

ELP: Elastic Lifetime Processors for Improving Server Energy Efficiency

Freddy Gabbay Jawad Haj-Yahya Firas Ramadan Majd Ganaïem Panagiota Nikolaou Yiannakis Sazeides
The Hebrew University Meta, Switzerland Technion Technion UCLan, Cyprus University of Cyprus

Abstract—Server processors are typically designed with highly conservative aging margins to tolerate worst-case lifetime (7–10 years) and junction temperature (up to 105 °C) assumptions. In practice, however, most servers operate under significantly milder thermal conditions and are decommissioned well before their nominal design lifetime, leaving substantial aging margins unused. This work identifies this over-provisioning as a key source of energy inefficiency and proposes a framework to reclaim excess aging margins to improve server energy efficiency. We introduce *elastic lifetime processors* (ELPs), a design paradigm for reclaiming excess aging margins. Under this paradigm, we develop a first-order, aging-aware margin reclamation model that estimates the timing degradation required to meet a target lifetime and mission profile, and quantifies the spare aging margin available beyond this requirement. The reclaimed margin can be safely reallocated to reduce operating voltage, increase clock frequency, or a combination of both, thereby improving energy–performance trade-offs without compromising reliability. We integrate this design paradigm into an evaluation framework targeting a 16 nm FinFET technology node and apply it to representative datacenter workloads. Our results show that reclaiming excessive aging margins enables significant gains in energy efficiency, delivering up to 20% improvement in queries per joule compared to a worst-case baseline design. These findings demonstrate that aging-margin reclamation is an effective and practical approach for reducing the energy cost of conservative lifetime provisioning in modern server processors.

Index Terms—Reliability, Chip aging, BTI, Lifetime management.

I. INTRODUCTION

In recent decades, the semiconductor industry has advanced VLSI technologies to unprecedented scales through aggressive transistor scaling, three-dimensional integration, and the adoption of new devices and materials. While these innovations have enabled higher performance and improved energy efficiency, they have also increased vulnerability to reliability challenges. One of the most critical issues is the gradual degradation of transistor performance over time, commonly referred to as “transistor aging” [22].

A dominant effect of transistor aging is an increase in threshold voltage, which leads to longer switching delays and degraded performance [23]. The extent of aging strongly depends on an IC’s *mission profile* [45], defined by the combined operating junction temperature and workload intensity over its lifetime. A common design-time mitigation strategy is to add timing margins [28], either by using faster logic elements, increasing the operating voltage (voltage overdrive), or a combination of both [25]. However, faster logic relies

on low-threshold voltage (low- V_t) cells that increase leakage current and energy consumption [34], while voltage overdrive raises both leakage and switching currents, further increasing power and energy consumption [26]. As a result, transistor aging has significant implications for processor design and operation, directly impacting energy efficiency. This challenge is increasingly critical for large-scale datacenters, where energy efficiency strongly affects operating costs, environmental impact, and system reliability [39].

Today, server processors are designed with conservative worst-case aging margins, assuming a maximum lifetime of 7–10 years [35]–[37] and worst-case junction temperatures of 105 °C [33]. This worst-case “aging margin tax” is paid by all processors, regardless of their actual mission profile or duration of use in the field. In practice, most datacenter servers are retired within 3–6 years and operate at temperatures well below 105 °C for most of their lifetime [30], [43]. Consequently, current design practices impose significant power and energy overheads, motivating approaches that both customize aging margins to the expected lifetime and mission profile and reclaim excessive, spare margins to improve processor energy efficiency.

Prior work has improved processor energy efficiency and reliability by adapting voltage, frequency, and resource management to reduce conservative design margins or mitigate aging effects. For example, NBTI-aware DVFS reclaims timing guardbands using aging models to reduce supply voltage while maintaining fixed performance targets [44], while other approaches exploit overclocking, enhanced cooling, or runtime reliability management to operate beyond nominal margins [26], [40], [41]. However, these techniques fundamentally assume fixed-lifetime targets and worst-case operating conditions provisioned by hardware vendors, and therefore do not exploit the substantial mismatch between conservative design-time lifetime assumptions and the expected operational lifetime of servers in practice (e.g., customer depreciation periods or refresh cycles). In contrast, this work explicitly exposes the expected target lifetime and mission profile specified by the customer as configurable parameters and optimizes processor operation by reclaiming over-designed aging margins, enabling a broader exploration of voltage–frequency operating points that more effectively trade off energy efficiency, performance, and longevity.

In this paper, we propose *elastic lifetime processors* (ELPs), a processor design paradigm that improves energy efficiency

by explicitly exposing processor lifetime as a configurable parameter and reclaiming excessive aging margins arising from conservative worst-case design assumptions. By leveraging the expected lifetime and mission profile of server workloads, ELPs enable processor operation to be tailored to customer requirements. Users specify their intended lifetime, workload characteristics, and power–performance preferences to the processor firmware, which estimates the available spare aging margins using our proposed design paradigm and dynamically selects voltage–frequency operating points that optimally trade off energy efficiency, performance, and lifetime.

To evaluate the benefits of ELP, we develop a first-order parametric *elastic lifetime design paradigm* that determines the aging margins required to meet a target processor lifetime and mission profile. The design paradigm comprises: (1) an application programming interface (API) for specifying lifetime and mission constraints; (2) a temperature aggregation model that maps time-varying junction temperatures to an effective temperature; (3) an aging degradation model that estimates delay increase due to transistor aging; and (4) a margin reclamation model that exploits spare margins through voltage reduction, frequency scaling, or both. We integrate this design paradigm into an evaluation framework that explores trade-offs among lifetime, performance, power, frequency, and temperature to improve energy efficiency. The evaluation targets a 16 nm FinFET technology node and is readily extensible to advanced process technologies.

Figure 1 illustrates the ELP exploration space, highlighting the range of feasible voltage–frequency–temperature operating points enabled by relaxing conservative lifetime and mission profile assumptions. The baseline processor operates at nominal voltage (1 V), a maximum temperature of 105 °C, and is provisioned for a 10-year lifetime. In contrast, ELP 1–3 represent dynamically configured instances of the same processor, tailored to customer-specified lifetime and mission profiles. For a 6-year lifetime at an effective temperature of 70 °C, the relaxed lifetime and thermal constraints expose spare aging margin that can be reclaimed either to increase frequency at nominal voltage (ELP 1) or to reduce supply voltage with a smaller increase in frequency (ELP 2). For a shorter 3-year lifetime at 90 °C, similar reclamation enables voltage–frequency operating points that prioritize a combination of performance and energy efficiency (ELP 3).

Through quantitative evaluation, we demonstrate that ELP can reclaim spare aging margins to deliver substantial performance and energy-efficiency gains while remaining compatible with realistic deployment constraints for processor vendors and customers. Experimental analysis for server processors running modern datacenter workloads confirms that ELP significantly improves energy efficiency. This paper makes the following primary contributions:

- 1) We propose *elastic lifetime processors* (ELPs), a processor design paradigm that enables in-field voltage–frequency configuration based on expected lifetime and mission profile to improve energy efficiency.




Fig. 1. IC lifetime, frequency, operating voltage, and temperature trade-offs.

- 2) We develop a parametric elastic lifetime model that estimates effective temperature and required aging margins from a target lifetime and mission profile, and reclaims spare margins for voltage reduction, frequency scaling, or both.
- 3) We evaluate ELP using five representative datacenter workloads, demonstrating up to 20% improvement in queries per joule.

II. BACKGROUND

This section provides background on transistor aging and reviews related work.

A. Transistor Aging

Transistor aging refers to the gradual degradation of transistor characteristics over time. In modern CMOS technologies, aging is primarily caused by bias temperature instability (BTI) and hot carrier injection (HCI) [19], [21]. In this work, we focus on BTI, which has emerged as the dominant aging mechanism in very large-scale integrated circuits [20], [28]. BTI occurs when a sustained gate bias is applied to a transistor over extended periods, leading to the generation of interface traps and oxide defects. This process increases the transistor threshold voltage V_{th} , which in turn degrades circuit performance by slowing transistor switching. The threshold voltage shift caused by BTI is commonly modeled using a reaction–diffusion model, in which the increase in threshold voltage ΔV_{th} is expressed as:

$$\Delta V_{th} = C_T \cdot e^{-E_a/(k_B T)} \cdot t^{1/n}, \quad (1)$$

where C_T is a technology-dependent constant, n is the time exponent, E_a is the activation energy, T is the operating temperature, k_B is the Boltzmann constant, and t denotes the elapsed stress time. Equation (1) highlights the strong dependence of BTI-induced degradation on temperature; for instance, operation at 105 °C results in a threshold voltage shift that is approximately 2.5× larger than at 90 °C.

An increase in V_{th} directly translates into longer gate and path delays. This relationship is captured by the α -power law

delay model [17], which relates the propagation delay t_{pd} to the supply voltage V_{DD} and threshold voltage V_{th} as:

$$t_{pd} \propto \frac{V_{DD}}{(V_{DD} - V_{th})^\alpha}, \quad (2)$$

where $\alpha \approx 1.3$ is the velocity saturation index. As BTI increases V_{th} over time, the effective overdrive voltage ($V_{DD} - V_{th}$) is reduced, leading to a nonlinear increase in propagation delay. Together, (1) and (2) establish a direct link between operating conditions, transistor aging, and performance degradation. This relationship forms the basis for the aging-aware delay modeling and margin analysis developed in Section IV.

Common techniques to mitigate transistor aging include increasing the supply voltage to compensate for delay degradation and tightening clock-cycle constraints during design time. While voltage overdrive reduces BTI-induced delay, it significantly increases both dynamic and static power consumption. Similarly, designing for tighter timing margins requires the use of fast, low- V_t cells, which incur exponentially higher leakage power, particularly at elevated temperatures [34].

B. Related Works

Prior work has improved server energy efficiency by operating beyond conservative design margins via voltage optimization, undervolting, overclocking with enhanced cooling, and runtime characterization [6], [7], [10], [11], [13], and by showing that operational energy dominates datacenter carbon emissions in the absence of renewables [4], [39]. We compare our ELP approach with the most relevant prior work, as summarized in Table I, along with their key contributions and limitations. BRAVO [42] integrates multiple reliability mechanisms and improves a Balanced Reliability Metric by 9%, though it is limited to early design stages and lacks runtime adaptability. The NBTI-aware DVFS technique [44] reclaims static design-time guardbands using predictive aging models to save up to 16% dynamic power and extend processor lifetime by two years, but assumes fixed performance and lifetime constraints. RAMP [26] dynamically adjusts reliability to improve performance by 6–16% under multiple failure mechanisms but assumes fixed temperature. Compared to prior work, ELP—aimed at reclaiming excess BTI margins—differs in several key respects: (1) it enables multi-objective optimization across power, performance, and lifetime; (2) it employs an effective-temperature model validated through aging simulations to capture diverse mission profiles; (3) it explicitly targets transistor-level aging mechanisms; and (4) it supports dynamic adaptation.

III. ELASTIC LIFETIME PROCESSORS DESIGN PARADIGM

ELP aims to reclaim excess BTI-induced aging margins and reallocate them to optimize energy efficiency, performance, and lifetime according to the intended deployment conditions. Hyperscale datacenters typically operate processors at lower temperatures and shorter service lifetimes than worst-case design assumptions [31], resulting in unused aging guardbands. ELP addresses this mismatch by translating user-specified mission profiles and depreciation targets [18], [32]




Fig. 2. ELP Design Paradigm.

into quantifiable reclaimable margins. These margins, which would otherwise remain as unused timing slack, can then be reallocated through voltage–frequency (VF) adaptation: voltage reduction improves energy efficiency, frequency scaling enhances performance, and hybrid policies enable balanced trade-offs. Accurately mapping aging margins to VF settings requires models detailed in Section IV.

To expose these capabilities, ELP provides an API through which users can specify constraints and optimization preferences (e.g., energy vs. performance bias) as illustrated in Fig. 2. ELP then computes appropriate VF settings using vendor-stored models. Similar interfaces already exist in current processors for tuning power-performance trade-offs [38]. ELP configurations and models are stored in the processor’s Vital Product Data (VPD), a non-volatile memory region accessible post-manufacturing. A vendor-supplied driver interfaces with the ELP logic, allowing VF control via standard mechanisms (e.g., Model-Specific Registers) [38]. To ensure reliability, ELP monitors the actual mission profile during operation using on-chip temperature sensors. These readings are aggregated and stored in VPD to reflect real usage over time. Discrepancies between expected and observed profiles are checked periodically (e.g., weekly), enabling two possible responses: (1) hardware autonomously adjusts settings to still meet lifetime goals, or (2) users are alerted for manual reconfiguration. Policy selection is left to future work.

IV. ELASTIC LIFETIME AGING MODEL

This section presents an aging-aware parametric model that enables reclaiming excessive timing margins provisioned at design time under conservative worst-case lifetime and temperature assumptions. By incorporating the target lifetime and mission profile, the model identifies spare aging margin that can be reallocated through voltage and frequency adaptation to improve energy efficiency, enhance performance, or extend operational lifetime. Unlike traditional design approaches that assume fixed operating conditions, our model captures realistic, time-varying thermal behavior and supports dynamic margin reclamation in ELP systems. Figure 3 illustrates the model flow used in ELP operation, which can be configured through the API.

The ELP model consists of the following components:

TABLE I
COMPARISON OF ELP WITH MOST CLOSELY RELATED WORKS.

Method	Reliability Mechanisms	Key Contributions	Limitations
BRAVO [42]	Soft errors, EM, TDDDB, NBTI	1) Improves Balanced Reliability Metric by 9%. 2) Combines multiple reliability mechanisms.	1) Limited to design-time use. 2) Lacks adaptability to diverse mission profiles or lifetime trade-offs.
RAMP [26]	EM, NBTI, TDDDB, thermal cycling	1) Runtime management of processor lifetime under various failure models. 2) Gains 6–16% performance while meeting reliability goals.	1) Uses an oversimplified temperature model. 2) Optimizes only the performance–reliability trade-off.
NBTI-Aware DVFS [44]	NBTI	1) Uses predictive aging to guide DVFS and workload mapping. 2) Saves up to 16% dynamic power, extends lifetime by 2 years.	1) Fixed lifetime and performance targets. 2) No runtime configurability. 3) Limited trade-off flexibility.
ELP (This Work)	NBTI	1) Dynamic lifetime configurability and mission-profile adaptation. 2) Optimizes energy, performance, and lifetime jointly. 3) Up to 20% improvement in queries-per-joule.	Limited to BTI degradation.




Fig. 3. ELP Model

- 1) **Effective Temperature Model (ETM)** – Computes an aggregate effective temperature, T_{eff} , that captures the aging impact of time-varying junction temperatures (Section IV-A).
- 2) **Aging Margins Model (AMM)** – Estimates timing degradation resulting from transistor aging as a function of target lifetime and T_{eff} (Section IV-B).
- 3) **Aging Margins Reclamation Model** – Identifies and reclaims spare aging margins—defined as the difference between conservative vendor-provisioned margins and the required aging margin—and reallocates them to voltage reduction, frequency scaling, or both (Section IV-C).

A. Effective Temperature Model (ETM)

IC vendors typically design processors for worst-case conditions, assuming lifetimes of up to 10 years and junction temperatures reaching 105°C [33], [35]–[37]. In practice, however, datacenter workloads and thermal environments vary significantly, often resulting in more relaxed operating conditions. These operating conditions are commonly captured using a mission profile, a compact representation of the passive and dynamic loads encountered over a device’s lifetime [33], [45]. Table II presents two representative examples of CPU mission profiles derived from hyperscale and high-performance computing (HPC) systems, based on data from recent studies [2], [3]. In the *Hyperscale* profile shows CPUs mostly below 60°C, with active workloads at 50–70°C and passive periods mainly at 40–50°C. The *HPC* profile represents compute-intensive systems with more activity at higher temperatures, especially 60–70°C, and significant passive time at cooler 30–40°C. These patterns illustrate how ELP can

TABLE II
REPRESENTATIVE CPU JUNCTION TEMPERATURE MISSION PROFILES.

T_j (°C)	Hyperscale		HPC	
	Active	Passive	Active	Passive
25	0%	2%	0%	2%
30	0%	0%	0%	38%
40	0%	50%	0%	60%
50	40%	45%	20%	0%
60	30%	3%	50%	0%
70	20%	0%	25%	0%
80	10%	0%	5%	0%
90	0%	0%	0%	0%

exploit deployment-specific thermal behavior to optimize for energy, performance, and lifetime.

Given the variability in datacenter operating conditions, we adopt a temperature aggregation model that has been validated through detailed aging simulations [1] to represent the mission profile’s temperature distribution using an equivalent constant temperature, T_{eff} (a similar model was presented in [45]). This temperature represents the fixed thermal condition that would induce the same BTI-related aging degradation as the actual time-varying profile. The model is formalized as:

$$T_{\text{eff}} = \left\{ \frac{1}{T_t} + \frac{k_B}{E_a} \cdot \ln \sum_{i=1}^n P_i \cdot \exp \left[\frac{E_a}{k_B} \cdot \left(\frac{1}{T_i} - \frac{1}{T_t} \right) \right] \right\}^{-1}. \quad (3)$$

Here, T_t denotes the reference temperature, E_a is the activation energy, k_B is Boltzmann’s constant, and P_i represents the fraction of time the integrated circuit (IC) operates at temperature T_i . Applying this model to the hyperscale and HPC mission profiles in Table II yields effective temperatures of $T_{\text{eff}} = 59.7^\circ\text{C}$ and $T_{\text{eff}} = 63.7^\circ\text{C}$, respectively, both significantly lower than the worst-case temperature of 105°C typically assumed by IC vendors. This gap can be exploited to reduce aging margins, enabling improved power efficiency, performance, and extended lifetime. Datacenter operators can typically predict future mission profiles using offline characterization, QoS requirements, and historical data. However,




Fig. 4. Logical cells delay shift (Δ_{Aging}) during a lifetime at different operating temperatures.

changes may still occur in the field, leading to higher or lower T_{eff} than planned. To address this, we propose *periodic mission profile tracking* and margin adjustment based on the actual T_{eff} , as outlined in Section III.

B. Aging Margins Model (AMM)

The AMM combines equations 2 and 1 to derive the *required clock-cycle aging margin*, denoted by Δ_{Aging} . This margin quantifies the relative increase in propagation delay due to aging and is defined as the normalized difference between the aged propagation delay and the fresh propagation delay:

$$\Delta_{Aging} = \frac{t_{pd-aged}}{t_{pd-fresh}} - 1 = \left(1 - \frac{C_T \cdot e^{-E_a/(k_B T)} \cdot t^{1/n}}{V_{DD} - V_{th}} \right)^{-\alpha} - 1. \quad (4)$$

Figure 4 shows the increase of Δ_{Aging} over time for fixed temperatures (60–105 °C). At 105 °C, delay degrades by 6.5% in 3 years and 10% in 10 years. At 90 °C, degradation is 2.2–2.6% after 2–3 years and nearly 4% after 10 years. The effective temperature T_{eff} , derived from Equation 3, can be substituted into Equation 4 to estimate the Δ_{Aging} corresponding to a given mission profile. For example, the T_{eff} values for the hyperscale and HPC scenarios (60–70 °C) yield a maximum degradation of 1%, highlighting the conservative guardbands and overdesign typically assumed by IC vendors targeting worst-case conditions.

C. Aging Reclamation Model

The aging reclamation model explores strategies to reclaim spare aging margins, defined as the difference between the conservative IC vendor-provisioned margin, Δ_{Vendor} , and the required aging margin, Δ_{Aging} , derived from the AMM model and tailored to the actual mission profile. This reclaimed margin is expressed as

$$\Delta_{reclaimed} = \Delta_{Vendor} - \Delta_{Aging}, \quad (5)$$

which enables the identification of multiple viable operating frequency–voltage (f, V) pairs. The reclaimed margin, $\Delta_{reclaimed}$, can be exploited to increase the operating frequency, reduce the operating voltage, or a combination of both. When

	Δ Aging											
Δ Vop/VDD	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%	
10%	0.0%											
9%	1.2%	0.2%										
8%	2.3%	1.3%	0.3%									
7%	3.3%	2.3%	1.3%	0.3%								
6%	4.4%	3.4%	2.4%	1.4%	0.4%							
5%	5.4%	4.4%	3.4%	2.4%	1.4%	0.4%						
4%	6.4%	5.4%	4.4%	3.4%	2.4%	1.4%	0.4%					
3%	7.3%	6.3%	5.3%	4.3%	3.3%	2.3%	1.3%	0.3%				
2%	8.2%	7.2%	6.2%	5.2%	4.2%	3.2%	2.2%	1.2%	0.2%			
1%	9.1%	8.1%	7.1%	6.1%	5.1%	4.1%	3.1%	2.1%	1.1%	0.1%		
0%	10.0%	9.0%	8.0%	7.0%	6.0%	5.0%	4.0%	3.0%	2.0%	1.0%	0.0%	

Fig. 5. Aging margin reclamation heatmap illustrating the remaining timing margin available for voltage–frequency optimization as a function of required aging margin and normalized voltage reduction.

a portion of this margin is allocated to voltage reduction, the resulting timing degradation consumes part of the reclaimed margin. In this context, we define the normalized timing degradation, denoted as η_{time} , induced by a reduction in the operational voltage, ΔV_{op} , from the nominal supply voltage V_{DD} . Using the α -power law (2), η_{V-time} is expressed as

$$\eta_{V-time} \triangleq \frac{\Delta t_{pd}}{t_{pd}} = \frac{t_{pd}(V_{DD} - \Delta V_{op}) - t_{pd}(V_{DD})}{t_{pd}(V_{DD})} = \frac{V_{DD} - \Delta V_{op}}{V_{DD}} \left(\frac{V_{DD} - V_{th}}{V_{DD} - V_{th} - \Delta V_{op}} \right)^{\alpha} - 1. \quad (6)$$

After accounting for the normalized timing degradation η_{V-time} caused by voltage reduction, the remaining reclaimed margin, $\Delta_{reclaimed} - \eta_{V-time}$, can be exploited for frequency scaling, allowing the clock period to be reduced. Figure 5 illustrates the aging margin reclamation model in the form of a heat map, parameterized by the required aging margin, Δ_{Aging} , and the normalized operating voltage reduction, $\Delta V_{op}/V_{DD}$. The heat map is generated by evaluating Equations (4) and (6) using technology parameters corresponding to our 16 nm FinFET process node. Each row of the heat map corresponds to a target percentage reduction in the nominal supply voltage, while each column represents the required aging margin expressed as a percentage of the clock cycle time. For a given pair of aging degradation and voltage reduction values, the corresponding heat map entry indicates the remaining margin (in percentage) that can be reclaimed. This remaining margin can subsequently be allocated to reduce the clock cycle time and increase the operating frequency. For example, when the required aging margin is 5%, a voltage reduction of 2% yields a remaining reclaimed margin of approximately 3.2%, whereas a more aggressive voltage reduction of 4% leaves only about 1.4% of margin available for frequency scaling.

V. SIMULATION ANALYSIS

Our simulation study evaluates trade-offs among IC lifetime, performance, operating conditions, and power efficiency by comparing multiple energy–lifetime–performance configurations against a baseline processor, as summarized in Table III. The static and dynamic power scaling coefficients (f_s and f_d) are derived from synthesis and place-and-route–based power analysis in a 16 nm FinFET process across multiple operating

TABLE III
TRADE-OFFS BETWEEN LIFETIME, OPERATING VOLTAGE, FREQUENCY,
AND TEMPERATURE FOR DIFFERENT ELP PRODUCTS.

Product	V_{op}	f [GHz]	f_s	f_d	Life [Y]	T_{eff} [°C]
Baseline	V_{nom}	2.0	1.00	1.00	10	105
ELP 1	V_{nom}	2×1.09	1.00	1.09	6	70
ELP 2	$0.92V_{nom}$	2×1.013	0.860	0.857	6	70
ELP 3	$0.95V_{nom}$	2×1.024	0.910	0.924	3	90

voltages and clock frequencies for key circuit modules, including the memory subsystem, execution units, register file, and control logic, and are reported as averages across these configurations. The baseline processor operates at 2 GHz and a nominal supply voltage $V_{nom} = 1$ V under a maximum temperature of 105 °C, assuming a 10-year lifetime and a conservative 10% clock-cycle aging margin. ELP configurations consider reduced lifetimes and lower effective temperatures. ELP-1 and ELP-2 target a 6-year lifetime at 70 °C and, according to the AMM (Equation 4), require approximately 1% aging margin. Using the aging reclamation model, we can identify multiple voltage–frequency trade-offs: ELP-1 operates at nominal voltage and reallocates the remaining 9% timing margin to increase clock frequency, while ELP-2 reduces the supply voltage by 8% and exploits a reclaimed 1.3% margin to reduce the clock period and increase frequency. ELP-3 targets a 3-year lifetime at an effective temperature of 90 °C, requiring a 3% aging margin (equation 4). In this suggested configuration, ELP-3 reduces the supply voltage by 5% and the clock period by 2.4%, enabling a higher operating frequency.

A. Simulation Environment and Benchmarks

To evaluate the ELP products listed in Table III, we use Intel Xeon Silver 4114 servers running five representative applications: Memcached [15], HDSearch, SetAlgebra, Recommend, and Router [5]. The servers are based on the Skylake-SP architecture and feature two CPU sockets, each with 10 physical cores (20 logical threads), 192 GB of DRAM across six channels, and a 27.5 MB last-level cache. Memcached is deployed on a six-node cluster, with one server node and five client nodes running the Mutilate load generator [16] to emulate Facebook ETC workloads [14]. The remaining applications follow a three-tier architecture (client, midtier, and leaf) and are executed on three nodes, one per tier. For all experiments, turbo mode, dynamic frequency scaling, and power-saving states are disabled, and both core and uncore frequencies are fixed at 2 GHz to ensure consistent power and performance measurements. CPU power is monitored using *turbostat*. Each application is executed five times, and the reported results are obtained by discarding the minimum and maximum values and averaging the remaining three. We collect 99th-percentile tail latency, throughput, and average CPU power, assuming QoS constraints of 400 μ s for Memcached and 1 ms for the other applications.

B. Simulation Results

To measure baseline active power (P_a), all applications are executed at a fixed 2 GHz core frequency, with CPU power




Fig. 6. Queries per Joule Normalized with Baseline for different applications monitored on the server node for Memcached and on the midtier for all other applications. Baseline idle power (P_i) is measured when the server is idle at the same frequency. Active and idle power for each ELP product are estimated using the static and dynamic power scaling factors from Table III as

$$P_{aELPx} = f_s P_s + f_d u P_d + f_d k P_d, \quad (7)$$

$$P_{iELPx} = f_s P_s + f_d k P_d, \quad (8)$$

where P_s and P_d denote static and dynamic CPU power, respectively, and u and k represent the percentage of active power when the machine is utilized and when idle, respectively. Static power is obtained by measuring CPU power in the C1 idle state to capture leakage. Total power consumption is computed as

$$P_{totalx} = u P_{aELPx} + (1 - u) P_{iELPx}. \quad (9)$$

For each ELP configuration, we evaluate throughput (QPS) and total power to derive Queries per Joule (Q/J) as the primary metric. Figure 6 presents normalized Q/J results relative to the baseline. Overall, ELP configurations substantially outperform the baseline across most applications. The only exception is ELP 1 for Memcached, where increased power consumption outweighs QPS gains due to the application's memory-intensive behavior. ELP 2 and ELP 3 provide the most favorable trade-offs, achieving up to 20% and 13% improvements in Q/J, respectively. The smaller gain of ELP 3 is attributed to limited frequency scaling relative to voltage reduction, which reduces benefits for workloads with significant idle time.

VI. CONCLUSIONS

This paper introduced *Elastic Lifetime Processors* (ELP), a design paradigm that improves server energy efficiency by exposing processor lifetime as a configurable parameter and tailoring aging margins to the expected mission profile. By identifying and reclaiming over-provisioned worst-case aging margins, ELP enables safe exploration of voltage–frequency operating points that trade off performance, power, and lifetime more effectively than conventional one-size-fits-all designs. We developed a first-order parametric aging model that translates lifetime and temperature profiles into required timing margins and a reclamation model that reallocates spare margins toward voltage reduction, frequency scaling, or a combination of both. Using a 16 nm FinFET technology and representative workloads, our evaluation demonstrates that ELP can significantly outperform baseline processors, achieving up to 20% improvement in queries per joule.

REFERENCES

- [1] Firas Ramadan, Maayan Ella and Freddy Gabbay. 2025. Mission Profile-Driven Transistor Aging Modeling and Simulation Flow. In *Proceedings of the 2025 IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC 2025)*.
- [2] Shashikant Ilager, Adel N Toosi, Mayank Raj Jha, Ivona Brandic, and Rajkumar Buyya. 2024. A Data-driven Analysis of a Cloud Data Center: Statistical Characterization of Workload, Energy and Temperature. In *Proceedings of the 16th IEEE/ACM International Conference on Utility and Cloud Computing (UCC '23)*. Association for Computing Machinery, New York, NY, USA, Article 12, 1–10.
- [3] De Chiara D, Chinnici M, Kor AL. Data Mining for Big Dataset-Related Thermal Analysis of High Performance Computing (HPC) Data Center. *Computational Science – ICCS 2020*. 2020 May 25;12143:367–81.
- [4] Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S. Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu, “Chasing carbon: The elusive environmental footprint of computing,” in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2021, pp. 854–867. IEEE.
- [5] Akshitha Sriraman and Thomas F. Wenisch, “ μ suite: a benchmark suite for microservices,” in *2018 IEEE International Symposium on Workload Characterization (IISWC)*, 2018, pp. 1–12. IEEE.
- [6] Majid Jalili, Ioannis Manousakis, Iñigo Goiri, Pulkit A. Misra, Ashish Raniwala, Husam Alissa, Bharath Ramakrishnan, Phillip Tuma, Christian Belady, Marcus Fountoura, and others, “Cost-efficient overlocking in immersion-cooled datacenters,” in *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, 2021, pp. 623–636. IEEE.
- [7] George Papadimitriou, Athanasios Chatzidimitriou, and Dimitris Gizopoulos, “Adaptive voltage/frequency scaling and core allocation for balanced energy and performance on multicore cpus,” in *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2019, pp. 133–146. IEEE.
- [8] Lev Mukhanov, Konstantinos Tovletoglou, Georgios Karakonstantis, George Papadimitriou, Athanasios Chatzidimitriou, Manolis Kaliorakis, Dimitris Gizopoulos, and Shidhartha Das, “Improving the Energy Efficiency by Exceeding the Conservative Operating Limits,” *Hardware Accelerators in Data Centers*, Springer, 2019, pp. 241–271.
- [9] Jayanth Srinivasan, Sarita V. Adve, Pradip Bose, Jude Rivers, and Chao-Kun Hu, “Ramp: A model for reliability aware microprocessor design,” *IBM Research Report*, 2003.
- [10] Anys Bacha and Radu Teodorescu, “Dynamic reduction of voltage margins by leveraging on-chip ECC in Itanium II processors,” in *Proceedings of the 40th Annual International Symposium on Computer Architecture*, 2013, pp. 297–307.
- [11] George Papadimitriou, Manolis Kaliorakis, Athanasios Chatzidimitriou, Dimitris Gizopoulos, Peter Lawthers, and Shidhartha Das, “Harnessing voltage margins for energy efficiency in multicore CPUs,” in *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*, 2017, pp. 503–516.
- [12] George Papadimitriou, Manolis Kaliorakis, Athanasios Chatzidimitriou, Charalampos Magdalinos, and Dimitris Gizopoulos, “Voltage margins identification on commercial x86-64 multicore microprocessors,” in *2017 IEEE 23rd International Symposium on On-Line Testing and Robust System Design (IOLTS)*, 2017, pp. 51–56. IEEE.
- [13] Panagiota Nikolaou, Yiannakis Sazeides, Alejandro Lampropoulos, Denis Guilhot, Andrea Bartoli, George Papadimitriou, Athanasios Chatzidimitriou, Dimitris Gizopoulos, Konstantinos Tovletoglou, Lev Mukhanov, and others, “On the evaluation of the total-cost-of-ownership trade-offs in edge vs cloud deployments: A wireless-denial-of-service case study,” *IEEE Transactions on Sustainable Computing*, vol. 7, no. 2, pp. 334–345, 2019. IEEE.
- [14] Berk Atikoglu, Yuehai Xu, Eitan Frachtenberg, Song Jiang, and Mike Paleczny, “Workload analysis of a large-scale key-value store,” in *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems*, 2012, pp. 53–64.
- [15] Brad Fitzpatrick and Anatoly Vorobey, “Memcached: a distributed memory object caching system,” 2011.
- [16] Jacob Leverich, “Mutilate: high-performance memcached load generator,” 2014.
- [17] Takayasu Sakurai and A. Richard Newton, “Alpha-power law MOSFET model and its applications to CMOS inverter delay and other formulas,” *IEEE Journal of Solid-State Circuits*, vol. 25, no. 2, pp. 584–594, 1990. IEEE.
- [18] Cinar Kilcioglu, Justin M. Rao, Aadharsh Kannan, and R. Preston McAfee, “Usage patterns and the economics of the public cloud,” in *Proceedings of the 26th International Conference on World Wide Web*, 2017, pp. 83–91.
- [19] D. S. Boudreaux, F. Williams, and A. J. Nozik, “Hot Carrier Injection at Semiconductor Electrolyte Junctions,” *Journal of Applied Physics*, 1980. American Institute of Physics.
- [20] Dieter K. Schroder, “Negative Bias Temperature Instability: What do we understand?” *Microelectronics Reliability*, 2007. Elsevier.
- [21] James H. Stathis and Sufi Zafar, “The Negative Bias Temperature Instability in MOS Devices: A Review,” *Microelectronics Reliability*, 2006. Elsevier.
- [22] Thomas Aichinger, Gerald Rescher, and Gregor Pobegen, “Threshold Voltage Peculiarities and Bias Temperature Instabilities of SiC MOS-FETs,” *Microelectronics Reliability*, 2018. Elsevier.
- [23] Mridul Agarwal, Bipul C. Paul, Ming Zhang, and Subhasish Mitra, “Circuit failure prediction and its application to transistor aging,” in *25th IEEE VLSI Test Symposium (VTS'07)*, 2007, pp. 277–286. IEEE.
- [24] Yongho Lee and Taewhan Kim, “A Fine-grained Technique of NBTI-aware Voltage Dcaling and Body Biasing for Standard Cell-based Designs,” in *ASP-DAC 2011*, 2011. IEEE.
- [25] Lide Zhang and Robert P. Dick, “Scheduled Voltage Scaling for Increasing Lifetime in the Presence of NBTI,” in *ASP-DAC*, 2009. IEEE.
- [26] Jayanth Srinivasan, Sarita V. Adve, Pradip Bose, and Jude A. Rivers A, “Lifetime reliability: Toward an Architectural Solution,” *IEEE Micro*, 2005. IEEE.
- [27] “OpenHW Group CV32E40P User Manual,” Available online: <https://docs.openhwgroup.org/projects/cv32e40p-user-manual/en/latest/index.html>.
- [28] Freddy Gabbay and Avi Mendelson, “Asymmetric aging effect on modern microprocessors,” *Microelectronics Reliability*, vol. 119, pp. 71–81, 2021. <https://doi.org/10.1016/j.microrel.2021.114090>.
- [29] Freddy Gabbay, Avi Mendelson, Basel Salameh, and Majd Ganaiem, “A Design Flow and Tool for Avoiding Asymmetric Aging,” *IEEE Design & Test*, vol. 39, no. 6, pp. 111–118, 2022. doi: 10.1109/MDAT.2022.3183552.
- [30] Jialun Lyu, Jaylen Wang, Kali Frost, Chaojie Zhang, Irvine Celine, Esha Choukse, Rodrigo Fonseca, Ricardo Bianchini, Fiodar Kazhami-aka, and Daniel S. Berger, “Myths and misconceptions around reducing carbon embedded in cloud platforms,” in *ACM 2nd Workshop on Sustainable Computer Systems (HotCarbon'23)*, 2023.
- [31] Yewan Wang, David Nörtershäuser, Stéphane Le Masson, and Jean-Marc Menaud, “Potential effects on server power metering and modeling,” *Wireless Netw.*, vol. 29, pp. 1077–1084, 2023. <https://doi.org/10.1007/s11276-018-1882-1>.
- [32] Sai Pilla, Vishnu, “Thermal and Energy Prediction for Energy-efficient Data Centers Using Machine Learning,” Dissertation, California State University, Northridge, 2023. <https://doi.org/10.1007/s11276-018-1882-1>.
- [33] Automotive Electronics Council, Component Technical Committee, “Failure Mechanism Based Stress Test Qualification for Integrated Circuit,” *AEC-Q100-REV-G standard*.
- [34] Tao Luo, David Newmark, and David Z. Pan, “Total power optimization combining placement, sizing and multi-Vt through slack distribution management,” in *2008 Asia and South Pacific Design Automation Conference*, 2008, pp. 352–357.
- [35] Shaoming Chen, Samuel Irving, and Lu Peng, Operational Cost Optimization for Cloud Computing Data Centers Using Renewable Energy, *IEEE Systems Journal*, vol. 10, no. 4, pp. 1447–1458, 2016. doi: 10.1109/JSYST.2015.2462714.
- [36] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers, “The case for lifetime reliability-aware microprocessors,” in *Proceedings. 31st Annual International Symposium on Computer Architecture*, 2004, 2004, pp. 276–287. doi: 10.1109/ISCA.2004.1310781.
- [37] Jean S. Bozman and Randy Perry, “Server Transition Alternatives: A Business Value View Focusing on Operating Costs,” White Paper 231528R1, 2012. https://www.ibm.com/smarter-computing/ap/pdf/IDC_white_paper_Server_Transition_Alternatives_A_Business_Value_View_Focusing_on_Operating_Costs.pdf.

- [38] Daniel Hackenberg, Robert Schöne, Thomas Ilsche, Daniel Molka, Joseph Schuchart, and Robin Geyer, “An Energy Efficiency Feature Survey of the Intel Haswell Processor,” in *2015 IEEE International Parallel and Distributed Processing Symposium Workshop*, 2015, pp. 896–904. doi: 10.1109/IPDPSW.2015.70.
- [39] Mueen Uddin, Yasaman Darabdarabkhani, Asadullah Shah, and Jamshed Memon, “Evaluating power efficient algorithms for efficiency and carbon emissions in cloud data centers: A review,” *Renewable and Sustainable Energy Reviews*, vol. 51, pp. 1553–1563, 2015. doi: <https://doi.org/10.1016/j.rser.2015.07.061>.
- [40] Charles R. Lefurgy, Alan J. Drake, Michael S. Floyd, Malcolm S. Allen-Ware, Bishop Brock, Jose A. Tierno, and John B. Carter, “Active management of timing guardband to save energy in POWER7,” in *2011 44th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2011, pp. 1–11.
- [41] Ke Huang, Md Toufiq Hasan Anik, Xinqiao Zhang, and Naghmeh Karimi, “Real-Time IC Aging Prediction via On-Chip Sensors,” in *2021 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, 2021, pp. 13–18. doi: 10.1109/ISVLSI51109.2021.00014.
- [42] Karthik Swaminathan, Nandhini Chandramoorthy, Chen-Yong Cher, Ramon Bertran, Alper Buyuktosunoglu, and Pradip Bose, “BRAVO: Balanced Reliability-Aware Voltage Optimization,” in *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2017, pp. 97–108. doi: 10.1109/HPCA.2017.56.
- [43] Luiz André Barroso, Urs Hölzle, and Parthasarathy Ranganathan, *The Datacenter as a Computer*, Springer Cham, 2019.
- [44] M. Basoglu, M. Orshansky, and M. Erez, “NBTI-aware DVFS: A new approach to saving energy and increasing processor lifetime,” in *Proceedings of the 16th ACM/IEEE International Symposium on Low Power Electronics and Design*, pp. 253–258, 2010.
- [45] G. Jerke and A. B. Kahng, “Mission profile aware IC design — A case study,” 2014 Design, Automation & Test in Europe Conference & Exhibition (DATE), Dresden, Germany, 2014, pp. 1–6.