INTELLIGENT BANDWIDTH COMPRESSION FOR SAR IMAGERY TRANSMISSION

Yoji Niho [1] and Robert Hummel [2]

[1] Raytheon Systems Company, P.O. Box 92426, Los Angeles, CA 90009
(310) 334-3334, ygniho@west.raytheon.com
[2] DARPA ISO, 3701 N. Fairfax Dr. Arlington, VA 22203
(703) 696-2303, rhummel@darpa.mil

ABSTRACT

The DARPA and DARO Intelligent Bandwidth Compression (IBC) project is developing technology to transmit synthetic aperture radar data, such as data collected by ASARS-2, JSTARS, Global Hawk sensors, or the ASARS-2 Improvement Program (AIP) sensor, through a comparatively narrow channel (a T1 transmission line). The challenge is to compress the complex imagery without destroying the ability to recognize and identify targets. The IBC software uses an automatic target recognition algorithm to locate potential target areas, and compresses the background regions more aggressively than the potential target regions. The detected targets are compressed by representing the Fourier coefficients as compactly as possible, while the background is compressed using a wavelet approach. The ability to use spatially varying compression in Wavelet Trellis Encoded Quantization methods is also under investigation. In a recent experiment, the real-time IBC concept was successfully demonstrated using a sequence of spotlight scenes collected in real time by the ASARS-2 system. Radar data were downlinked to the ASARS-2 processing facility in Palmdale, CA where complex SAR images were formed in near real time and immediately compressed using the IBC algorithms by the ASARS-2 ground processor. Background data was compressed at 128 to 1, and detected targets compressed at around 5 to 1. The compressed imagery was then transmitted over a T1 link to a remote site in El Segundo, CA, where they were reconstructed and displayed over the Sensor Control Workstation for real time subjective evaluation by radar image analysts. During the field demonstration, more than 125 ASARS-2 spotlight scenes were successfully processed in real time through the ASARS-2 image formation and IBC signal processing chain, and preliminary evaluation shows that all targets of interest were well preserved in the imagery.

1.0 BACKGROUND

Improved theater awareness drives a need for ever-increasing data rates. For example, a modern surveillance radar like the one onboard the Global Hawk has a nominal image creation rate of 3 Mpix/sec in the stripmap mode. At 32 bits/pixel, this implies a 96 Mbps data rate. How do we move that much data from the sensor to the users in a timely manner? Large bandwidth communications links are not widely available, and when they are, they are expensive.

The above question was the genesis of the Intelligent Bandwidth Compression (IBC) program. The goal of the IBC program has been to reduce the bandwidth required to transmit radar images from approximately 100 Mbps to 1.5 Mbps, the bandwidth of a T-1 connection. Such a large bandwidth reduction calls for a lossy compression strategy. The chosen strategy (and the reason for the "I" in IBC) is to allocate most of the available "bits" to those features in the radar image that are "interesting," i.e. targets. The rest of the image, the "not so interesting" part, receives a correspondingly smaller portion of the information carrying capacity. As a result, target areas are compressed less than non-target areas, or equivalently, target areas are compressed with less distortion than non-target areas. In the IBC program, the above strategy has been implemented in two different ways. One way is known as the "multi-scale" approach. The other way is known as the "single-scale" approach. This document only deals with the "single scale" approach.

Figure 1.1 shows a conceptual diagram of the "single scale" approach. The image is processed by a 128:1 compression module resulting in a moderately high distortion, compressed version of the original. Simultaneously, the image is processed by a target-region extraction and compression module. The output of the second module is a low-distortion, compressed version of the target regions. Reconstruction consists of decompressing the moderately-high-distortion version of the image and overlaying the decompressed low-distortion version of the target regions. Notice that the overall compression factor depends on the number of target-regions extracted and the ratio of "very bright" to "not so bright" targets. See references [1] and [2] for a compression technique used in the single scale approach.

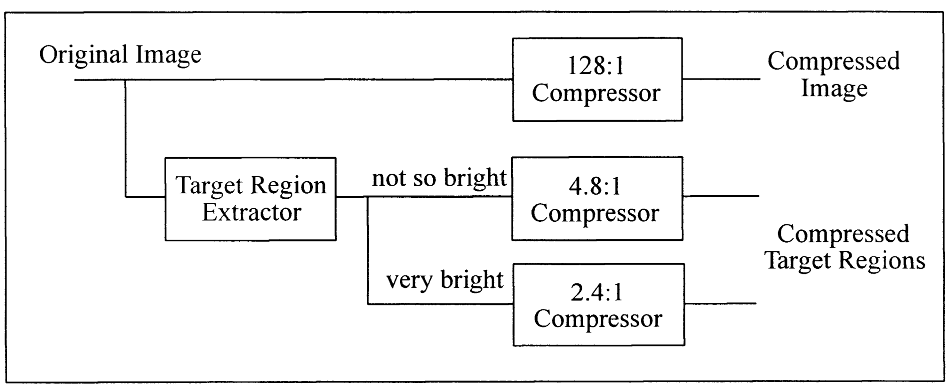

Figure 1.1 SS IBC concept.

The "single-scale" IBC (SS IBC) approach was demonstrated in the field on May 11, 1998. This paper documents the results of the demonstration.

2.0 OBJECTIVE

The objective of the field demonstration on May 11, 1998, was to show empirically that the single-scale (SS) IBC can be used to transfer real-time ASARS-2 Spot-3 images over a T-1 link.

3.0 METHODOLOGY

Figure 3.1 depicts the basic elements of the demonstration.

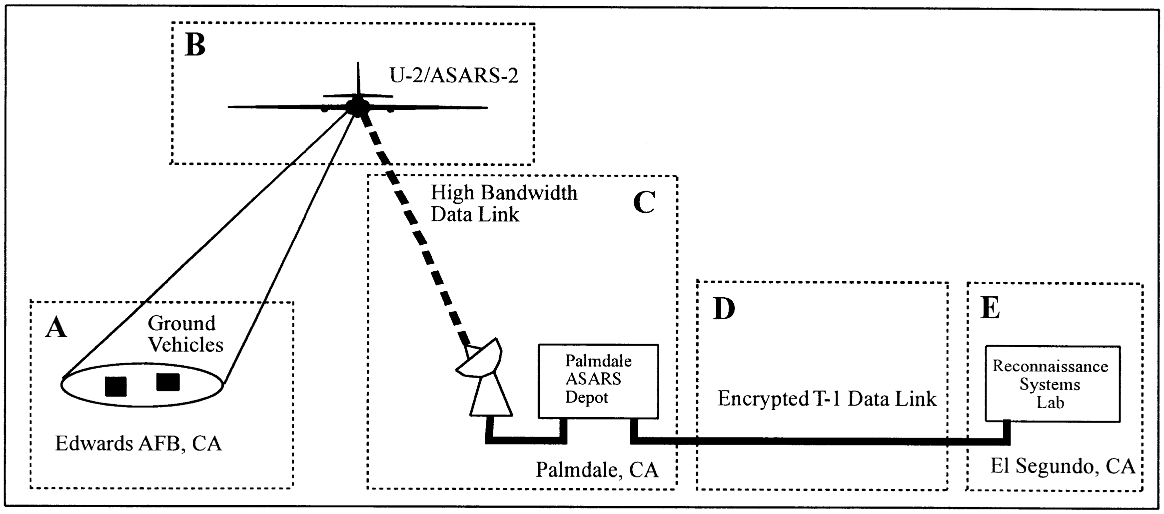

Figure 3.1 SS IBC demonstration set-up.

The demonstration proceeded as follows. A group of stationary vehicles located at Edwards AFB were imaged by an ASARS-2 radar. As the sensor data was collected, it was sent to the ground processing facility via a high-bandwidth data link. Once there, the sensor data was converted into complex SAR images and subsequently compressed using the SS IBC. Finally, the compressed images were sent via an encrypted T-1 link to the Reconnaissance Systems Lab (RSL) in El Segundo, CA where they were reconstructed and displayed.

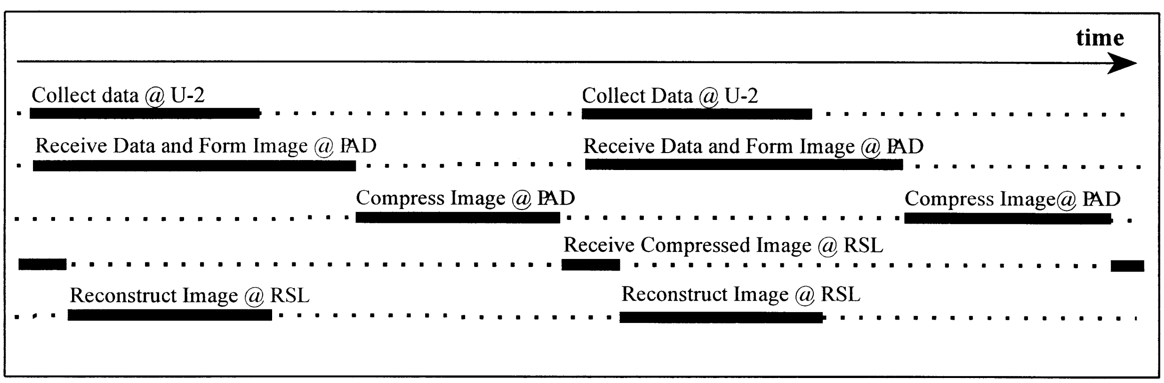

Figure 3.2. SS IBC demonstration timeline.

All the processing occurred in real-time. An observer in the RSL saw a continuous stream of radar images displayed at the radar data collection rate. Figure 3.2 shows a representative timeline. Notice the pipelined architecture.

There were five basic elements in the demonstration set-up: 1) targets, 2) aircraft, 3) ground processing facility, 4) data link, and 5) remote site.

3.1 Targets

The targets were located in an area within Edwards AFB known as the SCATT range. More specifically, they were located within a flat, barren circular area known as PB-11. The PB-11 circle has a radius of approximately 500 feet. Its center is located (WGS-84) at LAT: 34:52:51 North, LON: 117:35:51 West, ALT: 846 meters. At the time of the demonstration, there were 4 trucks, 8 tanks, and 5 armored personnel carriers within the circle. In addition, there were other miscellaneous targets. See Figure 3.3 for a detailed description.

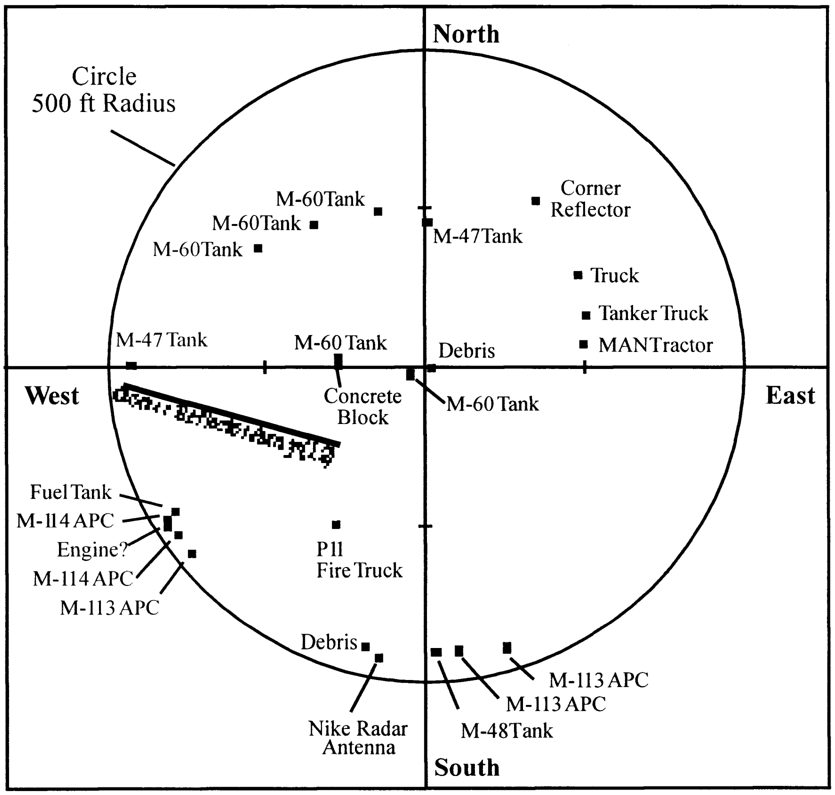

Figure 3.3 Distribution of targets on the day of the demonstration.

Outside the circle there was typical Mojave Desert vegetation: Joshua trees and desert scrub. At the time of the demonstration, desert scrub was beginning to appear within the usually-barren circular area. There were other potential targets in the vicinity of PB-11. Figure 3.4 shows the location of some of them.

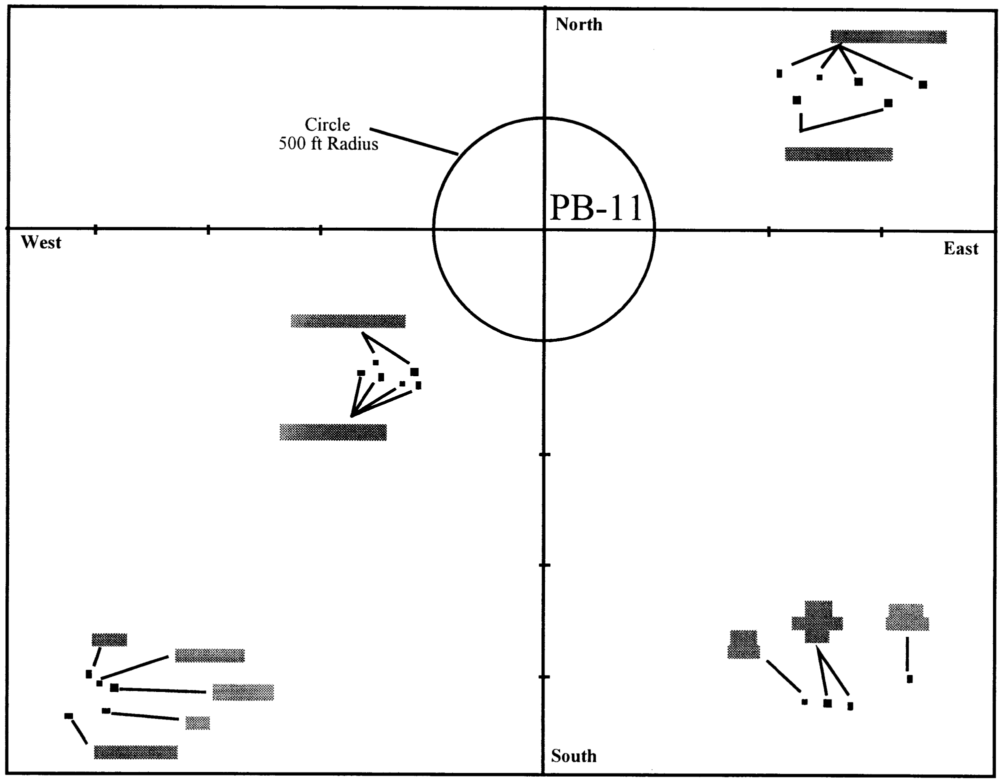

Figure 3.4. Other targets in the vicinity of PB-11.

3.2 Aircraft

A U-2 aircraft carrying an ASARS-2 radar was used for the demonstration. The flight plan, excluding the take-off/climb and descend/land segments, is illustrated in Figure 3.5. The arrows indicate the direction of flight (from segment 1, to segment 2, to segment 3, etc.) As Figure 3.5 shows, there is an inner orbit (segments 1-4) and an outer orbit (segments 5-9). From take-off to landing, the flight lasted approximately 3 hours.

The data collection phase lasted 2 hours and 3 minutes. The collection plan consisted of 142 images of PB-11 interleaved with 15 images of various areas: PB-12 (2), Fort Irwin Main Complex (2), Sewage Treatment Plant (2), Kramer Junction (2), Mojave Airport (2), PIRA Range Control Station (1), Edwards North Base (1), Edwards South Base (1), Santa Catalina Island (1), Haystack Butte (1). All the images were collected using the ASARS-2 Spot-3 mode.

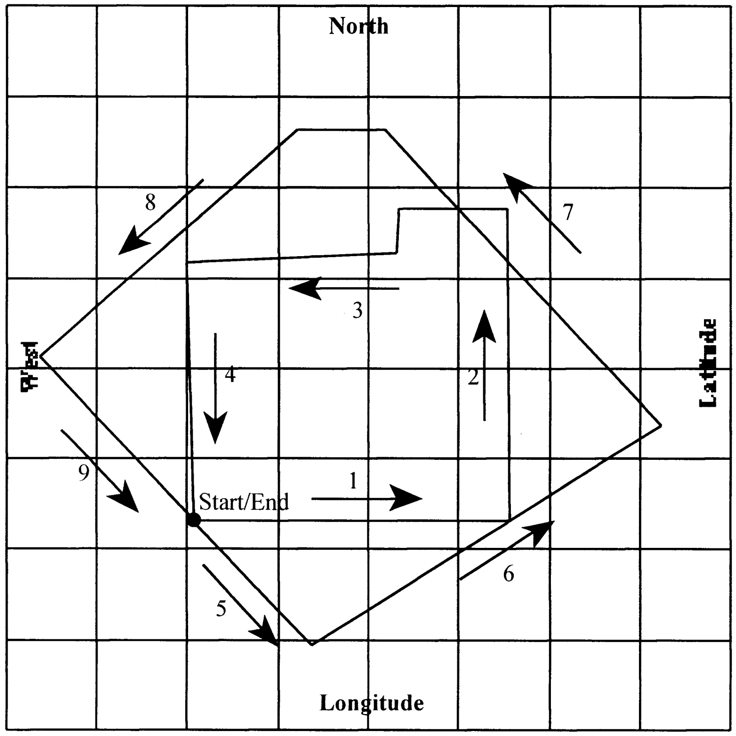

Figure 3.5. U-2 flight path.

3.3 Ground Processing Facility

The ground processing facility was located at the Palmdale ASARS-2 Depot (PAD) in Palmdale, CA, approximately 30 miles SW of the target site. The PAD is a (roughly) 10ft by 25ft shelter that houses the equipment necessary to generate radar images from the downlinked raw video phase history (VPH) data. From an IBC standpoint, the three main PAD components are: (1) the DCRSi high-bandwidth tape recorders used to preserve the VPH data from the U-2, (2) the 32-node COTS Based SAR Processor (CBSP) used to form and compress the images, and (3) the Sun SPARC-20 Sensor Control Workstation (SCWS) that buffers/relays the compressed images to the remote location. A simplified diagram of the PAD contents is shown in Figure 3.6.

3.4 Data Link

A T-1 data link connected the ground processing facility and the remote site. The nominal information rate of a T-1 link is 1.536 Mbps plus 0.008 Mbps of telephone company overhead. The network protocol used for this demonstration was TCP/IP.

Figure 3.7 illustrates the basic components of the link. Because the radar images are classified, the link was encrypted. With that in mind, KG-94A encryption/decryption devices were placed between the Data Service Unit/ Channel Service Unit (DSU/CSU) and the routers.

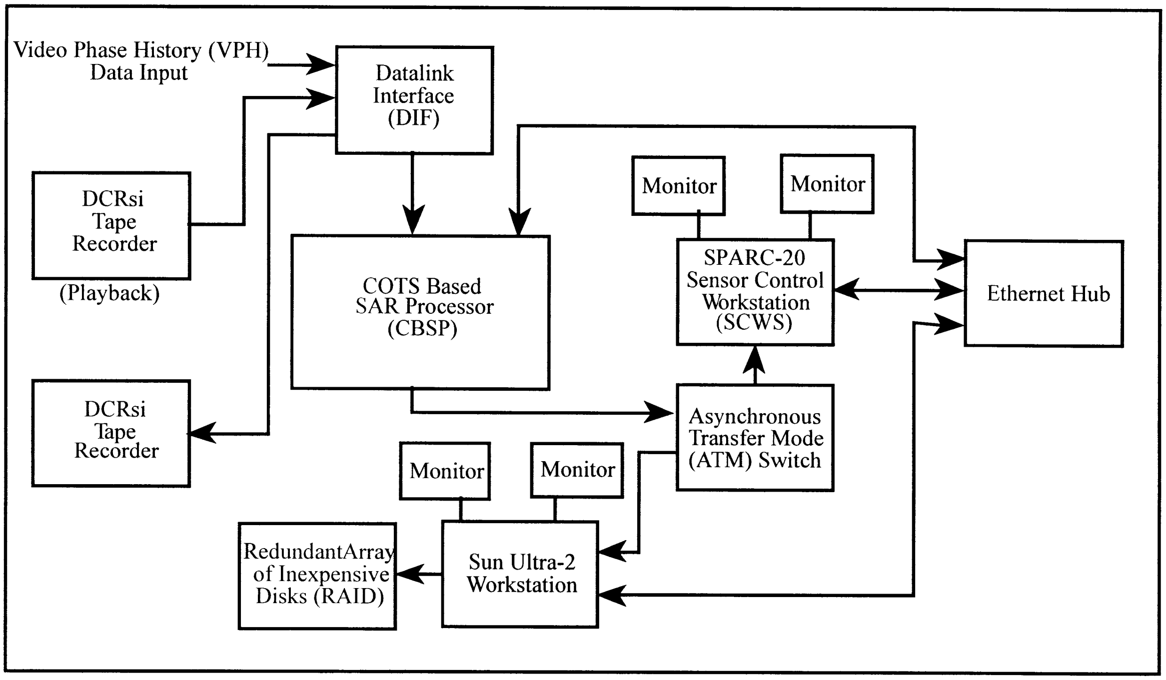

Figure 3.6. Simplified PAD diagram.

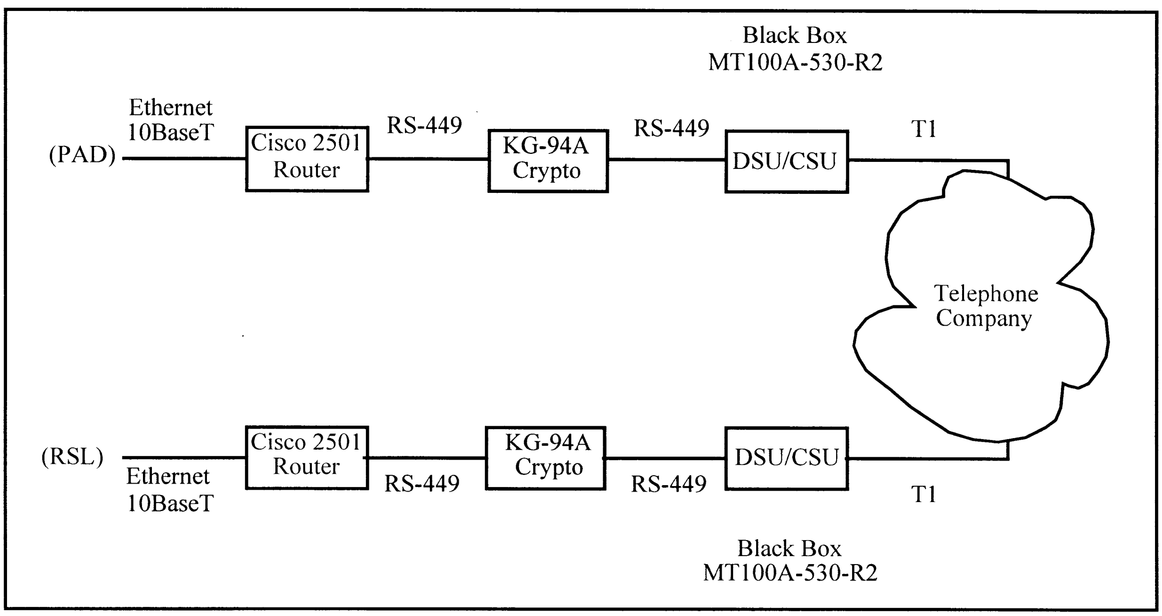

Figure 3.7. T-1 data link.

3.5 Remote Site

The remote site was located at the Reconnaissance Systems Lab (RSL) in El Segundo, CA, approximately 80 miles SW of the target site. At the end of the T-1 link, there was a Sun Ultra-2 workstation configured as a remote sensor control workstation. This workstation reconstructed the compressed images and displayed them to the user. The

Ultra-2 had two UltraSPARC-II processors running at 300 MHz. The published floating point and integer benchmarks for this processor are 20.2 SPECfp95 and 12.3 SPECint95, respectively. The workstation had 2 MB of cache and 512 MB of RAM.

4.0 RESULTS

4.1 Collection

A total of 160 scenes were collected: 157 scenes in the original mission plan plus 3 scenes that were added at the beginning of the collection. Of the 160 scenes, 129 were successfully compressed/reconstructed in real time. Of the remaining 31 scenes, 4 were aborted due to VPH errors, 3 "hung up" the CBSP during SS IBC background compression, 8 were lost while recovering from the "hang ups," and 16 were missed during SS IBC reconstruction due to threading errors. Except for the VPH errors, which are not related to IBC, all the other errors have been corrected in the current release of the SS IBC software.

Figure 4.1 shows the position of the U-2 (small circles) relative to the planned trajectory (solid line) for all the PB-11 scenes. The gaps indicate the U-2 positions for the non-PB-11 scenes.

Figure 4.2 indicates the relative time of collection for all the PB-11 scenes. Again, the gaps in the Figure indicate non-PB-11 targets. The total elapsed time was 2hr 2min 57sec. From the figure, the collection rate was fairly steady, averaging 47.3 sec from the start of data collection of one scene to the start of data collection of the following scene. That provides more than enough time to collect the data, form the image, and compress the data. For reference, the average time to form and compress an image was 16.2 sec.

Figures 4.3 and 4.4 show the measured speed of the U-2 and its distance to PB-11 during the collection. The aircraft speed remained within a 40-knot window while the PB-11 images were being collected.

4.2 Target Extractor

Figure 4.5 shows the number of target regions or "chips" for all the PB-11 scenes. Only scenes 149 and 150 had the maximum number of target chips (1750). Except for these two scenes, all the PB-11 scenes had the same number of Focus of Attention (FOA) pre-screener "blobs" and target chips. This indicates that most of the discrimination was done by the FOA pre-screener. It is possible that, even for those two scenes, all the "blobs" looked like targets to the detector but that it had no choice but to reject some of them because of the built-in limit on the number of target chips (1750). For non-PB-11 scenes, the pattern was the same: unless the number of "blobs" exceeded 1750, the number of target chips was equal to the number of "blobs."

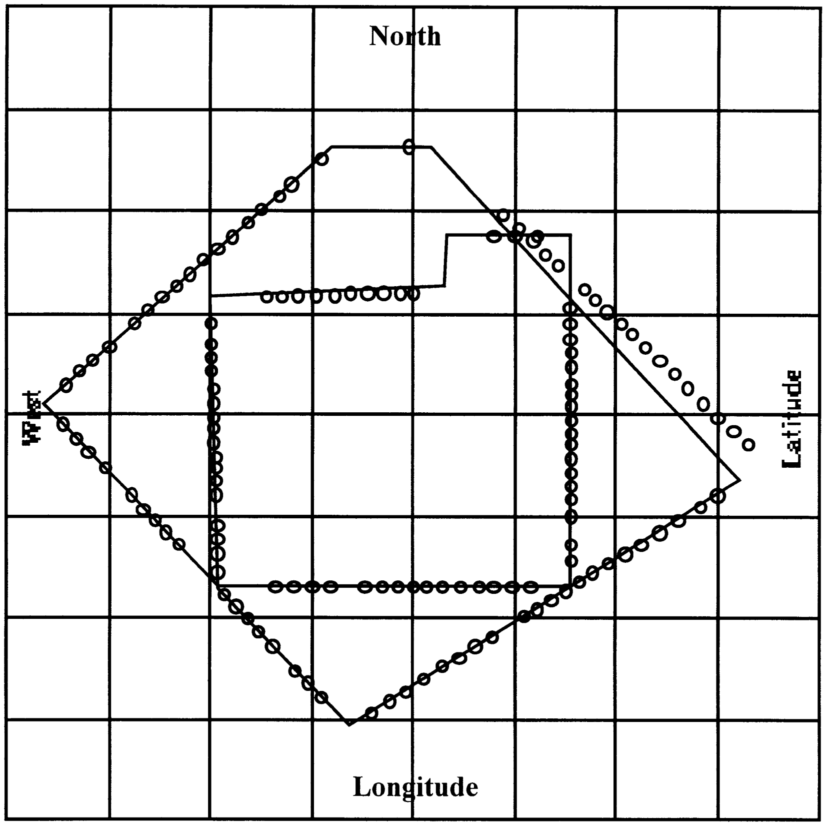

Figure 4.1. PB-11 scenes vs. planned flight path.

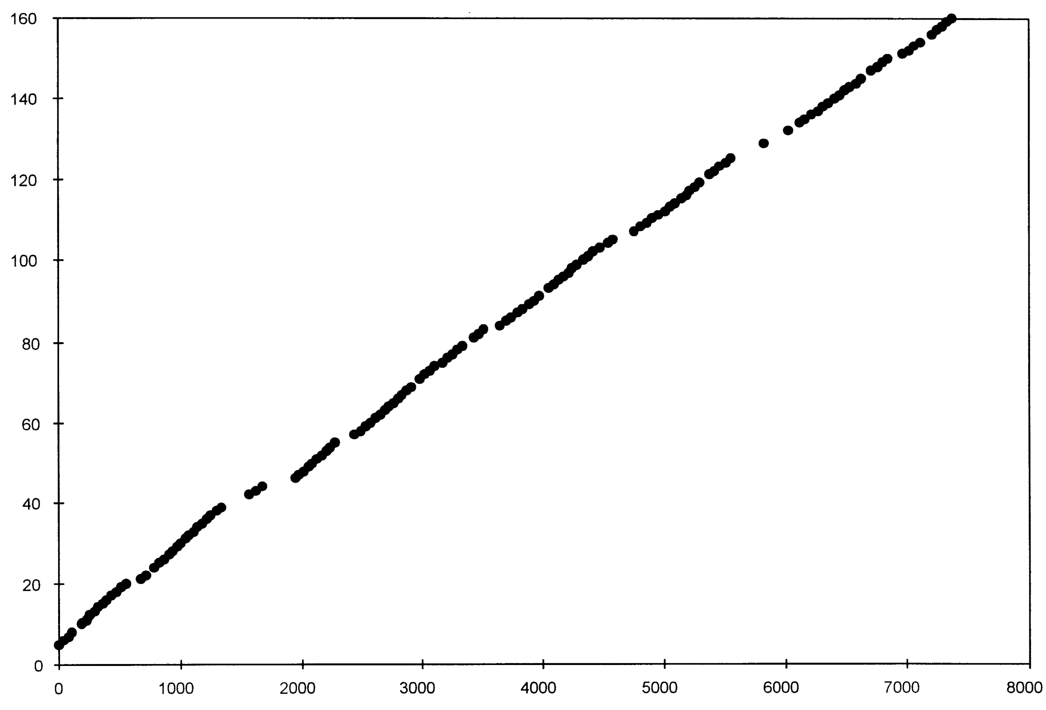

Figure 4.2. Relative time of collection for all the PB-11 scenes.

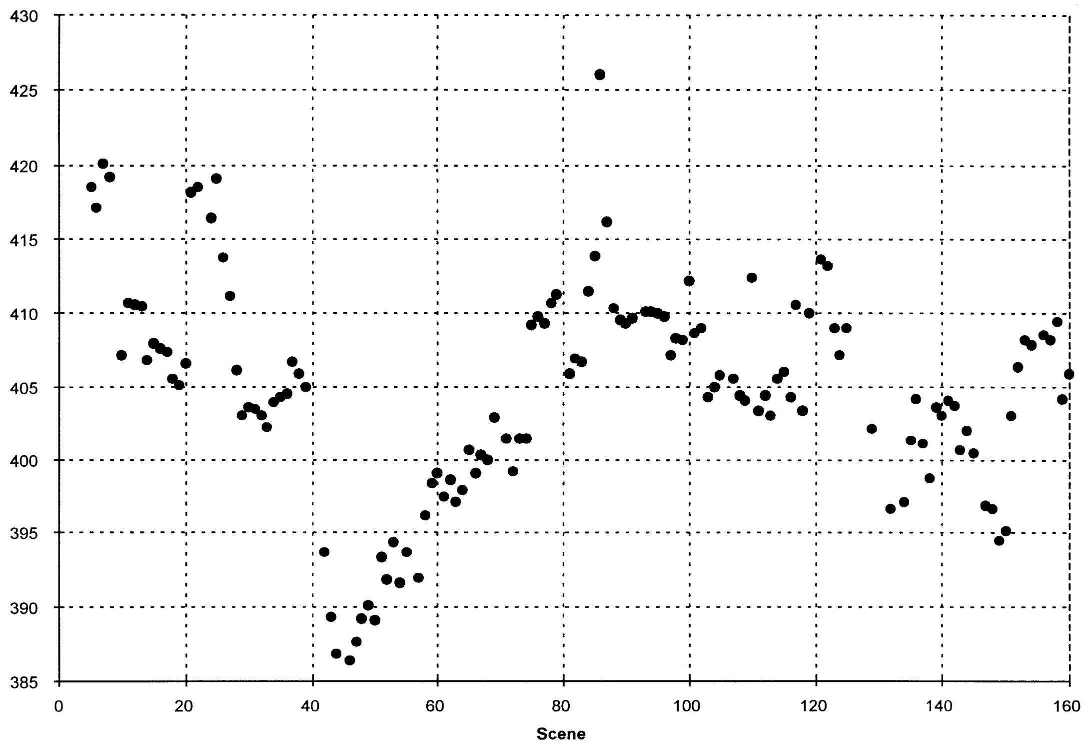

Figure 4.3. U-2 measured speed.

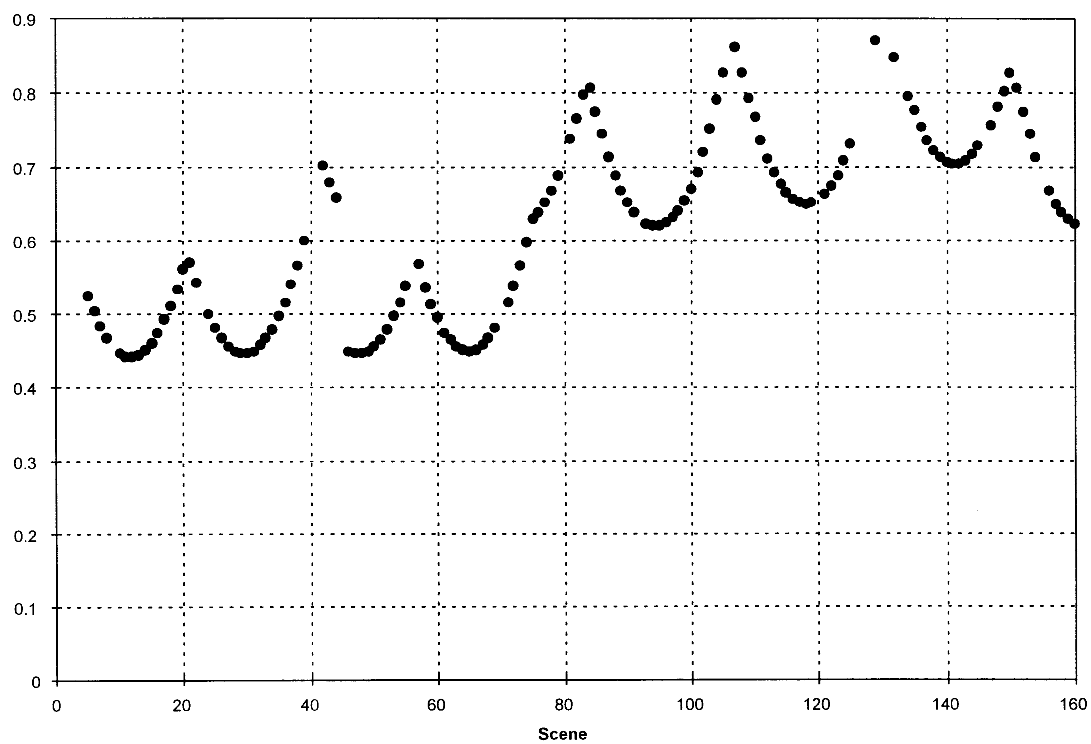

Figure 4.4. Normalized distance between the U-2 and PB-11.

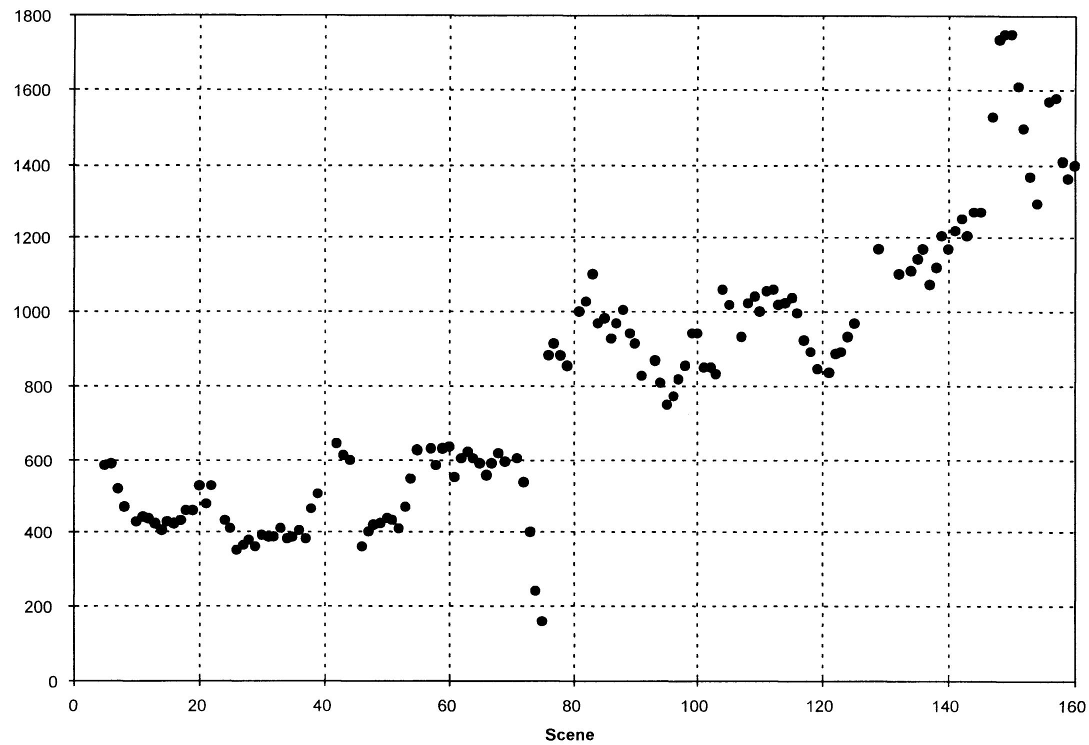

Figure 4.5 Number of target chips for PB-11 scenes.

Unexpectedly, there was a very strong correlation between the range and the number of target chips. Compare Figures 4.4 and 4.5. The range correlation was even evident as the range varied within a flight leg. For the inner loop (closer range), the number of target chips varied from 241 to 646 (the average was 483). For the outer loop (farther range), the number of target chips varied from 749 to 1750 (with an average of 1085). This excludes the first image of the outer loop that was blurred and as a result, the number of target chips was abnormally low (158 chips). Mysteriously, three inner loop scenes (42-44) with almost the same range as several outer loop scenes (but different squint angle) have about half the number of target chips of the corresponding outer loop scenes.

For the 17 vehicles within the PB-11 circle, the probability of detection was 98.1%. The probability of detection for the inner loop, was slightly higher than for the outer loop: 98.7% vs. 97.6% (even though the inner loop had fewer target chips). About two thirds of the 16/17 cases belonged to the outer loop. Since the FOA pre-screener did most of the discrimination, practically all the misses can be attributed to an FOA pre-screener rejection. Figure 4.6 shows the number of vehicle detections for every PB-11 scene.

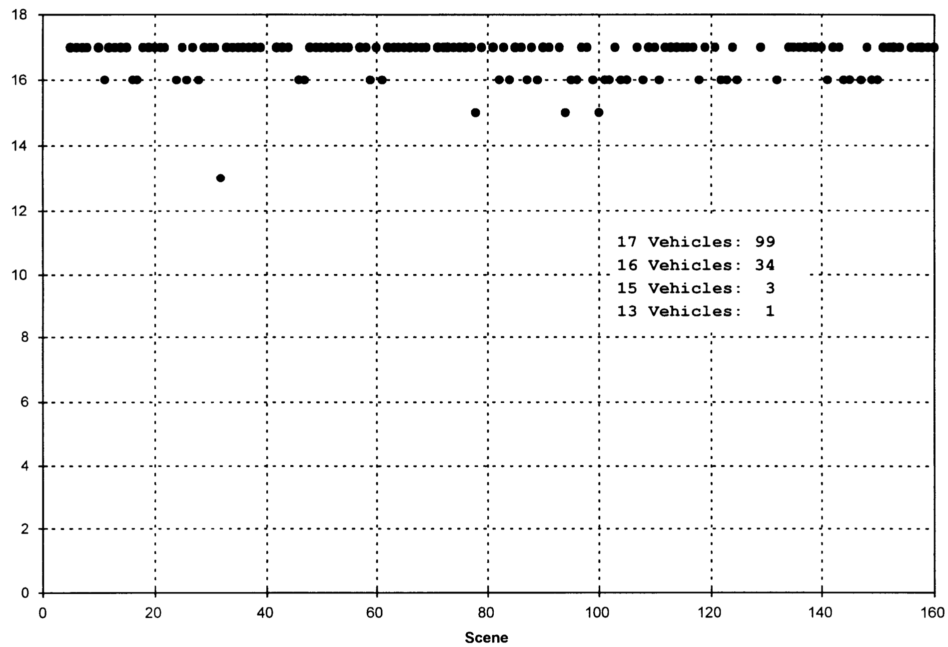

Figure 4.6. Number of vehicle detections for PB-11 scenes.

To determine the distribution of target chips over the PB-11 area, an 85 by 85 element imaginary grid was laid on top of the target area. After that, we counted the number of times that a detection fell within each grid element for all the PB-11 scenes. (This is essentially a two-dimensional histogram.) Clearly, a grid element with a high count indicates that targets were consistently detected in that area. Figure 4.7 shows the resulting grid. North is up and East is to the right. The darker grayshades are associated with low counts and the lighter grayshades with high counts. Notice that the targets within the PB-11 circle (central part of Figure 4.7) result in grid elements with high counts. Also notice that the distribution of target false alarms is more or less uniform. There are somewhat fewer target false alarms Northeast of the PB-11 circle. This agrees with the observation that the terrain in this region has slightly lower reflectivity, and therefore, is less likely to generate an FOA pre-screener "blob."

4.3 Compressor

The total time required to form and compress the images (batch time) varied from 15.9 sec to 17.1 sec. over the collection. The average batch time was 16.2 sec. Assuming 5.2 sec to form the image, that indicates an average compression time of 11 sec. See Figure 4.8. For batch times between 15.9 sec and approximately 16.5 sec, there was a strong correlation between the number of targets and the batch time. Beyond 16.5 sec., there was a weak correlation between the number of targets and the batch time. Possibly, that indicates that for the scenes with the longer batch times, the background compressor was the bottleneck. See Figure 4.9.

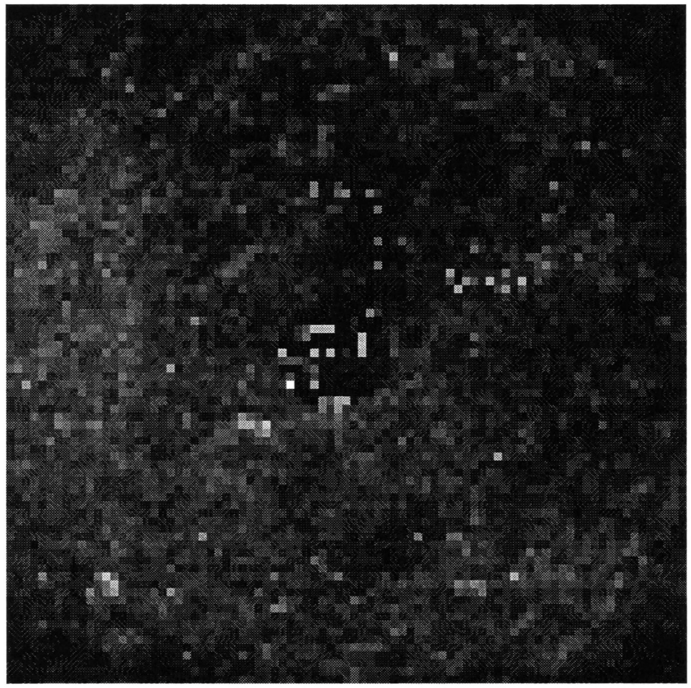

Figure 4.7. Spatial histogram of target detections.

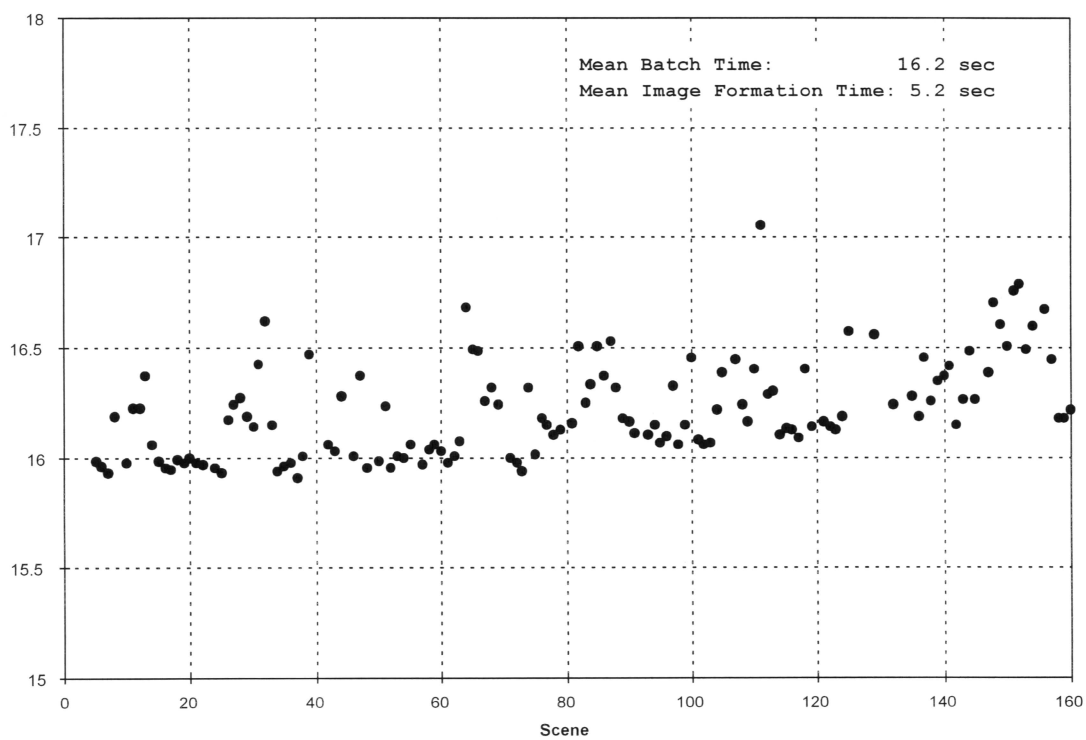

Figure 4.8. Total batch time for PB-11 scenes.

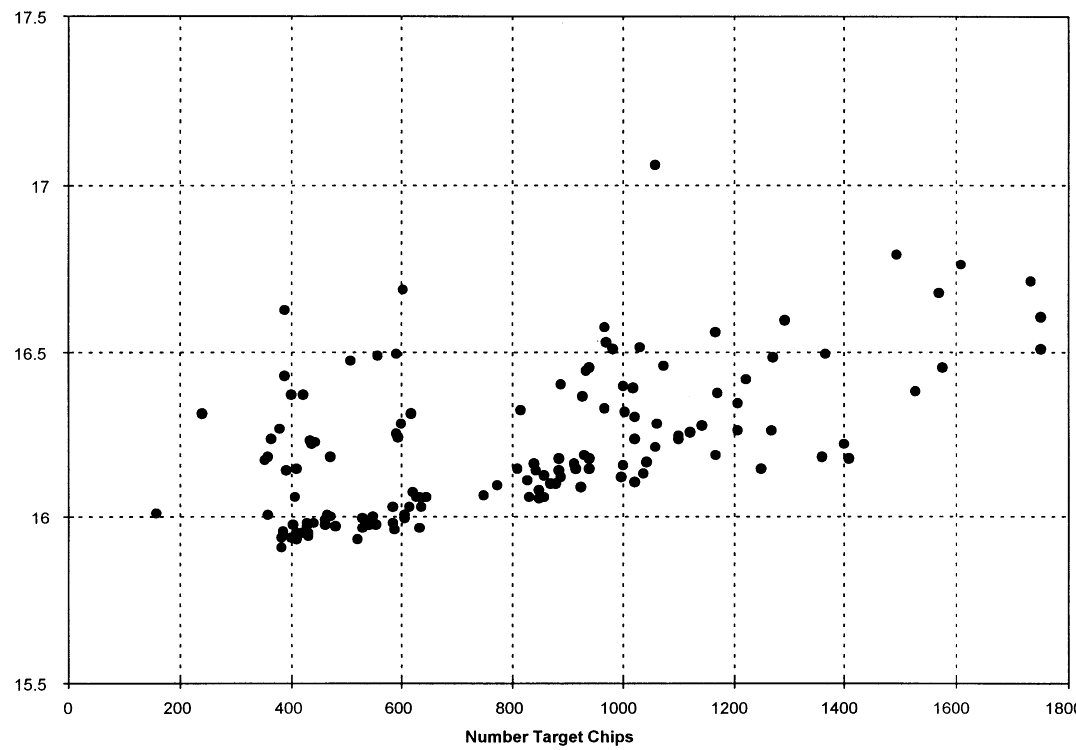

Figure 4.9. Total batch time vs. number of target chips for PB-11 scenes.

As mentioned before, the number of target chips affects the overall compression ratio. For the inner loop, the size of the compressed files varied from 0.55MB to 0.75MB (average: 0.62MB). For the outer loop, the size of the compressed files varied from 0.84MB to 1.66MB (average: 1.14MB). This implies compression factors between 38:1 and 51:1 (average 46:1) for the inner loop and between 17:1 and 34:1 (average 26:1) for the outer loop. See Figures 4.10 and 4.11.

Assuming a T-1 data rate of 1.536Mbps and negligible protocol overhead, it would take 3.4 sec to transmit the average inner loop compressed file and 6.1 sec to transmit the average outer loop compressed file. See Figure 4.12. In comparison, it would take 154.1 sec to send the uncompressed version of the image over the same communications link.

The ratio of "very bright" target chips (lower compression) to "not so bright" target chips (higher compression) also affects the overall compression factor. As Figure 4.13 shows, this ratio was fairly independent of the scene. Except for two scenes where the ratio was equal to 3.6% and 3.2%, most of the values were below 2% (average: 0.6%). This indicates that the lower compression factor of the bright chips did not significantly affect the overall compression factor.

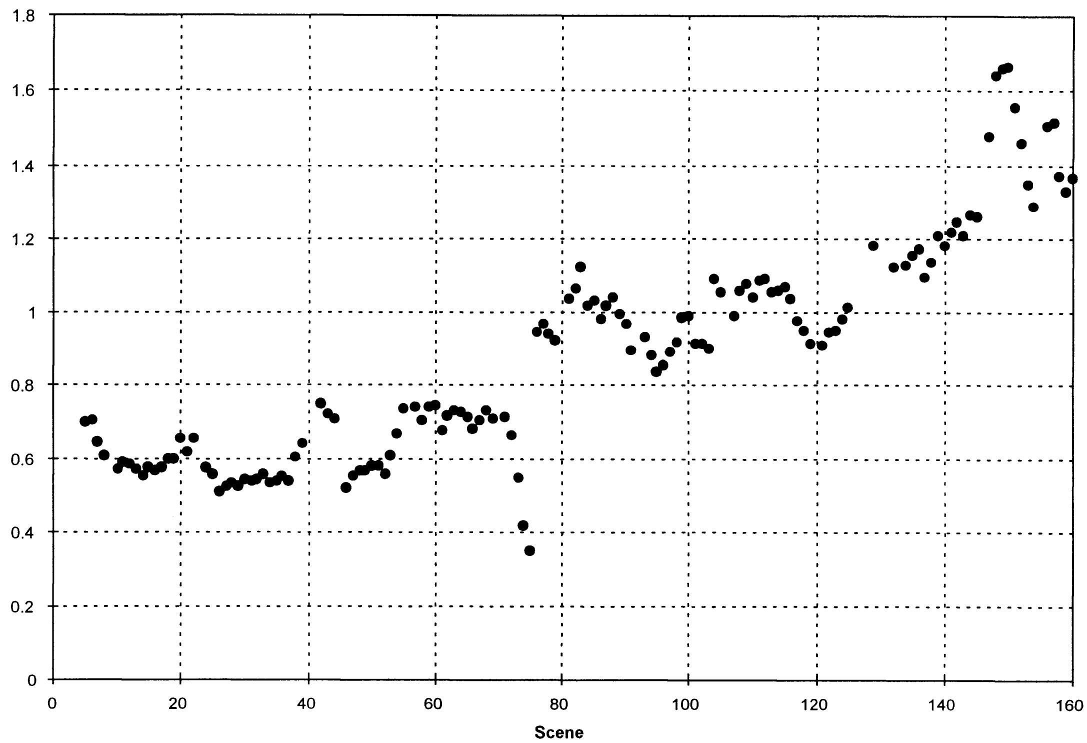

Figure 4.10. Compressed image size for PB-11 scenes.

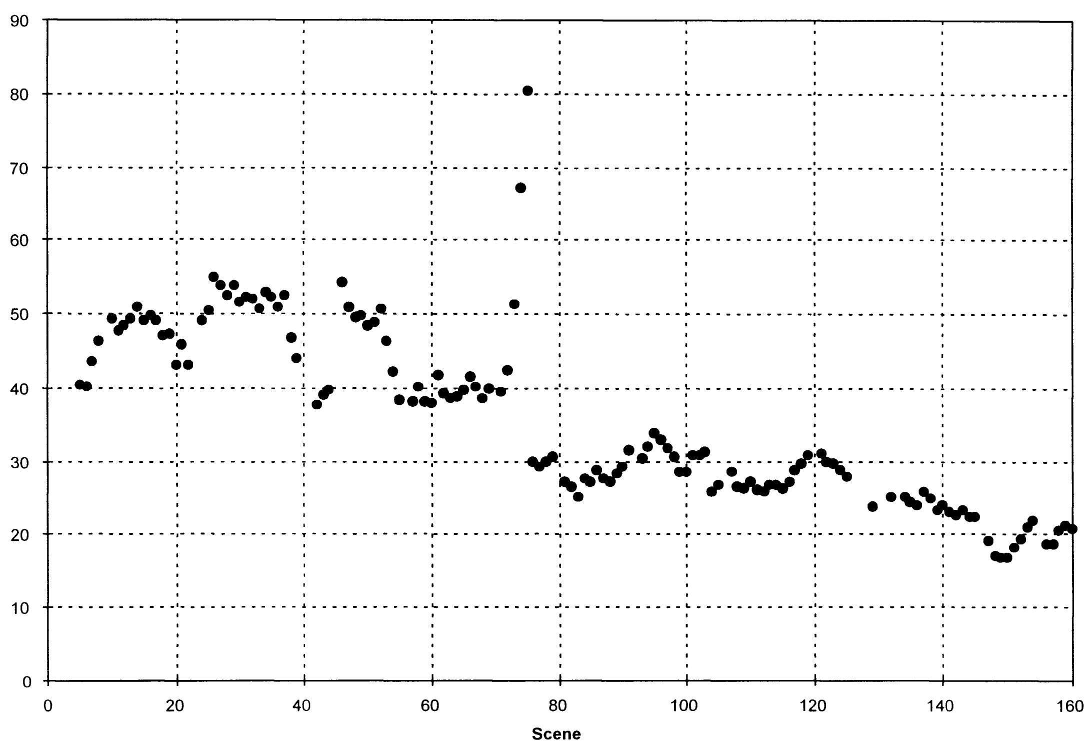

Figure 4.11. Compression factor for PB-11 scenes.

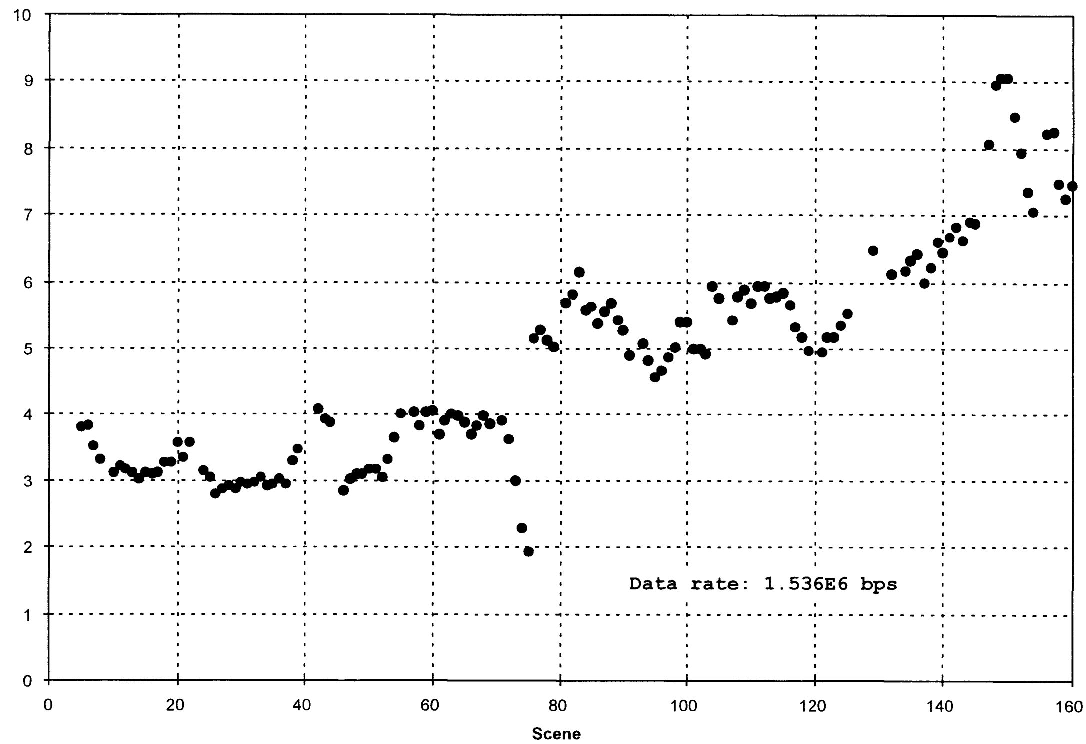

Figure 4.12. Projected time to transmit compressed images.

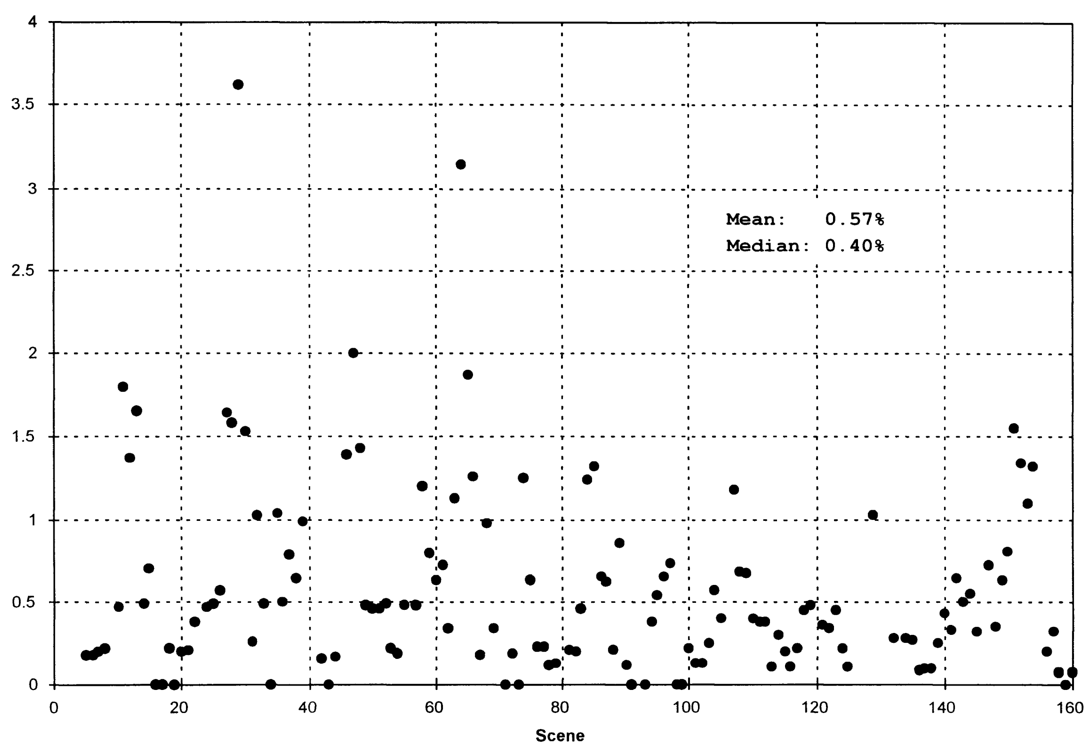

Figure 4.13. Percentage of bright chips (low compression) for PB-11 scenes.

5.0 CONCLUSIONS

The experiment successfully demonstrated that SS IBC can be used to transfer real-time, ASARS-2 Spot-3 images over a T-1 link. The experimental data showed that:

1. For the 17 vehicles within the PB-11 circle, the probability of detection was 98.1%.
2. Most of the discrimination was done by the FOA.
3. False alarms were distributed more or less uniformly over the target area.
4. There was a correlation between the range to the targets and the number of target chips.
5. The average compression time was 11 sec.
6. The average compression factor was 46:1 at close ranges and 26:1 at the far ranges.
7. The lower compression of the bright target chips did not significantly affect the overall compression factor.

6.0 REFERENCES

[1] Robert W. Ives, Paul Eichel and Neeraj Magotra, "Intelligent Low Rate Compression of Speckled SAR Imagery," IEEE 1998 National Radar Conference, 12-13 May 1998, Dallas Texas

[2] Paul H. Eichel and Robert W. Ives, "Very Low Rate Compression of Speckled SARA Imagery," SPIE Vol. 3391, Wavelet Applications V, 14 April 1998, Orlando, Florida