Signal and Power Integrity Design Methodology for High-Performance Flight Computing Systems

Nicholas G. Franconi, Alan D. George
NSF Center for Space, High-Performance, and Resilient Computing (SHREC)
University of Pittsburgh
Pittsburgh, PA, USA
nicholas.franconi@nsf-shrec.org
alan.george@pitt.edu

Alessandro D. Geist[1], Dennis Albaijes[2]
Science Data Processing Branch (Code 587) [1]
Flight Data Systems and Radiation Effects Branch (Code 561) [2]
NASA Goddard Space Flight Center
Greenbelt, MD, USA
{alessandro.d.geist, dennis.albaijes}@nasa.gov

2021 IEEE Space Computing Conference (SCC) | 978-1-6654-2400-4/21/$31.00 ©2021 IEEE | DOI: 10.1109/SCC49971.2021.00011

***Abstract*—Computing capabilities of space systems have increased onboard performance by orders of magnitude with the use of radiation-tolerant field-programmable gate arrays (FPGA) and processors. The incorporation of signal and power integrity analysis with printed circuit board (PCB) design in reliable computing architectures for space systems has become critical to enable future mission capabilities. Developers launch high-performance processors into a breadth of orbits and missions, running varying applications that create challenges for designing reliable computing hardware. Specifically, for these designs, academic and industry research has focused on component radiation performance, fault mitigation, and reliable architectures. However, other design parameters including electromagnetic interference (EMI), PCB stackup, signal integrity (SI), voltage regulator module (VRM) design, and power distribution network (PDN) are often deprioritized or disregarded as the design matures. Since these characteristics are becoming more significant in high-performance processor designs, this research presents a hardware design and analysis methodology for high-performance, space-computing systems that focuses on a holistic design approach and PDN reliability. While these challenges exist across all space hardware, the reduced PCB dimensions imposed by SmallSats and CubeSats introduce additional hurdles, specifically to VRM and decoupling design. By examining the relationship between the PDN and radiation performance, an analytical relationship is developed that incorporates Total Ionizing Dose and Single-Event Transients to ensure reliability throughout the mission duration. The presented design methodology is applied to the SpaceCube v3.0 Mini, an FPGA-based on-board science data processing system developed at NASA Goddard Space Flight Center.**

***Index Terms*—Power Distribution Network, Target Impedance, Signal Integrity, Printed Circuit Board, Space Computing**

I. Introduction

The increased accessibility to space provided by a multitude of launch opportunities has not led to more prominent advances in core spacecraft technologies (e.g. onboard computing). Due to unique mission architectures, innovations in sensor technology and more frequent launches, mission designers are able to propose more advanced concepts including spacecraft autonomy and machine learning that outpace the capabilities of onboard space processing. In the 2018 decadal strategy for Earth observations from space, the National Academies' Space Studies Board (SSB) highlighted the need for more advanced analysis methodologies to efficiently use the limited processing resources available to spacecraft [1]. NASA and AFRL have jointly identified improvements to onboard space computing as a "technology multiplier" and addressed a broad range of computing architectures requirements for future missions [2]. The next-generation of science and defense missions will require spacecraft processors to support a variety of computationally intensive tasks (e.g. multi core-processing), high-throughput sensor data processing (e.g. hyperspectral imagers), autonomous spacecraft and constellation operation (e.g. Blackjack "Pit Boss"), and real-time data generation and compression [3]. These systems require advanced machine-learning and deep-learning algorithms that have been developed and optimized for SmallSat and CubeSat applications to enhance these identified needs [4].

When evaluating the requirements for future missions, system designers must trade operational resiliency within the hazardous space environment, address application performance needs while meeting size, weight, power, and cost (SWaP-C) constraints [5]. Limitations of radiation-hardened (rad-hard) processors exacerbate the ever-increasing mission processing requirements and have led to the adoption of radiation-tolerant processors, field programmable gate arrays (FPGAs), and hybrid system-on-chip (SoC) devices that provide a suitable balance between processing capability, reliability, and SWaP-C [6], [7]. The increased availability of SmallSat and CubeSat launches has driven a transition from larger 6U (223×160 mm^2) printed circuit board (PCB) form factors to smaller 3U (100×160 mm^2) and even 1U (100×100 mm^2) dimensions. The reduction in PCB area combined with the increased processing capabilities from radiation-tolerant devices present a number of obvious advantages for mission designers, but also introduce significant engineering challenges during the design of architecture and layout, specifically relating to component density, reliability, power distribution networks (PDN), power integrity (PI), and signal integrity (SI). While advances in PCB manufacturing and assembly technologies are enabling for commercial processors, additional research is required to understand the impact of these advances on reliability [8]. When comparing devices incorporated into these PCB assemblies, researchers have historically focused on

978-1-6654-2400-4/21/$31.00 ©2021 IEEE
DOI 10.1109/SCC49971.2021.00011

the reliability and performance metrics of rad-hard processors and FPGAs [9]. However, these studies do not consider the radiation and reliability effects on complex and dense PDN designs with large step-load current transients required by state-of-the-art FPGAs and processors [10]. These transients can easily stress point-of load (PoL) converters, decoupling networks, and PCBs beyond their capabilities if not properly accounted for. The reliability requirements levied on spacecraft avionics and instrument processors combined with smaller PCB dimensions, high-density interconnects (HDI), limited radiation-hardened components, and mil-spec capacitors selection requires a re-evaluation of traditional spaceflight design methodologies for future missions.

In this publication, the current ad-hoc design process is discussed with commonly referenced rules-of-thumb for traditional larger spacecraft avionics. We incorporate this process with new research to present a design and analysis methodology that addresses the challenges of high-performance spaceflight processors design. Additionally, we present a novel methodology that enhances the traditional target impedance approach and incorporates total ionizing dose (TID) and single-event transient (SET) effects. Our hybrid approach provides a clear path to reliably evaluate high-performance spaceflight computing systems and tailor performance to mission specific orbits and operating environments. The proposed methodology is applied to the SpaceCube v3.0 Mini processor card that features the Xilinx Kintex UltraScale FGPA, and was developed at NASA Goddard Space Flight Center in the Science Data Processing branch. The remainder of this publication is organized as follows. Section II provides an overview of rad-hard processors, FPGAs, and hybrid SoCs for space applications, SmallSat reliability metrics, and power distribution networks for space hardware. Section IV describes our design and analysis methodology with the hybrid target impedance. Section V applies the design framework to the SpaceCube v3.0 Mini processor card including insights gained throughout the development process. Finally, Section VI provides concluding statements and insights for future work.

II. Background

This section provides an overview of rad-hard and rad-tolerant FPGAs, processors, and SoCs including their benefits and challenges for space computing. Radiation effects and reliability in SmallSat and CubeSat missions are presented as it pertains to the PDN and system-level design. Finally, we summarize challenges with space-grade capacitors.

A. High-Performance Processors for Space Systems

The landscape of spacecraft and mission development has shifted from large monolithic satellites to SmallSats and CubeSats. As the same time, the space industry has begun to develop systems with a combination of radiation-tolerant and rad-hard devices, using both processors and FPGAs to meet the processing and instrumentation needs of missions. Given their prevalence in the commercial world, hybrid SoCs are of particular interest as they provide the benefits of both fixed-logic processing and reconfigurable logic to enable resiliency and performance while adapting to the radiation environment [11]. The selection of an appropriate architecture for an intended application presents a number of difficulties and is non-trivial. To simplify the selection process, a framework for analyzing the performance of processor architectures based on the computational density (CD) and measured in gig-operations per second was developed by [9]. While the methodology was tested based on current onboard space computing devices at the time of publication, recent research has updated these metrics to include the current state-of-the-art processors, FPGAs, and SoCs [3], [12]. Figure 1 is a compilation of the previous metrics showing the continued exponential increasing trends in performance for radiation-tolerant SoC devices. Microsemi FPGAs prior to the RTG4 were not included in [9], but are still relevant to the comparisons and presented for continuity.

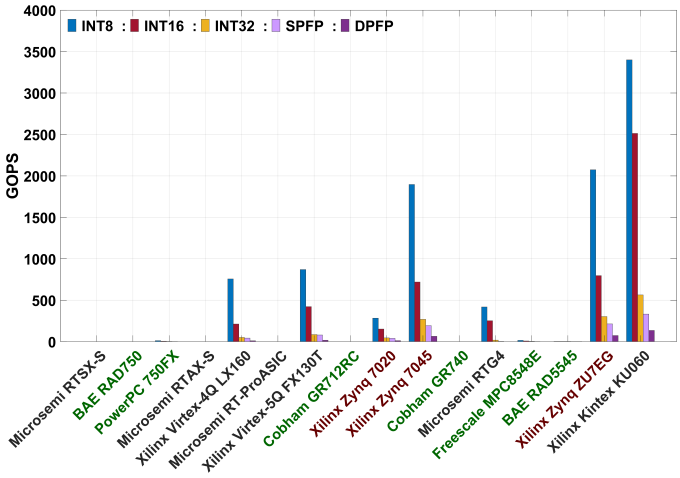

Fig. 1. Performance comparison between rad-hard and radiation-tolerant processors (green), FPGAs (black), and system-on-chip (red) devices based on the giga-operations per second (GOPS) metric compiled from [3], [9], [12]

While the performance increases are evident, the challenges levied to design and manufacture these systems has also increased exponentially. The combination of decreased transistor channel widths and the increase in operating frequencies of current architectures have led to higher power densities with core voltages below 1.0 V and tolerances less than $\pm 3\%$. Whereas processors 20 years ago could operate with core voltage tolerances of ± 125 mV, current state-of-the-art architectures requires regulation to within ± 25 mV [13]. Simultaneously, the maximum core current for FPGAs has increased beyond 30 A with load transients as fast 4 A/ns, and the next generation of processing architectures (e.g. Xilinx Versal) exceeding 165 A [14]. These core currents are composed of both static and dynamic components. The static current represents the continuous current drawn from a device with no activity, whereas dynamic current represents the additional current resulting from design activity. The static and dynamic currents are design dependent, but are typically derived from clock frequency, temperature, process variation, processor load, and FPGA fabric utilization. While the rela-

tionship between the static current, temperature, and transistor leakage is essential to switched converter design and thermal limitations, the dynamic current represents the dominant factor for PDN design. Fast transients will stress a poorly designed decoupling network and can lead to processor brownout and hardware reliability problems (e.g. premature capacitor failure) if not identified prior to flight. Figure 2 shows the core voltage with tolerance (left y-axis) and peak core current (right y-axis) that provides a comparison between processing capability and power for each device in Figure 1. As expected, the exponential increase in computing performance is mirrored by the peak core current.

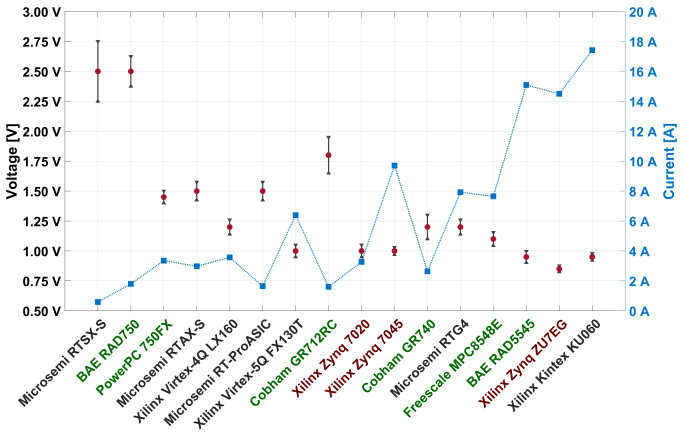

Fig. 2. Rad-hard and radiation-tolerant core voltages and currents for processor (green), FPGA (black), and system-on-chip (red) devices based on 90% fabric utilization for FPGAs and 100% processor load

To account for the variability in power estimation between devices, specific processor and FPGA designs were created to develop a real-world maximum current for each type of computing architecture. To provide a realistic use case, each architecture was configured with a mixture of high-performance non-volatile memory (e.g. DDR) in addition to low-speed (e.g. low-voltage differential signal (LVDS) and general purpose input output (GPIO)) and high throughput interfaces (e.g. multi-gigabit transceivers), where applicable. Additionally, each device was configured for maximum process variation and a 25°C temperature de-rating from the maximum junction temperature reported in the datasheet. For each processor, the core current at 100% processor utilization was estimated based on either the peak power in the datasheet or the processor's power estimator, if available. Since FPGAs require more complex PDNs, manufacturers provide the power estimator tools, which were obtained for each FPGA and SoC. For each FPGA, the power estimator was configured for 90% fabric utilization (i.e. look-up tables, random access memory, and digital signal processing slices) with a toggle rate of 25% and a system clock of 100 MHz. Finally, SoCs were configured with the combination of FPGA and processor settings. While the design configurations defined here would not be compliant with NASA GSFC-STD-1000G Goddard Open Learning Design (G.O.L.D) Rules for mission design, they represent an enveloping worst-case example for effective bounding of target impedance calculations.

B. Radiation Effects

Radiation effects relevant to processing architectures and power systems are broken into three categories: (1) Total Ionizing Dose (TID), (2) Displacement Damage (DD), and (3) Single-Event Effect (SEE). TID is the cumulative long-term damage resulting from ionization primarily caused by low energy protons and electrons. DD is permanent damage caused by the long-term displacement of lattice atoms after collisions with protons or neutrons. SEE are destructive and non-destructive events that occur when heavy ions or high-energy protons pass through a semiconductor and create excess charge carriers. Further categorization of relevant SEEs are provided below.

- Single-Event Transients (SET) are non-destructive temporary variations producing atypical circuit responses.
- Single-Event Upsets (SEU) are non-destructive recoverable events that induce a state change, typically associated with memory bit-flips.
- Single-Event Latch-up (SEL) is a potentially destructive event where a device current is driven out of specification.
- Single-Event Burnout (SEB) is a destructive event where a permanent short forms between the source and drain.
- Single-Event Gate Rupture (SEGR) is a destructive event where a permanent short forms between the gate & drain.

For processor and FPGA architectures, the response to SET, SEU, and SEL differ, but several dependable computing techniques exist for both architectures with a comprehensive overview covered in [15]. Unlike processing architectures, radiation effects in power systems are complicated by the wide range of analog and digital parameters that can be affected [16]. With the complexity related to space system design, radiation-beam testing is typically used to characterize device susceptibility to radiation and provide insight to the device's expected survival for a given orbit. Radiation effects and radiation-beam testing are covered in-depth in [17]. Additionally, NASA Electronics and Packaging Program (NEPP) developed the Radiation Hardness Assurance (RHA) program as a recommended multi-step approach to designing reliable space systems [18]. As a result, electronic part manufacturers now provide RHA data (e.g. screening, manufacturing, and radiation characterization) that simplifies part evaluation. Model-based risk management (e.g. Bayesian net analysis and fault-tree analysis) predict radiation induced effects throughout the mission duration, and is essential to understanding the coupling between hardware, software, and power systems in space systems.

C. Reliability in SmallSat Missions

While the RHA risk classification system provides insight to guide payload development based on national significance, the system does not address design difficulties associated with CubeSat and SmallSat design. Historically, it was accepted that CubeSats were high-risk endeavors; however, their growing

potential utility drove NASA to improve and quantify SmallSat mission risk with a public-private Small Satellite Reliability Initiative (SSRI) [19]. SSRI identified the potential benefits of CubeSat and SmallSat missions, especially as reliability increases. They identified the coupling between SmallSat and CubeSat missions with larger satellite buses for future deep space missions. In their concluding report, SSRI highlighted the trades required to increase mission confidence with respect to cost, schedule, performance, and risk through a consolidated set of tools, resources, best practices, and lessons learned. Designing systems for both SmallSat and higher-class missions is a complicated balance between part performance, screening, and cost. The designer must understand the root cause and failure mechanisms (e.g. part selection, manufacturing, and assembly), and incorporate that knowledge early in the design phase to ensure maximum reliability [20].

The Safety and Mission Assurance (SMA) division at NASA Goddard Space Flight Center is focused on characterizing and understanding risk in order to provide recommendations to missions. Since 2005, they have compiled electronic and electromechanical part failures for Class A to Class D missions that occurred during integration and testing (I&T) or throughout the mission duration. The data in Figure 3 shows each failure categorized by component type and failure mechanism, which provides a number of significant insights. First, integrated circuits represent 53% of the all recorded failures with space pedigree parts accounting for 7% of those failures and a nearly equal failure distribution between commercial, high-reliability, and MIL-spec components. The data indicates that part screening is only a portion of the overall system reliability equation, especially given the mission classes on which these failures were recorded. Second, capacitors and power systems together represent 40% of failures, which is especially concerning given the inherent complexity of power converters and decoupling networks in state-of-the-art processing architectures. Multi-layer ceramic capacitors (MLCC) are especially sensitive to differences in coefficient of thermal expansion, high-temperature manufacturing processes, and thermo-mechanical stresses associated with hand-soldering, all of which increase the probability premature failures [21]. Finally, nearly 1-in-5 failures can be attributed to the assembly process or rework. Optimal component placement is critical, especially with the push to smaller form-factor PCBs and complex PDNs that require hundreds of decoupling capacitors. If priority is not given to component placement early in the design process, ineffective PDNs have a high likelihood, and create reliability problems that require significant re-design work through the addition of decoupling capacitors to ameliorate mediocre designs. Additionally, cracked capacitors may not be identified during I&T, but instead fail on orbit. While the engineers designing space systems never expect to have re-work performed, the tight mission timelines and high costs associated with a PCB re-spin create an environment where board rework is the only option. As such, integrating design for manufacture and assembly (DFMA) practices are imperative to minimizing cost, meeting schedule, and maintaining reliability

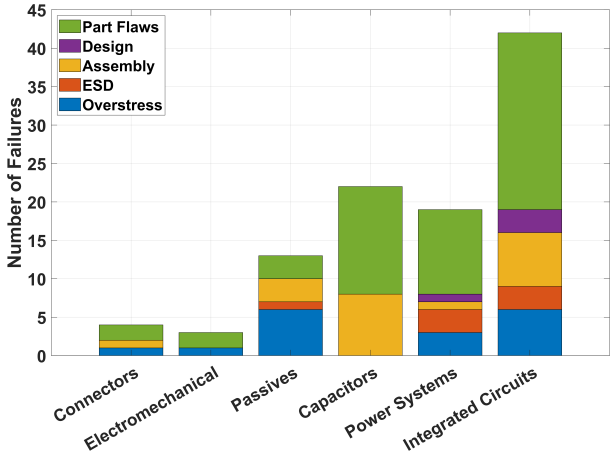

Fig. 3. I&T and on-orbit failures between 2005 and 2017 for Class A to Class D missions at NASA Goddard Space Flight Center

in next-generation space processors. Regardless of the mission class, the design of processing systems must be a holistic approach based on reliability.

D. Space and High-Rel Capacitors

Given the sensitivity of capacitors to external stresses, capacitor selection (e.g. Tantalum and MLCC) is especially important for reliable space systems. While tantalum capacitors provide a larger capacitance per unit area, the relatively high equivalent series resistance (ESR) and equivalent series inductance (ESL) limit their application when compared to MLCCs, which are better able to filter high-frequency transients. Commercial designs commonly use bulk tantalums with a mixture of chemistries, but outgassing and electrolytic field crystallization limit the wide availability and options for space applications. A comprehensive overview of capacitor stability across temperature, bias voltage, package size, and aging is covered in [22]. For avionics processors with low-voltage and high-current transient requirements, both polymer tantalum capacitors and MLCCs are primarily used. Compared to automotive and commercial grade components, additional challenges are introduced with the minimum ceramic thickness between plates required by space qualified capacitors standards (e.g. GSFC S-311, MIL-PRF-123, and MIL-PRF-55365). As a result, a significant reduction in capacitance is observed between automotive-grade and space-qualified capacitors.

Analysis of Table I shows that high-reliability automotive grade capacitors outperform the space rated capacitors by a factor of 4-to-1 for small MLCC 0402 packages and 10-to-1 for large MLCC 1210 packages. For example, each MLCC 0402 capacitor on a commercial processor would require four high-reliability flight capacitors to provide equivalent decoupling performance. Additionally, switched converters required large bulk capacitance with low ESR to minimize ripple, but the lack of MLCC options forces designers to select tantalum capacitors in larger packages with higher ESRs. The overall reduction in capacitance of space-qualified capacitors is especially evident when examining Xilinx's decoupling recommen-

TABLE I
MAXIMUM CAPACITANCE FOR AUTOMOTIVE AND SPACE QUALIFIED CAPACITORS WITH ESTIMATED ESR BASED ON MANUFACTURER DATA

	Automotive Grade					Space Qualified				
Type	*Package*	*Temperature Coefficient*	*Value*	*ESR*	*Voltage*	*Package*	*Temperature Coefficient*	*Value*	*ESR*	*Voltage*
MLCC	0201	X7S	0.10 µF	35 mΩ	10 V	0201	X7R	0.01 µF	75 mΩ	10 V
MLCC	0402	X7R	0.47 µF	15 mΩ	6.3 V	0402	X7R	0.10 µF	40 mΩ	10 V
MLCC	0603	X7R	2.2 µF	7 mΩ	6.3 V	0603	X7R	0.22 µF	30 mΩ	10 V
MLCC	0805	X7R	10 µF	6 mΩ	10 V	0805	X7R	0.47 µF	10 mΩ	10 V
MLCC	1210	X7R	22 µF	3 mΩ	16 V	1210	X7R	2.7 µF	4 mΩ	10 V
MLCC	1210	X7R	47 µF	2 mΩ	6.3 V	1812	X7R	4.7 µF	3 mΩ	10 V
Tantalum	7374	-	680 µF	23 mΩ	2.5 V	7374	-	220 µF	180 mΩ	6 V

dations for the radiation-tolerant Kintex UltraScale FPGA that requires an additional 110 capacitors for the core voltage rail when compared to the commercial recommendation [13], [14]. In addition to requiring more capacitors and PCB area, these space-qualified capacitors are substantially more expensive than their high-reliability counterparts further exacerbating the challenge of space affordability for lower budget SmallSats. Finally, the evaluation of decoupling requirements for current commercial processors from [10] shows a trend toward sub-milliohm target impedances that space-rated capacitors struggle to deliver. Given the variability between mission class, duration, and radiation environment, assessing the system performance, reliability, and cost is essential to the resilient design of next-generation avionics and instrument processors.

III. POWER DISTRIBUTION NETWORK AND RADIATION-BASED TARGET IMPEDANCE

The PDN is responsible for providing a stable low-noise supply from the voltage regulators through package interconnects, power planes, and decoupling capacitors to the die on each IC. Figure 4 shows an equivalent circuit model of a PDN with a simulated impedance waveform. As previously discussed, the complexity of power systems in high-performance processors has increased drastically over the past two decades, forcing designers to simultaneously evaluate and trade computing performance, functionality, and power requirements within the constraints of thermal and mechanical. While these trades do not represent new research for commercial industries, the limitations of space system design present a significant barrier, especially as industry drives towards SmallSat and CubeSat bus architectures that demands high-performance on-board processing. Even with the release of smaller rad-hard voltage regulator module (VRM), the limited area available on a 1U and 3U PCBs is exceptionally restrictive. When combined with the increased number of independent voltage rails and decoupling requirements on modern processors, designers are forced to make stark trade-offs between performance, reliability, and cost. With these increases in system complexity, PDN simulations are a necessity. Given the expense and schedule for complex layouts, simulation speed is often prioritized over model accuracy. Full 3D electromagnetic (EM) simulators provide greatest model accuracy, but the computational complexity and simulation duration limit their application. As such, many simulation tools use hybrid field solvers and equivalent circuit models to approximate the frequency and location dependence of vias, planes, and decoupling capacitors [23].

Larger and faster transient magnitudes place greater stresses on the PDN that manifest as processor brownout, power-plane crosstalk, radiated electromagnetic interference (EMI), SI problems, and premature component failures. VRM selection for high-performance processors determines power efficiency, PDN performance, and decoupling capacitor quantity. While linear VRMs generate low-noise voltage supplies in a small footprint, their low power efficiency limits functionality in an already power-constrained spacecraft environment. Conversely, switched converters are capable of high-current outputs at efficiencies greater than 90%, and able to respond quickly to large current transients induced by state-of-the-art FPGAs. The output voltage is generated by switching transistors at a specified duty cycle introduces noise that must filtered from the power rail. Since the focus of the proposed research relates to core voltages, the use of VRM in this publication refers to switched converters.

The PDN is affected by both DC voltage drop or AC transient responses. The DC voltage drop is proportional to the peak current and the power plane resistance. DC effects

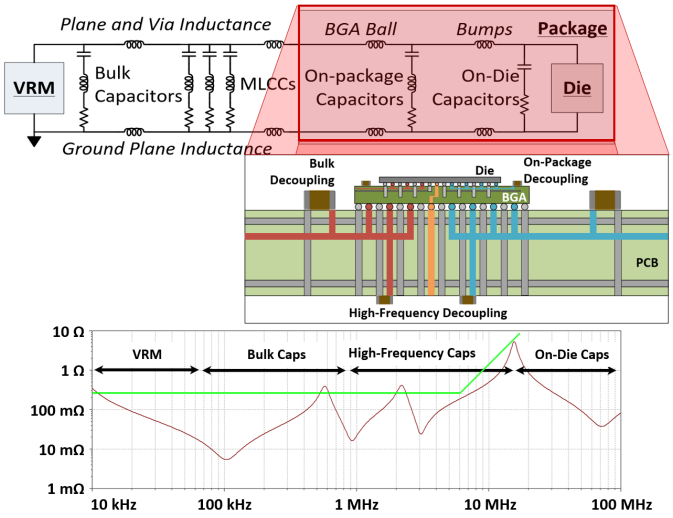

Fig. 4. Equivalent circuit model power distribution network with location and frequency dependent for high-performance processors

can be corrected by placing the VRM sense signal close to the processor to provide a closed-loop feedback mechanism that minimizes the effects of large DC currents and power plane resistance (IR) drop. However, AC transient response is more complicated as it depends on the PCB stackup, via stack, component location, capacitor selection, VRM design, and load-step current parameters. Given the significant coupling between the stackup and the PDN, accounting for stackup effects and component placement when simulating the VRM is essential to identifying problematic designs prior to manufacturing. Additionally, selecting the VRM switching frequency (f_{sw}) for a given topology, generally has the greatest impact on efficiency as it determines the inductor value (L), output capacitance (C_{out}), and voltage ripple (V_{ripple}) that is shown in Equation 1 where I_L is the inductor ripple current and R_{ESR} is the decoupling network's ESR. Traditionally, switching frequency has been limited by the gate capacitance of silicon field effect transistors (FETs). Combined with the inherent radiation tolerance of Gallium Nitride (GaN) transistors, recent advancements have enabled efficient power conversion while switching at frequencies in the megahertz range.

$$V_{ripple} = \frac{I_{ripple}}{(8 * f_{sw} * C_{out})} + R_{ESR} * I_L \tag{1}$$

A number of design methodologies exist for synthesizing PDN impedance profiles including multi-pole and distributed matched bypassing [24]. For voltage rails with low currents, deriving the output capacitance from the voltage ripple in Equation 1 may be sufficient. For higher and faster currents, a charge-based methodology provides a more accurate capacitance estimate [10]. Equation 2 relates the energy transfer between the VRM inductor, the decoupling capacitance, and the load.

$$C_{out} = \frac{Q}{\Delta V_{out}} = \frac{1}{2} \cdot \left[\frac{L_{PoL} \cdot I_{step}}{V_{out}} \right] \cdot \frac{I_{step}}{\Delta V_{out_{AC}}} \tag{2}$$

However, the charge-based estimate does not account for the loop controller bandwidth, and assumes a sufficiently slow current slew-rate such that the capacitor's ESL and ESR is negligible. While these methodologies provide initial estimates, the amount of additional capacitance is determined by modeling in the time (e.g. SPICE) and frequency (e.g. PDN Impedance) domains. Equation 3 shows the target impedance (Z_{target}) that represents the maximum PDN impedance at which the processor's voltage remains within specification for a given load current. The tol_{proc} and tol_{DC} represent the percent tolerance for the processor voltage rail and DC accuracy of the VRM. While the target impedance is constant across frequency, the harmonic content contained in the current transient waveform begins to decay at 40 dB per decade, which depends on the waveform shape. Above this frequency envelope, the low spectral content enables a 20 dB per decade increase in target impedance as indicated in Figure 4.

$$Z_{Target} = \frac{V_{core} * (tol_{core} - tol_{DC})}{I_{load} * \%_{step}} \tag{3}$$

While commercial VRMs have DC tolerances as low as 0.25%, the rad-hard design process introduces significant variability as indicated by DC tolerances upwards of 1.5%. The decreased processor voltage tolerances combined with the high-reliability required by space processors creates a difficult design trade space as the low target impedance requires significant capacitance to operate across the mission lifetime. Given these increased demands, a re-evaluation of the target impedance is proposed in Equation 4, which separates operating temperature (tol_{temp}), load regulation (tol_{load}), and radiation (tol_{TID}). Differences in manufacturing processes lead to increased tolerances at temperature extremes, typically greater at lower temperatures than higher temperatures. Overestimating the operational temperatures can constrains the PDN design which unnecessarily increases decoupling requirements. The proposed target impedance provides a realistic design metric to compare VRMs during the layout study. In addition to TID effects, SETs can induce both fast microsecond-duration and slow millisecond-duration changes to the output voltages that are not taken into account in Equation 4. Accounting for SET transients with target impedance is impractical, however the SET effects can easily be modeled with a SPICE simulator and an instantaneous ground fault of variable duration to determine the maximum SET that can be decoupled. Since the manufacturer radiation data is used to derive the maximum tolerance for RHA parts based on statistical data, no additional radiation analysis should be required.

$$Z_{Target} = \frac{V_{core} * [tol_{core} - tol_{load} - tol_{temp} - tol_{TID}]}{I_{core} * \%_{step-load}} \tag{4}$$

For example, a design team selects an FPGA with a core voltage of 1.0V $\pm$ 3% and worst-case current of 10 A load with a 25% load step, and trades two VRMs in Table II. Comparing the worst-case analysis, the target impedance for VRM_1 is 4 mΩ, and the target impedance for VRM_2 is 8 mΩ. To meet the worst-case target impedance, VRM_2 requires two additional bulk 1210 MLCCs when compared to VRM_1, which increases the required converter area by 35 mm^2. However, the orbit radiation (25 krad) and temperature range (-25°C to 85°C) shows the reference voltage tolerances are comparable leading to a 9.6 mΩ target impedance. This reduction in target impedance when compared to worst-case analysis signifies that less bulk output capacitors are needed, which enables additional processor functionality without loss of performance.

TABLE II
COMPARISON OF REPRESENTATIVE RAD-HARD BUCK CONVERTERS

	Area	Efficiency	tol_{worst}	tol_{rad}	tol_{temp}	tol_{load}
VRM 1	505 mm^2	88%	2.0%	0.1%	0.3%	0.3%
VRM 2	601 mm^2	82%	1.0%	0.2%	0.3%	0.2%

IV. PRINTED CIRCUIT BOARD DESIGN METHODOLOGY

The current state of space hardware and architecture design is decoupled from the layout, manufacturing, and assembly

as designs are outsourced to external vendors for a more turn-key experience. While this process was sufficient a decade ago with large ceramic components, the current state of the space industry has begun a shift to smaller ball-grid array packages with increased adoption of radiation-tolerant devices. The tight coupling between computational performance, PDN design, and software requires a multi-dimensional integrated team that profoundly understands all aspects and trades associated with design decisions. As engineering teams are often short-handed with schedule and budget constraints, the Science Data Processing Branch at NASA Goddard Space Flight Center has developed a design methodology that incorporates many complexities associated with space processor development and PCB design, shown in Figure 5. The foundation of the design process is decision making based on the simulation data, manufacturer recommendations, and experience in previous projects. While many signal and power integrity simulators exist (e.g. Siemens HyperLynx, Keysight ADS, and ANSYS SIWave), the cost and complexity is often a barrier to adoption but the ability to identify problems before manufacturing is essential to a fast engineering development cycle.

A. Processor Trade and Layout Study

Evaluating processor architectures is non-trivial since a desired performance metric is specific to a given engineering team and will vary depending on intended software application needs. Considering the complexity of the processor analysis, the comparison between processors, FPGAs, and SoCs is left to the reader, however a thorough comparison is presented in [5]. Throughout this analysis phase, the hardware design teams must be in close contact with the software and FPGA groups to receive input on architecture-level decisions that could impact performance, I/O interfaces, and future mission needs. A challenging, but critical part of the effort is determining how to decompose/partition a particular application between the disparate computing elements that informs the design trades for each of the processing elements. The design team must define and balance system functionality and computational performance which is constrained by the following parameters:

- Processor, memory, and input/output (I/O) constraints
- PCB stackup considerations
- Manufacturing and assembly limitations
- Maximum bus power and available voltage rails
- Mechanical mounting and thermal limitations
- Mission reliability requirements and component costs

While some of the listed variables are more flexible, developing a priority and understanding the future impact are especially important when performing the cost-to-benefit analysis of a specific design trade. During the component research process, it is beneficial to have a single document containing all relevant information for available components that include power supply requirements, package dimensions, thermal limitations, and layout constraints. Voltage tolerances, peak currents, and overshoot/undershoot limits can vary drastically even within a computing architecture family. Within the constraints of the PCB area, a layout study represents the first design metric for comparison. The package that a component is housed inside is tightly coupled with thermal performance and functionality (e.g. I/O constraint, performance). While large processing elements consume significant power and have priority in the layout study, even low-power components that are poorly placed during the layout study can quickly exceed maximum temperature ratings. As such, it is recommended that the hardware design team use the manufacturer-provided power estimator combined with processor and FPGA fabric utilization, I/O quantities, data-rates, and estimated junction temperature constraints. If power estimators are unavailable, the worst-case operational limits of the device within the intended environment is also acceptable. However, it should be noted that this could lead to over designing the PCB in situations where target applications are not stressing use cases for the device. Worst-case power estimates have negative impacts on cost and loss of functionality if far outside any realistic use case. Most state-of-the-art architectures require on-board volatile and non-volatile memory that have significant I/O requirements. Reference voltage rails with tight tolerances specific to certain memories (e.g. double data rate (DDR)) and termination schemes (e.g. emitter-coupled logic (ECL)) should be noted as specialized VRMs may be required to support those components. Any additional circuitry (e.g. watchdogs, transceivers) must be documented and evaluated to maintain system reliability and performance requirements.

The state-of-the-art FPGAs and processors have complex power systems requiring many voltage rails, and the limited number and large size of radiation-hardened VRMs presents a significant barrier. Commercial VRMs are available with a multitude of functionality (e.g. wide input voltage, pulse skipping, current sensing) and in a variety of configurations (e.g. integrated FETs, System-on-Module, Controller) that enable the designer to select an optimal configuration based on load current. However, the radiation, temperature, and reliability constraints limit functionality by combining the controller and FETs, which ultimately limit the input voltage range, current output, and efficiency. Given these constraints, a preliminary power system should be designed based on documented output voltages, expected load currents, and sequencing requirements. Compiling these values into a spreadsheet or program enables automatic calculation of currents and powers based on converter type and efficiency. The spreadsheet enables the identification of low-current voltage supplies that may be powered by linear VRMs, voltage rails that can be shared between multiple components, and identify high current rails that may require parallel synchronized switched VRMs. Given the overall complexity of the power system, a diagram documenting the entire system is recommended to prevent any accidental oversights. Finally, VRM simulations with estimated currents and step-load transients enable verification of bulk decoupling capacitors within voltage tolerance requirements.

Once preliminary components are selected for a given processing architecture, the layout study determines if the proposed components, passives, and connectors fit within the limited board area. A layout study is performed by arranging

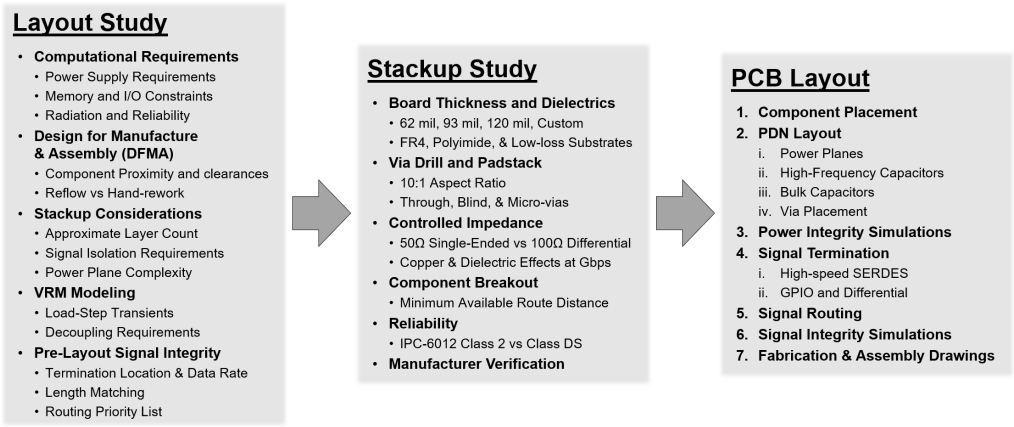

Fig. 5. PCB design and analysis methodology developed at NASA Goddard Space Flight Center and optimized for high-performance processors

shapes that are scaled to the package size of each component around the fixed board area. A variety of applications can be used (e.g. presentation, flowchart, PCB layout), but design for manufacture and assembly (DFMA) practices (e.g. component clearances) must be enforced. With increased complexity, layout studies benefit when they are performed in a PCB layout software that presents the clearest representation of the layout. Component placement should be optimized between the shortest path between devices, routing density, and the thermal performance. Areas with high routing density require greater component separation to balance length matching and crosstalk constraints. While reference designs released by manufacturers are convenient to reference, industry experts have noted they often do not conform to reasonable design practices and should be scrutinized before design decisions are implemented [25], [26]. At this stage, the designer should include sufficient margin based on additional passives and decoupling capacitors that may not be realistic to include in the layout study. Mechanical engineers play an important role throughout the design process, generating the Mechanical Interface Control Document (MICD) prior to the layout study. Additionally, thermal engineers should verify part placement based on estimated powers and packages parameters.

Once the layout study nears completion, a preliminary stackup should be developed based on the quantity of power planes and signal layers. Many methodologies exist for optimizing layer count and high-density interconnect (HDI) breakout in [14]. Pre-layout SI simulations with input/output buffer information specification (IBIS) models in conjunction with the estimated PCB trace lengths from the layout study provide insights into drive strength, slew rate, and termination value and location. All of these parameters have significant impact on the HDI breakout capability. Finally, different processing architectures can be compared based on layout complexity, total cost, power, and any reduction in capability resulting from power system and I/O limitations. Once a layout study is completed and a processor architecture is selected, a routing priority list should be created which is an ordered list that includes the voltage rails, signal interfaces, length matching requirements, data-rates, and trace-to-trace isolation requirements. There are two benefits to categorizing signals and power planes by their importance, intended routing layers, and proximity on the PCB: (1) development of an accurate stackup and (2) sufficient information for the layout engineer. It should be noted that a layout study is only successful after a stackup is approved by a vendor, which can lead to drastically different layout studies that depend on component selection.

B. Stackup Study

The foundation of the PCB is the stackup that is composed of alternating layers of copper foil and stacked dielectric materials. The stackup should be designed in parallel with the layout study as they are tightly coupled. The majority of challenges associated with complex PCB designs are attributed to improper stackup and via stack planning. The stackup study identifies and addresses problems associated with HDI breakout, SI, PI, and EMI prior to layout. Stackup design is a multi-faceted challenge with large solution space that consists of PCB thickness, dielectric material properties, via drill and padstack, reliability, manufacturability, and cost. The layer count estimate from the layout study provides a starting point for the trade space, but is subject to change throughout the layout process as unforeseen constraints arise. Since the layer count is primarily controlled by the signal breakout, each component package must be evaluated based on the pin pitch, minimum via diameter, trace width, and number of traces broken out through each via pair, represented in Figure 6.

Equation 5 provides an analytical representation of Figure 6 that is essential to the design of HDI breakouts. The available routing area drives the controlled impedance trace calculations, which is constrained by the package pin pitch, via drill/pad, clearances, and PCB thickness. For example, a device with a 1 mm pitch, a 24 mil via pad diameter, and a via clearance of 4 mil has a maximum route area of 7.37 mil. Depending on the mission reliability and IPC-6012 manufacturing class, these parameters can change significantly. From the desired impedance value, the maximum route area, a dielectric

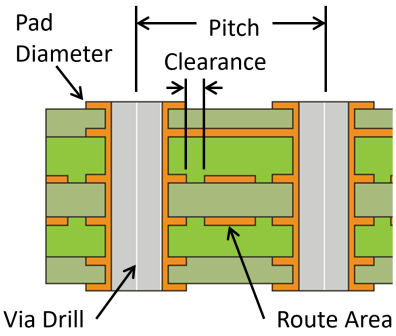

Fig. 6. Breakout definitions for high-density interconnects on PCBs

thickness, and copper weight can be derived from either approximations or 2D solvers that calculate transmission line impedance based on topology. While controlled impedance is essential to operating at high-data rates, trace impedance is often adjusted during HDI breakout to optimize cost with no impact on performance. These adjustments to trace width and impedance should be documented in the routing priority list.

$$L_{area} = L_{pitch} - 2 * \left[\frac{D_{pad}}{2} + L_{clearance} \right] \tag{5}$$

Once the dielectric thicknesses are verified with vendor data, the total board thickness can be calculated based on signal and power plane requirements. While manufacturing capabilities limit the minimum via drill diameter based on the PCB thickness, most manufacturers recommend a 10:1 minimum aspect ratio to ensure the stackup for manufacturability. If the drill diameter to board thickness ratio is less than the required ratio, either the layer count must be reduced or the stackup design should be adjusted for thinner dielectrics. However, the thickness reduction requires a re-calculation of the transmission line impedances and adjustments are not always feasible. By establishing a dialog with vendors early in the design process, the difficulties associated with stackup design (e.g. dielectric material, via padstack, and controlled impedance) and reliability are reduced while maintaining high manufacturing yields at the lowest possible cost.

As the trace geometry area decreases and rise-times drop below 100 picoseconds, copper effects (e.g. surface roughness, skin effect) and dielectric effects (e.g. dielectric losses, fiber weave) become crucial to stackup design. The skin effect is the tendency of current as low as 1 MHz to distribute on the surface of the conductor. Specific to each transmission line topology, the current distribution is non-uniformly distributed towards the closest adjacent reference plane. Discontinuities in the reference plane for both power and signal routes are the cause of many functional problems on PCBs as it leads to ringing on signal lines, radiated fields, and noisy power planes. The combination of the skin effect and surface roughness of copper foils increases the effective trace resistance as the distance traveled by the high-frequency currents is longer. This attenuation is manufacturing process dependent and typically becomes significant above 1 GHz. In addition to conductor effects, multiple dielectric effects can introduce functional problems. While manufacturers provide a single value for the dielectric constant, the dielectric is actually composed of an epoxy resin and fiber weave with different dielectric properties. As trace widths become comparable to a PCB's fiber weave, the propagating signal experiences different impedances and time-of-flights based on the continuously changing dielectric constants that induces reflections and reduces the channel operating margin. While low-loss substrates present a simple solution, the reduction in dielectric constant has significant impact on trace widths, component breakout, and length-matching requirements that may limit a design to a specific material. Additionally, increased costs associated with manufacturing and flight qualification of new low-loss dielectrics must not outweigh the performance benefits, especially with reduced area designs (e.g. 1U, 3U) and when simple solutions exist to mitigate these effects (e.g. zig-zag routing, spread-fabrics). The dielectric's frequency-dependence and stability become especially critical as line rates increase beyond 5 GHz.

As with signal routing consideration, the quantity of power planes in a stackup design is strongly dependent on the processing architecture and layout study. The design of an effective PDN is non-trivial with the performance dominated by two parameters: (1) the mounting inductance (2) spreading inductance. The mounting inductance is the series impedance seen by the capacitor when connecting to planes. As such, power and ground planes placed close to the surface layers and multiple vias placed close to each decoupling capacitor minimizes the mounting inductance. The spreading inductance is the series impedance observed between the decoupling capacitor and the IC pin, and placement of the decoupling capacitor as close to the IC power pin is essential to an effective PDN. While power planes near the surface layers provide improved PDN performance, a symmetric stackup with equal distributions of plane and signal layers on the top half and bottom half of the PCB is preferred by manufacturers as it prevents board warpage during assembly thermal cycles. These complications introduce additional complexity as board thickness, minimum via diameter, plane location, and layer count must be simultaneously evaluated. The decision to select a component without understanding and incorporating the manufacturing costs, often associated with complicated manufacturing practices (e.g. microvias, ¿10:1 aspect ratio, back drilling) that lead to lower yields, introduces significant risk to the project schedule and can lead to exponentially increasing costs. Once a stackup has been designed, manufacturer verification based on the reliability class (e.g. Class 2, Class 3, ES) is required. Small changes to the stackup can have catastrophic effects on the layout that require a complete re-design of the PCB. Given the additional costs associated with Class 3 and IPC-6012ES, a PCB can be designed to the higher reliability standard but manufactured to Class 2 as an affordable solution to prototype development. However, the reverse process is not as simple and will likely require changes to the PCB layout.

C. PCB Layout

The benefit of the proposed design methodology is that once the layout and stackup studies are completed, many of the common problems associated with complex designs have been evaluated and accounted for. By following the routing priority list, the goal of this methodology is prioritize the PDN and minimize plane layout complexity while maintaining processing functionality. As such, layout of high-current and low-noise planes should be prioritized over low-current planes. Since component locations are finalized in the layout study, the next step is the placement of high-frequency decoupling capacitors as close to the component pins as possible, and then bulk decoupling placement near each VRM. When capacitor placement is finalized, the placement of vias to connect decoupling capacitors to the respective planes is next. Optimal PDN performance is obtained when the distance between each decoupling capacitor and the via is minimized. Multiple vias are recommended for bulk capacitors to support the larger currents. After the preliminary power systems layout is complete, PDN simulations should be performed to identify any placement problems and if additional decoupling is needed. A decoupling capacitor mounting inductance report, which most simulators are capable of generating, easily identifies poorly mounted capacitors that would limit PDN performance. After which, impedance simulations for each power pin ensure the target impedance requirements derived during the VRM design are maintained. However, simulators output an impedance for each pad of the IC package, but the target impedance represents a single value required by the die. Since scattering parameters are not always available to model the package, by assuming each power pin has an equal current a first-order approximation is derived by multiplying the target impedance in Equation 4 by the quantity of power pins. Multiple design and simulation iterations are typically required before the PDN impedance is below the target impedance. Additionally, time-domain simulations for current transients ensure coupling between power planes does not affect voltage tolerances.

After the PDN has been optimized, signal termination is placed based on the routing priority list and the pre-layout SI analysis. Since routing and transmission line effects are exacerbated with increasing trace lengths, certain effects (e.g. crosstalk, termination location) can be ignored if the trace length is kept short. While skin effect and surface roughness are fixed at this point in the design, fiber weave effects can be mitigated with design practices (e.g. zig zag routing, panel rotation). After a preliminary routing of each interface is completed, board level SI simulations should be performed to verify that timing, cross-talk, and overshoot/undershoot do not exceed manufacturer recommended operating conditions for each device. At this stage in the design, SI and PI simulations may reveal additional problems that need to be addressed in the layout. This analysis leads to an iterative design process based on analysis versus rules of thumb. As such, the layout engineer must understand the impact of all parameters on the system design to make educated decisions based on performance and cost. Finally, fabrication and assembly drawings are developed with specifics relating to the stackup, controlled impedance trace widths, allowable tolerances, and any additional information relevant to the design that may not be apparent. While the developed PCB design methodology appears to be an ordered process, the approach is holistic and very iterative. Both SI and PI effects must be simultaneously evaluated within context of their impact to overall system performance as optimizing one facet of a design is often detrimental to another.

V. Case Study: SpaceCube v3.0 Mini

SpaceCube processors are a collection of FPGA-based onboard data processing systems developed in the Science Data Processing Branch at NASA Goddard Space Flight Center (GSFC). The goal of the SpaceCube program is to provide substantial increases to onboard computing while maintaining reliability and lowering relative power consumption and cost. In response to mission needs, the hybrid-processing design approach combines radiation-hardened and commercial components to form a novel architecture that combines the best capabilities of CPUs, DSPs, and FPGAs. As part of the third generation of processing architectures, the SpaceCube v3.0 Mini (SCv3M) was developed as a SWaP-C optimized 1U processor to the CubeSat card specification ($[CS]^2$) for the mission enabling capabilities and reusable box configurations [7]. The foundation of the SCv3M is the Xilinx Kintex UltraScale FPGA that provides orders of magnitude performance improvements when compared to rad-hard processors. Xilinx is committed to supporting the space-grade Kintex UltraScale (XQRKU060) in a column grid array (CGA) package with QML-V screening that features the same die as the industrial-grade part in a ball grid array (BGA) package [27]. While Xilinx designed the CGA package to be roughly pin compatible with the BGA package, significant differences between the impedance of the ball and column introduce difficulty when designing for a selective population scheme. The increased impedance of the CGA pins combined with the large currents induce greater voltage drop and require on-die voltage sensing for the VRM. Additionally, Xilinx adjusted the core voltage tolerance from 3% on the BGA to 4% and increased the decoupling capacitor recommendation by a factor of 1.7×. When compared to the BGA package, the increase requires an additional 100 space-qualified capacitors that consume approximately 1,000 mm^2 on the PCB. However, the core voltage pins under the part only represent 136 mm^2. Even with ideal placement, many decoupling capacitors will have marginal effect on the PDN while severely limiting processor functionality. Given the size constraints, the PDN was designed for the industrial-grade BGA package, while being able to accept the space-qualified package with a potential performance derating.

From a trade study early in the design process, the Microsemi RT ProASIC3 was selected as a radiation-hardened monitor (RHM) for programming, system monitoring, memory scrubbing, and radiation mitigation. The RT ProASIC3 provided the greatest functionality in a small package without

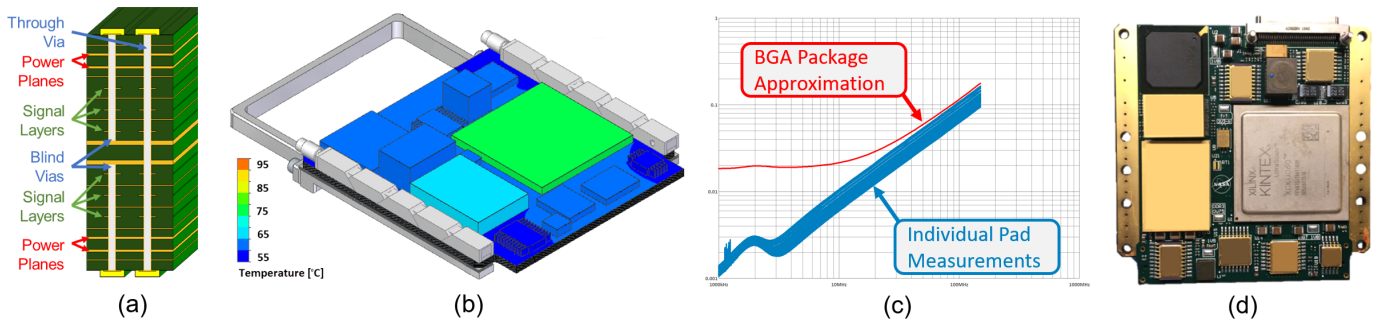

Fig. 7. (a) 22-layer stackup manufactured to IPC-6012DS from HyperLynx (b) Component placement and preliminary thermal analysis based on 15W case (c) Core voltage power integrity analysis at PCB vs BGA (d) Primary side of assembled SCv3M

the need for additional VRMs. For non-volatile memory, the Kintex and RHM have independent 16 GB NAND flash memory modules that store OS boot images and intermediate application data products. To support high-throughput instrument data buffering and softcore microprocessors, the SCv3M design includes 2 GB of 72-bit wide DDR3 SDRAM operating at 1,333 Mega Transfers per second. For external connectivity, the mass required to implement the large I/O count with connectors and cabling is prohibitive, whereas the backplane design provides a SWaP-C optimized solution. The SCv3M was designed with a 400-pin Samtec backplane connector to break out a multitude of LVDS, single-ended 3.3V GPIO, and multi-gigabit transceiver lanes. Additionally, an 85-pin Nano-D front-panel connector was incorporated that supports CameraLink and low-speed single-ended 3.3V interfaces. The 22-layer stackup in Figure 7 (a) was designed to the IPC-6012DS space and military avionics standard, and has 8 mil blind and 12 mil through vias that are essential to achieving high component density without sacrificing SI/PI performance. The stackup has four 1-ounce power planes split between the primary and secondary layers that are enclosed between grounds to minimize mounting inductance. Plane mirroring was selectively implemented to minimize DC voltage drop. Finally, the isolated signal layers enable routing without concern for adjacent layer crosstalk. Extensive pre- and post-layout simulations in Siemens HyperLynx were analyzed to characterize SI performance of the SCv3M. The drive strength, termination location, and termination values were optimized for each interface to balance high-data rates and provide sufficient margin to the recommended operating conditions. Post-layout crosstalk simulations with manufacturer-provided IBIS models showed an average crosstalk of 21 mV and a peak crosstalk of 115 mV across all interfaces at peak rates. Due to the PDN complexity, an Excel document was developed to track the combinations of target impedances, temperatures, and load currents. DC drop, AC impedance in Figure 7 (c), and current density simulations were performed for each power plane. The core voltage rail was shown to meet the required 3.6 mΩ target impedance based on the maximum converter tolerance.

During the course of the layout process, our design team documented four major issues during layout that are presented below. First, the high component density proved especially difficult as the initial layout studies in Microsoft Visio had insufficient component clearances. Second, the potential for high-power dissipation required placement of the Kintex FPGA near the card edge as determined by board level thermal analysis in Figure 7 (b). The placement became a limitation later in the layout as it became difficult to breakout and length match certain interfaces. While more time was spent in layout than expected, the potential for high power dissipation outweighed the extended layout times. Third, the complexity in routing the DDR3 interface drove placement of the DDR3 as far as possible from the Kintex to provide sufficient area for length matching. This extra distance proved beneficial as tight matching was easily achieved on the 72-bit interface. Fourth, a differential pair is unable to breakout between two through vias with a 1 mm pitch, but was not identified until after the component placement and routing priority list was developed. The design team worked with the manufacturer, and was able to develop a solution to decrease via clearance by removing non-functional pads without a stackup redesign.

When the assembled card arrived, a custom suite of automated tests developed by the embedded processing group were performed to verify both the functionality and performance of each interface on the SCv3M. Nearly all of the automated testing passed without problems, however the testing did identify three DDR3 data byte groups had failed calibration and were non-operational. Since the layout passed the PCB software design rule checks and the time-of-flight requirements for each data byte group were verified with HyperLynx DDRx analysis, the issue was not easily identifiable. Additionally, the data lines were routed with blind vias making oscilloscope probing of the problematic data byte groups during operation impossible. After much FPGA reconfiguration and software testing, the problem was identified as multiple signals shorted to ground. Further examination of the layout found that a trace had been shifted, creating a short between ground and three data lines, but was masked by a setting in the layout design rules. The DDRx simulator did not identify the shorted traces either as the removal of non-functional pads created a scenario where the simulator used the via pads to check for

shorts but does not check for shorts to drilled holes. With all other interfaces passing functional and performance tests, custom FPGA cores were developed to specifically test the PDN of the core voltage rail with large FPGA designs. Core dynamic currents of 12.3 A were tested with corresponding core voltage measurements taken to identify potential core voltage violations, and no PI issues were identified. As such, the relative success of the SCv3M design is attributed to design and analysis methodology presented in Section IV.

VI. Conclusions

The state-of-the-art commercial processing architectures require tight tolerance on voltage rails that feed large current transients inside high-density packages that breakout hundreds of IO to various interfaces. Conversely, traditional space systems with large PCB form-factors and high tolerance power systems have created an environment where strict design methodologies were not required for avionics and instrument processors. However, the next generation of systems demand a holistic design and analysis methodology. In this publication, we have reviewed the state-of-art for flight processors, the difficulty associated with powering these devices, and described the impact of the power system has on reliability. Classic decoupling strategies were presented including the ripple method, the charge-based approach, and the target impedance. We present the design methodology that was developed and tested in the Science Data Processing Branch at NASA GSFC. While the methodology is generic to optimizing all space-flight processors, the methodology is especially enabling to small form-factor 1U architectures, which must combine state-of-the-art FPGAs and processors that push the bounds of current PDN technologies. Given the variability in mission classes, a novel de-rated target impedance calculation was derived based on the operating temperature, load and line regulation, and radiation environment that enables designing power systems for the intended mission. Finally, the PCB design methodology was applied to the SpaceCube v3.0 Mini high-performance processor developed at NASA Goddard Space Flight Center. A target impedance of 3.6 mΩ was achieved with capacitor constraints. Trades between stackup and PDN were presented as it pertains to the layout and design.

VII. Acknowledgment

The authors would like to recognize the contributions and support of team members in the Science Data Processing Branch including Cody Brewer, Christopher Wilson, David Wilson, Sebastian Sabogal, Travis Wise, and David Petrick in addition to Thomas Cook from the University of Pittsburgh.

References

[1] National Academies of Sciences, Engineering, and Medicine, *Thriving on Our Changing Planet: A Decadal Strategy for Earth Observation from Space: An Overview for Decision Makers and the Public*. Washington, DC: National Academies Press, 2019.

[2] R. Doyle, R. Some, W. Powell, G. Mounce, M. Goforth, S. Horan, and M. Lowry, "High Performance Spaceflight Computing; Next-Generation Space Processor: A Joint Investment of NASA and AFRL," in *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (iSAIRAS 2014)*, Montreal, Canada, June, 2014.

[3] A. Geist, C. Brewer, M. Davis, N. Franconi, S. Heyward, and T. Wise, "SpaceCube v3.0 NASA Next-Generation High-Performance Processor for Science Applications," in *33rd Annual AIAA/USU Conference on Small Satellites*, ser. Advanced Technologies II, Logan, UT, August 3-8, 2019.

[4] S. Sabogal, P. Gauvin, B. Shea, D. Sabogal, A. Gillette, and C. Wilson, "SSIVP: Spacecraft Supercomputing Experiment for STP-H6," in *31st Annual AIAA/USU Conference on Small Satellites*, ser. Next on the Pad 2, Logan, UT, August 10-14, 2017.

[5] A. D. George and C. M. Wilson, "Onboard Processing With Hybrid and Reconfigurable Computing on Small Satellites," *Proceedings of the IEEE*, vol. 106, no. 3, pp. 458–470, 2018.

[6] D. Rudolph, C. Wilson, J. Stewart, P. Gauvin, A. D. George, H. Lam, G. Crum, M. Wirthlin, A. Wilson, and A. Stoddard, "CSP: A Multifaceted Hybrid Architecture for Space Computing," in *28th Annual AIAA/USU Conference on Small Satellites*, ser. Advanced Technologies I, Logan, UT, August 3-8, 2019.

[7] C. Brewer, N. Franconi, R. Ripley, A. Geist, T. Wise, and S. Sabogal, "NASA SpaceCube Intelligent Multi-Purpose System for Enabling Remote Sensing, Communication, and Navigation in Mission Architectures," in *34th Annual AIAA/USU Conference on Small Satellites*, ser. Advanced Technologies I, Logan, UT, August 7-12, 2020.

[8] J. Leitner, B. Sood, E. Isaac, J. Shue, N. Lindsey, and J. Plante, "Risk-based safety and mission assurance: Approach and experiences in practice," *Quality Engineering*, vol. 30, no. 4, pp. 648–662, 2018.

[9] 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.

[10] C. Parisi, "Voltage Regulator Design and Optimization for High-Current, Fast-Slew Rate Load Transients," in *2020 Texas Instruments Power Supply Design Seminar*. Texas Instruments, February 10, 2020.

[11] S. Sabogal, A. George, and C. Wilson, "Reconfigurable Framework for Environmentally Adaptive Resilience in Hybrid Space Systems," *ACM Trans. Reconfigurable Technol. Syst.*, vol. 13, no. 3, September 2020. [Online]. Available: https://doi.org/10.1145/3398380

[12] E. Gretok, "Comparative Benchmarking Analysis of Next-Generation Space Processors," Masters Thesis, University of Pittsburgh, 2019.

[13] *DS925: Zynq UltraScale+ MPSoC Data Sheet: DC and AC Switching Characteristics*, Xilinx, 2020, rev. 1.18.

[14] *UG583: UltraScale Architecture PCB Design*, Xilinx, 2021, rev. 1.2.

[15] I. Koren and C. M. Krishna, *Fault-Tolerant Systems, 2nd Edition*. San Francisco, CA, USA: Morgan Kaufmann, 2020.

[16] P. C. Adell and L. Z. Scheick, "Radiation effects in power systems: A review," *IEEE Transactions on Nuclear Science*, vol. 60, no. 3, pp. 1929–1952, 2013.

[17] National Academies of Sciences, Engineering, and Medicine, *Testing at the Speed of Light: The State of U.S. Electronic Parts Space Radiation Test Infrastructure*. Washington, DC: National Academies Press, 2018.

[18] M. J. Campola, "Radiation Hardness Assurance (RHA): Overview for Mission Success," in *NASA Electronic Parts and Packaging (NEPP) Program Electronics Technology Workshop (ETW)*, Greenbelt, MD, June 15-18, 2020.

[19] M. Johnson, "SmallSat Reliability:Issues and Solution Approaches," in *Pacific Operational Science and Technology Conference*, Honolulu, HI, USA, March 6-10, 2017.

[20] B. Sood, "Achieving Improved Reliability with Failure Analysis," in *IPC High Reliability Forum*, Linthicum, MD, May 15, 2018.

[21] A. Teverovsky, "Cracking Problems in Low-Voltage Chip Ceramic Capacitors," NASA Electronic Parts and Packaging Program, Greenbelt, MD, Tech. Rep., 2018.

[22] R. Deshpande, *Capacitors*. McGraw-Hill Education, 2015.

[23] E. Bogatin and L. Smith, *Principles of Power Integrity for PDN Design Simplified*. Prentice Hall, 2017.

[24] I. Novak, "Comparison of Power Distribution Network Design Methods: Bypass Capacitor Selection Based on Time Domain and Frequency Domain Performances," in *DesignCon 2006*, Santa Clara, CA, USA, February 6-9, 2006.

[25] E. Bogatin, "Keynote Speaker: Breaking Bad: A Downside of Open Source Designs," in *Altium Live 2019 Conference*, San Diego, CA, USA, October 12, 2019.

[26] R. Hartley, "Presentation: How to Achieve Proper Grounding," November 11, 2019, altium Expert Live Training.

[27] *DS882: Radiation Tolerant Kintex UltraScale XQRKU060 FPGA Data Sheet*, Xilinx, 2020, rev. 1.2.