

Comparative Benchmarking Analysis of Next-Generation Space Processors

Evan W. Gretok, Evan T. Kain, Alan D. George
NSF Center for Space, High-performance, and Resilient Computing (SHREC)
Department of Electrical and Computer Engineering
University of Pittsburgh
4420 Bayard Street, Suite 560
Pittsburgh, PA 15213
412-383-8122
{ewg13, etk18, alan.george}@pitt.edu

Abstract—Researchers, corporations, and government entities are seeking to deploy increasingly compute-intensive workloads on space platforms. This need is driving the development of two new radiation-hardened, multi-core space processors, the BAE Systems RAD5545™ processor and the Boeing High-Performance Spaceflight Computing (HPSC) processor. As these systems remained in the development phase, the Freescale P5020DS and P5040DS systems, based on the same PowerPC e5500 architecture as the RAD5545 processor, and the Hardkernel ODROID-C2, sharing the same ARM Cortex-A53 core as the HPSC processor, were selected as facsimiles for evaluation. Several OpenMP-parallelized applications, including a color search, Sobel filter, Mandelbrot set generator, hyperspectral imaging target classifier, and image thumbnailer, were benchmarked on these processing platforms. Performance and energy consumption results on these facsimiles were scaled to forecasted frequencies of the radiation-hardened devices in development. In these studies, the RAD5545 achieved the highest and most consistent parallel efficiency, up to 99%. The HPSC processor achieved faster execution times, averaging about half that of the RAD5545 processor, with lower energy consumption. The evaluated applications reached up to 3.9 speedup across four cores. The frequency-scaling methods were validated by comparing the set of scaled measures with data points from an underclocked facsimile, which yielded an average accuracy of 97% between estimated and measured results. These performance outcomes help to establish the capabilities of both the RAD5545 and HPSC processors for on-board parallel processing of computationally-demanding applications for future space missions.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. BACKGROUND.....	2
3. METHODOLOGY	5
4. RESULTS.....	6
5. CONCLUSIONS.....	10
APPENDICES.....	10
A. EXECUTION TIME	11
B. SPEEDUPS AND PARALLEL EFFICIENCY	12
C. ENERGY CONSUMPTION.....	14
ACKNOWLEDGEMENTS.....	14
REFERENCES.....	14

BIOGRAPHY.....	15
----------------	----

1. INTRODUCTION

Due to the harsh environment of space, the employment of radiation-hardened processors is essential to ensure success of many missions. Two of these processors currently in development, the BAE Systems RAD5545™ processor and the Boeing High Performance Spaceflight Computing (HPSC) processor, are the focus of this research. These new devices drastically improve performance compared to their predecessors and bring multi-core processor architectures to space-computing platforms. This additional performance enables computational loads that were previously deemed infeasible on radiation-hardened space platforms, including advanced sensor-data analysis, computer-vision applications, and autonomous spacecraft operations. These capabilities will equip a new generation of space systems to perform complex analysis on-board, effectively communicate key data, and make autonomous decisions for navigation and critical operations.

Many members of the scientific and aerospace research communities aim to employ sophisticated algorithms at larger scales for big-data processing in space. The increasing complexity and scope of systems and sensors push and often exceed the limits of current space-grade processors. The latest experiments often require larger data sets with long compute times and high resource requirements. If space is to continue to serve as a valuable domain for gathering scientific knowledge, the systems and tools employed must continue to evolve to allow for greater computational capability.

With radiation-hardened space processors crossing the boundary into multi-core architectures, shared-memory multiprocessing becomes a source of parallelism to exploit. The distribution of compute-intensive workloads across multiple processing cores can significantly reduce the impact of a lower clock frequency and achieve speedup over single-core execution. This approach enables the application of more advanced algorithms on larger data sets through on-board processing performance. This research seeks to investigate and compare parallel performance of the

RAD5545 and HPSC processors through application benchmarking. This exploration will provide insight into the advantages and disadvantages of each platform and elucidate the new capabilities emerging for on-board processing.

This paper is organized as follows. Section 2 provides background information on radiation-hardened processors, shared-memory parallelism, and the platforms and applications employed in this study. Section 3 outlines the methods and procedures that were conducted to realize a comparative analysis between the two competing platforms. Section 4 presents the results collected and incorporates observations and discussion. Section 5 presents conclusions and future work.

2. BACKGROUND

This section presents a cursory overview critical to the goals and motivations of this research. The fundamentals of radiation-hardened, space-grade processors are considered. Methods of enabling shared-memory parallelism through OpenMP are noted. Details on the platforms and applications investigated in this study are also shared in this section.

Radiation-Hardened, Space-Grade Processors

The latest space platforms for observation and science host a plethora of unique, sophisticated sensors. Some of these modern sensors can generate terabytes of raw data per day [1]. Transferring such large amounts of data would saturate even the highest bandwidth communication channels. This dilemma is compounded as the missions in need of the most radiation-hardened systems are typically probes or rovers with the farthest distance to travel and thus the lowest bandwidths over which to transmit. Previous research has considered the need for and benefit of on-board processing. Spaceborne high-performance computer systems facilitate applications of high computational complexity, such as sensor-data processing [2] or machine learning [3], which enable more innovative missions. For some distant missions, on-board processing and decision-making will become essential for even basic levels of operation [4].

Unfortunately, the harsh environment of space can be a difficult place for traditional computing devices to function. Impacts from particles like protons and heavy ions cause several types of single-event effects (SEEs). Temporary upsets or functional interrupts affect data or system integrity. More permanent and often destructive latch-ups, burnouts, and gate-ruptures can cause significant damage to the device [5]. Additionally, for long-term missions, the functional degradation of devices due to total ionizing dose (TID) of radiation becomes a serious issue. Typical radiation doses vary from as little as 0.1 krad per year in some low-Earth orbits to as much as 100 Mrad per pass for some Jupiter transfer orbits [6].

Some of the only computing systems capable of withstanding such harsh conditions are radiation-hardened, space-grade processors. The RAD6000TM radiation-hardened space processor was designed to handle a TID of greater than 1.0

Mrad(Si) with fewer than 7.4×10^{-10} upsets per bit per day. Unfortunately, it was only capable of up to 35 DMIPS (million Dhrystone 2.1 instructions per second) at 33 MHz [7], which is paltry compared to over 100,000 DMIPS for modern high-end processors [8]. Despite this lower performance, it achieved success in the Spirit and Opportunity Mars rovers as well as many other landers and probes [7]. The RAD750TM, a predecessor to the RAD5545, can withstand a TID of up to 1.0 Mrad(Si) while delivering consistent computation with fewer than 1.6×10^{-10} upsets per bit per day. However, it is limited to approximately 400 DMIPS at 200 MHz [9] [10]. This processor has been employed in the Lunar Reconnaissance Orbiter, the GPS III modernization effort, and the Curiosity Mars rover [11]. Some missions, though, require several of these processors to meet computational needs, adding to expense and complexity of the designed system [12]. The modern radiation-hardened processors explored in this study can withstand similar conditions while providing significantly higher computational capacity across multiple cores.

Previous research has been conducted to evaluate the performance of radiation-hardened processors. The study in [13] investigated the capabilities of the RAD5545 and several other CPU- and FPGA-based computing systems via performance metrics analysis. These metrics provide insight into performance characteristics such as computational density and memory bandwidth without requiring the device for analysis. That study is expanded in [14] to include kernel benchmarks on the same platforms. The research presented here focuses on application benchmarks to provide a more representative real-world assessment of the capability of these platforms.

Shared-Memory Parallelism with OpenMP

Parallel computing, once a niche discipline, is now ever expanding into a world of multi-core processors, massively parallel graphics processing units, and a myriad of hardware accelerators. This parallelization has allowed engineers to overcome the barriers that slowed performance gains in the processors of the past. Some complex algorithms and applications are now only realizable in given time constraints with sufficient parallelization [15]. The next generation of space processors has been equipped with immense capacity for multi-core data processing. Due to communication overhead and architectural limitations, performance does not necessarily scale with the increased number of cores. This variability in parallel performance presents the need for deeper study and analysis of different applications employed on these architectures, a need that this research is intended to help address.

There are many practical methods for parallelizing software across multiple processing units. The most commonly applied are the message-passing and shared-memory models [16]. Shared-memory models are used when compute nodes possess a common memory space, allowing operations to be conducted without the need for data transmission to and from separate nodes. The most widely used variant of this model

Table 1: Platform Specifications

Platform	PC	P5020DS	P5040DS	ODROID-C2
Processor	Intel Core i7-6700k	QorIQ P5020	QorIQ P5040	Amlogic S905
Architecture	x86-64	PowerPC e5500	PowerPC e5500	ARM Cortex-A53
Frequency (MHz)	3408.00	2000.00	2266.67	1540.00
Cores	4	2	4	4
Power Consumption (W)	65	28	49	5
Memory (GB)	16	4	4	2
Memory Type	DDR4	DDR3	DDR3	DDR3
Internal Memory Bandwidth (GB/s)	-	119.27	270.35	416
External Memory Bandwidth (GB/s)	34.1	21.3	25.6	1.5

is Open Multi-Processing (OpenMP), which allows for parallelism via compiler directives and multithreading [17]. The techniques involved in this study’s approach will be based primarily in OpenMP due to the multi-core, single-node architecture of the examined space processors.

Scheduling is another factor in parallelization that affects how a problem is divided and how well a parallel program performs. OpenMP’s default scheduling methodology is the static division of computation evenly across all cores at compile time. Dynamic scheduling refers to the process of OpenMP assigning small segments of the job to each core during run-time as pieces are completed. The dynamic approach allows processing to be split more evenly across time at the cost of some run-time scheduling overhead.

Platforms

To assess performance of the BAE Systems RAD5545 and Boeing HPSC processors during their development phases, platforms of similar architecture were selected as facsimiles upon which to perform comparative application benchmarking. For the RAD5545 processor, the PowerPC e5500-based Freescale P5020DS and P5040DS systems were selected. For the HPSC processor, the ARM Cortex-A53-based Hardkernel ODROID-C2 was employed. Applications were also run on a standard x86-64-based Intel Core i7 desktop workstation for a baseline performance comparison. Specifications of these platforms can be referenced in Table 1.

The RAD5545 is a radiation-hardened-by-design, space-grade processor. The device is designed for extreme reliability, with fewer than 2×10^{-9} upsets per bit per day, a TID rating of 1 Mrad(Si), and immunity to latch-up. This system is also specifically designed for on-board processing applications, equipped with four RAD5500™ Power Architecture processor cores to conduct computations in an efficiently parallel manner. This processor is capable of 5.6 GOPS (billions of operations per second), 3.7 GFLOPS (billions of floating-point operations per second), and up to 1398 DMIPS per core at 466 MHz [13], for a total of 5592 DMIPS. Its capability is aided by three levels of cache as well as the ability to interface with other devices via Serial

RapidIO for high-speed communication [18]. As this device was not yet available at the time of this study, its performance was approximated using commercially available processors.

The P5020 and P5040 systems served as useful facsimiles for the RAD5545 processor because, combined, they employ all the components of interest present in the RAD5545 processor. Although nearly identical, the P5020 and P5040 systems differ in both the number of processing cores available and the nature of their interconnects. The P5020 system only has two e5500 processor cores but possesses Serial RapidIO interconnects [19]. The P5040 system lacks Serial RapidIO but contains four e5500 processor cores [20].

Boeing’s HPSC processor is a similar radiation-hardened-by-design, space-grade processor currently in development. The HPSC processor was originally conceived to meet the mutual needs of the National Aeronautics and Space Administration (NASA) and the United States Air Force Research Laboratory (AFRL) for next-generation space processing capabilities. The device’s requirements aim for eight cores per “chiplet.” Overall, the processor is capable of a performance of up to 15 GOPS and attaining 1840 DMIPS per core at 800 MHz, 7360 DMIPS per ARM Cortex-A53 cluster, or 14,720 DMIPS per chiplet. Use of ARM’s single-instruction, multiple-data (SIMD) NEON accelerators can yield up to 100 GOPS. The system is intended to perform at 1×10^{-10} upsets per bit per day or fewer, exhibit a TID of 1 Mrad(Si), and incorporate latch-up immunity. HPSC chiplets are designed to be scalable via interconnection through several high-speed interfaces, including Ethernet, PCIe, and Serial RapidIO [21]. Integrated fault-tolerance will enable error detection and correction, checkpoint and rollback functionality, and n-modular redundancy [22]. The HPSC processor’s performance must also be approximated by commercial devices, as it is in even earlier development phases than the RAD5545 processor.

Many current devices employ a system-on-chip (SoC) containing the ARM Cortex-A53 processor architecture upon which the HPSC processor is based. The most accessible Cortex-A53 derivative to this research group is the Hardkernel ODROID-C2 platform, which boasts a quad-core

ARM Cortex-A53 processor equivalent to half of a chiplet in the HPSC processor. Space-processing solutions based on the HPSC processor are projected to scale to multiple chiplets, each including two ARM Cortex-A53 quad-core processor clusters coupled with an advanced microcontroller bus architecture (AMBA) interconnect for symmetric multi-processing operation. The ODROID-C2 only serves as a facsimile for a portion of a single chiplet. This approach is considered valid within the scope of this research due to the fair comparison it permits with the quad-core RAD5545 processor. Further, existing studies, such as [23], note that substantial overhead is incurred by parallelization over the AMBA interconnect. Simple applications parallelized across the HPSC processor's two Cortex-A53 clusters can experience significant reductions in speedup. It may be more effective to confine some applications to a single quad-core region to maximize parallel efficiency. For example, attaining speedups of up to 3.9 for two applications on each quad-core processor at once may be a significantly better use of resources than speedups that remain in the range of four to five across all eight cores.

Notable differences to highlight between the P5020 and P5040 systems and ODROID-C2 include the employment of an L3 cache in the P5020 and P5040 systems and differences in external and internal memory bandwidth. The ODROID-C2 excels with respect to internal memory bandwidth, with 416 GB/s compared to 119 and 270 GB/s for the P5020 and P5040 systems, respectively. The P5020 and P5040 systems are superior in external memory bandwidth, with 21.3 and 25.6 GB/s, respectively, compared to 1.5 GB/s for the ODROID-C2. The significant difference in thermal design power (TDP) between the P5020 and P5040 systems in comparison to the ODROID-C2 should also be noted, with the latter consuming significantly less than the P5040 system's energy needs.

Effectively comparing commercial-off-the-shelf (COTS) platforms and radiation-hardened derivatives is nontrivial. Despite the common architectures shared by the P5020, P5040, and RAD5545 processors or the ODROID-C2 and HPSC processors, the process of radiation hardening yields a device with substantial differences in performance and power characteristics. These discrepancies make final performance of the device difficult to predict. Due to architectural similarities, the performance data garnered from the facsimiles in this study is considered the best available basis for forecasting the performance of these radiation-hardened devices.

Applications

This research evaluated five applications in comparative benchmarking, including color search, hyperspectral imaging (HSI) linearly-constrained minimum variance (LCMV) beamforming, Mandelbrot set generation, Sobel filter, and image thumbnailer applications. Many of these applications were selected due to their relevance for numerous space mission scenarios. The test image used for most applications

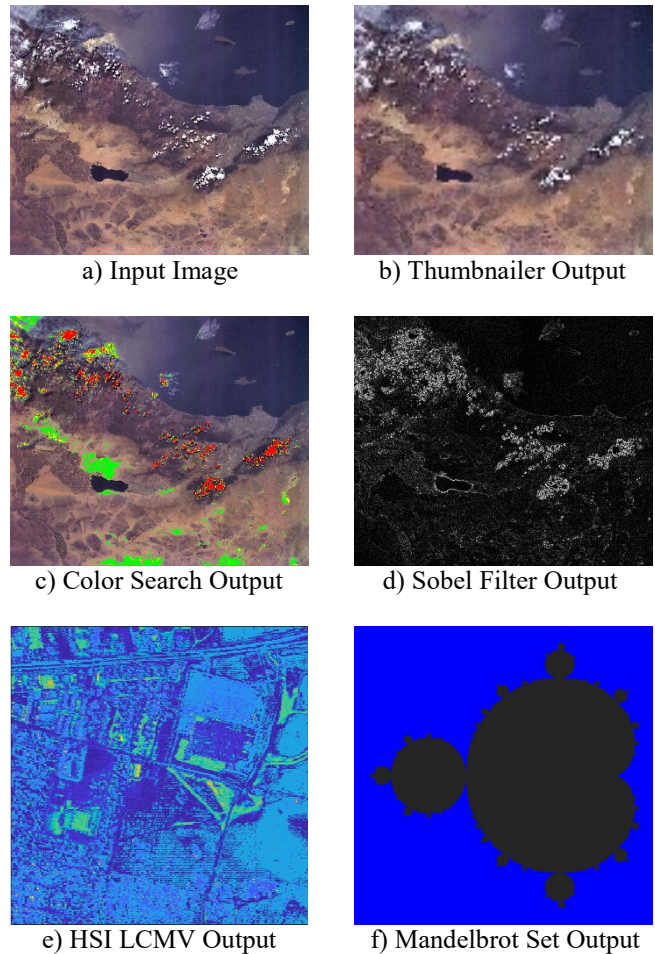


Figure 1: Application Output Images

as well as output images from each of the applications are visible in Figure 1.

The image thumbnailer application performs bilinear interpolation to resample an input image to an output of lower resolution, creating a thumbnail. These thumbnails are useful in space use cases for creating low-resolution versions of images for verification before downloading the full-resolution version, which takes much longer. A demonstration thumbnail for the previously presented input image can be referenced in Figure 1(b). The task of image thumbnailing could be parallelized simply by the horizontal lines of the image. While load balancing for the image thumbnailer was even, better performance was observed with the use of dynamic scheduling, and thus this modification was included in the employed thumbnailer application.

The color search application employed is a simple image-processing program that performs an exhaustive search of an image for a specified color value. The Euclidean distance between the color of each pixel in the image and a desired search pixel is calculated by the method described in [24]. If any pixel's distance is within a preset threshold, that pixel is highlighted in the output image to indicate a match. An example of the color search, a search for clouds in Earth-observing imagery, is depicted in Figure 1(c). It should be

noted that five, ten, and fifteen percent thresholds are denoted in this test as red, yellow, and green highlighting, respectively. The color search was parallelized via OpenMP with the image being evenly divided statically across the cores by horizontal lines.

The Sobel filter application performs edge detection on an image, which is computed in this case by performing a pair of two-dimensional convolutions with a window size of 3×3 . Calculations are performed on the intensity of each of the pixels within the window to determine a gradient for change in intensity in the x and y directions for each channel. The magnitude of these gradients highlights areas corresponding to edges, as observed in Figure 1(d). Parallelization of the Sobel filter also divides processing statically by horizontal lines in the image.

Hyperspectral imaging is the process of capturing images concurrently from many different spectral bands. The spectral profiles of the image are then used to identify objects and/or classify which materials are present at certain locations. This process can be used to build terrain maps or measure the advance of urbanization, deforestation, or glacial melt, among other object-sensing applications. The LCMV beamforming algorithm is a supervised-classification method that only requires spectral information for the targets to be detected. This application was developed as a benchmark in [25] and later parallelized. Most of the execution time consists of matrix multiplications, which are easily parallelized to provide a significant performance increase. A processed output image from the data set used in this study, colorized via the MATLAB `imagesc` function, is pictured in Figure 1(e).

The Mandelbrot set fractal generator application was included in this study due to its embarrassingly parallel nature and its use of intensive double-precision floating-point computations. Construction of a Mandelbrot set consists of checking points in a complex plane under the condition $z_{n+1} = z_n^2 + c$. If a point c yields a bounded sequence, that point is a part of the set. As the inclusion of one point is separate and does not depend on the inclusion of others, the problem can be easily parallelized and is considered embarrassingly parallel [26]. A fractal generated by this software can be referenced in Figure 1(f). The Mandelbrot set application was developed from examples accessible at [27] with the OpenMP parallelization verified by [26]. With its processing also divided by horizontal lines of the image, the Mandelbrot set demonstrated uneven load distribution. Greater computational density near the center “bulb” of the fractal was accounted for using dynamic scheduling.

3. METHODOLOGY

This methodology section conveys the steps performed in the realization of the goals of this research. The preparation of platforms, the methods of measurement and calculation, and the approach to transforming these results into an accurate prediction of performance for the RAD5545 and HPSC processors are conveyed.

Platform Preparation

Each of the platforms employed was prepared for application benchmarking by installing a lightweight operating system (OS) and the relevant libraries for program execution. PC benchmarks were conducted on an Ubuntu 16.04.4 LTS desktop installation. ODROID-C2 benchmarks were conducted within an Ubuntu MATE 16.04.4 installation. The P5020 and P5040 systems were both equipped with custom lightweight Linux images prepared via the Linux SDK for QorIQ processors. The GNU Scientific Library (GSL) packages required for execution of the hyperspectral imaging application were installed on each of these platforms.

Application Preparation and Input

Applications were garnered from their respective sources and parallelized for shared-memory multiprocessing using OpenMP. A goal of this study was to ensure optimal performance with consistent program code across platforms. Both the serial baseline and parallel variants of each application were optimized to -O2 levels during the compilation process. Program optimizations for the NEON SIMD accelerators of the ARM Cortex-A53 architecture are not used in this research.

For input to the color search and Sobel filter applications, a terrestrial image thumbnail acquired from the NSF SHREC Space Test Program - Houston 5 - CHREC Space Processor (STP-H5/CSP) experiment aboard the International Space Station (ISS) was scaled up to the standard pixel dimensions of a full-size image, 2448×2050 pixels [28]. For input to the thumbnailer, a 4256×2832 -pixel image of the Earth taken by an astronaut aboard the ISS was scaled down to typical “full HD” resolution (1920×1080). A different, larger thumbnail was created to ensure execution times of the thumbnailer reached within the same order of magnitude as most other applications. Finally, for the HSI LCMV beamforming application, the URBAN data set of HYDICE sensor imagery provided by the United States Army Corps of Engineering Geospatial Research Laboratory served as input [29].

Performance Measures

All applications recorded and output OpenMP wall-clock timing for the execution of the primary operation of the program, such as the convolutions of the Sobel filter or the fractal generation of the Mandelbrot set. Execution times for serial baseline and parallel variants of most applications were averaged over 1000 runs. The HSI LCMV beamforming application was run for only 100 runs due to its roughly two orders-of-magnitude longer execution time. Each parallel run collected execution times for one, two, three, and four cores for every platform except the P5020 system, which was limited to two cores.

Energy Consumption Measures

System power measurements were collected for each application and platform combination using a power meter. Measurements were taken at idle, serial load, and parallel load for one through four cores. Idle was defined as the

platform being fully booted and ready for application execution but with no foreground applications running. Serial load was determined as the peak power consumption while running the serial-baseline scripts. Parallel load was considered the peak power consumption for a certain number of cores while running the parallel energy-evaluation scripts. Calculations were performed for energy consumption of each combination of application, platform, and number of cores by multiplying the power consumption by the execution time.

Underclocking and Frequency Scaling

The RAD5545 and HPSC processors, being radiation-hardened, have a significantly lower clock frequency than the facsimile platforms assessed. Previous performance studies for radiation-hardened processors have employed frequency scaling as in [13]. This research proposes a hybrid approach to isolate and minimize scaling error by unifying two methods: underclocking and frequency scaling. Underclocking implies collecting actual results at a reduced device clock frequency. Frequency scaling implies projecting acquired results to a lower device frequency. Applying underclocking where feasible and frequency scaling where necessary generates an effective representation of the performance of these radiation-hardened space processors.

The P5020 and P5040 RAD5545 facsimiles employed could not be effectively underclocked without hardware reconfiguration and the generation of new boot images. It was therefore decided that frequency-scaling methods to the RAD5545 projected frequency of 466 MHz consistent with [13] would be applied. The ODROID-C2 HPSC facsimile was much more straightforward to underclock via a software frequency governor integrated as a component of the processor driver. For thorough HPSC processor analysis, two frequency values were targeted: 466 MHz and 800 MHz. The 800 MHz target is noted in [21] as the planned maximum frequency for the HPSC processor. The 466 MHz target is meant to equate to the RAD5545 processor frequency and allow more direct architecture comparison without variance in clock speed. Minimum and maximum frequencies were desired for the HPSC processor as further design and fabrication may reduce the target frequency. However, the ODROID-C2 was bound by hardware limitations to 500 MHz and 1000 MHz. In order to remedy this discrepancy from the target frequencies, a hybrid-scaling approach was adopted. The ODROID-C2 was underclocked to 500 MHz and 1000 MHz and results were gathered. These results were then scaled to the target frequencies of 466 MHz and 800 MHz, respectively using Eq. 1.

$$t_{scaled} = t_{exec} \times \frac{f_{target}}{f_{evaluated}} \quad (1)$$

While frequency scaling is a common and accepted practice for benchmarking predicted performance, parallel performance measures vary between frequencies, especially when non-static scheduling methods are employed, in manners not precisely predictable with scaling alone. The hybrid approach employed in this study allows a better

representation of parallel performance and scaling behavior by relying less on the frequency-scaling model and more on real data. In order to ensure that scaling was accurate in the cases where it was necessary, results acquired at 500 MHz and 1000 MHz frequencies were compared with the full-speed results scaled to those frequencies, directly comparing real and scaled versions of results.

4. RESULTS

This section displays all performance and energy consumption results and offers some discussion on the trends these results represent. The results include the execution times, speedups, parallel efficiencies, and energy consumptions for each combination of application, platform, and number of cores. A comparison of the validity of scaled versus underclocked benchmark results cements the legitimacy of the usage of frequency scaling where necessary.

Performance and Energy Consumption Results

The charts in Figure 2 and segments of discussion that follow reflect the execution time of each application on each platform. For the RAD5545 and HPSC processors, projected minimum and maximum execution times are shown. The front bars of typical coloration denote predicted minimum execution time while the back, darker bars denote predicted maximum execution time. For the RAD5545 processor, these were determined by taking the minimum and maximum-scaled values from the P5020 and P5040 results. For the HPSC processor, these denote results at 800 MHz as the fastest and 466 MHz as the slowest. The color search in Figure 2(a) and the Sobel filter in Figure 2(d) are the quickest of the five applications, an important consideration for later inspections of speedup and efficiency. The HSI LCMV beamforming application in Figure 2(b) is the slowest. Tabulated execution times for each application and platform averaged across runs can be referenced in Table 2 in Appendix A.

The P5040 system depicts execution times that are, on average, 10% faster than the P5020 system across all applications. This improvement is primarily due to its faster clock speed, 2266 MHz for the P5040 system compared to 2000 MHz for the P5020 system, a 13% increase. The P5020 system may reconcile some of this difference through more advanced scheduling and memory management schemes, as noted in [19]. These same advanced features are likely to be employed in the RAD5545 processor, improving its overall performance.

The ODROID-C2 depicts significantly faster execution times than the P5020 and P5040 systems. On average, it performs roughly 37% faster than the P5040 system. In some isolated cases, particularly involving memory-bound applications like the thumbnailer in Figure 2(e), the ODROID-C2 executed more than twice as fast. These faster times are notable considering its 1540 MHz clock speed compared to the P5040 system's 2266 MHz, 47% faster. This difference in performance is attributed to architectural advantages as well

as a higher internal memory bandwidth: 416 GB/s on the ODROID-C2 compared to 119 and 270 GB/s on the P5020 and P5040 systems, respectively. For the Mandelbrot set in Figure 2(c), a compute-bound application, the P5040 system outpaces the ODROID-C2 by up to 15%.

The RAD5545 and HPSC processors are projected to be capable of performance within the same order of magnitude for most applications. While the HPSC processor is consistently faster, much of this speed relies on the attainment of 800 MHz performance in a radiation-hardened package. The lower 466 MHz assessment, as visualized by the darker back bars, is still faster but less competitive for a few applications tested. Even so, the observed advantages of the ARM Cortex-A53 architecture will apply regardless of frequency.

The charts in Figure 3 and segments of discussion that follow reflect the speedups and parallel efficiencies of each application on each platform. For the RAD5545 and HPSC processors, projected minimum and maximum speedups and parallel efficiencies are shown. A common legend for these speedup and parallel efficiency charts is included. The circle and square markers denote the maximum speedups and parallel efficiencies, respectively. The triangle and diamond markers denote the minimum speedups and parallel efficiencies, respectively. For the RAD5545 processor, minima and maxima were determined by taking the minimum- and maximum-scaled values from the P5020 and P5040 results. For the HPSC processor, minima and maxima denote the minimum and maximum speedups and parallel efficiencies from results at 800 MHz and 466 MHz for each application. Tabulated speedups and parallel efficiencies for each application and platform averaged across runs can be referenced in Table 3 and Table 4 in Appendix B.

For applications that execute more quickly, the color search in Figure 3(a) and Sobel filter in Figure 3(d), a higher overhead is experienced. Preparing shared and private data for parallelization and forking or joining threads takes a larger portion of the overall execution-time of the program. This overhead results in lower speedups and parallel efficiencies for those applications. The Mandelbrot set in Figure 3(c) presents some of the most ideal trends, reaching speedups of up to 3.9 and an average efficiency of 98% across all platforms. To improve the performance of the Mandelbrot set and thumbnailer in Figure 3(e), dynamic scheduling is employed. These are two visible cases of trends where dynamic scheduling is more effective than statically dividing the workload at compile time.

Regardless of the scheduling methodology, the trends indicate significantly better speedups and parallel efficiencies from the P5020 and P5040 systems. These platforms average 5% improvement in speedup and parallel efficiency compared to the ODROID-C2. This average can be misleading, however, as consistency for larger numbers of cores is most critical. In some cases, such as the quad-core Sobel filter in Figure 3(d), the P5020 and P5040 systems boast more than 50% better speedup and parallel efficiency

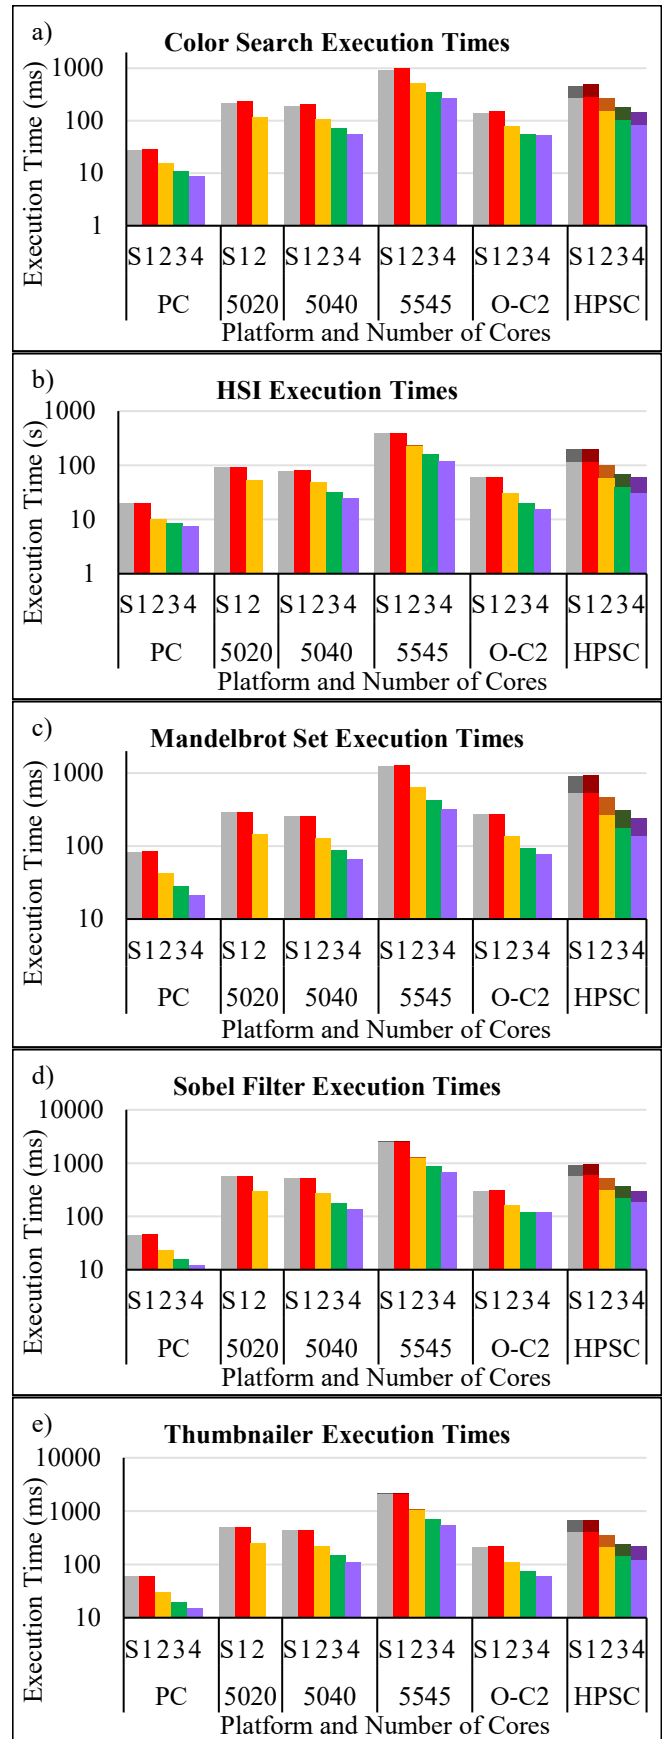


Figure 2: Parallel Application Execution Times

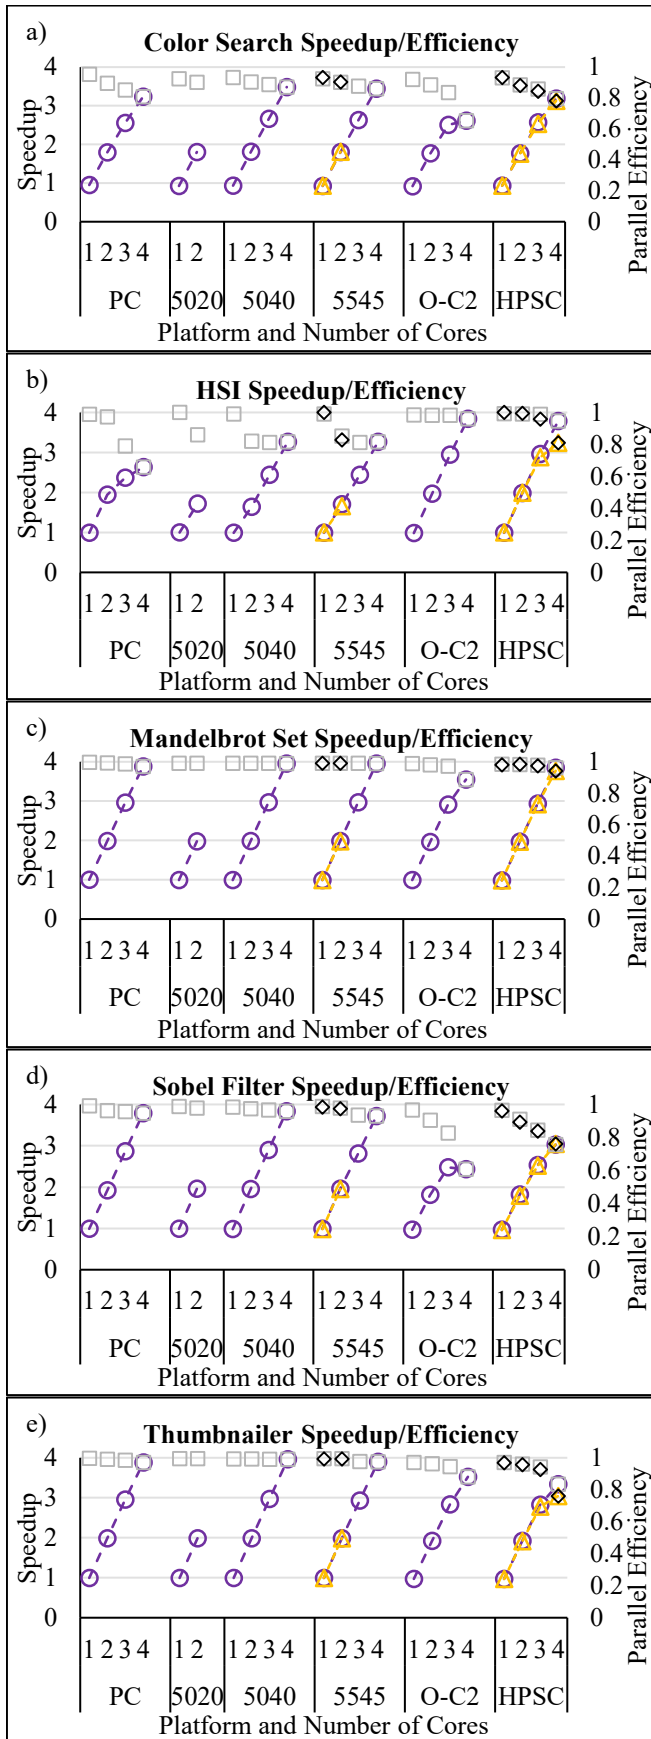


Figure 3: Parallel Application Speedups and Efficiencies

than the ODROID-C2. Despite falloff in efficiencies for the HSI application visible in Figure 3(b), likely due to the overhead of preparing large data structures for parallelism, efficiencies remain consistently high thereafter. The authors reason that the Data Path Acceleration Architecture (DPAA) and associated hardware accelerators present in the P5020 and P5040 systems for buffer, queue, and frame management allow significantly better performance in these cases.

In comparison, the ODROID-C2 exhibits large drops in speedup and parallel efficiency, occasionally even before using all four cores. Particularly large falloffs are experienced for four-core parallelization. However, OS overhead likely contributes to this lackluster quad-core performance, as one or more cores running the application must bear the overhead of running OS tasks. The P5020 and P5040 systems boast much lighter operating systems by comparison. Effectively parallelized applications, such as the HSI or Mandelbrot set, fare better on the ODROID-C2 and thus show more promise for further scaling across devices.

The charts in Figure 4 and segments of discussion that follow reflect the energy consumption of each application on each platform. While platform hardware contributes more to the energy consumption than the application, patterns of energy consumption relate highly to the patterns observed previously in execution time. The applications with the shortest execution times, the color search in Figure 4(a) and Sobel

filter in Figure 4(d), consumed less energy on all platforms. Due primarily to its significantly longer execution time, the HSI application in Figure 4(b) consumes the most energy, measured in kilojoules. Tabulated energy consumptions for each application and platform averaged across runs can be referenced in Table 5 in Appendix C.

The ODROID-C2, as a single-board computer, significantly bests the other platforms in energy consumption. In most tests, the P5020 and P5040 systems consumed more energy than the desktop workstation. System energy comparison with the ODROID-C2 is biased as the P5020 and P5040 systems include additional interfaces and peripherals, such as optical and hard disk drives as well as higher-rated power supplies, that are not present on the ODROID-C2 and will not be present on the RAD5545 or HPSC processors. It should be noted that the 17.7-Watt power consumption documented for the RAD5545 processor in [18] would significantly reduce its energy consumption in comparison to the P5020 and P5040 systems. However, the HPSC processor aims for significantly lower power consumption, below seven Watts per chiplet [21], which makes it favorable for many low-energy applications. In applications with larger energy budgets, room remains for additional devices, increasing the potential for scalability to higher computational capabilities. Due to the early development stages of the HPSC processor, no direct

system energy prediction or comparison between the RAD5545 and HPSC processors was conducted in this study.

Another key insight is the effect of parallelization on energy consumption. For applications that parallelize poorly, such as the color search and the Sobel filter due to their speed, the increased dynamic power of another core partaking in the workload results in higher energy consumption. For applications that parallelize well, especially visible in HSI, the reduction in processing time negates the additional dynamic-power overhead of another core, and energy consumption is significantly reduced. While HSI biases this assessment due to its long execution-time, the Mandelbrot set in Figure 4(c) and thumbnailer in Figure 4(e) also produce consistently reduced energy consumption. This energy-consumption trend is especially revealing considering many of the most critical applications for space-grade processors involve complex calculations and long execution-times. The knowledge that effective parallelization can further reduce energy consumption is significant motivation for the adoption of these multi-core platforms.

Frequency Scaling Versus Underclocking

To ensure that frequency scaling applied for these application and platform combinations with limited error, a validation method was devised. Full-speed results on the ODROID-C2 were scaled to 500 MHz and 1000 MHz, and then compared to results collected at those frequencies. Ratios were derived between the scaled and actual results and then averaged. This average was compared with the expected ratio of the device frequency to the target frequency, allowing a determination of the average scaling error for this set of applications on the ODROID-C2.

As expected, scaling error was minimal, averaging less than 2.70%. However, this scaling error does not merely result from subtle variations in underclocking the device. Much more impact is derived from how well the considered applications parallelize and what scheduling methodologies are used. This conclusion is supported by the tendency of faster applications, such as the color search and Sobel filter, to stray from the expected ratio due to relatively high parallel overhead relative to their total execution time. Again, this validation of scaling accuracy is also limited to one platform, the ODROID-C2, as underclocking the P5020 and P5040 systems would have required modifications to their underlying operating system images that were infeasible in the scope of this study. Thus, the full 2.70% scaling error applies to the projected RAD5545 processor measures. However, scaling was only applied to the HPSC processor measures for the transitions from 500 MHz to 466 MHz, a 6.8% change, and 1000 MHz to 800 MHz, a 20% change. Scaling error applied in this manner results in 0.18% and 0.54% scaling errors for 466 MHz and 800 MHz measures, respectively, averaging a 0.36% scaling error for predicted HPSC processor times.

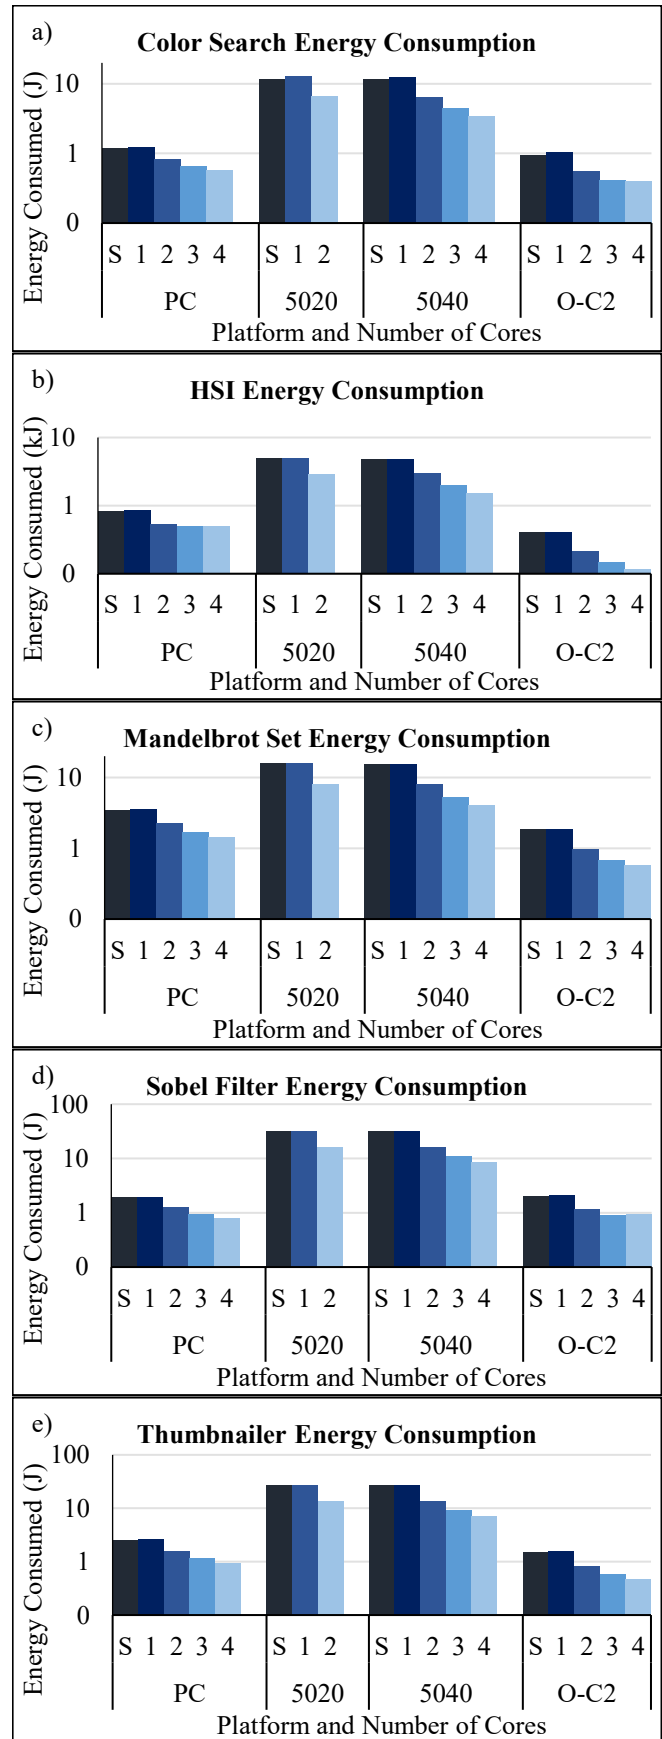


Figure 4: Parallel Application Energy Consumption

5. CONCLUSIONS

The primary focus of this study was the assessment of parallel application performance to predict and compare the capabilities of two next-generation space processors, BAE's RAD5545 processor and Boeing's HPSC processor. The primary platforms of consideration were the Freescale QorIQ P5020DS and P5040DS systems, which feature PowerPC e5500 architecture dual- and quad-core processors, respectively, and the Hardkernel ODROID-C2, which features a quad-core ARM Cortex-A53 processor. These platforms serve as COTS facsimiles for the space-grade RAD5545 and HPSC processors, respectively, currently in development. Several applications of relevance to space missions were benchmarked on these platforms to determine the expected performance as well as strengths and weaknesses of the facsimiles' space-grade counterparts. Facsimile power measures allowed comparison on the dimensions of system energy consumption.

Summary of Results

The high parallel efficiencies boasted by the PowerPC-e5500-based facsimiles indicate substantial scalability for parallel applications. Considering the RAD5545 processor's capacity for high-speed interconnect via Serial RapidIO, this high efficiency maximizes the effectiveness of parallelization over a network of interconnected processors. By comparison, the HPSC facsimile achieves significantly better performance at lower clock speeds and much lower energy consumption. This level of performance aids in ensuring the HPSC processor will remain competitive even after the decrease in clock speed and increase in power consumption inherent from the radiation-hardening process. Despite the HPSC processor's performance, efficiencies are projected to diminish for some applications even before parallelizing over all four cores. Performance may decrease further for parallelization over the AMBA interconnect for all eight cores of the chiplet as well as over multiple chiplets [23], even with the high-speed interfaces employed.

These results depict an effective forecast for the performance of the BAE Systems RAD5545 and Boeing HPSC next-generation space processors. Comparison of scaled versus underclocked performance results for the applications tested indicate validity of frequency scaling within 2.70% error. This scaling error applied in full to the expected RAD5545 results as frequency scaling alone was used for the P5020 and P5040 results. The error is further minimized to 0.36% error for the HPSC results by using a hybrid approach, underclocking the ODROID-C2 to the nearest supported frequency and scaling the remainder of the way to the target frequency of the final device.

Future Work

This research is easily extended to many additional applications and platforms for further analyses. With regard to the HPSC processor, the tests in this study only relate to the performance of one quad-core cluster of a single chiplet. Further exploration may investigate the performance of these

applications extended with parallelism across both quad-core clusters of a chiplet or with support of the SIMD NEON accelerators. For both the RAD5545 and HPSC processors, studies in scaling these applications across interconnects, such as Serial RapidIO, between chiplets or processors will yield intriguing results with respect to their large-scale employment on future space-computing platforms.

APPENDICES

The following appendices are included to better convey numerical data represented in the presented charts. The appendices are divided into execution time, speedup and parallel efficiency, and energy consumption categories.

A. EXECUTION TIME

Table 2: Execution Time for All Applications on All Platforms

Application	Platform	Serial	Number of Cores			
			1	2	3	4
Color Search (ms)	PC	27.86	29.22	15.54	10.91	8.61
	5020	214.31	231.70	118.67	-	-
	5040	191.34	205.23	105.82	71.93	54.90
	5545 Min	919.78	994.42	509.31	349.89	267.06
	5545 Max	930.71	998.27	514.71	-	-
	O-C2	138.37	150.24	78.22	55.12	52.99
	HPSC Min	266.77	286.27	150.68	103.67	83.80
	HPSC Max	463.27	496.72	263.01	182.86	148.14
HSI (s)	PC	19.81	20.04	10.19	8.37	7.54
	5020	90.67	90.70	52.78	-	-
	5040	79.25	80.06	48.35	32.54	24.32
	5545 Min	385.47	389.28	226.52	158.26	118.28
	5545 Max	389.16	389.42	235.19	-	-
	O-C2	59.58	60.59	30.34	20.25	15.52
	HPSC Min	114.76	115.65	58.01	38.75	30.35
	HPSC Max	196.52	196.61	98.86	68.35	60.73
Mandelbrot Set (ms)	PC	82.84	83.13	41.69	27.99	21.34
	5020	289.85	292.52	146.38	-	-
	5040	255.54	257.74	128.98	86.04	64.62
	5545 Min	1242.96	1253.67	627.36	418.52	314.31
	5545 Max	1244.01	1255.43	628.24	-	-
	O-C2	269.43	272.61	137.31	92.43	75.86
	HPSC Min	518.49	527.07	263.83	176.51	134.65
	HPSC Max	899.01	916.24	456.82	307.71	238.43
Sobel Filter (ms)	PC	45.35	45.79	23.59	15.82	11.98
	5020	571.12	578.69	291.99	-	-
	5040	519.70	528.88	266.99	179.55	135.72
	5545 Min	2451.17	2483.64	1253.17	873.36	660.17
	5545 Max	2527.87	2572.53	1298.64	-	-
	O-C2	296.97	307.33	164.16	119.90	122.24
	HPSC Min	560.75	581.87	307.98	221.50	185.50
	HPSC Max	915.97	954.99	513.06	364.21	301.50
Thumbnailer (ms)	PC	58.89	59.06	29.69	19.93	15.16
	5020	488.36	491.05	245.67	-	-
	5040	437.54	441.06	220.54	147.13	110.48
	5545 Min	2095.96	2107.53	1054.37	715.63	537.39
	5545 Max	2128.22	2145.37	1072.75	-	-
	O-C2	210.48	216.76	109.41	74.19	59.78
	HPSC Min	397.04	408.63	206.42	140.32	118.83
	HPSC Max	663.18	684.52	346.61	238.53	217.68

B. SPEEDUPS AND PARALLEL EFFICIENCY

Table 3: Speedup for All Applications on All Platforms

Application	Platform	Number of Cores			
		1	2	3	4
Color Search	PC	0.9537	1.7935	2.5551	3.2356
	5020	0.9249	1.8059	-	-
	5040	0.9323	1.8082	2.6600	3.4851
	5545 Min	0.9249	1.8059	2.6288	3.4442
	5545 Max	0.9323	1.8082	-	-
	O-C2	0.9210	1.7691	2.5103	2.6111
	HPSC Min	0.9319	1.7614	2.5335	3.1273
	HPSC Max	0.9327	1.7704	2.5733	3.1832
HSI	PC	0.9884	1.9448	2.3667	2.6279
	5020	0.9997	1.7180	-	-
	5040	0.9899	1.6390	2.4357	3.2590
	5545 Min	0.9902	1.6547	2.4357	3.2590
	5545 Max	0.9993	1.7017	-	-
	O-C2	0.9834	1.9637	2.9425	3.8393
	HPSC Min	0.9923	1.9782	2.8752	3.2358
	HPSC Max	0.9995	1.9879	2.9618	3.7818
Mandelbrot Set	PC	0.9965	1.9873	2.9599	3.8812
	5020	0.9909	1.9801	-	-
	5040	0.9915	1.9813	2.9699	3.9546
	5545 Min	0.9909	1.9801	-	-
	5545 Max	0.9915	1.9813	2.9699	3.9546
	O-C2	0.9884	1.9623	2.9151	3.5517
	HPSC Min	0.9812	1.9652	2.9216	3.7706
	HPSC Max	0.9837	1.9680	2.9374	3.8506
Sobel Filter	PC	0.9904	1.9224	2.8662	3.7861
	5020	0.9869	1.9560	-	-
	5040	0.9826	1.9466	2.8944	3.8291
	5545 Min	0.9826	1.9466	-	-
	5545 Max	0.9869	1.9560	2.8066	3.7130
	O-C2	0.9663	1.8091	2.4768	2.4293
	HPSC Min	0.9591	1.7853	2.5149	3.0229
	HPSC Max	0.9637	1.8207	2.5316	3.0381
Thumbnailer	PC	0.9971	1.9836	2.9550	3.8835
	5020	0.9945	1.9879	-	-
	5040	0.9920	1.9839	2.9739	3.9603
	5545 Min	0.9920	1.9839	-	-
	5545 Max	0.9945	1.9879	2.9288	3.9003
	O-C2	0.9710	1.9239	2.8372	3.5211
	HPSC Min	0.9688	1.9133	2.7803	3.0466
	HPSC Max	0.9717	1.9235	2.8295	3.3413

Table 4: Parallel Efficiency for All Applications on All Platforms

Application	Platform	Number of Cores			
		1	2	3	4
Color Search (%)	PC	95.37%	89.67%	85.17%	80.89%
	5020	92.49%	90.30%	-	-
	5040	93.23%	90.41%	88.67%	87.13%
	5545 Min	92.49%	90.30%	87.63%	86.10%
	5545 Max	93.23%	90.41%	-	-
	O-C2	92.10%	88.46%	83.68%	65.28%
	HPSC Min	93.19%	88.07%	84.45%	78.18%
HPSC Max	93.27%	88.52%	85.78%	79.58%	
HSI (%)	PC	98.84%	97.24%	78.89%	65.70%
	5020	99.97%	85.90%	-	-
	5040	98.99%	81.95%	81.19%	81.47%
	5545 Min	99.02%	82.73%	81.19%	81.47%
	5545 Max	99.93%	85.08%	-	-
	O-C2	98.34%	98.18%	98.08%	95.98%
	HPSC Min	99.23%	98.91%	95.84%	80.89%
HPSC Max	99.95%	99.39%	98.73%	94.54%	
Mandelbrot Set (%)	PC	99.65%	99.37%	98.66%	97.03%
	5020	99.09%	99.01%	-	-
	5040	99.15%	99.06%	99.00%	98.87%
	5545 Min	99.09%	99.01%	-	-
	5545 Max	99.15%	99.06%	99.00%	98.87%
	O-C2	98.84%	98.11%	97.17%	88.79%
	HPSC Min	98.12%	98.26%	97.39%	94.27%
HPSC Max	98.37%	98.40%	97.91%	96.26%	
Sobel Filter (%)	PC	99.04%	96.12%	95.54%	94.65%
	5020	98.69%	97.80%	-	-
	5040	98.26%	97.33%	96.48%	95.73%
	5545 Min	98.26%	97.33%	93.55%	92.82%
	5545 Max	98.69%	97.80%	-	-
	O-C2	96.63%	90.45%	82.56%	60.73%
	HPSC Min	95.91%	89.27%	83.83%	75.57%
HPSC Max	96.37%	91.04%	84.39%	75.95%	
Thumbnailer (%)	PC	99.71%	99.18%	98.50%	97.09%
	5020	99.45%	99.39%	-	-
	5040	99.20%	99.19%	99.13%	99.01%
	5545 Min	99.20%	99.19%	-	-
	5545 Max	99.45%	99.39%	97.63%	97.51%
	O-C2	97.10%	96.19%	94.57%	88.03%
	HPSC Min	96.88%	95.67%	92.68%	76.16%
HPSC Max	97.17%	96.17%	94.32%	83.53%	

C. ENERGY CONSUMPTION

Table 5: Energy Consumption for All Applications on All Platforms

Application	Platform	Serial	Number of Cores			
			1	2	3	4
Color Search (J)	PC	1.17	1.23	0.82	0.65	0.56
	5020	11.70	12.65	6.56	0.00	11.50
	5040	11.50	12.38	6.45	4.43	3.43
	O-C2	0.94	1.02	0.56	0.41	0.40
HSI (kJ)	PC	0.83	0.85	0.54	0.50	0.49
	5020	4.95	4.95	2.92	0.00	4.76
	5040	4.76	4.83	2.95	2.00	1.52
	O-C2	0.41	0.41	0.22	0.15	0.12
Mandelbrot Set (J)	PC	3.47	3.51	2.21	1.67	1.40
	5020	15.83	15.97	8.09	0.00	15.36
	5040	15.36	15.54	7.87	5.30	4.03
	O-C2	1.83	1.85	0.97	0.68	0.57
Sobel Filter (J)	PC	1.90	1.93	1.25	0.94	0.79
	5020	31.18	31.60	16.15	0.00	31.23
	5040	31.23	31.89	16.29	11.06	8.47
	O-C2	2.02	2.09	1.17	0.89	0.92
Thumbnailer (J)	PC	2.54	2.59	1.54	1.14	0.94
	5020	26.71	26.91	13.61	0.00	26.43
	5040	26.43	26.64	13.54	9.14	6.91
	O-C2	1.49	1.58	0.82	0.58	0.47

ACKNOWLEDGEMENTS

This research was funded by industry and government members of the NSF SHREC Center and the National Science Foundation (NSF) through its IUCRC Program under grant number CNS-1738783.

The authors wish to extend thanks to Mr. Joe Marshall at BAE Systems for providing the P5020DS system upon which many of these results were gathered, as well as offering feedback on early results.

The authors also wish to extend thanks to Dr. Robert Hodson at NASA Langley Research Center for providing insight into HPSC processor performance and reviewing results.

REFERENCES

- [1] P. J. Pingree, "Advancing NASA's on-board processing capabilities with reconfigurable FPGA technologies," *Aerospace Technologies Advancements*, 2010.
- [2] B. Schwaller, B. Ramesh and A. D. George, "Investigating TI KeyStone II and Quad-Core ARM Cortex-A53 Architectures for On-Board Space Processing," in *2017 IEEE High Performance Extreme Computing Conference (HPEC)*, Waltham, 2017.
- [3] J. Manning, D. Langerman, B. Ramesh, E. Gretok, C. Wilson, A. D. George, J. MacKinnon and G. Crum, "Machine-Learning Space Applications on SmallSat Platforms with TensorFlow," in *AIAA/USU Conference on Small Satellites*, Logan, 2018.
- [4] J. Krywko, "To make Curiosity (et al.) more curious, NASA and ESA smarten up AI in space," *Ars Technica*, 16 July 2018. [Online]. Available: <https://arstechnica.com/science/2018/07/from-oasis-to-noah-the-evolution-of-ai-in-modern-space-rovers/>. [Accessed 4 October 2018].
- [5] R. H. Maurer, M. E. Fraeman, M. N. Martin and D. R. Roth, "Harsh Environments: Space Radiation Environment, Effects, and Mitigation," *John Hopkins APL Technical Digest*, vol. 28, no. 1, pp. 17-29, 2008.
- [6] M. P. Petkov, "The Effects of Space Environments on Electronic Components," 2003.
- [7] BAE Systems, "RAD6000 Space Computers," 2006.
- [8] G. Shvets, "CPU-World," 2018. [Online]. Available: http://www.cpu-world.com/benchmarks/browse/910_80,965_61,993

_80,1035_96/?c_test=6&PROCESS=Show%20Selected. [Accessed 10 October 2018].

- [9] BAE Systems, "RAD750® radiation-hardened PowerPC microprocessor," 2008.
- [10] R. W. Berger, D. R. Bayles, R. D. Brown, S. E. Doyle, A. Kazemzadeh, K. R. Knowles, D. D. Moser, J. C. Rodgers, B. A. Saari, D. L. Stanley and B. E. Grant, "The RAD750(TM) - A Radiation Hardened PowerPC(TM) Processor for High Performance Spaceborne Applications," in *IEEE Aerospace Conference*, Big Sky, 2001.
- [11] TerraDaily Staff Writers, "BAE Computers To Manage Data Processing For Satellite Missions," *TerraDaily*, 23 June 2008.
- [12] R. Doyle, R. Some, W. Powell, K. Avery, G. Mounce, M. Lowry, M. Johnson, L. Bergman, W. Whitaker and M. Goforth, "Next Generation Space Processor (NGSP) High Performance Spaceflight Computing (HPSC) Next Steps at NASA and AFRL," NASA Jet Propulsion Laboratory, 2013.
- [13] T. M. Lovelly and A. D. George, "Comparative Analysis of Present and Future Space-Grade Processors with Device Metrics," *Journal of Aerospace Information Systems*, vol. 14, no. 3, pp. 184-197, 2017.
- [14] T. M. Lovelly, T. W. Wise, S. H. Holtzman and A. D. George, "Benchmarking Analysis of Space-Grade Central Processing Units and Field-Programmable Gate Arrays," *Journal of Aerospace Information Systems*, vol. 15, no. 8, pp. 518-529, 29 June 2018.
- [15] A. Grama, G. Karypis, V. Kumar and A. Gupta, Introduction to Parallel Computing, 2 ed., Boston, Massachusetts: Addison-Wesley Longman Publishing Co., Inc., 2002.
- [16] B. Barney, "Introduction to Parallel Computing," 25 June 2018. [Online]. Available: https://computing.llnl.gov/tutorials/parallel_comp/. [Accessed 6 October 2018].
- [17] B. Barney, "OpenMP," 28 June 2018. [Online]. Available: <https://computing.llnl.gov/tutorials/openMP/>. [Accessed 6 October 2018].
- [18] BAE Systems, "RAD5545™ multi-core system-on-chip power architecture processor," 2017.
- [19] Freescale Semiconductor, Inc., "QorIQ P5020 and P5010 communications processors," 2014.
- [20] Freescale Semiconductor, Inc., "QorIQ P5040 and P5021 communications processors," 2014.
- [21] R. F. Hodson, "The Next Generation of Spaceflight Computing," NASA, 2018.
- [22] J. Keller, "Boeing to develop next-generation radiation-hardened space processor based on the ARM architecture," *Military Aerospace Electronics*, 28 April 2017.
- [23] B. Schwaller, "Investigating, Optimizing, and Emulating Candidate Architectures for On-Board Space Processing," 2018.
- [24] R. Moss, "Finding Nearest Color Using Euclidean Distance," January 2017. [Online]. Available: <https://www.cyotek.com/blog/finding-nearest-colors-using-euclidean-distance>. [Accessed 18 April 2018].
- [25] C. Conger, A. Jacobs and A. D. George, "Multiparadigm Space Processing for Hyperspectral Imaging," in *2008 IEEE Aerospace Conference*, Big Sky, 2008.
- [26] J. Burkardt, "Mandelbrot Image Using OpenMP," 3 September 2012. [Online]. Available: https://people.sc.fsu.edu/~jburkardt/c_src/mandelbrot_openmp/mandelbrot_openmp.html. [Accessed 4 October 2018].
- [27] "Mandelbrot Set - C PPM Non-Interactive," [Online]. [Accessed 13 April 2018].
- [28] C. Wilson, J. Stewart, P. Gauvin, J. MacKinnon, J. Coole, J. Urriste, A. George, G. Crum, E. Timmons, J. Beck, T. Flatley, M. Wirthlin, A. Wilson and A. Stoddard, "CSP Hybrid Space Computing for STP-H5/ISEM on ISS," in *AIAA/USU Conference on Small Satellites*, Logan, 2015.
- [29] Geospatial Research Laboratory (U.S.), "Engineer Research and Development Center (ERDC)," US Army Corps of Engineers, 2015. [Online]. Available: <https://erdc-library.erdcresearch.dren.mil/xmlui/handle/11681/2925>. [Accessed 4 October 2018].

BIOGRAPHY



Evan W. Gretok earned his B.S. in Computer Engineering Technology from the University of Pittsburgh at Johnstown in 2017. He is currently pursuing a M.S. and Ph.D. in Electrical and Computer Engineering at the University of Pittsburgh. He is involved as a research task lead in the Autonomous Embedded Systems group at the NSF Center for Space, High-performance, and Resilient Computing (SHREC).



Evan T. Kain attained his B.S. in Electrical and Computer Engineering in 2017 from Cornell University. He is currently a Ph.D. student in Electrical and Computer Engineering at the University of Pittsburgh. He has worked extensively in architecture studies and distributed systems at the Air Force Research Laboratory Space Vehicles Directorate in Albuquerque. He is a research task lead at the NSF SHREC Center, also in the Autonomous

Embedded Systems group, with a focus on neuromorphic architecture studies.



***Alan D. George** is Department Chair, R&H Mickle Endowed Chair, and Professor of Electrical and Computer Engineering (ECE) at the University of Pittsburgh. He is Founder and Director of the NSF Center for Space, High-performance, and Resilient Computing (SHREC) headquartered at Pitt. SHREC is an*

industry/university cooperative research center (I/UCRC) featuring some 30 academic, industry, and government partners and is considered by many as the leading research center in its field. Dr. George's research interests focus upon high-performance architectures, applications, networks, services, systems, and missions for reconfigurable, parallel, distributed, and dependable computing, from spacecraft to supercomputers. He is a Fellow of the IEEE for contributions in reconfigurable and high-performance computing.