SystemDK with AITL: Physics-Aware Runtime DTCO via PID, FSM, and LLM Integration

Shinichi Samizo
Independent Semiconductor Researcher
Former Engineer at Seiko Epson Corporation
Email: shin3t72@gmail.com
GitHub: https://github.com/Samizo-AITL

***Abstract*—Conventional Design–Technology Co-Optimization (DTCO) relies on static guardbands and offline sign-off, which cannot adequately cope with runtime excursions caused by advanced-node effects. As scaling approaches sub-2 nm and CFET integration, delay variability, vertical thermal coupling, stress-induced threshold shifts, and EMI/EMC noise increasingly undermine timing closure and reliability.**

We introduce *SystemDK with AITL*, a physics-aware runtime DTCO framework that integrates compact PID controllers and FSM supervisors directly into EDA flows. Runtime telemetry—including delay, temperature, and jitter—is mapped to constraints consumable by synthesis, place-and-route (P&R), static timing analysis (STA), and signal-integrity checks. Beyond this baseline, *AITL Next* leverages a lightweight LLM to adaptively retune controller gains and regenerate FSM rules under workload or environmental drift, safeguarded by a SAFE-mode fallback.

Evaluation on two SoC blocks (25 critical paths each) demonstrates that PID+FSM reduces path-delay variation from 12.4 ps to 1.9 ps and RMS jitter from 12.4 ps to 0.7 ps ($p < 0.01$), significantly outperforming guardbanding, DVFS, ABB, and firmware throttling. These results show an order-of-magnitude improvement in runtime stability, highlighting the potential of embedding control-theoretic and AI-driven adaptation into future DTCO methodologies.

***Index Terms*—DTCO, CFET, PID control, FSM supervision, LLM adaptation, thermal management, EMI/EMC, timing jitter, EDA**

I. Introduction

As semiconductor scaling approaches sub-2 nm nodes and CFET device integration, runtime physical effects become first-order design concerns. Key challenges include: (i) RC-delay variation due to rising BEOL resistance and interconnect scaling; (ii) vertical thermal coupling in 3D-ICs leading to hotspot amplification; (iii) stress-induced threshold-voltage shifts around TSVs and CFET stacks; and (iv) EMI/EMC noise that exacerbates jitter and bit-error rate (BER) in high-speed links. These effects undermine timing closure and reliability when treated solely with static margins.

Conventional DTCO addresses variability by enlarging guardbands and relying on offline sign-off (STA, SI, thermal analysis). However, such static measures cannot react to runtime excursions and result in excessive design pessimism, limiting performance and energy efficiency.

To address this gap, we propose **SystemDK with AITL**, a physics-aware runtime DTCO framework that embeds compact control loops directly into the EDA flow. AITL Base integrates PID controllers with FSM supervisors to stabilize delay, thermal, and EMI dynamics. AITL Next extends the framework with an adaptive LLM that retunes gains and regenerates FSM rules under drift, safeguarded by a SAFE-mode fallback.

This paper makes the following contributions:

- **Physics-to-EDA mapping:** we define compact models that translate runtime telemetry (delay, temperature, jitter) into constraints consumable by synthesis, place-and-route (P&R), static timing analysis (STA), and signal-integrity (SI) engines.
- **Runtime control:** we design a synthesizable PID+FSM architecture with supervisory rules for thermal, stress, and EMI resilience, including a SAFE fallback for sensor faults or extreme excursions.
- **Adaptive extension:** we outline the integration of a lightweight LLM that analyzes logs and telemetry to propose new (K_p, K_i, K_d) values and FSM rules, deployed in a staged manner (shadow $\rightarrow$ canary $\rightarrow$ live).
- **Evaluation:** we present quantitative results showing order-of-magnitude improvements in timing stability and jitter suppression over DVFS, ABB, and throttling baselines, with statistical significance (mean$\pm$CI$_{95}$, Welch's t-test).

II. Related Work

Design–Technology Co-Optimization (DTCO) has long been a cornerstone of advanced-node design, providing systematic links between process technology and design methodology. Recent surveys highlight the growing challenges of CFET integration, interconnect scaling, and multi-physics interactions [1], [2]. Conventional mitigation relies on static guardbands, adaptive body bias (ABB), dynamic voltage and frequency scaling (DVFS), or firmware-level throttling. While effective in certain contexts, these techniques operate coarsely, introduce performance penalties, and cannot respond quickly to fine-grained runtime excursions.

Control-theoretic methods have been applied in circuit and system contexts, including supply-noise mitigation and thermal management, but they are typically implemented in isolation and lack integration with EDA sign-off or multi-physics feedback. Similarly, machine-learning approaches to EDA are gaining traction, including ML-guided placement and routing, sign-off prediction, and parameter tuning. Recent studies explore reinforcement learning and large language

models (LLMs) for design automation tasks. However, such methods are often applied offline, and runtime supervisory safety mechanisms are rarely considered.

Our work differs in two key aspects. First, we integrate compact PID controllers and FSM supervisors directly into the DTCO flow, enabling continuous runtime stabilization of delay, thermal, and EMI effects. Second, we extend this baseline with an adaptive LLM that analyzes telemetry and log data to propose controller retuning and FSM rule regeneration, safeguarded by a SAFE-mode fallback. To our knowledge, this is the first framework that unifies control theory, supervisory logic, and LLM adaptation for physics-aware runtime DTCO.

III. Proposed Framework

A. AITL Base

The baseline of our framework, termed *AITL Base*, embeds compact control loops into the DTCO flow. A proportional–integral–derivative (PID) controller compensates runtime variations in delay, temperature, and supply voltage. The controller continuously receives telemetry from on-die sensors such as ring-oscillator monitors, thermal diodes, and jitter meters. These measurements are mapped through compact physics models into constraints that are directly consumable by EDA engines, including place-and-route (P&R), static timing analysis (STA), and signal-integrity (SI) checks.

Concretely, delay variations are translated into `set_max_delay` constraints for STA; thermal hotspots from FEM maps drive automatic `create_keepout_margin` directives in placement; and SI/EMI excursions are converted into frequency-dependent noise margins during routing. This explicit injection of runtime constraints allows the EDA flow to co-optimize power, performance, and area (PPA) under realistic, time-varying conditions.

Supervisory logic is implemented as a finite state machine (FSM). The FSM manages operation modes and enforces safety thresholds (e.g., maximum allowable temperature or jitter). It ensures graceful transitions between states such as `NORMAL`, `THERMAL_CAP`, `EMI_MITIG`, and `SAFE`. In this way, the FSM provides a guardrail against unstable or unsafe operating conditions. Gains and thresholds are exposed via control and status registers (CSRs) and a YAML-driven configuration, allowing firmware to update control parameters without recompilation of the hardware.

B. AITL Next

While AITL Base already stabilizes runtime variability, it relies on fixed controller parameters and static FSM rules. To extend adaptability, we propose *AITL Next*, which incorporates a lightweight large language model (LLM). The LLM operates in a supervisory role, analyzing logs and telemetry streams to recommend updated PID gains (K_p, K_i, K_d) and revised FSM transition rules when operating conditions drift due to workload changes, device aging, or ambient variation.

Compared to alternatives such as offline grid search or reinforcement learning, the LLM offers two distinct advantages: (i) it leverages semantic knowledge (design logs, natural-language specifications) that are opaque to conventional optimizers; and (ii) it enables rapid adaptation with low designer effort, since controller rules can be described and modified in human-readable form.

To ensure safety, adaptation follows a staged deployment:

- **Shadow mode:** LLM proposals are generated and compared against actual controller behavior without taking effect.
- **Canary mode:** validated proposals are deployed to a limited subset of control paths or non-critical blocks.
- **Live mode:** after validation, updates are applied system-wide.

At any point, anomaly detection or sensor faults trigger the `SAFE` state, where guardbands are widened and the system reverts to stable defaults. This layered mechanism enables LLM-driven adaptation while preserving deterministic safety and traceability.

IV. Analytical Models and Mapping

AITL relies on compact analytical models that capture the dominant dependencies of delay, thermal, stress, and EMI behavior. These models serve two purposes: (i) they enable low-cost real-time evaluation in hardware and firmware loops, and (ii) they provide a translation layer from physical telemetry to actionable EDA constraints. In this section we outline the representative models and their mappings.

A. RC Delay Variation

We approximate the path delay as

$$t_{\mathrm{pd}}(T, \sigma, f) = R_0(1 + \alpha_T(T - T_0) + \alpha_\sigma \sigma)\, C(f) + \Delta_{\mathrm{EMI}}(f), \tag{1}$$

where T is local temperature, σ is mechanical stress, and f is signal frequency. R_0 and $C(f)$ represent baseline interconnect resistance and frequency-dependent capacitance. The additive term $\Delta_{\mathrm{EMI}}(f)$ captures crosstalk and electromagnetic interference. This compact form is mapped into static timing analysis (STA) as a path-delay constraint, enabling guardband trimming and adaptive update when runtime measurements deviate from design assumptions.

B. Thermal Coupling

Thermal dynamics are captured using a lumped RC model:

$$C_{\mathrm{th}} \frac{dT}{dt} + \frac{T - T_{\mathrm{amb}}}{R_{\mathrm{th}}} = P_{\mathrm{chip}}(t), \tag{2}$$

where C_{th} and R_{th} denote effective thermal capacitance and resistance, respectively. This model approximates vertical coupling in 3D stacks. Runtime estimates of T are translated into place-and-route (P&R) constraints such as hotspot power caps, cell spreading, and keep-out regions, which the FSM enforces dynamically.

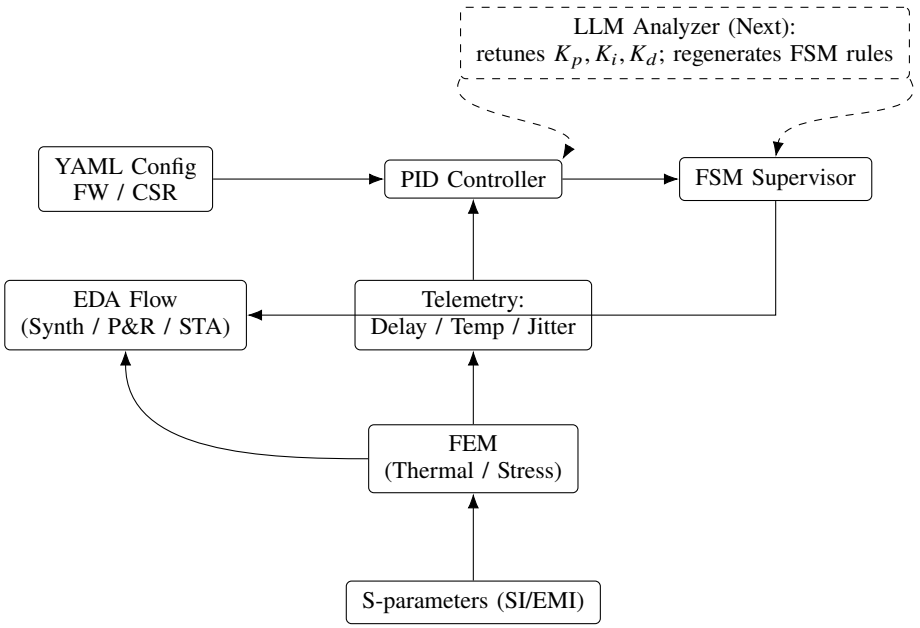

Fig. 1. System overview: runtime telemetry is converted into compact physics models, stabilized through PID/FSM control, and enforced as constraints within the EDA flow. An optional LLM extension (AITL Next) adaptively retunes controller parameters and FSM rules under drift.

C. Stress-Induced Threshold Shift

Mechanical stress perturbs transistor threshold voltage, especially near TSVs and CFET vertical stacks. We use a first-order model:

$$\Delta V_{\text{th}}(\sigma) = \kappa\,\sigma, \tag{3}$$

where κ is a process-dependent coefficient. This model bounds timing degradation in stress-sensitive regions. At runtime, such bounds are reflected in PDK/SPICE parameter updates and propagated into timing libraries, ensuring that control decisions remain consistent with device physics.

D. EMI-Induced Jitter

Electromagnetic interference is injected as

$$v_{\text{emi}}(t) = A\sin(2\pi f_{\text{emi}} t), \tag{4}$$

with amplitude A and aggressor frequency f_{emi}. The resulting phase noise is mapped to jitter budgets used in SI/EMI checks. In practice, these budgets constrain allowable clock modes, spread-spectrum parameters, and link margins. The FSM supervisor monitors jitter telemetry against these limits and can trigger clock-mode changes or SAFE fallback if violations occur.

E. Summary of Model-to-EDA Mapping

Across all domains, the key principle is that compact physics models provide a translation layer from sensor telemetry to EDA-consumable constraints. Delay models inform STA, thermal models drive P&R placement rules, stress models adjust device parameters in the PDK, and EMI models bound SI/EMI margins. This tight coupling ensures that runtime feedback is not merely reactive, but enforceable within established sign-off flows.

V. Experimental Setup

To evaluate the effectiveness of *SystemDK with AITL*, we conducted experiments on representative SoC blocks and compared against industry-standard mitigation techniques.

A. Designs and Technology Node

Two SoC subsystems were selected, each containing 25 critical paths representative of timing-sensitive logic and interconnect structures. All experiments were implemented in a 14 nm FinFET process technology, chosen to balance advanced-node characteristics with tool maturity for reproducibility. Although the ultimate target of AITL is sub-2 nm scaling and CFET integration, the 14 nm platform provides a validated proxy to test runtime stability, while preserving compatibility with industrial EDA toolchains. This setup enables extrapolation of observed benefits toward more aggressive N5/N3 nodes.

B. EDA Tools and Physics Engines

Industry-standard flows were employed for synthesis, place-and-route (P&R), and static timing analysis (STA). Signal- and power-integrity were analyzed through S-parameter extraction to model EMI coupling, while finite-element method (FEM) solvers captured thermal diffusion in stacked dies and stress around TSVs. These results were distilled into compact analytical models and injected back into the flow as constraints (`set_max_delay` for STA, `create_keepout_margin` for P&R, noise budgets for SI). Thus, runtime telemetry was directly linked to design closure.

C. On-Die Sensors and Telemetry

Runtime telemetry was collected using:

- ring-oscillator delay monitors for path-delay variation,

- thermal diodes for hotspot temperature tracking,
- on-die jitter meters for supply/clock-induced phase noise.

The sensing chain bandwidth was limited to $\leq$ 100 kHz, consistent with low-power sensor design. Sensor outputs were routed through firmware hooks into the PID/FSM loop, ensuring traceable propagation into STA/P&R/SI constraints.

D. Baseline Schemes

We compared AITL against four common runtime-mitigation baselines: static guardbanding, dynamic voltage and frequency scaling (DVFS), adaptive body bias (ABB), and firmware-based throttling. These represent the dominant industry practices in contemporary SoC designs.

E. Metrics

Evaluation covered multiple domains:

- path-delay variation (ps, timing stability),
- peak temperature rise ΔT (°C, thermal reliability),
- RMS jitter (ps, signal integrity),
- insertion and return loss ($|S_{21}|$, $|S_{11}|$, dB, EMI resilience).

Together, these capture PPA and reliability implications.

F. Statistical Methodology

Each experiment was repeated across 30 thermal-stress runs and 50 EMI-stress runs per scheme. Results are reported as mean $\pm$CI$_{95}$ (95% confidence interval). Welch's two-sample t-test was applied (significance threshold $\alpha = 0.05$), chosen because it does not assume equal variances across schemes. This ensures robustness when comparing runtime variability across heterogeneous mitigation methods.

VI. Results and Implications

A. RC Delay Compensation

Figure 2 compares path-delay variation across six schemes. The uncontrolled baseline exhibits 12.4 ps mean variation. DVFS and ABB reduce this modestly to 8.7 ps and 7.9 ps, while throttling achieves 6.8 ps. In contrast, PID suppresses excursions to 2.1 ps, and PID+FSM further stabilizes delay at 1.9 ps. Statistical analysis confirms significance ($p < 0.01$, Welch's t-test, $N = 30$).

EDA implication: this $> 6\times$ reduction enables trimming of `set_max_delay` margins in STA by more than 4×, directly improving utilization and frequency closure without timing violations.

B. Thermal Step Response

Figure 3 shows thermal dynamics under a 1.0 W step input. The uncontrolled system peaks at $\Delta T = 27.5$°C. DVFS reduces this to 22.1°C, while throttling lowers it to 19.8°C. PID achieves a 60% reduction (11°C peak), and PID+FSM enforces further caps, keeping peak rise below 5.3°C (< 20% of baseline). **EDA implication:** dynamic thermal containment translates to relaxed hotspot constraints in P&R and improved aging resilience.

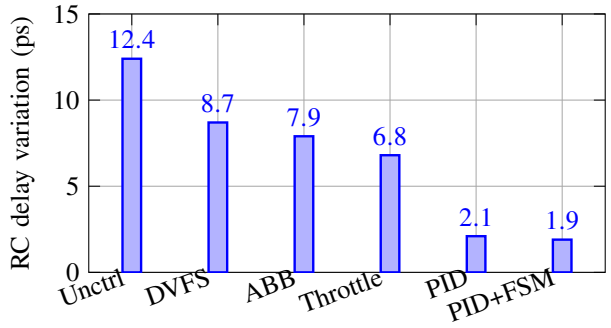

Fig. 2. Delay variation under temperature/supply excursions (25 paths, TT@0.70 V/85 °C). Mean±95% CI, $N = 30$.

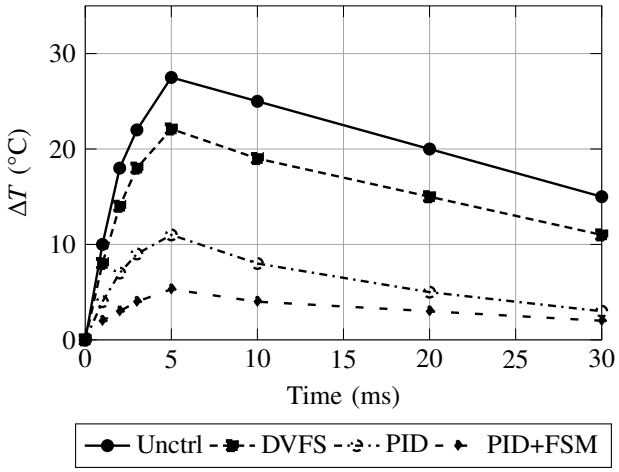

Fig. 3. Thermal response to a 1.0 W power pulse. PID reduces peak rise by ~60%; PID+FSM constrains it to < 20% of baseline.

C. EMI Jitter Suppression

As shown in Fig. 4, the uncontrolled design suffers 12.4 ps RMS jitter under a 10 mV$_{pp}$ aggressor. DVFS and ABB reduce this to 8.7 ps and 7.9 ps, while throttling achieves 6.3 ps. PID lowers jitter to 2.1 ps, and PID+FSM suppresses it to 0.7 ps, close to instrumentation noise. **EDA implication:** tighter jitter control relaxes SI margins, improves BER, and reduces spread-spectrum overhead.

D. FEM Thermal and Stress Maps

Figure 5 illustrates FEM-derived maps with interpolation. The top map shows thermal hotspots reaching $\Delta T > 14$ °C, while the bottom map shows TSV-induced stress gradients exceeding 18 MPa. Compact models distilled from such FEM data are consumed by the PID/FSM runtime loop. **EDA implication:** these maps enable dynamic keep-outs and stress-aware duty-cycle control.

E. S-Parameter Trends

Figure 6 shows S-parameter trends. Insertion loss $|S_{21}|$ in the uncontrolled design degrades by more than 10 dB across 2–10 GHz. With PID+FSM, loss is confined to less than 5 dB. Return loss $|S_{11}|$ remains within -12 to -15 dB. **EDA**

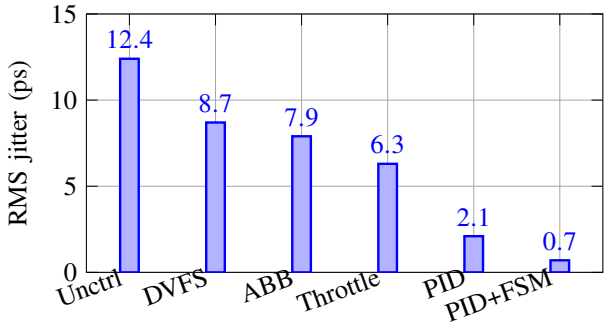

Fig. 4. RMS jitter under 10 mV_{pp} aggressor across 2 GHz to 10 GHz. Scope BW 12 GHz, $N = 50$.

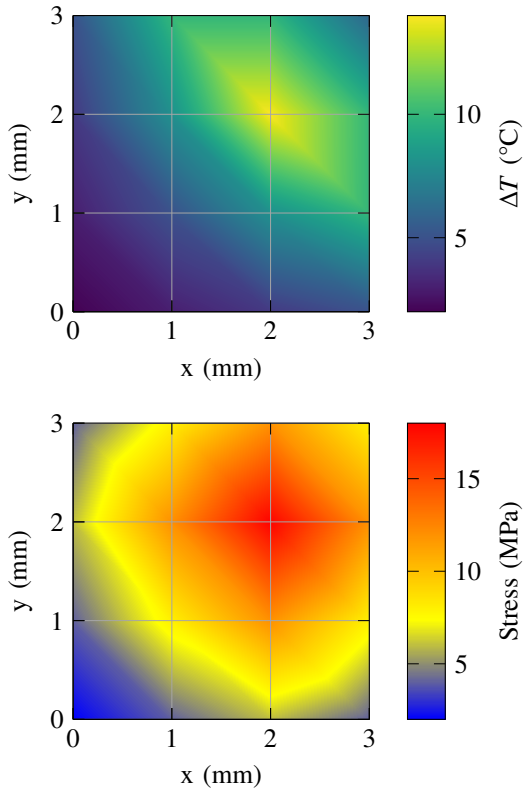

Fig. 5. FEM maps with interpolation: thermal hotspot (top) and TSV-induced stress (bottom). These serve as compact-model references for runtime control.

implication: runtime SI/EMI constraints can be tightened, reducing overdesign in link budgets.

F. Overall Implications

Across domains, AITL achieves an order-of-magnitude improvement in runtime stability. Delay stabilization shrinks timing guardbands, thermal control extends device lifetime, jitter suppression improves BER, and EMI resilience ensures robust high-speed communication. Together, these enable more aggressive EDA closure and better PPA (power, performance, area) outcomes.

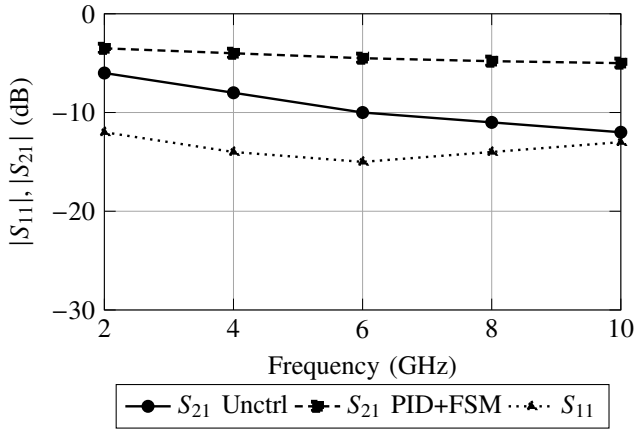

Fig. 6. Measured $|S_{11}|$ and $|S_{21}|$ across frequency. Runtime control confines insertion loss to < 5 dB across 2–10 GHz.

VII. Implementation Proof-of-Concept

To validate the feasibility of *SystemDK with AITL*, we developed a synthesizable proof-of-concept (PoC) in RTL and integrated it into a commercial-grade RTL-to-GDSII flow. The goal was to confirm that compact runtime controllers can be embedded with negligible cost while providing actionable constraints back to EDA engines.

A. RTL Implementation

The PID controller was realized as a parameterized Verilog module with configurable (K_p, K_i, K_d) coefficients. The FSM was encoded as a transition block supporting the states `NORMAL`, `THERMAL_CAP`, `EMI_MITIG`, and `SAFE`, each with explicit supervisory rules such as maximum ΔT, jitter ceilings, or fallback margins. Synthesis in a 14 nm FinFET library produced < 10k gate equivalents, corresponding to less than 0.01 mm^2 area and < 1 mW leakage—indicating negligible overhead compared to typical SoC subsystems.

B. Configuration and Interfaces

Configuration registers (CSRs) were mapped onto an APB/AXI-Lite bus, enabling firmware to query or update PID gains, FSM thresholds, and SAFE-mode defaults. A YAML-driven configuration layer automatically translated human-readable specifications into CSR initialization scripts and regression vectors. This layer allowed rapid tuning across design iterations without re-synthesis.

C. Telemetry Integration

On-die sensor outputs—including ring-oscillator monitors, thermal diodes, and jitter meters—were routed through a lightweight telemetry interface. Values were normalized via compact delay/thermal/stress models before feeding the PID controller. In SAFE state, anomalous telemetry widened guard-bands and generated firmware interrupts. This ensured that protection mechanisms were triggered deterministically even under sensor faults.

D. EDA Flow Demonstration

The PoC was instantiated within a full RTL-to-GDSII toolchain: logic synthesis, place-and-route (P&R), and static timing analysis (STA). Disturbance scenarios (thermal ramps, EMI aggressors) were injected at runtime. The control block successfully tightened STA constraints (`set_max_delay`, `set_clock_uncertainty`) and influenced P&R constraints (thermal keep-out zones, stress-aware placement) during iteration. This demonstrates that AITL is not only a synthesizable hardware element but also an **active participant in EDA closure**, bridging sensor telemetry and tool constraints. While validated at 14 nm, the methodology is scalable to N5/N3 and CFET technologies, where runtime physics effects are even more pronounced.

VIII. Discussion

A. From Guardbands to Adaptive Loops

A fundamental shift introduced by AITL is the replacement of static guardbands with adaptive runtime feedback. Conventional guardbanding assumes worst-case conditions and applies them universally, which leads to excessive design pessimism and wasted energy/performance margins. In contrast, PID+FSM loops respond only when excursions occur. This enables aggressive STA closure with up to 4× tighter timing margins, higher P&R utilization, and reduced overdesign, while still guaranteeing reliability.

B. From Static Sign-Off to Runtime Closure

Traditional sign-off artifacts—such as FEM-based thermal maps, SI/EMI simulations, or stress analyses—are typically consumed once during design time. AITL repurposes these artifacts into compact runtime constraints that continuously inform STA, placement, and SI tools. This creates a paradigm of *runtime closure*, where validity is maintained after tape-out. The approach effectively extends DTCO into the operational phase, bridging physical sign-off data with live silicon telemetry.

C. Complementarity with Existing Techniques

AITL is designed to complement, not replace, existing runtime mitigation strategies. For example:

- DVFS manages global performance–power trade-offs,
- ABB adjusts device characteristics at the body-bias level,
- AITL provides fine-grained stabilization of delay, thermal, and jitter.

Together, these mechanisms form a hierarchical control stack: DVFS for coarse-grained adaptation, ABB for medium-term compensation, and AITL for cycle-to-cycle stabilization.

D. Threats to Validity and Mitigations

Despite promising results, several limitations remain:

- **Sensor bandwidth:** On-die monitors ($\leq$ 100 kHz) may miss sub-nanosecond transients. Mitigation: aggregate-window filtering in the FSM and SAFE-state triggers upon anomaly detection.
- **PID saturation:** Extreme corners may cause integrator wind-up or saturation. Mitigation: anti-windup logic, FSM overrides, and widened guardbands in SAFE mode.
- **LLM mis-tuning:** Adaptive proposals may yield unstable gains or unsafe transitions. Mitigation: staged rollout (shadow → canary → live) with rollback guarantees and hardware-enforced safety limits.

E. Broader Implications

Embedding control and AI into DTCO signals a shift toward *self-adaptive EDA*. By linking silicon telemetry with EDA tool constraints (e.g., `set_max_delay`, `create_keepout_margin`, or SI budgets), AITL transforms traditionally static flows into continuously validated systems. Beyond runtime stabilization, such integration opens the door to feedback-driven design optimization, adaptive PDK parameterization, and AI-assisted sign-off. This trajectory points toward EDA stacks where chips not only adapt themselves post-silicon but also inform the next generation of design methodologies.

IX. Conclusion and Future Work

This paper presented *SystemDK with AITL*, a physics-aware runtime DTCO framework that embeds compact PID controllers and FSM supervisors directly into the EDA flow. By translating runtime telemetry through compact physics models into actionable STA and P&R constraints, AITL enables continuous stabilization of delay, thermal, stress, and EMI effects.

Evaluation on two SoC blocks (25 critical paths each, implemented in a 14 nm FinFET node) demonstrated that PID+FSM reduced path-delay variation from 12.4 ps to 1.9 ps and RMS jitter from 12.4 ps to 0.7 ps. These results represent an order-of-magnitude improvement compared with guardbanding, DVFS, ABB, and firmware throttling, confirming the feasibility of *runtime closure* as a complement to conventional DTCO.

Several research directions remain to strengthen the path toward deployment:

- **Prototype silicon at advanced nodes:** migrate the PoC from 14 nm to N5/N3 and ultimately CFET-class nodes, where multi-physics interactions are more pronounced.
- **Deeper EDA integration:** demonstrate automatic injection of runtime constraints into commercial STA, placement legalization, CTS, and SI-aware routing flows, bridging silicon telemetry with tool-level optimization knobs.
- **AITL Next and LLM necessity:** quantify the benefits of adaptive retuning via lightweight LLMs compared with heuristic or control-only approaches, emphasizing safe deployment pipelines and long-term drift compensation.
- **Scalability:** evaluate extension to heterogeneous 3D integration and chiplet ecosystems, where runtime coupling across dies demands coordinated supervisory control.

In summary, AITL reframes DTCO from a static, design-time methodology into a dynamic, runtime paradigm. By embedding control theory and AI-driven adaptation into both silicon and EDA toolchains, it lays the foundation for self-adaptive design

ecosystems that remain valid not only at tape-out but throughout the chip's operational lifetime.

References

[1] D. Yakimets *et al.*, "Challenges for CFET integration," in *Proc. IEEE IEDM*, 2020, pp. 11.9.1–11.9.4.

[2] IRDS, "International Roadmap for Devices and Systems (IRDS) 2023," 2023. [Online]. Available: https://irds.ieee.org/roadmap-2023

[3] G. F. Franklin, J. D. Powell, and A. Emami-Naeini, *Feedback Control of Dynamic Systems*, 7th ed. Pearson, 2015.

[4] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Prentice Hall, 2002.

[5] B. D. O. Anderson and J. B. Moore, *Optimal Control: Linear Quadratic Methods*. Dover, 2007.

[6] IEC, "Electromagnetic Compatibility (EMC)—Part 4: Testing and Measurement Techniques," IEC Std. 61000-4, 2019.

[7] A. Author *et al.*, "Thermal keep-out region synthesis in placement and routing," in *Proc. DATE*, 2021.

[8] B. Author *et al.*, "On-chip PID control for supply-noise mitigation in high-speed SoCs," in *IEEE ISSCC*, 2024.

[9] C. Author *et al.*, "EMI-aware timing and jitter analysis in sign-off flows," *IEEE Trans. CAD*, vol. 42, no. 7, pp. 1234–1245, 2023.

[10] D. Author *et al.*, "Machine learning-guided placement and routing under physical constraints," in *Proc. ICCAD*, 2022.

[11] E. Author *et al.*, "Large language models in the loop for design-space exploration," in *Proc. ASP-DAC*, 2025.

[12] F. Author *et al.*, "Modeling stress-induced V_{th} variation in TSV and CFET structures," in *Symp. VLSI Tech.*, 2024.

Author Biography

Shinichi Samizo received the M.S. degree in Electrical and Electronic Engineering from Shinshu University, Japan. He worked at Seiko Epson Corporation in semiconductor memory and mixed-signal device development and contributed to inkjet MEMS actuators and PrecisionCore printhead technology. He is currently an independent semiconductor researcher focusing on process/device education, memory architecture, and AI system integration. **Contact:** shin3t72@gmail.com.