Naval Research Laboratory

Washington, DC 20375-5320

A Stacked Analog-to-Digital Converter for Increased Radar Signal Processor Dynamic Range

VILHELM GREGERS-HANSEN

AEGIS Coordinator
Radar Division

STEVEN M. BROCKETT
PATRICK E. CAHILL
WIN-JOU CHEUNG

Advanced Radar Systems Branch
Radar Division

June 5, 2001

20010628 041

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (*Leave Blank*)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	June 5, 2001	Interim 1 October 1999 – 30 September 2000

4. TITLE AND SUBTITLE

A Stacked Analog-to-Digital Converter for Increased Radar Signal Processor Dynamic Range

5. FUNDING NUMBERS

N0002401WRA0256

6. AUTHOR(S)

Vilhelm Gregers-Hansen, Steven M. Brockett, Patrick E. Cahill, and Win-Jou Cheung

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Naval Research Laboratory
Washington, DC 20375-5320

8. PERFORMING ORGANIZATION REPORT NUMBER

NRL/FR/5306--01-9981

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Commander, Naval Sea Systems Command
Attn: PMS-400
2531 Jefferson Davis Highway
Arlington VA 22242-5170

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (*Maximum 200 words*)

The need for increased dynamic range at the input to the radar digital signal processor has increased steadily over the last decade or so. This has been the result of lower expected radar cross sections, the need for better clutter suppression, and the desire to operate without sensitivity time control (STC) in some radar applications. The analog-to-digital converter (ADC) is the most significant bottleneck in achieving this needed dynamic range performance. In this report, an approach for improving effective dynamic range using multiple ADCs is described. The ADCs are arranged in parallel channels with different gains and the approach is referred to as a "Stacked Analog-to-Digital Converter," or "Stacked ADC," in this report. Detailed results are presented for an experimental system assembled to demonstrate and evaluate the concept.

14. SUBJECT TERMS

Analog-to-Digital Converter	Dynamic range
A/D	Radar signal processing
ADC	

15. NUMBER OF PAGES

29

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

CONTENTS

A STACKED ANALOG-TO-DIGITAL CONVERTER FOR INCREASED RADAR SIGNAL PROCESSOR DYNAMIC RANGE

INTRODUCTION

In new radar systems, the desired dynamic range often exceeds the performance of available analog-to-digital converters (ADCs). To overcome this limitation, radar designers seek to accommodate the total required dynamic range through the addition of sensitivity time control (STC), automatic gain control (AGC), and use of intermediate frequency (IF) limiting amplifiers. Such IF limiters are usually adjusted to match the available instantaneous dynamic range of the ADC.

The advent of stealth technology for aircraft and missiles has increased sensitivity and clutter-suppression requirements by several orders of magnitude, and corresponding improvements in dynamic range and radar stability are needed. While significant advances in ADC technology have occurred, progress has not matched that seen in computer memory and computational capability. As a result, available ADCs do not meet many new advanced radar requirements.

This report describes an approach that effectively overcomes these limitations and provides accurate analog-to-digital conversion with a true instantaneous dynamic range not previously available. The technique uses multiple ADCs, each connected to the radar IF through amplifiers with different gain characteristics. After equalization and recombination, a composite dynamic range is realized that is much greater than that of each individual ADC. Individual ADC limitations on signal-to-noise ratio (SNR) and spurious responses will, however, not be improved by this approach, so that the main restriction of this new approach is an output signal-to-noise-and-distortion (SINAD) that can be no greater than that available from each individual ADC.*

To demonstrate the feasibility and performance of this concept, an experimental "Stacked ADC" was implemented and detailed test results are presented in this report.

RADAR DYNAMIC RANGE REQUIREMENTS

Table 1 is a list of radar needs for increased dynamic range at the input to the digital signal processor (DSP). Maintaining the best possible radar sensitivity at all radar detection ranges of interest is of particular importance in the case of the detection of low-cross-section targets where a significant detection margin may not be available, even at relatively short range. Often the use of STC and AGC will adversely affect the detection performance against such low-cross-section targets.

Clutter attenuation (or cancellation) is often an important radar requirement. For best performance, this requires linear processing over the full clutter dynamic range since any nonlinearity will result in the generation of intermodulation products and, therefore, lowered clutter attenuation [1].

Manuscript approved February 21, 2001.
*This approach will be referred to in this report as a "Stacked Analog-to-Digital Converter" or "Stacked ADC."

Table 1 — Radar Needs for Increased Dynamic Range

- Maintain optimum sensitivity in all range cells
  - Avoid global desensitization
- Maximize clutter attenuation
  - Clutter peaks near ADC top range
  - No clutter saturation or limitation
- Measure cross section accurately for target identification and adaptive sensitivity control
- Use zero-Doppler clutter level for residue control in moving target indicator (MTI) or pulse-Doppler filters
- Handle high-level electromagnetic interference (EMI) without desensitization or loss of data

The accurate measurement of target and clutter amplitude over their full dynamic range is frequently required for target recognition and discrimination and for the suppression of false alarms caused by clutter residues. Again, nonlinearity will result in erroneous measurements.

In a high-power long-range radar system, the dynamic range of input signals and clutter can be as high as 80 to 100 dB above thermal noise, even prior to pulse compression. Such values exceed the capability of currently available ADCs by 20 to 40 dB. The following sections discuss this shortfall, which is the focus of the Stacked ADC technique for dynamic range extension.

EXTENDING THE ADC DYNAMIC RANGE

Table 2 lists techniques that have been used, or proposed, for increasing the effective dynamic range of a radar signal processor beyond the inherent ADC capability. The conventional uses of STC and AGC were mentioned previously and have been used extensively in radar systems.

A more speculative approach uses a control signal, which originates in the signal processor. It sets an analog attenuator in the receiver to continuously ensure that signal and clutter levels fall within the dynamic range of the ADC. However, with this type of approach, it is difficult to maintain low switching transients and sufficient accuracy to support a high value of clutter suppression.

Provided that noise and spurious responses of individual ADCs are independent, an improvement in dynamic range can be achieved by connecting N ADCs in parallel and sum their outputs. However, since the theoretical improvement (in dB) is only $10 \bullet \log_{10}(N)$, a large number of ADCs would be required to realize a significant improvement.

The stacked ADC concept, as described in the next section, is similar to the above-mentioned paralleled ADC approach, except that ADCs are connected to the IF through amplifiers with different gains, and their outputs are combined by scaling and selection rather than by summing.

A "STACKED ADC" PRINCIPLE

Figure 1 shows the basic concept of the proposed "Stacked ADC" approach. This diagram assumes that direct sampling is used at the final IF frequency. These samples must first be converted to baseband by direct

Table 2 — Candidate Techniques for Improving DSP Dynamic Range

- ■ STC

- ■ Adaptive attenuation (switched gain)

- ■ Range-cell to range-cell analog AGC
    - • Constant across full coherent processing interval (CPI)
    - • Switching transients due to Doppler shift

- ■ Digital feedback to analog attenuator

- ■ N ADCs in parallel [$\Delta = 10 \log(N)$ dB]

- ■ N Stacked ADCs [$\Delta = 6(N - 1)$ dB]

- ■ Wait for better ADCs ($\Sigma\Delta$, Cryo)
    - • ADC6644 @ FS = 65 MHz provides 73 dB SINAD and about
      60 dB dynamic range [for 0.2 dB noise figure (NF) loss]

digital conversion, sometimes called digital product detection (DPD). For the direct sampling approach to avoid aliasing of positive and negative frequency spectral components, the ADC sampling rate must be related to the radar IF frequency according to

$$f_S = \frac{4 \cdot f_{IF}}{2 \cdot M - 1},\tag{1}$$

where M can take any value $M = 1, 2, 3$, etc. Thus, candidate sampling rates are $4 f_{IF}$, $1.3333 f_{IF}$, $0.8 f_{IF}$, $0.57 f_{IF}$, etc. The maximum unaliased bandwidth that can be supported is then $B_{max} = f_S/2$, and bandpass filters must be included prior to the ADC sampling to restrict the bandwidth of signals, clutter, and noise to this value, within specified error limits.

Prior to the ADC, the IF path is split into K parallel channels and incrementally amplified in steps of Δ dB. The actual gain increment to be used is a trade-off between the overlap between channels and the number of channels needed to realize the required dynamic range. Each channel is then individually analog-to-digital-converted and converted to complex baseband (I&Q) samples by Digital Product Detectors (DPD).
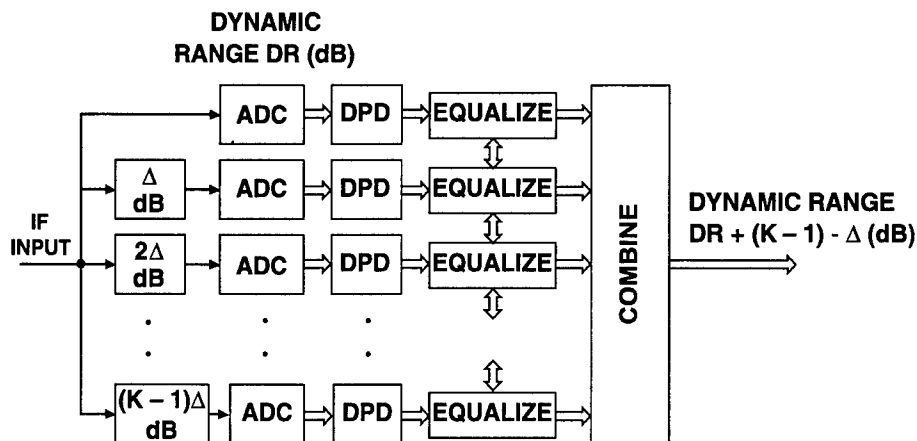

Fig. 1 — Basic "Stacked ADC" concept

Before the K channels can be combined into a single data stream, equalization in each channel is required to bring all outputs into a common number representation and correct for channel-to-channel gain and phase dispersion caused by hardware limitations. The equalizer can take the form of a transversal equalizer with adaptively computed coefficients or simply as a complex gain factor with a value determined by channel-to-channel comparisons. The latter approach was used in the experimental system described below, but a transverse equalizer might be effective for improved performance. The bidirectional arrows between adjacent equalizers suggest a strategy by which adaptive equalization weights are determined by pairwise comparisons of adjacent channels. This comparison may use dedicated pilot signals injected into the radar receiver, or actual target, clutter, and jamming returns received during radar operation. The latter strategy was used in the experimental system to be described.

Finally, the K channels are combined into a single data stream, but with a total dynamic range which is increased by $(K-1) \cdot \Delta$ dB over that of a single channel. The combination of channels is based on a strategy that, on a range-sample-to-range-sample basis, selects the output from the ADC with the highest nonlimiting output.

Figure 2 is a more specific example of an implementation of the Stacked ADC. The gain increments are set to $\Delta = 6$ dB, and the equalization is performed using a single complex constant. The complex equalization factors are obtained by pairwise comparison of nonsaturating data in adjacent channels. The experimental system investigated as part of this study closely follows the implementation shown in Fig. 2.

Figure 3 illustrates a somewhat more conservative combination strategy in which all data from a complete Coherent Processing Interval (CPI) is stored in memory. For each range sample, the output data are taken from the same ADC across the entire CPI. This selection is based on no ADC saturation for any of the received returns across the entire CPI. This approach will ensure that clutter suppression will not be degraded by any inaccuracies in the channel-to-channel equalization.

AN EXPERIMENTAL STACKED ADC SYSTEM

An experimental Stacked ADC system was constructed using available signal processing hardware and laboratory test equipment to the maximum possible extent. As a result, some of the characteristics do not represent those potentially achievable with state-of-the-art components. This approach was selected to
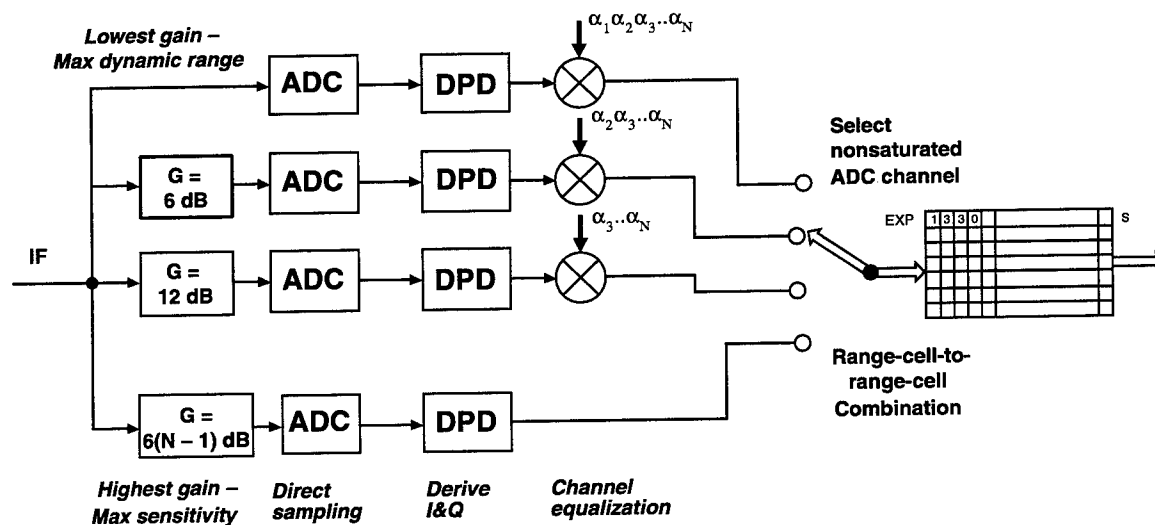

Fig. 2 — Stacked ADC implementation with range-cell-to-range-cell combination
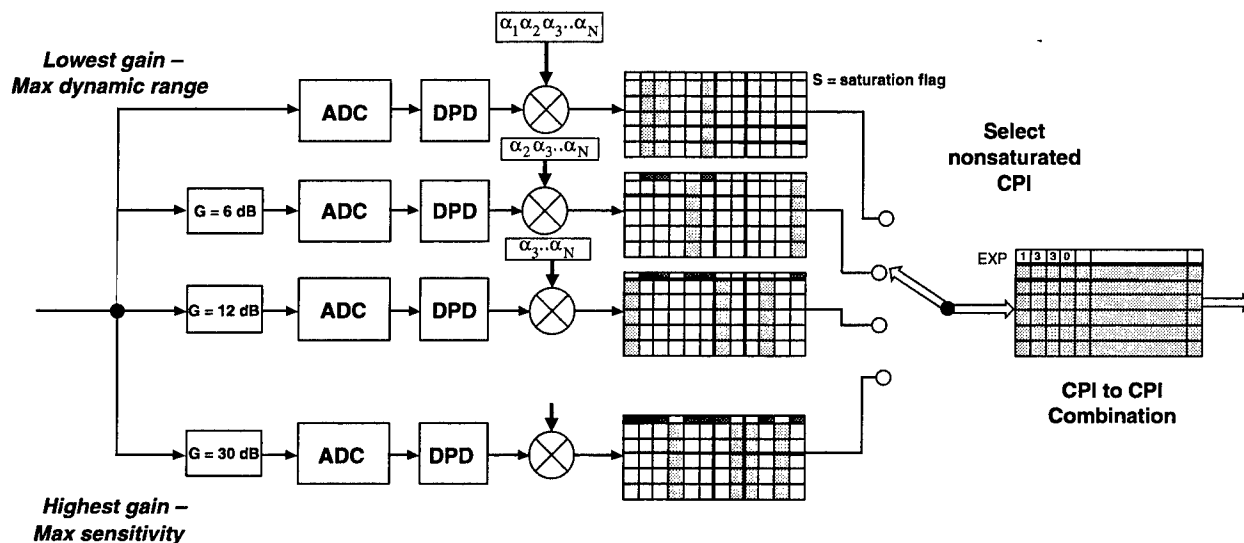
Fig. 3 — Stacked ADC implementation with CPI-to-CPI combination

obtain preliminary results of feasibility and performance in the shortest possible time. An effort is under way to develop an improved Stacked ADC system that will overcome some of the performance limitations of the current implementation.

Figure 4 is a block diagram of the experimental Stacked ADC system. With the ADC sampling rate selected at 20 MS/s (megasamples per second), the intermediate frequency input was defined to be $f_{IF} = 5$ MHz using $M = 2$ in Eq. (1). A total of four channels was implemented using a channel $\Delta = 6$ dB. A practical complication for the implementation of any ADC in a radar system is the need to maintain accurate linearity over the normal input range to the ADC, while limiting the input amplitudes to a value less than the maximum allowed value. In the experimental system, the ADC was protected against excess input voltage by a
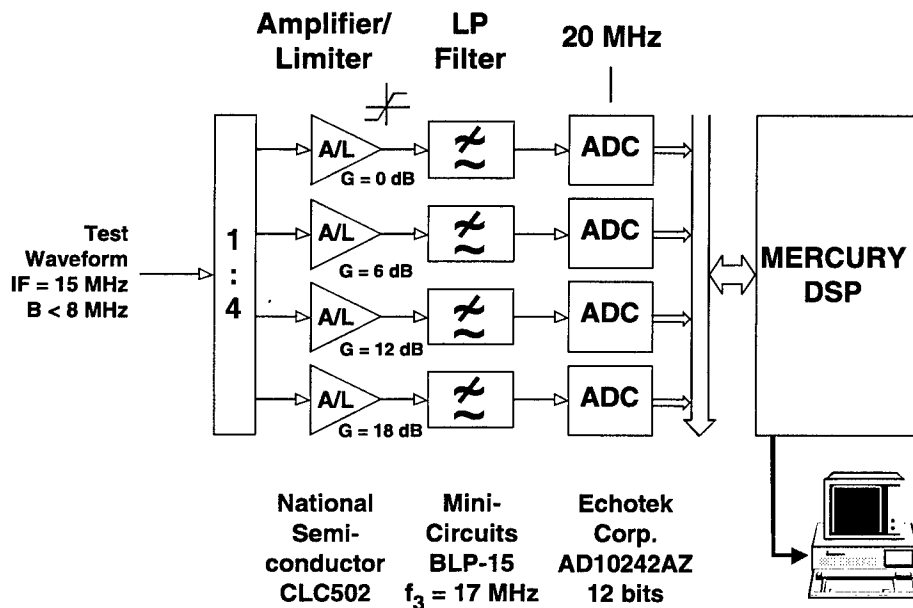

Fig. 4 — Experimental Stacked ADC implementation

National Semiconductor CLC502 limiting amplifier. This device permits setting the gain and limit level using external components. Freedom from harmonics and intermodulation products at the output of these amplifiers is of critical importance for output voltages within the ADC's linear dynamic range. For a required operating range at the ADC input of ±1 V and a maximum input that must be limited to ±2 V, the CL502 provides performance as summarized in Table 3. As seen in this table, harmonics and third-order intermodulation at the linear dynamic range limits show typical values of –50 dB and –60 dB, respectively. To reduce the output harmonics generated by the amplifier/limiter circuit, low-pass filters were inserted. These were catalog items from Mini-Circuits, Model BLP-15 with characteristics listed in Table 4. During system testing, it was discovered that inevitable mismatch in frequency responses between these filters degraded the performance of the Stacked ADC system and the filters were removed in later tests, which is discussed later.

Table 3 — Linear/Limiting Amplifier Specification

- **■ Rise and fall time <3.2 ns**

- **■ Second harmonic distortion at 2 V peak-to-peak output:**
  - **● –50 dB typical, –43 dB minimum**

- **■ Third harmonic intermodulation at 2 V peak-to-peak output:**
  - **● –60 dB typical, –53 dB minimum**

- **■ Settling time to 0.0025% <32 ns**

Table 4 — Lowpass Filter Specification

■ Mini-Circuits BLP-15	
● Passband loss	<1 dB to 15 MHz
● 3 dB point	17 MHz
● >20 dB loss	>23 MHz
● >40 dB loss	>32 MHz
● VSWR in-band	1:1.7 typical
■ Phase and amplitude flatness	Not specified

The ADCs and data-collection system were borrowed from an ongoing signal processor development for a high-power experimental millimeter-wave radar system currently under development by the Radar Division at NRL. The ADC cards, model AD10242AZ, were built by Echotek Corporation (Table 5). They provide 12-bit outputs at a conversion rate of 20 MHz. The Echotek ADC board is in a VME U6 format using an Analog Devices AD10242BZ ADC and provides four channels per card. The specified SINAD ratio for each channel is 61 dB (58 dB minimum) at a 19.5 MHz sampling rate. During testing in the NRL Radar Division, a SINAD of 58 dB was verified. The Echotek VME ADC card is part of a Mercury signal processing system using the Raceway architecture for data transfer. For the Stacked ADC data collection, the number of ADC samples that could be collected contiguously was limited to 8K by the buffer size on the Echotek card.

Table 5 — Analog-to-Digital Converter Specification

- Quad ADCs on Echotek VME 6U board
- ADC Analog Devices AD10242BZ
- Number of bits, 12
- SINAD at 19.5 MHz; 58 dB minimum, 61 dB typical
- Linear range input voltage, 2 V peak-to-peak
- Maximum input voltage, 4 V peak-to-peak

The test configuration for the experimental Stacked ADC system used two signal generators with their outputs summed as shown in Fig. 5. The harmonic level at the signal generator outputs was measured at –57 to –59 dBc, which is significantly better than the linear/limiting amplifiers. This setup allows two-tone intermodulation tests to be performed, providing an effective means of quantitative performance verification of the final Stacked ADC output. The oscilloscope trace in Fig. 5 shows the test input for equal-amplitude signal generator settings, using a frequency offset of 25 kHz.

Figure 6 shows corresponding outputs of the amplifier/limiters. The maximum linear range of the ADC is ±1 V, and effective limiting is ±2 V required to protect the individual ADCs against overvoltage.

Figure 7 shows that the level of the second harmonic generated by the amplifier/limiter for a ±2 V output signal is –49 dBc, which is in good agreement with the specification for the CLC502. Figure 8 shows third-order intermodulation performance of the amplifier/limiter for a two-tone input with peak amplitude of ±2 V. The measured third-order product was –62 dBc, which is slightly better than the typical value for the CLC502.

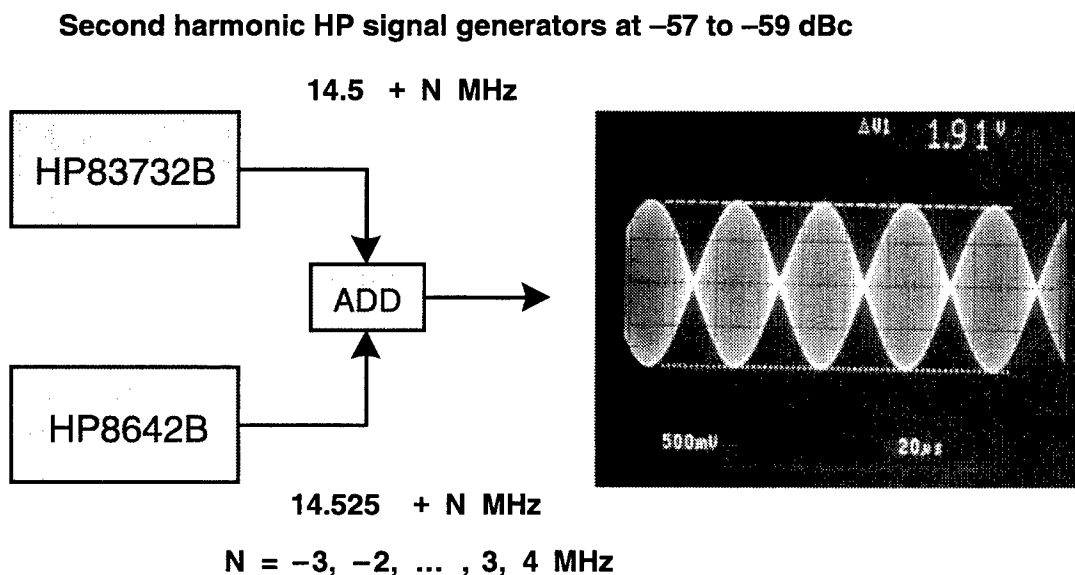**Second harmonic HP signal generators at –57 to –59 dBc**

Fig. 5 — Generation of test input using Hewlett-Packard signal generators
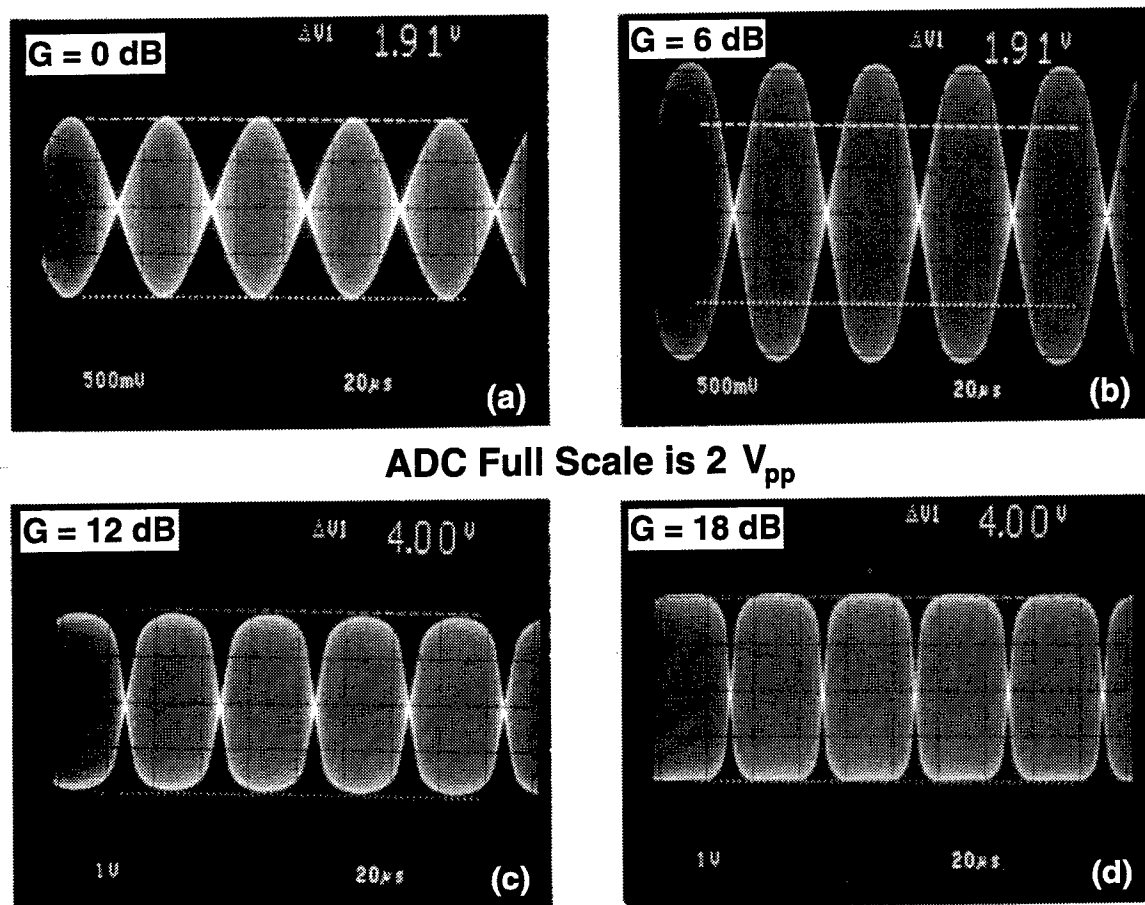
Fig. 6 — Outputs from CLC502 amplifier/limiters for pair of high-level input signals. Limit level is 4 V_{pp}.

■ Single-tone input
  @ 10.5 MHz

■ Amplifier output at full-scale A/D input
  (2 V_{pp})

■ Second harmonic at −48.9 dBc
  (specification: typical −50 dBc,
  minimum ≤ −43 dBc)

■ Low-pass filter added at output to
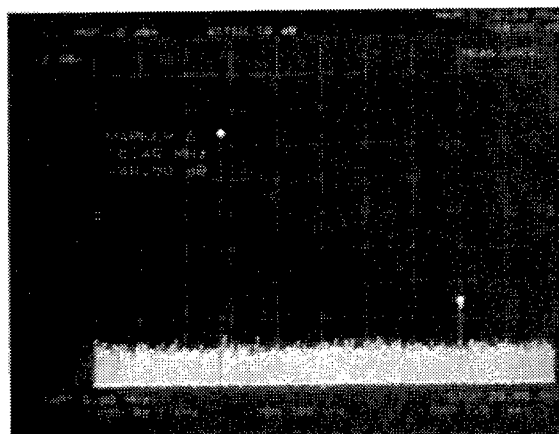 suppress second harmonic

Fig. 7 — Second harmonic at output of CLC502 amplifier/limiter (G = 0 dB) for full-scale single tone input

During each test, a total of 8096 data samples were collected at the ADC sampling rate of 20 MS/s and stored to memory.

- ■ Third-order products 62.2 dB down within linear range (specification: typical –60 dBc, minimum ≤ –53 dBc @ 2 V$_{pp}$)

- ■ Increases to –24.3 dBc for G = 6 dB in amplifier/limiters
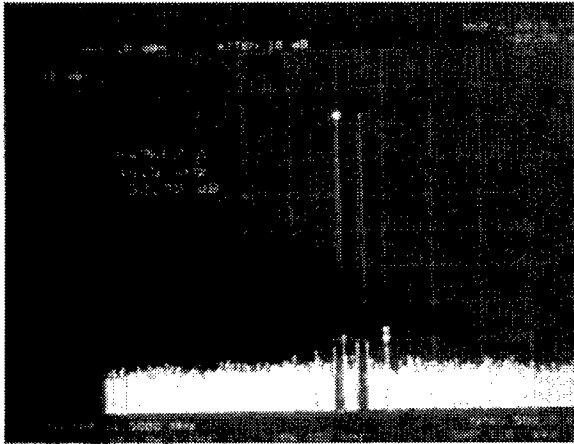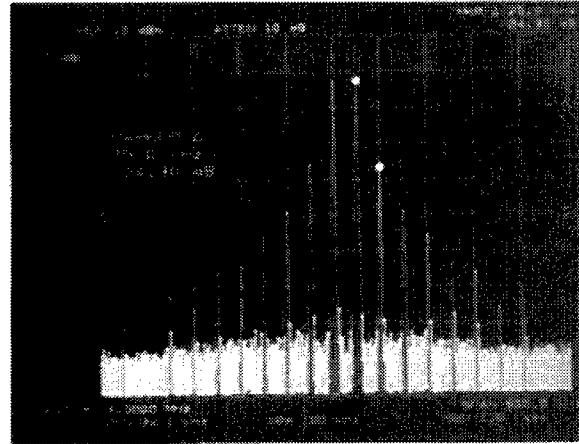
G = 0 dB   G = 6 dB

Fig. 8 — Third-order intermodulation products of two CLC502 amplifier/limiters for full-scale two-tone input

All subsequent data processing was performed off-line on a PC, based on these recorded samples. The analysis was performed in MATLAB and consisted of a number of discrete processing steps. The Appendix lists the MATLAB program. The first step in the processing reads the ADC samples from the data files into the MATLAB program. Figure 9 shows graphs of the four channels of ADC data using the two-tone test signal set as the input and with the amplitudes set close to the maximum linear range of the least-sensitive ADC, as described previously. The peak value shown is $2^{15} = 32767$ and should be interpreted as a signed 16-bit integer word with the 12-bit ADC data shifted to the most significant bit position.

The next step is to implement the DPD, which recovers the complex I&Q data from the direct samples. Figure 10 shows the basic principle of the DPD, or direct sampling technique. An IF filter at the input restricts the bandwidth of all input signals to be less than half the sampling rate of the ADC to be compatible with the DPD implementation, as discussed previously. The spectrum of the signal is illustrated at the right in Fig. 10 and consists of complex conjugate spectral components located at positive and negative frequencies. The complex envelope has the spectrum represented by $2 \bullet A_{+}(f - f_{IF})$, and it is the sampled representation of the corresponding complex signal at a rate equal to one-half the ADC sampling rate, which is sought. The real ADC digitizes the IF signal at the selected sampling rate. This results in an aliased spectrum as shown in Fig. 10 but, due to the carefully selected relationship between ADC sampling rate and the radar IF frequency (see Eq. (1)), no aliasing of positive and negative frequency components has occurred. By multiplication with the time series $e^{j \, \bullet \, i \, \bullet \, \pi/2}$, the spectrum is shifted up in frequency by an amount equal to $1/2 \bullet f_{IF}$. This results in a spectral representation in which the desired spectrum of the complex envelope is centered at zero frequency, but the signal still contains the unwanted negative spectral components (light shading). A finite impulse response (FIR) bandpass filter with a nearly rectangular response is then applied to filter out the negative-frequency components, as shown in the bottom graph on the right. The desired sampled complex envelope representation has now been realized but still at the original sampling
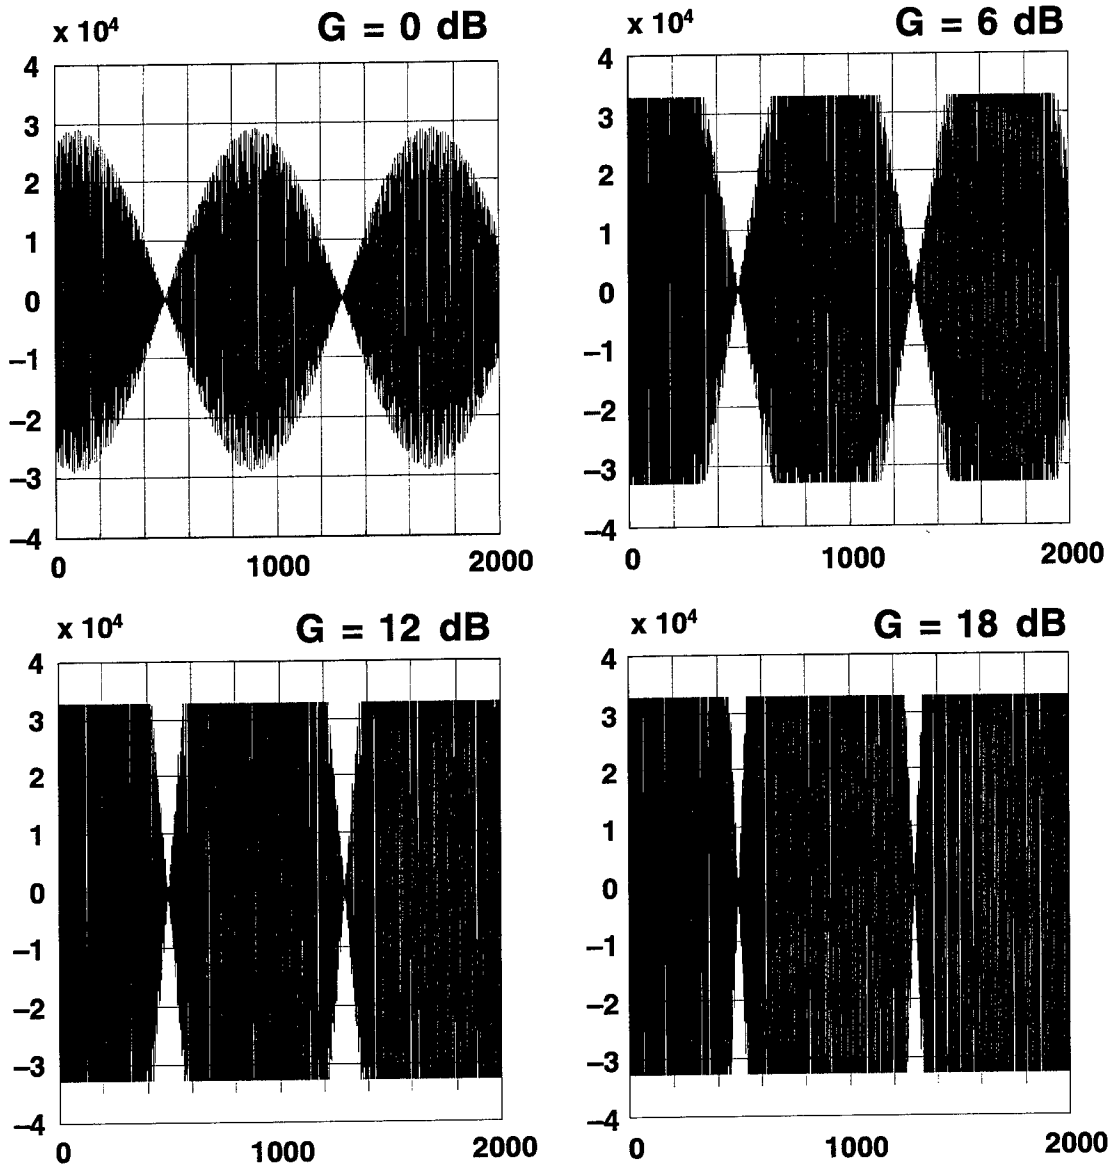
Fig. 9 — ADC direct samples at 20 MS/s in all four channels

rate of f_{AD}. The oversampling is finally removed through decimation by a factor of 2 to obtain the desired sampled complex envelope. Table 6 shows that the FIR filter used to implement the DPD function in this Stacked ADC experiment used 36 taps and real-valued coefficients. (This filter design was obtained by Lockheed Martin and is described in Ref. 2.) Figure 11 shows the amplitude response of the FIR filter. Note that combining FIR filtering and decimation in a multirate FIR filter would simplify the practical implementation of the DPD processing shown in Fig. 10.

Figure 12 shows the output magnitude of the DPD outputs in the four channels of the experimental Stacked ADC for the two-tone input signal. For each channel, a scale factor was applied to the ADC outputs as shown (×8, ×4, etc.) to approximately compensate for the different gains applied at the input. Only the least-sensitive channel correctly reproduces the signal input. In the other channels, the combination of limiting and the DPD operation produces a garbled output during the periods where limiting occurs. Such data would be useless for radar signal processing purposes.

IF
INPUT

IF
Filter

$A_-(f) = A_+^*(-f)$ $S(f)$ $A_+(f)$

$-f_{IF}$ 0 f_{IF}

f_{AD} → ADC

Direct
Sampling
@ f_{AD}

$f_{AD} = \frac{4}{3} \cdot f_{IF}$

$S(f)$

$A_+^*(-f)$ $A_+(f)$

$-f_{IF}$ 0 f_{IF}

$e^{j \cdot i \frac{\pi}{2}}$

Spectrum
Shifted Up
to Zero

FIR
FILTER

Negative
Spectrum
Suppressed

$S'(f)$

$A_+^*(-f)$ $A_+(f)$

$-f_{IF}$ 0 f_{IF}

DEC
By 2

Decimate
by 2

$S''(f)$

$A_+^*(-f)$ $A_+(f)$

$-f_{IF}$ 0 f_{IF}

Complex Envelope
Sampled @ $f_{AD}/2$

Fig. 10 — Basic principle of digital product detector performing direct ADC sampling of IF signal

Table 6 — DPD FIR Filter Coefficients

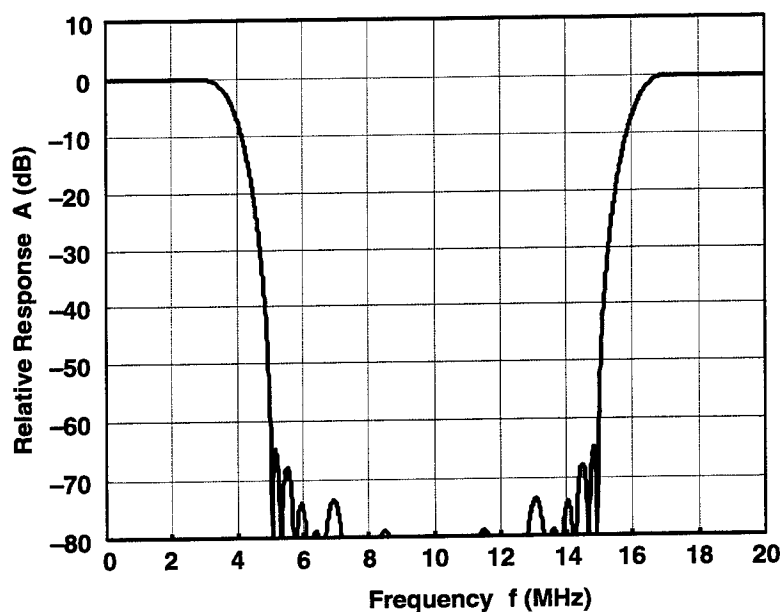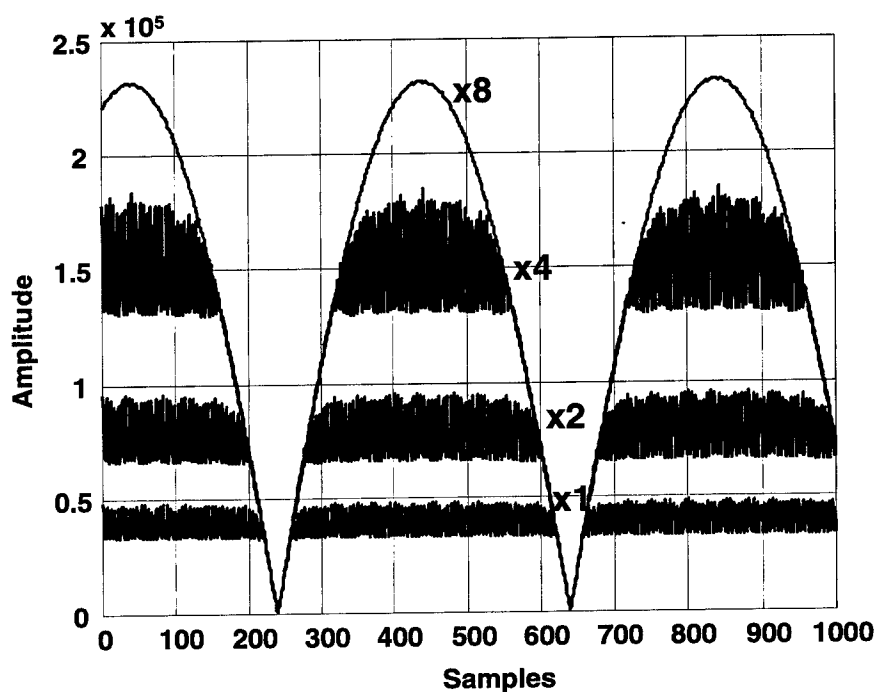| a1	a2	a3	a4	a5	a6	a7	a8	a9
11	32	11	−75	−109	53	272	151	−362
0	0	0	0	0	0	0	0	0
a10	a11	a12	a13	a14	a15	a16	a17	a18
−586	119	1103	745	−1273	−2507	171	6545	12092
0	0	0	0	0	0	0	0	0
a19	a20	a21	a22	a23	a24	a25	a26	a27
12092	6545	171	−2507	−1273	745	1103	119	−586
0	0	0	0	0	0	0	0	0
a28	a29	a30	a31	a32	a33	a34	a35	a36
−362	151	272	53	−109	−75	11	32	11
0	0	0	0	0	0	0	0	0

Fig. 11 — Response of DPD FIR filter

Fig. 12 — Scaled magnitudes in all channels after DPD

Figure 13 illustrates the fundamental principle of the Stacked ADC, where nongarbled data are taken successively from the most-sensitive through the least-sensitive channel to generate the combined magnitude output. For the two-tone input used here, the most-sensitive ADC contributes only to the output over very short periods of time.
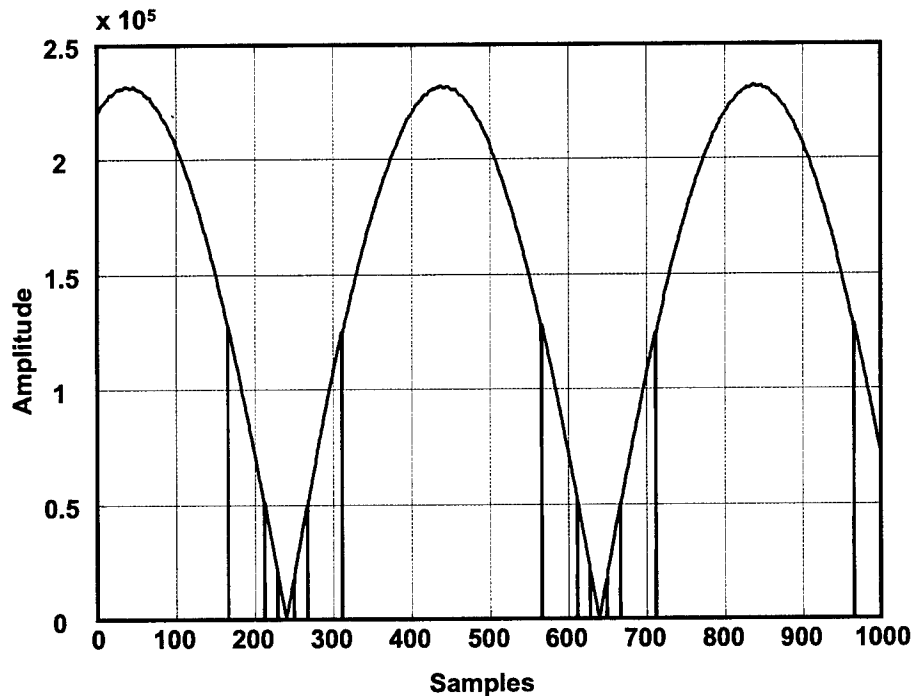
Fig. 13 — Magnitude of combined output from nonsaturating ADCs

To determine what data safely can be obtained from each channel, saturation of the ADC must be monitored sample-to-sample. For the experimental system, this was simply done by testing for full-scale positive or negative values at the ADC output. Sometimes ADCs will have a special output bit flagging the saturation condition. Additionally, the stretching of each saturated sample value by the length of the DPD FIR filter must be taken into account. For the FIR filter of length 36, which is effectively reduced by a factor of 2 due to the final decimation, each saturated sample affects a total of 18 successive samples at the DPD output. Any sample so designated is inhibited from the combined output shown in Fig. 13.

Finally, a strategy must be developed for dynamically obtaining equalization factors for the individual channels. In the experimental system, a simple strategy was implemented in the data analysis to obtain the equalization constant between any adjacent pair of ADC channels. The equalization constant is computed as the average of the complex ratio of all samples for which the most sensitive of the pair of ADCs has a magnitude between full scale and one-half of full scale. The byproduct of this simple strategy is that not all equalization constants are based on the same number of ratios averaged. It is, however, felt that any practical implementation of the Stacked ADC would derive the equalization factors based on a large number of samples. In summary, the data processing consists of the following steps:

1. Derive I&Q samples by DPD processing
2. Identify ADC saturation in each channel
3. Stretch saturation indication based on DPD FIR filter length
4. Perform channel equalization
5. Combine all channels into single-output channel
6. Perform spectral analysis to evaluate the Stacked ADC performance.

Results obtained using the final analysis step are described in the following section.

PERFORMANCE RESULTS FOR THE EXPERIMENTAL STACKED ADC

To evaluate the performance of the Stacked ADC concept using the two-tone input described in the previous section, a spectral analysis was performed on the final output. This spectral analysis was accomplished using a fast Fourier transform (FFT) on the output data while applying a 110 dB Chebyshev weighting to the data. A total of 4,096 complex samples were available at the output for this spectral analysis. Figure 14 shows the result of this spectral analysis for the case in which no attempt was made to equalize the four channels other than applying scale factors of two based on the approximate gain settings at the ADC input. The figure on the left shows the spectrum across the entire sampling interval; the figure on the right magnifies the portion of the spectrum around the input tones, clearly showing the third-order intermodulation products around the two input signals. The signal-to-interference-plus-noise ratio at the output (SINAD) was computed as the ratio between the power contained in the two spectral components of the signal to the total noise power in the remaining parts of the spectrum. A value of the SINAD of only 29.3 dB was achieved in this case without accurate equalization, as might be expected.

In Fig. 15, the equalization factors obtained through the self-calibration procedure were introduced prior to combining the four channels into a single output. A substantial improvement is noted, and the SINAD is now 53.3 dB. For comparison, Fig. 16 shows the corresponding result when only the least-sensitive ADC is used. Nearly identical performance is noted, demonstrating the inherent limitation of individual ADCs.

A critical issue with the Stacked ADC is how well a single set of equalization constants will work across the frequency band of interest. In Fig. 17, equalization was performed at 14.5 MHz, but the input tones were injected at 12.5 and 12.525 MHz. A significant increase in the close-in distortion products is seen, and the SINAD drops to 42.3 dB. While keeping the equalization constants obtained at 14.5 MHz constant, the two-tone input frequencies were varied across the band of interest and, at each input frequency, the SINAD was calculated. Figure 18 shows the resulting curve of SINAD vs input frequency. It is clear that significant channel-to-channel amplitude and/or phase variation is encountered as a function of frequency, resulting in significant performance degradation as a function of frequency. An examination of the hardware implementation of the Stacked ADC determined that the low-pass filters at the amplifier/limiter outputs were the most likely cause of this degradation, and the implementation was subsequently modified to eliminate these filters for a second set of tests, as described in the following section.

MODIFIED STACKED ADC AND REVISED EXPERIMENTAL RESULTS

Figure 19 shows the modified experimental Stacked ADC. The low-pass filters have been removed to improve channel-to-channel tracking in phase and amplitude. The penalty of this change is an increased level in the harmonic outputs of the amplifier/limiter. Based on the measurements performed on the amplifier/limiter, the level of second-harmonic level is –49 dBc. This value, therefore, becomes the expected upper bound on performance. Figure 20 shows results for the case where equalization is performed at the test frequency. Higher harmonic levels are noted, and the SINAD has dropped to 47.3 dB (from 53.3 dB in the original implementation). However, when the equalization constants from the 14.5 MHz test are used but the input two-tone frequencies are changed to 13.5 and 13.525 MHz, only a small degradation in the SINAD to 46.3 dB occurs, as shown in Fig. 21. A complete curve of SINAD vs input frequency, using equalization at 14.5 MHz, is shown in Fig. 22. With reference to the dashed curve copied from Fig. 18, performance obtained without the low-pass filters in the circuit is much more robust as the two-tone input frequency is varied across the band using a fixed value of the calibration constants.
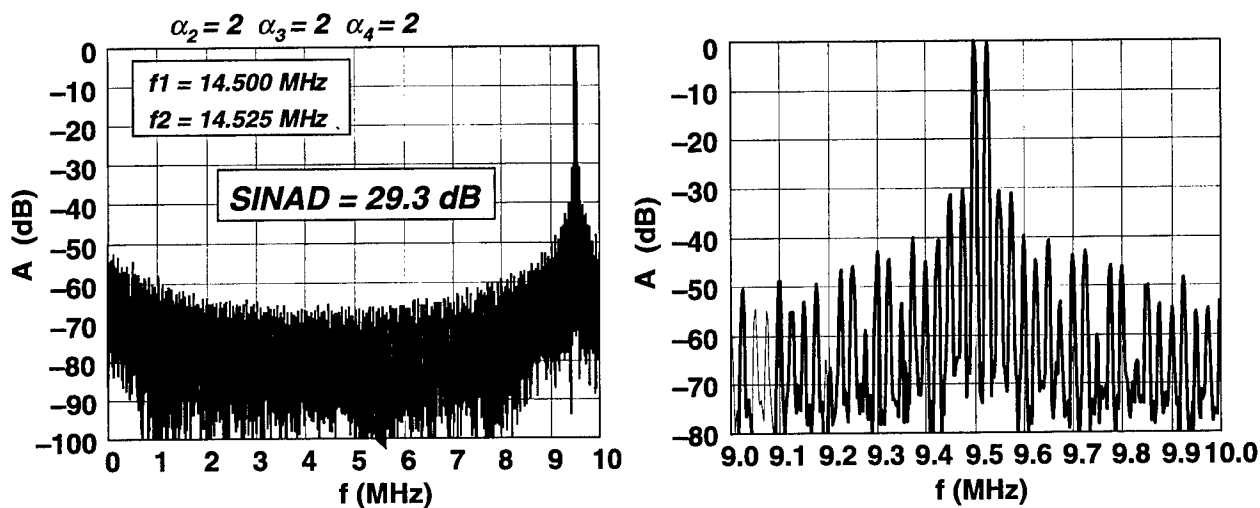
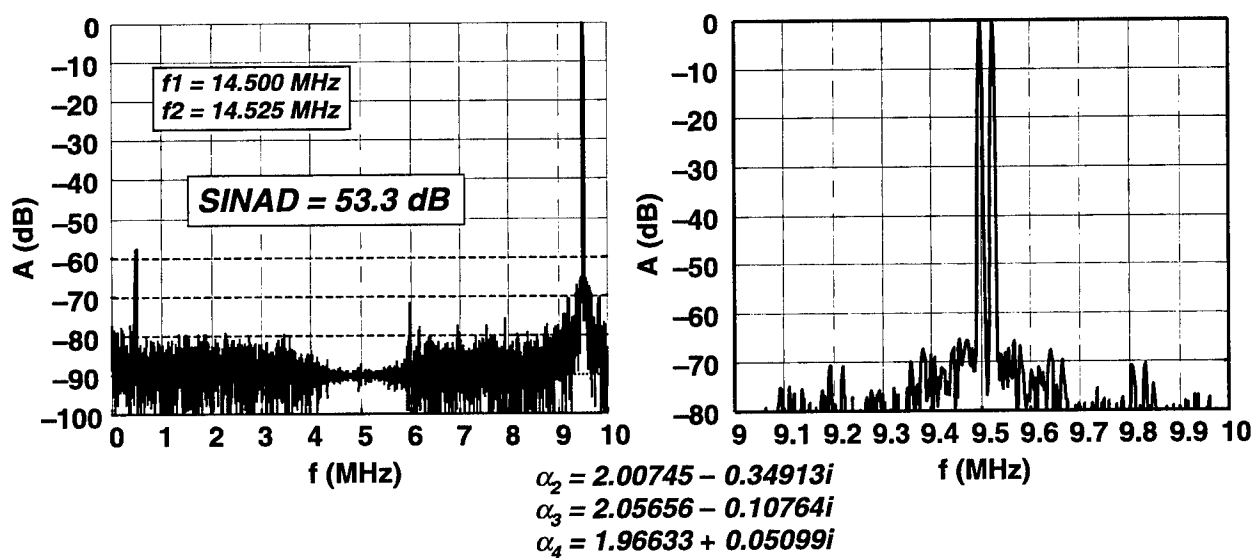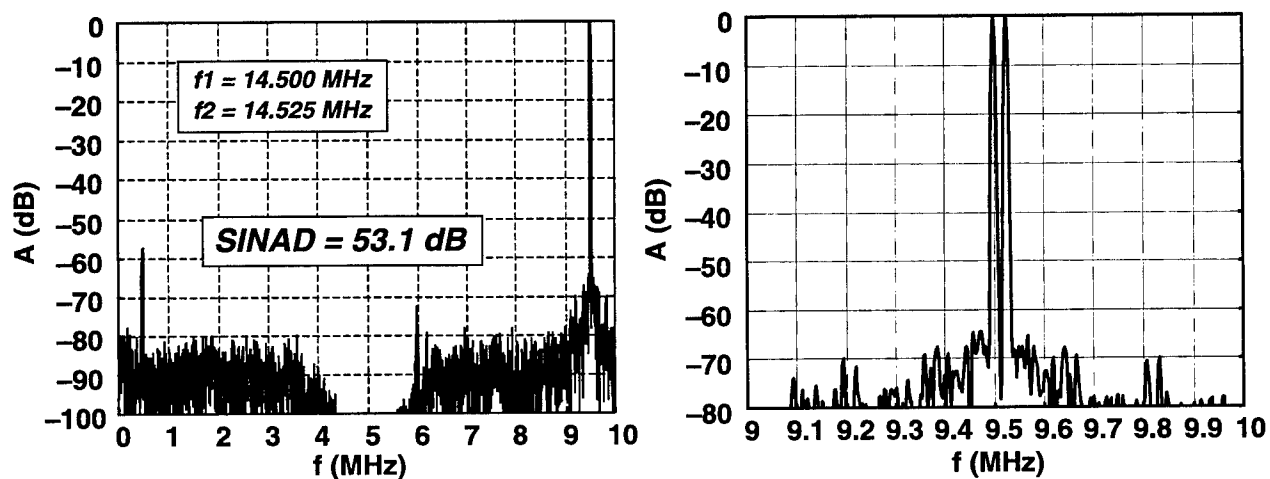Fig. 14 — Spectrum of combined output with no amplitude and phase correction applied

$$\alpha_2 = 2.00745 - 0.34913i$$
$$\alpha_3 = 2.05656 - 0.10764i$$
$$\alpha_4 = 1.96633 + 0.05099i$$

Fig. 15 — Spectrum of final output after gain and phase calibration

Fig. 16 — Spectrum at output of least-sensitive (G = 0 dB) channel only

$$\alpha_2 = 2.00745 - 0.34913i$$
$$\alpha_3 = 2.05656 - 0.10764i$$
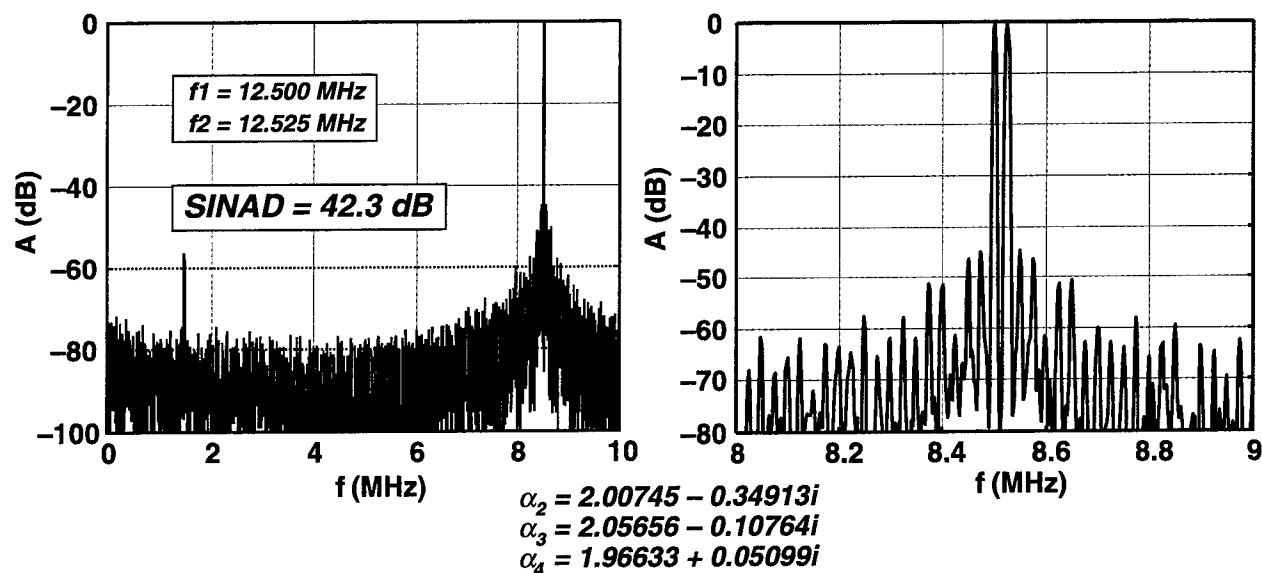$$\alpha_4 = 1.96633 + 0.05099i$$

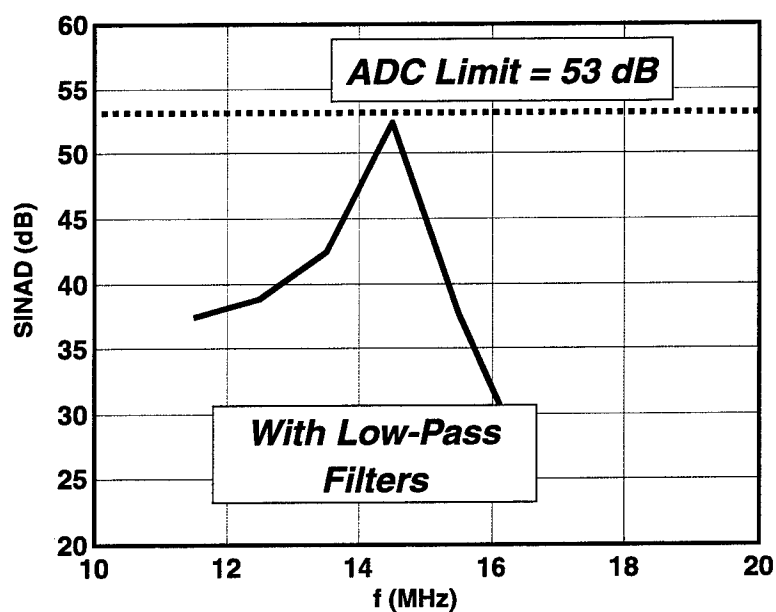Fig. 17 — Spectrum of final output after gain and phase calibration at 14.5 MHz

Fig. 18 — SINAD vs input frequency of two-tone signal for experimental
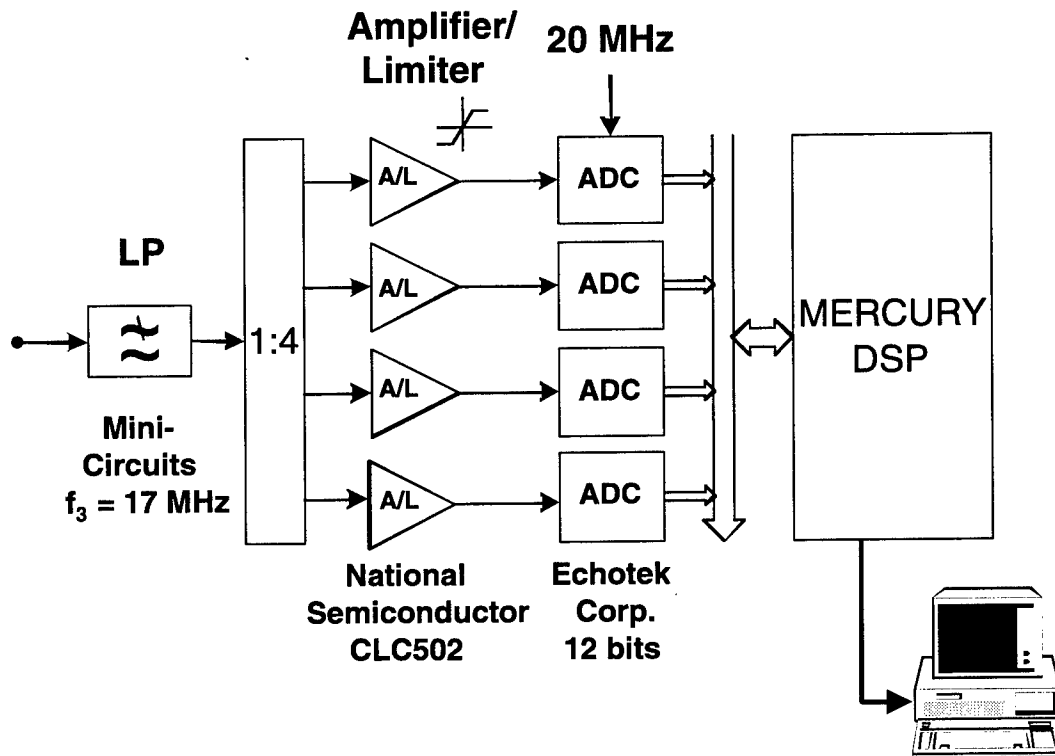Stacked ADC. Calibration is at 14.5 MHz only.

Expected Performance Limitation ~50 dB

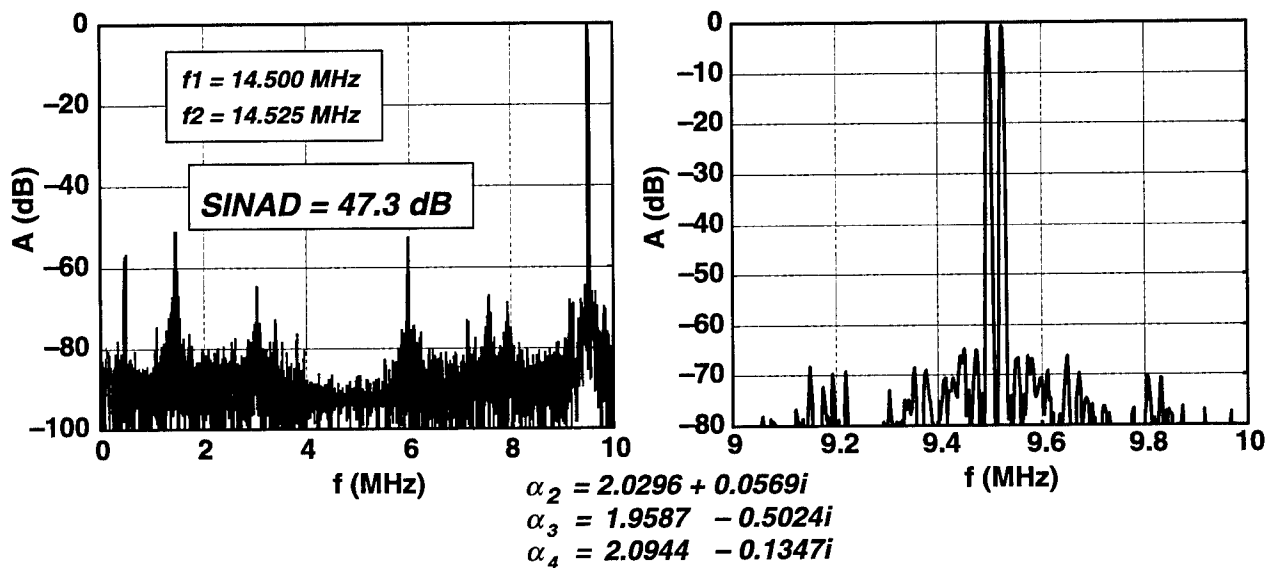Fig. 19 — Block diagram of modified experimental Stacked ADC implementation

$\alpha_2 = 2.0296 + 0.0569i$
$\alpha_3 = 1.9587 \ - 0.5024i$
$\alpha_4 = 2.0944 \ - 0.1347i$

Fig. 20 — Spectrum of final output of modified Stacked ADC after gain and phase calibration

$$\alpha_2 = 2.0296 + 0.0569i$$
$$\alpha_3 = 1.9587 - 0.5024i$$
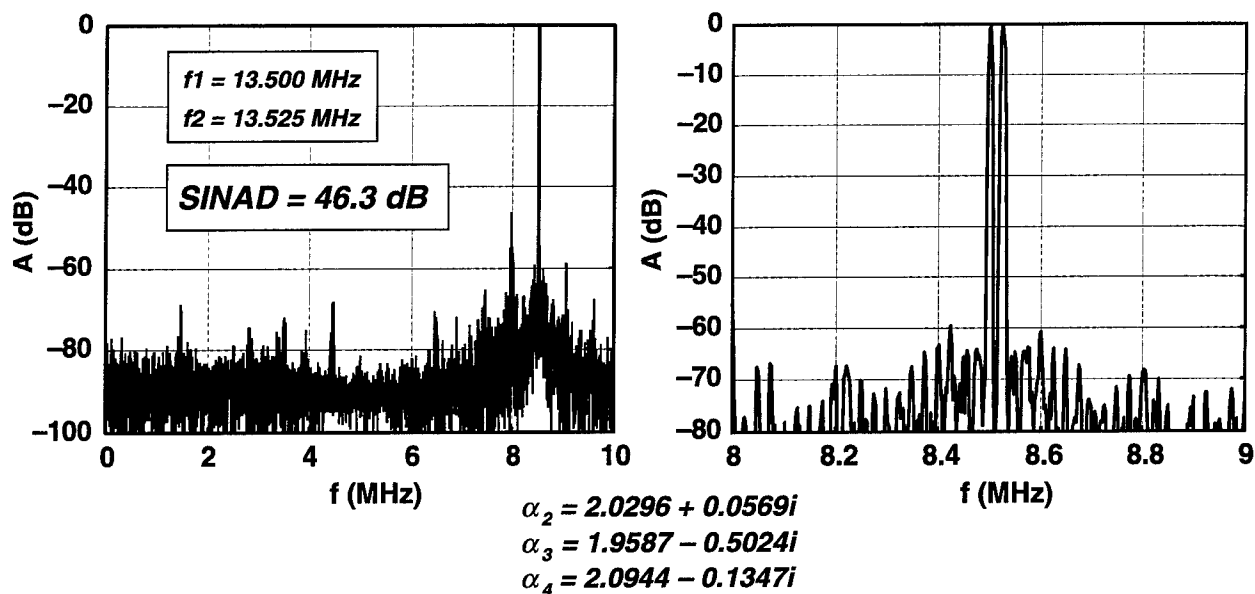$$\alpha_4 = 2.0944 - 0.1347i$$

Fig. 21 — Spectrum of final output of modified Stacked ADC
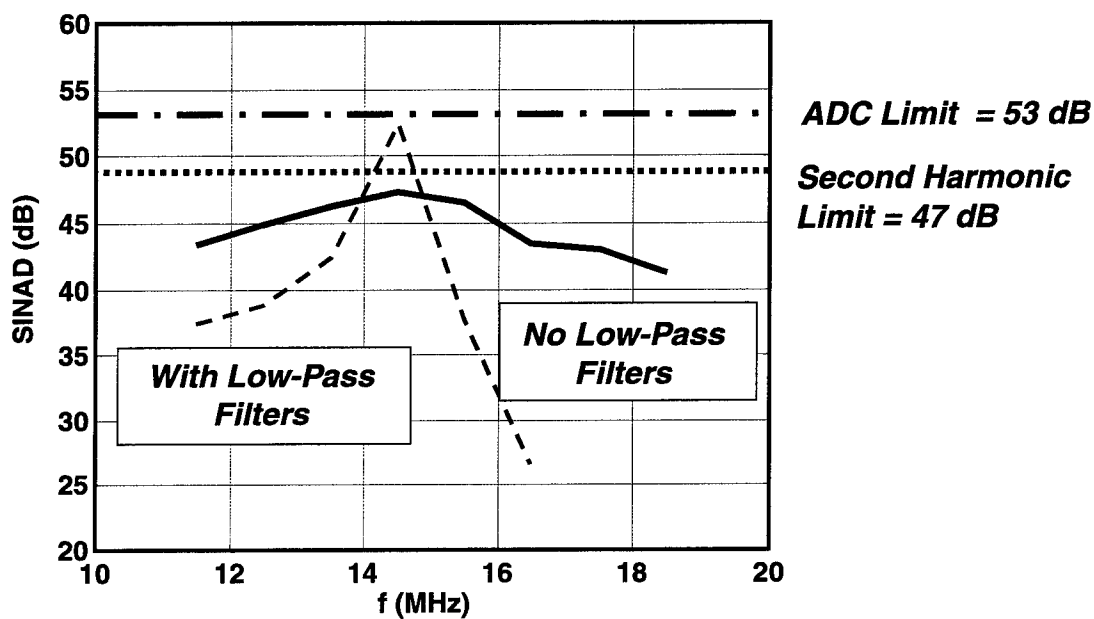after gain and phase calibration at 14.5 MHz

Fig. 22 — SINAD vs input frequency of two-tone signal for modified experimental
Stacked ADC. Calibration was performed at 14.5 MHz.

CONCLUSIONS

The performance realized with the experimental Stacked ADC described in this report has demonstrated the feasibility of increasing the effective dynamic range in a practical radar signal processor from 60 to 65 dB to 78 to 83 dB while maintaining a SINAD of 45 to 50 dB against a two-tone input signal. Major sources of this SINAD degradation, relative to that specified for individual ADCs, are the harmonics and intermodulation products at the output of the linear/limiting amplifier, which must be used prior to the ADC to prevent an excessive voltage from being applied at its input. It is suspected that many current implementations of the ADC interface in radar systems may have similar restrictions on SINAD for two-tone inputs. Also noteworthy is the fact that the results reported here assumed the most demanding Stacked ADC implementation, where ADC outputs are combined on a range-cell-by-range-cell basis. It is also expected that the suppression of strong but stable clutter at zero Doppler would be much closer to the ADC SINAD limit. An improved version of the Stacked ADC, using the latest Analog Devices 14-bit converters and better linear/limiting amplifiers, is currently under construction.

ACKNOWLEDGEMENT

REFERENCES

1. H.R. Ward and W.W. Shrader, "MTI Performance Degradation Caused by Limiting," IEEE EASCON Record (1968), pp. 168-174.

2. L.E. Pellon, "A Double Nyquist Digital Product Detector for Quadrature Sampling," *IEEE Trans. Signal Process.* **40**(7), 1670-1681 (1992).

Appendix

MATLAB PROGRAM LISTING

```
% Stacked A/D Processing Emulation
% A/D Direct Sampling at 20 MHz
% Input tones at equal amplitudes at 15.475 MHz and 15.5 MHz and at
%   offsets which are multiples of 1 MHz
% Vilhelm Gregers-Hansen
% 25 April 2000

clear all;

% Read data files recorded at A/D outputs
DAT1 = load('e:\matldat\stacked atod\Runs-05-05-00\AE005_1');
DAT2 = load('e:\matldat\stacked atod\Runs-05-05-00\AE005_2');

% Specify lowest frequency of input pair to program
fc=14.5;

% Find location of second harmonic
if fc<15

        fsecond=mod(-2*fc,20);
    fsecond=fsecond-15;
    fs2=mod(fsecond,10)+0.025
    fs2=mod(fs2,10)
    fs1=fs2-0.1
    fs1=mod(fs1,10)
else
    fsecond=mod(2*fc,20);
    fsecond=fsecond-15;
    fs1=mod(fsecond,10)-0.025
    fs1=mod(fs1,10)
    fs2=fs1+0.1
    fs2=mod(fs2,10)
end

fc=fc-15;

fb=floor(10+fc);
fb=mod(fb,10);
fe=fb+1;
NPts=length(DAT1);

%Define time scale
```

```
 fs=20e6;                     % A/D Sampling Frequency of 20 MHz
 t=(1:NPts)-1;
 t=t/fs;

%Separate into individual A/D channels

y1=DAT1(:,2);    % For data recorded 05-04-00 make this col. 3
y2=DAT1(:,3);
y3=DAT2(:,2);
y4=DAT2(:,3);

% Plot A/D data as specified
figure(10);
plot(y1(1:2000));
grid
axis([ 0 2000 -40000 40000]);
pause
figure(11);
plot(y2(1:2000));
grid
axis([ 0 2000 -40000 40000]);
pause
figure(12);
plot(y3(1:2000));
grid
axis([ 0 2000 -40000 40000]);
pause
figure(13);
plot(y4(1:2000));
grid
axis([ 0 2000 -40000 40000]);
pause
% Define A/D limit level (<32768 due to DC offset)
YLim=32765;

% Number of complex samples
i2=1:NPts/2;

% Create vectors of ones where A/D "limits"
y1lim=y1>YLim | y1<-YLim;
YL1=y1lim(2*i2) | y1lim(2*i2-1);
y2lim=y2>YLim | y2<-YLim;
YL2=y2lim(2*i2) | y2lim(2*i2-1);
y3lim=y3>YLim | y3<-YLim;
YL3=y3lim(2*i2) | y3lim(2*i2-1);
y4lim=y4>YLim | y4<-YLim;
YL4=y4lim(2*i2) | y4lim(2*i2-1);
```

```
% Direct Conversion FIR filter (D(V) & LSTP)


hi = [ 11 11 -109 272 -362 119 745 -2507 6545 12092 171 -1273 1103 ...
          -586 151 53 -75 32] /16393;
hq=fliplr(hi);

FIRLength=length(hi);

% Send A/D data through FIR filters
z1=dpd(y1,hi,hq);
z2=dpd(y2,hi,hq);
z3=dpd(y3,hi,hq);
z4=dpd(y4,hi,hq);

LTot=length(z2);

% Select valid (non-transient) data
vec1=z1(FIRLength:LTot);
vec2=z2(FIRLength:LTot);
vec3=z3(FIRLength:LTot);
vec4=z4(FIRLength:LTot);

L=length(vec2);

t2=2*(1:L)/fs;          % Time scale for output vectors

%Check magnitude of I,Q

mag1=abs(vec1);
mag2=abs(vec2);
mag3=abs(vec3);
mag4=abs(vec4);

figure(20);
hold off
plot(8*mag1(1:1000));
hold on
pause
%figure(21);
plot(4*mag2(1:1000));
pause
%figure(22);
plot(2*mag3(1:1000));
pause
%figure(23);
plot(mag4(1:1000));
grid
```

```
pause

% Identify bad data
% Define unit amplitude filter of same length as DDC FIR
href=ones(1,FIRLength);

% Do convolution
y1b=conv(href,YL1);
y2b=conv(href,YL2);
y3b=conv(href,YL3);
y4b=conv(href,YL4);

% Select bad data vector matching valid A/D data
y1bad=y1b(FIRLength:LTot)>0;
y2bad=y2b(FIRLength:LTot)>0;
y3bad=y3b(FIRLength:LTot)>0;
y4bad=y4b(FIRLength:LTot)>0;

%Calibration

% Find valid data for calibration
% Not bad (limiting) but greater than half scale
MMag=25000
set2=find(mag2 > MMag/2 & y2bad==0);
set3=find(mag3 > MMag/2 & y3bad==0);
set4=find(mag4 > MMag/2 & y4bad==0);

% If data found compute calibration constant otherwise use
%      constant
if length(set2)>1
      alpha21=vec2(set2)./vec1(set2);
      alpha2=mean(alpha21)
else
    alpha2=2;
end

if length(set3)>1
    alpha32=vec3(set3)./vec2(set3);
    alpha3=mean(alpha32)
else
    alpha3=2;
end

if length(set4)>1
    alpha43=vec4(set4)./vec3(set4);
    alpha4=mean(alpha43)
else
    alpha4=2;
```

```
end

disp([ abs(alpha2) angle(alpha2)*180/pi] );
disp([ abs(alpha3) angle(alpha3)*180/pi] );
disp([ abs(alpha4) angle(alpha4)*180/pi] );
% Change calibration constants if required

% Calibration based on data files AD011_1 and _2 (4 May 2000)
%alpha2= 2.0075 - 0.3491i;
%alpha3=2.0566 - 0.1076i;
%alpha4=1.9663 + 0.0510i;

% From data file AE005_1 and _2 (5 May 2000)
alpha2 =2.0296 + 0.0569i;
alpha3 =1.9587 - 0.5024i;
alpha4 =2.0944 - 0.1347i;

%alpha2=2.0;
%alpha3=2.0;
%alpha4=2.0;

% Disable an A/D channel by declaring all samples bad
y4bad=ones(L,1);        % never use most sensitive A/D
y3bad=ones(L,1);
y2bad=ones(L,1);

% Data to be used comes from most sensitive A/D not in limiting
pick1=y1bad==0 & y2bad==1;
pick2=y2bad==0 & y3bad==1;
pick3=y3bad==0 & y4bad==1;
pick4=y4bad==0;

vecfinal=zeros(L,1);
%vec1=0;            % No data from channel 1

vecf1=alpha4*alpha3*alpha2*vec1.*pick1;
figure(30)
hold off
plot(abs(vecf1(1:1000)))
hold on
pause

vecf2=alpha3*alpha4*vec2.*pick2;
plot(abs(vecf2(1:1000)))

hold on
pause
vecf3=alpha4*vec3.*pick3;
```

```
plot(abs(vecf3(1:1000)))
pause
vecf4=vec4.*pick4;
plot(abs(vecf4(1:1000)))
pause

vecfinal=vecf1+vecf2+vecf3+vecf4;

figure(1);
plot(real(vecfinal(1:300)));
hold on
plot(imag(vecfinal(1:300)));
%axis([ 1 300 -1000 1000]);
grid on;
pause

hold off

figure(7);
vecwgt=vecfinal.*chebwin(4000,110);
vecsp=fft(vecwgt' ,4096);
x=10*((1:4096)-1)/4096;
vdb=20*log10(abs(vecsp));
%vdb=fliplr(vdb);
vdbmax=max(vdb);
plot(x,vdb-vdbmax);
axis([ 0 10 -100 0])
grid
pause
figure(3)
plot(x,vdb-vdbmax);
axis([ fb   fe -80 0])
grid
pause

% Finally do SINAD calculation
% Convert to power spectrum
PSpec=abs(vecsp).^2;

f1=fb+0.45;                                    % Signal start
Nf1=round(4096*f1/10);
f2=fb+0.55;                                    % Signal end
```

```
Nf2=round(4096*f2/10);
PSig=sum(PSpec(Nf1:Nf2));  % Signal power

PSpec(Nf1:Nf2)=0.0;`                    % Exclude signal

Ptot=sum(PSpec(1:4096));
SINAD_All=10*log10(PSig/Ptot)

f3=10-f2;
Nf3=round(4096*f3/10);
f4=10-f1;
Nf4=round(4096*f4/10);
PSpec(Nf3:Nf4)=0.0;                     % Exclude image

Ptot=sum(PSpec(1:4096));
SINAD_NoImag=10*log10(PSig/Ptot)

Nf5=round(4096*fs1/10);
Nf6=round(4096*fs2/10);

if Nf5<Nf6
    PSpec(Nf5:Nf6)=0.0;                 % Exclude second harmonic
else
    PSpec(Nf5:4096)=0.0;
    PSpec(1:Nf6)=0.0;
end


Ptot=sum(PSpec(1:4096));
SINAD_NoIm_No_2ndH=10*log10(PSig/Ptot)


%figure(4)
veccens=vecsp;
vecsp(Nf1:Nf2)=0;
vecsp(Nf3:Nf4)=0;
vectime=ifft(veccens,4096);
%plot(20*log10(abs(vectime)));
```