General Disclaimer

One or more of the Following Statements may affect this Document

- This document has been reproduced from the best copy furnished by the organizational source. It is being released in the interest of making available as much information as possible.

- This document may contain data, which exceeds the sheet parameters. It was furnished in this condition by the organizational source and is the best copy available.

- This document may contain tone-on-tone or color graphs, charts and/or pictures, which have been reproduced in black and white.

- This document is paginated as submitted by the original source.

- Portions of this document are not fully legible due to the historical nature of some of the material. However, it is the best reproduction available from the original submission.

Produced by NASA STI Compliance and Distribution Services

77-1389

FUTURE AEROSPACE DIGITAL SIGNAL PROCESSING CONCEPTS

S. F. Hsueh, W. Vojir, and P. Burkhardt

Grumman Aerospace Corporation, Bethpage, N.Y.

Abstract

The increasing demands of digital signal processing in avionic real time applications along with technology advances have brought the aerospace industry to consider the utilization of a high-speed, modular, programmable signal processor for next generation aircraft systems. New architectural techniques, LSIC technologies, and innovative computational algorithms will be the primary factors for improving processing system performance. This paper explores these new concepts in view of future mission requirements, and processor design guidelines based upon actual experience are presented.

I. Introduction

The growing threats and mission requirements for military aircraft in the 1980-1990 time frame are expected to call for extensive processing of information contained in complex radar, ECM, communications and acoustic signals. Until recent times, analog equipments dominated signal processing. The conventional analog approach suffered from stability, flexibility, size, accuracy and dynamic range limitations. Reduction in the cost of digital hardware and particularly the advent of LSI is bringing digital hardware deeper into analog areas, as well as into areas where signal processing in any form has heretofore been impractical.

Signal processing in todays multisensor interactive avionics requires very high throughput machines capable of performing a wide variety of signal processing algorithms and simultaneous pre- and post-processing, display formatting and enhancement functions. Special-purpose hard-wired digital systems were used frequently in past years for these systems despite the unfavorable effects on life cycle costs that resulted from the design of a new processor for each application. Initially the hard-wired designs appear low in development time and risk but, when the overall avionic system is considered (including logistics, support, product improvement and next generation), the true development cost is higher. The hard-wired approach should be reserved for special applications with atypical signal processing requirements and limited potential distribution. Recent advances in computer architecture and LSI technology, DOD emphasis on standardization, and current mathematical developments are all challenging the aerospace industry to develop a programmable-digital-computer-family optimized for signal processing. Commercial programmable signal processors are being used in such diverse fields as seismic and mechanical stress analysis with some degree of success. Programmable signal processors are also being used with some success in avionic multimode applications. However, to achieve the degree of efficiency and flexibility that is necessary to prevent proliferation of avionic programmable and hard-wired designs, a carefully planned and specified family of programmable signal processors needs to be developed.

II. Airborne Signal Processing Requirements

Many digital signal processors, (e.g., the E-2C AEW radar processor), are currently implemented as special purpose hard-wired processors which perform a computational task dedicated to a particular system. The computational structure is usually based on special computation units - high-speed multiplier, adder, or multiply and add units - and the desired level of throughput is achieved by replication of these computation units in a combination of parallel, pipeline, or cascade architecture. For spectrum generation, the FFT butterfly has often been the arithmetic kernel implemented. Future aircraft signal processing requirements are reviewed briefly in the following paragraphs. In taking this look at future requirements the authors were reimpressed [1] with the desireability of achieving a multipurpose programmable signal processor family capable of effecitve reconfiguration to encompass the major areas of application.

Mission Requirements

When establishing future avionic digital signal processor requirements, one must begin by forecasting overall mission requirements of future aircraft. Grumman has considerable experience in various aircraft systems such as the E-2C Early Warning System, the A-6E All Weather Attack System, the EA-6B and EF-111 Electronic Counter Measures Systems and the F-14 Weapon System. At present, Grumman is participating in some R&D programs for Navy V/STOL aircraft and a next generation attack system [2] called Radar Guided Weapon System (RGWS). From the related overall future mission requirements we have attempted to ascertain signal processing capabilities that will contribute significantly to future aircraft missions.

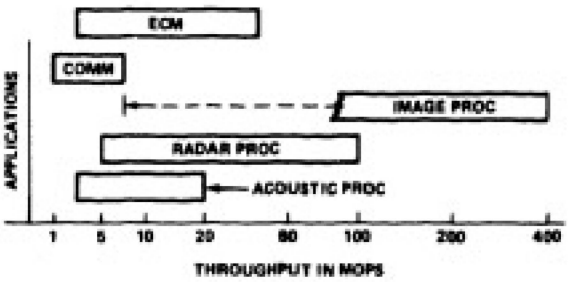
A representative group of aircraft mission modes requiring digital signal processing are shown categorized under radar air-to-air and air-to-ground and other systems in Figure 1. Figure 2 shows projected aerospace signal processing throughput requirements in these areas. Radar signal processing is the most widely distributed application in avionics and is relatively high in throughput requirements. Some simple image processing tasks such as brightness transfer are being performed in real time, as well as non-real time image enhancement on the airborne signal processors in radar and acoustic applications as indicated in the dotted portion. Applications already exist for real time image processing such as pattern recognition which require signal processing throughputs of 100 mega operations and up. These are currently being done only on the ground with large-scale systems on a non-real time basis. Certain signal processing tasks in passive ECM such as formation of pulse descriptors have been handled by special purpose hard-wired processors. These descriptors may be supplemented by pulse signature techniques such as Bessel transforms, etc., with some related tasks handled by macros on a programmable signal processor. Future ECM systems may handle several million pulses per second. Track of hundreds of emitters or radar targets involves association of signatures with large files. Demand for rapid update of track files is expected to require an associative processor or comparable associative macros and supportive architecture incorporated into a programmable array processor along with the conventional signal processing tasks.

The procedure of defining signal processor operational requirements consists of progressively defining the struc-

Copyright © American Institute of Aeronautics and Astronautics, Inc., 1977. All rights reserved.

RADAR		OTHERS
AIR-TO-AIR	AIR-TO-GROUND	
• PULSE DOPPLER SEARCH	• GROUND MAP SEARCH	• ECM
• RANGE-WHILE-SEARCH	• TERRAIN FOLLOWING/AVOIDANCE	• COMMUNICATION
• TRACK-WHILE-SCAN	• GRD MOVING TARGET INDICATION	• OPTICAL
• PULSE DOPPLER SINGLE-TARGET TRACK	• SEMI-AUTOMATIC TARGET TRACK	• ACOUSTIC AIR-TO-SURFACE
• PULSE SEARCH	• AUTOMATIC TRACK-WHILE-SCAN	– LOFAR PASSIVE SURVEILLANCE
• DOGFIGHT	• GROUND BEACON	– DIFAR PASSIVE SURVEILLANCE
• PULSE SINGLE-TARGET TRACK	• NAVIGATION UPDATE	– CASS ACTIVE DETECTION
	• LONG RANGE SAR/DBS SEARCH	– DICASS ACTIVE LOCALIZATION
	• MEDIUM RANGE SAR/DBS SEARCH/ ACQUISITION/TRACK	
	• HIGH RESOLUTION SAR/DBS	
	• SAR GUIDED WEAPON DELIV	

Fig. 1 Major Operational Modes of Next-Generation Aircraft Missions



MOPS = MILLION OPERATIONS PER SECOND
OPERATION = ONE EFFECTIVE COMPLEX MULTIPLY FOR 16 X 16 BITS

Fig. 2 Current and Projected Aerospace Signal Processing Requirements

ture of the problem from a mission requirement level down to specific algorithms and computer operations necessary to accomplish the mission. Figure 3 illustrates this top-down mission requirement to signal processor operations specification.

Operational Modes

A particular type of aircraft will of course be designed to perform a particular set of operational modes in compliance with specific vehicle capabilities and dedicated mission requirements. It is essential to specify or at least establish a preliminary framework of operational modes, mission performance parameters and accuracies commensurate with mission objectives and vehicle performance before going on to establish the processing techniques, iteration rates, algorithms, storage, data handling, programming, throughput and size, weight, power constraints, etc., of the weapon system. There are of course interfaces with other specification and design efforts not shown for simplification in the top-down flow diagram, including radar, navigation, displays, etc. These become increasingly important as the development effort proceeds.

Most of the modes of operation in avionic systems involve the use of a subset of the following major signal processing functions and techniques:

- Preprocessing
  - Presummation
  - Prefiltering and clutter cancellation
  - Pulse compression
  - Orthogonal ordering and windowing
  - Phase correction and motion compensation
- Signal Analysis
  - Weighting
  - Transform analysis
  - Time domain filtering
  - Matched filtering
  - Beam forming
- Post Processing
  - Detection, integration, and thresholding
  - Area phase angle averaging
  - Target/weapon angle measurement
  - Target association and track
- Display
  - Frequency centering and map correction
  - Pixel data compression and brightness transfer
  - Image enhancement

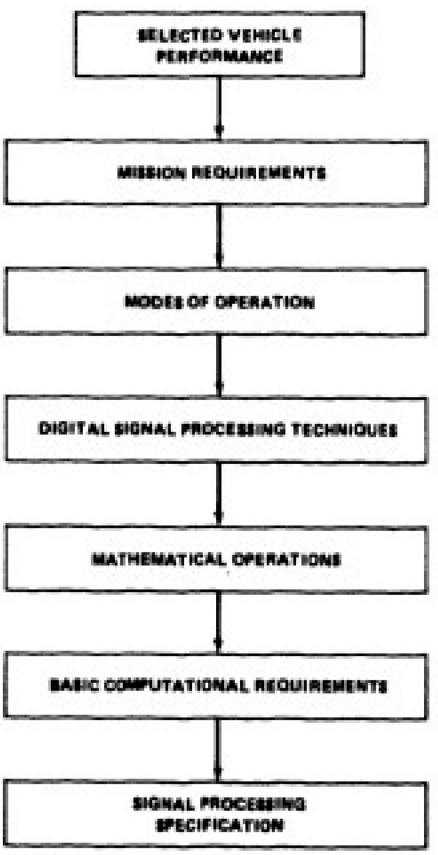


Fig. 3 Procedure for Defining Vehicle Signal Processing Requirements

Fundamental Operations

There are two key approaches associated with the above operational modes. One is the difference equation that describes the system in the time domain. For example, an MTI double delay canceller is implemented with second order difference equations to form a high-pass filter. The second is the time/frequency transform. This has typically been a fast Fourier transform (FFT) which in the frequency domain is used for spectrum analysis, signal convolution or correlation [3,4]. An example is pulse compression involving at least one forward transform, an array multiplication in the frequency domain, and an inverse transform to convert the correlated signal back into the time domain. The multiplier array is usually a stored transform table representing the compression algorithm. The signal processor not only has to perform spectrum analysis and digital filtering, it also supports pre- and post-processing operations performing such computations as presummation, signal to noise improvement, signal phase corrections, coordinate conversion, image correction, display formatting and avionic subsystem interfacing.

The mathematical and data handling operations involved in these applications can be categorized as:

- Fourier transforms
- Convolution & correlation
- Difference equations
- Vector operations
- Matrix operations

Each of these categories is a set of frequently used algorithms which are to be made available to the user of the multipurpose programmable signal processor as standardized macro instructions.

III. Future Computational Techniques

Background

Generally there are three approaches to improving the implementation of the fundamental mathematical operations outlined in the previous section. First, advances in device technology improve the inherent chip speed which affects all operations equally. Technology and architecture can be considered as having global effects and do not alter the problem from a theoretical standpoint. The third approach is to develop new computational techniques that will increase speed and accuracy and attempt to reduce the basic computer operational requirements. The effect of this theoretical approach on future signal processing systems is the subject of this section.

In practice, time domain digital filtering is used in real time applications on relatively short sample data blocks, typically of 64 or less. Transform domain processing is perferred for applications involving fine-resolution, narrow-band analysis on large blocks containing on occasion a thousand input data points. In both cases there is usually a continuous sequence of data blocks to be processed iteratively in real time. Thus it is expected that theoretical developments in these areas will have a marked impact on signal processor capabilities. Recently, new techniques have been developed which significantly reduce the amount of computations required to move between transform domains. These new forms are not bounded by data block sizes which are powers of 2 nor by the $N/2 \cdot \log_2 N$ complex multiplications of the FFT. The net effect is expected to provide a significant increase in the bandwidth of digital signal processors without the necessity for a corresponding increase in hardware speeds. The system designer is thus given greater freedom in choosing prefilter parameters, integration times and system response to be shaped in either the time or the frequency domain.

New Techniques

The new approaches are being evaluated and sifted for their areas of applicability. The ideal theorem is one (or a combination) which can be implemented with a common computational kernel that is readily reconfigured to realize the time and frequency domain transforms and other signal processing tasks with a marked optimization of the repetitive multiplications and additions.

The theoretical approach has historically advanced in three stages. The first stage involves the discovery of a more efficient organization of the basic calculations to reduce the required throughput. Thus the Goertzel Algorithm, the FFT Algorithm and more recently a multidimensional approach[5] and the Winograd Algorithm[6] all fit into the first-stage category. Although the computation has become faster due to the reduction of the number of multiplies (and, in fact, optimum via the Windgrad technique), any inherent undesireable attributes of the original process remain. Thus dynamic range, truncation/roundoff error, overflow, trigonometric function availability and multiplication are still problems that must be considered when designing the system.

The second stage involves taking advantage of the integral nature of digital devices and developing new techniques that alter the nature of the computation. These new approaches rely heavily upon number theory and advanced algebraic concepts such as ring theory and Galois field theory. The oldest of these techniques involves decomposing the original data into a residue representation, performing all calculations with respect to small integer moduli, and finally using the Chinese Remainder Theorem (CRT) to decode the results. This method appears promising due to the fact that parallel arithmetic operations utilizing small integers should enhance throughout. In addition, modular arithmetic eliminates the annoying problems of truncation/roundoff and overflow if the dynamic range is covered by the moduli set. The stumbling block has always been the reconstruction via the CRT. However, recent technology advances coupled with a new look-up-table approach to CRT decoding[7] has brought this technique back into contention for certain applications. Another promising new technique is the Number Theoretic Transform (NTT). This method is capable of making significant reductions in the number of multiplies required for some signal processing tasks. If the modulus is selected judiciously then the complex exponential of the DFT can be replaced by a power of 2 and still retain the cyclic convolution property (CCP)[8]. Utilizing such a transform reduces each DFT multiplication to a modular shift operation. Again the modular arithmetic eliminates the inherent errors of ordinary computer arithmetic. Unfortunately the preservation of the CCP leads to an interdependence of transform length, word length, modulus, and transform kernel complexity[9]. However, these drawbacks are minor irritations when the potential gain in speed and accuracy are considered. Greater flexibility may be possible by using pseudo NTTs[10]. The Galois Field Transform[11] is closely related to the NTT but is advantageous in handling complex data. This transform has all the accuracy advantages with the additional fact that a Mersenne prime is a good modulus to use thereby reducing the modular arithmetic complexity. The major drawback here is that multiplication cannot be replaced with shift operations.

The third and last stage of improvement involves analyzing each algorithm or transform and extracting its general properties to facilitate combining several methods into a single optimum technique. The rectangular transform[12] does this by combining the multidimensional, Winograd, and NTT (Specifically the Fermat number transform) methods. The result is a procedure that utilizes the CRT to reduce the calculation to short transforms in each dimension, minimizing the number of multiplications and replacing most multiplies with shift operations. At this time this method appears to offer the best computational solution although further effort could be expended in an attempt to minimize the number of additions required.

The advantages of these new processing techniques in terms of speed and accuracy will undoubtedly stimulate changes in future processor design. The major problem at this time appears to be the ability to perform residue arithmetic. The ideal modulus would be a Mersenne prime but the favored candidate is a Fermat number. A recent scheme involving a diminished-one representation tion[13] seems to offer a good approach. The effects on processor and system architecture are difficult to predict at this time. However, the increased algorithm speed coupled with technology advances will definitely make the programmable processor a major contender. The sacrifice of programmability would be made up by an ability to service more than a unique subsystem (among other advantages).

IV. System Architecture

Signal processors in the past could be characterized by the ways in which they implement a concert of simultaneous computations as opposed to general purpose computers which at any point in time are executing a single arithmetic add, shift, transfer, etc. And while new fast Von Neumann designs could handle some signal processing tasks, there are (as indicated in the first section of this paper) so many high speed signal processing tasks which are desireable on a real time basis in future avionics that the press for faster implementations will continue.

An advantage of hard-wired architectures over programmable architectures is speed or throughput with a minimum of hardware. A hard-wired system can be optimized to solve a specified problem in a specified time. Implicit in this advantage is its major weakness. It will handle only the limited number of algorithms, parameter variations and operational modes for which it has been designed, and can only be adapted to new problems, if at all, by hardware modification.

Programmable signal processors require additional hardware above that of equivalent hard-wired designs for program controlled decisions, accessions and steering of data through the arithmetic elements. Most programmable designs also require increased computational time to accomplish these operations. It is possible for a well wrought programmable design to approach equivalent throughput on the major applications as would be achieved on a hard-wired design with the the same arithmetic elements. But such a programmable design will definitely require additional hardware for program control. In addition it will require a full bag of architectural techniques to achieve the necessary program controlled simultaneous multiple accessions, address computations, arithmetic and I/O operations.

Early signal processors in radar, process control, etc., and for that matter some current and very powerful designs use a large special instruction set having perhaps 48 to 96-bit formats. These formats are broken up into multiple source/destination fields for simultaneous program controlled multiplexing of inputs and outputs to multiple arithmetic elements, registers, address and data buffers, etc. The wholly special purpose architecture and special programming tasks typical of these designs makes this approach unattractive for future multipurpose avionic signal processing applications.

There exists little commonality of hardware requirements for such designs that facilitate use of the building block modules developed for general purpose computer applications other than the usual memory expansion modules. Similarly, the packing of all I/O, addressing, arithmetic, executive, and logic tasks on one large instruction format has resulted in need for specially trained applications programmers where program execution times are to be optimized. This has left little opportunity for use of standard assembly or high order languages, or for the development of useable cross assemblers and efficient compilers. Commonality is an important factor that has not been adequately realized in the past. This factor is addressed in the context of architectural concepts as discussed below.

Multiprocessor Approaches

Some programmable signal processors have taken on features of multiprocessor architectures, including separately programmed processors for arithmetic, address, I/O and executive functions. However, these differ radically from the typical multiprocessor configuration in that some of the processor units have special purpose functions which are unique to signal processing and some have unique architecture as well.

A promising feature of the multiprocessor approaches to signal processing is that the address, I/O and executive processor units may be implemented using small general purpose computers. I/O processing and occassionally addressing or executive functions may even be implemented on dedicated microcomputers. A multiprocessor architecture such as is shown in Figure 4 is indicative of a way future signal processing could be implemented. By using small modular expandable general purpose computers for executive, I/O and address processing and standardized modular memories considerable commonality can be achieved with other avionic digital computer hardware.

Address and I/O Processing

The dedication of a separate processor to programmed compufation and control of multiple simultaneous memory accessions including multiarray address indexing, orthogonal accessions, loop counting, cueing, etc., is relatively unique to signal processing. Yet where standardized multibus control, memory interleave, and multiport shared memory or multimemories are available, a small general purpose computer with buffering for computed addresses can be used to implement these parallel accession tasks. Provision for modular changes throughput, number of ports, etc., will be necessary for efficient system configuration. A multiported interleaved memory is a minimum, with partitioning into separate memories likely for large systems.

I/O processing involves the more conventional sequential block transfers and is thus close to a conventional I/O controller. However, there are tasks of multipurpose signal processing such as data format conversion on the fly, occasional participation in orthogonal or other special accession/transfer ordering, or data truncation, presumming and scaling that are relatively unique. Thus a smart I/O controller with some buffering is necessary.

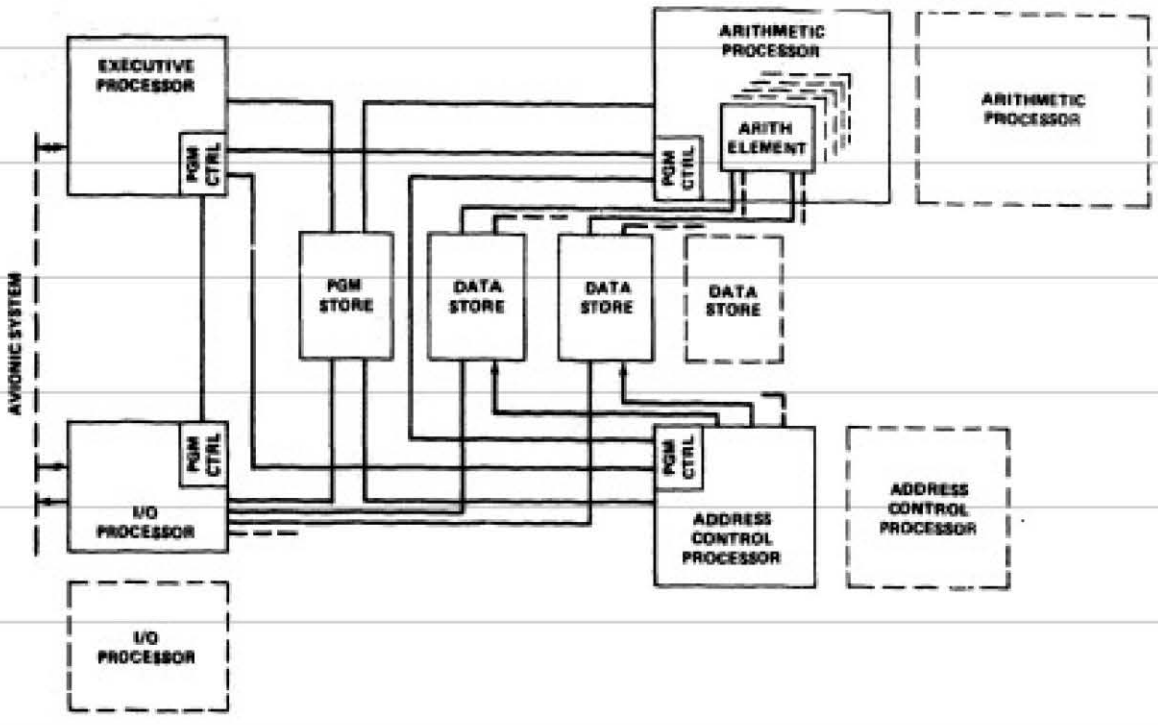


Fig. 4 Programmable Signal Processor System

Arithmetic Processing/Technology Trends

The arithmetic processor with its requirement to perform simultaneous multiplication, addition, and other arithmetic tasks in a program variable order at rates equivalent to 5 to 400 million complex multiplications per second is unique to signal processing. Careful application of parallel, pipeline and cascading techniques is necessary to achieve such rates of data crunching which must be accomplished in addition to the executive, I/O and addressing computations. The mix of multipliers, adders, complementing devices, etc, will be determined by evaluation of the new theoretic approaches to decimation, matrix computations, etc., as well as new technology developments. Multiplication in nanoseconds, reduction to simple shifts, switching in picoseconds will all speed up future implementations.

Today's technology offers a wide spectrum of logic families. Various MOS technologies (NMOS, CMOS, DMOS, VMOS) have been addressing the problem of improved speed and gate density while retaining low power leadership. Bipolar logic has become a true competitor to MOS with the maturing of I^2L, while ECL families are leaders in speed. The aerospace signal processor must consider specific requirements of circuit speed, power, cost and reliability for optimum component selection, with speed the dominant parameter for arithmetic elements. Bipolar devices allow greater throughput per LSI chip than comparable MOS devices but consume more power. Maturation of fast double diffusion (D^2) MOS or MOS/SOS could result in comparable speed. New oxide techniques applied to ECL have resulted in even greater increase in throughput. Improvement in lithography using electron beam techniques is expected to increase resolution in all technologies. Insulating substrates or cells will also reduce capacitance and increase speed. A move from 10 pj to .1 pj speed-power product might be achieved. New technology microprocessor cycle times may drop to 10—40 nanoseconds by the early 1980s.

The speed of the next generation signal processor is heavily dependent on the state of technology and the components selected. Figure 5 shows a projected time frame for development and estimated signal processor throughput. Technical evolution is expected to result in about 50 MOPS (million operations per second) for airborne signal processors by 1985 with the possibility of a technical breakthrough by the early 1980s resulting in even higher performance.

The arithmetic modules will be relatively unique to signal processing and should, by modular expansion and substitution, be able to cover perhaps two orders of magnitude in crunch capability. An interesting arrangement here is bit/byte slices allowing expandable, very

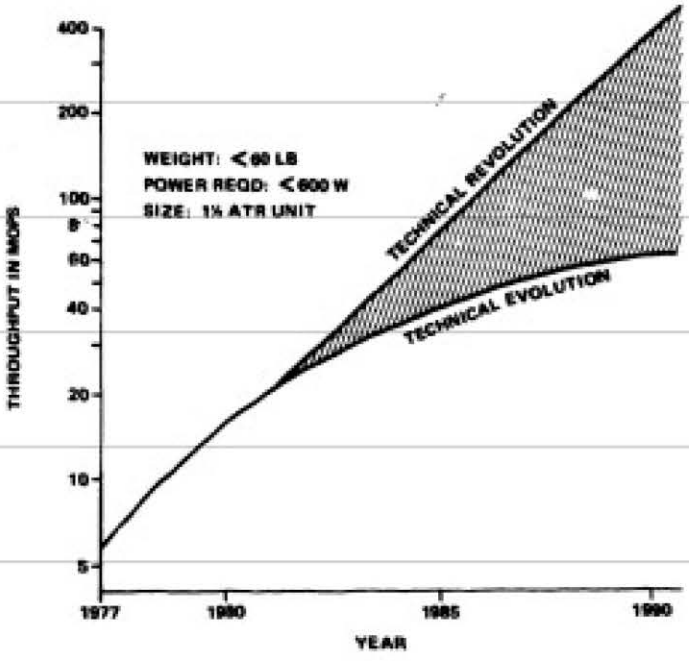


Fig. 5 Signal Processor Development Time vs Throughput

fast bit by bit multiplication [1, 14]. By retention of the multi source/destination approach to arithmetic control, programmed pipeline/parallel ordering of computations can be achieved. By separation of arithmetic control from executive, I/O and address control, common assembly language and psuedo assembly language programming techniques can be successfully applied.

At the user level a common high level language (e.g., SPL/1 or future DOD standard) and standardized signal processing macros can then be used for efficient development programming. The user may also be provided a simple procedure for combining standard macros into his own super macro to reduce subroutine transfer and initialization overhead.

Executive Processing/Task Synchronization

The executive functions have been handled with success by an external host computer in some cases. The signal processor in these implementations, though in actuality a multiprocessor system, is handled by the host as a peripheral device to be given task assignments from host program storage. The tasks are readily defined by a string of macro commands handled one or more at a time as subroutine calls from an application program executed in the host. A limitation here has been that of synchronization of program control of macro command transfers, data block transfers and signal processor executions to achieve efficient utilization of the signal processor over a variety of tasks. The overhead for Fortran calls alone on some implementations has idled the signal processor a significant portion of the time. Little program control can be exercised by the host on data transfers unless the data bus is shared by host and signal processor. These need to be program controlled in multipurpose applications as to source, destination, channel, block size, initialization, etc.

Another facet of system synchronization and control that puts a heavy load on an external host executive is accommodation of a complex stream of tasks, typically a stream of macros and sustaining tasks. Each processor unit will execute its own machine language routine representing its assignment for a given macro. A separate macro involving only the I/O units and data memory may be in progress, for instance, while an FFT macro is in execution on the arithmetic and address processors and the executive being otherwise idle is perhaps assigned computation of the motion compensation vector for the next data block until an interrupt indicates completion of an I/O transfer or transform task. A sequential stream of macros will be performed on a given data block while a new block is being input and formatted and the processed results of a previous block are formatted and output. A typical string of macros may start with preprocessing tasks such as presum and phase correction, proceed with weighting and spectrum generation, and continue with post integration and detection followed by pixel data compression. While all of this is to be coordinated by the executive processor, it will perform sustaining tasks such as servicing of displays and control requests for parameter and mode changes.

The problems of control synchronization within a signal processing system between the executive and the arithmetic, address and I/O processors are minimal only where dedicated control buses are provided. Thus a dedicated executive processor unit appears to be the better approach. The external host typically will function for data transfers, program loading and limited control via one or two buses. The limitation here is similar to that of real time multiprocessing systems comprised of several general purpose computers, where one processor assumes executive functions and iteration rates are limited to that of the processor having the longest task assignment plus hand shaking time to synchronize task initialization in each of the processors. A solution is a dedicated programmable executive processor unit with private control buses to minimize synchronization overhead. An interesting variant here is a data driven approach with the dedicated executive concerned with long term task assignments and synchronization. Short term task assignments are reduced by cued backlogs of control commands, data buffers and instructions. This asynchronous approach can be fast, but has some subtle requirements as to programming considerations and cueing.

V. Conclusions

This paper is an attempt to outline likely requirements for signal processing in avionics for 1985 and beyond, and to indicate some of the considerations that will influence the design and performance of future signal processing machines. Detailed discussion is not attempted in any of these areas but rather this paper serves to indicate some new as well as some currently successful techniques which should be exploited for multipurpose applications. In order to minimize the future need for relatively costly special purpose implementations, we conclude that a multipurpose programmable signal processor family is needed which meets the full diversity of major avionics applications. Functional modularity and software commonality are recommended as areas of standardization which will allow for growth in device technology and theoretic developments.

Acknowledgements

The authors would like to express their appreciation to Mr. L. Della Salle for his valuable comments and suggestions. They also wish to thank Mr. R. Hong for his support and helpful advice and Mr. G. McNiff for reviewing the manuscript.

References

1. P. E. Burkhardt and S. F. Hsueh, "Design of a General Purpose High Speed Programmable Fourier Processor," Grumman report ADR 03-09-70.1, March 1970, revised September 1974.

2. N. Speiss, "Standoff, Attack. . . .and Survive!," Horizons, Volume 13, Issue 2, pp. 8-21, Grumman, 1977.

3. L. R. Rabiner and B. Gold, "Theory and Application of Digital Signal Processing," Prentice-Hall, 1975.

4. A. V. Oppenheim and R. W. Shafer, "Digital Signal Processing," Prentice-Hall, 1975.

5. R. C. Agarwal and C. S. Burrus, "Fast One-Dimensional Digital Convolution by Multidimensional Techniques," IEEE Transactions on Accoustics, Speech, and Signal Processing, Volume ASSP-22, pp. 1-10, February 1974.

6. S. Winograd, "A New Method for Computing the DFT," IEEE International Conference on Acoustics, Speech & Signal Processing, 1977.

7. W. K. Jenkins and B. J. Leon, "The Use of Residue Number Systems in the Design of Finite Impulse Response Digital Filters," IEEE Transactions on Circuits and Systems, Volume CAS-24, No. 4, April 1977.

8. C. M. Rader, "Discrete Convolutions via Mersenne Transforms," IEEE Transactions on Computers, Volume C-21, pp. 1269-1273, December 1972.

9. R. C. Agarwal and C. S. Burrus, "Number Theoretic Transforms to Implement Fast Digital Convolution," Proc. IEEE Volume 63, pp. 550-560, April 1975.

10. H. J. Nussbaumer, "Digital Filtering Using Pseudo Fermat Number Transforms," IEEE Transactions on Acoustics, Speech, and Signal Processing Volume ASSP-25, pp. 79-83, February 1977.

11. I. S. Reed and T. K. Troung, "The Use of Finite Fields to Compute Convolutions," IEEE Transactions on Information Theory, Volume IT-21, pp. 208-213, March 1975.

12. R. C. Agarwal and J. W. Cooley, "New Algorithms for Digital Convolution," IBM Research Report RC 6446, IBM Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598, March 23, 1977.

13. L. M. Leibowitz, "A Simplified Binary Arithmetic for the Fermat Number Transform," IEEE Transations on Acoustics, Speech, and Signal Processing, Volume ASSP-24, pp. 356-359, October 1976.

14. L. Della Salle, W. Butler and T. Newman, "Aerospace Applications of Microprocessors," Proc. IEEE 1976 NAECON, pp. 307-313, May 1976.