A Full-Chip Temperature-Aware Multiphysics Framework for Electromigration, Thermomigration, and IR-Drop Analysis

Haotian Lu *Student Member, IEEE*, Sheldon X.-D. Tan *Fellow, IEEE*

***Abstract*—This paper proposes *EMSpice 3*, a full-chip multiphysics framework for coupled electromigration (EM), thermomigration (TM), and IR-drop analysis of practical power-grid (P/G) networks under realistic spatial thermal fields. To our knowledge, this is the first EM-IR analysis flow that jointly incorporates Joule heating, practical chip-level thermal maps, iterative resistance feedback, and Monte Carlo lifetime prediction for full-chip P/G designs. The framework integrates seamlessly with industry-standard Synopsys ICC and Fusion Compiler environments for practical full-chip EM/IR sign-off. It further avoids explicit tree splitting for interior-node void nucleation by leveraging the existing Robin boundary condition at the void node, and our validation demonstrates that this Robin treatment converges to the same stress and resistance-growth behavior as the explicit split-tree zero-flux formulation as the discretization is refined. The framework also checks early-failure conditions continuously during post-voiding resistance updates rather than assuming immediate open-circuit failure after nucleation. To make the coupled transient analysis more scalable, an extended rational Krylov subspace method is integrated into the solver, achieving a $1.21\times$–$1.42\times$ speedup in the six-design timing study, with nonzero but small deviations from the full-order backward-Euler finite-difference reference. The numerical results show that spatial temperature structure can dominate average temperature in determining network reliability: a higher-average-temperature profile can yield a longer TTF when its hotspots are less aligned with critical current paths, while mapped gradients can move the critical void location and change which resistor branches are degraded. Additional controlled hotspot sweeps separate the role of hotspot placement from map shape and show the limits of using average or peak temperature alone as sign-off metrics. Monte Carlo analysis further shows regime-dependent variation sensitivity: process variation can induce large lifetime spread in marginal designs, and extending the uncertain parameter set from two to four EM material variables can produce substantially stronger censoring and lifetime uncertainty. These results demonstrate that predictive full-chip EM-IR sign-off requires a unified treatment of spatial thermal fields, resistance feedback, void-growth physics, and stochastic material variation.**

***Index Terms*—Electromigration, Thermomigration, IR-drop analysis, Interconnect reliability, Model order reduction**

I. Introduction

Electromigration (EM) remains a major reliability concern in advanced copper-based VLSI circuits as technology scaling increases current density while shrinking interconnect dimensions. Traditional single-wire models such as Black's equation and Blech's limit are widely used but overly conservative for complex networks, while physics-based methods based on Korhonen's equation [1] still often neglect realistic spatial temperature distributions in coupled EM degradation and IR-drop analysis.

EM is driven by momentum transfer from conducting electrons to metal atoms. As atoms migrate away from the cathode, tensile stress builds there, while compressive stress accumulates near the anode. This stress gradient opposes the current-driven atomic flux, but if the tensile stress exceeds a critical threshold, void nucleation and growth can increase wire resistance and eventually cause circuit failure. Temperature gradients caused by Joule heating or active devices can also drive thermomigration (TM), whose magnitude can be comparable to EM [2], [3]. Although analytical and numerical EM solvers have been developed, including Krylov-accelerated approaches [4] and analytical stress solutions [5], practical chip-level thermal maps are still rarely incorporated into full-chip EM analysis.

Full-chip reliability analysis therefore requires coupled stress, void, and IR-drop modeling under realistic thermal fields, not just isolated Korhonen solves. Prior full-chip EM flows such as *EMSpice* [6] moved in this direction, yet important gaps remain: Joule heating and practical chip-level thermal maps are not fully incorporated together into the EM/IR-drop loop, and existing approaches offer limited guidance on when statistical EM analysis is most informative relative to deterministic prediction.

In this paper, we propose *EMSpice 3*, a full-chip multiphysics framework for coupled EM, TM, and IR-drop analysis of practical power-grid interconnects. The framework operates on extracted power-grid netlists, material and solver parameters, Joule-heating information, and practical chip-level thermal maps when available. It combines immortality screening, transient stress evolution, iterative resistance feedback, repeated IR-drop recomputation, and optional Monte Carlo lifetime prediction. Our results show that realistic spatial temperature fields can change both the mortal-tree population and network time-to-failure even when average temperature is unchanged. They also show that process-variation sensitivity depends strongly on how close a design is to the nucleation boundary. The key contributions are summarized as follows:

- We develop *EMSpice 3*, to our knowledge the first full-chip EM-IR analysis flow that jointly incorporates Joule heating, practical spatial thermal maps, coupled EM/TM stress evolution, iterative resistance feedback, and Monte Carlo lifetime prediction for power-grid network designs. It integrates seamlessly with industry-standard Synopsys ICC and Fusion Compiler environments for practical full-chip EM/IR sign-off and is validated on six industrial-scale designs. We also report material/solver parameters, benchmark-scale statistics, grid and time-step convergence, and full-order backward-Euler reference comparisons to make the numerical setup and accuracy claims checkable.
- We show that spatial temperature profiles can dominate average temperature in determining reliability: mapped gradients can move the critical void location, change which resistor branches are degraded, and produce different IR-drop failure outcomes even for profiles with

Haotian Lu and Sheldon X.-D. Tan are with the Department of Electrical and Computer Engineering, University of California, Riverside, CA 92521 USA

This work is supported in part by NSF grants CCF-2007135 and CCF-2305437.

similar average and maximum temperatures. This explains the observed non-monotonic relation between average temperature and TTF. We further add a controlled synthetic-hotspot sweep with fixed average temperature, peak temperature, and hotspot shape to isolate what can and cannot be inferred from hotspot placement alone.
- We integrate extended rational Krylov acceleration into the transient solver, achieving a $1.21\times$–$1.42\times$ speedup across six industrial-scale designs. The reduced-order path has small but nonzero deviations from the full-order backward-Euler finite-difference reference, which we quantify using final metrics and time-resolved stress/IR-drop errors.
- We avoid explicit tree splitting for interior-node void nucleation by leveraging the existing Robin boundary condition directly at the void node. We add both a derivation and a numerical split-tree convergence study showing that this Robin row treatment approaches the same stress distribution and resistance-growth response as the explicit split-tree zero-flux formulation under mesh refinement. We also check via-aware early-failure conditions continuously during post-voiding resistance updates rather than assuming immediate open-circuit failure after nucleation.
- We show that process-variation sensitivity is strongly regime-dependent. The Monte Carlo study reports estimator convergence and confidence intervals for the default two-variable material-variation model, and a separate four-variable study over D_0, critical stress, atomic volume, and activation energy shows that additional material uncertainty can substantially broaden the censored lifetime distribution.

The remainder of this paper is organized as follows. Section II reviews related EM, TM, and EM-aware IR-drop analysis work. Section III summarizes the multiphysics models used in this study. Section IV presents the proposed analysis framework and acceleration strategy. Section V reports the numerical results, and Section VI concludes the paper.

II. Review of Related Work

Early EM work relied on empirical models such as those of Black [7] and Blech [8], which are still widely used but overly conservative for network-level analysis and do not account for realistic thermal effects.

Korhonen et al. [1] introduced a physics-based PDE model for stress evolution in interconnect wires, enabling more accurate EM analysis at the segment level. As VLSI technology has advanced, however, it has become increasingly important to incorporate both electromigration and thermomigration, because power density and thermal resistance continue to rise in advanced ICs.

Several EM-aware IR-drop analysis frameworks have recently been proposed for full-chip EM/IR analysis [9], [6], [10], [11], [12]. The work in [9], [13] introduced the extended Korhonen model (EKM) to handle void nucleation at internal tree nodes through tree splitting. This approach also demonstrated the severity of early-failure effects on the total lifetime of power-grid networks. However, it requires a more complicated implementation because each split sub-tree must be simulated separately. It also assumes that voids quickly, or almost immediately, reach a steady state, which can reduce the accuracy of void-formation and lifetime prediction.

The *EMSpice* framework in [6] performs coupled IR-drop and EM analysis using finite-difference time-domain (FDTD) methods for full-chip EM-aware IR-drop lifetime analysis. However, this method ignores all vias and only considers void nucleation at terminal nodes, which can reduce the accuracy of void-formation and lifetime prediction. The work in [10] proposed a more accurate EM-aware IR-drop analysis flow based on a PDE-based void-growth model and explicitly considered the incubation phase of EM failure. However, it does not consider early-failure effects.

More recently, the open-source EM analysis tool PROTON was proposed in [11]. Like *EMSpice*, PROTON uses finite differences for stress analysis, but it relies on an external commercial simulator for IR-drop analysis and therefore does not capture the same integrated EM/IR-drop interaction. In addition, PROTON does not include an immortality screen to filter out EM-immune wires before transient simulation.

Overall, these EM-aware IR analysis frameworks offer more realistic EM analysis than classical empirical models, but they still do not fully account for Joule heating and the resulting thermomigration effects in both screening and transient stress analysis. The role of temperature in EM has historically been underestimated, and many models still fail to capture dynamic thermal conditions across an operating chip. Studies by Chen and Tan [3] and Abbasinasab et al. [2] highlight the importance of accurate temperature modeling for wire mortality and metal migration. More recent EM/TM work, including that of Kavousi and Chen [14] and Lamichhane et al. [15], [16], moves toward temperature-aware EM analysis, but practical chip-level thermal maps remain only partially integrated into the reliability loop, and the impact of Joule heating on EM/TM interactions in real chip designs remains insufficiently investigated and assessed. Moreover, no statistical analysis of the impact of temperature-aware EM/TM analysis on full-chip reliability has been reported.

The *EMSpice* family evolved as follows. The original *EMSpice* framework introduced the full-chip coupled EM/IR-drop loop with FDTD stress analysis and immortality filtering, but did not include Joule-heating or thermomigration effects. *EMSpice 2.0* added Joule-heating-driven temperature estimation and TM-aware stress/immortality analysis. *EMSpice 2.1* further integrated externally supplied chip thermal maps with Joule-heating analysis so that spatial operating-temperature profiles can be used in screening and lifetime prediction. *EMSpice 3* extends this line by adding real full-chip temperature-map sensitivity studies, via-retaining early-failure handling, Robin-boundary mid-tree nucleation treatment, rational-Krylov acceleration, and Monte Carlo lifetime analysis in one end-to-end EM/TM/IR-drop flow.

Compared with representative EM/IR flows [9], [10], [11], [12], *EMSpice 3* is positioned as an integrated full-chip analysis framework rather than a new EM physics model. Prior work has demonstrated extended Korhonen tree splitting, void evolution, redundancy-aware current redistribution, commercial-grade power delivery analysis, and silicon validation. Our complementary focus is the combination of full-chip thermal-map input, via-retaining netlist conversion, continuous via-aware early-failure checking, Robin-boundary mid-tree nucleation without runtime tree splitting, repeated EM/TM–MNA feedback, rational-Krylov transient acceleration, and Monte Carlo lifetime analysis on full-chip benchmarks.

III. Multiphysics Modeling for EM and IR-Drop Analysis

A. Temperature Modeling

Temperature strongly affects electromigration (EM) through both atomic diffusivity and thermomigration (TM) [17], [2].

Following the steady-state thermal model in [2], [18], the temperature profile along a single wire with boundary conditions $T(-L/2) = T_1$ and $T(+L/2) = T_2$ is

$$T(x) = T_0 + T_m \left[1 - \frac{\cos\left(\frac{x}{\Gamma}\right)}{\cos\left(\frac{L}{2\Gamma}\right)}\right] + T_n \left[1 - \frac{\sin\left(\frac{x}{\Gamma}\right)}{\sin\left(\frac{L}{2\Gamma}\right)}\right] \quad (1)$$

Here, $T_m = j^2\rho\Gamma^2/k_{\mathrm{Cu}}$ is the symmetric (Joule-heating) temperature rise, $T_n = (T_1 - T_2)/2$ is the antisymmetric boundary-asymmetry component, and Γ is the effective thermal length. Prior approaches often assume static boundary temperatures T_1 and T_2 for the whole chip [14], [15], which fail to capture the spatially varying thermal landscape of an operating die. *EMSpice 3* retains Eq. (1) but computes boundary conditions per tree: external thermal maps provide via-node temperatures when available, unconstrained endpoints are solved in tridiagonal thermal blocks, and Joule-heating-only mode uses the ambient reference T_0. The resulting per-tree temperature profile is used by both TM-aware screening and transient stress simulation.

B. Steady-State-Stress-Based Screening

In *EMSpice 3*, the primary immortality screen is a one-shot steady-state stress analysis on the original netlist. This screening step is optional and is performed before the transient EM/TM simulation begins. Its purpose is not to declare failure directly, but to identify the wires and trees that require transient follow-up.

As shown later in Section V, this screening approach does not account for the full interaction among trees over the entire EM aging process. As a result, it can identify only a subset of the truly immortal wires. Any wire that is not screened out still requires transient follow-up to determine its actual mortality. Even so, the screening remains useful in design exploration because it removes clearly immortal trees before the expensive transient stage.

For the non-TM-aware case, the screening flow first extracts the trees from the original power-grid netlist and runs the IR-drop solve to obtain the branch current densities. The original tree geometry is then combined with these current densities to build the steady-state tree set. For each tree, *EMSpice 3* solves the steady-state form of the stress equation and samples the resulting hydrostatic stress on the same branch-local stress grid used by the transient solver. A branch is classified as requiring transient follow-up when any finite sampled node stress on that branch exceeds the critical stress σ_{crit}. Otherwise, the branch is classified as immortal. A tree requires transient follow-up when at least one of its branches does; otherwise, the entire tree is screened out as immortal.

For the TM-enabled case, the flow first performs the one-time thermal analysis described above. The resulting temperature field is then used in the same steady-state stress solve, now with thermomigration-aware matrix stamping. The classification rule remains unchanged: branches and trees are screened according to whether their maximum finite steady-state stress exceeds σ_{crit}. In this sense, the TM flow differs only in the thermal field that enters the steady-state stress solve; the screening output is still reported as the set of immortal wires and trees and the set of wires and trees that require transient analysis.

This steady-state criterion is consistent with the implementation in *EMSpice 3* because it directly measures the stress state that drives void nucleation, while avoiding the need to infer mortality from an intermediate voltage criterion. The screen therefore acts as a physics-based preclassification step that reduces the transient workload without labeling all non-immortal branches as already failed.

C. IR-Drop in VLSI Interconnects

IR-drop is increasingly important in modern power grids because rising wire resistance degrades power delivery. Since EM aging evolves over days or months, the electrical potential can be treated as quasi-static [6]. Rather than solve the continuum Laplace problem directly at full-chip scale, we assume uniform current density along each wire segment, extract a resistance network from the power-grid layout, and apply modified nodal analysis (MNA):

$$M(t)u(t) = PI(t) \quad (2)$$

Here, $M(t)$ denotes the power-grid admittance matrix, which changes over time as wire resistances degrade due to EM. The matrix $P \in \mathbb{R}^{b\times p}$ maps the p current sources $I(t)$ to the nodal-voltage vector $u(t)$. The resulting nodal voltages and branch currents are then used in both the temperature-aware immortality screen and the transient EM/TM stress solve.

D. EM Stress Evolution Mechanism

In copper dual-damascene interconnects, the electron wind force drives atoms from the cathode toward the anode. Because the wire is encased in a rigid dielectric, this redistribution generates tensile hydrostatic stress at the cathode and compressive stress at the anode. Once the tensile stress exceeds a critical threshold, a void nucleates. Physics-based models [19], [20] characterize the subsequent EM failure process through three sequential phases: nucleation (t_{nuc}), incubation (t_{inc}), and growth (t_{growth}). During nucleation, stress accumulates, potentially leading to void formation at critical stress levels at cathode nodes. In the incubation phase, void growth occurs without resistance changes. The growth phase begins when the void reduces the cross-sectional area of the via, causing current to pass through a higher-resistance barrier, changing the wire resistance. This phase may result in early failure (open circuit) or progressive resistance increase up to a predefined threshold (e.g., 10% change). The total time to failure (TTF) is the sum of these three phases:

$$TTF = t_{\mathrm{nuc}} + t_{\mathrm{inc}} + t_{\mathrm{growth}} \quad (3)$$

1) Void nucleation phase: The nucleation phase lasts until a void forms. The transient hydrostatic stress evolution, $\sigma(x,t)$, is governed by the TM-aware Korhonen equation [21]:

$$\text{PDE: } \frac{\partial\sigma}{\partial t} = \frac{\partial}{\partial x}\left[\kappa(x)\left(\frac{\partial\sigma}{\partial x} - S - M\right)\right], \quad t > 0 \quad (4)$$

$$\text{BC: } \kappa(x_b)\left(\left.\frac{\partial\sigma}{\partial x}\right|_{x=x_b} - S - M\right) = 0, \quad 0 < t < t_{\mathrm{nuc}} \quad (5)$$

$$\text{IC: } \sigma(x,0) = \sigma_T \quad (6)$$

Here, $S = eZ\rho j/\Omega$ and $M = (Q/\Omega T)(\partial T/\partial x)$ represent the EM and TM driving terms. The position-dependent diffusivity is $\kappa(x) = D_a(T(x))B\Omega/(k_B T(x))$, with $D_a = D_0\exp(-E_a/(k_B T(x)))$; σ_T is the initial thermally induced residual stress, and x_b marks a block terminal.

2) Void incubation phase: When the stress reaches a critical threshold at t_{nuc}, a void forms, but the interconnect resistance remains largely unchanged until the void consumes the critical volume of the via cross-section. At the void boundary, the normal stress is approximately zero. Using an infinitesimal effective copper-void boundary thickness δ, the stress gradient

between the zero-stress void surface and the adjacent metal is modeled as [1], [22]

$$\left.\frac{\partial \sigma}{\partial x}\right|_{x=x_{\text{nuc}}} = \frac{\sigma(x_{\text{nuc}}, t)}{\delta}, \quad t_{\text{nuc}} < t < \infty \tag{7}$$

where x_{nuc} is the void-nucleation location. The void volume in a multi-segment wire follows atom conservation:

$$V_v(t) = -\int_{\Omega_L} \frac{\sigma(t)}{\mathcal{B}}\, dV \tag{8}$$

where $\mathcal{B}$ is the effective bulk elastic modulus and Ω_L is the remaining interconnect volume. Eq. (8) also applies during void growth.

The incubation interval $t_{\text{nuc}} < t < t_{\text{inc}}$ is defined by

$$V_v(t) < V_{\text{crit}} \tag{9}$$

where V_{crit} is set by the wire cross-section. In a one-dimensional model, the wire width W is used as the critical dimension because $W \geq H$ and a void spanning W consumes the cross-section. During incubation, the cathode stress rapidly drops to zero and the anode stress declines until saturation, driving void growth through Eq. (7). When the void reaches critical size, it may cause early failure by blocking a via-above connection, or late failure in a via-below structure where current can still traverse the barrier layer and increase resistance during growth [23], [6].

3) Void growth phase: During the void growth phase, the void continues to expand and wire resistance begins to change measurably, distinguishing this phase from the incubation phase. This resistance increase arises because the electrical current is forced to traverse a highly resistive barrier layer, often composed of Ta/TaN, whose resistivity is significantly greater than that of Cu. The change in wire resistance during this phase can be approximately expressed as [6]:

$$\Delta R(t) = \frac{V_\nu - V_{\text{crit}}}{WH}\left[\frac{\rho_{\text{Ta}}}{h_{\text{Ta}}(2H+W)} - \frac{\rho_{\text{Cu}}}{HW}\right] \tag{10}$$

Here, ρ_{Ta} and ρ_{Cu} represent the resistivities of the barrier material (Ta/TaN) and copper, respectively. W denotes the width of the wire, H refers to the thickness of the copper, and h_{Ta} is the thickness of the barrier layer.

4) Finite-difference transient formulation: Korhonen's equation, including the EM, stress migration (SM), and TM terms in Eq. (4), is solved by a finite-difference transient formulation on each multi-segment interconnect tree, similar to [6]. After finite-difference discretization, the PDE in Eq. (4), together with its boundary conditions (BCs) and initial condition (IC), is transformed into the linear descriptor system

$$\begin{aligned} C\dot{\boldsymbol{\sigma}}(t) &= A\boldsymbol{\sigma}(t) + B\boldsymbol{j}(t) - D \\ \boldsymbol{\sigma}(0) &= [\sigma_1(0), \sigma_2(0), \ldots, \sigma_n(0)]^\top \end{aligned} \tag{11}$$

In this formulation, $\boldsymbol{\sigma}(t)$ denotes the stress vector; C and A are $n \times n$ matrices; D is an $n \times 1$ vector associated with the TM term; and B is an $n \times p$ input matrix, where p is the number of current-density inputs. The initial condition $\boldsymbol{\sigma}(0)$ includes both the residual stress and, at later outer steps, the stress carried over from the preceding simulation state. Combining Eq. (2), Eq. (8), and Eq. (11), the coupled EM/TM/IR-drop system addressed in this work can be summarized as

$$\begin{aligned} C\dot{\boldsymbol{\sigma}}(t) &= A\boldsymbol{\sigma}(t) + B\boldsymbol{j}(t) - D \\ V_v(t) &= -\int_{\Omega_L} \frac{\sigma(t)}{\mathcal{B}} dV \\ M(t)u(t) &= PI(t) \\ \boldsymbol{\sigma}(0) &= [\sigma_1(0), \sigma_2(0), \ldots, \sigma_n(0)]^\top \end{aligned} \tag{12}$$

Eq. (12) summarizes the multiphysics system used to analyze coupled IR-drop, void volume, and stress evolution over time. The implementation details are presented in Section IV.

IV. Proposed EM/TM-Aware IR-Drop Analysis Framework with Accelerated Strategy

A. Overall Analysis Flow

Fig. 1 summarizes the overall analysis flow as implemented in *EMSpice 3*. The analysis starts from three inputs: an extracted power-grid netlist, a YAML parameter file containing the material and numerical settings, and an optional thermal map for TM-aware analysis. After parsing these files, the framework can perform an optional steady-state screening pass to prune trees that are already classified as immortal.

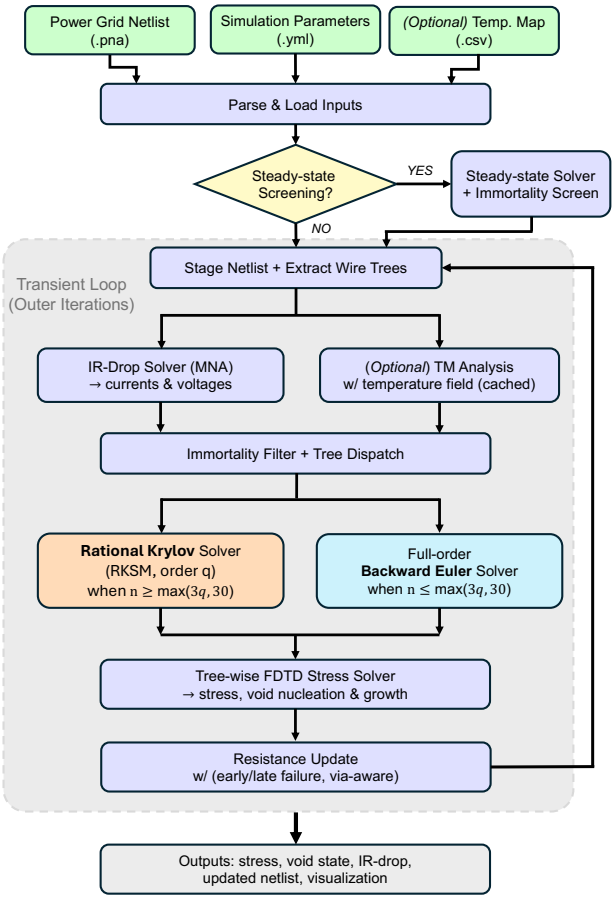


Fig. 1: Overall algorithmic flow of the proposed *EMSpice 3* framework

At each aging step the degraded netlist is partitioned into wire trees; MNA solves the IR-drop problem for nodal voltages and branch current densities, while the thermal field (when enabled) is loaded once and cached. Larger trees are accelerated by the rational Krylov path; smaller trees use backward Euler (BE) directly. Updated stress and void states are converted to resistance changes and written back to the netlist for the next outer step, closing the coupled electrical-thermal-reliability loop.

B. Finite-Difference Transient Stress Solver

The finite-difference transient stress solver combines the current-density data from the IR-drop solution with the temperature distribution to solve the discretized Korhonen equation over time. For each tree, the solver updates the hydrostatic

stress state, tracks void nucleation and growth, and computes the resulting change in resistance. These resistance updates are then written back to the staged netlist so that the next IR-drop solve reflects the current aging state of the power grid.

C. Early Failure Detection with Via-Aware Resistance Updates

Once the FDTD solver has advanced the stress state and updated void positions and lengths, the resulting resistance changes must be written back into the staged netlist before the next IR-drop solve. *EMSpice 3* distinguishes two physically distinct failure modes depending on whether the cathode node of the affected tree carries an explicit via connection to a higher metal layer.

In the common case, *late failure*, the void grows within a wire segment but does not fully block an inter-layer current path. The resistance increase is computed incrementally from the fraction of the wire cross-section consumed by the void, and the branch resistance is updated accordingly. When the cathode node of the tree is connected upward through an explicit via resistor to a strictly higher metal layer, however, the growing void can sever the only current path through that via if the void size is sufficiently large to cover the upper via, causing an abrupt open-circuit rather than a gradual degradation. This condition is termed *early failure*, and the affected branch resistance is set to an open-circuit sentinel value. Fig. 2 illustrates this via-above early-failure condition and contrasts it with the corresponding late-failure geometry.

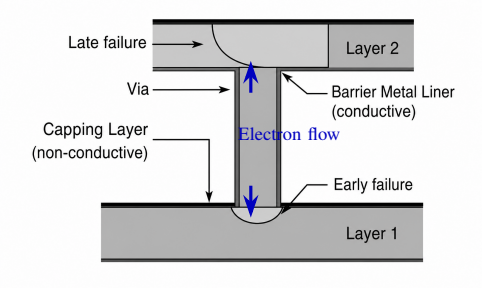


Fig. 2: Early- and late-failure mechanisms in a via-connected interconnect structure. In the via-above early-failure case, a void at the cathode can block the upward via connection and create an abrupt open circuit rather than gradual resistance degradation.

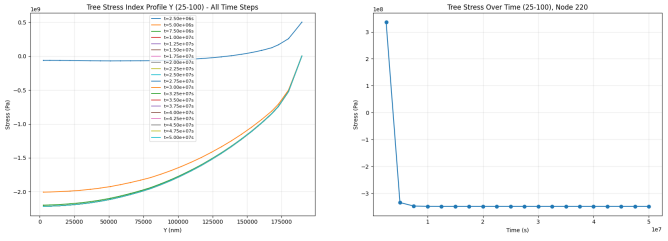
The distinction between early and late failure changes the electrical update applied after void nucleation. In the early-failure case, the electron flow drives depletion at a via-above cathode interface, so a sufficiently large void can cut the vertical current path and produce an immediate open circuit. In late (or normal) failure, the void does not fully isolate the via connection; current can still pass through the remaining liner path, so the failure appears as a gradual resistance increase during void growth rather than an abrupt topological disconnection.

Correctly identifying early failure requires explicit via information. *EMSpice 3* retains all via resistances during Synopsys PNA-to-SPICE conversion, then parses the netlist for two-terminal resistors connecting nodes on different metal layers. For each tree, the cathode node is resolved from the topology and its upward-layer connectivity is looked up in the resulting via map. When a via connection is detected, the early-failure update disconnects the upper via to create an open circuit.

In contrast, earlier *EMSpice* versions [6], [15] dropped vias during netlist conversion, making via disconnection unavailable and reducing early-failure accuracy when a cathode void severs a vertical connection.

Furthermore, in contrast to prior work [9], which assumes voids reach steady state immediately after nucleation for early-failure detection, we note that void growth depends on post-voiding stress evolution, which is very design-specific, as shown in this work. In *EMSpice 3*, the resistance update is performed at every time step, so the early-failure condition is checked continuously during void growth. This allows the framework to capture cases where a void nucleates but does not immediately cause an open circuit, instead growing over time until it eventually severs the via connection.

D. Treatment of Void Nucleation at Mid-Tree Junction Nodes



(a) Spatial stress profile (b) Junction-node stress history

Fig. 3: Mid-tree void nucleation in the dual-port RAM design: the spatial stress profile and junction-node history show interior-node nucleation and subsequent Robin-boundary saturation without dynamic tree splitting.

To handle void nucleation at mid-tree junctions, existing EM simulation frameworks apply the Extended Korhonen tree-splitting approach [13], [9], [10] to explicitly split the tree into two separate sub-trees at the nucleation point. A zero-flux (homogeneous Neumann) boundary condition is then imposed on each sub-tree at the split boundary, reflecting the physical fact that the void surface carries no net atomic flux. While this approach is physically correct, it requires dynamic bookkeeping: the tree topology must be restructured at every nucleation event, and the system matrices must be reassembled for the resulting sub-trees at each subsequent time step.

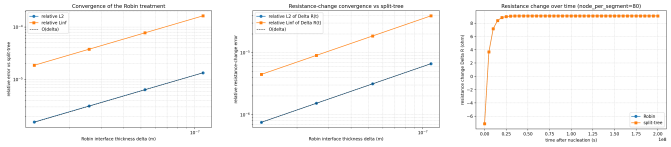
EMSpice 3 avoids tree splitting entirely through a natural application of the Robin boundary condition at the void node. The Robin condition can automatically enforce the zero-flux condition at the nucleation node within a very short time, without requiring an explicit split of the tree.

Specifically, once a void nucleates at stress-grid node v_i, the corresponding row of the assembled stiffness matrix is cleared and replaced with a Robin stencil stamped over every segment incident to v_i:

$$\left.\frac{\partial \sigma}{\partial x}\right|_{v_i} = \frac{\sigma(v_i,\, t)}{\delta}, \quad t > t_{\mathrm{nuc}}, \tag{13}$$

where δ is the infinitesimally thin void-surface thickness already used in the incubation-phase boundary condition (Eq. (7)). As $\delta \to 0$ the large diagonal coefficient $\kappa/(\Delta x \cdot \delta)$ holds the void surface at zero stress at v_i, so no net atomic flux crosses it, effectively decoupling the two sub-trees within a very short simulation time. The decoupling is therefore not instantaneous but is completed in a time period proportional to δ, which is kept negligibly small relative to the simulation time step.

This treatment significantly simplifies the implementation: no tree topology changes are performed at runtime, and the matrix assembly routine requires only a localized row replacement rather than a global remeshing. As illustrated in Fig. 3, the Robin-based approach reproduces the same physical



(a) Stress error. (b) Resistance error. (c) Resistance history.

Fig. 4: Robin-boundary validation against the explicit split-tree reference: stress and resistance errors converge with grid refinement, and the resistance-change history matches the split-tree result after nucleation.

boundary condition as the split-tree method while preserving the correctness of the EM stress dynamics.

To validate this equivalence numerically, we compare the Robin treatment with an explicit split-tree reference on a 20-segment tree whose first void nucleates at the internal node $n_{\text{void}} = 136$ at $t_{\text{nuc}} = 1.09 \times 10^1$ s. The grid is refined with $N_{\text{seg}} = 10, 20, 40, 80$, using $\delta = 0.10\Delta x_{\min}$ for each case.

Fig. 4 shows monotone convergence to the split-tree reference. From the coarsest to finest grid, the relative stress errors decrease from 1.34×10^{-5} to 1.55×10^{-6} in L_2 and from 1.64×10^{-4} to 1.87×10^{-5} in L_∞; the relative L_2 error in $\Delta R(t)$ decreases from 6.63×10^{-6} to 7.56×10^{-7}. On the finest grid, both treatments give the same final resistance change of $9.10\,\Omega$, confirming the Robin row replacement without dynamic topology changes.

E. Rational Krylov Subspace Acceleration

For power-grid trees with a large number of finite-difference stress nodes, the backward Euler solve dominates the per-time-step cost of EM/TM analysis. *EMSpice 3* therefore adopts the recently proposed extended rational Krylov subspace method [24], which projects the n-dimensional stress system onto a much smaller subspace of order q ($q \ll n$). This reduces the dominant linear-algebra cost on large trees while preserving high solution accuracy.

Compared with standard extended Krylov methods [4], the rational Krylov method requires solving a linear system with the shifted operator at each iteration, which naturally mitigates the singularity issue of the system matrix under Neumann-type boundary conditions in the EM nucleation formulation and ensures stable convergence. Furthermore, the shift parameter can be tuned to capture the dominant transient dynamics of interest, leading to improved accuracy compared with standard Krylov methods that do not use a shift [24].

However, implementing this method in a chip-level EM/IR analysis framework is not straightforward. The key challenge is that many trees in practical chip designs are small, and the overhead of the rational Krylov method can outweigh its benefits on such trees. To address this, *EMSpice 3* implements a dynamic dispatch strategy that applies the rational Krylov method only to trees above a certain size threshold, while smaller trees are solved directly with backward Euler. This allows the framework to leverage the efficiency of rational Krylov acceleration on large trees while avoiding unnecessary overhead on small trees, achieving an overall speedup across the full chip. We also provide three additional acceleration techniques, as outlined below, to further improve efficiency.

1) System formulation and extended rational Krylov subspace: After finite-difference spatial discretization of the Korhonen equation, the per-tree stress evolution is described by the linear time-invariant system

$$C\dot{\boldsymbol{\sigma}}(t) = A\,\boldsymbol{\sigma}(t) + B\,\boldsymbol{j} - D, \qquad \boldsymbol{\sigma}(0) = \boldsymbol{\sigma}_0, \tag{14}$$

which is the same descriptor form as Eq. (11). In this paper, $\boldsymbol{j}$ is assumed constant over one outer *EMSpice 3* update, but the thermomigration term satisfies $D \neq 0$ in general. Define the constant forcing vector

$$\boldsymbol{f} := B\,\boldsymbol{j} - D.$$

Applying the Laplace transform gives

$$(sC - A)\,\boldsymbol{\Sigma}(s) = C\boldsymbol{\sigma}_0 + \frac{\boldsymbol{f}}{s}.$$

Because the input is constant, define $\widetilde{\boldsymbol{\Sigma}}(s) := s\,\boldsymbol{\Sigma}(s)$, so that

$$(sC - A)\,\widetilde{\boldsymbol{\Sigma}}(s) = sC\boldsymbol{\sigma}_0 + \boldsymbol{f}.$$

We expand the solution around a shifted frequency $s_0 = 1/t_s$, where t_s is the selected expansion time scale. Define

$$K_{s_0} := s_0 C - A, \qquad R_{s_0} := K_{s_0}^{-1},$$

and let $s = s_0 + \Delta$. Expanding

$$\widetilde{\boldsymbol{\Sigma}}(s_0 + \Delta) = \sum_{k=0}^{\infty} \widetilde{\boldsymbol{m}}_k \Delta^k$$

and matching powers of Δ yields the compact recurrence

$$\widetilde{\boldsymbol{m}}_k = \begin{cases} R_{s_0}(s_0 C\boldsymbol{\sigma}_0 + \boldsymbol{f}), & k = 0, \\ R_{s_0}(C\boldsymbol{\sigma}_0 - C\widetilde{\boldsymbol{m}}_0), & k = 1, \\ -R_{s_0} C\,\widetilde{\boldsymbol{m}}_{k-1}, & k \geq 2. \end{cases} \tag{15}$$

Since $\boldsymbol{j}$ is constant here, no source-moment matching is needed; all higher-order shifted moments propagate only through the operator $R_{s_0}C$. This leads to the starting block

$$\mathcal{B}_{s_0} := \left[\, R_{s_0} C\boldsymbol{\sigma}_0,\ R_{s_0}\boldsymbol{f} \,\right]. \tag{16}$$

All shifted response moments therefore lie in the *block rational Krylov subspace*

$$\mathcal{K}_q(R_{s_0}C,\, \mathcal{B}_{s_0}) = \operatorname{range}\big(\big[\, \mathcal{B}_{s_0},\ (R_{s_0}C)\mathcal{B}_{s_0},\ \ldots,\ (R_{s_0}C)^{q-1}\mathcal{B}_{s_0} \,\big]\big). \tag{17}$$

An orthonormal projection basis $V_r \in \mathbb{R}^{n \times r}$ is constructed from this subspace via a block Arnoldi process up to order q. The implementation factors $K_{s_0} = s_0 C - A$, forms the shifted starting block from $K_{s_0}^{-1} C\boldsymbol{\sigma}_0$ and $K_{s_0}^{-1}\boldsymbol{f}$, orthonormalizes the resulting Krylov blocks, and then projects $(A, C, \boldsymbol{f})$ onto this basis. Because s_0 is chosen to match the dominant time scale of the simulation, the resulting basis captures the transient dynamics more accurately than a standard (unshifted) Krylov expansion for the same order q. Furthermore, factoring the shifted matrix $(s_0 C - A)$ rather than inverting A directly avoids the near-singularity issue of A under Neumann-type boundary conditions that has been reported for earlier Krylov EM solvers [25].

2) Reduced model and time integration: Using the final basis V_r, the reduced matrices are

$$A_h = V_r^\top A V_r, \qquad C_h = V_r^\top C V_r, \qquad \boldsymbol{f}_h = V_r^\top \boldsymbol{f}, \tag{18}$$

and the reduced descriptor system

$$C_h\dot{\hat{\boldsymbol{\sigma}}}(t) = A_h\hat{\boldsymbol{\sigma}}(t) + \boldsymbol{f}_h, \qquad \hat{\boldsymbol{\sigma}}(0) = V_r^\top \boldsymbol{\sigma}_0,$$

is integrated with backward Euler. The full-space stress field is recovered by the lift $\boldsymbol{\sigma}(t) \approx V_r\hat{\boldsymbol{\sigma}}(t)$.

3) Shift time selection: Rather than specifying s_0 directly, *EMSpice 3* uses the *shift time* $t_s = 1/s_0$, which identifies the transient time scale around which the rational Krylov moments are matched. In the present implementation, this value is fixed once per run as

$$t_s = 0.5\,T_{\text{total}}, \tag{19}$$

where T_{total} is the simulated time window. This global midpoint shift is independent of the output data and is shared by all trees in the run, which enables the same shifted factorization and basis cache strategy to be reused across time steps. The reduction-order sensitivity study in Section V therefore sweeps the Krylov order q while holding the shift rule fixed. Exposing a per-tree or user-specified shift factor would require a separate implementation option and is left outside the reported revision experiments.

4) Implementation improvements: Reduction techniques can be very effective for large trees, but they are not uniformly beneficial across all extracted power-grid trees. In practical PDN netlists, many trees are short, nearly straight segments with only one or two branches. These trees are typically very cheap to solve directly: empirical profiling in our implementation shows sub-millisecond runtime for tiny trees and only about 1–3 ms for small-to-medium trees. At that scale, the overhead of the reduction process may exceed the numerical solve cost. As a result, we propose three optimizations in *EMSpice 3* to accelerate the basic algorithm:

Basis caching. The Arnoldi basis V_r and reduced matrices $(A_h, C_h, \boldsymbol{f}_h)$ are determined by the per-tree descriptor system, the shifted starting block, and the shift s_0. *EMSpice 3* caches the basis in an LRU (least recently used) dictionary, where each entry is keyed by a fingerprint of the tree matrices together with s_0. Here, the fingerprint is a short, fixed-length hash value computed from the numerical data of the matrices and shift parameter; it uniquely identifies the input data so that cached results can be efficiently retrieved or invalidated if the inputs change. The expensive sparse LU factorization and Arnoldi iterations run only once per tree; subsequent outer steps retrieve the cached basis at negligible cost.

Loop-free time stepping for uniform grids. When the inner time grid is uniformly spaced with step Δt, the backward Euler recurrence $\hat{\boldsymbol{\sigma}}_k = F\hat{\boldsymbol{\sigma}}_{k-1} + \boldsymbol{g}$ (where $F = (C_h - \Delta t\, A_h)^{-1} C_h$ and $\boldsymbol{g} = \Delta t\,(C_h - \Delta t\, A_h)^{-1}\boldsymbol{f}_h$) is solved without a Python loop. F is eigendecomposed once as $F = E\,\mathrm{diag}(\lambda)\,E^{\top}$, and all time steps are computed simultaneously via vectorized power-series operations:

$$\hat{\boldsymbol{\sigma}}_k = E\left(\lambda^k \odot \boldsymbol{c}_0 + \textstyle\sum_{j<k} \lambda^j \odot \boldsymbol{c}_g\right),$$

where $\boldsymbol{c}_0 = E^{\top}\hat{\boldsymbol{\sigma}}_0$ and $\boldsymbol{c}_g = E^{\top}\boldsymbol{g}$. The inner cumulative sum is computed with a vectorized cumulative-sum primitive, making the entire trajectory a sequence of array operations with no per-step overhead. For nonuniform grids, the LU factorization of $(C_h - \Delta t\, A_h)$ is cached per unique Δt and reused across steps.

Batch lift. The $n_t \times r$ reduced-space trajectory is lifted to full space in a single BLAS DGEMM: series $= \hat{Z}\,V_r^{\top}$, replacing n_t individual $O(nr)$ matrix-vector products with one cache-friendly $O(n_t nr)$ matrix multiply.

5) Fallback to backward Euler and dynamic tree selection: Not all trees benefit from Krylov reduction. *EMSpice 3* therefore applies a *dynamic size threshold* before invoking the solver. In the implementation, this threshold depends not only on the state dimension n and reduction order q, but also on whether a Krylov basis is already cached and on the number of inner time steps to be advanced. On a cold solve, the basis construction and shifted sparse factorization are worthwhile only when n is sufficiently larger than q; specifically, the Krylov path is selected when

$$n \geq \max(3q,\ 30), \tag{20}$$

ensuring the full-order system is large enough relative to the reduction order that the setup cost is worthwhile. When the cache is warm, the threshold can be reduced substantially because the expensive setup is skipped. This policy is important in practice: many trees are small enough that direct backward Euler already costs only a few milliseconds, whereas the minority of larger trees benefit from repeated reduced-order solves across outer aging iterations. Small trees therefore fall through to the lightweight full-order path with negligible cost, while the expensive tail of the tree population is accelerated. The accelerated path can also be disabled globally, in which case full backward Euler is used throughout.

Even for trees that pass the size threshold, the solver may encounter numerical difficulty (e.g., rank deficiency in the Arnoldi basis or ill-conditioned shift). When this occurs, the tree is transparently re-solved with full backward Euler. Krylov reduction is therefore a *best-effort* performance optimization: it accelerates the common case without risk to correctness.

6) Computational complexity: Tab. I summarizes the per-run complexity of the baseline and accelerated *EMSpice 3* flows. Let N_g and E_g denote the number of nodes and nonzero stamps in the full-chip MNA problem, respectively. Let S be the number of outer aging steps, n_t the number of inner stress time steps per outer step, T the number of extracted interconnect trees, and $M \leq T$ the number of mortal trees after the one-shot immortality screen. For tree i, let n_i be the number of finite-difference stress nodes and $z_i = O(n_i)$ the number of nonzeros in the near-tridiagonal stress matrix. The sparse LU factorization and one triangular solve for the MNA matrix are denoted by $\mathcal{F}_g$ and $\mathcal{L}_g$, and the corresponding costs for tree i are denoted by $\mathcal{F}_i$ and $\mathcal{L}_i$. For the one-dimensional tree systems used here, $\mathcal{F}_i$ and $\mathcal{L}_i$ are close to linear in n_i; we keep them symbolic because fill-in depends on junctions, via branches, and post-void boundary conditions. The Krylov order is q, with effective reduced dimension $r \leq q \ll n_i$.

The resulting end-to-end runtime is therefore

$$O\Big(S(\mathcal{F}_g + \mathcal{L}_g) + S\sum_{i=1}^{M}(\mathcal{F}_i + n_t\mathcal{L}_i)\Big) \tag{21}$$

for the baseline solver, ignoring lower-order parsing and output costs. With Krylov acceleration and a warm basis cache after the first use of each tree, the corresponding cost becomes

$$\begin{aligned} O\Big(S(\mathcal{F}_g + \mathcal{L}_g) + \sum_{i=1}^{M}(\mathcal{F}_i + q\mathcal{L}_i + z_i q + n_i q^2) \\ + S\sum_{i=1}^{M}(q^3 + n_t q^2 + n_t n_i r)\Big). \end{aligned} \tag{22}$$

Since $r \leq q \ll n_i$, the accelerated solver replaces repeated full-order triangular solves with a small dense recurrence plus one batched lift. The full-chip MNA term remains unchanged, which is why the measured speedup is bounded by the fraction of runtime spent in the per-tree stress solver.

F. Monte Carlo Statistical Lifetime Analysis

Beyond deterministic EM/TM simulation, the framework also supports statistical reliability analysis through Monte Carlo sampling. In this mode, the full analysis pipeline is

TABLE I: Computational complexity of the baseline and accelerated *EMSpice 3* algorithms. The MNA cost is unchanged by Krylov reduction; the acceleration affects the per-tree transient stress solve and is amortized across outer aging steps by basis caching.

Stage	**Baseline full-order BE**	**Accelerated Krylov path**	**Notes**
Initial parsing, tree extraction, and matrix stamping	$O(E_g + \sum_{i=1}^{T} z_i)$	$O(E_g + \sum_{i=1}^{T} z_i)$	Sparse COO/CSC assembly is linear in the number of stamped entries.
One-shot immortality screening	$O(\mathcal{F}_g + \mathcal{L}_g + \sum_{i=1}^{T} z_i)$	$O(\mathcal{F}_g + \mathcal{L}_g + \sum_{i=1}^{T} z_i)$	Screening reduces the transient set from T to M trees, but is not itself a Krylov step.
Full-chip MNA solve per outer step	$O(\mathcal{F}_g + \mathcal{L}_g)$	$O(\mathcal{F}_g + \mathcal{L}_g)$	Resistance updates require a fresh DC solve at each aging step.
Transient stress solve for mortal tree i per outer step	$O(\mathcal{F}_i + n_t\mathcal{L}_i)$	Cold cache: $O(\mathcal{F}_i + q\mathcal{L}_i + z_i q + n_i q^2 + q^3 + n_t q^2 + n_t n_i r)$; warm cache: $O(q^3 + n_t q^2 + n_t n_i r)$	Baseline factors the full BE matrix and performs n_t full-order solves. Krylov pays shifted factorization and Arnoldi only when the basis is absent; later outer steps reuse V_r, A_h, and C_h.
Memory for mortal tree i	$O(z_i + n_i + n_t n_i)$	$O(z_i + n_i r + r^2 + n_t r + n_t n_i)$	The $n_t n_i$ term is present when the full stress history is retained for nucleation detection; otherwise final-state storage is $O(n_i)$.

re-executed for each sample, including immortality screening, transient stress evolution, resistance update, and IR-drop-based TTF evaluation. The output is therefore not a single lifetime estimate, but a distribution of TTF values together with the corresponding counts of censored and failed samples.

In the implementation used in this paper, the Monte Carlo model supports both the original two-variable material-variation setting and an extended four-variable setting. In addition to perturbing the diffusivity prefactor D_0 and the critical stress σ_{crit}, the MC flow can also sample the atomic volume Ω and activation energy E_a with prescribed coefficients of variation. Each accepted Monte Carlo draw is global for one full-chip sample: the sampled material parameters are applied consistently to all trees, branches, and finite-difference stress nodes in that run, rather than being resampled locally per tree, per branch, or per node. The sampled parameters are drawn independently from Gaussian distributions whose means are the nominal material values and whose standard deviations are specified by the user-provided CoV; nonphysical nonpositive samples are rejected and redrawn until all sampled parameters are positive. The numerical results in Section V use a 20% CoV for the sampled parameters. The 1000-sample baseline MC experiment perturbs D_0 and σ_{crit}, while the separate 100-sample four-variable experiment perturbs D_0, σ_{crit}, Ω, and E_a. Both use the *risc_core* 353 K Joule-heating condition and define failure by the same 10% IR-drop threshold as in the deterministic runs. This setup directly measures how process variation changes the predicted TTF distribution while keeping the electrical and thermal workloads fixed.

The sampling framework itself is more general than either the two-variable baseline or the four-variable material-variation study reported here. Any model input that can be mapped into the per-sample netlist, material file, or solver parameter set can in principle be treated as a random variable. Examples include resistivity, effective diffusivity, critical stress, activation-energy-related parameters, geometric dimensions, current loading, and even external temperature-map perturbations. From the solver viewpoint, these are simply alternate sampled inputs for the same coupled EM/TM/IR-drop flow, so no algorithmic change is required once the sampled values are stamped into the corresponding data structures before each run.

This separation between the sampling layer and the physics solver is important for extensibility. It allows the present paper to compare the two-variable baseline against the four-variable setting and to show that including Ω and E_a can broaden the lifetime distribution and increase censoring. It also keeps the door open to richer statistical reliability studies that incorporate additional process, thermal, electrical, or workload-induced uncertainty sources.

V. Numerical Results and Case Studies

In this section, we present numerical case studies and comparative evaluations performed using the proposed framework as implemented in *EMSpice 3*. All reported runs were performed on an Apple M3 macOS arm64 workstation with 8 CPU cores (8 logical CPUs) and 24 GB unified memory, using Python 3.13.5, NumPy 2.4.2, and SciPy 1.17.1. Unless explicitly stated otherwise, the CPU times reported in the benchmark tables are single-process baseline/Krylov runs rather than parallel-thread speedups.

The *risc_core* runs use $\sigma_{\mathrm{crit}} = 5.00 \times 10^8\,\mathrm{Pa}$, $t_{\mathrm{sim}} = 5.00 \times 10^7\,\mathrm{s}$, $N_{\mathrm{outer}} = 10$, $N_{\mathrm{inner}} = 20$, $N_x = 50$, and Krylov order $q = 5$. Key nominal material and geometry parameters are $D_0 = 7.56 \times 10^{-5}\,\mathrm{m^2/s}$, $E_a = 1.35 \times 10^{-19}\,\mathrm{J}$, $B = 1.00 \times 10^{11}\,\mathrm{Pa}$, $\Omega = 1.18 \times 10^{-29}\,\mathrm{m^3}$, $e = 1.61 \times 10^{-19}\,\mathrm{C}$, $Z = 10$, $Q = 1.29 \times 10^{-19}\,\mathrm{J}$, $\rho_{\mathrm{Cu}} = 3.00 \times 10^{-8}\,\Omega\cdot\mathrm{m}$, $\rho_{\mathrm{Ta}} = 1.76 \times 10^{-7}\,\Omega\cdot\mathrm{m}$, $\kappa_{\mathrm{Cu}} = 400.00\,\mathrm{W/(m\cdot K)}$, $h_{\mathrm{Cu}} = 9.00 \times 10^{-8}\,\mathrm{m}$, $h_{\mathrm{Ta}} = 2.00 \times 10^{-8}\,\mathrm{m}$, and $\delta = 1.00 \times 10^{-6}$. The primary MC study uses 1000 Gaussian samples, and the broader four-variable MC study uses 100 samples, both with 20% coefficient of variation and positive redraws.

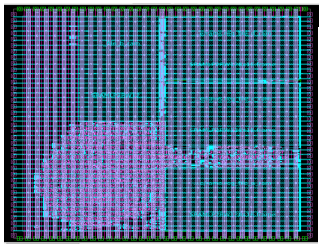
A. Case Study I: RISC-V Core

Our first case study focuses on the power grid of a RISC-V processor core, herein referred to as "RISC". This design was synthesized and placed-and-routed using the SAED32 (Synopsys 32/28nm Generic) library [26]. The design comprises 9 standard-cell blocks across 11 routing layers, with a total chip area of $104385.65\,\mu\mathrm{m}^2$ ($\approx 0.10\,\mathrm{mm}^2$). The RISC power grid spans four metal layers (M1, M2, M7, and M8, interconnected by VIA layers) with a total of 186 trees, 115 maximum branches per tree, and up to 5750 nodes per tree. The supply voltage is 0.95 V. Tab. II reports the extracted scale metrics for this benchmark. Here, the maximum stress-node count is the finite-difference stress discretization size, not the original circuit-node count; with 50 stress nodes per branch, the largest 115-branch tree contains 5750 stress unknowns. The finished layout and power grid structure are shown in Fig. 5. All simulations use a base temperature of 353 K (80 °C) and a total simulation horizon of 5×10^7 s ($\approx$1.60 years), with ten outer time steps of 5×10^6 s each.

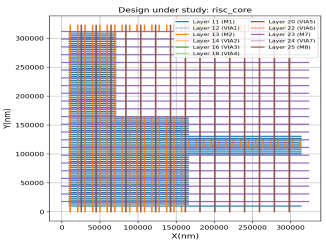
TABLE II: RISC-core benchmark statistics extracted from the deterministic run.

Design	Trees	Max br. /tree	Max stress nodes/tree	Init. IR (%)	Final IR (%)
risc_core	186	115	5750	6.63	17.75

We run two scenarios that differ only in their spatial temperature distribution, both with the same 353 K average temperature.



(a) RISC-V core: finished physical layout



(b) VDD power-grid structure

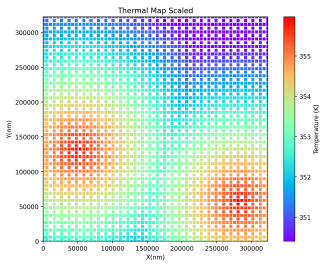
Fig. 5: RISC-V core: finished physical layout (left) and extracted power grid structure showing four metal layers (M1, M2, M7, M8) with VIA connections (right).

Case 1 — Baseline external temperature map. The first scenario uses a baseline spatial temperature distribution provided externally as input to *EMSpice 3*. As shown in Fig. 6a, this map features a localized hotspot concentrated in the lower-left region of the die, peaking at approximately 355 K while the rest of the chip stays near 351 K. The per-node Joule-heating temperature distribution computed by *EMSpice 3* under this external map is shown in Fig. 7(a). The steady-state immortality screening identifies 7 trees as immortal before the transient simulation begins. During the subsequent transient FDTD solve, 18 trees nucleate voids by the end of the simulation. The initial maximum IR-drop is 6.63% of V_{src}; it rises to 17.75% by the end of the simulation. In the deterministic sweep, the 10% IR-drop failure threshold is crossed at approximately 2.27×10^7 s (≈8.60 months).

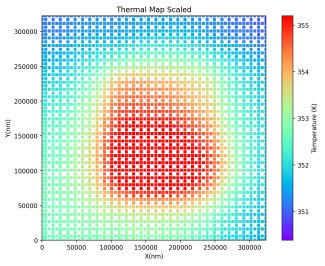
As a reference, a no-nucleation steady-state run produces continuously accumulating tensile stress, whereas void nucleation clamps the stress near the saturation threshold after flux interruption.

Case 2 — Qualcomm measured external temperature map. The second scenario uses a real spatial temperature distribution measured from a Qualcomm SM6225 Snapdragon 680 4G 8-core CPU chip, provided as an external input at the same 353 K average temperature [27]. As shown in Fig. 6b, this map features a broad central hotspot covering a large area of the die, rather than the localized corner hotspot of Case 1. The per-node Joule-heating temperature distribution computed by *EMSpice 3* under this external map is shown in Fig. 7(b). Despite the wider heated region, the transient simulation identifies only 16 mortal trees, two fewer than in Case 1. In the updated thermal sweep summarized in Tab. III, the Qualcomm 353 K case does not reach the 10% IR-drop threshold within the simulation horizon. This indicates a less severe degradation trajectory than Case 1 under the same average temperature.

Fig. 6 shows the two input spatial temperature maps. Despite sharing the same 353 K average, they differ fundamentally: the baseline map (Fig. 6a) has a compact lower-left hotspot peaking at ≈355 K, while the Qualcomm SM6225 measured map (Fig. 6b) distributes heat over a broad central region. Fig. 7 shows the resulting per-node Joule-heating distributions. In Case 1, the external baseline hotspot directly overlaps the highest-current wires in the lower-left, so the two gradients reinforce each other and compound the EM risk for those trees. In Case 2, the Qualcomm ambient peak is not aligned with the highest-current wires, so the combined thermal load on the most stressed segments is lower—explaining the smaller number of mortal trees (16 vs. 18) and the absence of a threshold crossing despite the identical average temperature.

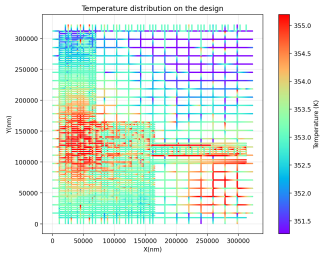


(a) Externally given baseline temperature map (Case 1): localized lower-left hotspot, peak ≈355 K, remainder ≈351 K

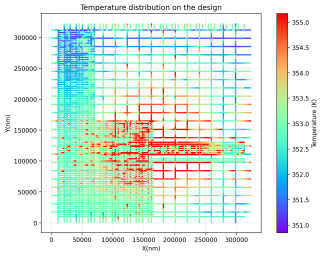


(b) Externally given Qualcomm SM6225 measured temperature map (Case 2): broad central hotspot, peak ≈355 K, same 353 K average

Fig. 6: Externally given spatial temperature maps provided as inputs to *EMSpice 3* for the two RISC-V core scenarios. Both maps share a 353 K die-average temperature but differ in spatial structure: the baseline map has a compact lower-left hotspot; the Qualcomm measured map spreads heat broadly across the central die area.



(a) Per-node Joule-heating temperature distribution (Case 1): intense heating along high-current wire segments in the lower-left; external baseline hotspot and Joule gradient reinforce each other



(b) Per-node Joule-heating temperature distribution (Case 2): similar wire-level pattern as Case 1 but modulated by the broad Qualcomm ambient, which raises base temperature centrally rather than over the highest-current wires

Fig. 7: Per-node Joule-heating temperature distributions on the RISC-V power grid computed by *EMSpice 3* for the two thermal scenarios. Individual wire segments are resolved, showing the spatial correlation between current density and local temperature. In Case 1 the external baseline hotspot directly overlaps the most current-stressed region; in Case 2 the Qualcomm ambient heat does not preferentially load the highest-current power-delivery wires.

The spatial pattern of the temperature distribution, rather than its average value, governs EM reliability: the baseline external map concentrates its hotspot over the highest-current trees near the power-delivery boundary, aligning the thermal and electrical stress peaks and producing the most severe Joule heating on those wires (Fig. 7a), while the Qualcomm map spreads ambient heat over a broader central region that does not preferentially overlap the most stressed wires, leading to fewer mortal trees and no 10% IR-drop threshold crossing for the 353 K and 373 K Qualcomm profiles within the simulated horizon.

Tab. III extends the comparison to seven thermal profiles. The mortal tree count varies from 7 (uniform 353 K) to 18 (baseline Joule-heating-derived external map at 353 K), confirming that the spatial structure—not the mean temperature—controls which trees fail. Among the 353 K profiles, the uniform case fails at 1.32×10^7 s, the baseline external map fails at 2.27×10^7 s, and the Qualcomm map does not reach the 10% IR-drop threshold within the simulation horizon. Although the uniform profile yields the fewest mortal trees, the spatially varying maps activate many more mortal trees, and the Qualcomm profile remains below the failure criterion

despite having 16 mortal trees. Notably, the two 353 K nonuniform profiles have similar maximum temperatures, 387.11 K for the baseline map and 387.57 K for the Qualcomm map, yet only the baseline map reaches the 10% threshold. At 373 K average temperature, the two maxima again remain close, 407.26 K and 407.72 K, while the Qualcomm case still does not reach the failure threshold. This indicates that hotspot alignment with critical current paths can dominate over both tree count and scalar maximum temperature.

Tab. III also exposes a more specific mechanism: the TTF of the power grid is governed not only by how many trees become mortal, but also by where the first void is placed inside a critical tree. In the uniform 353 K case, tree R11-135 nucleates at cathode node 0. The resulting void degradation first updates branch R11-135-120 and later branch R11-135-100, driving the final maximum IR-drop to 17.75% and producing a TTF of 1.32×10^7 s. Under the Qualcomm 353 K map, the same tree R11-135 nucleates instead at cathode node 9, so only branch R11-135-100 is degraded; the final maximum IR-drop remains 9.87% and the 10% threshold is not reached. The Qualcomm 373 K map follows the same void-location pattern, with a slightly higher final maximum IR-drop of 9.89%, still below the failure threshold. This comparison shows that the mapped temperature field changes the coupled EM/TM stress solution, including temperature-dependent diffusivity and thermomigration driving forces, so the peak-stress location selected for nucleation can move within the same tree. Once the void location is selected, the degraded resistor branch controls the network resistance perturbation and therefore dominates the IR-drop TTF. This branch-selection effect is a key reason why average temperature, maximum temperature, and mortal tree count alone are insufficient for sign-off.

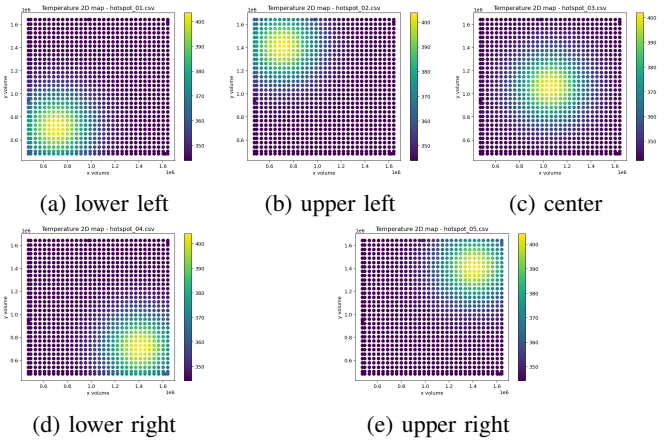
TABLE III: RISC-V core EM results under different thermal profiles (186 trees total). TTF of the power grid is the time at which the network-wide IR-drop exceeds the 10% failure threshold. A $>$ window entry indicates that the threshold was not reached within the 5×10^7 s simulation horizon.

Thermal Profile	Avg. Temp (K)	Max Temp (K)	Mortal Trees	TTF (s)
Uniform 353 K	353.00	353.00	7	1.32×10^7
Joule-heating (avg. 320 K)	320.67	354.38	18	4.71×10^7
Joule-heating (avg. 353 K)	353.40	387.11	18	2.27×10^7
Joule-heating (avg. 373 K)	373.55	407.26	18	1.45×10^7
Qualcomm map (avg. 309 K)	309.58	343.94	15	$>$ window
Qualcomm map (avg. 353 K)	353.21	387.57	16	$>$ window
Qualcomm map (avg. 373 K)	373.36	407.72	16	$>$ window

To isolate hotspot location from hotspot shape, we also generated five synthetic thermal maps on the same coordinate grid as the measured 353 K map. The original temperature values are replaced by the same analytic Gaussian hotspot in every case, and a constant offset is applied so all maps have the same average temperature. Thus the five cases in Fig. 8 and Tab. IV differ only in hotspot center location. For this *risc_core* benchmark, the five synthetic maps produce the same mortal-tree count and TTF, indicating that this particular controlled shape does not by itself produce a position-dependent TTF spread. The measured-map cases in Tab. III, however, still change the mortal-tree population and TTF because their hotspot location, shape, and spread jointly alter the thermal overlap with the critical current paths.

B. Case Study II: ARM Cortex-A Logic Core

The second test case focuses on the power grid of an ARM Cortex-A logic core design, herein referred to as "ARMcore". The design was synthesized and placed-and-routed using



(a) lower left (b) upper left (c) center

(d) lower right (e) upper right

Fig. 8: Synthetic thermal maps used for the RISC-core hotspot-position sweep. The hotspot shape, peak rise, and average temperature are fixed; only the hotspot center is moved.

TABLE IV: Synthetic hotspot-position sweep on *risc_core*.

Case	Avg. K	Max K	Mortal	TTF (s)
Hotspot 1	353.01	386.22	7	2.41×10^7
Hotspot 2	353.01	386.22	7	2.41×10^7
Hotspot 3	353.01	386.22	7	2.41×10^7
Hotspot 4	353.01	386.22	7	2.41×10^7
Hotspot 5	353.01	386.22	7	2.41×10^7

the Synopsys Design Compiler with the Synopsys 32/28nm Generic Library [26]. The ARMcore power grid spans three metal layers (M1/L11, M8/L18, and M9/L19, connected by VIA layers L21–L28) with a total of 208 trees, a maximum of 218 branches per tree, and up to 10900 nodes per tree. The supply voltage is 1.20 V. The power grid structure spanning a $450 \times 450\,\mu$m die footprint is shown in Fig. 9. All simulations use a base temperature of 353 K (80 °C) and a total simulation horizon of 5×10^7 s (≈1.60 years), with ten outer time steps of 5×10^6 s each.

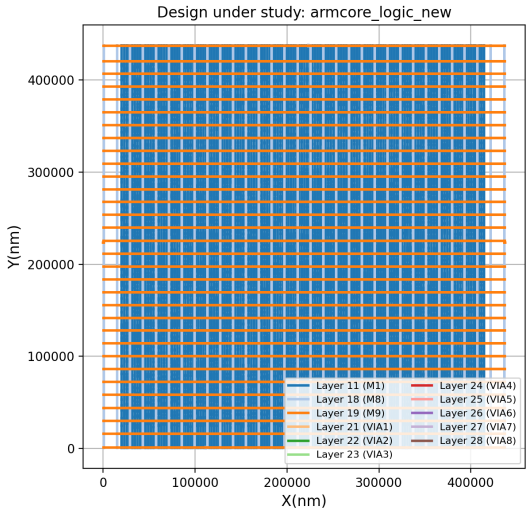


Fig. 9: Power grid structure of the ARM Cortex-A logic core (three metal layers M1, M8, M9 shown together with VIA connections)

We again run two scenarios that differ only in their spatial temperature distribution, both with the same 353 K average temperature.

Case 1 — Baseline external temperature map. The first scenario uses an updated baseline spatial temperature distribution provided externally as input to *EMSpice 3*. The resulting

per-node temperature distribution is shown in Fig. 10a. The rendered color scale emphasizes the bulk field near the ambient baseline, with warmer regions concentrated in the lower-left and lower-right portions of the die and the coolest region appearing in the upper-right corner. The full coupled run reaches a maximum temperature of about 385.40 K because localized wire-segment self-heating is superimposed on the external map.

As a diagnostic, steady-state screening reports 205 transient-follow-up trees and 3 immortal trees. In the full transient simulation used to generate Fig. 10, 206 of the 208 trees nucleate within the simulation horizon, confirming that nearly the entire network is vulnerable under this thermal field. The stress map in Fig. 10c shows a strong tensile band through the central region together with widespread nucleation across the design. The initial maximum IR-drop is 8.85% of V_{src}; it rises to 22.21% by the end of the simulation. The 10% IR-drop failure threshold is crossed at approximately 1.05×10^7 s ($\approx$4.00 months), as shown in Fig. 10e, where the maximum-drop node lies near the left boundary of the grid.

Case 2 — Qualcomm measured temperature map. The actual per-node temperature distribution derived from the Qualcomm SM6225 measured map is shown in Fig. 10b. The rendered color scale again emphasizes the bulk field near the 353 K baseline, but the warmer region now forms a broad central band spanning much of the middle of the die, while the outer edges remain cooler. This is a different spatial pattern from Case 1, whose warmer regions are split between the lower-left and lower-right portions of the die. In the coupled simulation, localized self-heating again raises the maximum wire temperature to about 385.40 K. Steady-state screening again reports 205 transient-follow-up trees and 3 immortal trees, and 206 of the 208 trees nucleate in the full transient simulation (Fig. 10d). The initial maximum IR-drop is 8.85% of V_{src}; it rises to 21.10% by the end of the simulation. The 10% failure threshold is crossed at 1.00×10^7 s ($\approx$ 3.90 months) (Fig. 10f).

1) Joule-Heating Effects on Wire Temperature: A separate ARMcore run with Joule heating only—no external thermal-map input—and a uniform 353.00 K ambient shows that self-heating is localized on loaded interconnect paths rather than forming a broad die-scale hotspot: tree 11-133 has the strongest spike, trees 18-122 and 18-128 show moderate local peaks, and tree 19-101 remains nearly flat.

Case 1 produces 206 nucleated trees and reaches the 10% IR-drop failure threshold at 1.05×10^7 s, while the Case 2 Qualcomm map also produces 206 nucleated trees and reaches failure at 1.00×10^7 s. Unlike the RISC-V core, where the thermal profile changed both the mortality regime and the dominant failure locations more dramatically, the ARMcore remains in the same nearly fully mortal regime under both 353 K external temperature maps. For this design, changing the map primarily redistributes hotspot locations and local stress patterns, while the overall IR-drop failure time remains clustered near 10^7 s.

Tab. V shows that the ARMcore remains close to a fully mortal regime across seven profiles: the mortal-tree count is 207 for the uniform 353 K case and 206 for all spatially varying thermal profiles. This weak dependence on the temperature profile occurs because the M1-layer current densities push nearly every tree far above the nucleation threshold. The TTF values also remain tightly grouped, from 1.00×10^7 s to 1.05×10^7 s, even though the maximum wire temperature ranges from 341.81 K to 405.59 K across the nonuniform maps. At 353 K average temperature, the Joule-heating and Qualcomm profiles have nearly identical maximum temperatures of 385.44 K and produce similar fully degraded behavior. Thus, for ARMcore, the thermal map changes the local hotspot and stress distribution, but the aggregate failure metric is dominated by the already deeply mortal current-density regime.

TABLE V: ARM Cortex-A logic core EM results under different thermal profiles (208 trees total). TTF of the power grid is the time at which the network-wide IR-drop exceeds the 10% failure threshold.

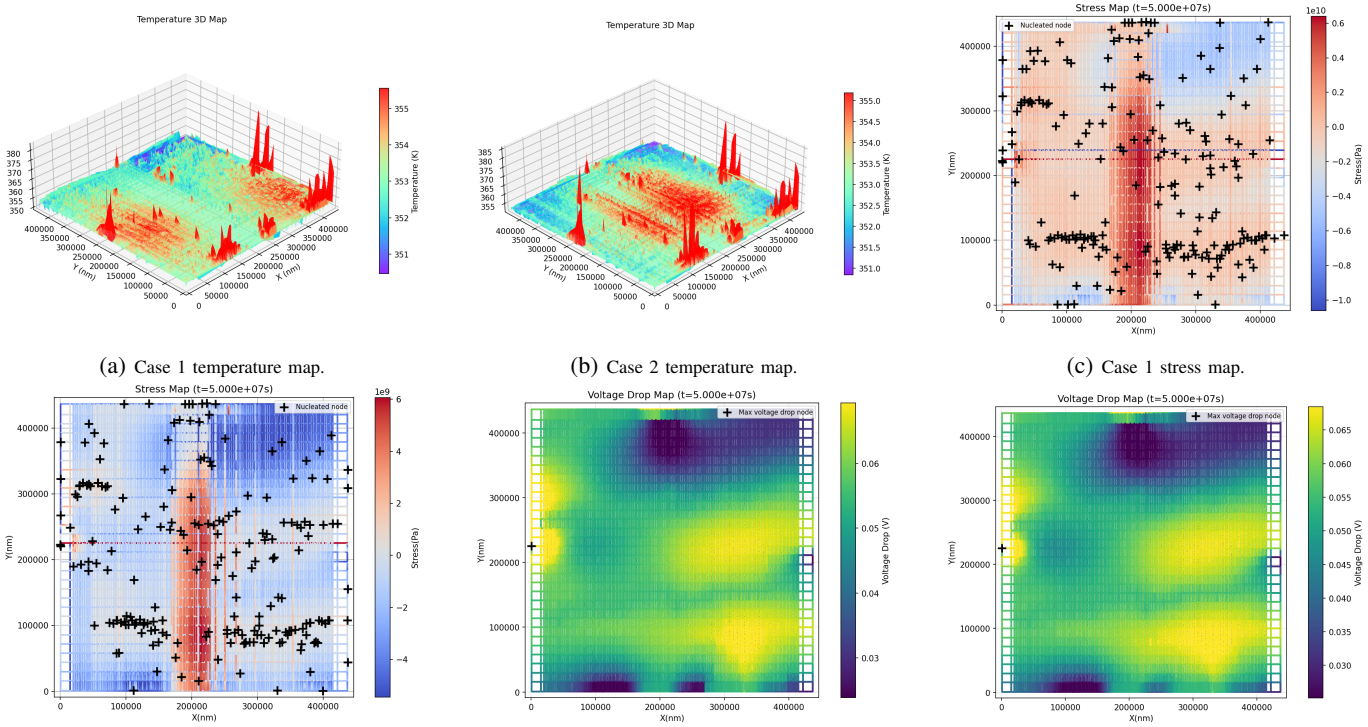
Thermal Profile	Avg. Temp (K)	Max Temp (K)	Mortal Trees	TTF (s)
Uniform 353 K	353.00	353.00	207	1.05×10^7
Joule-heating (avg. 320 K)	320.42	352.71	206	1.04×10^7
Joule-heating (avg. 353 K)	353.14	385.44	206	1.05×10^7
Joule-heating (avg. 373 K)	373.28	405.59	206	1.05×10^7
Qualcomm map (avg. 309 K)	309.55	341.81	206	1.01×10^7
Qualcomm map (avg. 353 K)	353.22	385.44	206	1.00×10^7
Qualcomm map (avg. 373 K)	373.40	405.59	206	1.00×10^7

C. Cross-Design Runtime and Accuracy Evaluation

To further evaluate computational cost and reduction effectiveness across different network structures, we tested the framework on a broader set of benchmark designs. Tab. VI summarizes the key statistics and runtime results for six designs, including the two case studies above. *AES_new* is an AES encoder design, *JPEG_new* is a JPEG encoder design, *armcore_logic* is an ARM Cortex-A logic core design, *armcore_pad* is an ARM Cortex-A pad design, *dual_ram* is a dual-port RAM design, and *risc_core* is a RISC-V core design. The PG networks of the six designs were obtained from Synopsys Fusion Compiler. All designs were synthesized and placed-and-routed using the SAED32 (Synopsys 32/28nm Generic) library [26].

For each design, we report the number of trees, the maximum number of branches and nodes per tree, the initial and final maximum IR-drop (as a percentage of V_{src}), the corresponding IR-drop time to failure (TTF), the mortal count, and the wall-clock CPU time (in seconds) with and without the Krylov-subspace acceleration described in Section IV-E. Here, the maximum node count is the finite-difference stress-node count used by the transient solver. For the Tab. VI timing runs, nonessential diagnostic text output is excluded so that the comparison measures solver and resistance-update cost rather than diagnostic-output I/O. This affects only CPU-time accounting, not the physical simulation procedure. For the two case-study designs, the listed TTF values use the same nominal 353 K Joule-heating operating point as Tabs. III and V; the final-IR and runtime columns are from the separate six-design timing batch used for Tab. VI. The "Mortal Seg. Count" column in Tab. VI reports degraded finite-difference segments, not the mortal-tree counts reported in the thermal sweep tables. The Krylov metric error uses the full-order backward-Euler finite-difference simulation as the reference: it compares TTF for designs that cross the 10% IR-drop threshold and final maximum IR-drop for designs that do not fail within the simulation horizon.

Several observations follow from Tab. VI. The Krylov-based solver consistently reduces runtime, with speedups from $1.21\times$ (*armcore_logic*) to $1.42\times$ (*risc_core*), while remaining close to the full-order backward-Euler finite-difference reference with small reported metric errors for all six benchmarks. The final IR-drop exceeds the initial value for all designs with mortal interconnects. Notably, *risc_core* shows an increase



(a) Case 1 temperature map. (b) Case 2 temperature map. (c) Case 1 stress map.

(d) Case 2 stress map. (e) Case 1 IR-drop map. (f) Case 2 IR-drop map.

Fig. 10: Temperature maps, stress distributions, and IR-drop maps for the ARM Cortex-A logic core under two thermal scenarios

TABLE VI: Performance comparison of *EMSpice 3* across benchmark designs. A TTF entry of $> 5.00 \times 10^7$ s indicates that the 10% IR-drop failure threshold was not reached within the simulated time window. The Krylov metric error is measured relative to the full-order backward-Euler finite-difference analysis, using TTF when failure occurs and the final IR-drop otherwise.

Design Name	No. Trees	Max Branch. per Tree	Max Nodes per Tree	Init. IR Drop (%)	Final IR Drop (%)	TTF (s)	Mortal Seg. Count	Baseline CPU (s)	Krylov CPU (s)	Speedup	Krylov Metric Error (%)
AES_new	97	69	3450	6.14	6.43	$> 5.00 \times 10^7$	688	2.34	1.77	1.32	< 0.01
armcore_pad	68	34	1700	0.60	0.62	$> 5.00 \times 10^7$	0	1.28	0.93	1.37	< 0.01
JPEG_new	178	126	6300	6.56	6.59	$> 5.00 \times 10^7$	1821	7.37	5.87	1.25	< 0.01
armcore_logic	208	218	10900	8.85	21.10	1.00×10^7	9283	21.69	17.87	1.21	< 0.01
dual_ram	55	29	1450	1.94	2.39	$> 5.00 \times 10^7$	179	1.09	0.79	1.38	< 0.01
risc_core	186	115	5750	6.63	10.53	2.27×10^7	19	2.78	1.96	1.42	< 0.01

from 6.63% to 10.53% with only 19 mortal wires in the separate six-design batch run, underscoring that a few high-current trees can dominate IR-drop degradation. *armcore_pad* records zero mortal interconnects, validating that temperature-aware screening efficiently prunes immortal trees before the transient FDTD stage. Runtimes scale with design complexity: the largest design, *armcore_logic* (208 trees, up to 10900 nodes), completes in about 17.87 s with Krylov acceleration, while the smallest, *dual_ram* (55 trees), finishes in under 0.80 s, confirming that the coupled analysis remains computationally tractable at the full-chip scale.

A stage-level *risc_core* breakdown separates kernel acceleration from end-to-end overhead: the transient stress solve accounts for 44.20% of the baseline runtime, and Krylov reduction lowers this stress-solve time from 1.27 s to 0.49 s ($2.61\times$). The complete EM/IR loop improves from 2.87 s to 2.12 s ($1.35\times$), consistent with the full-design trends in Tab. VI. Peak resident memory is also modest: the baseline and Krylov runs use 347.20 MB and 331.00 MB, respectively, indicating that the cached reduced bases do not dominate the process footprint for this benchmark.

Tab. VII shows the reduction-order sensitivity on *risc_core*. The paper uses $q = 5$ as the default setting. Lower order $q = 3$ is already accurate in final IR-drop but has a larger TTF deviation, whereas $q = 8$ and $q = 10$ further reduce the TTF deviation with modest runtime increase. No backward-Euler fallback occurs in these runs.

TABLE VII: Krylov reduction-order sensitivity on *risc_core* with the fixed midpoint shift $t_s = 0.5T_{\text{total}}$.

q	Runtime (s)	TTF (s)	TTF err. (s)	IR err. (pct.-pt.)
3	2.08	1.32×10^7	296.70	3.33×10^{-7}
5	2.07	1.32×10^7	14.30	3.57×10^{-8}
8	2.15	1.32×10^7	3.54×10^{-3}	8.88×10^{-9}
10	2.57	1.32×10^7	7.60×10^{-5}	8.87×10^{-9}

The small final-metric deviations do not imply mathematically zero solution error. Tab. VIII compares the Krylov path against the full-order backward-Euler finite-difference reference at every saved aging sample; the maximum relative L_∞ errors are 1.54×10^{-3} for stress and 9.22×10^{-6} for the IR-drop trajectory.

TABLE VIII: Time-resolved Krylov error on *risc_core*.

Metric	Max rel. L_2	Max rel. L_∞
Stress $\sigma(x,t)$	9.51×10^{-4}	1.54×10^{-3}
IR-drop trajectory	9.10×10^{-6}	9.22×10^{-6}

Finally, Tab. IX reports grid and time-step refinement for

risc_core. The finest internal reference uses $N_x = 100$, $N_{\text{outer}} = 40$, and $N_{\text{inner}} = 40$. Increasing the number of outer aging samples from 20 to 40 changes the final maximum IR-drop by 1.54×10^{-4} percentage points and leaves the IR-drop threshold-crossing time within 4.80×10^4 s of the finest reference. This convergence study also clarifies the TTF reporting resolution: IR drop is solved at outer aging samples, and the 10% threshold-crossing time is linearly interpolated between adjacent sampled maximum-IR-drop values that bracket the threshold.

TABLE IX: Grid and time-step convergence on *risc_core*.

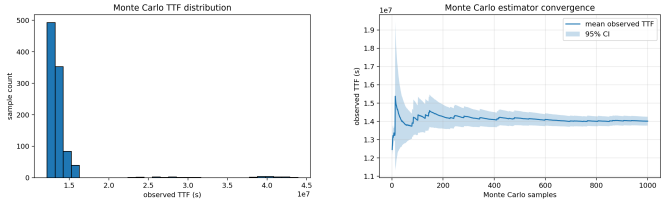
Case	N_x	N_{out}	N_{in}	TTF (s)
nx25	25	10	20	1.32×10^7
nx50	50	10	20	1.32×10^7
nx100	100	10	20	1.32×10^7
outer20	50	20	20	2.17×10^7
outer40	50	40	20	1.91×10^7
inner40	50	10	40	1.32×10^7

Component ablation on *risc_core* shows that thermal-map/TM modeling changes the physical answer by increasing the nucleated-tree count from 7 to 18 and moving the TTF from 1.32×10^7 s to 2.27×10^7 s. Krylov acceleration mainly changes cost, reducing the TM+Krylov runtime to 5.60 s while preserving the same TTF; without feeding EM-induced resistance changes back into MNA, the 10% IR-drop threshold is not reached within the simulation window.

D. Monte Carlo EM Reliability Analysis

To quantify the statistical variability of EM-induced time to failure (TTF), we perform Monte Carlo (MC) simulations using the same full-chip screening, transient stress, resistance-update, and IR-drop loop used in the deterministic analysis. Each accepted random draw is global for one full-chip sample rather than per tree or per discretization node, and sampled physical parameters are redrawn until positive. The TTF is defined by the 10% network-wide IR-drop threshold; samples that do not cross the threshold within the 5×10^7 s simulation horizon are treated as right-censored.

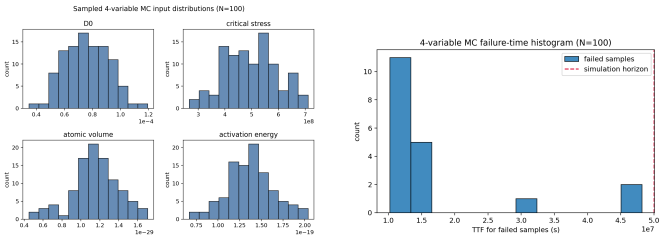
The primary MC study uses *risc_core* with independent 20% coefficient-of-variation Gaussian perturbations applied to D_0 and σ_{crit}. All 1000 samples reach the IR-drop failure threshold within the simulation horizon, with mean TTF 1.40×10^7 s, standard deviation 3.90×10^6 s, and 95% confidence-interval half-width 2.42×10^5 s; the running mean stabilizes as the sample count approaches 1000.



(a) TTF histogram. (b) Running mean and 95% confidence interval.

Fig. 11: Monte Carlo lifetime distribution and estimator convergence for the 1000-sample *risc_core* run.

To test sensitivity to a broader material-variation model, we also perform a separate 100-sample four-variable experiment in which D_0, σ_{crit}, atomic volume Ω, and activation energy E_a are independently perturbed with 20% coefficient of variation and positive redraws. This broader variation set produces substantially stronger censoring: 19 of 100 samples cross the 10% IR-drop threshold, while 81 samples remain right-censored at the simulation horizon. The observed-time mean is 4.37×10^7 s with a standard deviation of 1.39×10^7 s; considering only the 19 failed samples, the mean TTF is 1.70×10^7 s with a standard deviation of 1.16×10^7 s.



(a) Sampled input distributions. (b) Failure-time histogram.

Fig. 12: Four-variable 100-sample Monte Carlo run for *risc_core*. The sampled inputs are D_0, critical stress, atomic volume Ω, and activation energy E_a, all using independent 20% coefficient-of-variation Gaussian perturbations with positive redraws.

These results show that statistical EM analysis is most informative when the design sits near a marginal reliability boundary: uncertainty in EM transport and material parameters can shift both the threshold-crossing time and the censoring outcome. The two-variable 1000-sample study gives a stable estimate for the default MC model, while the four-variable study demonstrates that expanding the uncertainty set can qualitatively change the lifetime distribution and should be considered when the corresponding material statistics are available.

VI. Conclusions

We presented *EMSpice 3*, a full-chip multiphysics framework for coupled EM, TM, and IR-drop analysis of power-grid networks. We provide numerical evidence that the proposed Robin-boundary treatment for interior-node void nucleation reproduces the explicit split-tree zero-flux solution as the grid and void-interface parameters are refined, while avoiding explicit tree reconstruction in the full-chip flow. The results show that spatial thermal structure has a first-order impact on predicted reliability: equal-average thermal profiles can change the number of mortal trees, shift the void nucleation location and degraded branch inside a critical tree, and determine whether the 10% IR-drop failure criterion is reached within the simulation horizon. The controlled hotspot-position sweep further clarifies that hotspot alignment must be evaluated together with hotspot shape, spread, and current-path overlap, rather than through average or maximum temperature alone. The RISC-V results further show that the relationship between average temperature and TTF can be non-monotonic, because a higher-average-temperature profile can be less damaging when its hotspots are not aligned with the critical current paths. Monte Carlo analysis also shows that process-variation sensitivity is strongly design-dependent: the two-variable RISC-V study provides a stable lifetime-distribution estimate with confidence intervals, while the four-variable study shows broader censoring and stronger lifetime uncertainty when atomic volume and activation energy are also varied. Together with the consistent $1.21\times$–$1.42\times$ Krylov speedup and the measured nonzero transient-error bounds relative to the full-order backward-Euler finite-difference reference, these results indicate that predictive chip-scale EM analysis requires coupled treatment of spatial thermal fields, resistance feedback, and stochastic material variation.

ACKNOWLEDGMENTS

Generative AI tools, including Codex and Claude Code, were used solely to assist with language editing and proofreading of this paper. The authors reviewed and validated all AI-assisted edits and take full responsibility for the content of the final paper.

REFERENCES

[1] M. A. Korhonen, P. Borgesen, K. N. Tu, and C. Y. Li, "Stress evolution due to electromigration in confined metal lines," *Journal of Applied Physics*, vol. 73, no. 8, pp. 3790–3799, 1993.

[2] A. Abbasinasab and M. Marek-Sadowska, "RAIN: A tool for reliability assessment of interconnect networks—physics to software," in *Proc. Design Automation Conf. (DAC)*. New York, NY, USA: ACM, 2018, pp. 133:1–133:6.

[3] L. Chen, S. X.-D. Tan, Z. Sun, S. Peng, M. Tang, and J. Mao, "A fast semi-analytic approach for combined electromigration and thermomigration analysis for general multisegment interconnects," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 40, no. 2, pp. 350–363, 2021.

[4] C. Cook, Z. Sun, E. Demircan, M. D. Shroff, and S. X.-D. Tan, "Fast electromigration stress evolution analysis for interconnect trees using krylov subspace method," *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, vol. 26, no. 5, pp. 969–980, May 2018.

[5] H. Chen, S. X.-D. Tan, X. Huang, T. Kim, and V. Sukharev, "Analytical modeling and characterization of electromigration effects for multibranch interconnect trees," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 35, no. 11, pp. 1811–1824, 2016.

[6] Z. Sun, S. Yu, H. Zhou, Y. Liu, and S. X.-D. Tan, "EMSpice: Physics-Based Electromigration Check Using Coupled Electronic and Stress Simulation," *IEEE Transactions on Device and Materials Reliability*, vol. 20, no. 2, pp. 376–389, Jun. 2020.

[7] J. R. Black, "Electromigration-A Brief Survey and Some Recent Results," *IEEE Trans. on Electron Devices*, vol. 16, no. 4, pp. 338–347, Apr. 1969.

[8] I. A. Blech, "Electromigration in thin aluminum films on titanium nitride," *Journal of Applied Physics*, vol. 47, no. 4, pp. 1203–1208, 1976.

[9] F. N. Najm and V. Sukharev, "Efficient simulation of electromigration damage in large chip power grids using accurate physical models (invited paper)," in *2019 IEEE International Reliability Physics Symposium (IRPS)*, 2019, pp. 1–10.

[10] S. Torosyan, A. Kteyan, V. Sukharev, J.-H. Choy, and F. N. Najm, "Novel physics-based tool-prototype for electromigration assessment in commercial-grade power delivery networks," *Journal of Vacuum Science & Technology B*, vol. 39, no. 1, p. 013203, 12 2020. [Online]. Available: https://doi.org/10.1116/6.0000617

[11] O. Axelou, E. Tselepi, G. Floros, N. Evmorfopoulos, and G. Stamoulis, "Proton – a python framework for physics-based electromigration assessment on contemporary vlsi power grids," in *2023 19th International Conference on Synthesis, Modeling, Analysis and Simulation Methods and Applications to Circuit Design (SMACD)*, 2023, pp. 1–4.

[12] A. Kteyan, V. Sukharev, A. Volkov, J. H. Choy, F. N. Najm, Y. H. Yi, C. H. Kim, and S. Moreau, "Electromigration assessment in power grids with account of redundancy and non-uniform temperature distribution," in *Proceedings of the 2023 International Symposium on Physical Design*, ser. ISPD '23. Association for Computing Machinery, 2023, pp. 124–132.

[13] S. Chatterjee, "Fast and Scalable Physics-Based Electromigration Checking for Power Grids in Integrated Circuits," Ph.D. dissertation, University of Toronto, 2017.

[14] M. Kavousi, L. Chen, and S. X.-D. Tan, "Fast electromigration stress analysis considering spatial joule heating effects," in *2022 27th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2022, pp. 208–213.

[15] S. Lamichhane, M. Kavousi, and S. X.-D. Tan, "EMspice 2.0: Multiphysics electromigration analysis tool for beyond moore ics," in *2024 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, 2024.

[16] S. Lamichhane, H. Lu, and S. X. Tan, "Emspice 2.1: A coupled EM and IR drop analysis tool with joule heating and thermal map integration for VLSI reliability," in *21st International Conference on Synthesis, Modeling, Analysis and Simulation Methods, and Applications to Circuits Design, SMACD 2025, Istanbul, Turkiye, July 7-10, 2025*. IEEE, 2025, pp. 1–4. [Online]. Available: https://doi.org/10.1109/SMACD65553.2025.11092195

[17] L. Chen, S. X.-D. Tan, Z. Sun, S. Peng, M. Tang, and J. Mao, "A fast semi-analytic approach for combined electromigration and thermomigration analysis for general multi-segment interconnects," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pp. 1–1, 2020.

[18] M. Kavousi, L. Chen, and S. X.-D. Tan, "Electromigration Immortality Check considering Joule Heating Effect for Multisegment Wires," in *Proc. Int. Conf. on Computer Aided Design (ICCAD)*, 2020, pp. 1–8.

[19] S. X.-D. Tan, H. Amrouch, T. Kim, Z. Sun, C. Cook, and J. Henkel, "Recent advances in EM and BTI induced reliability modeling, analysis and optimization," *Integration, the VLSI Journal*, vol. 60, pp. 132–152, Jan. 2018.

[20] S. X.-D. Tan, M. Tahoori, T. Kim, S. Wang, Z. Sun, and S. Kiamehr, *VLSI Systems Long-Term Reliability – Modeling, Simulation and Optimization*. Springer Publishing, 2019.

[21] R. de Orio, H. Ceric, and S. Selberherr, "Physically based models of electromigration: From black's equation to modern tcad models," *Microelectronics Reliability*, vol. 50, no. 6, pp. 775–789, 2010, 2009 Reliability of Compound Semiconductors (ROCS) Workshop). [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0026271410000193

[22] V. Sukharev, A. Kteyan, and X. Huang, "Postvoiding stress evolution in confined metal lines," *IEEE Transactions on Device and Materials Reliability*, vol. 16, no. 1, pp. 50–60, 2016.

[23] S. M. Alam, C. L. Gan, C. V. Thompson, and D. E. Troxel, "Reliability computer-aided design tool for full-chip electromigration analysis and comparison with different interconnect metallizations," *Microelectron. J.*, vol. 38, no. 4-5, pp. 463–473, Apr. 2007. [Online]. Available: http://dx.doi.org/10.1016/j.mejo.2006.11.017

[24] H. Lu and S. X.-D. Tan, "Accelerating physics-based electromigration analysis via rational krylov subspaces," *arXiv preprint arXiv:2602.00330*, Feb 2026. [Online]. Available: https://arxiv.org/abs/2602.00330

[25] S. Chatterjee, V. Sukharev, and F. N. Najm, "Fast Physics-Based Electromigration Checking for On-Die Power Grids," in *Proceedings of the 35th International Conference on Computer-Aided Design*, ser. ICCAD '16. New York, NY: ACM Press, Nov. 2016, pp. 1–8. [Online]. Available: http://doi.acm.org/10.1145/2966986.2967041

[26] Synopsys, "Synopsys 32/28nm generic library for teaching ic design," Accessed: Feb. 2019, [Online]. Available. [Online]. Available: http://www.synopsys.com

[27] J. Lu and S. X.-D. Tan, "Commercial Multi/Many Cores Thermal Map Dataset," 2024, https://github.com/sheldonucr/commercial_thermal_map_dataset.



Haotian Lu received the B.E. degree in Microelectronics from Tianjin University, Tianjin, China in 2023. He is currently a Ph.D. student in VSCLAB, department of ECE at University of California, Riverside, advised by Prof. Sheldon X.-D. Tan.

His research interests mainly include VLSI reliability, electronic design automation, hardware-algorithm co-design, and hardware security.



Sheldon X.-D. Tan (M'99, SM'06, F'25) received the B.S. and M.S. degrees in electrical engineering from Fudan University, Shanghai, China, in 1992 and 1995, respectively, and the Ph.D. degree in electrical and computer engineering from the University of Iowa, Iowa City, IA, USA, in 1999. He is currently a Professor in the Department of Electrical and Computer Engineering, University of California, Riverside, CA, USA, and a cooperating faculty member in the Department of Computer Science and Engineering at UCR. He is a Fellow of IEEE and a Fellow of the Asia-Pacific Artificial Intelligence Association (AAIA). His research interests include agentic and automated design of VLSI chips and AI accelerators, AI-native multiphysics simulation and sign-off, and AI-driven agentic chiplet design. He has published more than 360 technical papers and has co-authored six books in these areas. Dr. Tan received the NSF CAREER Award in 2004. He received Best Paper Awards from ICSICT'18, ASICON'17, DAC'09, and ICCD'07, the ASP-DAC Prolific Author Award in 2020, and the Honorable Mention Best Paper Award from SMACD'18. He was a Visiting Professor at Kyoto University as a JSPS Fellow from Dec. 2017 to Jan. 2018. He served as an IEEE CEDA Distinguished Lecturer for 2022–2023, TPC Chair for ASP-DAC 2021, and TPC Vice Chair for ASP-DAC 2020. He served as Editor-in-Chief of Elsevier's *Integration, the VLSI Journal* from 2016 to 2025, and has served as Associate Editor for the *IEEE Transactions on VLSI Systems* (TVLSI), the *ACM Transactions on Design Automation of Electronic Systems* (TODAES), and Elsevier's *Microelectronics Reliability*.