

SiC Processors for Extreme High-Temperature Venus Surface Exploration

Heewoo Kim, Javad Bagherzadeh, and Ronald G. Dreslinski

University of Michigan, Ann Arbor

Ann Arbor, Michigan, USA

{heewoo, javadb, rdreslin}@umich.edu

Abstract—Being the ‘sister planet’ of the Earth, surface exploration of Venus is expected to provide valuable scientific insights into the history and the environment of the Earth. Despite the benefits, the surface temperature of Venus, at 450°C, poses a large challenge for any surface exploration. In particular, conventional Silicon electronics do not properly function under such high temperatures. Due to this constraint, the most prolonged previous surface exploration lasted only for 2 hours.

Silicon Carbide (SiC) electronics, which can endure and function properly in high-temperature environments, is proposed as a strong candidate to be used in Venus surface explorations. However, this technology is still immature and associated with limiting factors, such as slower speed, power constraint, limited die area, and approximately 1,000 times longer channel than the state-of-the-art Si transistors.

In this paper, we configure a computing infrastructure for high-temperature SiC-based technology, conduct design space exploration, and evaluate the performance of different SiC processors when used in Venus surface landers. Our evaluation shows that the SiC processor has an average 16.6× lower throughput than the RAD6000 Si processor used in the previous Mars rover. The Venus rover with SiC processor is expected to have a moving speed of 0.6 meters per hour and visual odometry processing time of 50 minutes. Lastly, we provide the design guidelines to improve the SiC processors at the microarchitecture and the instruction set architecture levels.

I. INTRODUCTION

A. Venus Exploration

Venus is often called the sister planet of the Earth because it has a similar size and composition as the Earth [1], [2]. As such, exploring Venus is expected to provide unique insights into 1) the history of planetary evolution, 2) the evolution of the atmosphere, liquid water, climate, and geological activity, and 3) the presence of organic materials and biosignatures. These insights inform us of the habitability of Venus and other planets [1]–[3]. Despite Venus’s importance, its surface is not amenable to exploration due to its harsh environment. The surface temperature is exceptionally high at 450°C, along with the surface atmospheric pressure of 92 atm under the dense atmosphere mainly comprising of CO₂ (96.5%). Due to these harsh environmental conditions, previous Venus landers that the Soviet Union carried out lasted only for a short period; Venera 7, which landed on Venus in 1970, operated for 23 minutes, and Venera 13, in 1982, functioned for only 2 hours.

To overcome these severe environmental factors, NASA proposes to conduct Venus explorations at different altitudes; employing orbiters, aerial platforms, probes, and landers [1]–

[4]. In particular, landers are categorized into short-term landers, long-term landers, and mobile surface platforms [1]–[4]. Short-term landers are targeted to operate for a few hours to conduct the mineralogical analysis of rocks and soils [1], [4]. Long-term landers are to operate for one Venus day, which is about 120 Earth days, analyzing the temperature, pressure, wind speed and direction, atmospheric chemistry, and seismicity [1], [2], [4]. Lastly, the mobile surface platform will conduct the mineralogical analysis on a regional scale, take high-resolution panoramic images and return the surface samples [1], [4]. For long-term landers and mobile surface platforms that operate for hundreds of days, Si electronics are not suitable as they lose their semiconducting property above 300°C [5]. This paper investigates silicon carbide (SiC) processors that can conduct applications under extremely high-temperature conditions.

B. High-Temperature SiC Electronics

SiC maintains its semiconductive property up to 600°C, whereas Si loses its above 300 °C as the bandgap of SiC is three times wider than Si (3.2eV versus 1.1eV) [5], [6]. The wider bandgap enables SiC to endure higher voltage, power, radiation, and lower off-state leakage current compared to Si.

However, SiC technology faces several challenges to be used in digital circuits. First of all, the SiC technology is inferior to the Si technology in terms of the wafer size, manufacturing cost, and transistor size [6]–[8]. The channel length for the current available SiC transistors is in the order of μm , which is comparable to the Si digital circuits in the 1980s [4]. Due to these technical constraints, deploying SiC in digital electronics will result in lower speed, smaller area, and a more limited transistor budget and hence processors with lower complexities and functionality. In addition, the imbalance between electron and hole mobility (1150 cm^2/Vs versus 120 cm^2/Vs) results in a larger W_p/W_n ratio (3:1 to 5:1) for SiC CMOS, which makes the area budget problem more severe [6], [8], [9].

While there exist other wide-bandgap semiconductors such as gallium nitride (GaN) and diamond, SiC is the most advanced and mature technology among the candidates [7], [10].

C. Contributions

There are multiple challenges associated with designing a SiC-based processor for high-temperature Venus surface exploration. With a long channel length, limited power, speed, area, and transistor numbers, SiC processors should conduct scientific measurements, analysis, and autonomous navigation. In

this paper, we study SiC processor design and implementation, conduct design space exploration, and provide design solutions and guidelines for improving the SiC-based processor design. The contributions of this paper can be described as follows:

1) Configure SiC-based high-temperature computing infrastructure in the transistor, circuit, and system-level. From SiC transistor models, we characterize the SiC standard cell library. Using this cell library, we synthesize different architectures and analyze power, clock frequency, and area. We deploy gem5 cycle-level simulator to run benchmark programs and model the system-level performance.

2) Design space exploration. We analyze how the SiC core design and complexity affect metrics such as delay, power, energy-delay-delay product (EDDP), and cache capacity. We choose EDDP over EDP to compare technologies with various voltage domains.

3) Performance evaluation to meet the requirements for a Venus exploration mission. Based on the previous Mars and future Venus missions data, we try to answer if and how SiC-based processor design can meet the constraints and workloads.

4) Provide guidelines and design tips to improve the SiC processors at the microarchitecture and the instruction set architecture (ISA) levels. We investigate what the pipeline design and its complexity should be, how a customized ISA would look, and how the microarchitecture and ISA level tunings affect the SiC processors' power, speed, and die area.

II. RELATED WORK

NASA provides a SiC n-type junction-gate field-effect transistor (JFET) model with a $3 \mu\text{m}$ channel length and $25 \text{ V } V_{dd}$ [11]. KTH Royal Institute of technology provides a SiC NMOS model that has a channel length equal to $5 \mu\text{m}$ and V_{dd} of 5 V [8]. The University of Arkansas has been providing a model for their SiC CMOS technology with a shorter channel length equal to $1.2 \mu\text{m}$, V_{dd} of 15 V , and W_p/W_n ratio of 5:1 [12].

NASA fabricated a SiC sub-MHz clock generator (175 JFETs) that can last for 60 days at Venus surface where the temperature is 500°C , and atmospheric pressure is 92 atm with a composition made of CO_2 , N_2 , etc [13]. KTH Royal Institute of technology developed a SiC process design kit and fabricated a 4-bit microcontroller and SRAM with 10,000 transistor-transistor logic (TTL) devices [10]. Ozarc IC and the University of Arkansas taped out a clock generator, a ring oscillator, and a D-flip flop that last for 100 hours in 470°C using Raytheon UK's SiC HiTSiC CMOS process [9]. From this, they simulated OpenMSP 16-bit 0.5 MHz microcontroller and SRAM for 470°C operation [9]. Previous works focus on fabrication and performance of the SiC technology on transistor and circuit levels and try to extend the operation hours in a high-temperature environment. Although previous works have synthesized and fabricated the processors, there has not been any architecture-level studies and design space exploration for Venus surface exploration workloads and constraints yet. In this work, we investigate how different options and choices will affect the processor performance, and suggest optimizations in different design levels to meet the design constraints and workload requirements.

TABLE I
THE RADIATION-HARDENED PROCESSORS AND THEIR USAGE IN MARS

	RAD750 [17]	RAD6000 [18]
Used in	Curiosity (2011) Perseverance (2020)	Spirit (2003) Opportunity (2003)
Technology	250~150 nm	500 nm
Clock frequency	133~200 MHz	2.5~20 MHz
Cache size	32kB L1 I/D	8kB L1
Power	5 W	5~20 W

III. METHODOLOGY

A. Implementation

We study the SiC transistor model and its fan-out of 4 (FO4) delay, power, area, and temperature dependency with Spectre (Cadence) and Hspice (Synopsys). Standard cell library characterization is done using Liberate (Cadence). We deploy Design Compiler (Synopsys) to synthesize different pipelines and analyze area, power, and clock period. We choose the gem5 cycle level simulator to run the benchmark programs instead of VCS (Synopsys) because gem5 simulation is much faster, and it is accurate enough for the purpose of this study [14]. The benchmark programs are compiled to RISC-V binaries with riscv-gnu-toolchain [15]. The dynamic binary analysis of the benchmarks is done with the rv8 simulator [16].

B. Design Specification

Among various SiC transistor models, we choose the University of Arkansas model as it is a CMOS process and has the shortest channel length ($1.2 \mu\text{m}$) [19]. The limitation of this model is that the upper limit of transistor modeling temperature is 300°C [12], not 450°C , so we conduct the high-temperature analysis at 300°C . The upper temperature limits of models by KTH and NASA are 450 , and 500°C , respectively, but the drawbacks include that they only have nfet, long channel length and high supply voltage. We choose INVX1, NAND2X1, NOR2X1, TBUF1, DFFSR, and LATCH cells for the baseline standard cell library design.

We also rely on RAD750 and RAD6000 processors as the Silicon counterparts that are shown in Table I. The RAD750 is a radiation-hardened version of the PowerPC750 microprocessor and widely used for space applications such as satellites, Mars 2011 curiosity rover, and Mars 2020 perseverance rover [20], [21]. We choose the second generation of RAD750 in our analysis, which is manufactured with 180 nm Si technology nodes with 32 kB for each of I/D cache and operates at 200 MHz where its power consumption is 5 W [17], [22].

The RAD6000, which is an earlier version of the radiation-hardened microprocessors compared to RAD750, was used for Mars 2003 spirit and curiosity rovers [23]. RAD6000 has a 32-bit architecture and is manufactured with a $0.5 \mu\text{m}$ technology with 8 kB L1 cache, 2.5~20 MHz operation frequency, and 5 ~ 20 W power consumption [18].

RAD6000 is used in earlier stages of Mars exploration and is a more suitable target for comparison in this work. Based on Mars exploration mission history, the earlier missions are more conservative and usually contain more basic tasks to gather information for paving the road for future and more

TABLE II
SiC PROCESSORS DESIGN SPECIFICATIONS AND CONSTRAINTS

SiC transistor	SiC die area	Power capacity
Channel length 1.2 μm	49, 150.5, 250 mm^2	Day time : 4.216 W
$W_p/W_n = 5:1$		Night time : 0.8785 W

complex rovers. Moreover, the natural risk factors and high costs of each mission for planetary exploration entail a more robust and simple design to lower the risk factors. Comparing to Mars, Venus has a much harsher environment and is much less studied. Designs similar to Mars's early-stage exploration that benefited from RAD6000 are a better fit for estimating the required design specifications and constraints in this study.

To meet the power and area budgets of SiC technology and considering the architectures and workloads of RAD750 and RAD6000, we use four different 32-bit architectures that are VeriSimple Alpha (classic 5-stage), ARM Cortex-M0, TPISA core (24-bit), and Vanilla-5 [24]–[26]. We consider two options for the Vanilla-5 core, one equipped with the floating-point unit [26] and one that has only the integer unit [25]. These simple processors consume less power and area compared to more complicated ones. We choose RISC-V ISA because it is open-source, small and simple, to avoid over-architecting and ease of extensions [14], [27].

We choose nine benchmark programs including Dhrystone, Whetstone, Coremark, and six benchmarks from the MiBench suite (basicmath, bitcount, dijkstra, fft, qsort, and stringsearch) [22]. The benchmarks are open-source and are already used for several studies of aerospace processors [22], [28], [29]. They include several workloads for integer and floating-point arithmetic, sorting and searching, graph traversal, trigonometric calculation, and state machine implementation, which are used for automotive, control, communication, and networking [22].

C. Design Constraints

Immaturity of the SiC technology and the harsh environment of Venus entail restrictions on the die area and power. 49 mm^2 is a reported SiC die area without any yield issue [10], [30]. However, another research that fabricated SiC CPU with TTL logic achieved a chip size equal to 150.5 mm^2 [10]. The other group has been able to synthesize a SiC microprocessor with an area equal to 250 mm^2 [9]. In this work and based on these reported die areas, we consider three cases for the SiC die area of 49 mm^2 , 150.5 mm^2 , and 250 mm^2 to study how increasing the SiC die area can affect the performance.

The Venus lander will generate its power using a wind turbine and solar panels [31]. The wind turbine generates 3.18 W day and night, and the solar panel generates 10.27 W of power only during the daytime [31]. They estimated that 35% of the generated power is wasted in transmission loss, 0.31 W is continuously required for the communication system [31], and a 50% power margin should be added [31] for system stability. Based on these numbers and our calculation, the Venus lander's processor's power capacity is 4.216 W for the daytime and 0.8785 W for the nighttime. The SiC processor design specifications and constraints are summarized in Table II.

TABLE III
FO4 ANALYSIS

INVX1 FO4	Delay	Power	Area
Si AMI 500 nm [32]	0.158 ns	93.34 μW	144 μm^2
Si TSMC 180 nm [32]	0.075 ns	6.86 μW	16 μm^2
SiC 1.2 μm 25°C [19]	2.17 ns	4373 μW	1152 μm^2
SiC 1.2 μm 300°C [19]	2.63 ns	4493 μW	1152 μm^2

TABLE IV
DC SYNTHESIS RESULTS OF PIPELINES IN 25°C

SiC1.2um 6cells	Stages	Clock Period	Power	Cell Count	Area
VeriSimple Alpha	5	627 ns	4232 mW	79,340	127.77 mm^2
CortexM0	3	499 ns	153 mW	55,185	88.62 mm^2
TPISA core	1	124 ns	315 mW	1,390	2.19 mm^2
Vanilla-5 (w/ FPU)	5	498 ns	1527 mW	120,122	204.2 mm^2
Vanilla-5 (w/o FPU)	5	498 ns	435 mW	51,007	91.36 mm^2

IV. EVALUATION

A. SiC Transistor Performance

As shown in Table III, we compare the fan-out of 4 (FO4) delay, power, and area of the Si INVX1 at room temperature, the SiC INVX1 at room as well as a high temperature (300°C). The input signal is a square pulse with the period of 20 ns. For Si, we used AMI 500 nm and TSMC 180 nm technologies since the RAD6000 and RAD750 are manufactured with 500 nm and 180 nm technology nodes. The feature size of the SiC model is 1.2 μm . Based on the analysis in Table III, we can make two observations:

1) **A considerable performance gap between the Si and SiC electronics at gate-level.** Comparing to the AMI 500 nm (TSMC 180 nm) technology, the SiC INVX1 is 13.7 (28.9) times slower, consumes 46.9 (637) times more power in the same frequency, and requires 8 (72) times more area.

2) **Temperature dependency of SiC technology.** At 300°C, SiC INVX1 FO4 delay and power consumption increases by 21.2% and 2.74% compared to the room temperature.

B. DC Synthesis Results of SiC Pipelines

The frequency of the SiC-based cores ranges from 1.59 MHz to 8.06 MHz , which is calculated from the clock period ranges from 627 ns to 124 ns in Table IV. Power consumption also ranges from 153 mW to 4232 mW , so all pipelines except for the VeriSimple Alpha and the Vanilla-5 with FPU satisfy the nighttime power budget. The VeriSimple Alpha pipeline consumes slightly more power than the daytime budget. The area of SiC pipeline varies from 2.193 mm^2 to 204.2 mm^2 . Only TPISA core satisfies the 49 mm^2 area constraint. The cache size for each core is calculated based on the remaining SiC die area considering the area budget in Table II, which is the total SiC die area - pipeline area. In our calculations, we consider the area for a 1 kB SRAM memory array with sense amplifiers to be 19.3 mm^2 in SiC 1.2 μm technology [9].

C. Figure of Merit (FOM) Analysis

Fig. 1 shows the clock period delay of the evaluated cores discussed in Section IV-B. The TPISA core that is a simpler pipeline is only 2 times slower in terms of clock delay compared to the RAD6000, but the more complex pipelines such




Fig. 1. Pipeline delays in 25°C.




Fig. 2. FOM graph. FOM values are normalized to the RAD6000 values. Inside the same pipeline category, the left most one is the cache metric for 49 mm^2 die area, the middle one is for 150.5 mm^2 die area, and the right most one is for 250 mm^2 die area. Negative cache metric means that the pipeline area is bigger than the SiC die area.

as the VeriSimple Alpha are more than 12 times slower than the RAD6000.

To further evaluate the performance of SiC processors, we define two FOMs which are the cache metric as shown in equation (1), and the energy-delay-delay product (EDDP) metric as described in equation (2). Both of the metrics are normalized to RAD6000's value, and the higher the metric value is, the better the performance of the processor would be.

$$\text{Cache Metric} = \frac{\text{SiC Processor Cache Capacity}}{\text{RAD6000 Cache Capacity}} \quad (1)$$

The cache metric shows the tradeoff between the pipeline area and the cache area. Deploying a more complex architecture with more area potentially can improve performance. But if we consider a constant budget for the overall area, it will reduce the cache capacity and deteriorate system performance. Cache metric is also related to the SiC technology maturity. Advanced SiC technology will increase the SiC die area and decrease the SiC cache unit area, which will improve the cache capacity and cache metric.

$$\text{EDDP Metric} = \frac{\text{RAD6000 EDDP}}{\text{SiC Processor EDDP}} \quad (2)$$

We choose EDDP instead of EDP for energy and delay analysis since we are comparing pipelines with various voltage domains. EDDP metric also depends on SiC technology maturity. According to Dennard scaling [33], reducing the channel length by the factor of $s > 1$ will reduce the power by s^2 and the delay by s , which will reduce the EDDP by s^5 .

TABLE V
CPU SPECIFICATIONS IN GEM5 SIMULATION.

Processor	Clock Freq	I/D Cache Size	Gem5 CPU Model
Si RAD750	200 MHz	32 kB each	DerivO3 (Out-of-Order)
Si RAD6000	20 MHz	4 kB each	DerivO3 (Out-of-Order)
SiC T4	8 MHz	4 kB each (250 mm^2)	TimingSimple (1 Stage, ALU only)
SiC V1	2 MHz	1 kB each (250 mm^2)	Minor (In-Order)
SiC V2	2 MHz	2 kB each (testing)	Minor (In-Order)
SiC V4	2 MHz	4 kB each (RAD6000)	Minor (In-Order)
SiC V4 O3	2 MHz	4 kB each (RAD6000)	DerivO3 (Out-of-Order)




Fig. 3. Throughput of benchmarks in SiC processors. SiC T4 is the TPISA core based SiC processor for a 250 mm^2 die area with I/D cache size of 4 kB. SiC V1 is the Vanilla-5 with FPU based SiC processor for 250 mm^2 die area with I/D cache size of 1 kB. SiC V2 and SiC V4 are conducted to see the cache capacity dependency of the performance. SiC V4 O3 is the Out-of-Order version of the SiC V4.

As shown in Fig. 2, the TPISA core, which is the smallest and least complicated design, has a higher EDDP metric compared to the RAD6000. Cortex-M0 has about 10% of the RAD6000's EDDP metric value, and Vanilla-5 has about 1~5% of the RAD6000's EDDP metric value. The VeriSimple Alpha pipeline, which is the most complicated design, has the lowest EDDP metric value.

By considering a budget of 250 mm^2 as the die area, we can assign equal or even larger cache capacity for TPISA core, CortexM0, and Vanilla-5 w/o FPU compared to the RAD6000. However, CortexM0, Vanilla-5, and VeriSimple Alpha pipeline are too large to meet the 49 mm^2 SiC die budget. Vanilla-5 with FPU is also too large to fit within a 150.6 mm^2 die area.

TPISA core has the best EDDP and cache metric scores compared to the other processors, and its EDDP metric is even higher than the RAD6000. TPISA core is a customized ISA and microprocessor for low-cost disposable printed microelectronics and is optimized for small area constraints, low power budget, and printed applications such as food temperature sensors and smart bandages [24]. TPISA core's FOM result shows that ISA and microarchitecture customization can significantly improve the power, delay, and cache capacity by designing a simple pipeline that has efficient data and control flows. In section V-B, we discuss the method to implement a Venus Exploration Specific ISA(VES-ISA).

D. Application-Level Analysis in Gem5

Table V shows the parameters used for the gem5 simulation. We use DerivO3CPU to model both RAD750 Si Processor as well as RAD6000. We use DerivO3CPU for RAD750 since it is an out-of-order superscalar pipeline [22]. RAD6000 processes fixed point operations out-of-order and floating point interactions in-order [34]. However, the control instructions for floating point operations, such as 'for loop' or 'if' are

executed in out-of-order fixed-point units. Therefore, we use DerivO3CPU setting for RAD6000 and model its performance in a more aggressive position.

We refer to the TPISA core and the Vanilla-5 with FPU for SiC processors. Considering that the TPISA core is a single-stage integer-only pipeline, we use the TimingSimpleCPU model for TPISA core-based SiC processor simulation. Both the TimingSimpleCPU and TPISA core compute non-memory instructions in a single cycle, but only the TPISA core has detailed microarchitecture such as the ALU, reg files. SiC T4, for which we use this specification, is tailored for a 250 mm^2 SiC die and has a 4kB I/D cache capacity.

Since Vanilla-5 is a single-issue in-order pipeline [35], we use the MinorCPU model with a 4-stage in-order pipeline setting for Vanilla-5 with FPU based SiC processor simulation. Both the MinorCPU and the Vanilla-5 are in-order pipelines, but the microarchitectural details are different (e.g., MinorCPU has four stages, two of which are fetch stages, but Vanilla-5 has five, only one of which is a fetch stage). We use this specification for SiC V1, V2, and V4. SiC V1, with 1 kB I/D cache capacity, is tailored for 250 mm^2 SiC die. SiC V2 and SiC V4, with 2 kB and 4 kB as the I/D-cache capacity respectively, are simulated to study how cache size affects the performance.

We compile the benchmark with RV64I ISA for SiC T4, and RV64GC ISA for SiC V1, V2, V4, and V4 O3. We apply the -O3 compiler optimization. We use RV64 instead of RV32 because RV32 does not work correctly in gem5. SiC V4 O3 has an identical specification to the SiC V4 but has an out-of-order pipeline and is simulated for our analysis in section V-A.

For RAD750, We compare the obtained gem5 throughput results with the numbers from [22] by running the same benchmarks. The comparison shows that the gem5 average throughput is 49% higher than the reported values. The mismatch in throughput is due to the architectural details, such as the branch predictors, superscalar width, ISA, etc.

Figure 3 shows the throughput for the modelled SiC processors for each benchmark. The results show that **the SiC T4 and SiC V1 are 23.1 \times and 16.6 \times slower on average in terms of throughput compared to the RAD6000.**

Comparing SiC T4 (TPISA core, 4 kB cache) with the SiC V1 (Vanilla-5, 1 kB cache) shows the following: on average, the throughput of SiC T4 is 0.7 \times that of SiC V1. However, in-depth analysis shows that the performance differs significantly by the benchmarks; while SiC T4 is 1.7 \times better on integer benchmarks, it suffers at floating-point benchmarks by approximately 10 \times . For the floating-point calculation, the low binary code density in TPISA core has a larger negative impact on throughput than that of the small cache and low clock frequency in Vanilla-5.

To estimate and evaluate the actual Venus lander performance that will deploy the SiC processor, we refer to the data available from Mars Spirit and Opportunity rovers that were equipped with 20 MHz RAD6000 [23]. When the rovers used autonomous navigation (AutoNav) in obstacle-laden terrain, the driving speed was up to 10 meters/hour [23]. Each step of the Visual Odometry (VisOdom) processing required three

minutes, limiting the speed of the rovers [23]. Since the bottleneck is the computation time [23], the 16.6 \times lower SiC throughput translates to a Venus lander with a Vanilla-5 based SiC processor having **AutoNav speed of 0.6 meters/hour and VisOdom processing time of 50 minutes.**

Figure 3 also shows that increasing the cache capacity from 1 kB to 2 kB and 4 kB improves the throughput by 2.6% and 4.3%, respectively. The throughput improvement comes from the decrease in the cache miss rate as the cache capacity increases. But the throughput gain is not significant because the 1 kB cache is enough for the benchmarks used in this study.

V. WAYS TO IMPROVE SiC PROCESSORS

Beyond the obvious scaling of the transistor, there are also techniques that can be taken at above the device level to improve performance. This section focuses on two of those, one at the microarchitecture and one at the ISA.

A. Microarchitecture

The difference between the daytime and the nighttime power capacities of the Venus lander provides an opportunity for more power optimizations. Dynamic voltage and frequency scaling (DVFS) can be adopted so that during the daytime, tasks are done faster, and during the nighttime, the tasks are done at a lower speed. Big/little cores can also be adopted once the SiC die area increases or SiC transistor size scales down. Similar to DVFS, an out-of-order big core can be used during the daytime, and an in-order small core is deployed during the night. As shown in figure 3, we observe that an out-of-order SiC processor with 4 kB I/D cache has an average 59.4% higher throughput compared to an in-order SiC processor, which shows that a big core can boost the Venus lander exploration performance when the power and area are not constrained.

B. Instruction Set Architecture

Optimizing a specific ISA for the Venus surface lander applications (VES-ISA) is another way to boost the performance, similar to the TPISA core [24] mentioned in Section IV-C. Venus exploration specific ISA can reduce the code size, area, power and improve the throughput and program runtime. To be specific, VES-ISA decreases the compiled assembly code size and the instruction data size, which reduces the number of instructions fetched from memory and executed in the pipeline. Deploying a VES-ISA based microarchitecture avoids redundant hardware modules to give an opportunity to create additional units that can boost performance.

Customizing the number and size of registers and operands helps to configure the program-specific VES-ISA [24]. VES-ISA will be configured based on the RISC-V ISA to reduce the design overhead. VES-ISA that is a subset of the RISC-V ISA can include custom instructions such as square root and transcendental functions [24]. The number of registers can be customized to match the workloads [24], such as RV32E (RISCV32 Embedded ISA), which has only 16 registers instead of 32. The size of each register can also be reduced by minimizing the address and instruction size to match the application program [24], such as the 24-bit TPISA core. A smaller ISA

set and instruction size can also reduce the decoder size and delay. However, the trade-off between the code density and the pipeline implementation simplicity should be considered when minimizing the program size.

We conduct dynamic binary analysis of the benchmarks with the rv8 simulator [16]. Since MULH, MULHU, MULHSU, DIVU, DIV, REM, and REMU instructions are used infrequently (less than 0.5 %), we remove these, thereby simplifying the multiplier and removing the divider. This modification saves 2.86 % power, 3.35 % standard cell counts, and 3.61 % total die area at the cost of 9.11 % more instructions to run. The small observed gain is likely due to the wide range of general-purpose benchmarks. However, specific application benchmarks for Venus landers would allow us to tune the instruction sets further to save more area and power with smaller runtime costs, or reduce the runtime with small power and area overhead.

VI. FUTURE WORK

First of all, we plan to search and analyze the benchmarks for Venus lander-specific applications. Then we will configure a VES-ISA and design the customized microprocessors discussed in sections V-A and V-B. We also plan to consider the main memory. There are reports that low-power high-temperature data storage is not imminently available [2], but are in development [4]. Periodically broadcasting the data to an orbiter is suggested to solve the memory issue [4]. Some ferroelectric materials maintain the ferroelectricity at high temperatures [30], which shows the possibility for high-temperature Fe-RAM or Fe-FET as the main memory.

VII. CONCLUSION

The durability of SiC electronics in high temperatures makes them suitable for Venus surface exploration. However, research on both the SiC electronics and the Venus surface exploration is in their early stages. According to our evaluation, SiC processor has $16.6\times$ lower throughput in average compared to RAD6000, and the Venus lander is expected to have the AutoNav speed of 0.6 meters/hour and the VisOdom processing time of 50 minutes. To remedy an order of magnitude gap in the performance compared to RAD6000, we suggested improvements in customized microarchitecture and ISA. Venus lander that is equipped with the improved SiC processor based on our studies would be able to successfully conduct Venus surface exploration and have a similar performance compared to the Spirit and Opportunity on their Mars missions.

REFERENCES

- [1] J. A. Cutts, M. Amato *et al.*, "Roadmap for Venus exploration," 2019.
- [2] R. Grimm, M. Gilmore *et al.*, "Venus bridge summary report," 2018.
- [3] M. Gilmore, P. Beauchamp *et al.*, "2020 Venus flagship mission study," 2020.
- [4] G. Hunter, J. Balcerski *et al.*, "Venus technology plan," 2019.
- [5] P. G. Neudeck, R. S. Okojie *et al.*, "High-temperature electronics-a role for wide bandgap semiconductors?" *Proceedings of the IEEE*, vol. 90, pp. 1065–1076, 2002.
- [6] H. Kim, A. Amarnath *et al.*, "A survey describing beyond Si transistors and exploring their implications for future processors," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 17, pp. 1–44, 2021.
- [7] P. G. Neudeck, S. L. Garverick *et al.*, "Extreme temperature 6H-SiC JFET integrated circuit technology," *physica status solidi (a)*, vol. 206, pp. 2329–2345, 2009.
- [8] D. Kimoto, "Characterization and modeling of SiC integrated circuits for harsh environment," 2017.
- [9] J. Holmes, A. M. Francis *et al.*, "Extended high-temperature operation of silicon carbide CMOS circuits for venus surface application," *Journal of Microelectronics and Electronic Packaging*, vol. 13, pp. 143–154, 2016.
- [10] M. Shakir, "Process design kit and high-temperature digital ASICs in silicon carbide," Ph.D. dissertation, KTH Royal Institute of Technology, 2019.
- [11] P. G. Neudeck, "Technical primer on design and SPICE modeling of circuits for NASA Glenn SiC JFET IC Version 12 prototype wafer run part 1: SiC JFET behavior and SPICE modeling," 2019.
- [12] S. Ahmed, "Modeling and validation of 4H-SiC low voltage MOSFETs for integrated circuit design," 2017.
- [13] P. G. Neudeck, L. Chen *et al.*, "Operational testing of 4H-SiC JFET ICs for 60 days directly exposed to Venus surface atmospheric conditions," *IEEE Journal of the Electron Devices Society*, vol. 7, pp. 100–110, 2018.
- [14] A. Roelke and M. R. Stan, "Risc5: Implementing the RISC-V ISA in gem5," in *First Workshop on Computer Architecture Research with RISC-V (CARRV)*, 2017.
- [15] SiFive, "GNU toolchain," <https://www.sifive.com/software>, 2020.
- [16] M. Clark and B. Hoult, "rv8: a high performance RISC-V to x86 binary translator," in *First Workshop on Computer Architecture Research with RISC-V (CARRV)*, 2017.
- [17] D. Rea, D. Bayles *et al.*, "PowerPC™ RAD750™-a microprocessor for now and the future," in *2005 IEEE Aerospace Conference*. IEEE, 2005, pp. 1–5.
- [18] "RAD6000 Space Computers," https://montcs.bloomu.edu/~bobmon/PDFs/RAD6000_Space_Computers.pdf.
- [19] U. of Arkansas, "BSIMSIC: BSIM Models Compatible for SiC Low-Voltage MOS Devices," <https://mscad.uark.edu/compact-models/>, 2019.
- [20] "The Mars 2020 Rover's Brains," <https://mars.nasa.gov/mars2020/spaceraft/rover/brains/>.
- [21] "NASA Seeks High-Performance Spaceflight Computing Capabilities," <https://www.nasa.gov/topics/technology/features/spaceflight-comp.html>.
- [22] M. R. Gardiner, "An evaluation of soft processors as a reliable computing platform," 2015.
- [23] M. Maimone, J. Biesiadecki *et al.*, "Surface navigation and mobility intelligence on the Mars exploration rovers," *Intelligence for space robotics*, pp. 45–69, 2006.
- [24] N. Bleier, M. H. Mubarak *et al.*, "Printed microprocessors," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2020, pp. 213–226.
- [25] "Vanilla Bean," https://bitbucket.org/taylor-bsg/bsg_manycore/src/master/v/vanilla_bean/.
- [26] "Vanilla Bean with FPU," https://github.com/bespoke-silicon-group/bsg_manycore/tree/master/v/vanilla_bean.
- [27] A. Waterman, Y. Lee *et al.*, "The RISC-V instruction set manual. Volume 1: User-level ISA, Version 2.0," California Univ Berkeley Dept of Electrical Engineering and Computer Sciences, Tech. Rep., 2014.
- [28] S. Nakabeppu, Y. Ide *et al.*, "Space responsive multithreaded processor (SRMTP) for spacecraft control," in *2020 IEEE Symposium in Low-Power and High-Speed Chips (COOL CHIPS)*. IEEE, 2020, pp. 1–3.
- [29] C. H. Koo and H. Kim, "Measurement of cache-related preemption delay for spacecraft computers," in *2018 IEEE 24th International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*. IEEE, 2018, pp. 234–235.
- [30] C. F. Wilson, C.-M. Zetterling *et al.*, "Venus long-life surface package," *arXiv preprint arXiv:1611.03365*, 2016.
- [31] J. Sauder, E. Hilgemann *et al.*, "Automation rover for extreme environments," 2017.
- [32] "OSU Standard cells," https://vlsiarch.ecen.okstate.edu/flows/MOSIS_SCMOS/osu_stdcells_v2.4, 2017.
- [33] R. H. Dennard, F. H. Gaensslen *et al.*, "Design of ion-implanted MOS-FET's with very small physical dimensions," *IEEE Journal of Solid-State Circuits*, vol. 9, pp. 256–268, 1974.
- [34] C. R. Moore, D. Balsler *et al.*, "IBM single chip RISC processor (RSC)," in *Proceedings 1992 IEEE International Conference on Computer Design: VLSI in Computers & Processors*. IEEE Computer Society, 1992, pp. 200–201.
- [35] S. Davidson, S. Xie *et al.*, "The Celerity open-source 511-core RISC-V tiered accelerator fabric: Fast architectures and design methodologies for fast chips," *IEEE Micro*, vol. 38, pp. 30–41, 2018.