

The Use of the Philips TM1300 for Machine Industrial Inspection

Imaging systems employed in demanding industrial and military applications, such as computer vision and automatic target recognition, typically require real-time high-performance computing resources. While these systems have traditionally relied on proprietary architectures and custom components, recent advances in high-performance general-purpose microprocessor technology have produced an abundance of low cost components suitable for use in high-performance computing systems.

A common pitfall in the design of high-performance imaging systems, particularly systems employing scalable multiprocessor architectures, is the failure to balance computational and memory bandwidth with I/O. The recent introduction of microprocessors with large internal caches and high-performance external memory interfaces make it practical to design high-performance imaging systems with balanced computational and memory bandwidth. In these systems, both the board level memory and I/O architecture, as well as the microprocessor memory, have significant performance impact. Systems that do not scale the memory bus bandwidth as processors are added do not typically improve in performance, or they reach a limiting value after some initial small gains in performance. In addition, the I/O bandwidth plays a significant role in overall performance.

Therefore, the most important factors in the selection of a programmable imaging system are: (a) will the system be "native" or will it use a co-processor model, (b) which processor has the correct balance of processing and I/O for the application, and finally (c) what is the overall cost? A real-world example of such designs will be presented, and its performance and cost will be analyzed. We conclude by outlining the advantages of the Philips TM1X00 series processors for machine vision applications.

Description of the Application

In order to characterize the advantages of the various state-of-the-art microprocessors, we will present a typical application. We will calculate the performance of the latest microprocessors from Analog Devices, Intel®, Philips Semiconductors and Texas Instruments¹.

The important properties of the selected processors are presented in the table below. Cluster mode designs will reach

¹ The Motorola PPC750-AV (AltiVec) and the Hitachi/Equator MAP1000 are not presented since, at the time of this writing, there is not sufficient public information available to adequately present them.

bus saturation very quickly, to the point that adding additional processors will not improve the rate at which calculations can be performed. The table assumes that the processors are isolated from each other when they are performing the calculations. This assumption is made because most multiprocessor co-processor boards have isolated or local memory designs due to scalability issues. The only non-local memory design is the "native" PIII-450 design, which is added for comparison purposes.

SPECIFICATION	PIII-	Philips	TI C6701	ADI 21160
Architecture		VLIW	VLIW	
FPU	Yes		Yes	yes
	250	720		600
MIPS (Peak)		900	1336	
MOPS (Peak)	900		1336	800
Bandwidth (MB/s)	800		332	400
	300	106		90
1K 16 bit cfft (µsec)		63	108	
1K FP dot product (µsec)	7.38		3.07	5.12
	2.32	1.42		5.32
8x8 MACs (MMAC/s)		6.55	7.11	
512 x8 bit Conv3x3 (msec)	5.34		7.11	11.80
³ x8 bit Erosion/Dilation	7.34	1.42		3.93
"Glue" Logic Cost (\$/CPU)		\$3	\$65	
CPU Price (\$)	\$300		\$150	\$100

Industrial Inspection of Components on a Production Line

As a product passes the inspection station it is examined for color, and shape, and accepted or rejected based on those factors. The components are placed at any angle on the conveyer while they are imaged from above. The inspection algorithm detects connected regions of color and shape, computes statistics on those regions, and applies the resulting features to a neural network discriminator. The product is examined as it moves rapidly down the line. The camera collects images at a rate of 30 per second to be sure there is at least 50% overlap in each image, ensuring each part will be imaged completely, at least once. The image is 512x512 by 24 bit (RGB) color.

² & ³ Herman, GT. "Image reconstruction for projections", Academic Press. 1980.

The algorithm used to detect defects is shown in the diagram above. The raw image is filtered to remove noise and to smooth the variation in color and brightness created by the part being at different angles on the conveyor.

The image is converted to HLS so it can be converted to monochrome by two 16 bit input 8 bit output LUTs (look up tables). HLS was selected as the H (hue) and S (saturation) components determine the color, while the L (brightness or lightness) component is primarily determined by the orientation of the surfaces of the part relative to the illumination. These two components are encoded in the resulting 8 bit monochrome image (M in the diagram). This step is important as it is used to detect color errors and gross orientation errors.

The thresholded image is analyzed for connectivity, after being de-speckled to reduce noise. The resulting connective is used to select regions from the monochrome and color images for additional feature extraction. This step reduces the number of pixels being considered by 70% typically.

The regions are measured in monochrome, and color developing features for each region. The features extracted are the color, corrected by the monochrome image, the bonding box, the bounding circle, the perimeter, the convex hull, and the area.

These features are used as inputs to a neural network discriminator. A neural network discriminator was selected because of the complexity of the decision region. Statistical discriminators were found to be too difficult to calculate, and were very susceptible to noise. The neural network used has 100 inputs, 200 nodes in the first hidden layer, 100 nodes in the second layer, and one node in the output layer (pass / fail).

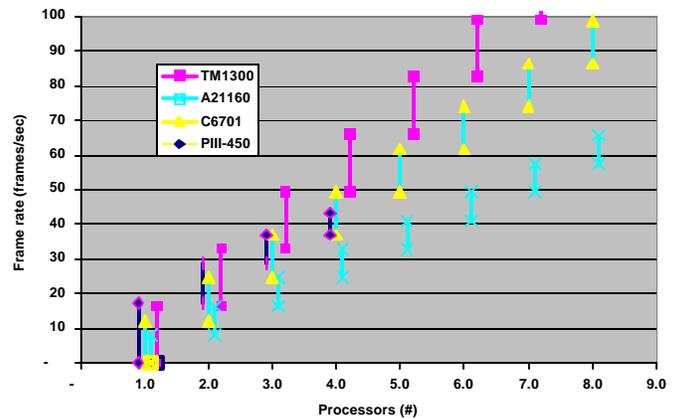
In an effort to select the processing system needed to perform this inspection application, five different processors were examined. Two new Philips TriMedia TM1300 processors were examined. The Analog devices Hammerhead 21160 was also examined, as well as the TI TMS320C6701, and the Intel PIII-450. In each case, code was optimized on each processor and the resulting execution time for a single processor of each type was measured. The table below gives the results of processor on the steps of the algorithm profiles collected.

		Philips	ADI	TI	Intel
Step	Function	TM1300	A21160	C6701	PIII-450
1	Read image from Video Interface	9.83	1.97	1.97	0.98
2	Filtering of the Image	13.54	58.98	35.53	26.65
3	Conversion to HLS	4.81	19.66	11.84	8.74
4	LUT	2.75	3.93	3.93	1.97

	Conversion to Monochrome				
5	Threshold to bi-tonal	0.516	0.737	0.737	0.369
6	De-speckle (3X3)	1.49	1.97	1.97	1.06
7	Connectivity (Blobs)	1.432	1.393	1.393	1.024
8	Masking of Original Color Image	0.842	1.204	1.204	0.602
9	Masking of Monochrome Image	0.43	0.418	0.418	0.307
10	Measurement of regions	1.1	1.58	1.58	0.79
11	Neural Network Evaluation	1.11	2	1.21	1.78
12	Display of defective image	9.83	1.97	1.97	0.98
	Total	47.68	95.81	63.75	45.25

All times are in milliseconds

Upon examining the table above it can be seen that each processor has advantages over the other. The best result overall was obtained by the PIII-450 MHz. Two processors are required to keep up with the image rate from the camera. In all cases an actual system would required an additional processor, to supply operating system support, like disk drives and user interface. Examining the table, the PIII is seen to be an unimpressive computation engine but an excellent processor in applications limited by memory bus performance. The PIII memory bus is twice as fast as the other processors, except the TM1300, which has a memory bus about 1.4 times slower than the PIII (572 MB/s vs 800 MB/s).



The TM1300 processor also requires two processors to keep up with the image rate, but the cost of the processors (on an auxiliary board in the PC), is significantly less than that of a PIII-

450 implementation. The TM1300 product can be obtained for around \$1500 complete, while the multiprocessor PIII system cost several thousand dollars (~\$3000) – more than the basic single processor PC used in the TriMedia solution.

In addition, the efficiency of the PIII begins to suffer as additional processors are added. The SMP (Shared Multiprocessor) bus used in multiprocessor PIII systems reduces the performance of memory intensive applications due to bus contention between the processors. As additional processors are added, the improvement becomes less because of the contention. Systems are not routinely available with more than 4 PIII processors.

Because of the special hardware needed to interface to the camera, the cost of solutions based on the 21160 and the TMS320C6701 are significantly higher than the TriMedia solution and approaching that of the PIII solution. The 21160 and C6701 fall short because of their performance in the compute intensive parts of the application (filtering and conversion). These two processors are limited by the number of processing elements they contain.

The graph above demonstrates the frame rates achievable as the number of processors increases. As can be seen in the graph, the TM1300 performs about as well as the PIII-450 in this application. The TMS320C6701 is a good deal slower, while the ADI21160 is by far the slowest. These two processors both suffer from a lack of processing elements.

Conclusions

An example application was presented illustrating how memory bandwidth limits the performance of multiprocessor systems – whereas memory bus saturation severely limits the scalability of

cluster based architectures (i.e. PIII-450), and local memory architectures allow throughput to scale linearly with the number of processors.

Notably an excellent processor, the Intel PIII is limited by its surrounding logic (the PC) and is unable to perform in some applications. Although the use of the AGP bus would improve the situation, its SMP design will ultimately limit its scalability. Therefore the most practical solution for demanding application remains a co-processor board that is more scalable, has higher throughput, and ultimately is cheaper than the native solution. Among these the Philips TM1300 stands out in the performance versus price curves.

*Correspondence: Email: joe@alacron.com;
<http://www.alacron.com>; Fax (603) 891-2750



71 Spit Brook Road, Suite 200
Nashua, NH 03060

Tel: 603.891.2750 Fax: 603.891.2745

Web: www.alacron.com E-mail: sales@alacron.com

© Alacron, Inc. All Rights Reserved.