

A Finite-Element Analogy for Distributed Computing Resilience: Predictive, Non-Invasive Resiliency Engineering Beyond Chaos Testing

Anand Sunder

Capgemini technology solutions Limited

Abstract- We present a predictive, telemetry-compatible finite-element (FEA) analogy for distributed computing resilience. Traffic is modeled as load vectors, latency/error/saturation as strain components, and capacity & coupling as a stiffness matrix. We derive node and system resilience scores, a von-Mises-style fragility metric, closed-form critical-load predictions, and cascade propagation conditions. This revision addresses reviewer requests: explicit problem statement, a concise state-of-the-art section, narrative bridging before mathematics, telemetry-based parameter estimation, a fully worked 4-node numeric example with embedded TikZ plots, and rigorous proofs (critical load, modal fragility, cascade).

Keywords - Distributed Computing, Resilience Metrics, Finite Element Analysis (FEA) Analogy, Telemetry-Based Modeling, Network Traffic Modeling.

I. INTRODUCTION

Distributed systems—microservices, serverless platforms, hybrid cloud—must tolerate load surges and partial failures. Chaos engineering has advanced resilience but remains reactive [1, 2].

Problem Statement. Current practice lacks a predictive, non-invasive framework to identify fragile components and modal failure channels without injecting failures. This paper proposes a finite-element analogy that integrates telemetry with matrix mechanics to compute resilience metrics proactively.

State of the Art (SOTA)

Chaos engineering and SRE (SLIs, SLOs, error budgets) are dominant approaches but require experimentation. Analytical methods (queueing, reliability block diagrams, cascading models [3, 4]) provide insights but lack a unifying predictive model calibrated from telemetry. Our work fills this gap.

FEA Analogy and Telemetry-Compatible Model Equilibrium model

We model response u and load f as $Ku = f$, $K = K(0) + \text{diag}(\mu_i)$. Here μ_i is node capacity, and $K(0)$ encodes inter-service couplings.

Telemetry estimation

k_{ij}
 $A_{ij} + A_{ji}$
 $= 2(1 + RA_{ij})$
 $\mu_i = \alpha_{\text{cpu}} \text{CPU}_i + \alpha_{\text{mem}} \text{Mem}_i$,
 CPU_{norm}
 $f_{0,i} = p_{95}(\text{requests})_i$,
 Mem_{norm}
 $\sigma_{y,i} = \phi(\text{SLA}_i, \text{error budget}_i)$.

Resilience and fragility metrics

Node resilience:

System resilience: Von-Mises fragility:

$R_i = w_L(1 - \epsilon_{L,i}) + w_E(1 - \epsilon_{E,i}) + w_S(1 - \epsilon_{S,i})$,

$$R_{sys} = \min R_i.$$

$$I_{\sigma m,i} = q \frac{1}{(\epsilon_L - \epsilon_E)^2 + (\epsilon_E - \epsilon_S)^2 + (\epsilon_S - \epsilon_L)^2}.$$

Theorems and Proofs

[Critical Load] For scaled load $f(s) = sf_0$, the response is $u(s) = su_0$. Fragility scales linearly:

$$\sigma_{m,i}(s) = \sigma(0)$$

. Critical load factor:

$$s_{\sigma,i}$$

$$= \min$$

crit

$$i(0)$$

$$v_{m,i}$$

[Modal Fragility] With $K = V \Lambda V^T$, response decomposes as

$$u = \sum_j \lambda_j^{-1} (v_j^T f) v_j.$$

$$j$$

$$j$$

Modes with small λ_j and large projection $v_j^T f$ dominate fragility.

[Cascade Condition (sketch)] If node j yields and its load redistributes via $R(j)$, then any neighbor k with $\sigma_{m,k} + \Delta \sigma_{m,k} \geq \sigma_{y,k}$ will fail next. Fragility amplifies if $(v_{TR}(j)^T f) / \lambda_j$ is large.

Worked 4-Node Example

We illustrate with nodes: Frontend, API, Worker, DB.

Baseline matrices

Coupling from telemetry:

$$\begin{bmatrix} 0 & 50 & 20 & 0 \\ 100 & 30 & 5 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad \mu = [100, 150, 80, 200].$$

$$A = \begin{bmatrix} 100 & 30 & 5 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad \mu = [100, 150, 80, 200].$$

$$0 \quad 0 \quad 0 \quad 0$$

Assembled stiffness:

$$\begin{bmatrix} 140 & -30 & -10 & 0 \\ -30 & 200 & -17.5 & -2.5 \\ -10 & -17.5 & 127.5 & -20 \\ 0 & -2.5 & -20 & 222.5 \end{bmatrix}.$$

Baseline load:

$$\begin{bmatrix} -30 & 200 & -17.5 & -2.5 \\ -10 & -17.5 & 127.5 & -20 \\ 0 & -2.5 & -20 & 222.5 \end{bmatrix}.$$

$$K = \begin{bmatrix} -30 & 200 & -17.5 & -2.5 \\ -10 & -17.5 & 127.5 & -20 \\ 0 & -2.5 & -20 & 222.5 \end{bmatrix}.$$

$$0 \quad -2.5 \quad -20 \quad 222.5$$

$$f_0 = [200, 50, 20, 5]^T.$$

Response and fragility

Solving $Ku = f_0$ yields

Fragility:

$$u_0 \approx [1.565, 0.517, 0.360, 0.061]^T.$$


$$\sigma(0) \approx [0.977, 0.323, 0.225, 0.038]^T.$$

With thresholds $\sigma_y = [1.5, 0.8, 0.5, 0.4]$, critical scales:

$$s_{crit,i} \approx [1.53, 2.48, 2.22, 10.56],$$

global $s_{crit} \approx 1.53$ (frontend yields first).

5TikZ plots



Scale s

Figure 1: Maximum fragility growth vs load scaling. Threshold crossing ≈ 1.53 signals first yield.

II. DISCUSSION AND CONCLUSION

This framework predicts fragility from telemetry without injecting failures. It provides inter- pretable metrics (resilience scores, fragility, scrit, modal channels) and guidance for mitigation (capacity increase, rerouting). Limitations: linear approximation, telemetry noise. Future work: nonlinear extensions, stochastic calibration, hybrid validation with chaos testing.




Figure 2: Eigenvalues (bars, left axis) and modal amplifications (red line, right axis).

REFERENCES

1. A. Basiri et al., "Chaos Engineering: Building Confidence in System Behavior through Controlled Experiments," IEEE Software, 2016.

2. N. Beyer et al., Site Reliability Engineering, O'Reilly, 2016.
3. A. E. Motter and Y.-C. Lai, "Cascade-based attacks on complex networks," Phys. Rev. E, 2002.
4. I. Dobson et al., "Complex systems analysis of series of blackouts," Chaos, 2007.