates spotlight reformatting, where the range to the first point in a record is adjusted to keep the target in the center of a short record. Basically, the centroid of a bright point is maintained in the center of the collection window, and the range shift performed on each record to keep it centered is captured so that later stages know how to interpret the data. The per pulse SNR is high enough to do single pulse detection and to allow correlations between records to provide the time lags (and phase rotations) required to estimate the trajectory. Trajectory estimation is also allows computation of the Stolt interpolator and phase multiply matrix.

To illustrate the data reduction, assume we want an 18m window of data surrounding the target. Assume a range bin is 7.5 cm. Then the number of output range bins is $18\text{m}/7.5\text{cm} = 240$ points rather than the 1366 points we started with.

In addition to allowing computation of the Stolt interpolator and phase multiply matrix, trajectory estimation also allows presum of spotlight centered range records. This pre-summing would drastically cut the amount of data that would need to be exfiltrated and/or processed by FPGA-2. For example, the aperture for a target at maximum range contains 1600 records. If we can presum by a factor of 8, the output could be reduced to 200 records. Order $3H \log_2(H)$ reduction in *total* processing from could be gained by presuming by H. In other words, the computational cost of performing spotlight reformatting and trajectory estimation is small compared to the savings it could achieve. Figure 32 depicts spotlight processing and trajectory estimation, where the output record has been reduced to 200 by 240 points.




Figure 31: Spotlight formatting




Figure 32: FPGA-1 Pulse Compression Operation Count

5.6.2 FPGA-2 2-D αk Image Former

Figure 33 shows the major processing blocks in both FPGAs where the coloring matches the pipeline sequencing shown in Figure 28—green for FPGA-1 and purple, brown, and blue for FPGA-2. Figure 34 is a timing diagram showing how the various functions are done in parallel and how the 8k FFTs dominate the time line. Figure 35 shows how the metric computations are so minor they can be ignored for the purposes of power calculations.




Figure 33: Wavenumber Migration Processing Partitioned Around Corner-Turns




Figure 34: Example of Timing on FPGA-2 During Image Formation




Figure 35: Metrics Negligible Processing Load

5.7 Power Consumption Analysis

Summary power numbers for FPGA-1 are shown in Figure 36. Summary power numbers for FPGA-2 are shown in Figure 37. The numbers are based on back-to-back targets continuously appearing at the farthest range. At closer ranges the power will be lower because an aperture has fewer pulses. While the FPGA-2 static power is drastically reduced by power cycling, the remaining static power dissipation is caused by the configuration time each time the FPGA is powered up. FPGA-1 static power is high because there is not enough time to power it down given the time it takes for configuration before it must be ready to operate again. Using a process like Altera's "Hardcopy™" would eliminate the configuration time and thereby eliminate this configuration time power, and in the case of FPGA-1, the time it remains powered up, when it could be powered down.

Element	Avg Power mW	Element mW	Sub Processing Block Duty		Processing Block Duty		Power On/Off Duty config + processing/ duration of data	
Static	150.00	150	1	1	100	1	1	1
512k FFT	0.17	50	41/160	0.26	83/164	0.51	164/6400	0.026
8k FFT	3.84	300	1	1	82/164	0.5	164/6400	0.026
I/O	2.56	200	1	1	82/164	0.5	164/6400	0.026
misc	1.28	100	1	1	82/164	0.5	164/6400	0.026
TOTAL	157.85							

Figure 36: FPGA-1 Power Summary Calculation

Element	Avg Power mW	Element mW	Sub Processing Block Duty		Processing Block Duty		Power On/Off Duty config + processing/ duration of data	
Static	5.25	150	1	1	100	1	(164+60)/6400	0.035
4k FFT	0.66	200	41/160	0.26	83/164	0.51	164/6400	0.026
8k FFT	3.84	300	1	1	82/164	0.5	164/6400	0.026
1k FFT	1.28	100	10/20	0.5	1	1	164/6400	0.026
VO	2.56	200	1	1	82/164	0.5	164/6400	0.026
Interp/Windows/Mag/etc	1.28	100	1	1	82/164	0.5	164/6400	0.026
TOTAL	14.88							

Figure 37: FPGA-2 Power Summary Calculation

Another option to reduce power further would be to add a smaller FPGA to just manage a buffer between the radar and the A/B RAM. The task would be storing N=64 records (64*4ms=256ms) and bursting them to RAM-A/B. FPGA-1 could now operate in a burst mode as follows:

- 60ms wake/config, 10ms to process 64 pulses & go/stop test, Power off for 186 ms
- 70/256 duty cycle would reduce static power by 1/4 -- i.e. from 150mW to 41 mW

Summary power numbers for the memory is shown in Figure 38. Summary power numbers for the ARM processor is shown in Figure 39.

Element	Avg Power mW	Element mW	Duty Cycle	
A/B Hold	0.35	0.360	rest	0.97
A/B Dyn	1.61	63.000	2*82/6400	0.03
C/D Hold	0.34	0.360	rest	0.95
C/D Dyn	3.23	63.000	2*164/6400	0.05
E/F Hold	0.34	0.360	rest	0.95
E/F Dyn	3.23	63.000	2*164/6400	0.05
ARM Hold	0.36	0.360	rest	1.00
ARM Dyn	0.06	63.000	1/1000	0.00
TOTAL	9.53			
ARM Total	0.42			

Figure 38: Memory Power Summary Calculation

Element	Avg Power Detect	Avg Power Imaging	Element mW	Sub Processing Block Duty	
Sleep	0.01	0.01	0.013	rest	0.97
Detect	4.05		135	30/1000	0.03
Write Flash		2.91	300	62/6400	0.01
TOTAL	4.06	2.92			

Figure 39: ARM Power Summary Calculation

5.8 Digital Hardware Power Summary

Figure 40 sums the power drawn by the individual parts.

Chip	1-Target back-to-back Continuously (per 6.4sec) mW	Detect Only mW
ARM	2.9	4.1
FPGA-1	157.9	0.0
FPGA-2	14.9	0.0
Memory	9.5	0.4
TOTAL	185.2	4.5

Figure 40: Total Power Summary Calculation

We have shown that the processing can be done in a small ARM processor, 2 FPGAs, and 12 cell-phone RAM chips. We have shown that the power drawn in the detection mode is ~5 mW. For imaging, we have shown that a total power draw of <200 mW in a “worst case” scenario where a maximum-range target is in scene continuously. For ~\$200k NRE (non recurring engineering cost), the “HardCopyTM” process can be applied to harden FPGA-1, resulting in a total power draw of <100mW.

6 Antenna Design

Wide bandwidth antennas are often dispersive. Figure 41 illustrates most of the classic wideband antenna designs. The fat dipole and the TEM horn are the only ones on this chart that are not dispersive. In order to obtain high resolution SAR and ISAR imagery, the antenna must have low dispersion, and have low dispersion across the wide integration angle of the SAR or ISAR system. If the antenna does not inherently have low dispersion, then the dispersion can be theoretically be processed out by deconvolution—but only if the dispersion is consistent over the wide integration angle of the SAR or ISAR. Figure 42 illustrates this deconvolution processing. In this case, the antenna has a transfer function of $H(w)$ and the deconvolution filter $Q(w) = 1/H(w)$. The key is that the resulting impulse must be stationary relative to the antenna pointing angle. In other words, the two impulses pictured on the oscilloscope must not move when the antenna is rotated about its phase-center.





Figure 41: Wide bandwidth antennas are often dispersive



We desire τ and waveform shape
to be independent of rotation angle (θ, ϕ).

Figure 42: Setup to measure angle dependent dispersion

Since the FCC adopted its UWB regulations, significant effort has gone into building small non-dispersive ultra wide bandwidth antennas. Figure 43 is a picture of a set of these. The third from the right on the top row is a meanderline antenna. The majority are variations of a fat dipole or "bowtie" antenna with different feeds. Some are electrically small dielectric loaded elements that sit on a large "ground plane." In reality, the so-called ground-plane is actually the main radiator and essentially the other half of a dipole formed with the dielectrically loaded element. Indeed, a major problem with these antennas was the cable radiation caused by the antenna using "ground" currents on the shield of the coax. A pair of dielectrically loaded elements connected with a balun resulted in a small radiator with good efficiency and low cable radiation. The two antennas on the lower left can be thought of as a folded slot, or a top loaded monopole, again with different feed structures. These antennas had the highest and flattest gain of all the antennas pictured, as well as reasonably low shield currents when the coaxial feed connection was placed on the side, rather than bottom of the antenna (as shown on the bottom left antenna).



Figure 43: UWB antennas developed since the FCC issued UWB rules

Figure 44 is an expanded view of this antenna showing several feeds. This antenna has the following features that are important for this application:

- It is electrically small $\sim \frac{1}{4} \lambda$ at lowest frequency, where the return loss is only 6 dB
- It is Ultra Wide Bandwidth - The -6 dB bandwidth is about the same as the center frequency. Figure 45 shows the transfer function of the antenna.
- It can have a tuning element that creates a narrow notch in the spectrum if needed to reject interference (US Patent 6,590,545 and 6,914,573 and 7,352,333)
- It has a null in the plane of the antennas that will enhance isolation between Tx and Rx antennas. This is important because Tx Rx coupling is effectively a “jamming” term that stresses the dynamic range of the radar.
- It can be conformed to a curved shape if necessary – it does not have to be planar.
- It has a broad beam, which is needed to support wide SAR and ISAR integration angles. Its peak gain is only about 5 dB on boresite at the peak frequency.
- It has low dispersion –its impulse response has low time sidelobes. This is important in that battery draining processing is not required to deconvolve the antenna response. Down-range-resolution will be close to ideal without deconvolution. If deconvolution corrections are used, they will be mild and are not likely to cause serious artifacts. Figure 46 shows the impulse response of a pair of antennas – in other words, the round-trip radar response. This data is the derivative of step response taken with a TDR
- The impulse response is stable across wide integration angles. This is important because the ISAR integration gain and cross-range-resolution will be close to ideal. Figure 47

shows the impulse from several orientations overlaid to illustrate the stability. This data is the real part of the inverse Fourier transform (IFFT) of S21 measured with a vector network analyzer (VNA).

- It can be backed by a reflector. The first test antenna used in the MTRI field tests used a transmit and receive element backed with an aluminum plate. The back plate enhanced the midband gain and gave the antenna a good front to back ratio.
- A cavity backed version, with a modified shape, is also viable. This version was used in MTRI's second field test. The cavity backed version improved the Tx-Rx isolation.




Figure 44: Folded Slot




Figure 45: Facc-to-Face Two-Antenna S21 (nose-to-nose S21, NOT gain)




Figure 46: Face-to-Face Two-Antenna S21




Figure 47: Tight Impulse Response Versus Orientation

To further understand the antenna characteristics, chamber measurements were performed on a scaled antenna – scaled up in frequency by a factor of 5.7. Figure 48 shows the orientation nomenclature used in plotting the results. Figure 49 shows a co-pol E-plane and H-plane frequency vs angle energy plot. Note the natural isolation in the plane of the antenna at 90 and 270 degrees in the H-plane cut data. Note also the broad flat radiation pattern in the $\pm 30^\circ$ from boresite in both elevation & in azimuth. Figure 50 shows a cross-pol E-plane and H-plane frequency vs angle energy plot. Figure 51 shows beam patterns at 4, 5, and 8 GHz which scale to 700MHz, 877MHz, and 1.4 GHz. These patterns show excellent wide angle wide bandwidth imaging capability.




Figure 48: Measurement Orientation Nomenclature




Figure 49: Co-pol E-plane and H-plane Frequency vs Angle energy plots

H-Plane Cut

E-Plane Cut




Figure 50: Cross-pol E-plane and H-plane Frequency vs Angle energy plots




Figure 51: Beam Patterns Scaled to 700 MHz, 877 MHz, & 1.4 GHz

Two antennas were fabricated for MTRI. Figure 52 is a picture showing the first antenna being mounted to a tree. It consisted of a pair of the folded-slot antennas, one for transmit and the other for receive, mounted in front of a reflector plate. Figure 53 shows a picture of the second antenna with the cavity backing. The cavity backing that improves the isolation between the Tx and Rx antenna elements. The antenna was slightly larger than the first to move the bandwidth to lower frequencies that would penetrate wet wood better.



Figure 52: First Antenna Delivered To MTRI



Figure 53: Second Antenna Showing Cavity Backing

Figure 54 shows data collected from the cavity backed antenna delivered to MTRI for its second field test. The reflection coefficient measured with a TDR shows a relatively smooth taper from 50 ohms to an open that approximates a Klopfenstein taper. As a result, the plots of return loss and VSWR show the antenna match improves with frequency and remains good as the frequency goes up. The nose-to-nose impulse response shows low sidelobes, and the transfer function plot shows ± 1 dB of ripple through the passband of the antenna. This ripple, and the sidelobes in the impulse response could be improved by loading the cavity, but at the expense of reduced gain.

In conclusion, an antenna small enough to fit the intended application was fabricated and proved have all the performance metrics needed to perform the high resolution UWB ISAR function.




Figure 54: Trimmed ASI390-T2.5-TG1.8