Apple Silicon Performance in Scientific Computing

Connor Kenyon and Collin Capano, Center for Scientific Computation and Data Science Research

Abstract—With the release of the Apple Silicon Systemon-a-Chip processors, and the impressive performance shown in general use by both the M1 and M1 Ultra, the potential use for Apple Silicon processors in scientific computing is explored. Both the M1 and M1 Ultra are compared to current state-of-the-art data-center GPUs, including an NVIDIA V100 with PCIe, an NVIDIA V100 with NVLink, and an NVIDIA A100 with PCIe. The scientific performance is measured using the Scalable Heterogeneous Computing (SHOC) benchmark suite using OpenCL benchmarks. We find that both M1 processors out perform the GPUs in all benchmarks.

Index Terms—Scientific Computing, Embedded Devices, Accelerators, Prallel Computing, Supercomputing, GPU, GPGPU, OpenCL, SHOC, Physics

I. INTRODUCTION

RECENT trends in high performance computing (HPC) have shown a significant shift in types of processors and architectures being utilized in clusters. There is a transition from discrete CPU and GPU systems to systems that include heterogeneous architectures. This shift is most visible within the Top500 supercomputer list, in which the top 3 as of June 2022 - Frontier, Fugaku, and Lumi - all take advantage of high core-count heterogeneous processors [1]. The main benefit of heterogeneous architectures is the ability to utilize shared memory in order to minimize or entirely eliminate data transfer times. This is done through zerocopy algorithms, or algorithms that allow for memory access from both the CPU and GPU without requiring data transfer [2]. Removing communication bottlenecks in this way has been shown to improve performance by as much as a factor of 30 in data-center GPUs [3].

With the introduction of the Apple line of System-ona-Chip (SoC) processors, Apple Silicon, there is now an additional manufacturer of heterogeneous architectures [4]. The competitive processors from Apple include their M1 series, as well as the newly announced M2 series. The first M1 processor was released in November 2020, followed by the M1 Pro and M1 Max in October of 2021, and the M1 Ultra in March of 2022 [5]. Initial CPU and GPU benchmarks showed significant promise in the overall performance capabilities of the M1 [6]–[8]. With their relatively low cost, low power consumption, and the general trend of HPC toward SoC processors, the potential for using these new Apple Silicon processors in scientific computing applications is manifold.

This paper is organized as follows: in Section II specific details about the tested hardware are discussed; in Section III the benchmarks are introduced and explained; results are shown in Section IV; concluding remarks are in Section V.

II. HARDWARE DETAILS

The hardware being compared in this work represents current state-of-the-art HPC GPUs alongside two of the Apple silicon processors. This allows for a comparison to be made between performance that would be obtainable on a standard compute node in an HPC cluster and a Mac computer featuring one of the Apple Silicon processors.

For data-center GPUs, we use two NVIDIA Tesla V100s [9] — one with PCIe interconnect and one with the SXM2 interconnect — along with an NVIDIA Tesla A100 [10]. These represent the previous and current generations of NVIDIA data-center GPUs, respectively. They have been a staple in HPC clusters since their releases. The specifications for these are outlined in Table I.

We benchmark two Apple Silicon processors: the firstgeneration M1 [11] and the M1 Ultra [12], the latter of which is the the most recent and top-performing processor in the M1 series. Specs for both are shown in Table II.

This article was submitted for peer-review on July 15, 2022. This work was supported in part by NSF Grant No. DMS-19127165.

Connor Kenyon is a member of the Physics Department and the Center for Scientific Computing and Data Science Research at the University of Massachusetts Dartmouth, North Dartmouth, MA 02747 USA (e-mail: ckenyon@umassd.edu).

Dr. Collin Capano is member of the Physics Department and the Center for Scientific Computing and Data Science Research at the University of Massachusetts Dartmouth, North Dartmouth, MA 02747 USA, and a senior scientist with the Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), D-30167 Hannover, Germany (e-mail: ccapano@umassd.edu).

TABLE I HARDWARE DETAILS FOR NVIDIA V100 AND A100 GPUS.

	V100	A100
Memory	32 GB HBM2	80 GB HBM2
Memory Bandwidth	900 GB/s	1,935 GB/s
TDP	300W	300W
Interconnect	PCIe 3.0 / NVLink	PCIe 4.0

 TABLE II

 HARDWARE DETAILS FOR APPLE M1 AND M1 ULTRA.

	M1	M1 Ultra
Memory	16 GB LPDDR4	32 GB LPDDR5
Memory Bandwidth	4,226 MB/s	800 GB/s
TDP	15 W	60 W
Interconnect	N/A	N/A

III. SHOC BENCHMARKS

The SHOC benchmark suite is an open source collection of benchmark programs selected to measure the performance of a variety of computing devices with a focus on multi-core performance tests. The benchmarks are available in multi-core MPI, OpenCL, and CUDA, which provides support for most modern heterogeneous computing hardware [13]. In this work, we benchmark the Nvidia GPUs using the CUDA benchmarks, and the Apple M1s using the OpenCL benchmarks.

To subdivide the particular benchmarks, there are three different performance measurement levels: level 0 measures the low level behavior; level 1 measures higher level behavior with commonly used algorithms; level 2 measures realistic application kernels. Specifically, level 0 benchmarks include:

- Bus Speed Download;
- Bus Speed Readback;
- Device Memory;
- Kernel Compile;
- Max Flops;
- Queue Delay.

Level 1 benchmarks are:

- Breadth-First Search (BFS);
- Fast Fourier Transform (FFT);
- Molecular Dynamics (MD);
- MD5Hash;
- Reduction;
- Single-Precision Matrix Multiplication (SGEMM);
- Scan;
- Sort;
- Sparse matrix-vector multiplication (SPMV);
- Stencil2D;
- Triad.

Level 2 benchmarks include Quality Threshold clustering (QTC) and S3D.

IV. RESULTS

The performance of all processors being compared in various benchmarks allows for a broad comparison of the capabilities of each device for scientific computing. We show the results of SHOC benchmarks in single precision for all five processors. Single precision benchmarks are selected because of the lack of double precision capability in the Apple M1s' GPU.

For the level 0 Max FLOPs benchmark, we find peak single precision performance is lowest with the V100 with PCIe $(1.40 \times 10^4 \text{ GFLOPs})$, followed by the V100 with NVLink $(1.55 \times 10^4 \text{ GFLOPs})$. The A100 outperforms both V100s at $1.93 \times 10^4 \text{ GFLOPs}$. However, both the Apple M1 and M1 Ultra dwarf this figure, at 2.97×10^8 and $6.56 \times 10^8 \text{ GLOPs}$ peak performance, respectively. The reported performance for the M1s, while impressive, seems significantly larger than plausible.

In Figure 1, a subset of the level 1 SHOC benchmarks are shown. This subset includes Matrix Multiplication (GEMM), Sparse matrix-vector multiplication (SPMV), and Fast Fourier Transform (FFT). The performance shown in Figure 1 includes hatched bars for the M1 and M1 Ultra because they are able to utilize shared memory. The hatched portion indicates the performance with transfer time, while the portion without hatches shows the performance with zero-copy algorithms being used. The real performance when optimally using the hardware is reflected by the bars without hatches. These hatched bars are used similarly in Figures 2 and 3

The NVIDIA GPUs performance varies widely between each of the SHOC level 1 benchmarks shown in Figure 1, with the V100 with NVLink consistently outperforming both the V100 and the A100 with PCIe. The performance is strongest on matrix multiplication, where the NVIDIA GPUs all perform competitively with the Apple M1, while still falling an order of magnitude behind the M1 Ultra. For FFTs, both the Apple M1 and M1 Ultra maintain approximately an order of magnitude of separation in performance, while outperforming all three GPUs. In SPMV this difference is seen even more dramatically, with a separation of two orders of magnitude between the V100 with NVLink and the M1, with the M1 Ultra performing beyond that.

Benchmark results for S3D are shown in Figure 2. This is the largest benchmark problem featured in the SHOC benchmarks, and reflects the realistic performance of the processors when used with typical scientific code. Figure 2 compares all processors, including the NVIDIA V100 and A100 as well as the Apple M1 and M1 Ultra.

The results for S3D shown in Figure 2 show a similar pattern to the SPMV benchmark seen in Figure 1. The most notable change is that the A100 performs slightly



Fig. 1. Set of SHOC Level 1 benchmark results showing Generic Matrix Multiplication (GEMM), Sparse Matrix Vector Multiplication (SPMV), and Fast Fourier Transform (FFT) on all devices. Hatched lines indicate performance with shared memory that includes transfer time.



Fig. 2. Set of SHOC Level 2 benchmark results showing only S3D on all devices.

better than the V100 with NVLink, but performance still falls significantly behind both the M1 and the M1 Ultra. The performance of both Apple processors maintains a significant margin above all three Nvidia GPUs. The difference between the hatched and unhatched bars for the M1 and M1 Ultra show nearly no difference. This reflects a heavily operations-bound curve for S3D, meaning shared memory independently provides no significant benefit to this algorithm.

A pertinent question is how much each processor costs relative to its performance. Comparing costs between the GPUs and the M1s is complicated by the fact that the M1 processors can only be purchased as a part of a complete Mac desktop (which comes with storage) while the Nvidia GPUs can be purchased on their own (without storage). However, as a rough approximation we use the cheapest available Mac Mini with the M1 tested here, and the cheapest available Mac Studio with the M1 Ultra tested here. For the Nvidia GPUs, current available GPU prices are used. Specifically, we estimate the price of each processor as:

- Apple M1: \$899.99 [14]
- Apple M1 Ultra: \$4,999.99 [15]
- Nvidia V100 PCIe: \$3,899.00 [16]
- Nvidia V100 NVLink: \$2,990.00 [17]
- Nvidia A100: \$11,157.19 [18]

The cost per GFLOP is shown in Figure 3 (here, smaller is better).

The cost per GFLOP shown in Figure 3 has the M1 and M1 Ultra performing comparably in each benchmark. In GEMM, the M1 and M1 Ultra cost an order of magnitude less per GFLOP. For SPMV and FFT the difference is more substantial, with the difference between the M1 and M1 Ultra and the Nvidia GPUs increasing to over two orders of magnitude. The Nvidia GPUs maintain a relatively fixed difference in price per GFLOP relative to one another, with the V100 with NVLink consistently the lowest cost per GFLOP, followed by the V100 with PCIe, and then the A100 at the largest cost per GFLOP in all benchmarks shown. A significant highlight of these benchmarks is that by optimally utilizing shared memory with data transfer-dominated algorithms, such as FFT, there is a two order of magnitude decrease in price per GFLOP of performance.

V. CONCLUSIONS

Recent trends in high performance computing have shown a significant shift towards heterogeneous system architectures, often including ARM processors as an integral component of the system. Previously, this market has been dominated by Nvidia and AMD. The Apple silicon processors introduce new competition.

While there are some features that the Apple M1 and M1 Ultra lack, such as double precision GPU computing

978-1-6654-9786-2/22/\$31.00 ©2022 IEEE



Fig. 3. The cost per GFLOP of performance in each processor (lower cost per GFLOP is better), shown for matrix multiplication (GEMM), sparse matrix-vector multiplication (SPMV), and fast Fourier transform (FFT). Hatched lines indicate that the performance when data transfer time is included for a processor with shared memory available.

capabilities, the single precision performance shows an incredible amount of promise. Though currently only available in Mac computers and not sold individually, these processors show impressive performance that is applicable to many machine learning or single precision-dominated research in scientific computing. Even when considering cost of purchase, the price per GFLOP of an M1 or M1 Ultra comes at a fraction of the cost of a data-center GPU and accompanying server node.

Heterogeneous system architectures are an extremely promising path forward in high performance computing. Despite the lack of double precision GPU computing in the Apple M1 series, the performance is promising and may see use in single-precision applications in research computing. While there is no sign of these becoming available to use in clusters, they will be able to provide robust single precision computing on a workstation.

The authors would like to thank Gaurav Khanna, Scott Field, and Geoffrey Cowles for useful comments and discussions. This work was supported in part by NSF Grant No. DMS-19127165.

REFERENCES

[1] *Top500 june 2022*. [Online]. Available: https://top500.org/lists/top500/2022/06/.

- [2] J. Song and J. Alves-Foss, "Performance review of zero copy techniques," *International Journal of Computer Science and Security (IJCSS)*, vol. 6, no. 4, p. 256, 2012.
- [3] C. Kenyon, G. Volkema, and G. Khanna, "Overcoming limitations of gpgpu-computing in scientific applications," in 2019 IEEE High Performance Extreme Computing Conference (HPEC), IEEE, 2019, pp. 1–9.
- [4] Apple announces mac transition to apple silicon, Jun. 2020. [Online]. Available: https://www.apple. com/newsroom/2020/06/apple-announces-mactransition-to-apple-silicon/.
- [5] K. Haslam, M1 chip guide: M1, m1 pro, m1 max and m1 ultra, Jun. 2022. [Online]. Available: https://www.macworld.com/article/676321/m1pro-max-ultra-guide.html.
- [6] Apple m1 tested: Performance benchmarks and thermal throttling, explained, Apr. 2022. [Online]. Available: https://www.androidauthority.com/ apple-m1-test-benchmark-performance-thermal-1185988/.
- [7] *M1 ultra cpu geekbench*. [Online]. Available: https://browser.geekbench.com/v5/cpu/14664498.
- [8] M1 ultra gpu geekbench. [Online]. Available: https://browser.geekbench.com/v5/compute/ 4757238.
- [9] N. Oh, Nvidia formally announces pcie tesla v100: Available later this year, Jun. 2017. [Online]. Available: https://www.anandtech.com/show/ 11559/nvidia - formally - announces - pcie - tesla v100-available-later-this-year.
- [10] Nvidia a100 gpus power the modern data center. [Online]. Available: https://www.nvidia.com/enus/data-center/a100/.
- [11] M. Ehrhardt, Apple m1 chip: Specs, performance, everything we know, Jun. 2021. [Online]. Available: https://www.tomshardware.com/news/ Apple-M1-Chip-Everything-We-Know.
- [12] Apple unveils m1 ultra, the world's most powerful chip for a personal computer, Jun. 2022. [Online]. Available: https://www.apple.com/newsroom/ 2022/03/apple - unveils - m1 - ultra - the - worlds most-powerful-chip-for-a-personal-computer/.
- [13] A. Danalis, G. Marin, C. McCurdy, et al., "The scalable heterogeneous computing (shoc) benchmark suite," in Proceedings of the 3rd Workshop on General-Purpose Computation on Graphics Processing Units, 2010, pp. 63–74.
- [14] Mac mini. [Online]. Available: https://www.apple. com/shop/buy-mac/mac-mini/apple-m1-chipwith-8-core-cpu-and-8-core-gpu-256gb.
- [15] *Mac studio*. [Online]. Available: https://www.apple.com/shop/buy-mac/mac-studio/20-core-

978-1-6654-9786-2/22/\$31.00 ©2022 IEEE

cpu-48-core-gpu-32-core-neural-engine-64gbmemory-1tb.

- [16] Newegg nvidia v100 pcie. [Online]. Available: https://www.newegg.com/nvidia-900-2g500-0000-000/p/1DH-00RJ-00001.
- [17] *Newegg nvidia v100 nvlink*. [Online]. Available: https://www.newegg.com/p/1FT-0004-002M1.
- [18] Newegg nvidia a100 pcie. [Online]. Available: https://www.newegg.com/nvidia-900-21001-0000-000/p/N82E16814132090.

APPENDIX A NVIDIA V100 PCIE SHOC SINGLE PRECISION OUTPUT

result	for	bspeed_download:	12.3233	GB/sec
result	for	bspeed_readback:	13.1718	GB/sec
result	for	maxspflops:	14039.2000) GFLOPS
result	for	gmem_readbw:	934.3160	GB/s
result	for	gmem_readbw_strided:	404.3400	GB/s
result	for	gmem_writebw:	725.9060	GB/s
result	for	gmem_writebw_strided:	18.6136	GB/s
result	for	lmem_readbw:	10286.9000) GB/s
result	for	lmem_writebw:	11825.5000) GB/s
result	for	tex_readbw:	1528.0900	GB/sec
result	for	ocl_kernel:	0.0019	sec
result	for	ocl_queue:	0.0034	ms
result	for	bfs:	15.4394	GB/s
result	for	bfs_pcie:	7.5798	GB/s
result	for	bfs_teps:	376338000	.0000 Edges/s
result	for	fft_sp:	2293.4700	GFLOPS
result	for	fft_sp_pcie:	35.2146	GFLOPS
result	for	ifft_sp:	2279.9500	GFLOPS
result	for	ifft_sp_pcie:	35.2113	GFLOPS
result	for	sgemm_n:	9353.9900	GFLOPS
result	for	sgemm_t:	9233.5700	GFLOPS
result	for	sgemm_n_pcie:	3795.6600	GFLOPS
result	for	sgemm_t_pcie:	3776.0300	GFLOPS
result	for	md_sp_flops:	998.9960	GFLOPS
result	for	md_sp_bw:	765.6000	GB/s
result	for	<pre>md_sp_flops_pcie:</pre>	57.1860	GFLOPS
result	for	md_sp_bw_pcie:	43.8256	GB/s
result	for	md5hash:	31.0101	GHash/s
result	for	reduction:	302.1470	GB/s
result	for	reduction_pcie:	11.8165	GB/s
result	for	scan:	85.5936	GB/s
result	for	<pre>scan_pcie:</pre>	5.9156	GB/s
result	for	sort:	1.8335	GB/s
result	for	sort_pcie:	1.4230	GB/s
result	for	<pre>spmv_csr_scalar_sp:</pre>	56.2969	Gflop/s
result	for	<pre>spmv_csr_scalar_sp_pcie:</pre>	1.1735	Gflop/s
result	for	<pre>spmv_csr_scalar_pad_sp:</pre>	57.7069	Gflop/s
result	for	<pre>spmv_csr_scalar_pad_sp_pcie:</pre>	1.0237	Gflop/s
result	for	<pre>spmv_csr_vector_sp:</pre>	88.8624	Gflop/s
result	for	<pre>spmv_csr_vector_sp_pcie:</pre>	1.1824	Gflop/s
result	for	<pre>spmv_csr_vector_pad_sp:</pre>	93.5053	Gflop/s
result	for	<pre>spmv_csr_vector_pad_sp_pcie:</pre>	1.0307	Gflop/s
result	for	<pre>spmv_ellpackr_sp:</pre>	55.6213	Gflop/s
result	for	stencil:	623.3620	GFLOPS
result	for	triad_bw:	12.2925	GB/s
result	for	s3d:	447.3900	GFLOPS
result	for	s3d_pcie:	144.7300	GFLOPS

6

APPENDIX B NVIDIA V100 SXM2 SHOC SINGLE PRECISION OUTPUT

result	for	bspeed_download:	37.5824 GB/sec
result	for	bspeed_readback:	38.8395 GB/sec
result	for	maxspflops:	15516.8000 GFLOPS
result	for	gmem_readbw:	888.3320 GB/s
result	for	gmem_readbw_strided:	479.0020 GB/s
result	for	gmem_writebw:	742.7190 GB/s
result	for	gmem_writebw_strided:	59.8676 GB/s
result	for	lmem_readbw:	9453.9000 GB/s
result	for	lmem_writebw:	10179.5000 GB/s
result	for	tex_readbw:	1512.2300 GB/sec
result	for	bfs:	10.5773 GB/s
result	for	bfs_pcie:	7.3547 GB/s
result	for	bfs_teps:	378866000.0000 Edges/s
result	for	fft_sp:	2278.6600 GFLOPS
result	for	fft_sp_pcie:	175.6960 GFLOPS
result	for	ifft_sp:	2260.2600 GFLOPS
result	for	ifft_sp_pcie:	176.1480 GFLOPS
result	for	sgemm_n:	14643.4000 GFlops
result	for	sgemm_t:	14347.2000 GFlops
result	for	sgemm_n_pcie:	8729.9400 GFlops
result	for	sgemm_t_pcie:	8623.8000 GFlops
result	for	md_sp_flops:	912.5020 GFLOPS
result	for	md sp bw:	699.3130 GB/s
result	for	md sp flops pcie:	132.5160 GFLOPS
result	for	md sp bw pcie:	101.5560 GB/s
result	for	md5hash:	34.7245 GHash/s
result	for	nn learning:	BenchmarkError
result	for	nn_learning_pcie:	BenchmarkError
result	for	reduction:	325.9520 GB/s
result	for	reduction pcie:	34.3715 GB/s
result	for	scan:	198.6420 GB/s
result	for	scan pcie:	17.0488 GB/s
result	for	sort:	21.3906 GB/s
result	for	sort pcie:	9.9923 GB/s
result	for	spmv csr scalar sp:	68.7036 Gflop/s
result	for	spmv csr scalar sp pcie:	6.3301 Gflop/s
result	for	spmv csr scalar pad sp:	77.8411 Gflop/s
result	for	spmv csr scalar pad sp pcie:	6.6088 Gflop/s
result	for	spmv csr vector sp:	164.1990 Gflop/s
result	for	spmy csr vector sp pcie:	6.6855 Gflop/s
result	for	spmv csr vector pad sp:	172.6370 Gflop/s
result.	for	spin csr vector pad sp pcie:	6.9219 Gflop/s
result	for	spmv ellpackr sp:	90.2127 Gflop/s
result	for	stencil:	690.8430 GFLOPS
result	for	triad bw:	36.0896 GB/s
result	for	s3d:	462.9920 GFLOPS
result	for	s3d pcie:	363.5560 GFLOPS

APPENDIX C NVIDIA A100 PCIE SHOC SINGLE PRECISION OUTPUT

result	for	bspeed download:	26.8296 GB/sec
result	for	bspeed readback:	27.1204 GB/sec
result	for	maxspflops:	19309.1000 GFLOPS
result	for	gmem readbw:	2865.4300 GB/s
result	for	gmem_readbw_strided:	601.8820 GB/s
result	for	gmem_writebw:	2214.4100 GB/s
result	for	gmem_writebw_strided:	160.3150 GB/s
result	for	lmem_readbw:	14217.0000 GB/s
result	for	lmem_writebw:	15983.8000 GB/s
result	for	tex_readbw:	1560.7800 GB/sec
result	for	ocl_kernel:	0.0004 sec
result	for	ocl_queue:	0.0050 ms
result	for	bfs:	23.4141 GB/s
result	for	bfs_pcie:	13.4664 GB/s
result	for	bfs_teps:	556110000.0000 Edges/s
result	for	fft_sp:	4006.2800 GFLOPS
result	for	fft_sp_pcie:	74.3269 GFLOPS
result	for	ifft_sp:	3941.3600 GFLOPS
result	for	ifft_sp_pcie:	74.3042 GFLOPS
result	for	sgemm_n:	13458.9000 GFLOPS
result	for	sgemm_t:	12928.0000 GFLOPS
result	for	sgemm_n_pcie:	5807.9600 GFLOPS
result	for	sgemm_t_pcie:	5717.3800 GFLOPS
result	for	md_sp_flops:	1558.6600 GFLOPS
result	for	md_sp_bw:	1194.5100 GB/s
result	for	<pre>md_sp_flops_pcie:</pre>	91.7757 GFLOPS
result	for	md_sp_bw_pcie:	70.3341 GB/s
result	for	md5hash:	42.8308 GHash/s
result	for	reduction:	239.9240 GB/s
result	for	reduction_pcie:	23.7619 GB/s
result	for	scan:	74.4782 GB/s
result	for	scan_pcie:	11.3414 GB/s
result	for	sort:	1.8775 GB/s
result	for	sort_pcie:	1.6457 GB/s
result	for	<pre>spmv_csr_scalar_sp:</pre>	82.0996 Gflop/s
result	for	spmv_csr_scalar_sp_pcie:	2.3338 Gflop/s
result	for	<pre>spmv_csr_scalar_pad_sp:</pre>	69.8097 Gflop/s
result	for	<pre>spmv_csr_scalar_pad_sp_pcie:</pre>	2.4474 Gflop/s
result	for	<pre>spmv_csr_vector_sp:</pre>	137.9340 Gflop/s
result	for	<pre>spmv_csr_vector_sp_pcie:</pre>	2.3610 Gflop/s
result	for	<pre>spmv_csr_vector_pad_sp:</pre>	145.6470 Gflop/s
result	for	<pre>spmv_csr_vector_pad_sp_pcie:</pre>	2.4929 Gflop/s
result	for	<pre>spmv_ellpackr_sp:</pre>	83.6452 Gflop/s
result	for	stencil:	1303.0500 GFLOPS
result	for	triad_bw:	25.2434 GB/s
result	for	s3d:	827.5860 GFLOPS
result	for	s3d_pcie:	407.1100 GFLOPS

APPENDIX D APPLE M1 SHOC SINGLE PRECISION OUTPUT

result	for	bspeed_download:	846.8210 GB/sec
result	for	bspeed_readback:	848.3620 GB/sec
result	for	maxspflops:	297487000.0000 GFLOPS
result	for	gmem_readbw:	2773.7900 GB/s
result	for	gmem_readbw_strided:	2172.6800 GB/s
result	for	gmem_writebw:	2931.5800 GB/s
result	for	gmem_writebw_strided:	1661.5100 GB/s
result	for	lmem_readbw:	25617.3000 GB/s
result	for	lmem_writebw:	35432.5000 GB/s
result	for	tex_readbw:	3091.7900 GB/sec
result	for	ocl_kernel:	0.0000 sec
result	for	ocl_queue:	0.0018 ms
result	for	bfs:	47.7882 GB/s
result	for	bfs_pcie:	46.1705 GB/s
result	for	bfs_teps:	33420200.0000 Edges/s
result	for	fft_sp:	4021.3500 GFLOPS
result	for	fft_sp_pcie:	52.5703 GFLOPS
result	for	ifft_sp:	3983.4400 GFLOPS
result	for	ifft_sp_pcie:	52.5638 GFLOPS
result	for	sgemm_n:	20437.2000 GFLOPS
result	for	sgemm_t:	21263.8000 GFLOPS
result	for	sgemm_n_pcie:	18435.0000 GFLOPS
result	for	sgemm_t_pcie:	19708.6000 GFLOPS
result	for	md_sp_flops:	3706.0700 GFLOPS
result	for	md_sp_bw:	2840.2200 GB/s
result	for	<pre>md_sp_flops_pcie:</pre>	2279.1000 GFLOPS
result	for	md_sp_bw_pcie:	1746.6300 GB/s
result	for	md5hash:	76.5948 GHash/s
result	for	reduction:	2347.8900 GB/s
result	for	reduction_pcie:	807.0470 GB/s
result	for	scan:	19.66/4 GB/s
result	for	scan_pcie:	19.06/3 GB/s
result	IOT	sort:	0.3223 GB/s
result	for	sort_pere:	0.3221 GB/S
result	for	spmv_csr_scalar_sp:	103.2920 GIIOP/S
result	for	spmv_csr_scalar_sp_pcle:	/3.4816 GIIOP/S
result	for	spmv_csr_scalar_pad_sp:	74 4469 Gflop/S
result	for	spmv_csr_scalar_pad_sp_pcle:	74.4468 GIIOP/S
result	for	spmv_csr_vector_sp:	176 4760 GILOP/S
result	for	spmv_csr_vector_sp_pcre:	1/6.4/60 GIIOP/S
result	for	spmv_csr_vector_pad_sp:	195 9000 Cflop/s
rogult	for	spmv_csi_vector_pad_sp_pere:	$388 \ 9170 \ Cflop/s$
result	for	spmv_errpackr_sp;	64 3946 CFIODS
rogult	for	triad hw.	713 6200 CR/c
regult	for	s3d•	1590 5500 GFLOPS
rogult	for	sid noie.	1554 4700 GELOPS
TESUIC	TOT	DONTHOTE.	TODI I ION GLIOLD

APPENDIX E APPLE M1 ULTRA SHOC SINGLE PRECISION OUTPUT

result	for	bspeed_download:	2237.6300 GB/sec
result	for	bspeed_readback:	2024.0300 GB/sec
result	for	maxspflops:	655714000.0000 GFLOPS
result	for	gmem_readbw:	10938.8000 GB/s
result	for	gmem_readbw_strided:	37415.7000 GB/s
result	for	gmem_writebw:	10237.2000 GB/s
result	for	gmem_writebw_strided:	13399.4000 GB/s
result	for	lmem_readbw:	301038.0000 GB/s
result	for	lmem_writebw:	358733.0000 GB/s
result	for	tex_readbw:	28706.6000 GB/sec
result	for	ocl_kernel:	0.0000 sec
result	for	ocl_queue:	0.0027 ms
result	for	bfs:	1009.6800 GB/s
result	for	bfs_pcie:	919.6110 GB/s
result	for	bfs_teps:	73887900.0000 Edges/s
result	for	fft_sp:	37314.0000 GFLOPS
result	for	fft_sp_pcie:	123.9950 GFLOPS
result	for	ifft_sp:	36082.6000 GFLOPS
result	for	ifft_sp_pcie:	123.9810 GFLOPS
result	for	sgemm_n:	248243.0000 GFLOPS
result	for	sgemm_t:	258108.0000 GFLOPS
result	for	sgemm_n_pcie:	188918.0000 GFLOPS
result	for	sgemm_t_pcie:	195467.0000 GFLOPS
result	for	<pre>md_sp_flops:</pre>	18996.3000 GFLOPS
result	for	md_sp_bw:	14558.1000 GB/s
result	for	<pre>md_sp_flops_pcie:</pre>	12922.9000 GFLOPS
result	for	md_sp_bw_pcie:	9903.7200 GB/s
result	for	md5hash:	620.4690 GHash/s
result	for	reduction:	7491.7000 GB/s
result	for	reduction_pcie:	4660.4200 GB/s
result	for	scan:	57.8603 GB/s
result	for	scan_pcie:	57.3959 GB/s
result	for	sort:	0.6722 GB/s
result	for	sort_pcie:	0.6721 GB/s
result	for	<pre>spmv_csr_scalar_sp:</pre>	1054.3300 Gflop/s
result	for	<pre>spmv_csr_scalar_sp_pcie:</pre>	620.2710 Gflop/s
result	for	<pre>spmv_csr_scalar_pad_sp:</pre>	1074.9700 Gflop/s
result	for	<pre>spmv_csr_scalar_pad_sp_pcie:</pre>	747.3490 Gflop/s
result	for	<pre>spmv_csr_vector_sp:</pre>	2657.9100 Gflop/s
result	for	<pre>spmv_csr_vector_sp_pcie:</pre>	961.8070 Gflop/s
result	for	<pre>spmv_csr_vector_pad_sp:</pre>	2916.4600 Gflop/s
result	for	<pre>spmv_csr_vector_pad_sp_pcie:</pre>	1332.3300 Gflop/s
result	for	<pre>spmv_ellpackr_sp:</pre>	694.4130 Gflop/s
result	for	stencil:	432.3320 GFLOPS
result	for	triad_bw:	1614.1800 GB/s
result	for	s3d:	11276.0000 GFLOPS
result	for	s3d_pcie:	10913.0000 GFLOPS