

Syracuse University

SURFACE

Electrical Engineering and Computer Science -
Technical Reports

College of Engineering and Computer Science

1-1992

Benchmarking the CM-5 for Image Processing Applications

Ravi V. Shankar
Syracuse University

Ravi Ponnusamy
Syracuse University

Sanjay Ranka
Syracuse University

Follow this and additional works at: https://surface.syr.edu/eecs_techreports



Part of the [Computer Sciences Commons](#)

Recommended Citation

Shankar, Ravi V.; Ponnusamy, Ravi; and Ranka, Sanjay, "Benchmarking the CM-5 for Image Processing Applications" (1992). *Electrical Engineering and Computer Science - Technical Reports*. 180.
https://surface.syr.edu/eecs_techreports/180

This Report is brought to you for free and open access by the College of Engineering and Computer Science at SURFACE. It has been accepted for inclusion in Electrical Engineering and Computer Science - Technical Reports by an authorized administrator of SURFACE. For more information, please contact surface@syr.edu.

SU-CIS-92-01

***Benchmarking the CM-5
for Image Processing Applications***

Ravi Shankar, Ravi Ponnusamy, and Sanjay Ranka

January, 1992

*School of Computer and Information Science
Syracuse University
Suite 4-116, Center for Science and Technology
Syracuse, New York 13244-4100*

Benchmarking the CM-5 for Image Processing Applications

Ravi Shankar Ravi Ponnusamy Sanjay Ranka
School of Computer and Information Science
Syracuse University, Syracuse, NY 13244-4100
e-mail: rshankar, ravi, ranka@top.cis.syr.edu

January 16, 1992

Abstract

This paper presents benchmarking results for image processing algorithms on the Connection Machine model CM-5 and compares them with the results from the CM-2 and the Sun-4. Image processing algorithms with varying communication and computational requirements were implemented, tested and timed. The performance and the scalability of the CM-5 were analyzed and compared with that of the CM-2.

Keywords – Benchmarking – Image Processing – Connection Machine – Performance Analysis – Scalability

1 Introduction

This paper presents benchmarking results for common image processing algorithms on the Connection Machine¹ model CM-5 and compares them with the results from the CM-2 and the Sun-4.

2 The Machines

The Connection Machine model CM-5 [1] is a scalable parallel machine. Each node on the CM-5 can operate at 32 Mips and is accelerated by four optional vector pipes with a peak performance of 32 MFlops. Each of these vector pipes is connected by a 64-bit path to the 32 Mbyte memory.

¹The Connection Machine, CM-5, and CM-2 are registered trademarks of Thinking Machines Corporation

Each CM-5 node is thus capable of 128 MFlops of peak 64-bit performance. The nodes can be organized into a single partition or multiple partitions. The partition manager manages the allocation of parallel resources.

The Connection Machine model CM-2 [2] is an SIMD machine based on a hypercube architecture. Each hypercube node has a cluster of bit-serial processors and a floating point unit. The total number of processors in the CM-2 is usually between 8192 and 65536. The CM-2 provides two forms of communication: the router which allows any processor to communicate with any other processor (Random Access Read/Random Access Write) and the NEWS grid which allows processors to pass data according to a regular rectangular pattern. The advantage of this mechanism over the router is that the overhead of explicitly specifying destination addresses is eliminated.

The image processing algorithms were coded in CM Fortran (the only language currently available on the CM-5 at MSC). The CM implementation used routines from the CM FORTRAN Utility library to perform Random Access Write with collisions.

Timing results from the Sun 4, a sequential machine, are included for comparison. The same amount of effort was spent while programming all three machines (The Sun 4 implementation required a few modifications, because of the unavailability of a Fortran 90 compiler). In particular, the program code was not optimized by hand for either CM architecture. The timings from the Sun 4 are NOT the timing results from the best sequential implementation of the above image processing algorithms. We propose to add these timings later.

While comparing results between the two Connection Machines it should be noted that the CM-2 is an older machine than the CM-5.

The CM-5 at MSC can be configured as two partitions with 512 and 32 nodes, or as three partition with 256, 256, and 32 nodes. The CM-2 at NPAC has four sequencers with 8k processors each (a total of 32k processors) and supports fieldwise computation. Users can attach to either one, two, or all four sequencers at a time. Both the CM-5 as well as the CM-2 support timesharing. Fast file reading/writing are provided on both systems by the DataVault.

3 Benchmarking Results

A variety of image processing algorithms from [3-4] were implemented, tested, and benchmarked on the CM-5, the CM-2, and the Sun4. Timing results from similar algorithms are presented together in the following sections.

3.1 Convolution Based Algorithms

Convolution of an image I of size $N \times N$ with a template W of size $M \times M$ is expressed by the equation

$$C[i, j] = \sum_{u=0}^{M-1} \sum_{v=0}^{M-1} I[(i+u) \bmod N, (j+v) \bmod N] \times W[u, v], \quad 0 \leq i, j < N$$

2-D convolution on a processor array involves repeated NEWS communication.

Sobel edge detection is a special case of 2-D convolution. It is done by convolving the input gray-level image with the following windows:

-1	-2	-1
0	0	0
1	2	1

-1	0	-1
-2	0	-2
-1	0	-1

In order to get the gradient image the convolved images are combined using the absolute value function and addition.

The gradient image returned by the Sobel edge detector has to be thresholded to get edge points. The time taken by thresholding is listed separately for completeness and for comparison purposes. Thresholding does not involve any communication between processors.

Benchmarking results for convolution based algorithms are presented in tables 1 and 2. The time taken by these algorithms is independent of the input image.

3.2 Histogramming Based Algorithms

The histogram H of a given gray-level image I with N gray-levels is an array of size N such that its i -th entry ($0 \leq i < N$) equals the number of pixels in image I with gray-level value i . The histogram indicates the utilization of gray-level values in an image. Two implementations of the histogramming algorithm were benchmarked. The first operated using a single Random Access Write algorithm, where each pixel in the image voted for the bucket labeled with its gray-level value. Collisions occur when more than two or more pixels have the same gray-level value. Collisions should be resolved by addition. The second histogramming implementation sorted the pixels in the image based on the gray-level value, and performed a segmented prefix scan to count the number of pixels with the same value. Results from a single pixel in each segment were combined using a Random Access Write algorithm with no collisions.

Histogram equalization is the process of modifying the histogram of an image to improve the utilization of gray-level values. The equalized histogram is used to enhance the image. This is done by modifying the gray-level values in the image based on the new entries in the histogram following equalization.

Table 1: Convolution based algorithms on a 256x256 image

Algorithm		CM-2 8K	CM-2 16K	CM-5 256P	CM-5 32P	Sun4
3x3 Convolution	Elapsed Time	0.0135	0.0070	0.009	0.062	1.800
	CM Busy Time	0.0134	0.0067	0.009	0.061	
5x5 convolution	Elapsed Time	0.0439	0.0221	0.024	0.163	4.950
	CM Busy Time	0.0437	0.0218	0.023	0.162	
Sobel Edge Detection	Elapsed Time	0.0077	0.0043	0.005	0.028	0.520
	CM Busy Time	0.0077	0.0042	0.005	0.028	

Table 2: Convolution based algorithms on a 512x512 image

Algorithm		CM-2 8K	CM-2 16K	CM-5 256P	CM-5 32P	Sun4
3x3 Convolution	Elapsed time	0.0467	0.0240	0.031	0.237	7.190
	CM busy time	0.0465	0.0239	0.031	0.237	
5x5 Convolution	Elapsed time	0.1443	0.0737	0.082	0.627	19.790
	CM busy time	0.1442	0.0735	0.082	0.627	
Sobel Edge Detection	Elapsed time	0.0252	0.0135	0.015	0.107	2.060
	CM busy time	0.0251	0.0135	0.015	0.107	

Table 3: Histogramming on a 256x256 image

Algorithm		CM-2 8K	CM-2 16K	CM-5 256P	CM-5 32P	Sun4
<i>Image1</i> Histogramming using RAW	Elapsed time	0.0106	0.0060	0.175	0.193	0.090
	CM busy time	0.0104	0.0058	0.175	0.193	
Histogramming using sort	Elapsed time	2.8054	0.1782	0.019	0.118	-
	CM busy time	2.8048	0.1778	0.019	0.117	
<i>Image2</i> Histogramming using RAW	Elapsed time	0.0339	0.0195	0.0050	0.0360	0.100
	CM busy time	0.0339	0.0195	0.0050	0.0360	
Histogramming using sort	Elapsed time	0.0368	0.0208	0.0040	0.0220	-
	CM busy time	0.0362	0.0205	0.0040	0.0220	

Table 4: Histogramming based algorithms on a 512x512 image

Algorithm		CM-2 8K	CM-2 16K	CM-5 256P	CM-5 32P	Sun4
<i>Image1</i>						
Using RAW on 512x512 image	Elapsed time	0.0383	0.0198	2.892	0.796	0.390
	CM busy time	0.0380	0.0196	2.892	0.796	
Using sort on 512x512 image	Elapsed time	7.8439	4.7016	0.071	0.609	-
	CM busy time	7.8426	4.7000	0.068	0.606	
<i>Image2</i>						
Histogramming using RAW	Elapsed time	0.1309	0.0748	0.0190	0.1520	0.380
	CM busy time	0.1309	0.0748	0.0190	0.1520	
Histogramming using sort	Elapsed time	0.1406	0.0811	0.0150	0.0870	-
	CM busy time	0.1398	0.0806	0.0150	0.0870	

Benchmarking results from the histogramming based algorithms are presented in tables 3 and 4. Synthetic images were used as input to the histogramming algorithms since the timings are dependent on the distribution of gray-level values in the input image. Histogramming using RAW and using sort were timed with two images *Image1* and *Image2*. In *Image1* all pixels had the same gray-level value. In *Image2* the pixels had random gray-level values in the range [0,512). Image enhancement and thresholding were also timed using *Image2*. These results are presented in table 5.

The time taken by the histogramming algorithm using RAW and image enhancement was found to vary with the size of the histogram. The dependence on histogram size is shown in the plots in figure 1 to 4. Histogramming using sorting was stable even when the histogram size was changed. Tables 15 through 20 in the appendix give details.

3.3 Image transformations

The following image transformation routines were implemented and benchmarked: scaling, translation, and rotation. Image scaling is done by allowing each pixel in the scaled image to compute the location of the pixel in the original image whose gray-level value it should receive. This computation is followed by a Random Access Read (assuming that the scale factor is ≥ 1) where multiple pixels could read from a single pixel. Image translation is done using NEWS communication alone. Image rotation is similar to image scaling. Each pixel computes the location of the pixel in the original image which supplies its new gray-level value, and receives the new value through a Random Access Read.




Figure 1: Histogram: Dependence on Histogram size for a 256x256 image




Figure 2: Histogram: Dependence on Histogram size for a 512x512 image

Table 5: Image thresholding and enhancement algorithms

Algorithm		CM-2 8K	CM-2 16K	CM-5 256P	CM-5 32P	Sun4
Image size: 256x256						
Image enhancement	Elapsed time	2.7667	0.2313	0.0260	0.1360	0.090
	CM busy time	2.7655	0.2308	0.0250	0.1360	
Thresholding	Elapsed time	0.0011	0.0007	0.000	0.003	0.090
	CM busy time	0.0010	0.0005	0.000	0.003	
Image size: 512x512						
Image enhancement	Elapsed time	8.0935	4.8384	0.0830	0.6410	0.360
	CM busy time	8.0924	4.8374	0.0830	0.6370	
Thresholding	Elapsed time	0.0040	0.0021	0.001	0.010	0.410
	CM busy time	0.0038	0.0019	0.001	0.010	




Figure 3: Enhancing: Dependence on Histogram size for a 256x256 image




Figure 4: Enhancing: Dependence on Histogram size for a 512x512 image

Table 6: Image transformation algorithms on a 256x256 image

Algorithm		CM-2 8K	CM-2 16K	CM-5 256P	CM-5 32P	Sun4
Translation by 100 units	Elapsed time	0.0122	0.0062	0.0010	0.0060	0.080
	CM busy time	0.0121	0.0060	0.0010	0.0060	
Rotation by 75 degrees	Elapsed time	0.0401	0.0219	0.0050	0.0380	0.320
	CM busy time	0.0396	0.0218	0.0050	0.0380	
Scaling by 1.5	Elapsed time	0.0579	0.0304	0.0060	0.0360	0.320
	CM busy time	0.0573	0.0298	0.0060	0.0360	

Benchmarking results for the image transformation algorithms are presented in tables 6 and 7. The time taken by the image transformation algorithms is independent of the input image.

However, the variations for translation and rotation are different compared to scaling. Rotation and translation are one-to-one mappings or nearly one-to-one mappings but for errors due to traslation), while scaling is one-to-many(for scale factor > 1). Thus the variation in translation and rotation is due to different random permutations. The variation in scaling is largely due to the scaling factor. The larger the scaling factor, the lesser the number of sources resulting in more collosions(hot spots). Thus the time increases as the scaling factor increases. This dependence of scale factor in scaling is shown in the plots in figures 5 and 6. Tables 21 through 26 in the appendix give timing results for the image transformation algorithms in detail.

Table 7: Image transformation algorithms on a 512x512 image

Algorithm		CM-2 8K	CM-2 16K	CM-5 256P	CM-5 32P	Sun4
Translation by 100 units	Elapsed time	0.0459	0.0231	0.0030	0.0240	0.340
	CM busy time	0.0458	0.0229	0.0030	0.0240	
Rotation by 75 degrees	Elapsed time	0.1604	0.0859	0.0200	0.1810	1.270
	CM busy time	0.1599	0.0853	0.0200	0.1810	
Scaling by 1.5	Elapsed time	0.2338	0.1205	0.0220	0.1680	1.370
	CM busy time	0.2327	0.1200	0.0220	0.1680	




Figure 5: Scaling: Dependence on scale factor for a 256x256 image




Figure 6: Scaling: Dependence on scale factor for a 512x512 image

Table 8: Time taken by Relaxation

Algorithm		CM-2 8K	CM-2 16K	CM-5 256P	CM-5 32P	Sun4
Relaxation 256x256 image	Elapsed time	0.7106	0.3861	0.745	5.108	1.230
	CM busy time	0.7105	0.3860	0.742	5.108	
Relaxation 512x512 image	Elapsed time	2.4817	1.3036	2.691	32.801	4.890
	CM busy time	2.4816	1.3035	2.690	32.666	

3.4 Relaxation

Relaxation is an iterative algorithm that is used in image processing and numerical analysis for solving a wide variety of problems. Examples of image processing problems that could be solved using relaxation are image segmentation, image labeling, thresholding, edge and curve detection [4]. Each iteration in these image processing algorithms is characterized by data collection from neighboring image points. A relaxation algorithm for image labeling with a fixed number of iterations was benchmarked. Only NEWS communication was used in each iteration. Benchmarking results for relaxation are presented in table 8.

Table 9: Convolution – Scalability of the CM for different problem sizes

Algorithm	256x256 image		512x512 image	
	8K/16K	32p/256p	8K/16K	32p/256p
3x3 Convolution	2.000	6.778	1.946	7.645
5x5 convolution	2.005	7.043	1.962	7.646
Sobel Edge Detection	1.833	5.600	1.859	7.133
Ideal machine	2	8	2	8

Table 10: Convolution – Scalability of the problem for different machine sizes – Time taken by 512x512 image/Time taken by 256x256 image

Algorithm	CM-2		CM-5	
	8K	16K	256p	32p
3x3 Convolution	3.470	3.567	3.444	3.885
5x5 convolution	3.300	3.372	3.565	3.870
Sobel Edge Detection	3.260	3.214	3.000	3.821
Ideal problem	4	4	4	4

4 Analysis of Timing Results

In this section we present an analysis of the timing results presented in the previous sections. The timing results in all the tables in this paper are in seconds. The CM elapsed time and CM busy time are included for all Connection machine timings.

Processing Speed: The convolution based algorithms perform intensive computations in addition to communication. Hence they provide the best base for analyzing the CM processing speed.

Table 9 indicates the speedup observed when the number of CM processors was increased. Table 10 gives the factor by which computation was speeded up when input image size was decreased. Results for the ideal machine and the ideal problem are included in both tables. Table 11 directly compares the performance of the two Connection Machines. All the ratios in these three tables were computed from the CM busy time.

The 256 processor CM-5 was found to be comparable in processing speed with the 16K CM-2 processors (based on convolution, and relaxation). For image processing algorithms that require

Table 11: Convolution – Relative CM performance

Algorithm	32p/8K	32p/16K	256p/8K	256p/16K
256x256 image				
3x3 Convolution	4.552	9.104	0.672	1.343
5x5 convolution	3.707	7.431	0.526	1.055
Sobel Edge Detection	3.636	6.667	0.649	1.190
512x512 image				
3x3 Convolution	5.097	9.916	0.667	1.297
5x5 Convolution	4.348	8.531	0.569	1.116
Sobel Edge Detection	4.263	7.926	0.598	1.111

minimal communication (convolution), CM-5 with 256 processors was found to be approximately 200 times faster than the Sun 4 (tables 1 and 2).

In tables 9 to 11 8K, 16K, 256p, and 32p indicates the time taken by the algorithm on a CM-2 with 8K and 16K processors, and a CM-5 with 256 and 32 processors respectively. The notation 8K/16K therefore indicates the speedup when moving from 8K processors to 16K processors on the CM-2.

Random Communication: For applications which required random communication (scaling, rotation, translation, histogramming, enhancement) the 256 processor CM-5 was found to be a factor of three to ten faster than the 16K processor CM-2. The communication time for scaling, rotation and translation can be estimating by subtracting the time when there is no data movement (e.g. scaling by a factor of 1, translation of 0, and rotation by 0°). The CM-5 was faster than CM-2 by a factor of five to ten (tables 12, 13 and 14). Further, the time for random communication on the CM-5 decreases as the number of collisions decreases. Thus, in the case of histogramming the total time decreased considerably as the number of bins (of the histogram) increased. The same effect was seen for scaling where the time increased because of increase of scaling factor – the larger the scaling factor the larger the number of pixels which read from the same pixel. On the other hand, for operations like translation or rotation (one-to-one mapping), the communication times were relatively stable (with minor variations). Similar effects were seen on the CM-2 although to a lesser extent.

Random Communications on the CM-5 were found to be scalable. The applications using random communications (translation, rotation, scaling, histogramming) consistently gave a factor of 5 to 8 improvement from 32 processors to 256 processors. Further, the scalability improves when the granularity is increased (going from a 256×256 image to a 512×512 image).

Table 12: Image Transformation – Scalability of the CM for different problem sizes

Algorithm	256x256 image		512x512 image	
	8K/16K	32p/256p	8K/16K	32p/256p
Translation	2.018	3.000	1.995	–
Rotation	1.582	7.500	1.691	7.625
Scaling	1.880	3.333	1.891	4.889

Table 13: Image Transformation – Scalability of the problem for different machine sizes – Time taken by 512x512 image/Time taken by 256x256 image

Algorithm	CM-2		CM-5	
	8K	16K	256p	32p
Translation	3.761	3.804	1.000	4.000
Rotation	4.111	3.846	4.000	4.067
Scaling	4.115	4.090	3.000	4.400

Table 14: Image Transformation – Relative CM performance

Algorithm	32p/8K	32p/16K	256p/8K	256p/16K
256x256 image				
Translation	0.265	0.536	0.088	0.179
Rotation	1.042	1.648	0.139	0.220
Scaling	0.321	0.602	0.096	0.181
512x512 image				
translation	0.282	0.563	0.024	0.047
Rotation	1.030	1.743	0.135	0.229
Scaling	0.343	0.648	0.070	0.133

Specialized Communication: For applications using NEWS communication, the 256 processor CM-5 was found to be comparable to the 16K processor CM-2 (based on the results of relaxation and convolution). Further, NEWS communication seems to be scalable (comparing the results for these applications for the CM-5 with 256 processors and 32 processors).

General Scalability: Communication (both random and specialized) appear to be scalable. For most experiments presented in this paper, the CM-5 was found to be scalable. More experiments need to be performed to verify the scalability of the CM-5 for other image processing applications.

5 Conclusion

Our main intention in this work was to demonstrate the relative power of the CM-5 and the CM-2 for image processing applications. This was motivated by the widely varying architectures of these two machines. Image processing algorithms with varying communication and computational requirements were implemented, tested and timed. The performance and the scalability of the CM-5 were analyzed and compared with that of the CM-2.

Acknowledgements

The authors would like to acknowledge the computing resources provided by the Minnesota Supercomputer Center and the Northeast Parallel Architectures Center at Syracuse.

References

- [1] *CM-5 Technical Summary*, Thinking Machines Corporation, Cambridge, Massachusetts, 1991.
- [2] *CM-2 Technical Summary*, Thinking Machines Corporation, Cambridge, Massachusetts, 1988.
- [3] Pavlidis, T., *Algorithms for Graphics and Image Processing*, Computer Science Press, 1982.
- [4] Rosenfeld, A., and Kak, A., *Digital Image Processing, Vol I-II*, Academic Press, 1982.

APPENDIX

Table 15: Histogramming using RAW for a 256x256 image

Histogram size		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
128	CM elapsed time	0.0377	0.0208	0.0100	0.0690
	CM busy time	0.0377	0.0208	0.0100	0.0690
256	CM elapsed time	0.0365	0.0208	0.0070	0.0470
	CM busy time	0.0365	0.0208	0.0070	0.0470
384	CM elapsed time	0.0351	0.0204	0.0050	0.0400
	CM busy time	0.0351	0.0204	0.0050	0.0400
512	CM elapsed time	0.0339	0.0195	0.0050	0.0360
	CM busy time	0.0339	0.0195	0.0050	0.0360
640	CM elapsed time	0.0328	0.0194	0.0050	0.0340
	CM busy time	0.0328	0.0194	0.0050	0.0340
768	CM elapsed time	0.0322	0.0189	0.0040	0.0320
	CM busy time	0.0322	0.0189	0.0040	0.0320
896	CM elapsed time	0.0321	0.0186	0.0040	0.0310
	CM busy time	0.0321	0.0186	0.0040	0.0310
1024	CM elapsed time	0.0311	0.0187	0.0040	0.0300
	CM busy time	0.0311	0.0187	0.0040	0.0300

Table 16: Histogramming using RAW for a 512x512 image

Histogram size		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
128	CM elapsed time	0.1470	0.0803	0.0380	0.2860
	CM busy time	0.1470	0.0803	0.0380	0.2860
256	CM elapsed time	0.1428	0.0778	0.0260	0.1980
	CM busy time	0.1428	0.0778	0.0260	0.1980
384	CM elapsed time	0.1372	0.0781	0.0210	0.1670
	CM busy time	0.1372	0.0781	0.0210	0.1670
512	CM elapsed time	0.1309	0.0748	0.0190	0.1520
	CM busy time	0.1309	0.0748	0.0190	0.1520
640	CM elapsed time	0.1286	0.0733	0.0180	0.1430
	CM busy time	0.1286	0.0733	0.0180	0.1430
768	CM elapsed time	0.1262	0.0725	0.0170	0.1370
	CM busy time	0.1262	0.0725	0.0170	0.1370
896	CM elapsed time	0.1236	0.0714	0.0160	0.1340
	CM busy time	0.1236	0.0714	0.0160	0.1340
1024	CM elapsed time	0.1219	0.0697	0.0160	0.1290
	CM busy time	0.1219	0.0697	0.0160	0.1290

Table 17: Enhancing a 256x256 image

Histogram size		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
128	CM elapsed time	0.0407	0.0229	0.0110	0.0490
	CM busy time	0.0402	0.0225	0.0110	0.0490
256	CM elapsed time	0.0391	0.0228	0.0070	0.0310
	CM busy time	0.0388	0.0224	0.0070	0.0310
384	CM elapsed time	0.0378	0.0224	0.0050	0.0250
	CM busy time	0.0373	0.0222	0.0050	0.0250
512	CM elapsed time	0.0368	0.0208	0.0040	0.0220
	CM busy time	0.0362	0.0205	0.0040	0.0220
640	CM elapsed time	0.0350	0.0205	0.0030	0.0190
	CM busy time	0.0347	0.0202	0.0030	0.0190
768	CM elapsed time	0.0357	0.0203	0.0030	0.0180
	CM busy time	0.0353	0.0200	0.0030	0.0180
896	CM elapsed time	0.0336	0.0203	0.0030	0.0160
	CM busy time	0.0332	0.0200	0.0030	0.0160
1024	CM elapsed time	0.0324	0.0200	0.0020	0.0150
	CM busy time	0.0321	0.0197	0.0020	0.0150

Table 18: Enhancing a 512x512 image

Histogram size		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
128	CM elapsed time	0.1576	0.0886	0.0420	0.1960
	CM busy time	0.1570	0.0880	0.0420	0.1960
256	CM elapsed time	0.1528	0.0862	0.0270	0.1260
	CM busy time	0.1521	0.0856	0.0270	0.1260
384	CM elapsed time	0.1479	0.0852	0.0190	0.1010
	CM busy time	0.1471	0.0847	0.0190	0.1010
512	CM elapsed time	0.1406	0.0811	0.0150	0.0870
	CM busy time	0.1398	0.0806	0.0150	0.0870
640	CM elapsed time	0.1374	0.0796	0.0130	0.0800
	CM busy time	0.1367	0.0790	0.0130	0.0800
768	CM elapsed time	0.1357	0.0785	0.0120	0.0720
	CM busy time	0.1352	0.0779	0.0120	0.0720
896	CM elapsed time	0.1313	0.0764	0.0110	0.0670
	CM busy time	0.1308	0.0758	0.0110	0.0660
1024	CM elapsed time	0.1301	0.0760	0.0090	0.0600
	CM busy time	0.1296	0.0755	0.0090	0.0600

Table 19: Histogramming using sort/count for a 256x256 image

Histogram size		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
128	CM elapsed time	2.8017	0.2286	0.0280	0.1660
	CM busy time	2.7627	0.2271	0.0280	0.1660
256	CM elapsed time	2.7687	0.2278	0.0260	0.1450
	CM busy time	2.7570	0.2257	0.0260	0.1440
384	CM elapsed time	2.7713	0.2318	0.0260	0.1410
	CM busy time	2.7692	0.2312	0.0260	0.1410
512	CM elapsed time	2.7667	0.2313	0.0260	0.1360
	CM busy time	2.7655	0.2308	0.0250	0.1360
640	CM elapsed time	2.7663	0.2291	0.0250	0.1400
	CM busy time	2.7635	0.2286	0.0250	0.1400
768	CM elapsed time	2.7981	0.2300	0.0250	0.1380
	CM busy time	2.7626	0.2296	0.0250	0.1380
896	CM elapsed time	2.7626	0.2283	0.0250	0.1370
	CM busy time	2.7614	0.2277	0.0250	0.1370
1024	CM elapsed time	2.7637	0.2290	0.0250	0.1370
	CM busy time	2.7627	0.2287	0.0240	0.1360
Mean CM busy time		2.7631	0.2287	0.0255	0.1423
Std dev of CM busy time		0.0011	0.0006	0.0004	0.0033

Table 20: Histogramming using sort/count for a 512x512 image

Histogram size		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
128	CM elapsed time	9.0123	5.0414	0.0990	0.7740
	CM busy time	8.0966	4.8386	0.0980	0.7670
256	CM elapsed time	9.1276	4.8679	0.0880	0.6840
	CM busy time	8.0546	4.8292	0.0880	0.6810
384	CM elapsed time	8.0960	4.8427	0.0850	0.6570
	CM busy time	8.0948	4.8414	0.0850	0.6530
512	CM elapsed time	8.0935	4.8384	0.0830	0.6410
	CM busy time	8.0924	4.8374	0.0830	0.6370
640	CM elapsed time	8.0795	4.8363	0.0870	0.6370
	CM busy time	8.0783	4.8349	0.0870	0.6330
768	CM elapsed time	8.1868	4.8380	0.0870	0.6330
	CM busy time	8.0777	4.8333	0.0860	0.6270
896	CM elapsed time	8.0706	4.8335	0.0860	0.6290
	CM busy time	8.0694	4.8308	0.0860	0.6230
1024	CM elapsed time	8.0892	4.8336	0.0850	0.6210
	CM busy time	8.0741	4.8324	0.0850	0.6210
Mean CM busy time		8.0797	4.8347	0.0873	0.6552
Std dev of CM busy time		0.0048	0.0014	0.0015	0.0163

Table 21: Time taken to translate a 256x256 image

Translation		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
0	CM elapsed time	0.0010	0.0006	0.0010	0.0030
	CM busy time	0.0008	0.0004	0.0000	0.0030
20	CM elapsed time	0.0072	0.0037	0.0010	0.0050
	CM busy time	0.0070	0.0035	0.0010	0.0050
40	CM elapsed time	0.0081	0.0042	0.0010	0.0060
	CM busy time	0.0080	0.0040	0.0010	0.0060
60	CM elapsed time	0.0123	0.0062	0.0010	0.0060
	CM busy time	0.0121	0.0060	0.0010	0.0060
80	CM elapsed time	0.0106	0.0054	0.0010	0.0060
	CM busy time	0.0105	0.0052	0.0010	0.0060
100	CM elapsed time	0.0122	0.0062	0.0010	0.0060
	CM busy time	0.0121	0.0060	0.0010	0.0060
120	CM elapsed time	0.0082	0.0042	0.0010	0.0060
	CM busy time	0.0081	0.0040	0.0010	0.0060
140	CM elapsed time	0.0097	0.0050	0.0010	0.0060
	CM busy time	0.0095	0.0047	0.0010	0.0060
160	CM elapsed time	0.0105	0.0053	0.0010	0.0060
	CM busy time	0.0104	0.0051	0.0010	0.0060
180	CM elapsed time	0.0099	0.0050	0.0010	0.0060
	CM busy time	0.0097	0.0048	0.0010	0.0060
200	CM elapsed time	0.0131	0.0066	0.0010	0.0060
	CM busy time	0.0129	0.0064	0.0010	0.0060
Mean CM busy time		0.0092	0.0045	0.0009	0.0056
Std dev of CM busy time		0.0010	0.0005	0.0001	0.0003

Table 22: Time taken to translate a 512x512 image

Translation		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
0	CM elapsed time	0.0034	0.0018	0.0020	0.0120
	CM busy time	0.0033	0.0016	0.0020	0.0120
20	CM elapsed time	0.0167	0.0084	0.0030	0.0160
	CM busy time	0.0165	0.0083	0.0030	0.0160
40	CM elapsed time	0.0282	0.0142	0.0030	0.0190
	CM busy time	0.0280	0.0140	0.0030	0.0190
60	CM elapsed time	0.0423	0.0213	0.0030	0.0230
	CM busy time	0.0421	0.0211	0.0030	0.0230
80	CM elapsed time	0.0319	0.0161	0.0030	0.0240
	CM busy time	0.0318	0.0159	0.0030	0.0240
100	CM elapsed time	0.0459	0.0231	0.0030	0.0240
	CM busy time	0.0458	0.0229	0.0030	0.0240
120	CM elapsed time	0.0484	0.0243	0.0030	0.0240
	CM busy time	0.0482	0.0242	0.0030	0.0240
140	CM elapsed time	0.0330	0.0166	0.0030	0.0240
	CM busy time	0.0329	0.0165	0.0030	0.0240
160	CM elapsed time	0.0420	0.0211	0.0030	0.0250
	CM busy time	0.0418	0.0209	0.0030	0.0250
180	CM elapsed time	0.0561	0.0282	0.0030	0.0260
	CM busy time	0.0560	0.0280	0.0030	0.0260
200	CM elapsed time	0.0483	0.0242	0.0030	0.0240
	CM busy time	0.0481	0.0241	0.0030	0.0240
Mean CM busy time		0.0359	0.0180	0.0029	0.0219
Std dev of CM busy time		0.0045	0.0022	0.0001	0.0013

Table 23: Time taken to rotate a 256x256 image

Rotation		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
0	CM elapsed time	0.0254	0.0129	0.0030	0.0230
	CM busy time	0.0252	0.0127	0.0030	0.0230
15	CM elapsed time	0.0378	0.0203	0.0050	0.0320
	CM busy time	0.0373	0.0202	0.0050	0.0320
30	CM elapsed time	0.0396	0.0205	0.0050	0.0340
	CM busy time	0.0391	0.0203	0.0050	0.0340
45	CM elapsed time	0.0414	0.0211	0.0050	0.0370
	CM busy time	0.0409	0.0210	0.0050	0.0370
60	CM elapsed time	0.0392	0.0214	0.0050	0.0380
	CM busy time	0.0386	0.0213	0.0050	0.0380
75	CM elapsed time	0.0401	0.0219	0.0050	0.0380
	CM busy time	0.0396	0.0218	0.0050	0.0380
90	CM elapsed time	0.0358	0.0179	0.0040	0.0350
	CM busy time	0.0354	0.0177	0.0040	0.0350
105	CM elapsed time	0.0421	0.0210	0.0050	0.0380
	CM busy time	0.0416	0.0208	0.0050	0.0380
120	CM elapsed time	0.0399	0.0209	0.0050	0.0380
	CM busy time	0.0395	0.0208	0.0050	0.0380
135	CM elapsed time	0.0419	0.0216	0.0060	0.0370
	CM busy time	0.0413	0.0215	0.0060	0.0370
150	CM elapsed time	0.0397	0.0200	0.0050	0.0360
	CM busy time	0.0391	0.0198	0.0050	0.0360
165	CM elapsed time	0.0430	0.0219	0.0050	0.0340
	CM busy time	0.0424	0.0218	0.0050	0.0340
180	CM elapsed time	0.0366	0.0184	0.0040	0.0280
	CM busy time	0.0360	0.0182	0.0040	0.0280
Mean CM busy time		0.0382	0.0198	0.0048	0.0345
Std dev of CM busy time		0.0012	0.0007	0.0002	0.0012

Table 24: Time taken to rotate a 512x512 image

Rotation		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
0	CM elapsed time	0.1012	0.0509	0.0120	0.1200
	CM busy time	0.1007	0.0503	0.0120	0.1200
15	CM elapsed time	0.1515	0.0791	0.0200	0.1560
	CM busy time	0.1510	0.0786	0.0200	0.1560
30	CM elapsed time	0.1590	0.0802	0.0210	0.1670
	CM busy time	0.1587	0.0797	0.0210	0.1670
45	CM elapsed time	0.1638	0.0830	0.0220	0.1790
	CM busy time	0.1630	0.0825	0.0220	0.1790
60	CM elapsed time	0.1553	0.0863	0.0220	0.1820
	CM busy time	0.1548	0.0857	0.0220	0.1820
75	CM elapsed time	0.1604	0.0859	0.0200	0.1810
	CM busy time	0.1599	0.0853	0.0200	0.1810
90	CM elapsed time	0.1407	0.0705	0.0170	0.1670
	CM busy time	0.1402	0.0700	0.0170	0.1670
105	CM elapsed time	0.1631	0.0842	0.0200	0.1810
	CM busy time	0.1625	0.0837	0.0200	0.1810
120	CM elapsed time	0.1575	0.0810	0.0220	0.1840
	CM busy time	0.1569	0.0806	0.0220	0.1840
135	CM elapsed time	0.1629	0.0831	0.0230	0.1810
	CM busy time	0.1622	0.0825	0.0230	0.1810
150	CM elapsed time	0.1573	0.0796	0.0210	0.1770
	CM busy time	0.1567	0.0791	0.0210	0.1770
165	CM elapsed time	0.1685	0.0861	0.0210	0.1690
	CM busy time	0.1680	0.0855	0.0210	0.1690
180	CM elapsed time	0.1447	0.0723	0.0150	0.1400
	CM busy time	0.1440	0.0718	0.0150	0.1400
Mean CM busy time		0.1522	0.0781	0.0197	0.1680
Std dev of CM busy time		0.0046	0.0026	0.0008	0.0051

Table 25: Time taken to scale a 256x256 image

Scale factor		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
1.0	CM elapsed time	0.0263	0.0133	0.0030	0.0260
	CM busy time	0.0261	0.0132	0.0030	0.0260
1.1	CM elapsed time	0.0422	0.0216	0.0040	0.0280
	CM busy time	0.0417	0.0214	0.0040	0.0280
1.2	CM elapsed time	0.0477	0.0238	0.0050	0.0300
	CM busy time	0.0472	0.0237	0.0050	0.0300
1.3	CM elapsed time	0.0538	0.0269	0.0050	0.0330
	CM busy time	0.0533	0.0268	0.0050	0.0330
1.4	CM elapsed time	0.0557	0.0290	0.0060	0.0350
	CM busy time	0.0552	0.0289	0.0060	0.0350
1.5	CM elapsed time	0.0579	0.0304	0.0060	0.0360
	CM busy time	0.0573	0.0298	0.0060	0.0360
1.6	CM elapsed time	0.0570	0.0300	0.0060	0.0370
	CM busy time	0.0565	0.0298	0.0060	0.0370
1.7	CM elapsed time	0.0600	0.0314	0.0060	0.0390
	CM busy time	0.0595	0.0311	0.0060	0.0390
1.8	CM elapsed time	0.0641	0.0331	0.0060	0.0400
	CM busy time	0.0635	0.0327	0.0060	0.0400
1.9	CM elapsed time	0.0678	0.0373	0.0070	0.0410
	CM busy time	0.0673	0.0368	0.0070	0.0410
2.0	CM elapsed time	0.0729	0.0404	0.0070	0.0410
	CM busy time	0.0724	0.0398	0.0070	0.0410

Table 26: Time taken to scale a 512x512 image

Scale factor		CM-2 8K	CM-2 16K	CM-5 256p	CM-5 32p
1.0	CM elapsed time	0.1048	0.0527	0.0130	0.1240
	CM busy time	0.1043	0.0521	0.0130	0.1240
1.1	CM elapsed time	0.1678	0.0841	0.0170	0.1310
	CM busy time	0.1672	0.0835	0.0170	0.1310
1.2	CM elapsed time	0.1925	0.0961	0.0190	0.1430
	CM busy time	0.1920	0.0955	0.0190	0.1430
1.3	CM elapsed time	0.2250	0.1111	0.0200	0.1560
	CM busy time	0.2244	0.1106	0.0200	0.1560
1.4	CM elapsed time	0.2286	0.1164	0.0220	0.1630
	CM busy time	0.2279	0.1159	0.0220	0.1630
1.5	CM elapsed time	0.2338	0.1205	0.0220	0.1680
	CM busy time	0.2327	0.1200	0.0220	0.1680
1.6	CM elapsed time	0.2329	0.1235	0.0230	0.1730
	CM busy time	0.2322	0.1230	0.0230	0.1730
1.7	CM elapsed time	0.2371	0.1260	0.0240	0.1800
	CM busy time	0.2364	0.1254	0.0240	0.1800
1.8	CM elapsed time	0.2525	0.1303	0.0250	0.1850
	CM busy time	0.2518	0.1297	0.0250	0.1850
1.9	CM elapsed time	0.2743	0.1440	0.0260	0.1900
	CM busy time	0.2738	0.1434	0.0260	0.1900
2.0	CM elapsed time	0.2968	0.1560	0.0250	0.1850
	CM busy time	0.2964	0.1555	0.0250	0.1850