SANDIA REPORT

SAND95–2050 • UC–706
Unlimited Release
Printed September 1995

Hardware Architecture Study for Color Printed Pattern Defect Detection

P. Eichel, A. Yee, G. Sleefe

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550
for the United States Department of Energy
under Contract DE-AC04-94AL85000

SF2900Q(8-81)

* incl. cover (both sides) 18 p. *

SAND 95-2050
Unlimited Release
Printed September 1995

Distribution
Category UC_706

Hardware Architecture Study for Color Printed Pattern Defect Detection

P. Eichel, A. Yee and G. Sleefe
Sandia National Laboratories
Albuquerque, NM 87185

Abstract

This paper presents the real-time hardware architecture study for the on-line color printed pattern task of the Computer Aided Fabric Evaluation (CAFE) project. Two types of hardware architectures are considered. The first is a hybrid architecture using a combination of pipeline processors and Digital Signal Processors (DSP) and the second study is of a pure DSP architecture. For each hardware architecture the rational for the architecture, system block diagram, implementation issues, advantages and disadvantages, and system modularity and extendability are presented.

1.0 Introduction

The following is the result of hardware architecture study performed at Sandia National Laboratories for the Color Printed Pattern Defect Detection project. Two basic approaches were considered. The hybrid approach incorporate both pipeline processors and DSP boards to take advantage of current developments. The Digital Signal Processor (DSP) approach considers the feasibility of a solution using off the shelf DSP boards. This study included some justification for each approach, possible system layout, and a list of advantages and disadvantages of the two systems.

2.0 Hybrid Hardware Architecture

2.1 Introduction

A real-time system for the color printed pattern project must be able to reliably collect, process and display the defect detection results of the real-time camera data from the web. It is proposed that a hybrid system consisting of pipeline and DSP processors can be used for the implementation of the defect detection algorithm. Such mapping would provide the necessary processing elements for acquiring, processing and displaying of the real-time data.

2.2 Rationale

Specialized pipeline hardware provided by companies such as Datacube consists of architectures tailored toward a specific family of problems. The systems are in general hard wired to perform a sequence of functions efficiently on large arrays of data. Although ideal for algorithms within the assumed scope of the design, these systems often lack the necessary flexibility to implement functions as required by a general algorithm. In addition, the required overhead for the accumulated pipeline latency may become unreasonable for small data sets. Therefore, as an algorithm progresses beyond the scope of the pipeline system design, it becomes necessary to supplement the system with general purpose processors in order to achieve optimum utilization of the pipeline processors.

The current data collection system provided by ORNL utilizes some functionality provided by the Datacube board, MV200, for data collection, storage, and display. This board consists of different modules of specialized pipeline processors which can be connected via software selected register settings to perform desired functions. Although this is ideal for steps involving sequential operations at the pixel level, it has been determined that for the calculation of the warp coefficients using only the pipeline architecture would require an unreasonable amount of overhead on the MV200 board.

The calculation of warp coefficients requires correlation of various regions of interest between the ideal and the camera image. Such operation over the required size of region of interest is not directly supported by the MV200 hardware. Although the functionalities can be implemented, the required number of operations becomes prohibitive in real-time. The alternative method for calculating correlation peaks is using Fourier transforms, which is well supported by DSP boards. The DSP boards are general purpose and many are capable of performing the required operations in real-time. Datacube supports DSP boards based on the i860. Furthermore, this board has already been successfully integrated into the data acquisition system although currently only acting as an additional memory module.

For the alpha test, the hybrid system option seems ideal because it builds upon previous efforts and does provide the necessary real-time processing requirements.

2.3 System Block Diagram

The proposed system consists of two boards. The pipeline processor (MV200) has the required input and output capabilities for data acquisition and display. Furthermore, the processors can perform the sequential detection step efficiently. The Max860 (i860 based DSP board) is used for the general purpose processing.

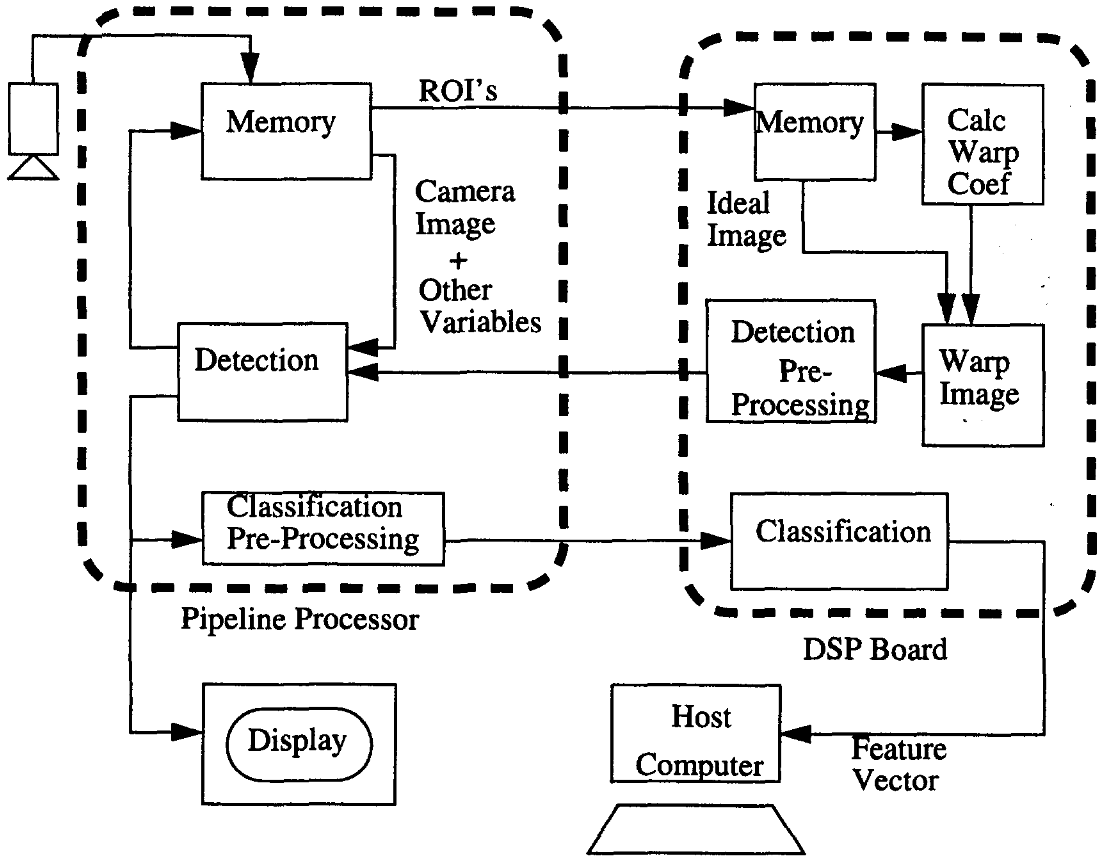

Figure 1: Possible algorithm mapping onto the proposed hybrid system architecture.

Although the image warping is shown in the diagram to be done on the DSP board, there exist a specialized warping module which may be purchased should the data throughput requirement be increased. Furthermore, the Max860 board comes with either one or two i860's and could easily be placed in parallel to gain processing speedup.

The elementary components of the proposed hybrid system are as follows along with a brief description of each device contained in the system.

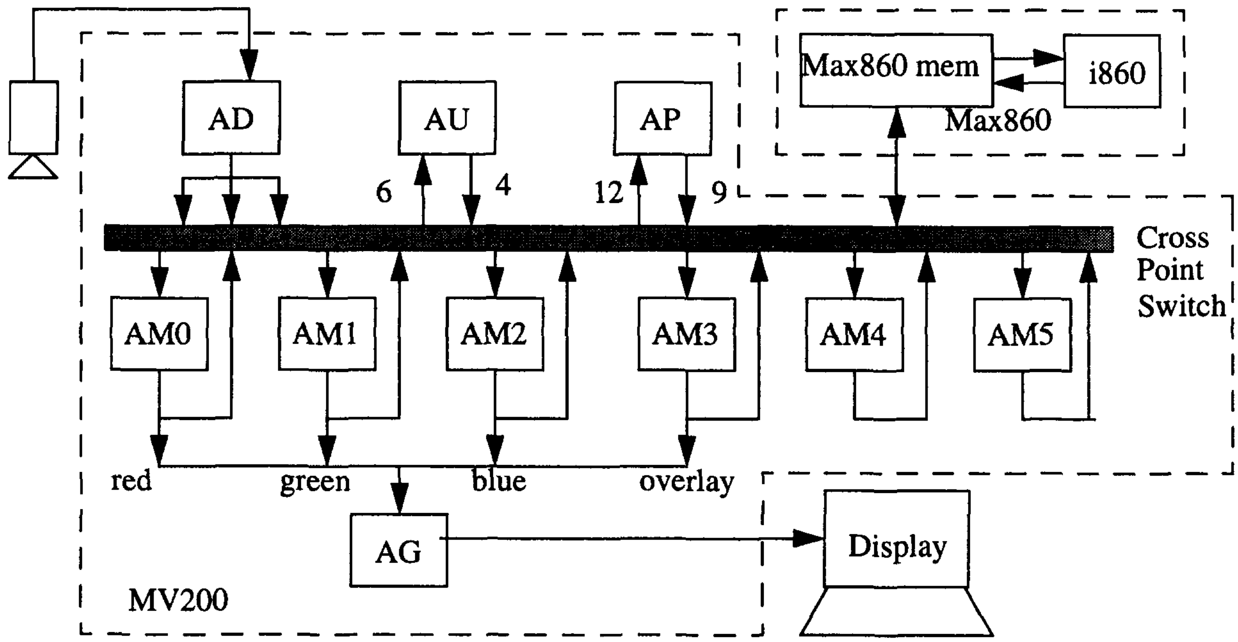

Figure 2: Processor layout of the MV200.

- The AD device is the 24bit digital camera interface with internal address generator which puts the acquired RGB data onto the crosspoint switch (three bytes per image pixel).
- The AU device is the set of pipeline processors which performs linear, non-linear, and statistical operations. It can receive up to six bytes as input and output up to four bytes.
- The AP device performs more advanced processing. It contains various modules to perform statistical operations, up to 8x8 neighborhood multiply and accumulate, morphological operations, and an additional slot for a daughter card such as the 1Kx1K mini-warper or the rank order filter.
- The AG device is the memoryless display module which selects the desired input configuration from the memory outputs and converts the data to various types of analog signal to conform to the display terminal.
- The AM device is the memory module provided in the MV200. This device contains either 4 or 16 MB of allocatable memory in 1, 2, 4, or 8 bit format. There is a look-up table, an ALU capable of performing simple operations and statistical operators native to this device. The receiving gate can be set to subsample an image in both the X and Y direction and it can also selectively write protect a pixel location. The transmit gate can duplicate pixels in both the X and Y direction. Since all incoming data must pass through one receiving gate and all outgoing data must pass through one transmitting gate there can only be one stream of input and one stream of output at any given time for a memory module.

- The Max860 memory is a 16MB allocatable memory. It can either receive or transmit data from the external ports at a given time but not both. The on board i860(s) can only access data within this memory. Memory can be allocated in 1, 2, or 4 bytes and with some boundary restrictions.

2.4 Advantages and Disadvantages

Due to the nature of the algorithm, the pure pipeline architecture has become cost ineffective to perform the required operations. On the other hand, a pure DSP architecture is very cost effective, but considered exploratory at this time. A hybrid system for the alpha test would provide the required functionalities within the required time frame, and such a system has already been demonstrated for the RGB data collection system. This system consists of one MV200 and one Max860 (one i860) board. Although currently the Max860 board is used only as a memory buffer, the necessary mathematical libraries and software development software are available and the data connection between the two boards has already been proven in the RGB data collection implementation.

The hybrid system provides the path for good utilization of all processors for their intended usages including some parallelization for intermediary processing steps due to the memory set up of the MV200. The memory modules allow easy manipulation of data (downsampling, translation for camera compensation, and region extraction). The i860 provides sufficient processing capabilities to both calculate the warp coefficients and extract the warped and downsampled image variables from the ideal image data. This system can also be expanded without much effort to include parallel boards should parallel processing be necessary to accomplish classification or to attain speed up from the current process in order to accommodate either higher image resolution or faster web speeds.

One disadvantage of such a hybrid system is the use of the pipeline processors for sequential detection. The hardware does not allow for flexibility in changes within the algorithm, and additional boards may be necessary should the complexity of the algorithm increase causing a possible need to re-map the algorithm onto the new hardware configuration. Furthermore, porting of the software onto new technologies or other hardware platforms may require a major overhaul of the software.

The use of the MV200 board may also require either a frame grabber or an additional memory module to properly display the acquired images while using the pipeline processors for sequential detection due to its memory and display configuration. The requirement for multiple boards increases the cost of the final system and may not be acceptable.

2.5 Implementation

The inspection algorithm would be implemented to reflect the following order. First the images are acquired continuously. Then the host processor polls the registers to see how many rows of the data have been acquired. When a sufficient number of rows have been acquired, regions of interest are fired to the i860 for processing and a subregion of full resolution image is downsampled in order to prevent overwriting of the buffers by the subsequent frame. The warp coefficients are calculated by the i860 and the ideal image and its associated variables are warped using the computed coefficients. These warped images are then used for sequential defect detection and classification.

2.6 Modularity and Extendibility

The proposed hybrid system is extremely modular and extendable. These boards are VME based with an additional specialized bus available which can be connected in series for faster data transfer. Furthermore, there exist software development tools and optimized libraries in support of board operations and timing controls for parallel operations. Additional MV200, Max860 or a variety of other boards are supported by the software development environment which allows for ease of data transfer and control between parallel boards performing other required operations.

2.7 Areas of Future Work

Explore other hardware combinations which would provide either performance or cost benefits and also address portability issues to take advantage of emerging technologies. One possible direction would be to use the MV200 only as a data acquisition and display system and port all processing to DSP boards. This would take advantage of existing I/O capabilities and provide more portable software for the future.

The classification algorithm in all likelihood would be mapped onto a general purpose DSP board. The detection algorithm acts as a cueing system and the classifier would determine the error type based on the ideal image information and the associated accumulated deviations.

Currently the data acquisition terminates before a full repeat length is acquired and the misalignment of each plane from the camera are adjusted by discarding misaligned pixels at the beginning and end of each frame. This misalignment adjustment causes a loss of 32 rows per frame. The problem needs to be addressed once full repeats are acquired in order to prevent any gaps in defect detection.

3.0 A DSP--Based Hardware Architecture Introduction

In this Section, we describe a potential hardware architecture for real time color printed pattern defect detection based on a parallel Digital Signal Processor (DSP) structure. The technique is proposed for application to the Color Print task of the Computer Aided Fabric Evaluation (CAFE) project. The hardware described herein is intended to implement the sequential defect detection and MAP classification algorithms described in companion papers [1][2]. This approach is one of several under consideration for the prototype architecture. It is intended to take advantage of modern commercial off the shelf (COTS) DSP boards featuring high processing bandwidths, fast inter-processor communications, rich instruction sets, and low cost.

3.1 Rationale

No commercially available single-IC DSP module today boasts the computational throughput required to implement the CPP algorithms in their entirety and meet the real time requirements of this project. In these situations, it has become standard practice to divide up the processing requirements into pieces that may be implemented on a single DSP device and then network the various DSPs together to accomplish the whole job. The separate computational units then work in *parallel*, each with its proscribed piece of the problem. Broadly speaking, this parallelization may be accomplished either by dividing the input data at the outset and having each DSP implement the entire processing chain on its share of the data, or by assigning each DSP a specialized portion of the algorithm, which it applies to all the data. A mixture of both of these methods can also be used, the precise choice often depending on the particulars of the algorithms themselves.

Despite the fact that rather high *computational* throughputs have been available for some time with DSP technology, a task such as CPP defect detection could not until recently be economically mapped to such hardware. This is because of the very high inter-processor data rates demanded by this problem. More recently, however, DSP processing boards have become available with specialized hardware to handle the inter-processor communication bottleneck. These boards combine high communication bandwidth, memory bandwidth, and processing throughput to realize a truly potent parallel processing architecture. This together with the usual DSP advantages of low cost, plug-and-play modularity, and sophisticated development environments suggest that this architecture is a potential contender for the CPP problem.

3.2 Algorithm Implementation

The algorithms to be implemented on the proposed hardware are described in detail in references [1] and [2]. They may be roughly divided into four subtasks as follows:

Preprocessing This subtask involves interfacing with the digital RGB camera, data input, camera compensation, subimaging (to support image-to-exemplar registration processing), and subsampling. This step requires few operations per sample, but the data rate is highest here.

Image Registration Accurate registration of the camera image data to an exemplar is a crucial step in the defect detection process. This subtask involves relatively little data, but is very intense computationally. The computations are principally floating point operations (to compute Fourier transforms).

Sequential Detection This is the heart of the defect detection algorithm. It falls between the above two subtasks in that a moderately high number of data samples must be processed per pattern frame with medium number of operations per sample. Most operations are integer arithmetic, comparisons, and data moves.

MAP Classification Since this subtask operates only on detected defect samples, the amount of data samples processed per frame ought to be far below the input data rates. Most of the operations involve table lookups and accumulations.

Of these algorithm subtasks, only the first is not particularly well suited to a DSP implementation. DSPs do not represent the best choice for the preprocessing subtask, since it is one mainly of moving data around. If an all-DSP architecture is desired, however, we have identified several DSP boards that could accomplish this subtask as well. Two features, zero wait state access by the CPU to input digital data FIFOs and concurrent I/O processing, enable these boards to keep up with the real time requirements of this function. One processor would be required.

Since image registration involves fast Fourier transforms, correlations, filtering, and linear regression, this subtask is ideally suited to DSPs. In fact, they are the clear best choice for implementing this function. Using commercial third party software, published benchmarks indicate that a number of standard DSP chips can easily perform the 9 128x128 FFT/correlations per frame required. The linear regression has already been implemented by SNL on a DSP using only high level (C) programming and commercial DSP libraries.

Sequential detection is likewise a natural choice for DSP processors. This task involves relatively few data moves, but can take advantage of the rich instruction sets of current DSP processors. The type of operations required by this algorithm are varied and involve floating point, comparisons, arithmetic, and logical manipulations of the data. While certainly not the only means of accomplishing this subtask, DSPs represent a particularly cost effective implementation. Analysis indicates two high performance processors per camera could perform this function. The classifier represents a rather small percentage of the whole processing task. It may be easily implemented in a DSP or perhaps even on the host computer.

It remains to consider inter-processor communication. Modern DSP boards accomplish large communication bandwidths in three ways. One method is to endow the processor chip itself with ports dedicated to inter-processor communications. These ports are backed up by DMA coprocessors that handle the data flow independently of and concurrently with the cpu. A second method is to provide an off chip DMA engine and dual port shared memory with the cpu. In this way, the DMA processor can access the memory on a non-blocking basis. Finally, the third and most sophisticated method is to use non-blocking crossbar switches between processors, memory, and the system bus. This scheme allows point to point communications between all the system assets.

3.3 System Block Diagram

An all--DSP system to implement the detection and classification of color printed pattern defects would require four DSP processors and a host. A representative system built around an industry standard system bus is depicted in Figure 1. This architecture utilizes two dual processor DSP boards and a system host board. The display requirements, which would typically call for a true color downsampled image refreshed at about 1 Hz., are not very stringent. Thus the host could provide supervisory control, display, and classification, and any other high level processing. The data is shared between boards using one of the methods discussed above to avoid contention on the system bus.

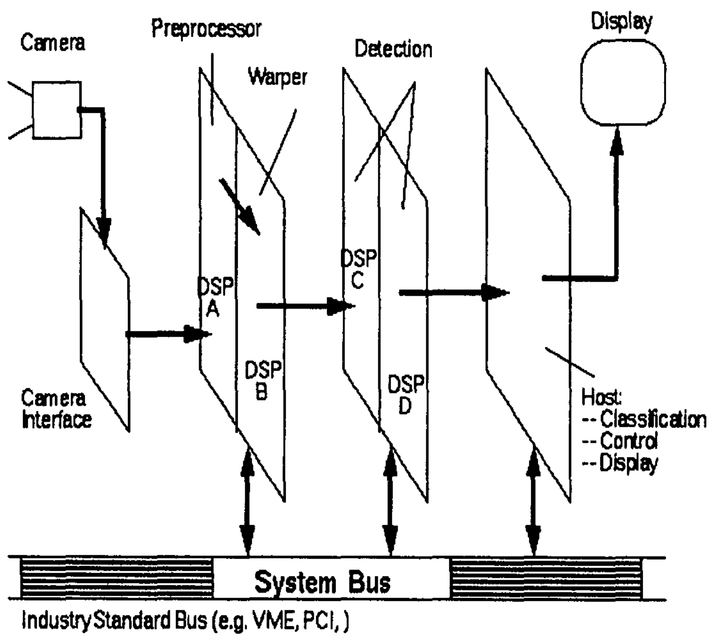

Figure 3:A representative DSP--based hardware architecture for CPP defect detection.

3.4 Advantages and Disadvantages

The advantages of a DSP--based architecture are many. First and foremost, any application implemented on DSPs can ride a very broad technology curve. There are a very large number of commercial vendors competing for ever higher performance / cost ratios. Although there is currently a wealth of VME based DSP boards available at this point in time, the industry is moving very quickly in the direction of PC hosted, PCI bus technology. This will lead to even broader implementation avenues at yet lower cost. A second advantage in a development program such as this, is the development environment itself. Practically all DSP hardware is supported by comprehensive suites of tools including high level cross compilers, simulators, emulators, and standarized realtime operating systems. These are available from third party vendors as well. All this makes code development, portability, configuration control, and maintenance much more tractable. Finally, a DSP approach is highly modular and scalable. Obtaining more computational power can generally be accomplished simply by sliding more cards in the cardcage or even upgrading the existing cards themselves.

The disadvantages of such an architecture include the fact that DSP boards exhibit limited pixel-rate processing capability, especially compared to pipelined architectures, and today's boards have a limited support of digital camera I/O and video output. These latter two might be expected to improve as the technology matures. The first may only be ameliorated through the use of parallelism. This fact increases the technical risk of this approach, since the parallelism required entails high speed communications between DSP processors not only on the same board but also at the board level. The overall management at the system level is no trivial task.

3.5 Modularity and Extensibility

Of all three designs discussed in this report (pipelined, hybrid, and DSP), the DSP approach represents the most modular and extensible approach. With the pipelined approach, the algorithm complexity is limited by the computational elements, the bus throughput, and the memory constraints. As such, adding additional pipelined processors does not necessarily allow for extension of the baseline set of algorithms. While the hybrid approach allows for more general purpose algorithms, there is as of yet no industry standard for connecting the pipelined processors with the DSP's.

Several families of DSP's have been developed which naturally allow for a modular and extendable system using parallel arrangements of DSP's. Two such popular DSP's are the Texas Instruments TMS320C40 family and the Analog Devices 21060 SHARC family. Both of these devices allow for interprocessor communications and plug-and-play arrangements of local and global memory. This allows the system to be expanded simply by adding additional DSP processor modules. For example, there exists an industry standard TIM-40 module which defines plug-in C40 processors with a wide variety of memory and i/o configurations. Furthermore, a common development system and excellent 3rd party support assures that these systems can be maintained and grow as system requirements change.

There is naturally a limit to the extendability of the all-DSP approach. The approach is based on the assumptions taken in the sequential detection described in reference [1]. Should more extensive pixel-rate processing be required, then the all DSP system would need to be extended into a hybrid approach.

3.6 Summary

The all DSP based hardware architecture described here represents a high performance, low cost approach to the color printed pattern defect detection problem. DSP's provide unparalleled algorithm flexibility, modularity, and a steeply improving cost / performance ratio from a wide variety of commercial vendors. This is balanced by a higher level of technical risk associated with integrating a system with a number of DSPs implementing a complex algorithm suite in parallel and in real time.

3.7 References

1.) P. Eichel, ``A sequential detection algorithm for color printed pattern defects,'' AMTEX CAFE Publication, 1995.

2.) P. Eichel, ``An optimal MAP classification algorithm for color printed pattern evaluation,'' AMTEX CAFE Publication, 1995.

4.0 Summary

Architecture Study This document describes two architecture approaches which can be applied to the Color Printed Pattern inspection task. The approaches have various trade-offs in terms of complexity, modularity, and cost. The first approach, referred to as the Hybrid System, adds a general purpose digital signal processor (DSP) to a state-of-the art pipelined image processing system. The pipelined system is based on the DataCube MV-200, and the DSP is based on the Intel I860. The hybrid approach allows for real-time execution of various general purpose algorithms (such as 2D FFT) and permits more flexibility in the development of the proposed CPP algorithms. The down-side of this approach is that the interconnection of these different architectures is currently costly and non-standardized.

A system based on an all-DSP architecture has also been proposed and has the potential advantage of low cost and modular expansion. DSP's can implement a wide variety of signal/image processing algorithms and offer an industry standard development environment. It is a high risk approach, however, since a high level of parallelism is required to achieve a large number of pixel-level operations on large images.

Based on the above considerations, the SNL team has decided to take a phased approach to the system development. Extensive work has already been committed to the pipelined approach and has resulted in the Bobbin Show Demo (LLNL) and the RGB Test System (ORNL/SNL). While many of the proposed algorithms can be implemented on this system, a hybrid system may be required to test and evaluate certain functions, such as real-time warp coefficient computations.

As such, we propose that the Alpha Test system will be a hybrid system which enables a wide range of algorithms to be tested. The hybrid system will be completely compatible with the pipeline (MV-200) system already being utilized by the laboratory partners. Should our research indicate that the required level of parallelism and pixel throughput can be achieved using DSP technology, then the hybrid system could be converted over to an all-DSP system. The conversion is simplified by the fact that much of the DSP programming performed for the hybrid system will be directly transportable to the all-DSP system. Therefore, this approach represents a phased, cost-effective system development effort with minimal technical risk.

Distribution:

1 AMTEX Program Office
Pacific Northwest Laboratory
Attn: Doug Lemon
P.O. Box 999, MS K7-80
Richland, WA 99352

1 Burlington Industries Inc./Klopman Fabrics
Department of Quality Assurance
Attn: Andy Russell
P.O. Box 2000
Hurt, VA 24863

1 Cone Mills Corporation
CAD Department
Attn: Pamela Richard
P.O. Box 8
Carlisle, SC 29015

1 E.I. du Pont de Nemours and Company
Central Research and Development Department
Attn: Arun Prakash
Experimental Station
P.O. Box 80357
Wilmington, DE 19880-0357

1 Fieldcrest Cannon, Inc.
Department of Research/Quality Control
Attn: William Hamlett
One Lake Circle Drive
Kannapolis, NC 28081

1 Fruit of the Loom
Department of Product Development
Attn: David Allen
Route 1, Box 1
Rabun Gap, GA 30568

1 Institute of Textile Technology
CAFE Department
Attn: Dan McCreight
2551 Ivy Road
Charlottesville, VA 22903

1 Lawrence Berkeley Laboratory
Department of Engineering
Attn: Craig Fong
1 Cyclotron Road, Building 90/2148
Berkeley, CA 94720

1 Lawrence Livermore National Laboratory
Engineering Research Division
Attn: Jose Hernandez
P.O. Box 808, L-156
Livermore, CA 94551

1 Milliken Research Corporation
Attn: Louis Adams
P.O. Box 1927
920 Milliken Road
Spartanburg, SC 29304-1927

1 Oak Ridge National Laboratory
Attn: Glenn Allgood
Bethel Valley Road
Oak Ridge, TN 37831

1 Oak Ridge National Laboratory
Instrumentation and Controls Division
Attn: Bryan Allen
P.O. Box 2009
Oak Ridge, TN 37831-8066

1 Oak Ridge National Laboratory
Instrumentation and Controls Division
Attn: David Sitter
Building 3546, MS 6011
Oak Ridge, TN 37831

1 Oak Ridge National Laboratory
Instrumentation and Controls Division
Attn: Ken Tobin
Bethel Valley Road
Oak Ridge, TN 37831-6011

1 Spartan Mills
Cherokee Finishing Company/Color Shop
Attn: Benji Mills
210 Chandler Drive
Gaffney, SC 29341

1 Spartan Mills
R&D Department, Information Services
Attn: Patrick Chong
805 Spartan Boulevard
Spartanburg, SC 29301

1 Thomaston Mills, Inc.
Finishing Department
Attn: Joe Perdue
P.O. Box 311
Thomaston, GA 30286

1	MS0842	Carolyne Hart, 9100
1	MS0842	Larry Hostetler, 9133
1	MS0843	Terry Stalker, 9136
20	MS0843	Paul Eichel, 9136
5	MS0843	Gerard Sleefe, 9136
1	MS0844	Amy Yee, 9133
1	MS9018	Central Technical Files, 8523-2
5	MS0899	Technical Library, 13414
1	MS0619	Print Media, 12615
2	MS0100	Document Processing, 7813-2, for DOE/OSTI

Org.	Bldg.	Name	Rec'd by	Org.	Bldg.	Name	Rec'd by

Sandia National Laboratories