Implementing Split Spectrum Processing with the TMS320C26 DSP

APPLICATION REPORT: SPRA197

Author: Scott E. Bailey
Advisor: Dr. T. L. Stewart
Electrical Engineering Department
Bradley University
Peoria, Illinois

Digital Signal Processing Solutions
June 1997

TEXAS INSTRUMENTS

IMPORTANT NOTICE

Texas Instruments (TI) reserves the right to make changes to its products or to discontinue any semiconductor product or service without notice, and advises its customers to obtain the latest version of relevant information to verify, before placing orders, that the information being relied on is current.

TI warrants performance of its semiconductor products and related software to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are utilized to the extent TI deems necessary to support this warranty. Specific testing of all parameters of each device is not necessarily performed, except those mandated by government requirements.

Certain application using semiconductor products may involve potential risks of death, personal injury, or severe property or environmental damage ("Critical Applications").

TI SEMICONDUCTOR PRODUCTS ARE NOT DESIGNED, INTENDED, AUTHORIZED, OR WARRANTED TO BE SUITABLE FOR USE IN LIFE-SUPPORT APPLICATIONS, DEVICES OR SYSTEMS OR OTHER CRITICAL APPLICATIONS.

Inclusion of TI products in such applications is understood to be fully at the risk of the customer. Use of TI products in such applications requires the written approval of an appropriate TI officer. Questions concerning potential risk applications should be directed to TI through a local SC sales office.

In order to minimize risks associated with the customer's applications, adequate design and operating safeguards should be provided by the customer to minimize inherent or procedural hazards.

TI assumes no liability for applications assistance, customer product design, software performance, or infringement of patents or services described herein. Nor does TI warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right of TI covering or relating to any combination, machine, or process in which such semiconductor products or services might be or are used.

Copyright © 1997, Texas Instruments Incorporated

TRADEMARKS

TI is a trademark of Texas Instruments Incorporated.

Other brands and names are the property of their respective owners.

Contents

Figures

Tables

Implementing Split Spectrum Processing with the TMS320C26 DSP

Abstract

The ultrasonic evaluation of specimens with large microstructures can be difficult because of interference caused by the grains in the material. The split spectrum processing (SSP) filtering technique effectively reduces this type of noise by employing a bank of filters followed by a nonlinear detector that increases signal-to-noise ratios.

This application report describes the benefits of implementing SSP using the Texas Instruments (TI™) TMS320C26 digital signal processor (DSP). The high performance capability of the TMS320C26 DSP offers a quick and efficient means of conducting SSP, which means researchers can now freely experiment with SSP filter parameters and obtain almost instantaneous results.

In addition, previously time-consuming filtering applications (such as MATLAB simulations and two- and three-dimensional SSP) are now more practical because they can be accomplished much faster.

Ultrasonic Non-Destructive Evaluation and Testing

The process of ultrasonic nondestructive evaluation uses high frequency acoustic pulses to identify the characteristics of a particular test specimen. A broadband ultrasonic transducer produces the acoustic pulses at the resonant frequency of the transducer crystal. Figure 1 illustrates a typical ultrasonic test configuration.

Figure 1. Ultrasonic Test Configuration

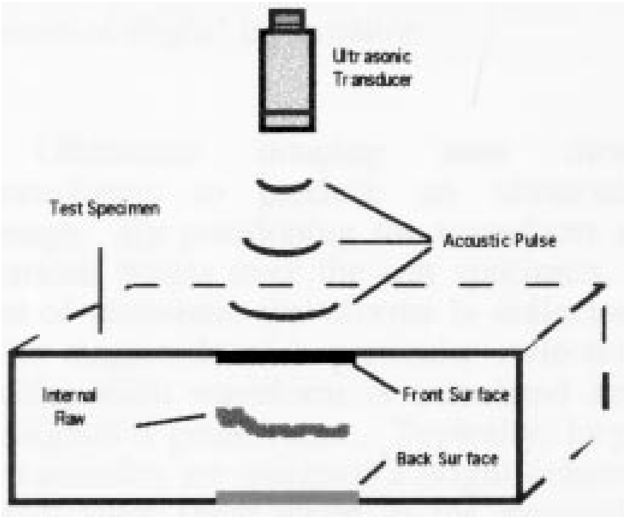

The transducer must be coupled to the specimen to provide a path through which the pulse can travel. Coupling is accomplished by immersing the transducer and specimen in water, which serves as the couplant.

Each time an interface is encountered as the acoustic wave propagates through the water, part of the signal reflects the object and part of the signal penetrates the object. As the pulse travels through the object, reflections occur at each interface until the pulse reaches the back surface, at which point the last reflection occurs. The result is a collection of pulse reflections similar to that shown in Figure 2.

The first interface the pulse reaches is known as the front surface. Reflections that follow are produced by any internal interfaces. Internal interfaces are caused by cracks, disbonds, pores, or other flaws. If the energy of the pulse is not completely dissipated as it travels through the test specimen, a final reflection will be produced upon reaching the back surface. The reflections are picked up by the transducer, which not only initiates the original pulse, but acts as a receiver as well. Each type of reflection is represented in Figure 2.

Figure 2. Ultrasonic Pulse Reflections

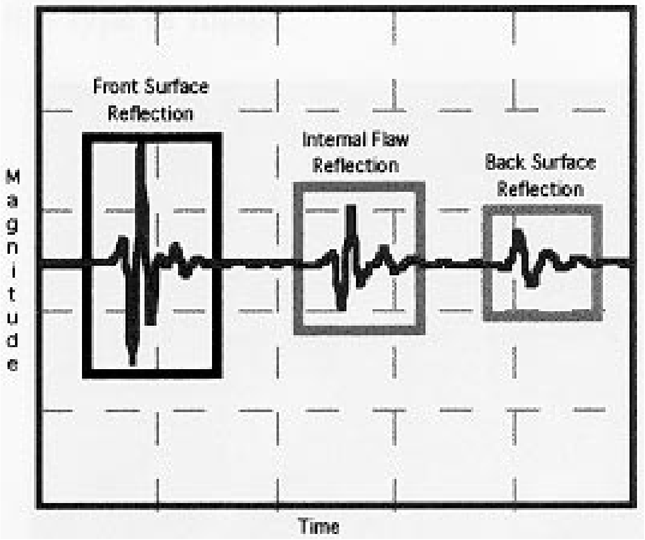

Analyzing ultrasonic waveforms similar to the ones shown in Figure 2 results in the identification of many test specimen characteristics. Surface characteristics can be determined by monitoring the magnitude of the front surface reflection. A large magnitude indicates a flat smooth surface. Small magnitudes indicate a rough or angled front surface caused by scattering of the acoustic pulse.

Similarly, monitoring the magnitude of an internal reflection indicates the significance of a flaw. A large magnitude results from flaws of large size; a small magnitude indicates a small flaw.

Specimen thickness can be determined by measuring the amount of time it takes for the waveform to pass from the front surface to the back surface of a material with a known thickness. The velocity of sound for the material is then calculated. The thickness can be calculated using this value and the time interval required for the waveform to pass from the front to the back surface at an unknown thickness. This method of calculating specimen thickness is useful in situations where material thickness is not easily determined by conventional measurement techniques.

The depth of a flaw within a specimen can be determined based on the *time of flight*, which is the amount of time between the front surface reflection and an internal flaw reflection and the known velocity of sound.

Ultrasonic imaging uses these waveforms to produce an ultrasonic image. Positioning the transducer at various points over the test specimen provides a set of ultrasonic waveforms. The magnitude of a particular reflection within each waveform is calculated and assigned a pixel value. Typically, large magnitudes are assigned a bright colored pixel value (near white in the grayscale case), and small magnitudes are assigned a dark colored pixel value (near black in the grayscale case). Arranging these pixel values in software results in an ultrasonic image. Figure 3 shows an ultrasonic image in which the magnitude of the front surface reflection was monitored.

Figure 3. Front Surface Ultrasonic Image

Bright areas indicate a large magnitude resulting from a relatively smooth and flat front surface. The dark areas are the result of rough, uneven surfaces that scatter the ultrasonic pulse, causing reflections of small magnitude.

Figure 4 shows an image of a failed integrated circuit. By examining this specimen and comparing it to ultrasonic images of chips that have not failed, trends and consistencies can be determined, resulting in a method of identifying potentially defective specimens.

Figure 4. Ultrasonic Image of Integrated Circuit

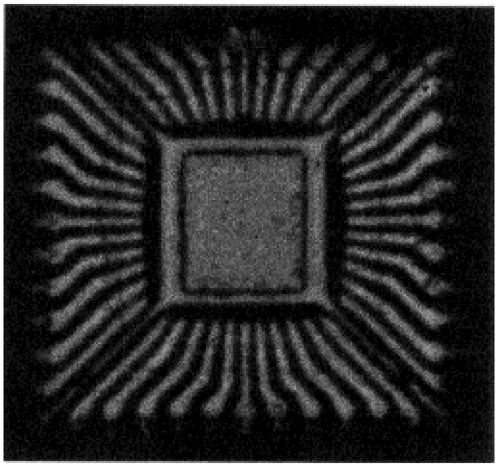

Many types of materials can be evaluated ultrasonically to determine material characteristics without changing the physical or chemical composition. This non-destructive approach offers many benefits, including time and cost savings.

Split Spectrum Processing

Recent improvements in hardware technology offer many more applications of ultrasonic evaluation. Transducers operate at higher frequencies with smaller focal spot sizes so that even very small flaws can be found. However, when transducers with these increased capabilities are used on materials with a large microstructure, such as cast iron, ceramics, or porous materials, the ultrasonic pulse is reflected not only by the flaws of the material but by the pores, or *grains*, as well. Figure 5 illustrates the effects of a specimen with a large microstructure.

Figure 5. Specimen with a Large Microstructure

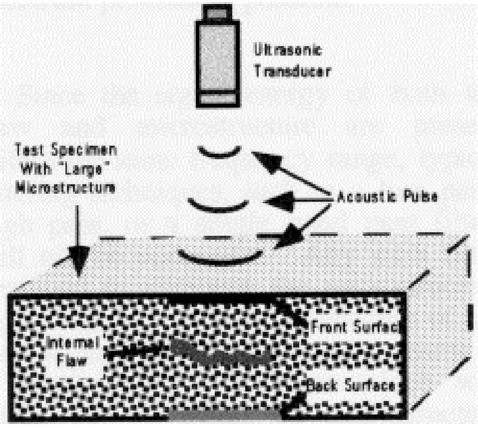

Microstructure reflections can interfere with the reflections from the flaws and back surface. This *microstructural noise* makes the imaging of small flaws difficult, if not impossible. Figure 6 shows an example of an ultrasonic waveform from a specimen with a large microstructure. As seen in this ultrasonic waveform of a cast iron specimen, the grain reflections completely obscure any internal flaw reflection as well as the back surface reflection. SSP provides a way to filter out unwanted microstructural noise.

Figure 6. *Microstructure Reflection*

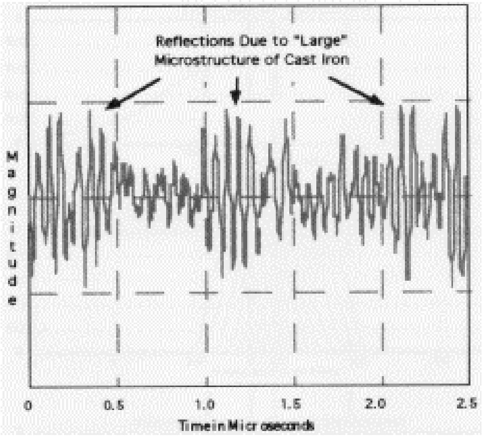

Converting the ultrasonic reflection of the flaw shown in Figure 2 into the frequency domain yields the top signal spectrum plot shown in Figure 7.

Figure 7. *Ultrasonic Signal Spectra*

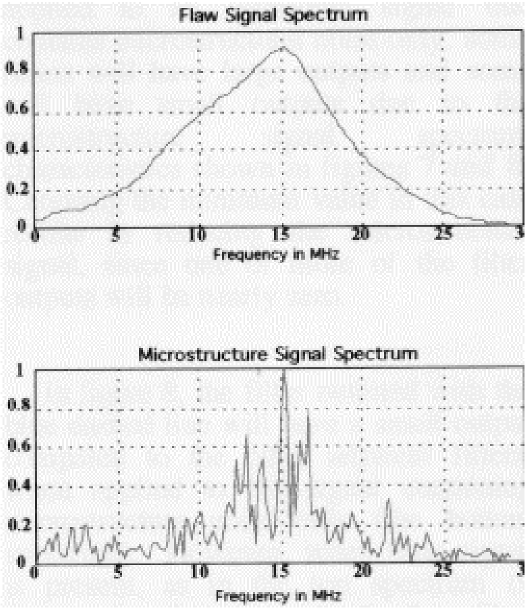

As the top graph of Figure 7 shows, the signal energy of the flaw reflection is distributed in a finite frequency band centered around the resonant frequency of the transducer (15 MHz in this case). Most of the signal energy is relatively evenly distributed within an 8- to 10-MHz band.

Analysis of the signal spectrum of reflections caused by microstructure shows that most of the signal energy lies within the same frequency band as the flaw signal energy. However, as shown in the bottom graph of Figure 7, the energy is not evenly distributed but instead exhibits voids or nulls within the 8- to 10-MHz frequency band. This characteristic of microstructure signal spectrum makes SSP possible.

Since the signal energy of both the flaw and the microstructure occurs within the same frequency range, typical filtering techniques (such as a low pass, high pass, or a single band pass filter) are not successful. Any such filter designed to eliminate the microstructure signal also eliminates the flaw signal. SSP offers a signal filtering technique that allows the reduction of microstructure signal while retaining most of the flaw signal.

The SSP method evenly divides the frequency band containing most of the flaw signal energy into a number of smaller frequency bands. A bank of bandpass filters can represent these frequency bands. Applying the bank of filters to an ultrasonic signal and selecting the minimum output at each point in time reduces the microstructure signal, leaving only any flaw reflections that may be present. Figure 8 illustrates this process.

If the flaw signal is present, applying the bank of filters results in filter outputs that all have significant and relatively equal values. Choosing the minimum of several nearly equal output values causes little flaw signal loss. When the filter bank is applied to an ultrasonic signal containing microstructural noise only, some filters will have large outputs and some will have small outputs. This is the result of the microstructure signal spectrum characteristics shown in Figure 7 and Figure 8. In this case, choosing the minimum value reduces the microstructure signal because one or more of the filter outputs is nearly zero.

Figure 8. Split Spectrum in Frequency Domain

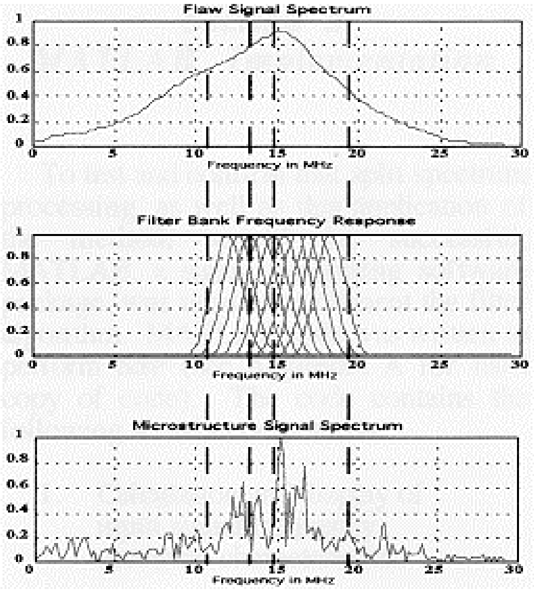

The filter outlined with the blue dashed line in Figure 8 will have a small output (compared to the other adjacent filters) when applied to the signal containing microstructure noise only (the bottom spectrum). However, when a flaw signal is present, as in the top spectrum of Figure 8, the output of the filter highlighted in blue is relatively equal to the other adjacent filters. Therefore, choosing the minimum output of the bank of filters at each point in time reduces the microstructure signal but retains the flaw signal.

Figure 9 shows an SSP block diagram. The input *x(n)* is presented to the bank of filters. At each point in time, the corresponding input sample is applied to each filter. The *y'(n)* values represent the individual filter outputs. Once the filter outputs are calculated for the current input value, the minimum *y'(n)* is selected, which then becomes the final output value, *y(n)*, for the corresponding point in time. The process is repeated for each input sample and results in the split spectrum processed signal, which should contain flaw signal reflections only.

Figure 9. SSP Block Diagram

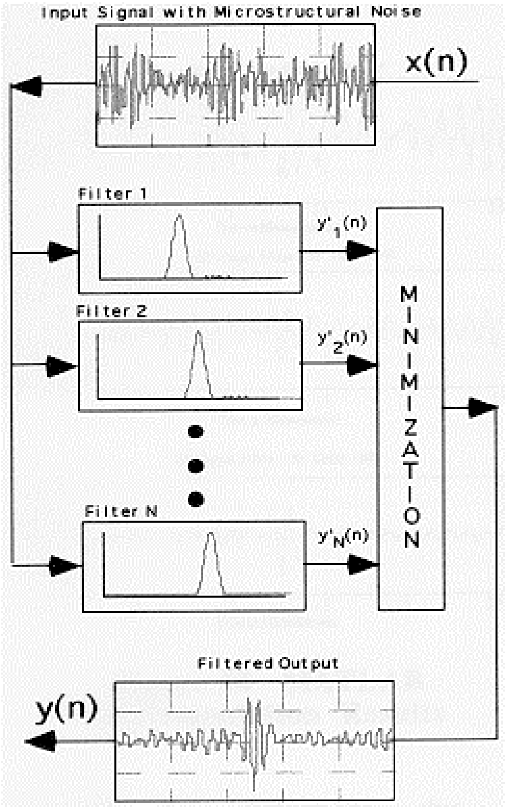

MATLAB Implementation

MATLAB is a signal processing software package used to implement the filter algorithm required to test and confirm the successful application of SSP. MATLAB code was written to perform SSP and contains the following four main components (Appendix A provides an example of the MATLAB code):

- ☐ Calculation and display of input signal frequency spectrum characteristics

  First, the input signal spectrum characteristics are calculated and displayed to determine the frequency range in which most of the signal energy is located.

  For this stage, an input signal containing a front surface reflection produced by the same transducer that generated the input data is analyzed in the frequency domain. This analysis allows the user to select the proper frequency band in which to design the bank of bandpass filters.

- ☐ User input of filter bank parameters and calculation of filter coefficients; display of filter frequency response

  Second, the user specifies the various filter parameters used to create the bank of bandpass filters, including:

  - Type of bandpass filters
  - Number of filters
  - Order of the filters
  - Overall bandwidth
  - Individual bandwidth

  MATLAB uses these filter parameters to create the filter coefficients, then displays their frequency response so the user can verify the design. Initially, finite impulse response (FIR) filters were chosen for this project but infinite impulse response (IIR) filters were used in the DSP implementation (see Section 0 for more information).

- ☐ Calculation of individual filter outputs

  Third, the input signal is applied to each filter and individual filter outputs are calculated.

- ☐ Minimization and display of final output

  Fourth, the minimum filter output at each point in time (minimization) is chosen and the final output is displayed.

Figure 10 shows the results of the MATLAB code for a cast iron sample containing a seeded defect (that is, a flat-bottomed hole is drilled from the back surface). The top graph of Figure 10 represents the original ultrasonic waveform, which is the input. The flaw reflection is completely obscured by reflections from the microstructure.

Figure 10. MATLAB Implementation Results

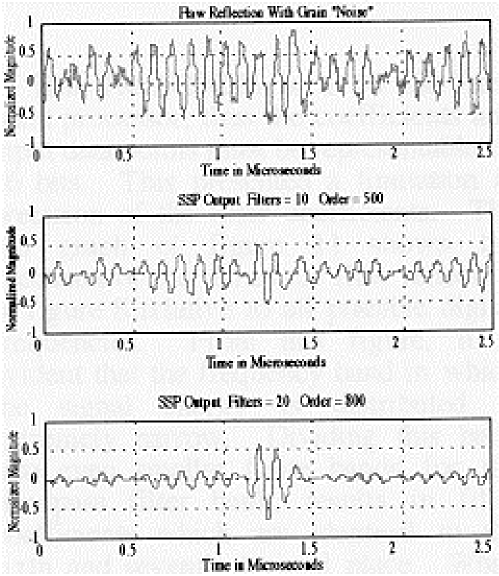

The middle graph of Figure 10 shows the SSP output using 10 FIR filters of order 500. Although microstructural noise is reduced, the flaw reflection is still difficult to isolate.

The bottom graph of Figure 10 shows the results of SSP using 20 FIR filters of order 800. The flaw reflection is now clearly identifiable. Because the defect is seeded, the depth of the flaw within the specimen and the location of the flaw within the waveform are known. The output plots of Figure 10 correspond with this known flaw location.

TMS320C26 DSP Implementation

Implementing an SSP application using the TMS320C26 DSP ensures that the application will benefit from the speed and efficiency with which the DSP calculates filter difference equations. In a single instruction cycle, the typical DSP can perform a multiply, accumulate the previous product, move data, and modify auxiliary registers. Many DSPs are capable of performing 10 to 50 million instructions per second. These high performance calculations offer an ideal way to conduct SSP.

Increasing the speed of SSP facilitates applications that would be impractical otherwise. The MATLAB simulations discussed in Section 0 took from 1 to 5 minutes to perform. Using two- and three-dimensional SSP (in which an array or matrix of waveforms is processed rather than a single waveform) can increase simulation time to 5 hours. Because the DSP significantly reduces this computation time, such applications become more practical. In addition, you can adjust the filter parameters and experiment with various configurations without wasting time waiting for the results.

This project used the TI TMS320C26 DSP Starter Kit to perform SSP on ultrasonically acquired data. One goal of this project was to adhere to on-chip memory and programming constraints to minimize interfacing complexities and cost.

TheTMS320C26 DSP starter kit includes the following components:

- ☐ On-board power supply
- ☐ RS-232 interface
- ☐ Analog interface chip (AIC)

Since the AIC is incapable of the sampling rates needed for ultrasonic data, the input data is post-processed (that is, the data is sampled on external hardware and then loaded into DSP memory). However, the AIC is used to output the results of SSP to an oscilloscope through the D/A (digital-to-analog) portion of the AIC.

The TMS320C26 is a 16-bit, fixed-point processor; thus, the coefficients and input data points must be capable of being represented by 16 bits, which limits the precision of the filter coefficients. The top graph of the ultrasonic flaw reflection shown in Figure 8 is relative to all possible digital frequencies. As shown in Figure 8, the frequency band in which the signal energy is distributed is extremely narrow. Dividing this band into even smaller finite bands for the bandpass filter bank results in filter coefficients that are identical to the sixth and seventh decimal place. When these coefficients are represented with 16 bits, no difference is distinguished from one filter to the next. As a result, all filter outputs are the same and the spectrum is not split.

To account for this limitation, the input data must be down-sampled. This was accomplished for this project by low pass filtering the input data, then selecting every *n* samples, where *n* is the down-sampling rate. The low pass filter must have a gain of 1 and a cutoff of *pi/n*.

The bottom plot in Figure 11 shows the results of down-sampling the input data by 5. The center frequency of the signal energy is the same, but now the bandwidth of the signal energy is larger compared to the possible digital frequencies. Using these bandwidth characteristics, the filter coefficients can be more accurately represented with 16 bits.

Figure 11. Digital Frequency Spectrum

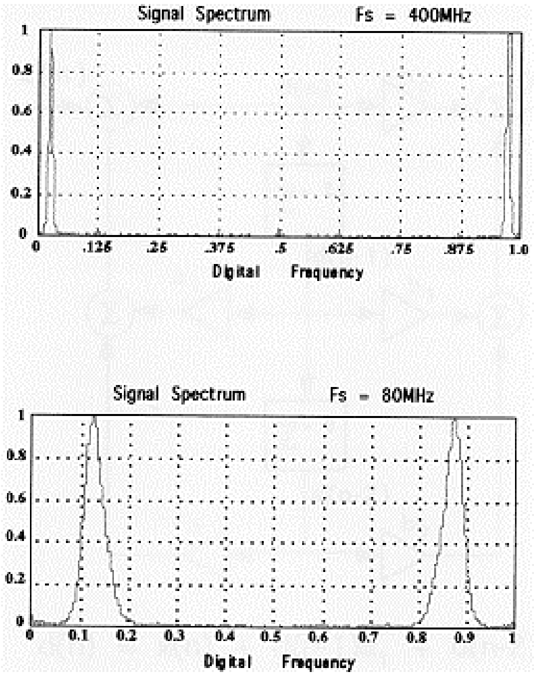

The adherence to on-chip memory constraints also resulted in limitations on the type of filter bank parameters that could be selected. The DSP has three blocks of 512 words (16 bits) of memory. Some must be reserved for program memory, leaving the rest for input data and filter coefficients. Using the CONF 1 instruction, block B0 is configured as program memory, leaving 1024 words of data in blocks B1 and B3 for data RAM. An additional 32 words of data RAM are also available in block B2. Figure 12 shows the memory map for the TMS320C26 after the CONF 1 instruction is executed.

Figure 12. TMS320C26 Memory Map

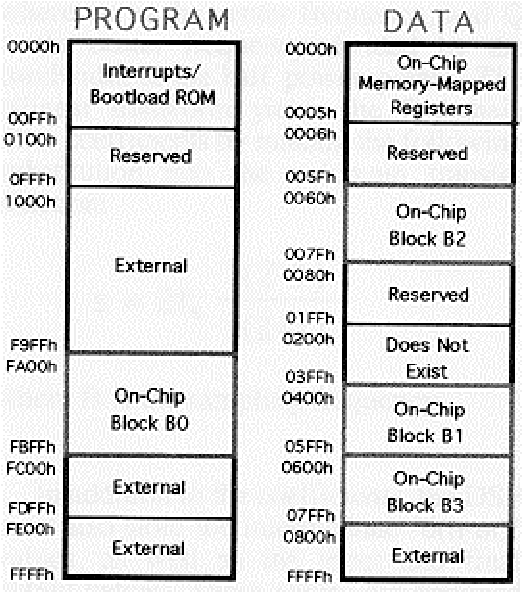

As mentioned in Section 0, FIR filters were used initially to implement the filter bank. However, the high order (500th to 800th or higher) needed to create the very narrow band FIR filters meant that the available on-chip data RAM was insufficient to accommodate the large number of filter coefficients. Even after the data was down sampled by 5, the FIR filters still needed to be at least 100th order.

Typically, SSP requires at least 10 filters in the filter bank. This requirement results in 1000 or more FIR filter coefficients – still too many coefficients. Additional down sampling caused an excessive loss of information from the original signal. IIR filters were designed in response to this limitation.

After downsampling the input data by 5, it was determined that a second-order IIR filter design would result in a bank of filters with satisfactory bandwidth characteristics. Figure 13 shows the Direct Form II block diagram representation of a second-order IIR filter.

Figure 13. Second Order IIR Filter Block Diagram

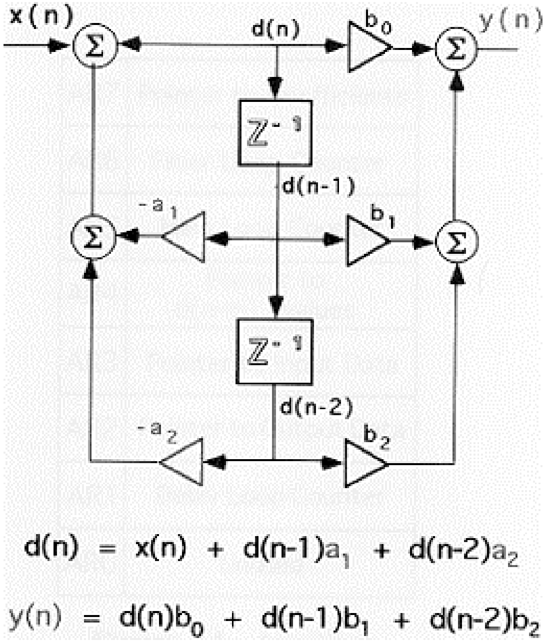

Each IIR filter within the filter bank has five coefficients (b1 is actually zero because the filter is a symmetric bandpass filter). The DSP memory easily accommodates this number of filter coefficients.

The coefficients for the bank of digital IIR filters are designed in MATLAB using bilinear transformation. Specify the overall filter bank bandwidth, the number of filters, and the individual filter bandwidth at the half power point. The S-domain transfer function is determined by the following equation:

$$H(s) = \frac{(w_c / Q)s}{s^2 + (w_c / Q)s + (w_c)^2}$$

where:

w_c = center frequency
Q = center frequency divided by bandwidth at the half-power point

The bilinear transform yields the Z-domain filter coefficients by making the following substitution into the S-domain transfer function:

$$s = 2\,fs\,\frac{Z-1}{Z+1}$$

where:

fs = sampling frequency

In addition to the coefficients, the DSP must also store the intermediate *d(n-m)* values as well as the input and final output values. These values are assigned within the 1024 words of data RAM as shown in Table 1.

Table 1. Intermediate, Input, and Final Output Data Values

Data Type	**Max. Words**	**Address**
Coefficient	250	0496H
Input data	250	0600H
Output data	250	06FAH
d(n-m) values	150	0400H

The coefficients all have an absolute value of less than 2. The input data is normalized to a maximum absolute value of 1. The *Q14* format represents these fractional decimal numbers using 16 bits. The Q14 format assigns the most significant bit as the sign bit and the least significant 14 bits as the fractional portion of the number.

Eight multi-purpose auxiliary registers are available on the TMS320C26 for use as pointers in the indirect addressing mode as well as for count values for looping purposes. The code that performs SSP uses the auxiliary registers as shown in Table 2.

The DSP assembler code was written, tested, and debugged. Appendix B includes a list of the assembler code. An ultrasonic reflection containing no microstructure noise was applied to the DSP to test the code design (the test results are shown in Figure 14). Because the flaw reflection is the signal to be recovered, the output contains the original flaw.

Table 2. Auxiliary Register Map

AR7	Pointer to Coefficients
AR6	Filter Loop Counter
AR5	Main Loop Counter
AR4	Pointer to d(n–m) Values
AR3	Pointer to Input Data
AR2	Pointer to Output Data
AR1	Delay Loop Counter
AR0	Unused

Figure 14. Split Spectrum Processing of a Noise Free Signal

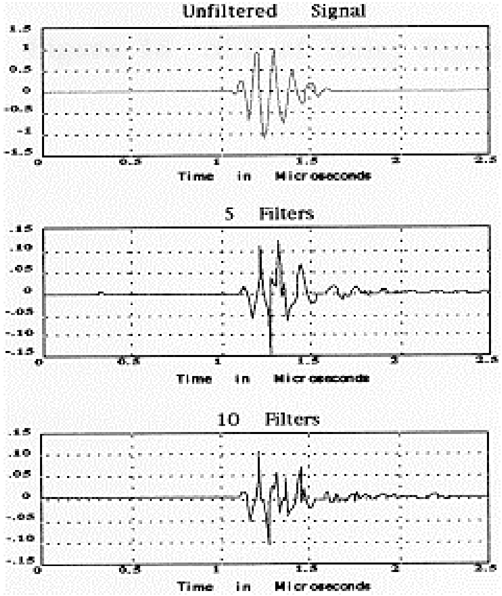

The next step was to test the system with ultrasonic data containing microstructural noise. The output in Figure 15 shows the results of applying the noisy data to the system.

Figure 15. Split Spectrum Processing of a Noisy Signal

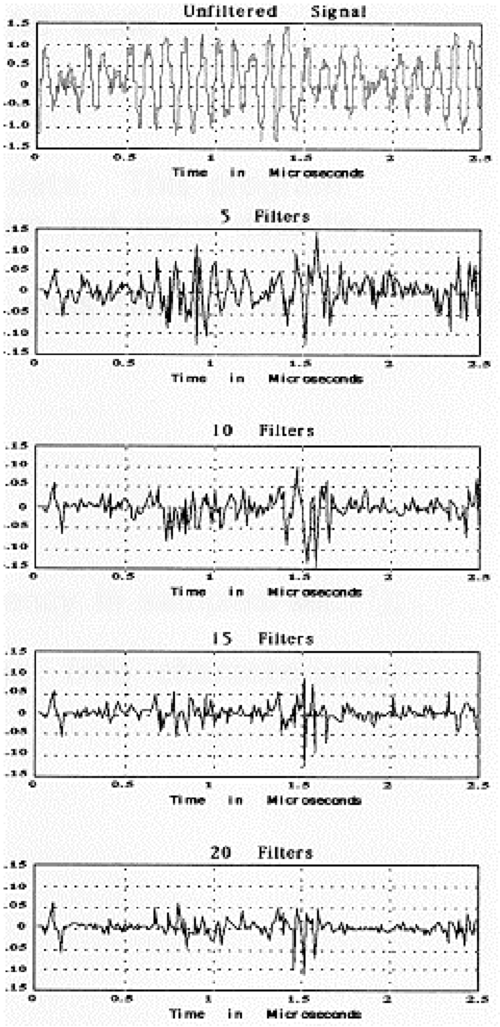

The input data originated from an ultrasonic reflection of a cast iron specimen containing a seeded flaw. The original sampling rate of 400 MHz was reduced to 80 MHz by down sampling. The center frequency of the transducer was 10 MHz. The overall low frequency cutoff of the filter bank was selected at 7 MHz. The high frequency cutoff was 13 MHz.

Figure 15 shows the effect of changing the number of IIR filters in the filter bank. In each case, the bandwidth selected was 0.8 MHz at the half-power point. As the number of filters is increased, the flaw reflection becomes more easily distinguishable. The TMS320C26 code was successful.

Conclusion

The TMS320C26 DSP performed the SSP application extremely quickly. No perceptible delay was discerned from the time the code was invoked to the time that the DSP output the results to the oscilloscope. SSP allows the ultrasonic non-destructive evaluation of many materials that previously either could not be evaluated or had to be evaluated at such low frequencies that small flaws were rendered undetectable.