

**Project Report
ATC-236**

**Description of Radar Correlation and
Interpolation Algorithms for the ASR-9
Processor Augmentation Card (9-PAC)**

**O.J. Newell
J. R. Anderson**

27 October 1995

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LExINGTON, MASSACHUSETTS




Prepared for the Federal Aviation Administration,
Washington, D.C. 20591

This document is available to the public through
the National Technical Information Service,
Springfield, VA 22161

This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof.


1. Report No. ATC-236	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle Description of Radar Correlation and Interpolation Algorithms for the ASR-9 Processor Augmentation Card (9-PAC)		5. Report Date 27 October 1995	6. Performing Organization Code
7. Author(s) Oliver J. Newell and John R. Anderson		8. Performing Organization Report No. ATC-236	
9. Performing Organization Name and Address Lincoln Laboratory, MIT 244 Wood Street Lexington, MA 02173-9108		10. Work Unit No. (TRAI5)	11. Contract or Grant No. DTFA01-93-Z-02012
12. Sponsoring Agency Name and Address Department of Transportation Federal Aviation Administration Washington, DC 20591		13. Type of Report and Period Covered	
15. Supplementary Notes This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology under Air Force Contract F19628-95-C-0002.		14. Sponsoring Agency Code	
16. Abstract MIT Lincoln Laboratory, under sponsorship from the Federal Aviation Administration (FAA), is conducting a program to replace/upgrade the existing ASR-9 array signal processor (ASP) and associated algorithms to improve performance and future maintainability. The ASR-9 processor augmentation card (9-PAC) replaces the ASP four-board set with a single card containing three TMS320C40 processors and 32 Megabytes of memory. The resulting increase in both processing speed and memory size allows more sophisticated beacon and radar processing algorithms to be implemented. The majority of the improvement to the radar correlation and interpolation (C&I) function lies in the area of geocensoring and adaptive thresholding, where the larger memory capacity of the 9-PAC allows more detailed maps to be maintained. A dynamic road map mechanism has been implemented to reduce the need for manual tuning of the system when the radars are first installed or when new road construction occurs. The map is twice the resolution of the original geocensor map, resulting in a decrease in total area desensitized to radar-only targets. In addition, the new geocensor mechanism makes use of target amplitude information, allowing aircraft with amplitudes significantly greater than the road traffic returns at a particular cell to pass through uncensored. The adaptive thresholding cell geometry has been modified so that adaptive map cells now overlap one another, eliminating the false target breakthrough that occurs in the present system when regions of false alarms due to birds or weather transition from one cell to the next. The entire C&I function has been recoded in a high-level language (ANSI-C), allowing it to be easily ported between platforms and better facilitating off-line analysis.			
17. Key Words ASR-9 algorithm 9-PAC geocensor radar road		18. Distribution Statement This document is available to the public through the National Technical Information Service, Springfield, VA 22161.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 98	22. Price



ABSTRACT

M.I.T. Lincoln Laboratory, under sponsorship from the Federal Aviation Administration (FAA), is conducting a program to replace/upgrade the existing ASR-9 Array Signal Processor (ASP) and associated algorithms to improve performance and future maintainability. The ASR-9 Processor Augmentation Card (9-PAC) replaces the ASP four-board set with a single card containing three TMS320C40 processors and 32 Megabytes of memory. The resulting increase in both processing speed and memory size allows more sophisticated beacon and radar processing algorithms to be implemented.

The majority of the improvement to the radar correlation and interpolation (C&I) function lies in the area of geocensoring and adaptive thresholding, where the larger memory capacity of the 9-PAC allows more detailed maps to be maintained. A dynamic road map mechanism has been implemented to reduce the need for manual tuning of the system when the radars are first installed or when new road construction occurs. The map is twice the resolution of the original geocensor map, resulting in a decrease in total area desensitized to radar-only targets. In addition, the new geocensor mechanism makes use of target amplitude information, allowing aircraft with amplitudes significantly greater than the road traffic returns at a particular cell to pass through uncensored. The adaptive thresholding cell geometry has been modified so that adaptive map cells now overlap one another, eliminating the false target breakthrough that occurs in the present system when regions of false alarms due to birds or weather transition from one cell to the next. The entire C&I function has been recoded in a high-level language (ANSI-C), allowing it to be easily ported between platforms and better facilitating off-line analysis.



ACKNOWLEDGMENTS

The authors wish to acknowledge the contributions of Ed Hall, Jim Pieronek, and Greg Rocco in designing and building the 9-PAC prototype boards and associated hardware. Gabe Elkin and Walter Heath wrote real-time recording and flash card file system software used in the C&I analysis. Jenifer Evans wrote the 9-PAC radar tracker and performed analysis on the new C&I outputs. Bill Goodchild contributed to the design of the geocensor map. Wesley Johnston and Paul Wenz aided the Albuquerque field site testing effort.




TABLE OF CONTENTS


<u>Section</u>	<u>Page</u>
ABSTRACT	iii
ACKNOWLEDGMENTS	v
LIST OF ILLUSTRATIONS	ix
LIST OF TABLES	xi
1. INTRODUCTION	1
2. ASR-9 BACKGROUND	3
2.1 ASR-9 Front-End Radar Processing	4
2.2 9-PAC Background	7
3. C&I REQUIREMENTS	11
4. 9-PAC C&I ALGORITHM DESCRIPTION	13
4.1 Initialization/Reset	14
4.2 Input Parsing and Geo/MTI Flagging	14
4.3 Saturation and ZVF Overflow Processing	15
4.4 Correlation of Radar Primitives	15
4.5 RFI Processing	18
4.6 Interpolation	18
4.7 Target Reformatting	25
4.8 Delay Processing	25
5. GEOCENSORING AND ADAPTIVE THRESHOLDING	27
5.1 Overview	27
5.2 Geocensoring	28
5.3 Multi-Grid Adaptive Thresholding	34
6. TWO-LEVEL WEATHER	39
6.1 Temporal Smoothing	39
6.2 Contouring	39
7. VARIABLE SITE PARAMETERS	41
7.1 RFI Processing	41
7.2 Interpolated Doppler Processing	41
7.3 Geocensoring and Adaptive Thresholding	42

TABLE OF CONTENTS
(Continued)

<u>Section</u>	<u>Page</u>
7.4 Two-Level Weather	45
7.5 Miscellaneous	46
8. PERFORMANCE MONITORING/ALARMS	47
8.1 Performance Monitoring	47
8.2 Alarms	48
9. CURRENT STATUS AND FUTURE WORK	49
APPENDIX A. C&I STATE AND DATA FLOW DIAGRAMS	51
APPENDIX B. INPUT/OUTPUT FORMATS	69
APPENDIX C. C&I CONSTANTS	79
GLOSSARY	83
REFERENCES	85

LIST OF ILLUSTRATIONS

<u>Figure</u>		<u>Page</u>
1.	ASR-9 range/azimuth/PRF layout.	3
2.	ASR-9 front-end signal processing.	4
3.	ASR-9 STC attenuation vs. range (typical).	5
4.	Doppler filter bank (high PRF).	5
5.	9-PAC block diagram.	8
6.	C&I processing block diagram.	13
7.	Detection of multiple targets in a primitive range rroup.	16
8.	Antenna azimuth beam pattern (low beam).	21
9.	Azimuth correction vs. CPI amplitude difference.	22
10.	Multigrid thresholding stages.	27
11.	Survey map word format.	30
12.	Geomap amplitude histogram.	31
13.	Hit history decay curve, $\alpha = 321$.	32
14.	Unmodified single pole filter thresholding time lag.	33
15.	Target exclusion processing.	35
16.	Counting and thresholding cell geometry.	36
17.	Adaptive map histogram-based threshold setting technique.	37
18.	C&I processing top level block diagram.	51
19.	Input processing state transition diagram.	52
20.	Front-end geocensor processing.	53
21.	ZVF overload processing.	54
22.	Saturation preprocessing.	55
23.	Target correlation.	56
24.	Association / range resolution processing.	57
25.	Primary RFI processing.	58
26.	Azimuth centroiding.	59
27.	Doppler interpolation.	60
28.	Doppler averaging.	61
29.	Supplemental RFI processing.	62
30.	Geocensoring and adaptive thresholding.	63
31.	Geomap threshold update.	64
32.	Geo enable map update.	65
33.	Adaptive map threshold update process.	66
34.	Geo/adaptive map target count update process.	67
35.	CPIP input data layout.	69




1. INTRODUCTION

The ASR-9 Processor Augmentation Card (9-PAC) provides increases in processing speed, memory size, and programming ease when compared with the existing ASR-9 Array Signal Processor (ASP). These increased capabilities allow for some of the known problem areas of the ASR-9 to be addressed with improvements to the existing radar and beacon processing algorithms. This report describes the improvements to the radar Correlation and Interpolation (C&I) process, as well as providing implementation details for the C language version of the entire C&I function.

Unlike the new 9-PAC beacon algorithms [1], which are fundamentally different from the original ASP version in many respects, a substantial fraction of the C&I algorithms remains essentially intact, with the major 9-PAC enhancement consisting of porting the ASP assembly language to a high-level language (C). The exceptions to this are the areas of geocensoring and adaptive thresholding, where the large memory capacity of the 9-PAC allows for more detailed geocensoring and thresholding 'maps' to be maintained, allowing for more intelligent rejection of false targets (roads, weather, birds) and a corresponding gain in system sensitivity to aircraft targets.

Sections 2 and 3 of this document provide background information on the ASR-9, 9-PAC, and C&I requirements. Section 4 describes the subset of C&I algorithms that were ported to the C language with only minor enhancements. Section 5 describes the new geocensoring and adaptive thresholding algorithms. Section 6 supplies details of the two-level weather processing algorithms, which are not part of the C&I processing per se, but are grouped with C&I because the weather data is embedded in the radar data stream. Sections 7 and 8 describe the Variable Site Parameters (VSPs) and Performance Monitoring/Alarm outputs. Section 9 summarizes the work done to date and discusses possible future enhancements. Lastly, appendices contain algorithm flowcharts, report formats, and constant tables.



2. ASR-9 BACKGROUND

To better understand the algorithm descriptions that follow, it is useful to be familiar with some basic characteristics of the ASR-9. Table 1 lists the most relevant system parameters.

TABLE 1: ASR-9 System Parameters

Antenna	Cosecant ² Fan Beam. Single transmit beam, Dual receive beams (1 passive)
Rotation Rate	12.5 RPM +/- 10%
Frequency	2.7-2.9 GHz
Transmit Power	1.12 MW (Peak)
Pulse Width	1.03 μ s
Azimuth Beamwidth	1.4 deg.
Elevation Beamwidth	4.8 deg. (min)
PRF	Block Staggered. Average PRF between 1059-1172
A/D Converters	12-bit
Range Gate Size	1/16th NMI
System Range	60 NMI

The ASR-9 transmits using a block staggered PRF to help eliminate blind speeds. Ten pulses are transmitted at the high PRF, followed by eight pulses at the low PRF. Additional 'fill' pulses (usually one or two, but dependent on PRF and antenna wind loading) are transmitted at the low PRF until the antenna reaches the next 1.4 degree sector, when the sequence is repeated. Each series of ten or eight pulses is termed a 'Coherent Processing Interval' (CPI), with the pair referred to as a 'CPI Pair.' An entire antenna scan contains 256 1.4 degree CPI Pairs. Range cells in each CPI are spaced at 1/16th nm intervals, with a total of 960 gates providing the full 60 nm coverage. This range, azimuth layout is depicted in Figure 1.




Figure 1. ASR-9 range/azimuth/PRF layout.

2.1 ASR-9 FRONT-END RADAR PROCESSING

The ASR-9 performs several front-end signal processing/thresholding operations prior to sending radar data to the 9-PAC, illustrated in Figure 2 [2]. High- and low-beam data first passes through separate digital 0-63 dB STC attenuators to prevent receiver saturation at close range. Data from both beams are then input to a high-speed waveguide switch capable of switching between the high and low beam at some point during each pulse. The crossover point is typically 15 nm, with the high beam enabled at close range to reduce ground clutter contamination while the low beam is utilized at further ranges to provide adequate low-altitude coverage. The composite low-/high-beam data is fed to the radar receiver, which outputs 12-bit I and Q A/D samples. Individual radar pulses are then checked for RFI and receiver saturation and flagged if necessary. The A/D samples are passed through a Doppler filter bank, power combined, and subjected to Constant False-Alarm Rate (CFAR) and geocensoring operations before being output to the 9-PAC.




Figure 2. ASR-9 front-end signal processing.

2.1.1 STC

The STC attenuation vs. range used at a typical ASR-9 site is shown in Figure 3. The curve falls as $1/R^4$, providing for relatively constant point target amplitudes out to the range where the beam switch occurs or the STC attenuation reaches zero. Although the STC curves shown are typical, the STC decay rate can, in fact, be set independently for 12 range regions for adaptability to a variety of clutter environments. In general, because target amplitudes may vary with range and/or elevation, the C&I algorithms do not depend on absolute amplitude values, but instead utilize magnitude ratios during the centroiding process.

2.1.2 RFI/Saturation Processing

Range gates in a single CPI that contain one or two pulses with significantly more power than expected, given the power in the other pulses and the azimuth beam pattern, are most likely the result of

RFI interference from another radar. If any range gate within a single CPI exhibits these characteristics, the entire CPI is flagged to indicate the RFI condition. Individual range cells where receiver saturation occurred are also flagged at this stage.




Figure 3. ASR-9 STC attenuation vs. range (typical).

2.1.3 Doppler Filter Bank

Following RFI/saturation processing, the A/D samples are input to a Doppler filter bank which produces 10 outputs for the high PRF CPI and eight outputs for the low PRF CPI. The filter bank bandpass characteristics are shown in Figure 4 (side lobes omitted)[3]. The filter outputs are referred to throughout C&I by the numbers shown in the figure, ranging from (-4) to (+4) for the high PRF and (-3) to (+3) for the low PRF. Filter numbers (+0,-0) are also referred to as the 'Zero-Velocity Filters' (ZVF), while all others are also referred to as the 'Non-Zero-Velocity Filters' (NZVF). The filtered I and Q time series data is converted to a power estimate (dB), and passed on to the CFAR function.




Figure 4. Doppler filter bank (high PRF).

2.1.4 CFAR

The CFAR acts to keep the number of radar primitives output to the ASP/9-PAC to a manageable level. The CFAR logic is different for the ZVF and NZVF primitives due to the fundamentally different phenomena that produce the ZVF and NZVF returns.

ZVF CFAR. The majority of ZVF primitives are the result of stationary ground clutter and can be effectively filtered using a dynamic clutter map containing a smoothed history of the recent amplitudes at each individual range, azimuth cell. The ASR-9 uses the following smoothing function at each clutter map cell:

$$\text{SmoothedValue} = 7/8 * \text{PrevValue} + 1/8 * \text{NewValue}$$

The smoothed value is summed with a Variable Site Parameter (VSP) (typically set to approximately 12 dB) to produce the final ZVF+ or ZVF- threshold. Because nearby ground clutter can be strong enough to cause residual power in the NZVF filters, a 'residue' map is maintained and a threshold generated using the same logic as above for all NZVF filters within the first 2 nm. Depending on the environment at a specific site, a VSP can control whether the residue map threshold or the NZVF CFAR threshold is used for NZVF returns within 2 nm.

NZVF CFAR. NZVF returns, by their very nature, change locations from scan to scan and cannot be thresholded utilizing a multiple scan history at a given cell. Instead, a sliding range window technique is used whereby the threshold for a given cell is determined by examining the 14 range gate window ending with the cell of interest (leading window) and the 14 range gate window starting with cell of interest (trailing window). Independent thresholds are determined for the leading and trailing windows, with the final threshold being the greater of the two. For each window, the threshold is determined by the following procedure:

1. Window Editing.

Certain cells in the window are excluded from contributing to the threshold generation. The cell of interest and its neighboring cell are always excluded. Cells tagged with RFI or Saturation flags are excluded. The remaining cells are then scanned, and the cell with the maximum amplitude is also excluded, as are the two closest cells to the maximum amplitude cell. This prevents a second aircraft target in the CFAR range window from raising the threshold.

2. Threshold Generation.

The amplitudes from the remaining cells in the window are then averaged and scaled by a 'Desired False Alarm Rate' VSP value to produce an appropriate threshold for the range window. Separate thresholds are maintained for all filter outputs.

At system start-up, the clutter map history is not present, and large numbers of ZVF targets pass through the thresholding step and are sent to the ASP. To prevent data overruns, the front end limits the number of ZVF primitives to 50 (nominal VSP setting) per CPI. If the count exceeds this number, then no more ZVF primitives are output for the CPI in question, and a flag set in the *next* CPIP azimuth header indicates that a ZVF overflow occurred.

2.1.5 Geocensoring

The unmodified ASR9/ASP also performed a front-end geocensoring operation which removed or flagged primitives occurring over roadways. This function has been moved to the 9-PAC in the new configuration, but a brief description of the original geocensoring logic is historically useful, as some of the features have been included in the new design. During site optimization, a 'Geo Map' was built utilizing a time-space history of radar-only targets and downloaded to EEPROM in the signal processor. Once installed, incoming primitives at a mapped location that fell below a selectable VSP threshold were rejected while those that were above the threshold were simply flagged. Each geocell could be specified as

'shaped' or 'flat,' with a set of five individual thresholds for the shaped category (for the +/-0, +/-1, +/- 2, +/-3, and +/- 4 filter 'classes') and a single threshold for the flat category that was used for all filters. The 'shaped' designation was originally intended for regions of strong ground clutter where the amplitudes (and the corresponding thresholds) would typically be higher in the low velocity filters. The 'flat' designation was intended for road traffic cells where target amplitudes are more evenly distributed across all the Doppler filters. During the early FAA ASR-9 evaluation period, these intended meanings were superseded, and the designations "flat" and "shaped" are now used to represent geo cells at ranges less than 3nm and ranges greater than 3 nm, respectively. The baseline flat threshold is set to 42.8 dB; geocensored primitives that are within 3 nm and below 42.8 dB (typically clutter breakthrough, not automobile traffic) will be rejected while all others will be flagged. The five shaped thresholds are normally set to zero, resulting in all primitives at ranges > 3 nm coinciding with a shaped geocell being simply flagged.

2.1.6 Two-Level Weather

The ASR-9 Target Channel is required to provide a backup weather detection capability in the event that the dedicated (non-redundant) weather channel should malfunction. Instead of the six-level representation provided by the weather channel, the target channel produces a simpler two-level output, where the two levels are user selectable (usually two and four). The levels are the standard NWS levels, with the following level-to-dBZ correspondence:

<u>Level</u>	<u>dBZ</u>	<u>Level</u>	<u>dBZ</u>
1	>18	4	>46
2	>30	5	>51
3	>41	6	>57

The two-level weather detection algorithm compares the average signal level (summed across all Doppler filters) at each range gate/CPI to predetermined weather data thresholds. The outputs from the thresholding step are then integrated spatially to form the tentative 0.5-resolution output detections for each CPI. The spatial integration logic declares a tentative detection at each 0.5 nm (eight range gate) interval if at least eight of the previous 16 range gates crossed the weather threshold. To minimize the effects of second trip weather, tentative detections are required to be present in both the high PRF CPI and an adjacent low PRF CPI before being output as a single valid detection for the CPI pair. To prevent cells containing ground clutter from producing false weather detections, a clear day clutter map is created at site optimization time, and ZVF returns from flagged cells are excluded from the average signal level calculation. To provide for the two weather levels, a separate set of thresholds is used on alternate scans.

2.2 9-PAC BACKGROUND

The 9-PAC replaces one of the ASR-9 dual-port memory boards with a card that contains, in addition to the original 64K of dual-port memory, three TMS320C40s, 35MBytes of memory, a 20 MByte Flash Memory Card (in a PCMCIA slot), 4 MBytes of on-board flash memory, and 4 ASYNC/SYNC serial ports. In its Phase 2 configuration, the 9-PAC completely replaces the ASP board set; in fact, the ASP boards must be removed to avoid bus conflicts. The dual-ported memory, originally used to transfer the radar/beacon primitive data between the High-Speed Interface Buffer (HSIB) and the ASP, now provides the equivalent data path between the HSIB and the 9-PAC. The 9-PAC also has access to the ASR-9's second dual-port memory board via the ASR-9 backplane, allowing the 9-PAC to communicate with the Message Interface Processor (MIP) in place of the ASP. A block diagram of the 9-PAC, including the primary tasks that run on each processor, is shown in Figure 5.





Figure 5. 9-PAC block diagram.

As shown in the diagram, all three C40 processors have 1 MByte of zero wait-state static RAM. Processors #1 and #3 have 8 MBytes of single wait-state dynamic RAM, while processor #2 has 16 Mbytes (providing for the large C&I adaptive thresholding maps). Processor #1 is used to communicate with all the peripherals (serial ports, flash memory card), as well as the on-board and external dual-port RAM. All three C40s are connected together via their high-speed (20 MBytes/sec) communications ports. The 9-PAC has no global memory available for interprocessor communication, so all communication is accomplished via the high-speed ports.

Although the C&I code is written in a portable fashion (and in fact runs on UNIX platforms as well as the 9-PAC), the 9-PAC hardware influenced the way in which some features were implemented. Of particular note is the fact that the smallest addressable word on the C40 processor is 32-bits ('char,' 'short,' and 'int' data types are all 32-bits), requiring that large data 'maps' wishing to conserve memory by using smaller data types be manually packed into the larger 32-bit words. Also, because a 'char' is the same size as an 'int,' some of the C&I structures use 'ints' where only eight bits of precision is necessary. The data structures of this type are relatively small in number and the memory wasted has not been significant.




3. C&I REQUIREMENTS

The C&I implementation is required to handle a maximum of 700 aircraft targets plus 300 non-aircraft targets per scan (max. of 31000 primitives), with the following additional peak loading characteristics:

- A peak of 250 total targets (max. of 11000 primitives) uniformly distributed across eight contiguous 11.25 degree sectors (90 degrees of antenna scan).
- A peak of 100 total targets (max. of 4400 primitives) uniformly distributed across two contiguous 11.25 degree sectors.
- A short-term peak of 16 targets (max. of 1200 primitives) in a 1.4 degree azimuth wedge, lasting for not more than two contiguous wedges.

The ASR-9 specification requires that the maximum C&I boresight delay is 0.14 seconds (109 ACPs at the slowest antenna rotation rate), where the delay is defined as the difference in azimuth between a given target's azimuth centroid and the current antenna boresight position at the time of actual output to the 9-PAC or Mode-S Merge process.

Where target loadings exceed the peak loadings stated here and the allowable boresight delay is exceeded, the C&I processing is required to reduce the processing range starting from the outer range limit until the delay returns to an acceptable level.



4. 9-PAC C&I ALGORITHM DESCRIPTION

The C&I algorithms are responsible for correlating the raw radar primitives into groups and interpolating between the primitives in each group to produce target centroid, amplitude, and Doppler estimates. A block diagram of C&I is shown in Figure 6. Each of the separate functions is explained in detail below.

Radar primitives from the ASR-9 are passed to the input parsing module which validates the data, flags primitives as geocensored if necessary, and groups the primitives into contiguous range groups for subsequent processing. The range groups for each CPIP are checked for a saturation or ZVF overload condition, followed by a correlation of the radar primitives across multiple CPIPS to form a radar target. Targets are then subjected to a filtering step that removes highly probable RFI targets. Surviving targets are scanned to locate the highest quality radar primitives, and interpolated values of the range and azimuth centroids, as well as the target's doppler velocity, are produced. A supplemental RFI test is used to flag (not delete) targets that are most likely due to RFI and were not detected by the primary RFI test. Following a reformatting operation, the targets are subjected to a geocensoring and multi-grid adaptive thresholding process to flag/remove false detections due to roads, birds, and weather. Prior to final output to the Merge process, targets are checked for excessive boresite delays (due to heavy CPU loading). Excessive delays cause the C&I processing range to be reduced until the delay drops back to an acceptable level.

To avoid including aircraft in the geocensoring/multi-grid map statistics (and possibly raising the thresholds unnecessarily), the thresholding counting process is not performed on the report stream flowing through C&I but is deferred until after the merge to allow radar reports associated with a beacon return to be excluded from the count. In the 9-PAC implementation, the counting/update process is run as a separate task that taps into the merge output stream to facilitate this.




Figure 6. C&I processing block diagram.

4.1 INITIALIZATION/RESET

Before any processing takes place, C&I must be initialized. Initialization consists primarily of allocating memory for all the C&I data structures, which is done up front to avoid memory fragmentation problems.

The processing range is set to the full range of the radar, 60 nm. The geocensor map, if present, is loaded from either the flash card (9-PAC) or disk file (UNIX). The fine and coarse adaptive map thresholds are zeroed, as are their target count arrays. A set of default VSPs is loaded, which is mainly of use when executing the code off line (under UNIX). The 9-PAC startup code always waits to receive C&I VSPs to avoid any possible confusion. The final initialization step is to call the C&I reset routine.

A reset of C&I resets the input state machine, flushes the input buffer, and clears the active and mature target lists. Current processing range, geocensoring thresholds, and adaptive thresholds are not affected, to prevent loss of information which must be preserved across input-error-induced resets (see below).

4.2 INPUT PARSING AND GEO/MTI FLAGGING

The input processing function serves three purposes: synchronization/validation of the input stream, a first-stage grouping of contiguous range cells, and the flagging of radar primitives that are likely due to road traffic and MTI (parrot) reflectors.

4.2.1 Input Parsing

The input stream consists of azimuth headers/data, range headers/data, CPI headers/data, and two-level weather headers/data. Data is not present for all range gates - only those containing at least one radar primitive that exceeded the front-end CFAR thresholds. Likewise, at each range gate only the filter amplitudes that exceeded their corresponding CFAR thresholds are present. Each type of header (az, range, CPI, WX) has a unique three-bit code, allowing for a simple state-machine to be implemented to verify the incoming data. A detailed description of the input data format is provided in Appendix B.

As the input data are parsed, range cells that are contiguous are grouped together. The maximum number of range cells in a group is nine (sufficient range extent to deal with the two-target case) after which the current group is terminated and a new group started. At the end of each CPIP (signalled by the arrival of the next CPIP's azimuth header), the groups are passed to the subsequent target formation routines.

The input state machine contains a number of checks to ensure data validity. In general, an unexpected sequence of data resets the state machine to wait for the next azimuth header, and the remainder of the current CPIP is discarded. Azimuth values of neighboring CPIPs are also tested. A jump of more than 32 ACPs is considered an error and the CPIP is discarded. If three consecutive azimuth errors occur, an alarm is set and the C&I reset routine is executed.

4.2.2 Geo/MTI Flagging

The 9-PAC's geocensoring strategy relies primarily on target centroids, both to produce the geomap and to perform the actual censoring operation, in order to keep the geomapped regions as confined as possible (See the geocensoring section for more details). Although the geomap is generated and used at a later stage in C&I, it is also available at the front-end of C&I and can be used to flag data at the primitive level.

This provides useful information to the downstream centroiding algorithms, particularly the beamshape match algorithm, which makes use of the information to prevent undesirable target splits over roads.

Two types of geocensoring are provided for in the 9-PAC implementation: the new adaptive version, and a limited form of the original ASR-9 shaped/flat geocensoring, providing some level of backward compatibility to address certain pathological cases. It is anticipated that the new algorithm will perform sufficiently well to negate the need for the old, in which case the old will be disabled.

If a given cell is tagged as a new-style adaptive geocell and primitives fall in that cell with amplitudes less than the geocensoring threshold, then the cell data (for a single PRF) is flagged as geocensored. Radar primitives are never removed by this scheme.

If a given cell is tagged as an old-style fixed threshold geocell, primitives with amplitudes falling below the fixed threshold are deleted, and primitives with amplitudes exceeding the threshold are flagged. Unlike the original geomap which distinguished between 'shaped' and 'flag' categories, the new map supports two identical fixed threshold sets, with each set consisting of five threshold values, corresponding to filters -0/+0, -1/+1, -2/+2, -3/+3, -4/+4. The thresholds are configurable via VSPs.

Lastly, a single bit in the geocensoring cell is used to identify radar primitives from MTI reflectors and flag them appropriately.

4.3 SATURATION AND ZVF OVERFLOW PROCESSING

When receiver saturation occurs, it is undesirable to permit radar primitives in the vicinity of the saturated cell from initiating a new target, as an inaccurate range centroid may result. To avoid this, all radar primitives within -2,+4 range counts of a saturated cell on the same CPI are flagged to prevent target initiation.

The number of ZVF primitives is usually maintained at a small number (a few hundred per scan at most), but when the ZVF CFAR clutter map is not fully initialized, such as can happen when the radar is just turned on, then large numbers of ZVF primitives can reach the 9-PAC. The front-end signal processor detects an overflow when the number of ZVF primitives in any given CPI exceeds 50 and removes subsequent detections for the remainder of the CPI, but the ZVF primitives prior to the overflow being detected are output to the 9-PAC and are a potential source of false alarms. The 9-PAC, like the ASP, discards these excess ZVF primitives, utilizing the ZVF overflow bits in the following CPIP's azimuth header (which indicate an overflow condition existed on the *previous* CPIP). Because the 9-PAC does not begin processing the current CPI until the header for following CPIP is detected, this is simple to implement. Note that the removal of ZVF primitives can result in the removal of range cells, and in some cases entire range groups, if there are no NZVF detections in the range cell or group.

4.4 CORRELATION OF RADAR PRIMITIVES

The C&I correlation process groups the radar primitives into targets. As successive CPIPs of data are processed, radar primitives that correlate with an active target at the same range are incorporated into the existing target. Primitives that fail to correlate with an active target initiate a new target. Those primitives active targets that fail to get updated with additional primitives over the course of a CPIP, or reach seven CPIPs in azimuth extent, are declared to be a "mature" target group and are forwarded to the next step in the processing chain (RFI Processing).


4.4.1 Target Initiation

A new target is initiated whenever a series of one or more primitives is encountered that fails to correlate with an existing target. The first three cells in the range group are searched to find the maximum amplitude. Cells that are inhibited from initiating a target due to nearby saturation are excluded from the maximum amplitude search. The maximum amplitude cell becomes the initial range centroid (R_c) of the target. The primitives on adjacent range cells are examined (if existing), and the cell with the maximum amplitude declared the adjacent cell, R_{adj} . As new CPIPs of data are integrated into the target, detailed data is maintained for only these two key cells. Note that if no adjacent data exists at the CPI responsible for target initiation, then at each subsequent CPIP an attempt is made to establish an adjacent cell. Normally an adjacent cell will be present in the first or second CPIP, although for weak single range gate targets it is possible to complete the target without an adjacent cell. Following the determination of R_c , the target update routines are executed as for subsequent CPIPs to add the appropriate data to the target data structure.

4.4.2 Target Updating

As the primitives from successive CPIs are grouped into radar targets, tests are performed to detect targets that are closely spaced in range and are producing overlapping range groups. In either case, primitives from R_{c-1} to R_{c+1} are associated with the target centered at R_c and are used to update the active target structure. In the single-target case, primitives at R_{c-2} and R_{c+2} are also grouped into the target at R_c , although they are not used to update the target data fields. If a primitive at R_{c+3} exists, it is not grouped with the target at R_c , but it is inhibited from initiating a new target. If there are any additional contiguous primitives in the raw range group and a new target is initiated, then the primitive at R_{c+3} will be included in the second target. An example of a single-target primitive range group is shown in Figure 7a. Note that the majority of aircraft returns do not span the range extent shown in the example but are typically only three or four range cells in length.


a) Single-Target Example



- $A_c/A_{c-2} \geq \text{Predicted}$ - No leading second target

- $A_c/A_{c+2} \geq \text{Predicted}$ - No trailing second target

b) Two-Target Example



- $A_c/A_{c-2} \geq \text{Predicted}$ - No leading second target

- $A_c/A_{c+2} < \text{Predicted}$ - Trailing second target declared.

Figure 7. Detection of multiple targets in a primitive range group.

When two aircraft are in close proximity, the amplitude vs. range relationship will not exhibit the same rise/fall characteristics as a single target but will instead appear more like Figure 7b. The presence of multiple targets can be detected by examining the amplitudes surrounding the cell at R_C . If data exist at R_C , R_{C-1} , and R_{C-2} (leading range split test) or R_C , R_{C+1} and R_{C+2} (trailing test), and none of the data is flagged as saturated or geocensored, then the amplitude ratio A_C/A_{C-2} or A_C/A_{C+2} is tested to see if it falls within the limits expected for a single target. If not, a target split is declared and the data at R_{C-2} or R_{C+2}, R_{C+3} is not associated with the current target but is allowed to associate with other targets or initiate a new target. Multiple target splits are disallowed. If a split is declared at the leading edge of the group (R_{C-2}), then the test for a split at R_{C+2} is not performed. In addition, once a target assimilates a batch of data that trigger a target split, subsequent target splits are inhibited over the entire lifetime of the active target. The specific tests used for the range split determination, using log domain amplitude data (3/32 dB), are:

Leading range split: if $((A_C - A_{C-2}) < 200)$ Declare leading range split

Trailing range split: if $((A_C - A_{C+2}) < 117)$ Declare trailing range split

The threshold values are compile-time constants and cannot be changed via VSPs.

One special case exists to limit the interference effects of two closely spaced (in azimuth) targets that are separated by a sufficient amount in range to trigger the range split logic (St. Louis parallel approach problem). In this case, if a range split occurs but the target is already at least two CPIs in run length, then the existing active target is terminated without integrating any of the primitives on the current CPI and the primitives are used to generate a new, separate target. This new target is pre-initialized to a range split to prevent it from being split a second time.

During the association process, a certain amount of processing and data reduction takes place to reduce the overall storage requirements and to reduce the amount of computation that must be performed once the target group has been completed. As mentioned above, filter data from associated cells at ranges other than R_C or R_{adj} are discarded. The filter magnitudes at each CPIP are distilled into six categories, three at each PRF. The three categories are:

- Zero velocity data at R_C
- Non-zero velocity data at R_C
- Non-zero velocity data at R_{adj}

The maximum filter amplitude in each category (if existing) is selected and stored in the target data structure. All other filter magnitudes are discarded.

In addition to the combined filter magnitudes, an interpolated Doppler value is determined for the cell at R_C using the peak filter magnitude and the adjacent filter magnitude.

4.4.3 Active Target Termination

An active target is considered to be complete when any of the following occur:

1. No primitives update the target for two consecutive CPIs. In the simplest case, there is no new data for either CPI in a CPIP, and the active target can be terminated on that CPIP. The second

case is when there is a miss on the low PRF CPI in one CPIP followed by a miss on the high PRF CPI on the following CPIP (even though there is a hit on the low PRF CPI). This is handled as a split case, with the most recent CPIs low PRF primitives possibly initiating a new target.

2. A hit/miss/hit pattern exists for either PRF. Although the association algorithm allows an all-miss scenario for either PRF to accommodate blind speeds, a hit/miss/hit pattern on either PRF is not allowed and is treated as a two-target split. The current CPIP's data is not associated with the active target (which is terminated) and is permitted to associate with another target or initiate a new target. Note that for this scenario to be invoked, the target must already have at least two CIPs worth of data prior to the CPIP being added.
3. The target runlength exceeds seven CIPs.

Once an active target is declared finished, it is placed on the completed target list which is passed to the next stage of processing.

4.5 RFI PROCESSING

When RFI occurs in the same range cell as ground clutter of comparable magnitude, the front end pulse-to-pulse RFI filter is rendered ineffective. After the Finite Impulse Response (FIR) filtering and CFAR operations, the clutter will be eliminated but the broad spectrum RFI signal will typically result in the generation of multiple NZVF primitives. Two mechanisms exist in C&I to reduce the effects of the RFI 'breakthrough'. The first, termed 'Primary RFI Processing,' removes targets that are only a single CPI (one PRF) in azimuth extent and contain more than RFI_HI_THR (5) or RFI_LO_THR (5) NZVF primitives. Analysis has shown that targets with these characteristics have a very high probability of being RFI false alarms. Primary RFI processing occurs immediately following radar primitive association to prevent false alarms of this type from reaching the centroiding process (reduces CPU utilization).

In environments with substantial amounts of RFI, some false targets will still leak through the front-end test and the C&I test described above. A second test, 'Supplemental RFI Processing,' counts the number of single-CPI targets in a given five-degree wedge, and if there are more than SUPP_RFI_SINGLE_CPI_THR (5), flags them as RFI targets. No targets are deleted by this test but all single CPI targets are delayed by up to five degrees to accomplish the counting process. Later on in the tracker process, targets flagged as confidence 2 (RFI) are not allowed to initiate a new track, only update an existing track. Note that this test utilizes target azimuth centroids, necessitating that it be performed following the centroiding function.

4.6 INTERPOLATION

Interpolation techniques are used to produce the final range and azimuth centroids as well as high and low PRF Doppler estimates.

4.6.1 Range Interpolation

A target's initial range centroid is determined by picking the range cell with the greatest amplitude and is therefore accurate to within 1/16 nm, the range gate size. This estimate can be improved upon by comparing the amplitude at R_c with the amplitude at the adjacent cell (the next largest amplitude by definition). If the two amplitudes are sufficiently close, then a range 'straddle' is declared and the target range is corrected by $\pm 1/32$ nm. (This is the final range accuracy of C&I even though the range is output in units of 1/64 nm).

The range straddle amplitude comparison is performed on the CPIP at which the adjacent cell is established. Data from both PRFs are checked - a range straddle condition can be declared if either PRFs data indicates it. At each PRF, the peak filter amplitude is selected for both R_C and R_{adj} . If there are no data at R_C on the current CPI, then a range straddle is immediately declared. If R_C data are present and $(Amp_C - Amp_{adj}) < 49$, a straddle is declared (amplitudes and threshold in units of 3/32 dB). The threshold of 49 is a compile-time constant and is the same value used in the original ASP implementation.

In addition to the range straddle correction, a constant range bias of 1/32 nm is added to compensate for the ASR-9 time-to-first range gate.

4.6.2 Azimuth Interpolation

The azimuth interpolation process is more complex than the range interpolation in part because the azimuthal resolution of the radar has been somewhat reduced by the front end FIR filter (pulse integration), and additional computation is required to recover the lost accuracy where possible. First, a search is performed across all CPIs in the target to find the highest quantity/quality data set that is available, and then one of five algorithms is chosen for azimuth determination.

Azimuth Centroiding Data/Algorithm Selection. A scoring procedure is used to select the highest quality set of data from among the six different types of filter data stored in a target report. (Three at each PRF: ZVF_RC, NZVF_RC, and NZVF_ADJ)

The following general rules are used to select the optimal set of data.

- NZVF data is preferred over ZVF data.
- Longer runlengths are preferred.
- The use of data from the same PRF is preferred over data from different PRFs.
- R_C data is preferred over adjacent cell data.
- Data with the beamswitch or saturation flag always scores lower than data without, regardless of runlength.

The meaning of the score values as implemented in the C code is as follows:

0	No hits for the data type	4	Runlength = 1 CPI
1	Long runlength (≥ 7 CPIs)	5	Runlength = 2 CPIs
2	Beamswitch condition	6	Runlength = 3 CPIs (and less than 7)
3	Saturation condition		

One special case exists - if there is any NZVF data type with a runlength of two or greater, then the ZVF data is not used, even when it has a longer runlength than the NZVF data.

Once the scores have been established for each data type they are compared, and a six-bit mask is produced containing a 1 for every data type that ties the maximum score. The mask value is used in combination with the score value as an index into a two-dimensional lookup table that selects the centroiding algorithm and input data set based on the general rules above.

There are five fundamental algorithms used for azimuth centroiding. They are:

1. Single CPI
2. Two-PRF Interpolation (two CPIs, different PRFs)
3. Single-PRF Interpolation (two CPIs, same PRF)
4. Beamshape Match (\geq three CPIs, same PRF)
5. Beamsplit (\geq seven CPIs - long runlength)
- 6.

The individual algorithms are described in detail in the following sections.

Single CPI. This algorithm is used when data is available at only a single CPI. This is the trivial case since no interpolation is needed. The target azimuth is simply set to the azimuth of the single CPI.

Two-PRF Interpolation. In this case, data is available from two adjacent PRFs, i.e., a high/low or low/high sequence. This is not the optimal case because of the filter magnitude fluctuations that can occur due to the differing PRFs, so azimuth beamshape information is not utilized. A more complicated algorithm is probably not warranted in any case because the azimuth correction is limited by the azimuth span of the data to approximately eight ACPs. A simple center of mass algorithm is employed:

$$\text{Centroid} = ((\theta_1 A_1 + \theta_2 A_2) / (A_1 + A_2))$$

This can be rewritten as:

$$\text{Centroid} = \theta_1 + K(\theta_2 - \theta_1)$$

where:

$$K = A_2 / (A_1 + A_2)$$

This algorithm assumes that A_1 and A_2 are linear target amplitudes that are consistent from one PRF to another. Prior to the centroiding computation, the amplitudes are converted to linear units and a 1 dB correction is added to the low PRF magnitude to compensate for high-low filter gain difference.

To protect against antenna north crossings, the more complete equation is:

$$\text{Centroid} = \text{Modulo4096}(\theta_1 + K(\text{Modulo4096}(\theta_2 - \theta_1))) \quad (1)$$

Single PRF Interpolation. The algorithm is used when two CPIs of data at the same PRF are available. Because the azimuth gap between successive CPIs at the same PRF is fixed at 16 ACPs and the antenna's azimuth beam pattern is known, the ratio of the amplitudes on the two CPIs can be used in conjunction with the beamshape to generate an interpolated azimuth. The equation used in this case is:

$$\text{Centroid} = \theta_1 + \frac{1}{2}(\theta_2 - \theta_1) + K_{\text{Beam}}(A_1 - A_2)$$

where K_{Beam} is a correction constant determined separately for each beam, and A_1, A_2 are log amplitudes. This process is illustrated in Figure 8. The true azimuth centroid for a two-CPI target with amplitudes proportional to A_1 and A_2 is clearly not at the midpoint between the two CPIs but lies closer to A_2 by the angular quantity that provides the best fit between the actual amplitude difference and that predicted by the beamshape pattern (four ACPs in this case).




Figure 8. Antenna azimuth beam pattern (low beam).

The predicted amplitude difference vs. azimuth offset is highly linear in the +/- 16 region of interest, and the constant K_{Beam} can be easily obtained. The relationship for the low beam is shown in Figure 9. A least squares fit of the data results in a K_{Beam} value of -0.307 for the low beam and 0.388 for the high beam.

To again account for the north crossing, the final form of the equation is:

$$Centroid = Modulo4096 \left(Modulo4096 \left(\theta_1 + \frac{1}{2} (\theta_2 - \theta_1) \right) + K_{Beam} (A_1 - A_2) \right) \quad (2)$$




Figure 9. Azimuth correction vs. CPI amplitude difference.

Beamshape Match. When three or more CPIs of data are available, the beamshape match algorithm is used. This algorithm computes the error between the known beamshape and the target amplitudes for a range of possible target azimuths and selects the azimuth with the minimum error as the target azimuth. If the minimum error exceeds a threshold, it is assumed that multiple targets are present and a target split is performed.

Only three data points are used by the match processing. If there are more than three CPIs of data available, the most 'central' set of data is selected for use. The three possible cases are:

1. Four CPIs a, b, c, d
If $a > d$, (a,b,c) are used, otherwise (b,d,c) are used.
2. Five CPIs a, b, c, d, e
(b,c,d) are always used.
3. Six CPIs a, b, c, d, e, f
If $b > e$ (b,c,d) are used, otherwise (c,d,e) are used.

The following discussion uses A, B, and C to refer to the central data set, with A representing the counterclockwise 'edge'.

The beamshape match process makes the assumption that the true target centroid lies within ± 8 ACPs of the azimuth of the center amplitude (θ_B). This requires that the antenna beamshape pattern for a ± 24 ACP window be stored in a table and used in the comparison process. (If the true centroid azimuth lies at the maximum bound of $\theta_B + 8$ ACPs, then θ_A is located 24 ACPs earlier than the centroid due to the 16 ACP spacing of successive CPIs at the same PRF). To test all possible azimuths in the ± 8 ACP window, 17 iterations are required. The error for each iteration is determined by the following equation:

$$Err = \left(\left(\frac{M_A}{M_B} \right) - \left(\frac{P_A}{P_B} \right) \right)^2 + \left(\left(\frac{M_C}{M_B} \right) - \left(\frac{P_C}{P_B} \right) \right)^2 \quad (3)$$

where M_A, M_B, M_C are the voltage equivalents of the measured target amplitudes, and P_A, P_B, P_C are the predicted voltage values based on the stored antenna pattern. The error computation is performed in the voltage domain because the maximum error between the actual antenna pattern (whose azimuth beam-shape varies as a function of frequency and elevation angle) and the ideal stored pattern is more constant across the azimuth region of interest when voltage values are used in place of power [4].

Once the minimum error has been determined, it is checked to determine if it is within an acceptable range for a single target. If so, the azimuth is set to the azimuth that produced the minimum error value. If not, the target is split into two targets if it has not already been split earlier (a range split) or flagged as geocensored. The method used for target splitting depends on the target runlength. If the runlength is four or five, then the single PRF interpolation is used on the first two hits and the last two hits to produce the azimuths for the split targets. If the runlength is three or six, the azimuths for the two targets are simply set to the leading edge azimuth plus 1/3 or 2/3 of the target runlength. Note that the new target created by the split process is simply a copy of the old target, sharing the same hit history, runlength, etc. No attempt is made to go back and intelligently split any of the fields (the hit history, for example) between the two targets. Most fields that would be candidates for splitting are never used following centroiding, so it is unnecessary.

The 9-PAC implementation of the beamshape match is substantially simpler than the ASP version, which utilized multiple layers of lookup tables and approximations to compensate for lack of floating-point capabilities. The original ASP 3/32 dB beamsplit error thresholds, along with the linear equivalents used in the 9-PAC, are:

Low Beam: ASP value (3/32 dB): 175	9-PAC linear equivalent: 43.71
High Beam: ASP value (3/32 dB): 150	9-PAC linear equivalent: 25.48

If the beamshape match process indicates a two-target case and some of the target primitives are flagged as geocensored, no split is performed (to avoid excessive numbers of split targets over roads), and the azimuth centroid is simply set to the azimuth of the largest amplitude in the central data set. This prevents the undesirable case of a target azimuth somewhere in between the true azimuth of the two collocated targets. If one of the targets is an aircraft and it produces higher amplitude returns than the collocated road traffic (the normal case for large jet aircraft), this approach results in the most consistently reasonable azimuth.

If a target was already split at some point, a second split is disallowed, and the target azimuth is produced using the single-PRF interpolation algorithm using the central pair of data in the target (with amplitude resolving which two are considered the central pair for odd runlengths).

Beamsplit Algorithm. The beamsplit algorithm is used when the beamswitch or saturation flags are present or the target runlength reaches seven CPIPs. In these cases, the data quality is not high enough to utilize the antenna beamshape. The target azimuth is obtained by simply bisecting the leading edge and trailing edge azimuths, taking both low and high PRF data into account.

Algorithm ID Tagging. In order to facilitate subsequent data analysis, each target is tagged with an ID ranging from 0 -57 to identify the algorithm and data set used to calculate the target centroid. These IDs are documented in Appendix B.

Azimuth Correction for Sampling and Signal Propagation Errors. Before final output, two minor corrections are made to the target azimuth. The first correction is necessary to correct for the azimuth sampling error. Ideally, for the high PRF (10 pulses), the azimuth sample would represent the azimuth halfway between pulse five and pulse six. Instead, the ASR-9 samples at the beginning of pulse six of the ten-pulse sequence (and pulse five of the eight-pulse low PRF sequence), so a correction of 0.5 ACP is subtracted from the target azimuth to improve the estimate. Next, a correction is made to compensate for the round-trip signal propagation delay. The antenna scans at a rate of 13.0 RPM (actual timing of two ASR-9s), or 890.4 ACPs/sec. The time it takes for a signal to return from a target at the full 60 nm range is $(60\text{nm} \cdot 2) \cdot (1852\text{m/nm}) / (3.0\text{e}8\text{m/s}) = 7.41\text{e-}4$ seconds, or 0.66 ACPs. To correct for both conditions, the following equation is used:

$$Az = Az - 0.5 + (0.66) \left(\frac{R}{960} \right) \quad (4)$$

where Az is in ACPs and R is in 1/16 nm range gates. This differs slightly from the original ASP implementation, which approximated 0.66 as 0.75 and 960 as 1024 in the calculation due to lack of floating-point support.

4.6.3 Doppler Interpolation and Smoothing

Each target contains a Doppler estimate for each PRF. These estimates are produced using a two-tiered approach. As each CPI of data is incorporated into a target, a Doppler estimate is produced by interpolating between the maximum filter magnitude and the magnitude of the adjacent filter (if present). This operation is performed during the target correlation phase to preclude the need to store all the filter information for each primitive. The intermediate value for each CPI is stored in the target. When the target correlation is completed, the interpolated Doppler values at each CPI are averaged together to produce the final smoothed Doppler value for each PRF.

Doppler values for each PRF are stored as integers that range from 0-63, where 0-32 represent positive Doppler quantities from zero out to the Nyquist interval, and 33-63 represent negative Doppler quantities from the Nyquist interval back to zero (63 is the smallest negative Doppler value). This Doppler scale is used for backward compatibility, and is the same as used in the final C&I output reports.

Doppler Interpolation. At each CPI there may be multiple filter crossings, especially true when the true target velocity lies somewhere between two of the filters. In such a case there will typically be two 'adjacent' filter crossings, i.e., +2 and +3, and the amplitude ratio of the two can be used to determine an improved estimate of the target Doppler. This is very similar to the single PRF interpolation method used by the azimuth centroiding process. The interpolated Doppler value is given by:

$$Dop = Dop_{avg} + K_n (A_n - A_{(n-1)}) \quad (5)$$

K_n is a table of interpolation constants for each pair of adjacent filters (see Appendix C), A_n and $A_{(n-1)}$ are log magnitudes in 3/32 dB, and the Doppler values use the aforementioned 0-63 folded Nyquist scale. Note that if the amplitude difference is larger than the predicted maximum, this equation can result in a Doppler value that is outside the interval in question. To correct for this, the interpolated Doppler value is hard limited so that it never falls outside of the two-filter interval.

Doppler Smoothing. The interpolated Doppler data from each CPI of data in the target is averaged together to produce the final Doppler estimate for each PRF. To prevent outliers from being included in the average, a simple filtering step is performed that eliminates any Doppler values that are further than 12 Doppler 'counts' away from the Doppler value of the max filter (for each PRF) of the target.

4.7 TARGET REFORMATTING

Following the interpolation process, the radar targets are reformatted into the output report format. The format is nearly identical to the ASP output format, with the only difference being the addition of adaptive thresholding information in a formerly unused word. During the reformatting operation, the filter magnitudes for the high and low PRFs are normalized to account for the 10 vs. eight pulse integration in the front end as well as the small differences in attenuation of the various filters. The reformatted reports are then output to the new geocensoring/adaptive thresholding process, described in Section 5.

4.8 DELAY PROCESSING

Following the geocensoring/adaptive thresholding process and just prior to output to the Merge process, the targets are checked for boresight delay. As stated in Section 3, the allowable boresight delay in the non-Mode-S configuration is determined by the 9-PAC Merge window, which is set to a minimum of 176 ACPs. Subtracting 24 ACPs to allow for communications latency between C&I and Merge (a generous amount -- it will normally run from four to eight ACPs) results in an delay threshold of 152 ACPs. If more than four targets per scan exceed this threshold, a processing overload is assumed and range reduction occurs. Range reduction takes place in fixed size steps, which vary in size from 12 nm at full range to 1 nm at very short range. Table 2 shows the step sizes for all ranges. Assuming that the targets causing the delay are distributed fairly evenly over the current processing range interval, these step sizes result in a minimum load reduction of 20 percent at all ranges.

TABLE 2: Range Reduction Step Sizes

Range (nm)	Reduction Step Size (nm)
51-60	12
41-50	10
31-40	8
21-30	6
11-20	4
2-10	2

When the maximum boresight delay for all targets in a scan falls back below 136 ACPs, the processing range is allowed to recover back to the full 60 nm range at a rate of 2 nm/scan. The stricter require-

ment of 136 ACPs here instead of 152 prevents the processing range from continually 'hunting' when range reduction is in effect.

It should be noted that the capacity tests are designed to be sufficiently strenuous to ensure that range reduction never occurs under normal circumstances. It is most likely to be triggered by abnormal situations, such as inadvertent radar jamming or an STC malfunction in the ASR-9 front-end.

5. GEOCENSORING AND ADAPTIVE THRESHOLDING

5.1 OVERVIEW

All digital radar systems require some sort of post-signal-processing adaptive thresholding to eliminate false alarms which survive the thresholding process in the target extraction system. In MTD systems such as the ASR-9 these false alarms fall into two basic classes. The first class consists of false alarms which occupy a significant amount of area and usually do not have fixed geographic locations. These false targets include large concentrations of birds or "angels" as well as weather returns which break through the signal processor thresholds [5]. The breakthrough weather returns result from a reflectivity structure which is sharper than the range CFAR window after being processed through the Doppler filter bank, or zero velocity filter returns which are moving faster than the time CFAR algorithms used by the zero velocity filter thresholding can react. In some cases false ground clutter targets associated with anomalous propagation or "ducting" have sufficient spatial extent to fall into this class. Targets of this sort are best removed through the use of a time-area-Doppler amplitude CFAR system. The original MTD and the ASR-9 employ a relatively simple form of such a system. The multi-grid adaptive system which is described here is designed to provide improved rejection of targets of this sort by taking advantage of the additional computational and storage capability available in the 9-PAC system.

The second class of radar false alarms includes false targets caused by ground vehicles on visible sections of roads, large fixed clutter returns, and other features with localized geometry such as trains, windmills, etc. These targets differ from the first class in that they often have small spatial extents, and while they may exhibit a large amount of temporal variability, the geographic locations involved are fixed and cover a fairly small fraction of the radar surveillance space. In the 9-PAC system these targets are removed using an adaptive geomap system which has very high spatial resolution and uses very long observation times to identify regions where false alarms of this sort are likely to occur. The geocensoring process then closely monitors those locations to determine the appropriate threshold level.

One problem common to all time-space CFAR schemes is the conflicting requirements between the need for fast response vs. the desire to have as high a spatial resolution as possible. In general one finds that as the size (in range-az-Doppler space) of the cells decreases, a longer integration period is necessary to achieve a stable set of false alarm statistics. Conversely, larger cell sizes can result in desensitization of an unnecessarily large fraction of the radar coverage area. In the 9-PAC thresholding implementation, this issue is addressed by the use of multiple thresholding layers, one very high resolution geocensoring layer followed by two successively lower resolution, faster responding adaptive thresholding layers. This topology is illustrated in Figure 10.




Figure 10. Multigrid thresholding stages.

The geomap has a resolution of 1 CPI x 1 range gate x 1 Doppler class and a time constant on the order of days. The fine and coarse-grained adaptive maps utilize cartesian grids of 1/2 x 1/2nm, and 2 x 2 nm, and have time constants in the neighborhood of 10 and 2 minutes (VSPs), respectively. The adaptive maps maintain separate cell statistics for seven separate Doppler 'classes,' as targets filtered by these stages tend to exhibit more uniform Doppler characteristics than the road traffic returns handled by the geomap.

Although the geocensoring implementation differs somewhat from that of the adaptive map due to its much greater resolution and time constant, the overall thresholding approach can be thought of as a three-layered process, starting with the highest resolution/slowest-reacting layer and moving down through successively lower resolution/faster-reacting layers. Each layer consists of a counting phase, during which eligible targets are used to compute an amplitude histogram and threshold level for each cell, and a thresholding phase, where the targets are flagged/deleted if necessary. Targets that are thresholded by an upper layer are prevented from updating a lower layer's target counts. False alarms therefore impact only the thresholds in the highest resolution map that they 'trigger', minimizing the total area of reduced radar sensitivity. Note that because the higher resolution maps react more slowly, this behavior is not achieved instantaneously. When a set of false alarms with a small-scale spatial structure first appear, the coarser thresholds rise quickly to control the output false alarm rate, followed by a gradual increase in the finer cell thresholds and a subsequent decrease in the coarser thresholds as enough statistics build up to define the spatial structure.

Certain targets are excluded from the counting and/or thresholding processes, to prevent the threshold levels from being falsely raised, and to prevent targets that are almost certainly aircraft from being incorrectly flagged/deleted. Targets flagged as RTQC, MTI, or RFI are excluded from both the counting and thresholding phases. Targets that merge with a beacon target are excluded from the counting process. This is absolutely critical for correct operation of the geocensoring logic; otherwise, airport approaches would eventually become mapped as roads. (The counting phase is performed as a separate task following the Merge to allow the merge status of the target to be used -- see Figure 6). In addition, for the case of the fine and coarse adaptive thresholding, targets meeting a certain quality requirement, and targets whose velocity exceeds that expected of birds or weather, are excluded from the counting and thresholding phases. This is not a valid approach for the geocensoring logic because road traffic returns will often produce high-quality targets and target Doppler characteristics which are not predictable due to combinations of reflections from multiple vehicles. Additional details regarding the exclusion criteria are provided in the following sections and in flowcharts in Appendix A.

5.2 GEOCENSORING

The adaptive geomap processing is designed to actively detect and remove false radar targets that are closely localized in space and are produced by processes such as returns from automobile traffic and large specular clutter reflectors. The 9-PAC algorithm represents a significant change from the original ASR-9 implementation which used a static map calculated from observed target records and loaded by the person performing the initial site optimization. It is expected that the automation of this process will result in a significant reduction in the site setup and optimization time. It will also eliminate the need for human intervention when changes such as the construction of new roads occur.

The spatial resolution of the geothresholding process is the fundamental radar resolution, 1 CPI x 1 range gate, or twice the resolution of the original ASR-9 clutter map in both range and azimuth. This higher resolution allows for the mapping of a smaller percentage of the overall radar cells and should provide increased sensitivity to radar-only targets flying in close proximity to roads. The geomap count/update period is fixed at 512 scans (about 40 minutes), an interval long enough to gather meaningful statis-

tics for the small geocells as well as being a convenient number from an implementation standpoint (it matches the number of CPIs in the map, so 1 CPI can be updated per scan to distribute the processor load)

The geocensoring process is itself split into four phases: active cell determination, target counting and threshold determination, geomap enable/disable, and target censoring. Details of each are provided in the following sections.

5.2.1 Active Cell Determination and Survey Map Description

Since there are nearly 500,000 resolution cells it is not practical to maintain a complete geomap threshold record (about 100 bytes) for each cell. Instead, a 'survey map' is created which contains a single 32-bit word per cell. The survey map consists of raw target counts and is used to identify which cells have large enough target densities to be observed more closely. If a given range, azimuth cell contains more than GEO_ACTIVE_COUNT_THR (5) hits over the 512 scan update period, the cell is declared to be an 'active cell', and more detailed statistics are maintained. The ACTIVE_CELL bit is set in the survey map word to indicate the condition, and the original 'count' field in the survey map is changed to contain an offset pointing to the active cell data structure. When an active cell has no significant activity for an extended period of time, it is returned to inactive status, and the storage used is freed up for later use by another active cell. The current 9PAC implementation allows for a maximum of 25000 active cells. Analysis of data from Albuquerque, a fairly active road traffic site, showed that roughly 6000 active cells were being used when the map reached the steady state, an indication that 25000 is most likely sufficient for all ASR-9 sites (the same analysis has yet to be performed for other high-road traffic sites, however).

Four bits in the survey map are used to handle special situations. The AIRPORT_MASK bit signifies that the cell should never be considered active. This is used to prevent geocensoring over runways (where beacons are disabled), both at the local airport and satellite airports. At system start-up, a series of rectangular airport-mask regions are passed to C&I from the RMS. During C&I initialization, the AIRPORT_MASK bit is set for all geomap cells that fall within a specified airport-mask region.

The MTI_MASK bit is used to indicate regions where MTI reflectors are present. There are typically one or two MTI reflectors at each airport, situated near the end of a runway. Primitives occurring in a survey map cell whose MTI_MASK bit is set are flagged as MTI but are not deleted. MTI reflectors generate echoes in only a very small area, so in this case, polar wedges (azimuth extent, range extent) are passed to C&I from the RMS to delineate the region of interest.

At the FAA's request, logic to emulate the original SHAPED/FLAT thresholding capabilities of the ASR-9 (see section 2.1.5) has been included in the 9PAC. A two-bit field in the survey map word is used to indicate that this alternate thresholding logic be used. If the field is a 1, fixed threshold set #1 (consisting of five threshold values, one for each of the filter classes +/-0, +/-1, +/-2, +/-3, +/-4) is used to delete/flag the radar primitives. If the field is a 2, fixed threshold set #2 is used. A value of 3 is currently illegal. Like the MTI regions, the regions where the fixed filter approach will be used are expected to be quite small (to solve certain pathological cases), and wedges specified in polar coordinates are used to mark their range/azimuth extent. Up to four polar wedges can be specified as VSPs. The actual threshold values for sets #1 and #2 are also set via VSPs.

A copy of the current threshold in units of 3/32 dB is also included in the survey map word. This is done to allow efficient implementation of C&I's front end geoflagging operation. The final format of the survey map word is shown in Figure 11.




Figure 11. Survey map word format.

5.2.2 Geomap Target Counting and Threshold Determination .

When a cell is first flagged as active, the following data structure is allocated.

```

struct {
    int    range ;                /* Cell range (0-960 gates) */
    int    azimuth ;            /* Cell azimuth (0-512 CPI's) */
    int    threshold ;          /* In dB */
    int    countHistogram[10] ; /* Target count vs amplitude histogram */
    float  hitHistory[10] ;      /* Single-pole filtered hit history */
    float  hitHistoryNormalizationFactor ;
}

```

Since a primary mission of the geomapping stage is the reduction of automobile reports, it is necessary to maintain statistics for cells which may be observed only during peak traffic loads, twice per day, with potential periods of up to 72 hours between observations to account for weekend intervals. At the same time, transitory events that cause a large number of hits during two or three update intervals but are not present on multiple days should not result in geocensoring but should instead be handled by the adaptive thresholding fine and coarse maps. Both of these issues are addressed by equipping the geomap with a long time constant.

In order to determine the recent ground traffic activity for a particular cell, a count vs. amplitude histogram is created for each 512-scan update cycle. An example is shown in Figure 12. The histogram bins have a fixed width of 8 dB, starting at a conservative minimum detectable signal figure of 10 dB (For radars tested to date, the actual minimum detectable signal is closer to 20dB). For each amplitude bin, a 'hit' is declared for the current update cycle if the number of targets falling in the amplitude bin exceeds GEO_HIST_COUNT_THR (5).




Figure 12. Geomap amplitude histogram.

Because of the need for the geocensoring function to respond slowly to external conditions, a hit/miss on a single update cycle is not sufficient to change the threshold level. Instead, the hit/miss status is passed through a single-pole filter of the following form:

$$\text{if (hit declared)} \\ \text{hitHistory[bin]} = (\text{hitHistory[bin]} + 1) * (\alpha - 1)/\alpha \quad (6)$$

$$\text{else} \\ \text{hitHistory[bin]} = \text{hitHistory[bin]} * (\alpha - 1)/\alpha \quad (7)$$

where α is a filter constant chosen to provide a time constant on the order of days. This filter has the characteristic that the equilibrium value of $\text{hitHistory[bin]}/(\alpha-1)$ represents the fraction of update cycles that were hits. For example, given a α value of 300, a hitHistory equilibrium value of 150 indicates that there were hits on 50 percent of the update cycles. Only when the percentage of update cycle hits for a given amplitude bin exceeds GEO_HIT_THR (10 percent) is the cell's threshold raised. The final threshold for the cell is set to the ceiling value of the highest amplitude bin with a hitHistory > GEO_HIT_THR, plus an additional 'boost' factor of GEO_THR_BOOST(4 dB).

5.2.3 Determination of Geomap Single-Pole Filter Coefficient (α)

As stated above, the desired characteristics of geocensoring are to maintain a thresholding capability for up to 72 hours after false-alarm activity dies down, even for those regions that are normally only present during peak traffic periods. Analysis of Albuquerque data has shown that the great majority of road cells, even those that appear only during traffic peaks, will cause hits on at least 14 percent of the 36 daily update cycles; the value of 14 percent is therefore used as the "three-day minimum detectable signal." The second factor driving the choice for α is the desired bottom limit for thresholding, the hit rate below which the cell will no longer actively threshold targets. This was chosen as 10 percent. A value lower than 10 percent begins to run the risk of allowing events of a transitory nature to cause thresholding. The determination of α , then, can be restated as: What is the required value for α that allows a cell with the minimum detectable hit history (14 percent) to decay no further than the bottom threshold (10 percent) after 72 hours (108 update periods)?

Modifying Equation 7 to reflect 108 update cycles without a hit results in:

$$0.14 * ((\alpha - 1)/\alpha)^{108} = 0.10$$

Solving for α yields 321.

Figure 13 shows the persistence of geocells based on the equilibrium hit percentage and an α value of 321. As can be seen on the plot, the hit history for a geocell that is normally active 100 percent of the time (Hit History = 1.0) and then suddenly becomes completely inactive will be maintained above the 10 percent threshold (and actively censor targets) for 20 days. And as the choice for α dictates, a geocell that is present 14 percent of the time will persist for three days before dropping below the bottom threshold. Note that all three of the quantities used to determine α (14 percent, 10 percent, and 72 hours) are configurable via VSPs.




Figure 13. Hit history decay curve, $\alpha = 321$.

Left unmodified, the use of a single-pole filter results in a significant lead time before the hit history value has risen to the point where thresholding can begin. As shown in Figure 14, a cell with a hit rate of 20 percent will not begin to threshold until the six-day mark. Cells with higher hit percentages do mature faster, but still take longer than is desirable from the point-of-view of site personnel responsible for system optimization. To reduce this time lag as the hit_history is updated, a normalized version of the value is computed that better represents the hit ratio even early on in the cells history, and this version is used to make the thresholding determination. For example, a cell that had hits on 10 of the first 50 update cycles (1.4 days) would have a normalized value of hit_history that was very close to 20 percent, and would be actively thresholding at that time instead of needing the six days of history shown in the figure.




Figure 14. Unmodified single pole filter thresholding time lag.

The normalization process has one drawback in that transient events can be counted as significant normalized hit percentages very early in the accumulation process (one hit in the first update cycle is normalized to a hit rate of 100 percent). To circumvent this problem, cells are prevented from thresholding until they have reached `GEO_CHILDHOOD_THR`, a VSP currently set to 36 update cycles. By that time, a two-hit transient event only represents only a 5.6 percent hit rate, which is insufficient to trigger thresholding.

5.2.4 Geocensoring Enable Map

Areas of roads that are visible on a small fraction of the total days (~20 percent) can result in unnecessarily aggressive geocensoring. This is especially true in the case of roads which become visible only during (repeated) periods of anomalous propagation in which case a geocell can potentially be completely devoid of road traffic hits for many consecutive days. To reclaim some of this lost sensitivity to radar-only targets, a separate 'enable' map is used which allows the geocensoring to be enabled/disabled over a relatively short time scale without disturbing the normal geocensoring counting/thresholding process.

The enable map is a coarse-grained version of the normal geomap with a resolution of four range cells by four CPIs. Two values are stored at each location, a hit count and an 'age' value representing the number of scans since the hit count was last exceeded are stored. The enable map counting period is the same as for the thresholding geo map, 512 scans. When the hit count in an enable cell drops below `GEO_ENABLE_COUNT_THR` (7) for more than `GEO_ENABLE_AGE` (3) 40-minute update cycles, then the cell is declared disabled and fine-grained geocensoring is essentially turned off for all geo cells that lie within the coarse enable map cell. If road traffic reappears in the enable map cell, as soon as the hit count exceeds the threshold, the cell is re-enabled, not requiring the full 512-scans cycle to complete. (The multi-gridded thresholding described in the next section has a similar fast-update capability).

Care must be taken when setting the enable map count threshold. A value that is too high will disable a large number of cells during the overnight period, and the number of false alarms at the start of the morning rush hour period will be unacceptably high. Analysis of Albuquerque data has shown that a target density threshold of seven in combination with an enable age of three is a conservative enough setting to prevent excessive false alarms.

5.2.5 Geocensoring of Targets

The actual geocensoring of radar targets is straightforward when compared with the determination of the threshold level. First, each target is checked to see if it qualifies for exclusion from the geocensoring process because the RTQC/MTI or RFI flag is set. If not, and the target's amplitude is less than the corresponding geomap threshold and the geocell is enabled (see above), then the confidence field of the report is set to zero to indicate a geocensored target.

As described in section 4.1, the geomap information is also used to tag radar primitives as geocensored at the front-end of C&I. That information is used during the centroiding process to improve accuracy, but unlike the original ASR-9 implementation, the inclusion in a target of a single radar primitive flagged as geocensored is not sufficient to flag a target as confidence zero -- that is based solely on the geomap threshold for the cell that coincides with the target's centroid.

5.2.6 Geocensor Map Archival

In the 9PAC implementation, the geocensor map is archived to the flash card once/day to allow a valid map to be loaded at system start-up. Two files are maintained on the card and used in a ping-pong manner to ensure that in the event of an inadvertent system shutdown, at least one valid copy will be present when the system is restarted. A checksum mechanism is used to determine file integrity.

5.3 MULTI-GRID ADAPTIVE THRESHOLDING

In the original ASR-9 system the adaptive thresholding process is accomplished with a simple adaptive filter. For each target that appears in a thresholding cell (range-azimuth-Doppler) whose amplitude is greater than the current threshold value, the threshold for that cell is implemented by some fraction of a dB (a VSP value). Every scan, all of the cells which are above receiver noise are decremented by another logarithmic value (VSP). This approach has the advantages of being simple to implement and easy to understand but it also has several disadvantages. One of the disadvantages is that it is hard to increase the thresholds in a cell fast enough to control the false-alarm rate without making them also sensitive to improperly rising when a small concentration of real primary radar returns are localized in a cell. A second problem with the original scheme is that the decrement rate is fairly small and it takes a relatively long time for the threshold value to decay back to the noise level once the false alarms have left the cell. In essence, both of these limitations arise from the fact that the amplitude information in the target returns is not being fully utilized. Only one bit of information, whether the target amplitude is above or below the current threshold value, is extracted from each target report. In the 9-PAC approach, this situation is improved through the generation and use of a target amplitude histogram, similar to that used by the geocensoring process, for each thresholding cell.

As stated earlier, the new adaptive thresholding utilizes two layers, a fine-grained layer and a coarse-grained layer, to achieve the desired balance between spatial resolution and response time. The fine-grained layer has a cell size of $1/2 \times 1/2$ nm and a update period of 120 scans (VSP). The coarse-grained map has a cell size of 2×2 nm and an update period of 20 scans (VSP). At each map cell, an amplitude histogram and current threshold level are maintained for each of seven Doppler classes. The Doppler class

of any given target is determined by the filter number of the maximum amplitude return. The filter-to-class correspondence is shown in Table 3.

TABLE 3: Doppler Class / Filter Number Correspondence

Class	Filter(s)	Class	Filter(s)
0	-4,-3	4	+1
1	-2	5	+2
2	-1	6	+3,+4
3	-0,+0		

5.3.1 Adaptive Map Target Counting and Thresholding Determination.

As with the geomap, the first stage in each layer of the adaptive thresholding process is the target counting/ threshold-setting stage. During this stage eligible targets are used to create a target count vs. amplitude histogram and the cell thresholds are established using the histogram information. To prevent the histogram counts from being unnecessarily biased, certain targets are excluded from the counting process. The first reason for exclusion is a “technical exclusion.” These targets are false alarms which can be identified before adaptive thresholding and currently include RTQC/MTI targets, single CPI targets with flat spectra (RFI bit set), and targets which have been flagged as failing the geocensor processing. A second class of targets excluded from the counting phase are ‘star targets’ and ‘semi-star’ targets. The targets identified as star targets have a very high confidence that they are actual aircraft. In addition to not being counted, these targets also automatically pass the thresholding stage. The current implementation considers targets with ambiguous Doppler returns (higher velocities than birds or weather) as star targets in agreement with current ASR-9 practice. Targets that merge with a beacon target are also considered to be star targets. Semi-star targets are “probable aircraft” and include targets whose quality and confidence values exceed a VSP set threshold that is range dependent. Note that semi-star targets do not automatically pass the thresholding process and do not have to be guaranteed to be real aircraft, merely that they represent a small fraction of the false targets so that the histogram statistics are not compromised. The exclusion processing is summarized in Figure 15.




Figure 15. Target exclusion processing.

One major difference from the geocensoring algorithm arises from the fact that the false target regions are often moving, such as when a streak of bird targets or a weather system is traversing the coverage area. When the region of false targets moves into a previously 'clear' adaptive map cell, a significant number of false alarms can pass through before the threshold rises to the appropriate level. To deal with this problem, the 9-PAC adaptive thresholding logic uses two separate definitions of each cell. One of these is the "thresholding zone," which is the region to which a particular set of thresholds are applied. The second region is the counting or "influence zone," which is a larger area that includes the thresholding zone and a guard band which extends to an area one half of the cell size in each direction (see Figure 16). The influence zone is the region where targets are counted to compute the amplitude histograms for a particular cell. Because of this geometry, problems with moving false alarm regions are effectively eliminated since the threshold calculation logic "feels" the false alarms before they enter the thresholding zone. An effect of the geometry choice is that each target can influence the threshold values in more than one cell. For the current implementation, each target influences the threshold values for four adjacent cells due to the half cell-width guard bands.




Figure 16. Counting and thresholding cell geometry.

The amplitude histogram is used to directly establish the cell's threshold. Basically, the threshold is set to the ceiling value of the highest amplitude bin that exceeds the trigger target count. This threshold setting process is depicted in Figure 17. In this example, because the count of targets with amplitude ≥ 27 dB is greater than the trigger target count of 10, the threshold is set to 31 dB (the floor value of the next bin is equivalent to the ceiling value of the current bin). If the target count for the maximum amplitude bin exceeds the trigger value, the threshold is set to the maximum allowable threshold level for the map (VSP). Note that the amplitude histogram bin spacing can also be configured vis VSPs.

In order to allow the adaptive maps to react as quickly as possible, a "fast update" capability is provided to allow the threshold to be increased as soon as the target count for a bin exceeds the trigger count, without waiting for the end of the update cycle. In general, the higher the false-alarm load, the faster the thresholds will be raised, a feature providing similar functionality to the original ASR-9 adaptive thresholding scheme. A decrease of the threshold, however, is never considered until the end of the update cycle, when a complete set of statistics has been accumulated. If the false targets dissipate, this process still ensures that any desensitization will be removed following the first full update period, providing for a more rapid recovery than the original ASR-9 thresholding scheme.





Figure 17. Adaptive map histogram-based threshold setting technique.

5.3.2 Adaptive Map Thresholding

As was the case for geocensoring, targets flagged as RTQC/MTI are excluded from the adaptive thresholding process. Targets flagged as geocensored and targets with a high-quality Doppler estimate exceeding that expected for birds and weather are also excluded. In accordance with current ASR-9 practice, several VSPs are supplied to provide some flexibility as to what constitutes a 'high-quality' Doppler value.

Targets that do not meet any of the above criteria are separated by Doppler class based on the filter number of the target's maximum amplitude return. Targets whose amplitude falls below the current threshold for their Doppler class are flagged as thresholded. Targets thresholded by the fine map are excluded from consideration by the coarse map threshold (redundant).

In the current implementation, thresholded targets are not immediately deleted, as in the original ASR-9 implementation, but are simply flagged and passed on to Merge. This is done simply to provide a communications path to the adaptive thresholding counting process which is delayed until after the merge (the flagged targets are never considered for a merge with a beacon report). This method, as opposed to providing a separate feedback path to the update task for the flagged targets, simplifies the implementation and does not increase the CPU utilization by more than one or two percent. (If performance becomes an issue, a more efficient implementation will be pursued)



6. TWO-LEVEL WEATHER

Two-level weather data is output by the ASR-9 front end on an alternating-scan basis, with each scan containing either level 1 or level 2 weather detections. The resolution of the weather data is 0.5nm by 1.4 degrees. The ASR-9 specification requires that the incoming data be temporally and spatially smoothed, with a new map output to the SCIP every 30 seconds.


6.1 TEMPORAL SMOOTHING

Temporal smoothing is performed at each individual range/az cell. Because the required time between weather map outputs is 30 seconds, three scans of low-level and three scans of high-level weather are available for smoothing. The algorithm is simple: if the cell contains a detection on two out of three scans, then it is declared to be a detection. The output from the temporal smoothing operation is then passed to the contouring step.

6.2 CONTOURING

Contouring counts the number of weather detections in each 3x3 group of cells and declares a hit at the center cell if the total count exceeds WX_CONTOUR_THR (5). Cells at 0 range and full range, where only six neighboring cells are available for contouring, use $0.66 * WX_CONTOUR_THR$ (3) as a threshold. Following the contouring operation, the completed map is output to the MIP.

In order to prevent clutter from causing false WX detections, the ASR-9 front-end utilizes a clear day map generated by this process. In this case, the contouring algorithm recognizes the CLEAR_DAY_MAP mode bit in the weather data header, and contouring is not performed to allow an accurate representation of the ground clutter to be generated. Once completed, the clear day map is output via the normal mechanism to the ASR-9, where it is stored in non-volatile memory.



7. VARIABLE SITE PARAMETERS

C&I is configurable via a set of Variable Site Parameters (VSPs), that are input to the 9-PAC via the ASR-9 RMS. Note that in the current design, each time a VSP message is processed, C&I must reset in order to update internal lookup tables, thereby causing a brief interruption in the output data. This is comparable with the original design and is not a problem since the ASR-9 does not allow the VSPs to be updated while a channel is on-line.

The following description of the VSPs is broken down by function group. For each VSP, the nominal value is shown in parentheses.

7.1 RFI PROCESSING

RFI_ENABLE	Enable RFI filter crossing test. Default is Enabled (1)
RFI_HI_THR	High PRF filter crossing threshold. Single CPI (High PRF) targets containing RFI_HI_THR or more filter crossings are deleted. The default value is 5.
RFI_LO_THR	Low PRF filter crossing threshold. Single CPI (Low PRF) targets containing RFI_LO_THR or more filter crossings are deleted. The default value is 5.
SUPP_RFI_ENABLE	Enable supplemental RFI single CPI target density test. The default is enabled (1).
SUPP_RFI_SINGLE_CPI_THR	Single CPI target threshold. If the number of single CPI targets in a single 4 CPIP (5 deg.) wedge reaches or exceeds this number, all single CPI targets in the wedge are flagged (not deleted) as RFI targets. The default density value is 5.

7.2 INTERPOLATED DOPPLER PROCESSING

ID_TEST_QUAL	Interpolated Doppler test quality. To be considered for exclusion from adaptive thresholding based on interpolated Doppler, the report must meet or exceed this quality. The default value is 3.
ID_TEST_CONF	Interpolated Doppler test confidence. To be considered for exclusion from adaptive thresholding based on interpolated Doppler, the report confidence must equal ID_TEST_CONF. Two special values are used to provide variations. A value of -2 (the default) signifies that the interpolated Doppler test should always be performed regardless of the confidence. A value of -1 indicates that the exclusion test should be performed only if the confidence is either 3 or 5 (probable aircraft). The default value is -2
ID_TEST_DOP_THR	Doppler velocity, expressed in knots, above which a target is considered to be an aircraft, and is excluded from adaptive thresholding. The default value is 70 knots.

7.3 GEOCENSORING AND ADAPTIVE THRESHOLDING

ADAPT_ENABLE	Enable/disable all adaptive thresholding, including geocensoring, and the fine and coarse adaptive thresholding. Default is enabled (1)
GEO_ENABLE	Independent enable/disable switch for geocensoring. Default is enabled (1)
ADAPT_FINE_ENABLE	Independent enable/disable switch for fine map thresholding. Default is enabled (1).
ADAPT_COARSE_ENABLE	Independent enable/disable switch for coarse map thresholding. Default is enabled (1).
GEO_ACTIVE_COUNT_THR	Geomap active cell count threshold. If the number of target reports in a given geo cell over the 512 scan integration period reaches or exceeds this number, the cell is declared an 'active' cell and more detailed statistics are maintained. The default is 5 targets.
GEO_HIST_COUNT_THR	Geomap histogram bin count threshold. If the target count in a given histogram bin exceeds this threshold (default of 5), the hit history is incremented for the amplitude bin.
GEO_HIT_THR	Geomap hit threshold, expressed as a percent. The nominal value of 10 results in cells that exceed the count threshold on more than 10 percent of the 512-scan update cycles having their amplitude thresholds set. Cells with lower hit percentages will have thresholds of 0 but will still be maintained as active cells until the hit fraction drops below GEO_DROP_THR.
GEO_MIN_DET_THR	Geomap minimum detectable hit threshold, expressed as a percent, that will result in the hit history remaining above GEO_HIT_THR for up to GEO_PERSIST_PERIOD hours without a hit. The nominal value is 14 percent. This VSP, in conjunction with GEO_PERSIST_PERIOD and GEO_MIN_THR, are used to describe the desired behavior of the geomap, and the C&I code internally calculates the appropriate time constant. See section 5.2.3 for details.
GEO_PERSIST_PERIOD	Geomap persistence period, expressed in hours. The default is 72 hours, resulting in a geomap time constant sufficient to maintain cells with the minimum detectable hit percentage (14 percent) above the minimum value for active thresholding (10 percent) for up to three days (a long weekend with little rush hour traffic).
GEO_DROP_THR	Geomap drop threshold, expressed as a percent. If the long-term average hit percentage is below this value, the cell will no longer be considered an active cell, freeing it up for use at another location. This value should always be set to less than ADAPT_GEO_HIT_THR to prevent cells at the minimum

GEO_CHILDHOOD_THR

detectable level from being constantly dropped and then later re-quired. The nominal value is 4 percent.

Geomap childhood age threshold, in units of 512-scan update periods (40 minutes each). Newly active geocells are considered to contain insufficient history for thresholding purposes (or potential discarding of the active cell) until they have aged beyond childhood. There are 36 512-scan update periods per day. The default setting is 36 (1 day).

GEO_THR_BOOST

Geomap threshold boost value, in dB. When the geo threshold is set for an active cell, it is set to the floor value of the first histogram bin whose target count is less than the count threshold. To provide an extra thresholding margin for targets that lie near the boundary, this 'boost' value is added to the base threshold determined from the histogram. The default value for this parameter is 4 dB, which is one-half the histogram bin width.

GEO_ENAB_COUNT_THR

Geo enable map count threshold. Coarse resolution geo cells that have fewer hits than this number in a single update cycle are 'aged' by 1 cycle. If the hit count exceeds the threshold, the enable map's age is reset to 0. The coarse map cell 'age' then is a measure of how many update cycles have passed since there was significant target activity in the cell. The default value for this parameter is 7, which because of the increased cell size of the enable map is significantly less in terms of target density than the high-resolution geomap count threshold.

GEO_ENABLE_AGE_THR

Geo enable map age threshold, in units of 512-scan update intervals. If the enable map cell age exceeds the specified age, the enable map cell is flagged as disabled, thereby inhibiting all geocensoring in that region. The default setting for this parameter is 3 (120 minutes)

GEO_N_AIRPORT_MASKS

Number of airport mask regions. The maximum number of regions is 20. The default is 0. (Very site dependent)

GEO_AIRPORT_MASKS[20][4]

Airport mask regions specifying areas to disable geocensoring. Airport runways and final approaches, both at the main airport and any satellite airports, are normally mapped out using this mechanism. Each region consists of the rectangular area bounded by two corner points, specified in polar coordinates. For each region, the four values are:

- R₁ Range of 1st corner, in range gates
- A₁ Azimuth of 1st corner, in ACP
- R₂ Range of 2nd corner, in range gates
- A₂ Azimuth of 2nd corner, in ACP

GEO_N_MTI_REGIONS

Number of MTI reflector regions. The maximum number of MTI regions is four. Most airports use one or two MTI reflectors posi-

tioned near the end of a runway (a known position on the radar scope) as a system confidence check.

GEO_MTI_REGIONS[4][4]

MTI reflector locations, specified in polar coordinates. Unlike the airport mask, these regions are very small and well defined, so they are specified in normal polar space. For each region, the four values are:

- R₁ Minimum range, in range gates
- A₁ Minimum azimuth, in ACP (the counter-clockwise 'edge')
- R₂ Maximum range, in range gates
- A₂ Maximum azimuth, in ACP (the clockwise edge)

GEO_N_FIXED_REGIONS

Number of fixed thresholding regions. The maximum number is 4.

GEO_FIXED_REGIONS[4][5]

Coordinates of fixed thresholding regions, along with the fixed threshold choice, (0 or 1). Like the MTI regions, these will be small in size as they are intended to solve certain pathological problems at road intersections, etc. The five values for each region are:

- R₁ Minimum range, in range gates
- A₁ Minimum azimuth, in ACP (counter-clockwise 'edge')
- R₂ Maximum range, in range gates
- A₂ Maximum azimuth, in ACP (the clockwise edge)
- F Fixed Threshold Set (0 or 1)

GEO_FIXED_THR_SET_0[5]

First set of five thresholds, one for filter class, where the classes are defined as filters +/-0, +/-1, +/-2, +/-3, and +/-4. Threshold values are specified in dB.

GEO_FIXED_THR_SET_1[5]

Second set of five thresholds, specified as above.

FINE_UPDATE_INT

Fine map update interval, in scans. The default value is 120 scans.

FINE_HIT_THR

Fine map hit threshold. If the number of hits in any amplitude bin at any given cell exceeds this number, the fine threshold is raised for that cell and the four closest neighboring cells. The default threshold is 10 hits.

FINE_AMP_HIST_FLOORS[10]

Adaptive thresholding amplitude histogram bin floor values. The fine and coarse adaptive thresholds both utilize the same amplitude histogram bin spacing. The default set of bin floor values is:

<u>Bin</u>	<u>Floor Amp (dB)</u>	<u>Bin</u>	<u>Floor Amp (dB)</u>
1	15	6	31
2	17	7	35
3	20	8	39
4	23	9	43
5	27	10	47

FINE_MAX_THR

Fine map maximum threshold value, in dB. When the hit count in the uppermost histogram bin exceeds its hit threshold, the amplitude for that cell is raised to this maximum value (default value of 50 dB). Additionally, targets with amplitudes exceeding this value will not be subjected to the fine map counting process.

COARSE_UPDATE_INT

Coarse map update interval, in scans. The default value is 20 scans.

COARSE_HIT_THR

Coarse map hit threshold (See above). The default value is 10 hits.

COARSE_AMP_HIST_FLOORS[10]

Coarse map adaptive thresholding amplitude histogram bin floor values. The default set of bin floor values (currently identical to fine map defaults) is:

<u>Bin</u>	<u>Floor Amp (dB)</u>	<u>Bin</u>	<u>Floor Amp (dB)</u>
1	15	6	31
2	17	7	35
3	20	8	39
4	23	9	43
5	27	10	47

COARSE_MAX_THR

Coarse map maximum threshold value, in dB. Default is 50.

ADAPT_EXCL_RANGE_BOUND

Adaptive thresholding exclusion test range boundary. Target at shorter range than ADAPT_EXCL_RANGE_BOUND uses the first set of exclusion table entries (see next VSP) while targets at greater range use the second set.

ADAPT_EXCL_TABLE[2][6]

Adaptive thresholding exclusion table, indexed by confidence and range ring (see above). The table values represent the quality necessary at the corresponding confidence/range to exempt the report from adaptive thresholding. The default setting requires that the quality is 3 for all combinations of confidence and range, so the table is simply filled with the value of 3.

7.4 TWO-LEVEL WEATHER

WX_HI_NWS_LEVEL

High-level weather threshold (2-6). The levels represent NWS weather levels. This value is not so much a parameter as an

informational input (from the RMS) that is made necessary because the incoming WX data stream does not contain the NWS levels represented by the high and low WX detections. Most sites use 4 as the high level, and 2 as the low level, and RMS passes the information to C&I via the VSP structure.

WX_LO_NWS_LEVEL

Low level weather threshold. See above.

WX_9_CELL_SPATIAL_THR

Nine-cell spatial filter threshold for contouring of weather data. If the number of WX detections in a 3x3 cell region reaches the threshold, then the center cell in the 3x3 region is declared a detection; otherwise, it is filtered out. The default value is 5. At the range boundary where only six spatial cells are available, 2/3 of this number is used as the threshold.

7.5 MISCELLANEOUS

ARTSIII_QUAL_TABLE[4][6]

ARTSIII Quality Table. This table provides a mechanism for converting the C&I quality, confidence values to the desired ARTSIII Quality values. The default settings for the table are shown below:

	Confidence					
Quality	0	1	2	3	4	5
0	0	1	1	3	3	4
1	0	2	0	5	4	5
2	0	3	0	6	5	6
3	0	4	0	7	6	7

8. PERFORMANCE MONITORING/ALARMS

For every scan, C&I accumulates the following set of performance monitoring and alarm information and outputs it to the ASR-9 RMS via the MIP.

8.1 PERFORMANCE MONITORING

8.1.1 Radar Input Primitives

FILTER_COUNTS[10] Number of hits in the last scan for each of the Doppler filters.
TOTAL_FILTER_COUNT Number of hits in the last scan, summed over all Doppler filters.

8.1.2 RFI

RFI_DEL_RPTS Number of reports deleted by primary RFI thresholding.
RFI_FLAG Flag that is set if the supplemental RFI thresholding detected RFI on the previous scan.

8.1.3 Adaptive Thresholding

ADAPT_INPUT_RPTS Number of reports input to the adaptive thresholding process.
GEO_EXCL_RPTS Number of reports excluded from geocensoring, based on interpolated Doppler and RTQC/MTI tests.
GEO_FLAGGED_RPTS Number of reports flagged by geocensoring.
GEO_PASSED_RPTS Number of reports passed by geocensoring.
GEO_ACTIVE_CELLS Count of the number of active cells in the current geocensor map.
ADAPT_MG_EXCL_RPTS Number of reports excluded from multi-grid thresholding due to quality/confidence/interpolated Doppler.
ADAPT_MG_DEL_RPTS Number of reports deleted by multi-grid thresholding.
ADAPT_MG_FINE_DEL_RPTS Number of reports deleted by fine map thresholding.
ADAPT_MG_COARSE_DEL_RPTS Number of reports deleted by coarse thresholding.
ADAPT_MG_PASSED_RPTS Total number of reports passed by fine/coarse multi-grid thresholding.

8.1.4 Report Statistics

TOTAL_REPORTS Total number of reports output by C&I.
REPORTS_BY_QUAL[4] Number of reports output for each of the four quality values.
REPORTS_BY_CONF[6] Number of reports output for each of the six confidence values.

8.1.5 Two-Level Weather

WX_HI_INPUT_COUNT	Number of raw high-level weather detections input to C&I.
WX_LO_INPUT_COUNT	Number of raw low-level weather detections input to C&I.
WX_HI_OUTPUT_COUNT	Number of contoured high-level weather detections output by C&I.
WX_LO_OUTPUT_COUNT	Number of contoured low-level weather detections output by C&I.

8.2 ALARMS

C&I sets a number of alarms to alert the RMS when something unusual occurs. Depending on the seriousness of the alarm, the RMS can elect to switch to the standby channel or simply log the alarm in a history buffer to inform the site technician of the event. The alarms are sent at the same interval as the performance monitoring messages, once per scan.

8.2.1 Input Data Integrity

INPUT_RANGE_SEQ_ALARM	This alarm is set if the input data contains range data that is in non-increasing order.
INPUT_AZ_SEQ_ALARM	This alarm is set if the input data contains azimuth data that is not monotonically increasing from one CPIP to the next.


8.2.2 Adaptive Thresholding

GEO_ACTIVE_OVERFLOW	The number of active geo cells exceeded the maximum allowed, resulting in a newly active geocell not being added to the map, but simply being dropped.
----------------------------	--

9. CURRENT STATUS AND FUTURE WORK

Like the other 9-PAC algorithms, the C&I algorithms are now implemented 100 percent in ANSI-C and can now be run off line under UNIX as well as in the 9-PAC board. The algorithms have undergone initial field testing at the Lincoln Laboratory ASR-9 TRDF Site in Albuquerque, New Mexico. When run in conjunction with the new 9-PAC Tracker, performance exceeded that of the original ASR-9 C&I/Tracker combination, with a false-alarm rate of less than one target/scan.

In the short term, future efforts will be directed towards passing the Independent Verification and Validation (IV&V) test phase, which will be conducted at Westinghouse Corporation during the coming months. In addition, the code will be integrated and tested with the Mode-S sensor, either at the FAA Technical Center or at Lincoln Laboratory's ASR-9 in Lexington, MA. As time permits, the algorithms will be run off line with an increasing variety of data sets from ASR-9s around the country, primarily to monitor the effect of different environments on the new geocensoring/adaptive thresholding algorithms.



APPENDIX A. C&I STATE AND DATA FLOW DIAGRAMS

A top-level block diagram of C&I is shown in Figure 18. Detailed flowcharts of all major components are shown in Figures 19-35.




Figure 18. C&I processing top level block diagram.




Figure 19. Input processing state transition diagram.




Figure 20. Front-end geocensor processing.




Figure 21. ZVF overload processing.




Figure 22. Saturation preprocessing.




Figure 23. Target correlation.




Figure 24. Association / range resolution processing.




Figure 25. Primary RFI processing.




Figure 26. Azimuth centroiding.




Figure 27. Doppler interpolation.




Figure 28. Doppler averaging.




Figure 29. Supplemental RFI processing.




Figure 30. Geocensoring and adaptive thresholding.






Figure 31. Geomap threshold update.



* Geo Enable Map Update is Spread out over the Update-Cycle to reduce peak processor loading (as with the Primary Geo Map).

Figure 32. Geo enable map update.



Notes:

1. Cells to Update Spread Across Update Cycle to Reduce Peak Processor Loading
2. Identical Logic Used by Fine and Coarse Adaptive Maps

Figure 33. Adaptive map threshold update process.





Figure 34. Geoladaptive map target count update process.



APPENDIX B. INPUT/OUTPUT FORMATS

INPUT FORMAT

The expected sequence of data types for a single CPIP is shown below. Detailed descriptions of the individual word formats are provided on following pages.

C & I Input Format

Az Header
Az Header 1's Complement
Az Data Word 0 - High PRF CPI
Az Data Word 1 - Low PRF CPI
Range Group(s) (0-960)
.
.
WX Header
WX Data (8 words)

Range Group Data Block

Range Header
Range Data Word
CPI Data Block(s) (1 if data exists on only 1 CPI, 2 if data for both CPI's)

CPI Data Block

CPI Header (High or Low PRF)
CPI Filter Magnitude Data (0-10 words)

Figure 35. CPIP input data layout.

Az Header

X	X	X	X	X	X	X	X	X	X	F ₆	F ₅	F ₄	F ₃	0	0	1
---	---	---	---	---	---	---	---	---	---	----------------	----------------	----------------	----------------	---	---	---

F₃ High PRF ZVF Overload Flag. This flag is set (1), if a ZVF overload was detected on the *previous* high PRF CPI.

F₄ Low PRF ZVF Overload Flag. This flag is set (1), if a ZVF overload was detected on the *previous* low PRF CPI.

F₅ Short CPI Pair. If set, there were fewer than the required 18 pulses in the previous CPI Pair.

F₆ Coast Flag. If set, the radar transmitter is off.

Az Header Complement

1's Complement of Az Header Word (above)
--

The 1's complement of the azimuth header word immediately follows the header word itself and is used to validate the header

Az Data Word

12-bit Azimuth Position (ACP)	0	0	0	0
-------------------------------	---	---	---	---

Two azimuth data words follow the az complement word, the first containing the antenna azimuth for the high PRF CPI and the second containing the azimuth for the low PRF CPI. Antenna position is given in terms of antenna control pulses (ACPs), with 12-bit accuracy. An ACP of 0 always corresponds to magnetic north. The ACP for the high PRF CPI is sampled at pulse 6, and for the low PRF CPI at pulse 15.

Range Header

11-bit Range Data (LSB = 1/16 NMI)	GZON	0	1	0
------------------------------------	------	---	---	---

GZON 0 = No geocensoring at this range cell
 1 = Shaped geocensoring
 2 = Flat geocensoring
 3 = MTI reflector

This ASR-9 generated GZON field is ignored by the 9PAC, which instead uses its own internal geocensoring.

Range Data

X	X	X	X	X	X	X	X	X	X	X	Block Count	0
---	---	---	---	---	---	---	---	---	---	---	-------------	---

Block Count Number of words of CPI headers + data in the current range block.

CPI Header

X	X	X	X	X	F ₁₀	F ₉	F ₈	F ₇	F ₆	F ₅	F ₄	F ₃	1	0	1
---	---	---	---	---	-----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	---	---	---

F₃ PRF (0 = High PRF CPI, 1 = Low PRF CPI)

F₄ Saturation Flag (1 = Saturation present)

F₅ RTQC Flag. If this bit is set to 1, the data for this CPI is from a test

- F₆ Short CPI Pair (1 = Yes, 0 = No)
- F₇ FIR Filter Selection (1 = Heavy Clutter Filters, 0 = Normal)
- F₈ Beam (1 = High, 0 = Low)
- F₉ -ZVF Delta Crossing
- F₁₀ +ZVF Delta Crossing

The ZVF Delta crossing flags signify that the clutter exceeded the clutter map threshold plus a 'delta' value (VSP). This has the potential of being used to differentiate 'high confidence' ZVF returns from those that just barely exceed the clutter map threshold. Current ASR-9 practice is to set the 'delta' VSP to 0, so all ZVF primitives will have the flag set.

CPI Data

10-bit Filter Magnitude (LSB = 3/32 dB)	Filter Number	F ₁	0
---	---------------	----------------	---

F₁ Peak filter tag. If set, this is the peak filter output for the CPI

Filter Number The filter number correspondence is as follows:

- 1 = (-3) 6 = (+1)
- 2 = (-2) 7 = (+2)
- 3 = (-1) 8 = (+3)
- 4 = (-0) 9 = (+4)
- 5 = (+0) 10 = (-4)

WX Header

X X X X X X X X X X X X	F ₄	F ₃	1	1	1
-------------------------	----------------	----------------	---	---	---

F₃ Weather Level (0 = Low, 1 = High). The level alternates from low to high on successive scans.

F₄ Clear Day Map Mode Bit (0 = Normal Mode, 1 = Clear Day Map Mode). If this bit is set, the smoothing of weather data is inhibited to allow generation of an accurate clear day WX map.

WX Data

WX Detections (15 range cells worth)	0
--------------------------------------	---

Each word of WX data contains weather detections for 15 1/2 NMI range cells (15 bits where 1 = WX, 0 = No WX). A total of 8 words are used to pass the entire 120 range cells of WX data for each CPI. Cell Range increases from LSB to MSB, and from word 0 to word 7.


OUTPUT FORMAT

The C&I output report format closely resembles the original ASP output format, with a few exceptions. The first is that the C&I code outputs the reports as a list, requiring the presence of the linked list fields at the start of the report. These words are actually stripped off by the 9-PAC 'glue' code (see ~9pac/comm/src), which sends the reports to Merge using an ordinary flat buffer. Second, the ASP's 'Initial Quality and Confidence' field is no longer used and is now a spare (word 20 below). One new field, 'Multigrid Adaptive Thresholding Info' has been added in a previously unused word (21).

0	Linked List Forward Link
1	Linked List Backward Link
2	Report Type
3	Range
4	Azimuth
5	Runlength
6	Track Eligibility
7	Quality
8	Confidence
9	Max Amplitude
10	Low PRF Max Amplitude Filter Number
11	High PRF Max Amplitude Filter Number
12	ARTSIII Quality
13	High PRF Interpolated Doppler
14	Low PRF Interpolated Doppler
15	Max Amplitude Filter Number
16	Hit/Miss History
17	Flags Word 1
18	Flags Word 2
19	Azimuth Centroiding Algorithm ID
20	Spare
21	Multigrid Adaptive Thresholding Info
22	Azimuth Degrade Flag
23	Max Filter Amplitude for (-3) Filter
24	Max Filter Amplitude for (-2) Filter
25	Max Filter Amplitude for (-1) Filter
26	Max Filter Amplitude for (-0) Filter
27	Max Filter Amplitude for (+0) Filter
28	Max Filter Amplitude for (+1) Filter
29	Max Filter Amplitude for (+2) Filter
30	Max Filter Amplitude for (+3) Filter
31	Max Filter Amplitude for (-4/+4 combined) Filter


Note that in the 9-PAC Phase 2 Implementation, this format is used internally by the 9-PAC, but the actual output format to the MIP is a separate MODE-S compatible format. The translation is done by the I/O processor (not the C&I processor) just prior to outputting the reports to the MIP.

Type




F₃ RTQC Flag. Set to 1 if the report is an RTQC target


Range




Azimuth



Runlength




Track Eligibility



ELIG 0 = No Initiate, no update. 1 = No initiate, update ok.
2 = Initiate ok, update ok. 3 = No initiate, update only if (original ASP) Tracker in states 4, 6-9.


Note: This field is actually not used by the new 9-PAC Tracker but is currently set for backwards compatibility reasons. It may be freed up for other uses at some point.

Report Quality



QUAL 0 = 1 CPI Report (High or Low PRF).
1 = 2 CPI Report (High/Low or Low/High PRF combination).
2 = 2+ CPI report (one PRF).
3 = 3+ CPI (two PRF's)

Report Confidence



CONF 0 = Geocensored (Adaptive Method)
1 = Geocensored (Fixed threshold method)
2 = RFI Flagged Target
3 = NZVF target in zone 1
4 = ZVF target
5 = NZVF target in zone 2
6,7 illegal confidence values.

Max Amplitude 0 0 0 0 0 0 10-bit Filter Amplitude (LSB = 3/32 dB)

This word contains the maximum filter amplitude contained in the report, using the information stored for both PRFs.

Low PRF Max Amplitude Filter 0 0 0 0 0 0 0 0 0 0 0 0 FILTNUM

High PRF Max Amplitude Filter 0 0 0 0 0 0 0 0 0 0 0 0 FILTNUM

- FILTNUM
- 1 = (-3) 6 = (+1)
 - 2 = (-2) 7 = (+2)
 - 3 = (-1) 8 = (+3)
 - 4 = (-0) 9 = (+4)
 - 5 = (+0) 10 = (-4)

ARTSIIIA Quality 0 0 0 0 0 0 0 0 0 0 0 0 ARTSQUAL

ARTSQUAL
 ARTSIIIA quality value based on a combination of the 'normal' report confidence and quality. Values range from 0 to 7, with the lower numbers (0,1,2) generally representing geocensoring and RFI and higher numbers representing higher quality/confidence data.

High PRF Interpolated Doppler 0 0 0 0 0 0 0 0 0 0 6-bit Doppler Value

Low PRF Interpolated Doppler 0 0 0 0 0 0 0 0 0 0 6-bit Doppler Value

The Doppler value is stored as a six-bit signed number scaled to the Nyquist value for the particular PRF.

- 0 Zero doppler
- 31 Max positive doppler
- 31 Max negative doppler
- 32 Nyquist doppler (ambiguous)

**Max Amplitude
Filter Number**

0	0	0	0	0	0	0	0	0	0	0	0	0	FILTNUM
---	---	---	---	---	---	---	---	---	---	---	---	---	---------

FILTNUM

- 1 = (-3) 6 = (+1)
- 2 = (-2) 7 = (+2)
- 3 = (-1) 8 = (+3)
- 4 = (-0) 9 = (+4)
- 5 = (+0) 10 = (-4)

Hit/Miss History

0	0	14-bit High PRF/Low PRF CPI Hit Bit Mask
---	---	--

This field contains a bit mask representing the hit history of a target. As each CPIP of data is incorporated into the target, its previous history is left-shifted by two, and the hit information for the current CPIP is OR'ed into the lowest two bits (Bit 0 is low-PRF, Bit 1 is high PRF). 14-bits provide enough storage for the maximum runlength target (seven CPIs). Examples:

History value: 00000000000001

Single CPI target with hit on low PRF

History value: 00000000000111

3 CPI target with hits on low-high-low PRF

Flags Word 1

F15	F14	F13	F12	F11	F10	F9	F8	F7	F6	F5	F4	F3	F2	F1	F0
-----	-----	-----	-----	-----	-----	----	----	----	----	----	----	----	----	----	----

- F0 RTQC Flag. Set to 1 if report is an RTQC target
- F1 Geocensor MTI Flag
- F2 Range Straddle Flag.
- F3 Geocensor (Fixed) Flag
- F4 Geocensor (Adaptive) Flag
- F5 Adjacent Cell (Rc-1) Flag
- F6 Adjacent Cell (Rc+1) Flag
- F7 Beam Flag (0 = Low, 1 = High)
- F8 Low PRF (Rc) Saturation Flag
- F9 High PRF (Rc) Saturation Flag
- F10 Low PRF (Adj) Saturation Flag
- F11 High PRF (Adj) Saturation Flag
- F12 Low PRF (Rc) Beamswitch Flag
- F13 High PRF (Rc) Beamswitch Flag
- F14 Low PRF (Adj) Beamswitch Flag
- F15 High PRF (Adj) Beamswitch Flag

Flags Word 2

F ₁₅	F ₁₄	F ₁₃	F ₁₂	F ₁₁	F ₁₀	F ₉	F ₈	F ₇	F ₆	F ₅	F ₄	0	0	0	0
-----------------	-----------------	-----------------	-----------------	-----------------	-----------------	----------------	----------------	----------------	----------------	----------------	----------------	---	---	---	---

- F4 Target Range Split Flag
- F5 Target Azimuth Split Flag
- F6 Target Beamshape Split Flag
- F7 Antenna Speedup Flag
- F8 Low PRF ZVF- Delta Flag
- F9 High PRF ZVF- Delta Flag
- F10 Low PRF ZVF+ Delta Flag
- F11 High PRF ZVF+ Delta Flag
- F12 High PRF Interp Dop Test Qualified
- F13 Low PRF Interp Dop Test Qualified
- F14 RFI Flag
- F15 Special Range Split Flag (St. Louis Mod)

Azimuth Centroid Algorithm

0	0	0	0	0	0	0	0	0	0	0	ALG ID
---	---	---	---	---	---	---	---	---	---	---	--------

ALG ID Algorithm ID Ranging from 0-57. See Table 4 in Appendix B for specific ID meanings

Adaptive Thresh Info

F ₁₅	7-bit coarse threshold (dB)	F ₇	7-bit fine threshold (dB)
-----------------	-----------------------------	----------------	---------------------------

- F₇ Flag bit indicating that this report failed the fine threshold. If set, the actual threshold used is stored in bits 0-6
- F₁₅ Flag bit indicating that this report failed the coarse threshold. If set, the actual threshold used is stored in bits 8-14. Note that if the fine threshold is set, the coarse is never set

Azimuth Degrade Flag Word

0	0	0	0	0	0	0	0	0	0	F ₆	0	0	0	0	0	0
---	---	---	---	---	---	---	---	---	---	----------------	---	---	---	---	---	---

- F₆ Single flag used to indicate the data used to form the azimuth estimate was less than optimal. This includes range, azimuth, and beamshape match splits, saturation and beamswitch conditions and targets where the beamshape match split was inhibited due to geocensoring.

Max Filter Amplitudes (words 23-31)

0	0	0	0	0	0	0	10-bit Filter Amplitude (LSB = 3/32 dB)
---	---	---	---	---	---	---	---

Max amplitude target return for each filter number. For each filter, the maximum value for all data in the target matching the filter number (including both PRF's) is determined. Filters -4/+4 are combined into 1 value (the maximum value is chosen)

TABLE 4: Centroid Algorithm IDs

ID	Algorithm/ Data Type(s)
1	Single CPI (One of the Three High PRF Data Types)
2	Two PRF Interpolation (NZVF_ADJ_HI, NZVF_ADJ_LO)
3	Two PRF Interpolation (NZVF_HI, NZVF_ADJ_LO)
4	Two PRF Interpolation (NZVF_ADJ_HI, NZVF_LO)
5	Two PRF Interpolation (NZVF_HI, NZVF_LO)
6	Single PRF Interpolation (NZVF_ADJ_HI)
7	Single PRF Interpolation (NZVF_ADJ_LO)
8	Single PRF Interpolation (NZVF_HI)
9	Single PRF Interpolation (NZVF_LO)
10	Beamshape Match (NZVF_ADJ_HI)
11	Beamshape Match (NZVF_ADJ_LO)
12	Beamshape Match (NZVF_HI)
13	Beamshape Match (NZVF_LO)
14	Beamsplit (NZVF)
15	Beamsplit (Saturated Data of any type)
16	Single CPI (One of the 3 Low PRF Data Types)
17	Two PRF Interpolation (ZVF_HI, ZVF_LO)
18	Single PRF Interpolation (ZVF_HI)
19	Single PRF Interpolation (ZVF_LO)
20	Beamshape Match (ZVF_HI)
21	Beamshape Match (ZVF_LO)
22	Beamsplit (ZVF)
23	Unused
24	Two PRF Interpolation (NZVF_HI, ZVF_LO)
25	Two PRF Interpolation (ZVF_HI, NZVF_LO)
26	Two PRF Interpolation (NZVF_ADJ_HI, ZVF_LO)
27	Two PRF Interpolation (ZVF_HI, NZVF_ADJ_LO)

TABLE 4: Centroid Algorithm IDs (Continued)

ID	Algorithm/ Data Type(s)
28	Beamsplit (Long runlength)
34	Beamshape Match Interpolated Beamsplit (NZVF_ADJ_HI)
35	Beamshape Match Interpolated Beamsplit (NZVF_ADJ_LO)
36	Beamshape Match Interpolated Beamsplit (NZVF_HI)
37	Beamshape Match Interpolated Beamsplit (NZVF_LO)
44	Beamshape Match Interpolated Beamsplit (ZVF_HI)
45	Beamshape Match Interpolated Beamsplit (ZVF_LO)
46	Beamshape Match 1/3,2/3 Beamsplit (NZVF_ADJ_HI)
47	Beamshape Match 1/3,2/3 Beamsplit (NZVF_ADJ_LO)
48	Beamshape Match 1/3,2/3 Beamsplit (NZVF_HI)
49	Beamshape Match 1/3,2/3 Beamsplit (NZVF_LO)
56	Beamshape Match 1/3,2/3 Beamsplit (ZVF_HI)
57	Beamshape Match 1/3,2/3 Beamsplit (ZVF_LO)

APPENDIX C. C&I CONSTANTS

Tables 5-7 contain the C&I constants.

TABLE 5: High PRF Doppler Interpolation Constants

Filter n-1	Filter n	Dop _n	Dop _{avg}	K _n
-4	-3	40	37	5.191e-2
-3	-2	45	42	4.935e-2
-2	-1	50	47	4.666e-2
-1	-0	59	54	2.618e-2
-0	+0	5	0	2.188e-2
+0	+1	14	10	2.618e-2
+1	+2	19	17	4.666e-2
+2	+3	24	22	4.935e-2
+3	+4	29	27	5.191e-2
+4	-4	35	32	4.459e-2
Modifications for Heavy Clutter Filters				
-2	-1	50	47	4.630e-2
-1	-0	59	54	2.408e-2
+0	+1	14	10	2.408e-2
+1	+2	19	17	4.630e-2

TABLE 6: Low PRF Doppler Interpolation Constants

Filter (n-1)	Filter n	Dop _n	Dop _{avg}	K _n
-3	-2	40	38	7.819e-2
-2	-1	47	44	6.766e-2
-1	-0	58	52	3.439e-2
-0	+0	6	0	2.466e-2
+0	+1	17	12	3.439e-2
+1	+2	23	20	6.766e-2

TABLE 6: Low PRF Doppler Interpolation Constants (Continued)

Filter (n-1)	Filter n	Dop _n	Dop _{avg}	K _n
+2	+3	29	26	7.819e-2
+3	-3	35	32	6.766e-2
Modifications for Heavy Clutter Filters				
-2	-1	47	44	5.472e-2
-1	-0	58	52	3.009e-2
+0	+1	17	12	3.009e-2
+1	+2	23	20	5.472e-2

TABLE 7: ASR-9 Low/High Beam Patterns


ACP From Boresight	Low Beam Gain (3/32dB units)	High Beam Gain (3/32 dB units)
0	0	0
+/- 1	-2	-2
+/-2	-4	-4
+/-3	-9	-8
+/-4	-15	-14
+/-5	-23	-21
+/-6	-34	-31
+/-7	-45	-41
+/-8	-60	-55
+/-9	-76	-68
+/-10	-96	-86
+/-11	-116	-104
+/-12	-141	-114
+/-13	-168	-135
+/-14	-197	-157
+/-15	-232	-183

TABLE 7: ASR-9 Low/High Beam Patterns (Continued)

ACP From Boresight	Low Beam Gain (3/32dB units)	High Beam Gain (3/32 dB units)
+/-16	-268	-209
+/-17	-313	-240
+/-18	-359	-272
+/-19	-418	-308
+/-20	-481	-346
+/-21	-557	-390
+/-22	-647	-437
+/-23	-725	-492
+/-24	-769	-566

GLOSSARY

9-PAC	ASR-9 Processor Augmentation Card
A/D	Analog to Digital
ACP	Azimuth Correction Process
ANSI-C	American National Standards Institute C programming language
ASP	Array Signal Processor
ASR-9	Airport Surveillance Radar, 9th generation
ASYNC	Asynchronous
BTD	Beacon Target Detector
C&I	Correlation and Interpolation
CFAR	Constant False-Alarm Rate
CPI	Coherent Processing Interval
CPIP	Coherent Processing Interval Pair
CPU	Central Processing Unit
DRAM	Dynamic Random Access Memory
EEPROM	Electrically Erasable Read-Only Memory
FAA	Federal Aviation Administration
FIR	Finite Impulse Response
HSIB	High Speed Interface Buffer
I and Q	In phase and Quadrature
I/O	Input/Output
IV&V	Independent Verification and Validation
MIP	Message Interface Processor
MTD	Moving Target Detector
MTI	Moving Target Indicator
nm	nautical mile
NZVF	Non-Zero Velocity Filter
PCMCIA	Personal Computer Memory Card International Association
PRF	Pulse Repetition Frequency
R/B	Radar/Beacon target
RAM	Random Access Memory
RFI	Radio Frequency Interference
RMS	Remote Monitoring System
RPM	Revolutions Per Minute
RTQC	Radar Target for Quality Control
SCIP	Surveillance and Communications Interface Processor
STC	Sensitivity Time Control
TRDF	Terminal Radar Development Facility
VSP	Variable Site Parameter
ZVF	Zero Velocity Filter



REFERENCES

1. J. Gertz and G. Elkin, "Documentation of 9-PAC Beacon Target Detection Processing Function," Lexington, MA, M.I.T. Lincoln Laboratory Project Report ATC-220, 26 July 1994.
2. "ASR-9 System Radar Receiver/Processor Instruction/Field Maintenance Manual," Westinghouse Electric Corporation, Baltimore, Maryland, 1989.
3. "Moving Target Indicator and Detector Systems - Resident Course Manual 40392-4/2," Federal Aviation Administration Academy, Department of Transportation, Oklahoma City, Oklahoma, 1989.
4. "ASR-9 System Definition Document for C&I Algorithms," Westinghouse Electric Corporation, Baltimore, Maryland, 1985.
5. D. Karp and J.R. Anderson, "Moving Target Detector (Mod II) Summary Report," Lexington, MA, M.I.T. Lincoln Laboratory Project Report ATC-95, 3 November 1981.