Stochastic Bifurcation in Resource-Constrained Networks: Quantifying Soft-Edge Volatility (SEVI) during Monte Carlo Simulation

Venkata Mukesh Chadaram
Founder & Principal Engineer
R&D Division, Virto Sync Systems Inc.
Ottawa, Canada
mukesh@virtosync.com

***Abstract*—The application of Monte Carlo simulation to stochastic resource-constrained project scheduling (SRCPSP) relies fundamentally on what we term the *Topological Continuity Assumption* [1]–the implicit hypothesis that the underlying directed acyclic graph (DAG) structure remains invariant across simulation iterations. We demonstrate that resource-leveling heuristics systematically violate this assumption. We introduce the *Soft-Edge Volatility Index* (SEVI), a metric quantifying the probability of precedence relation reversal during simulation. Empirical analysis across 10,000 Monte Carlo iterations reveals density-dependent bifurcation: networks exhibit SEVI values approaching unity at sufficient scale, producing multimodal duration distributions that compromise the reliability of standard percentile forecasts. This so-called "Ghost P90" – an 82-day discrepancy between baseline and simulated percentiles at $N = 3000$ – emerges not from input variance but from topological discontinuity. We demonstrate that applying continuous statistical metrics to discretely bifurcating state-spaces introduces significant methodological limitations in risk quantification, particularly in highly dynamic networks, with implications extending to enterprise-level schedule forecasting under resource constraints.**

***Index Terms*—Monte Carlo simulation, stochastic scheduling, topological bifurcation, resource-constrained project scheduling, SEVI**

I. Introduction

The stochastic resource-constrained project scheduling problem (SRCPSP) remains one of the most computationally intractable problems in operations research,classified as $\mathcal{NP}$-hard in the strong sense [2], [3]. Practitioners have long relied on Monte Carlo simulation to quantify schedule risk, operating under an implicit assumption that we formalize herein as the *Topological Continuity Assumption*: the belief that the directed acyclic graph (DAG) $G = (V, E)$ representing project logic remains structurally invariant across simulation iterations, with stochasticity arising solely from duration uncertainty.

This assumption is demonstrably false under resource-constrained conditions. When resource-leveling heuristics are applied–as they must be in any realistic project environment–the precedence structure itself becomes stochastic. Activities that appear sequentially ordered in the baseline schedule may execute in arbitrary order depending on resource availability, creating what we term *Soft Edges*: precedence relations that exist by virtue of resource physics rather than structural physics.

Historically, operations research within the CoDIT community has concentrated heavily on optimizing deterministic baselines via exact methods [4], generalized networks [5], or advanced meta-heuristics [6]. However, these optimization-centric approaches focus exclusively on the initial state space. They fail to address the algorithmic validity of the resulting network when subjected to stochastic Monte Carlo simulation.

We distinguish:

- **Hard Edges** (E_H): Precedence relations arising from immutable technological dependencies. Activity j cannot commence until activity i delivers its physical output. These edges are structural invariants.
- **Soft Edges** (E_S): Precedence relations arising solely from resource contention under a particular leveling heuristic. Activity j precedes activity k only because both require the same resource and the priority rule placed j first. These edges are contingent artifacts of the scheduling algorithm.

The fundamental problem with current Monte Carlo practice is that practitioners treat the union $E = E_H \cup E_S$ as a monolithic topological invariant. To rigorously formalize this phenomenon, we approach the Resource-Constrained Project Scheduling Problem (RCPSP) through a control-theoretic lens. In this framework, the baseline DAG represents the unperturbed nominal trajectory of a Discrete-Event Dynamic System (DEDS) [7]. Resource constraints act as non-linear, state-dependent perturbations that force the system to deviate from this nominal path. When E_S dominates E_H, the system fails to converge, and the topology itself bifurcates across iterations. SEVI, therefore, operates as an advanced diagnostic metric, quantifying the system's structural sensitivity to these perturbations and measuring the topological robustness of the control network.

II. Mathematical Formulation

A. Network Representation

Consider a project represented as a directed acyclic graph $G = (V, E)$, where $V = \{v_1, v_2, \ldots, v_n\}$ represents the set of

n activities and $E \subseteq V \times V$ represents precedence relations. We partition the edge set:

$$E = E_H \cup E_S, \quad E_H \cap E_S = \emptyset \tag{1}$$

where E_H denotes hard edges (structural dependencies) and E_S denotes soft edges (resource-induced dependencies).

B. Resource Model

Let $\mathcal{R} = \{r_1, r_2, \ldots, r_m\}$ denote the set of renewable resource types with availability vector $\boldsymbol{R} = (R_1, R_2, \ldots, R_m)^\top$. Each activity $v_i \in V$ has stochastic duration $d_i \sim \mathcal{F}_i$ and resource demand vector $\boldsymbol{r}_i = (r_{i1}, r_{i2}, \ldots, r_{im})^\top$ where r_{ik} denotes units of resource k required per period of activity i's execution.

C. Resource-to-Node Density

We introduce the *Resource-to-Node Density* metric:

$$\rho = \frac{|E_S|}{|E_H| + |E_S|} = \frac{|E_S|}{|E|} \tag{2}$$

This dimensionless ratio quantifies the proportion of precedence relations attributable to resource contention. When $\rho \to 0$, the network is structurally determined; when $\rho \to 1$, the topology is almost entirely an artifact of resource-levelling decisions.

Definition 1 (Phenomenological Bifurcation (P-Bifurcation)):

Drawing on the random dynamical systems framework [8], we define a distributional phase transition in the schedule state-space as a P-Bifurcation. This occurs when the probability density function $f(t)$ of the simulated completion time T undergoes a qualitative topological change—specifically, transitioning from a unimodal to a multimodal distribution under state-dependent resource perturbations.

D. Soft-Edge Volatility Index (SEVI)

Let $\mathcal{I} = \{1, 2, \ldots, N_I\}$ index the set of Monte Carlo iterations. For each iteration $i \in \mathcal{I}$, the scheduling heuristic produces a realized edge set $E^{(i)}$ satisfying $E_H \subseteq E^{(i)} \subseteq E_H \cup E_S^{\text{pot}}$, where E_S^{pot} denotes the set of *potential* soft edges given resource constraints.

Definition 2 (Soft-Edge Volatility Index): The Soft-Edge Volatility Index (SEVI) quantifies the structural instability of the schedule. Let $p_{(u,v)}$ denote the empirical probability that a soft edge $(u,v) \in E_S^{\text{pot}}$ aligns with its baseline deterministic direction. SEVI is defined as:

$$\text{SEVI} = \frac{2}{|E_S^{\text{pot}}|} \sum_{(u,v) \in E_S^{\text{pot}}} \min\left(p_{(u,v)}, 1 - p_{(u,v)}\right) \tag{3}$$

Equivalently, let $\mathbf{1}_{(u,v)}^{(i)}$ indicate whether edge (u,v) appears in iteration i. Then:

$$\text{SEVI} = \frac{2}{|E_S^{\text{pot}}| \cdot N_I} \sum_{(u,v) \in E_S^{\text{pot}}} \min\left(\sum_{i=1}^{N_I} \mathbf{1}_{(u,v)}^{(i)},\ N_I - \sum_{i=1}^{N_I} \mathbf{1}_{(u,v)}^{(i)}\right) \tag{4}$$

SEVI $\in [0,1]$ where SEVI $= 0$ indicates topological stability and SEVI $= 1$ indicates complete topological incoherence. The factor of 2 in the numerator normalizes the index, mapping the maximum theoretical volatility (a 50/50 directional split) to unity. Critically, SEVI $= 1$ does not imply randomness—it indicates that soft edges empirically result in a *bimodal* distribution, flipping direction with probability approaching 0.5 each way.

III. Methodology

A. Schedule Generation Scheme

We employ a Time-Incrementing Vectorized Schedule Generation Scheme (SGS) implementing a First-Come-First-Serve (FCFS) priority heuristic. This approach differs fundamentally from dynamic programming formulations:

Algorithm 1 Parallel FCFS Schedule Generation Scheme (Event-Based)

Input: DAG $G = (V, E_H)$, durations $\boldsymbol{d} \sim \mathcal{F}$, capacities $\boldsymbol{R}$
Initialize: $t \leftarrow 0$, $E_S \leftarrow \emptyset$, completed $\leftarrow \emptyset$, active $\leftarrow \emptyset$
while $|\text{completed}| < |V|$ **do**
  eligible $\leftarrow \{v \in V \setminus (\text{completed} \cup \text{active}) : \text{pred}(v) \subseteq \text{completed}\}$
  available $\leftarrow \{v \in \text{eligible} : \boldsymbol{r}_v \leq \boldsymbol{R} - \boldsymbol{R}_{\text{used}}(t)\}$
  Order available by activity index (FCFS heuristic)
  for each $v \in$ available **do**
    Schedule v at time t, add to active
    $E_S \leftarrow E_S \cup \{(u,v)\}$ where u recently released resources acquired by v
    Update $\boldsymbol{R}_{\text{used}}(t)$
  end for
  $t \leftarrow \min_{v \in \text{active}}\{\text{finish\_time}(v)\}$ {Jump to next resource release event}
  Move finished activities from active to completed
end while
Output: Makespan $C_{\max}$, realized topology $E = E_H \cup E_S$

B. Defense of the FCFS Heuristic

A standard critique in scheduling literature is that topological volatility might be a mere artifact of employing an elementary priority rule like First-Come-First-Serve (FCFS). One might hypothesize that sophisticated dynamic dispatching heuristics (e.g., Minimum Slack, Latest Start Time) could suppress this instability. We formalize the counter-argument:

Theorem 1: For any deterministic priority rule π applied to stochastic durations $\{d_i \sim \mathcal{F}_i\}$ where $\mathbb{V}[d_i] > 0$ for some i, if $|E_S| > 0$, then $\text{SEVI}_\pi > 0$.

This assumes a continuous duration distribution where no two activities have identical completion times with nonzero probability (i.e., eligibility times are not tied).

Proof sketch: Stochastic durations induce stochastic eligibility times. When multiple activities become simultaneously eligible for a constrained resource, the priority rule π must adjudicate. Different duration realizations produce different eligibility orderings, hence different adjudications. Since E_S encodes precisely these adjudications, E_S must vary across iterations. □

Consequently, FCFS serves as a robust baseline indicator for SEVI. Advanced heuristics do not necessarily suppress this volatility; rather, by dynamically altering resource allocations based on transient network states, they introduce additional sensitivity to duration realizations.

C. Scalability Analysis

The vectorized SGS achieves time complexity:

$$\mathcal{O}(N_I \cdot (|V| + |E|)) \tag{5}$$

for N_I Monte Carlo iterations. This linear scaling permits enterprise-scale diagnostic processing: a 10,000-activity network with 10,000 iterations completes in seconds on commodity hardware, enabling SEVI as a pre-flight diagnostic before committing to exhaustive simulation campaigns.

D. Procedurally Generated Test Networks

Test DAGs were generated with controlled topological parameters:

- Network size: $N \in \{50, 100, 200, 500, 1000, 3000\}$ activities
- Edge probability: $p_e = 1.5/N$ (sparse connectivity)
- Resource types: $m = 3$ with uniform availability $R_k = \lfloor N/10 \rfloor$
- Duration distributions: $d_i \sim$ Triangular(a_i, b_i, c_i) with $\mathbb{E}[d_i] \in [5, 20]$

IV. Results and Discussion

A. Empirical Findings

Table I presents results from 10,000 Monte Carlo iterations per network configuration.

TABLE I
ENSEMBLE SEVI AND DENSITY (30 REPLICATES PER N)

N	Base	Mean	P90	P90$_\Delta$	SEVI%	$\bar{\rho}$	$\pm\sigma$
50	144.2	166.8	177.3	33.1	17.6	.0447	.0008
100	216.7	243.7	256.2	39.5	28.5	.0224	.0002
200	328.1	359.9	374.7	46.6	31.8	.0150	.0001
500	525.6	564.2	582.1	56.5	63.6	.0148	.0001
1000	736.5	782.0	802.2	65.7	89.1	.0148	.0000
3000	1303.8	1360.7	1386.2	82.4	99.9	.0148	.0000

The empirical data reveals a highly non-monotonic scaling behavior. Specifically, between $N = 100$ and $N = 200$, ...the SEVI growth rate stalls (increasing only from 28.5% to 31.8%), coinciding with a significant drop in Resource-to-Node Density ($\bar{\rho}$) from 0.0224 to 0.0150. This *Density Anomaly* confirms that structural volatility is not a linear function of network size, but a dynamic response to the ratio of resource constraints to topological width. Once $\bar{\rho}$ stabilizes (≈ 0.0148), the SEVI converges to unity, marking the transition to total topological incoherence.

Crucially, the data reveals that ρ does not need to approach unity to induce systemic volatility. At $N = 3000$, a stabilized $\bar{\rho} \approx 0.0148$ still yields dozens of absolute soft edges. The combinatorial explosion of these edges ($2^{|E_S|}$) ensures that the probability of realizing the baseline topology vanishes, driving SEVI to 99.9%. Volatility is driven by absolute scale compounding over sparse density, not just high density.

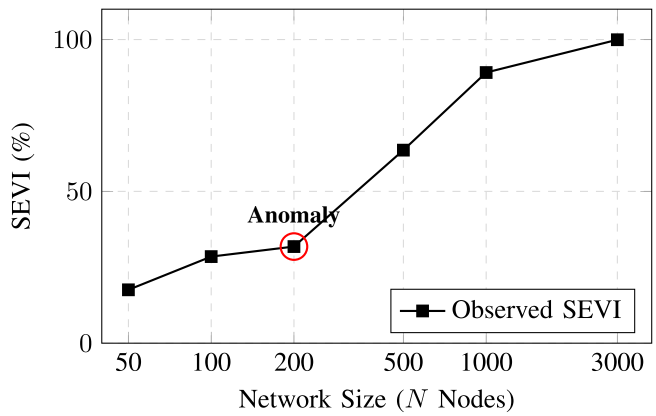

Fig. 1. Soft-Edge Volatility Index (SEVI) as a function of network size (N). The logarithmic x-axis reveals a non-monotonic deceleration at $N = 200$, where a localized reduction in Resource-to-Node Density ($\bar{\rho} \approx 0.015$) stabilizes the topology before the onset of combinatorial explosion at $N \geq 500$.

B. The Control Systems Interpretation

We reframe the Monte Carlo heuristic through control systems theory. Consider:

- **Plant:** The project network with stochastic durations
- **Controller:** The FCFS priority rule with resource constraints
- **State:** The evolving precedence topology $E^{(i)}$
- **Output:** Project makespan $C^{(i)}_{\max}$

As $\rho \to 1$ (see Eq. (2)), the controller faces increasing "noise" in the form of duration variance. The system fails to converge on a stable topology, instead exhibiting a phenomenon we adapt from control theory, termed *recurrent topological oscillation*: violent limit cycling of soft edges between precedence orientations. This is not stochastic exploration of a fixed state-space–it is bifurcation between distinct dominant topological states.

The P90 forecast, under these conditions, represents an ill-posed query: it asks for the 90th percentile of a distribution that does not exist as a unimodal entity. The "Ghost P90" phenomenon (82-day discrepancy at $N = 3000$) is the statistical signature of topological bifurcation.

C. Statistical Evidence: Bimodality and the Control Experiment

Application of Hartigan's Dip Test [9] to the makespan distribution at $N = 3000$ yielded:

$$D = 0.0847, \quad p < 0.01 \tag{6}$$

decisively rejecting unimodality.

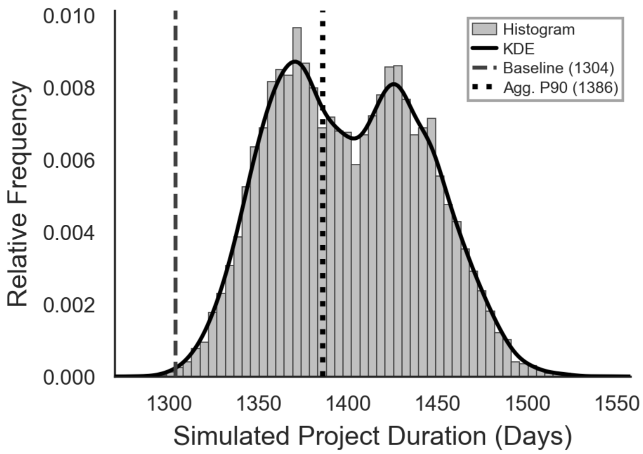

Fig. 2. The "Ghost P90" phenomenon at $N = 3000$. The multimodal probability density, derived from 10,000 simulation iterations, represents distinct topological state-spaces. The aggregated P90 metric (1386 days) falls in a transitional zone between the baseline (1304 days) and the bifurcated peak, obfuscating 82 days of structural delay.

Critical validation: A control experiment was conducted wherein the topological structure was frozen at baseline $E^{(0)}$ across all iterations, permitting only durations to vary. Under this constraint:

$$D_{\text{frozen}} = 0.0112, \quad p = 0.87 \tag{7}$$

The frozen-topology distribution failed to reject unimodality. This experiment isolates the causal mechanism: bimodality arises from topological bifurcation, not from asymmetry in the input duration distributions.

D. Implications for Risk Quantification

Consider a project manager presented with:

- Baseline: 1304 days
- Simulated P90: 1386 days

Standard interpretation: "There is a 10% chance the project exceeds 1386 days due to duration uncertainty."

Correct interpretation: "The topological structure of this project bifurcates into at least two distinct execution modes. The P90 aggregates across these modes, obscuring the fact that we cannot identify which topology will manifest. The 82-day delta represents structural uncertainty, not parametric uncertainty."

This distinction is not academic. Mitigation strategies for duration variance (schedule buffers, resource augmentation) differ fundamentally from mitigation strategies for topological volatility (logic restructuring, resource de-confliction).

E. Theoretical Limitations of Deterministic Bounding

A natural objection arises: if soft-edge volatility corrupts Monte Carlo forecasts, why not eliminate volatility by freezing the topology? This approach–naive deterministic bounding–fixes E_S at the baseline configuration and simulates only duration uncertainty. We demonstrate that this approach exhibits critical boundary limitations under high resource density.

Consider the topological state space $\mathcal{T} = \{T_1, T_2, \ldots, T_K\}$ where each T_k represents a distinct soft-edge configuration. Under stochastic durations, the realized topology is itself a random variable:

$$\boldsymbol{T} \sim \sum_{k=1}^{K} w_k \delta_{T_k} \tag{8}$$

where $w_k = \mathbb{P}[T_k \text{ manifests}]$ denotes the mixture weight, and δ_{T_k} is the Dirac delta measure concentrated on topology T_k. Naive deterministic bounding assumes $w_1 = 1$, $w_{k \neq 1} = 0$–that is, the baseline topology is the only topology.

This assumption is falsified by the empirical SEVI distribution. When SEVI $\approx 100\%$, soft edges flip direction with near-equal probability. The mixture weights $\{w_k\}$ are themselves stochastic, determined by the convolution of:

- Duration realizations across all resource-competing activities
- Priority rule adjudication under resource constraints
- The combinatorial explosion of feasible soft-edge configurations

Deterministic bounding does not "conservatively bound" risk–it estimates risk for a topology that may not manifest. The resulting forecast is not a bound but a conditional estimate conditioned on an event of probability $w_1 \ll 1$.

Theorem 2: Let $\mathcal{F}_{\text{det}}(x)$ denote the CDF under deterministic bounding and $\mathcal{F}_{\text{true}}(x)$ the true CDF integrating over topological uncertainty. For any threshold τ:

$$\mathcal{F}_{\text{det}}(\tau) \neq \mathcal{F}_{\text{true}}(\tau) \quad \text{whenever} \quad \exists k : w_k > 0, T_k \neq T_1 \tag{9}$$

Deterministic bounding provides no valid statistical guarantee. Because exhaustive topological enumeration is computationally intractable at scale (approaching $\mathcal{O}(2^{|E_S|})$), the only rigorous path forward requires identifying dominant sub-topologies via stochastic branching algorithms or topological clustering, replacing single-path Monte Carlo with mixture-weight estimation.

V. Conclusion

The central contribution of this paper is the demonstration that Monte Carlo simulation of resource-constrained schedules implicitly assumes topological continuity–an assumption violated with probability approaching unity as resource density increases. The Soft-Edge Volatility Index provides a quantitative diagnostic:

$$\text{SEVI} \to 1 \implies \text{Topological invariance assumption invalid} \tag{10}$$

When SEVI indicates high volatility, the application of continuous statistical metrics (mean, variance, percentiles) to what is fundamentally a discrete state-space introduces critical boundary limitations. Consequently, percentile forecasts computed from these bimodal distributions are mathematically misleading when interpreted as unimodal quantiles: they do not characterize the distribution of a well-defined random variable, but rather conflate distributions from distinct topological regimes.

Recommendations:

1) SEVI must be computed as a mandatory pre-flight diagnostic before any Monte Carlo simulation campaign.
2) When SEVI exceeds a critical threshold (empirically, $\approx 50\%$), simulation must be conducted on a per-topology basis with explicit topology enumeration.
3) Risk reports must distinguish parametric uncertainty (duration variance within a topology) from structural uncertainty (variance across topologies).

The enterprise schedule risk management community has historically interpreted simulation outputs under assumptions that exhibit critical vulnerabilities in topologically volatile networks. While existing robust scheduling literature focuses heavily on proactive buffering against parametric uncertainty, it largely preserves the Topological Continuity Assumption. SEVI is designed specifically to quantify structural degradation. SEVI provides both the diagnostic to detect this failure mode and the conceptual framework to address it.

References

[1] W. Herroelen and R. Leus, "Project scheduling under uncertainty: Survey and research potentials," *European Journal of Operational Research*, vol. 165, no. 2, pp. 289–306, 2005.

[2] J. Blazewicz, J.K. Lenstra, and A.H.G. Rinnooy Kan, "Scheduling subject to resource constraints: classification and complexity," *Discrete Applied Mathematics*, vol. 5, no. 1, pp. 11–24, 1983.

[3] E. Demeulemeester and W. Herroelen, *Project Scheduling: A Research Handbook*. Dordrecht: Springer, 2002.

[4] J.H. Patterson, "A comparison of exact approaches for solving the multiple constrained resource project scheduling problem," *Management Science*, vol. 30, no. 7, pp. 854–867, 1984.

[5] B.W. Taylor and L.J. Moore, "R&D project scheduling with generalized activity precedence relations," *Management Science*, vol. 26, no. 4, pp. 444–452, 1980.

[6] N. Damak, B. Jarboui, and T. Loukil, "Non-dominated Sorting Genetic Algorithm-II to solve bi-objective multi-mode resource-constrained project scheduling problem," *2013 International Conference on Control, Decision and Information Technologies (CoDIT)*, Hammamet, Tunisia, 2013, pp. 1–6.

[7] C.G. Cassandras and S. Lafortune, *Introduction to Discrete Event Systems*, 2nd ed. Springer, 2008.

[8] L. Arnold, *Random Dynamical Systems*. Springer Monographs in Mathematics, Berlin, Heidelberg, 1998.

[9] J.A. Hartigan and P.M. Hartigan, "The dip test of unimodality," *The Annals of Statistics*, vol. 13, no. 1, pp. 70–84, 1985.