CPU-Power Optimization in Harsh Environments Using Temperature-Aware Scheduling

Nico Rumpeltin*, Tim-Marek Thomas†, Malte Rücker‡, Malte Hawich*, Daniel Lohmann†, Holger Blume*
*Institute of Microelectronic Systems, Leibniz University Hanover, Hanover, Germany
†Institute of Systems Research and Architecture, Leibniz University Hanover, Hanover, Germany
‡Drilling Services, Baker Hughes, Celle, Germany
Accepted Version/Authors copy. Final Version: http://dx.doi.org/10.1109/ASAP71856.2026.00037

***Abstract*—Deep drilling systems must process sensor data directly downhole because the communication channel to the surface is too limited for continuous offloading. This incorporates the processor into the control loop: it executes sensing, signal processing, and actuator tasks locally while operating from a restricted power budget and under strongly varying ambient temperature.**

In the target high-temperature 180 nm SOI technology, static leakage is not the dominant energy term; the main opportunity is therefore to reduce dynamic power by selecting voltage-frequency operating points more carefully. At the same time, circuit timing depends on temperature, so safe and energy-efficient operation requires temperature-aware control. We propose hardware interfaces for a RISC-V processor that expose temperature information, stable timing, and safe operating points to the operating system. The scheduler can then use task slack and ambient temperature to choose the lowest safe operating point at runtime.

***Index Terms*—RISC-V, Harsh Environment, Energy-aware scheduling.**

I. Problem and Motivation

Deep drilling systems must make decisions far below the surface. The communication channel to the surface can provide less than 10 bit/s, which is insufficient to stream raw sensor data or outsource time-critical decisions [1]. Consequently, sensing, signal processing, control, and actuator management have to be executed downhole on a embedded processor itself.

This local processing must operate under a strict power budget. Energy may be supplied by a mud-flow generator during regular operation; during generator downtimes, the system relies on batteries. Reducing processor energy directly increases autonomous operating time and preserves safety margins. The electronics are also exposed to pressure, vibration, and high ambient temperature. Temperatures of up to 175 °C are part of the intended operating range, while temporary spikes can exceed this range [2].

The target platform uses a high-temperature 180 nm SOI technology in which static leakage is comparatively small, whereas dynamic power dominates during signal processing. Therefore for the processor, the relevant optimization target is dynamic rather than static power. Since dynamic power scales quadratically with supply voltage and linearly with frequency, energy can be reduced by executing tasks at the lowest voltage-frequency pair that still satisfies their deadlines. This favors slack-aware *Dynamic Voltage and Frequency Scaling (DVFS)* over a pure race-to-completion strategy [3].

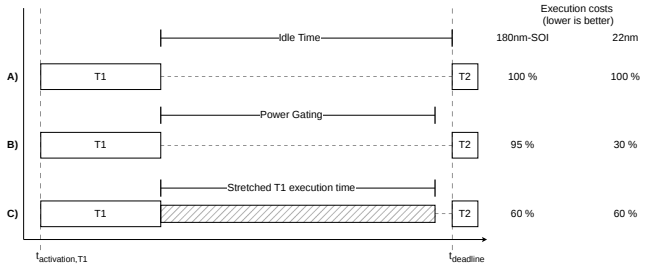

Fig. 1: Illustrative example of how different execution strategies affect the relative energy required to complete the same workload: A) standard execution, B) race-to-completion strategy with power gating, C) deadline stretching with DVFS.

The harsh thermal environment makes this optimization temperature-dependent. The chip temperature follows the ambient temperature closely, and the maximum safe clock frequency for a given voltage changes with temperature. This predictable dependence can be exploited: if the current ambient temperature allows a lower supply voltage for the required frequency, the processor can save dynamic energy without violating timing constraints, see Figure 1.

A practical system therefore needs temperature-aware operating points and an operating system that can use them. Safe voltage-frequency pairs can be obtained from offline characterization or compact timing models [4][5]. Selecting among them requires task-level information, such as deadlines, slack, execution history, and interrupt behavior, which is available only to the operating system. This motivates a hardware-operating-system interface that exposes temperature, stable timing, and controllable operating points, while leaving the energy-management policy to the operating system scheduler.

II. Hardware Considerations

Temperature-aware DVFS depends on operating points that are valid for the actual thermal state of the processor. For the hardware, the main challenge is therefore to define which voltage-frequency pairs are timing-safe under the current thermal conditions, as illustrated for our characterization in Figure 2. A purely nominal operating-point set would either be unsafe at some temperatures or overly conservative across most of the operating range. The operating system can exploit

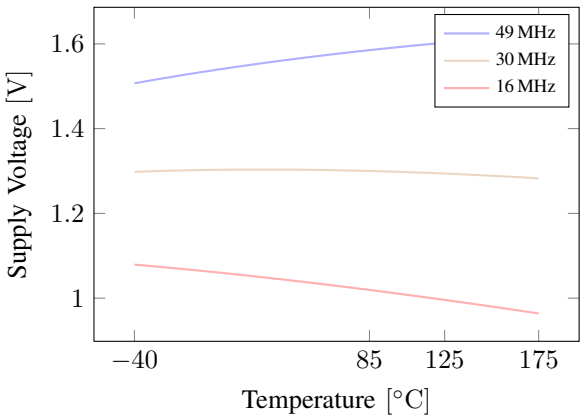

Fig. 2: Supply voltage required to sustain three fixed frequencies over temperature. The voltage can be reduced by up to 0.12 V over the total temperature range [5].

slack only if the safe set is available at runtime; otherwise, lowering voltage risks violating timing constraints.

The required information can be derived from offline characterization or from a compact timing model that maps supply voltage and temperature to the maximum safe clock frequency. Both forms produce a temperature-dependent operating-point description that defines the safe voltage-frequency pairs for the relevant temperature range. This description can be stored as a table or evaluated as a fitted model, depending on the available memory and the desired interpolation granularity. Operating-point validity is therefore a requirement for safe operation, not just an energy-optimization hint.

The hardware modifications needed to enable temperature-aware DVFS are mostly the same as in conventional systems. The processor must expose interfaces to control clock dividers and voltage regulators, provide a frequency-independent time base for scheduling and deadline accounting, and be characterized to operate correctly across the supported operating points. In addition, temperature-aware DVFS requires software-visible temperature information, so that the operating system can select the currently valid operating point. A key prerequisite for this approach is that the ambient temperature changes predictably enough for the selected operating point to remain valid between temperature updates, which is the case in the deep-drilling scenario addressed in this paper.

III. Operating-System Requirements

While the hardware interface defines which operating points can be used safely, it does not determine whether using a lower point is worthwhile. This decision is made by the operating system, which considers task execution, deadlines, slack, interrupts, ambient temperature and the current hardware configuration. Therefore, at each task switch, the scheduler has to make two decisions: which task runs next and which operating point should be used for the next execution interval.

For the task to be dispatched, the operating system estimates the required runtime, compares it with the time until the next activation or deadline, and selects the lowest safe voltage-frequency pair that still allows the task to complete in time. If no slack is available, it falls back to the highest safe operating point. This deliberately reclaims slack: in a technology where dynamic power dominates, stretching execution can save more energy than finishing early and entering a light sleep state, while sleep remains useful for remaining slack.

The operating system must also account for reconfiguration costs, since short tasks may not justify a voltage or frequency transition. For predictable periodic workloads, modest bookkeeping is sufficient: task execution histories are stored in task-local metadata and updated on each context switch. Sporadic interrupts and communication events require additional care; the scheduler can avoid aggressive slowdown for tasks likely to be preempted.

IV. Evaluation Scope

Figure 2 indicates the potential of temperature-aware operating points: for the characterized processor, supply voltage can be reduced by up to 0.12 V, corresponding to an additional dynamic-power reduction of up to approximately 20 % over conventional DVFS. As these values represent the upper bounds of individual operating points, a full evaluation must quantify the potential savings across entire drilling process, taking into account scheduler overheads, reconfiguration costs and workload-dependent slack.

V. Conclusion

Harsh-environment processing requires a hardware-software contract that exposes temperature-aware control to the operating system. The hardware must provide runtime voltage and frequency control, a frequency-independent timer, temperature observability, safe operating-point information, voltage-domain visibility, and low-power or clock-gating mechanisms.

We presented such a co-design for a RISC-V processor in a high-temperature 180 nm SOI technology. By keeping mechanisms in hardware and policy in the operating system, the scheduler can use task slack and ambient temperature to select the lowest safe operating point at runtime. Temperature awareness is therefore not only an optimization, but a prerequisite for safe and energy-efficient operation in harsh environments.

References

[1] T. Stuckenberg, M. Rücker, M. Gottschlich, N. Rother, R. Nowosielski, F. Wiese, and H. Blume, "Design and evaluation of a 180 nm powerline communication ASIC for harsh environment," *Microprocess. Microsystems*, vol. 97, p. 104742, 2023.

[2] M. Hawich, M. Rücker, T. Stuckenberg, and H. Blume, "The HOT-RISC-V: A high-operating-temperature RISC-V architecture for extreme conditions," in *Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS) – Proceedings of the 25th International Conference*, ser. Lecture Notes in Computer Science. Springer, 2025.

[3] H. Aydin, P. Mejía-Alvarez, D. Mossé , and R. G. Melhem, "Dynamic and aggressive scheduling techniques for power-aware real-time systems," in *Proceedings of the 22nd IEEE Real-Time Systems Symposium (RTSS 2001), London, UK, 2-6 December 2001*. IEEE Computer Society, 2001, pp. 95–105.

[4] M. Bao, A. Andrei, P. Eles, and Z. Peng, "On-line thermal aware dynamic voltage scaling for energy optimization with frequency/temperature dependency consideration," in *Proceedings of the 46th Annual Design Automation Conference*, 2009, pp. 490–495.

[5] N. Rumpeltin, T.-M. Thomas, M. Rücker, M. Hawich, D. Lohmann, and H. Blume, “A hardware/operating system co-design approach for energy optimization in harsh environments,” in *2025 Kleinheubach Conference*, Sep. 2025.