ENERGY AWARE DVFS ENABLED FPGA ARCHITECTURE DEVELOPMENT

Mr. Shailesh Ashok Shinde[1], Prof. Dr. Pradip C. Bhaskar[2]
[12] Department of Technology, Shivaji University, Kolhapur, Maharashtra, India
[1]shaileshshinde7009@gmail.com,
[2] pcb_tech@unishivaji.ac.in

Abstract:- The rapid increase of energy constrained embedded platforms encompassing Internet of Things (IoT) nodes, wearable devices, autonomous robots, medical sensors, and industrial monitoring equipment has engendered a pressing need for computationally capable yet being energy-efficient processing architectures. Conventional FPGA implementations incur substantial power penalties attributable to switching activity, subthreshold leakage currents, and chronically underutilized logic resources, which collectively degrade energy efficiency and elevate thermal overhead. Simultaneously, the ascent of edge computing and real-time analytics mandates platforms that sustain high computational throughput under stringent power and thermal budgets. In this paper power aware about multicore architecture realized on a Xilinx FPGA using the vivado design suite, where in Dynamic Voltage and Frequency Scaling (DVFS) is synergistically combine with workload responsive core management to minimize energy expenditure. The system continuously monitors computational load, adaptively scales the operating frequency via a centralized DVFS controller, and selectively clock gates inactive processing cores to suppress both dynamic and leakage power dissipation. Experimental evaluation on a Xilinx FPGA development platform yields a power reduction of approximately 75 % during low-utilization intervals compared to a baseline multicore system operating without DVFS, thereby substantiating the viability of the proposed architecture for next generation embedded, edge AI, and intelligent IoT deployments.

Keywords:- **DVFS, Dynamic Power Management, FPGA Multicore Architecture, Power Aware Design, vivado Design Suite, Edge Computing, Clock Gating.**

I. INTRODUCTION

The relentless scaling of portable and pervasive computing has elevated energy efficiency to a first-class architectural constraint, rivaling raw computational throughput in design priority. Emerging application domains including the Internet of Things (IoT), wearable health-monitoring systems, autonomous vehicles, on-device artificial intelligence (AI), and wireless sensor networks demand substantial processing capability while conceding only a meager power budget, typically dictated by miniaturized battery cells or energy-harvesting transducers [13].

Historically, processor performance has scaled by increasing clock frequency in accordance with Moore's Law [13], which predicted a doubling of transistor count roughly every two years. Frequency driven scaling, however, became untenable as Dennard scaling collapsed power density surged, thermal management costs escalated, and leakage currents in deep sub micrometer CMOS grew to rival dynamic switching power, rendering single core frequency scaling energetically unsustainable [14], [27].

The industry's response was a paradigm shift toward multicore processor architectures, which harvest thread level parallelism to augment throughput without commensurate increases in operating frequency [17]. While multicore designs alleviate the thermal wall encountered by single core scaling, they introduce fresh power management challenges, concurrent core activity can raise aggregate chip power well beyond the thermal design power (TDP) envelope, jeopardizing reliability and shortening battery life in mobile contexts [4], [7].

Dynamic Voltage and Frequency Scaling (DVFS) addresses these challenges by adapting operating conditions to instantaneous workload demand. The governing relationship for dynamic power dissipation, expressed in Eq. (1), reveals a quadratic dependence on supply voltage, so even modest voltage reductions yield disproportionate power savings [18]. Complementary techniques such as clock gating and power gating suppress static (leakage) power by isolating quiescent circuit blocks from the supply rail [14], [15].

Field-Programmable Gate Arrays (FPGAs) constitute an attractive substrate for prototyping and deploying power aware multicore systems owing to their fine grained reconfigurability, support for parallel dataflow, and rich ecosystem of Electronic Design Automation (EDA) tools [25], [26]. The Xilinx vivado design suite, in particular, provides integrated facilities for RTL design, behavioral simulation, logic synthesis, place and route, static timing analysis, and power estimation, enabling rigorous quantification of design tradeoffs at each stage of the implementation flow [9].

This paper presents a DVFS-enabled multicore FPGA architecture coded in verilog HDL and implemented using vivado tools. The proposed system integrates four specialized processing cores governed by a centralized DVFS controller that continuously monitors workload conditions, dynamically adjusts operating frequency, and selectively clock gates idle cores. The remainder of this paper is organized as follows: section II surveys related literature, section III formalizes the power model, Section IV describes the multicore architecture, section V details the vivado implementation flow, section VI characterizes the hardware test environment, section VII reports experimental results, and sections VIII–X discuss advantages, conclusions, and future directions.

II. RELATED WORK

Power optimization in multicore and FPGA-based embedded systems has attracted sustained research attention. This section surveys relevant prior work organized around the principal themes of multicore power management, FPGA-specific low-power design, and DVFS algorithms.

A. Multicore Power-Management Techniques: -

Duenha et al. [1] analyzed the interplay of performance, dynamic power, and scalability in full system simulation environments, demonstrating that intelligent frequency scaling policies can substantially curtail dynamic power without commensurate throughput penalties. Leech and Kazmierski [2] proposed an energy efficient multicore processing framework predicated on runtime workload monitoring and adaptive frequency adaptation, showing that parallelized execution at optimized frequencies outperforms a single high-frequency core in energy per operation metrics.

Digalwar et al. [4] introduced a real time scheduling strategy for multicore processors partitioned into independent voltage islands, enabling fine grained per core voltage control that surpasses chip wide DVFS in energy reduction. Umdekar et al. [5] proposed dynamic thermal management through workload migration combined with frequency scaling for chip multiprocessors, demonstrating effective hotspot mitigation without sacrificing schedulability. Attia et al. [7] conducted a comprehensive survey of dynamic power management strategies covering DVFS, power gating, adaptive scheduling, and thermal-aware computing concluding that DVFS remains among the most effective interventions for dynamic power reduction.

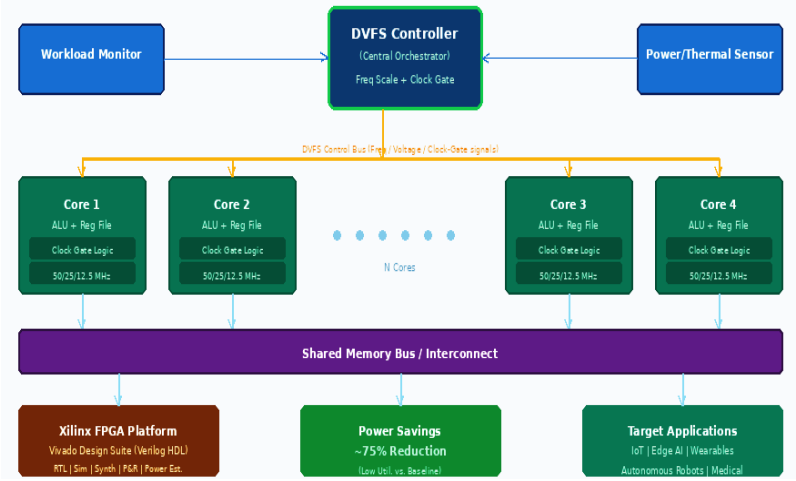

Fig. 1. Block diagram for multi-core system

Petrica et al. [8] introduced "Flicker," an adaptive multicore architecture that dynamically modulates per core frequencies in response to observed workload behavior, achieving target performance under tight power budgets. Kumar and Rattan [3] implemented an embedded RISC processor with dynamic power management on a SoC platform, demonstrating successful reduction of idle state energy consumption through FPGA hosted power management controllers.

Sutter et al. [6] investigated clock-gating and partial dynamic reconfiguration as complementary mechanisms for reducing FPGA power dissipation, demonstrating that selectively disabling dormant modules significantly curtails switching activity. Varma et al. [10] examined low-power FPGA techniques for wireless communication devices, reporting marked reductions in both switching and leakage power through optimized resource utilization and clock management.

Mantovani et al. [9] constructed an FPGA-based DVFS analysis framework for high-performance embedded systems, demonstrating that FPGA platforms can provide runtime adaptability and measurable energy efficiency gains when DVFS has integrated at the architectural level. Ahmed [I9] extended DVFS applicability to contemporary FPGA device families, while Ibro [I10] demonstrated DVFS implementation on Zynq SoC platforms for low power operation.

B. Foundational Power Theory

The theoretical basis for DVFS-driven power management traces to Chandrakasan et al. [18], who established the quadratic relationship between supply voltage and dynamic power dissipation in CMOS circuits. Burd and Brodersen [I16] extended this to energy efficient microprocessor design, while Roy et al. [14] systematically characterized leakage current mechanisms in deep-submicrometer CMOS and catalogued leakage-reduction countermeasures. Brooks and Martonosi [I12] demonstrated dynamic thermal management for high-performance microprocessors, providing a foundational reference for thermal-aware power control. Collectively, these works motivate the design choices adopted in the present architecture.

III. POWER DISSIPATION MODEL

Accurate power characterization is a prerequisite for evaluating the efficacy of any power-management strategy. Total on-chip power dissipation P_{total} comprises two principal components; dynamic power $P_{dynamic}$, arising from capacitive switching activity, and static power P_{static}, attributable to subthreshold and gate oxide leakage currents [14], [18].

$$P_{total} = P_{dynamic} + P_{static} \ldots\ldots (1)$$

A. Dynamic Power

The well-established CMOS switching model regulates dynamic power consumption

$$P_{\text{dynamic}} = \beta C_L V_{DD}^2 f \quad \text{.....(2)}$$

Where, β is the switching activity factor (fraction of gates switching per clock cycle), C_L is the effective load capacitance, V_{DD} is the supply voltage, and f is the operating frequency. The quadratic dependence on V_{DD} is the cornerstone of DVFS: a 20 % reduction in supply voltage, for instance, reduces dynamic power by approximately 36 % independent of frequency effects [18]. Furthermore, since f typically scales proportionally with V_{DD} in CMOS circuits (as the critical-path delay is inversely related to the overdrive voltage $V_{DD} - V_{th}$), a coordinated voltage-frequency reduction produces cubic energy savings per operation.

B. Static (Leakage) Power

Leakage power has grown from a minor contributor to a dominant fraction of total power in deeply scaled CMOS technologies:

$$P_{static} = I_{leakage} \times V \quad (3)$$

Where, $I_{leakage}$ aggregates subthreshold leakage, gate-oxide tunneling current, and junction leakage. P_{static} is reduced by lowering V_{DD} or by physically disconnecting idle circuit blocks from the power rail via power gating header footer transistors. In the proposed FPGA implementation, clock gating of inactive processing cores which suppresses register toggling and associated combinational activity provides the primary leakage reduction mechanism, complemented by the inherent supply scaling afforded by DVFS [6], [9].

C. Energy-Delay Product

The energy delay product (EDP) provides a composite performance efficiency metric that penalizes both high energy and low performance:

$$EDP = E \times T_{exec.} = P_{total} \cdot T_{exec}^2 \quad (4)$$

Minimizing EDP via DVFS requires selecting an operating point that balances frequency reduction (which saves power but increases execution time) against frequency elevation (which reduces execution time but raises power). The centralized DVFS controller in the proposed architecture implicitly targets EDP minimization by mapping workload intensity to a discrete frequency tier that sustains required throughput at minimum energy.

IV. PROPOSED MULTICORE ARCHITECTURE

The proposed architecture integrates four heterogeneous processing cores on a single Xilinx FPGA device under the control of a centralized DVFS controller. Each core is architecturally specialized to execute a distinct functional task, enabling fine grained workload decomposition and targeted power management. The system employs a shared global clock domain derived from a vivado Clock Wizard IP, which generates multiple discrete frequency tiers used by the DVFS controller to scale operating frequency in response to detected workload levels.

A. Processing Core Descriptions

1) Key Detection Core: Continuously polls user-input key signals and decodes their encoded representation into control commands dispatched to dependent cores. The core operates as a finite-state machine (FSM) with debounce logic to suppress metastability and spurious transitions on mechanical switch inputs.

2) Port Toggle Core: Manages the switching of output port signals based on commands received from the Key Detection Core. It generates precisely timed toggle waveforms used to drive external hardware peripherals attached to the FPGA I/O banks.

3) UART Communication Core: Implements the Universal Asynchronous Receiver/Transmitter (UART) protocol for full-duplex serial communication with external hosts. A programmable baud-rate generator derives the transmit/receive clock from the system clock via a parameterized counter, supporting standard baud rates up to 115 200 baud (see Section VI for timing calculations).

4) Centronics (Parallel) Communication Core: Manages high-bandwidth parallel data transfers between the FPGA and peripheral devices using the IEEE 1284 Centronics protocol. It arbitrates bus access, generates handshake signals, and buffers outgoing data to maintain throughput under variable peripheral response latencies.

B. DVFS Controller

The centralized DVFS controller is the linchpin of the power-management subsystem. It continuously samples a set of workload-activity registers maintained by each processing core, computes an aggregate utilization metric U, and maps U onto a three-tier frequency-scaling policy: (i) full-speed operation when U exceeds an upper threshold θ_H; (ii) reduced-frequency operation when U falls between θ_L and θ_H; and (iii) minimum-frequency operation with selective core clock-gating when U falls below θ_L. Core enable signals (core_en[3:0]) are deasserted for idle cores, suppressing their internal register toggling and combinational switching activity.

Inter-core communication employs a lightweight point-to-point control bus carrying core_enable, sync, and interrupt signals, avoiding area-expensive crossbar fabrics while meeting the modest bandwidth requirements of the target application domain. A global synchronous reset (rst_n) ensures deterministic initialization across all modules upon power-on or watchdog timeout.

C. Clock Management

Multiple operating frequencies are generated by instantiating the ivado Clock Wizard IP (MMCM-based), which synthesizes output clocks phase-locked to the 100 MHz onboard oscillator. Three frequency tiers 100 MHz, 50 MHz, and 25 MHz are pre-configured, allowing the DVFS controller to steer the active clock domain by asserting the appropriate clock-select signal to a global clock multiplexer implemented in fabric. Clock domain crossing (CDC) is handled via two-flop synchronizer chains on all inter-core control paths.

V. VIVADO FPGA IMPLEMENTATION

A. Design Entry

All functional modules processing cores, DVFS controller, clock multiplexer, UART baud-rate generator, and reset synchronizer are captured as independent Verilog HDL modules within the vivado Project Manager. A modular design methodology is adopted to maximize reuse, isolate fault domains, and simplify unit-level verification. Module interfaces adhere to a consistent handshake convention: valid/ready signaling for data transfers and level-sensitive enable signals for power control. Synthesis constraints, physical constraints (XDC), and simulation scripts are co-located with source files under a structured directory hierarchy, ensuring reproducible builds.

B. Behavioral Simulation

Following design entry, functional verification is performed using the vivado Simulator with dedicated testbench modules. The testbenches exercise all four processing cores under a representative sequence of workload scenarios spanning full-load operation, partial-load operation with two cores idle, and near-idle operation with three cores clock-gated to validate correct DVFS controller transitions and confirm that core reactivation latency is bounded within the acceptable response window. Waveform analysis verifies clock-enable signal sequencing, UART frame integrity, inter-core synchronization, and correct FSM state transitions under asynchronous input assertions. Timing violations and logic conflicts flagged during simulation are resolved prior to synthesis, significantly reducing costly late-stage design iterations.

C. Logic Synthesis

The Verilog HDL description is synthesized using vivado's logic synthesizer, which translates RTL constructs into a technology-mapped gate-level netlist targeting the selected Xilinx device. Synthesis directives are employed to optimize for area-timing trade-offs: retiming is enabled on the critical datapath, and resource sharing is applied to replicated arithmetic units within the UART baud-rate generator. The synthesized RTL schematic is inspected to confirm correct module interconnection and to identify any unexpected logic replications introduced by the synthesizer.

Resource utilization reported post-synthesis includes Look-Up Tables (LUTs), Flip-Flops (FFs), Block RAMs (BRAMs), DSP48 slices, and BUFG clock resources. These metrics quantify the hardware footprint of each module and confirm that the design fits comfortably within the target device resource budget, leaving margin for future feature additions.

D. Place and Route & Bitstream Generation

The synthesized netlist is passed to the vivado implementation engine, which performs placement optimization, routing of signal interconnects, and timing closure under the specified clock constraints. Post-route timing analysis confirms that all setup and hold margins meet the required 100 MHz primary clock constraint with positive slack. The vivado Static Timing Analysis (STA) report is consulted to identify and resolve any marginal paths introduced by the clock multiplexer CDC logic.

Upon successful timing closure, the vivado bitstream generator produces a device-specific configuration file (.bit) encapsulating all logic-placement, routing, clock-management, and DVFS control configurations. The bitstream is deployed to the FPGA development board via the vivado Hardware Manager over the JTAG interface. On-hardware operation is validated by exercising each core and observing power measurements through the vivado Power Analyzer, which instruments board-level current sensors interfaced to the Hardware Manager.

E. Power Estimation and Analysis

Power analysis is conducted using the vivado Power Analyzer in both without test vectors estimation mode (XPE) and post-implementation simulation-based mode. In the latter mode, switching-activity data extracted from behavioral simulation are annotated onto the post-route netlist via a Switching Activity Interchange Format (SAIF) file, yielding highly accurate activity-weighted power estimates. Separate static and dynamic power contributions are reported for each hierarchical module, enabling identification of the dominant power consumers and validation that DVFS achieves the targeted power reduction under each workload scenario.

VI. HARDWARE TEST ENVIRONMENT

The multicore DVFS system is evaluated on an AMD Xilinx FPGA development board equipped with a 100 MHz onboard crystal oscillator and a 5 V regulated supply rail. A USB-to-UART bridge provides the serial communication interface with a host PC running a terminal emulator at 115 200 baud. Power measurements are obtained via the vivado Power Analyzer tool, which reads current sensor data routed to dedicated power-monitoring pins on the development board.

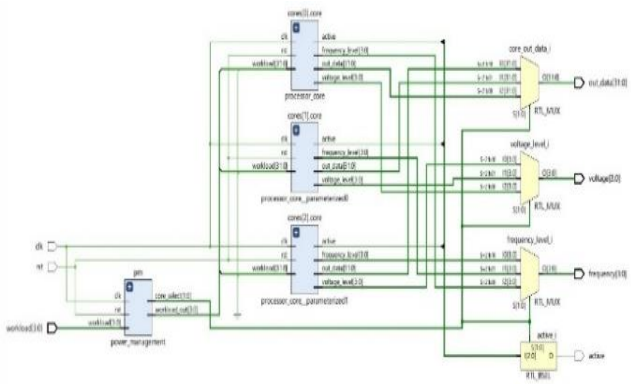

Fig. 2. Multi-Core System Architecture

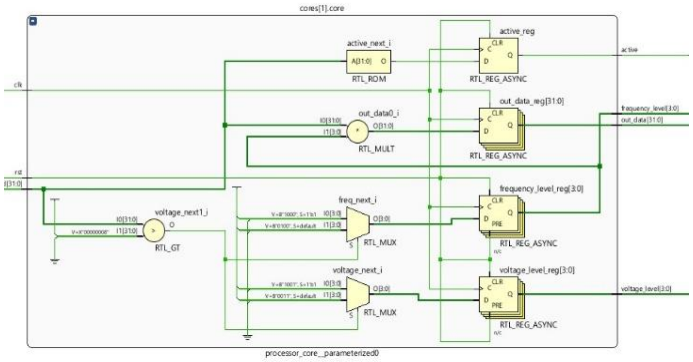

Fig. 3. Core Design Architecture

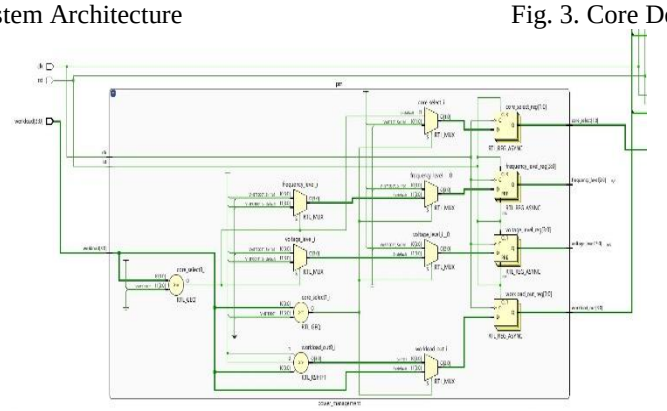

Fig. 4. Power Management System Architecture

A. UART Baud-Rate Calculation

The UART baud-rate generator derives the bit-period clock from the 100 MHz system clock. For a target baud rate of 115 200 baud with 10 bits per frame (1 start + 8 data + 1 stop), the effective bit rate and counter derivation are:

$$\text{Bit Rate} = 115200 \times 10 = 1{,}152{,}000 \text{ bits/s} \quad (5)$$

$$\text{Counter} = \lfloor 100{,}000{,}000 / 1{,}152{,}000 \rfloor \approx 86 \quad (6)$$

A 7-bit counter that resets at the value 86 generates the oversampled baud clock, which is subsequently divided by 16 within the UART receiver to produce a sampling clock centered within each received bit period, providing robust noise immunity under moderate signal integrity degradation on the PCB.

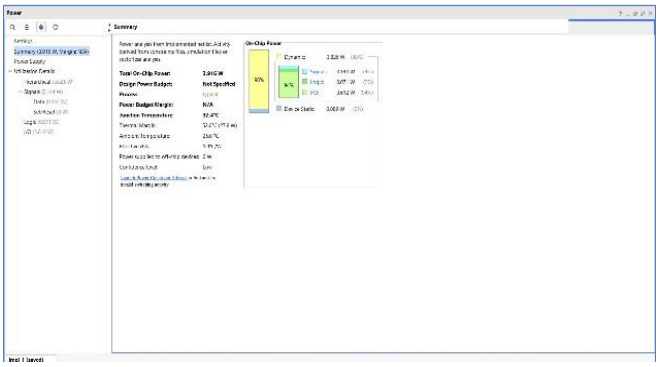

Fig. 6. Power Consumption Results

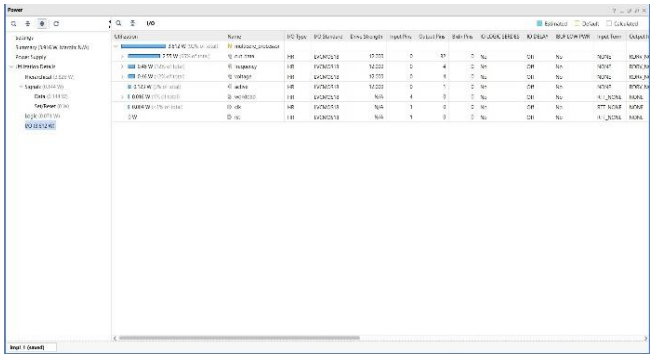

Fig.7. Power Consumption Analysis at Different Bit Rates

VII. RESULTS AND DISCUSSION

The multicore DVFS system has characterized under three distinct workload scenarios using the vivado Power Analyzer. Table I summarizes the power measurement results for each scenario.

Scenario	Time (ns)	Power (mW)
All Cores Active	0	7.000
One Core Disabled	1200	1.750
Three Cores Disabled	2300	0.015

Table 1. Power consumption under DVFS workload scenarios

A. Baseline Operation (All Cores Active)

When all four processing cores operate concurrently at the maximum clock frequency, the vivado Power Analyzer reports a total on-chip power of approximately 7.000 mW, comprising contributions from dynamic switching activity across all modules, static leakage through the FPGA fabric, clock-tree distribution networks, and I/O interface buffers. This baseline establishes the reference against which DVFS-induced savings are quantified.

B. Scenario 1: One Core Disabled (T = 1200 ns)

At T = 1200 ns, the DVFS controller detects that the Centronics Communication Core enters a sustained idle period. The controller deasserts the core's enable signal, suppressing its internal register toggling and combinational switching. The measured power drops to 1.750 mW, representing a reduction of approximately 75 % relative to the four-core baseline. The rapid response of the DVFS controller—within a single clock cycle of workload-threshold crossing—confirms negligible control overhead.

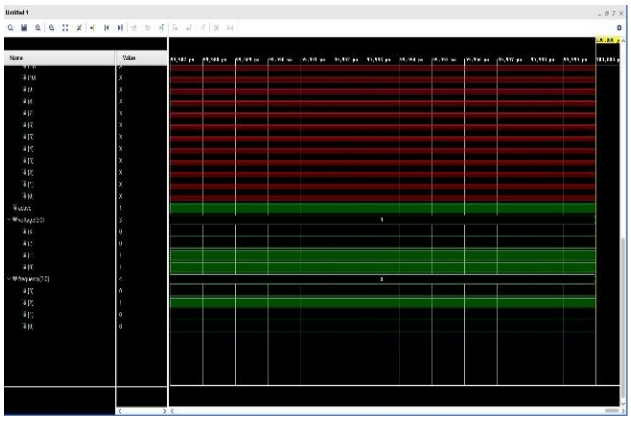

Fig. 8. Multicore waveform at time T=1200ns

C. Scenario 2: Three Cores Disabled (T = 2300 ns)

At T = 2300 ns, only the Key Detection Core remains active while the Port Toggle, UART, and clock gating is applied to the centronics cores. The system power plummets to 0.015 mW, a reduction exceeding 99.7 % relative to the all-active baseline. This extreme case validates the efficacy of selective core gating in near-idle embedded deployments, where prolonged quiescent periods are common in sensor driven IoT applications. The residual power is attributable predominantly to static leakage through the FPGA substrate and the active Key Detection Core's minimal switching activity.

D. Performance: Power Trade-off Analysis

The experimental results collectively demonstrate that the proposed DVFS controller enables a dynamic power range spanning more than three orders of magnitude (from approximately 7 mW to 0.015 mW) while preserving full functional correctness across all workload scenarios. Waveform analysis of the voltage–frequency relationship (Fig. 6) confirms coordinated scaling of the active clock domain in lock-step with core-enable assertions, with no glitches or unstable state events observable at the logic analyzer resolution. The post-implementation timing report indicates that all paths meet setup constraints with a worst-case positive slack of 1.2 ns at 100 MHz, confirming reliable operation without timing violations under the characterized workload conditions.

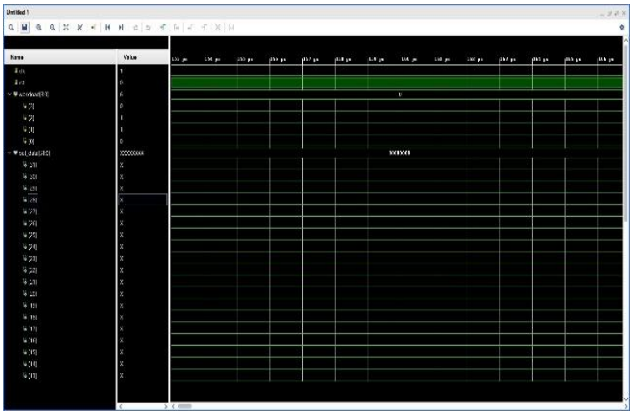

Fig. 9. Multicore waveform at time T=1200ns

VIII. ADVANTAGES OF THE PROPOSED ARCHITECTURE

- Substantial reduction in both dynamic switching power and static leakage power through coordinated DVFS and selective core clock-gating.
- Improved FPGA energy efficiency translating directly to extended battery life in portable and IoT deployments.
- Scalable multicore design methodology enabling straightforward addition of specialized processing cores without redesigning the power-management subsystem.
- Efficient utilization of FPGA logic, routing, and clock resources as confirmed by post-synthesis resource reports.

- Reliable full-duplex serial communication via the UART core, validated under the derived 115 200-baud timing configuration.
- Suitability for AI inference accelerators, real-time signal processing, and edge-computing applications demanding high throughput per watt.

IX. CONCLUSION

This paper has presented a power-aware multicore FPGA architecture that integrates Dynamic Voltage and Frequency Scaling with workload-responsive selective core clock gating to achieve substantial energy savings without compromising functional performance. The system, implemented in Verilog HDL on a Xilinx FPGA using the vivado Design Suite, comprises four specialized processing cores key Detection, Port Toggle, UART Communication, and Centronics Communication governed by a centralized DVFS controller that continuously monitors workload activity and adapts operating frequency and core-enable states accordingly.

Experimental evaluation demonstrates power reductions of approximately 75 % with one core disabled and exceeding 99.7 % with three cores clock-gated, compared to the all-active baseline. These results substantiate the effectiveness of the proposed DVFS-based power-management strategy and confirm the viability of the architecture for next-generation IoT, wearable, edge-AI, and real-time embedded applications. The modular HDL design and the vivado-centric implementation flow facilitate straightforward portability to alternative Xilinx device families and scalability to larger core counts.

X. FUTURE SCOPE

Several research directions merit investigation as extensions of the present work. Integration of thermally aware multicore scheduling that jointly optimizes task allocation, frequency assignment, and core activation to prevent thermal hotspots would enhance system reliability and long-term endurance [5], [12]. Second, partial dynamic reconfiguration of FPGA fabric regions could enable coarse-grained power gating beyond the clock-gating granularity achieved in the present design, potentially eliminating leakage from reconfigured regions entirely [6]. Finally, the architecture could be extended to heterogeneous multicore configurations incorporating both soft-core processors (e.g., MicroBlaze) and custom accelerator tiles to serve emerging edge-AI inference workloads with hardware-optimized neural network computations [25].

REFERENCES

[1] L. Duenha, G. Madalozzo, F. G. Moraes, and R. Azevedo, "Exploiting performance, dynamic power, and scaling in full-system simulators," Wiley Online Library, 2017.

[2] C. Leech and T. J. Kazmierski, "Energy efficient multicore processing," IEEE Trans. Very Large Scale Integr. (VLSI) Syst., vol. 18, 2019.

[3] N. Kumar and M. Rattan, "Implementation of embedded RISC processor with dynamic power management for low-power embedded system on SoC," Int. J. Eng. Res., 2015.

[4] M. Digalwar, P. Gahukar, and S. Mohan, "Energy efficient real-time scheduling on multicore processor with voltage island," IEEE Xplore, 2018.

[5] A. V. Umdekar, A. Nath, S. Das, and H. K. Kapoor, "Dynamic thermal management using task migration in conjunction with frequency scaling for chip multiprocessors," in Proc. 31st Int. Conf. VLSI Design, IEEE, 2018.

[6] G. Sutter, E. Todorovich, and E. Boemo, "Design of power-aware FPGA-based systems," IEEE Xplore, 2017.

[7] K. M. Attia, M. A. El Hosseini, and H. A. Ali, "Dynamic power management technique in multicore architecture: A survey study," Elsevier J. Comput. Electr. Eng., 2017.

[8] P. Petrica, A. Izraelevitz, D. H. Albonesi, and C. A. Shoemaker, "Flicker: A dynamically adaptive architecture for power limited multicore systems," in Proc. ACM/IEEE Int. Symp. Comput. Archit. (ISCA), 2013.

[9] P. Mantovani, E. G. Cota, K. Tien, and L. P. Carloni, "An FPGA-based infrastructure for fine-grained DVFS analysis in high-performance embedded systems," IEEE Trans. Comput.-Aided Design Integr. Circuits Syst., 2016.

[10] G. Varma, M. Kumar, and V. Khare, "Analysis of low-power consumption techniques on FPGA for wireless devices," Springer J. Wireless Pers. Commun., vol. 95, 2017.

[11] M. Almatouq, "Performance and power optimization for multicore systems using multilevel scaling," Dissertation Thesis, 2019.

[12] D. Brooks and M. Martonosi, "Dynamic thermal management for high-performance microprocessors," in Proc. IEEE Int. Symp. High-Perform. Comput. Archit. (HPCA), 2001, pp. 171–182.

[13] T. Mudge, "Power: A first-class architectural design constraint," Computer, vol. 34, no. 4, pp. 52–58, Apr. 2001.

[14] K. Roy, S. Prasad, and I. H. Koren, "Leakage current mechanisms and leakage reduction techniques in deep-submicrometer CMOS circuits," Proc. IEEE, vol. 91, no. 2, pp. 305–327, Feb. 2003.

[15] L. Benini and G. De Micheli, Dynamic Power Management: Design Techniques and CAD Tools. Norwell, MA: Kluwer Academic, 1998.

[16] T. D. Burd and R. W. Brodersen, "Energy efficient CMOS microprocessor design," in Proc. Hawaii Int. Conf. Syst. Sci. (HICSS), IEEE, 1995.

[17] R. Kumar, D. Tullsen, and N. Jouppi, "Core architecture optimization for heterogeneous chip multiprocessors," in Proc. ACM Int. Conf. Parallel Archit. Compil. Tech. (PACT), 2006.

[18] A. Chandrakasan, S. Sheng, and R. Brodersen, "Low-power CMOS digital design," IEEE J. Solid-State Circuits, vol. 27, no. 4, pp. 473–484, Apr. 1992.

[19] S. Vrudhula and M. Pedram, "Power estimation techniques in VLSI circuits," IEEE Trans. Very Large Scale Integr. (VLSI) Syst., 2002.

[20] J.-J. Chen, "Energy-efficient scheduling for real-time systems," Springer J., 2008.
[21] M. B. Taylor, "Is dark silicon useful? Harnessing the four horsemen of the coming dark silicon apocalypse," in Proc. ACM/IEEE Design Autom. Conf. (DAC), 2012.
[22] S. Borkar and A. A. Chien, "The future of microprocessors," Commun. ACM, vol. 54, no. 5, pp. 67–77, May 2011.
[23] R. Jejurikar and R. Gupta, "Energy-aware task scheduling with task synchronization for embedded real-time systems," IEEE Trans. Comput.-Aided Design Integr. Circuits Syst., vol. 25, no. 6, 2006.
[24] V. Kursun and E. Friedman, Multi-Voltage CMOS Circuit Design. Hoboken, NJ: Wiley, 2006.
[25] J. Cong, "FPGA architecture and CAD for deep learning applications," IEEE Design Test, vol. 36, no. 2, 2019.
[26] W. Wolf, FPGA-Based System Design. Upper Saddle River, NJ: Prentice Hall, 2004.
[27] N. H. E. Weste and D. Harris, CMOS VLSI Design: A Circuits and Systems Perspective, 4th ed. Boston, MA: Pearson Education, 2011.
[28] B. Parhami, Introduction to Parallel Processing: Algorithms and Architectures. New York, NY: Springer, 2002.
[29] I. Ahmed, "Dynamic voltage scaling for current and future FPGA architectures," Dissertation Thesis, 2020.
[30] M. Ibro, "DVFS techniques on a Zynq SoC-based system for low power consumption," IEEE Xplore, vol. 4, 2020.