Inference-Conditioned Stochastic N-k Interdiction: Extending Fixed-Scenario Security Analysis with TERRA

Ricardo Marquez, Ph.D. rmarquez123@gmail.com

***Abstract*—We show that TERRA reproduces the stochastic N-k interdiction benchmark of Sundar *et al.* on RTS-GMLC *exactly*, and then generalizes the *workflow*: it treats the uncertain system state as an inferential object — inferred, conditioned on evidence, and updated — rather than a fixed external scenario file. Sundar's objective is recovered as the special case in which that distribution is pinned to his empirical 200-scenario measure; through the identical certified-expectation code path we reproduce all ten published optima to $3.5e-05$ MW. We then let evidence move the decision: conditioning on a severe, correlated storm over a two-tie corridor relocates the optimal interdiction out of the largest generator `g74@121` and *into* the islanded area (`g67@318`), a *topological* reordering absent from capacity ranking and from Sundar's published fixed empirical measure — and re-derived at *zero* incremental physical solves. Where the reliability tail is enumerable (an Area-1 subsystem), we additionally issue a *deterministic* certified interval and an unconditional optimality certificate at k=1. The distinction we claim is at the pipeline level, not one of formulation: a fixed-sample program *consumes* a scenario distribution; TERRA infers, conditions, and updates it.**

I. Introduction

Stochastic N-k interdiction seeks the k component removals that maximize expected load shed under a distribution of background contingencies,

$$A^\star = \arg\max_{|A|\le k} \; \mathbb{E}_Z\big[Q(A,Z)\big], \tag{1}$$

where $Q(A,Z)$ is the DC minimum-load-shed response to the combined outage $A \cup Z$. It is the adversarial counterpart of resource-adequacy analysis: rather than asking which contingencies are probable, it asks which deliberate removals are most damaging under a backdrop of random failures. Sundar *et al.* fix Z to a sample of M=200 equally weighted background-outage scenarios and solve (1) as an extensive-form program.

A companion Phase-1 study reproduced that benchmark exactly and established a sharp negative result: on this operating point the outer objective is effectively a knapsack in aggregate reserve margin, so single-component ranking already recovers the published optima and a smarter combinatorial search cannot win. The natural response is not to search harder but to change what is *modelled*. This paper does so on the axis where TERRA is distinctive — the inner expectation and its *inference* — rather than the outer search.

A. Thesis, stated carefully

We do *not* claim a weighted-scenario stochastic program is mathematically incapable of $\sum_s p_s\, Q(A,s)$; it plainly is, and one could update the weights and re-solve. The defensible distinction is at the level of the *decision pipeline* (Fig. 1):

> *Sundar's objective is a special case of the TERRA pipeline $E \to P(Z \mid E) \to A^\star(E)$; the fixed-sample workflow* consumes *a scenario distribution, whereas TERRA infers, conditions, and updates it.*

Throughout, "the fixed-sample method does not do X" means the published Sundar workflow does not — never that no mixed-integer program could.

B. Related work

Deterministic N-k interdiction is classically a bilevel max–min solved by duality or Benders decomposition; the stochastic variant averages the inner value over a scenario ensemble and is typically attacked with sample-average approximation and Wood-type covering cuts. Those cuts are *supervalid*: they can cut optimal integer points, so the resulting optimality guarantee is conditional on assumed dual bounds. Our certified variant instead brackets the inner expectation with a *deterministic* truncation interval from probability-ordered state enumeration, and prunes only on genuine two-sided bounds. The regime layer is a discrete common-cause reliability model conditioned by Bayes' rule; the novelty here is not the model class but its placement *inside* the interdiction pipeline, so that evidence re-derives the decision without re-solving the physics.

II. Experiment A: reduction is exact

Setting $Z \in \{s_1,\dots,s_{200}\}$ with $P(Z{=}s_m){=}1/M$ and residual mass 0, the general certified-expectation operator $J(A) = \sum_i p_i\, Q(A \cup z_i)$ collapses to Sundar's sample average, with a *zero-width* certified interval. Routed through the *same* plugin path Experiments B and C

Companion to study directory `inference_conditioned_nk_interdiction`, Phase 2 of the Sundar N-k program. Every number is read from committed, runtime-asserted artifacts. Evaluation counts are the portable quantity; wall-clock is one workstation's.

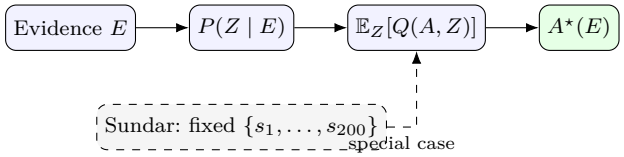


Fig. 1. The TERRA decision pipeline. Sundar's fixed 200-scenario sample enters as the special case in which $P(Z \mid E)$ is pinned to the empirical measure (Experiment A); relaxing that pin lets evidence move the decision (Experiments B, C).

TABLE I
EXPERIMENT A — TEN PUBLISHED OBJECTIVES z_k THROUGH THE GENERALIZED CERTIFIED-EXPECTATION PATH (INTERVAL WIDTH 0, EXACT).

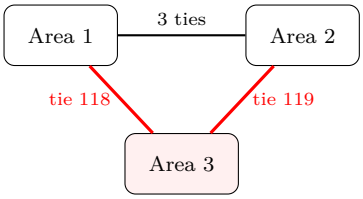
k	z_k published (MW)	z_k TERRA path (MW)	$\lvert\Delta\rvert$ (MW)
1	151.8247	151.8247	3.5e-05
2	481.3227	481.3227	1.5e-05
3	835.3350	835.3350	9.1e-13
4	1190.3350	1190.3350	4.5e-13
5	1545.3350	1545.3350	2.0e-12
6	1900.3350	1900.3350	3.0e-12
7	2255.3350	2255.3350	6.4e-12
8	2605.3350	2605.3350	2.3e-12
9	2942.5850	2942.5850	2.3e-12
10	3274.5100	3274.5100	3.2e-12

use — not a bespoke re-implementation — it reproduces every published optimum (Table I) to the rounding of the reference files, and the search's own argmax over all 216 singletons independently recovers the published A_1. This is the "equivalent to Sundar" half: same network, same scenarios, same inner problem, same numbers, obtained as a special case of the more general operator.

III. Experiment C: evidence moves the decision

We model uncertainty as a discrete regime $W \in \{\text{calm}, \text{fuel}, \text{storm}\}$; evidence updates $P(W \mid E)$ by Bayes. Each regime's objective is the finite empirical config-sum $J_w(c) = \sum_{\text{cfg}} P(\text{cfg} \mid w) \frac{1}{200} \sum_s Q(\{c\} \cup s \cup \text{cfg})$ — evaluated directly over the 200-scenario measure, never enumerated over the full fleet. The storm regime is an explicit *bivariate common mode* over the two Area-3 corridor ties, with $P(\text{both}) = m^2 + \rho\, m(1-m)$; at m=0.3, ρ=0.5 the ties fail together with probability only 0.195. Positive ρ supplies within-regime tie dependence that a product-Bernoulli model ($P(\text{both})$=m^2) cannot represent; it is the physical common mode of one weather system tripping both ties together. Because the objective is the finite empirical sum, the regime 'q' is unconstrained and the tail-mass wall that blocks full-fleet enumeration (Section IV) never arises here.

The corridor and the reason it matters are topological (Fig. 2): Area 3 is a razor-balanced net exporter joined to the rest of RTS-GMLC by exactly two ties. That two-tie boundary is asserted from the case data at run time, and it is what makes the area's *internal* generators — worthless to attack while imports flow — decisive the moment a storm islands them.



load 2850 ≈ gen 2875 (**+25 MW export**)

Fig. 2. Area 3 is a razor-balanced net exporter joined to the grid by exactly two ties. Degrading one leaves the optimum unmoved; the correlated *double*-tie event islands the area, turning an in-area 355 MW removal into ∼330–350 MW of unavoidable local shed.

TABLE II
EXPERIMENT C — STORM-SEVERITY BOUNDARY (ρ=0.5): THE WITHIN-STORM OPTIMUM RELOCATES INTO AREA 3 ONLY ABOVE $m \approx 0.25$.

m	P(both)	within-storm optimum
0.15	0.086	Area 1 (hold)
0.20	0.120	Area 1 (hold)
0.25	0.156	Area 3
0.30	0.195	Area 3
0.40	0.280	Area 3
0.50	0.375	Area 3

The move. Under calm the optimal single interdiction is the largest unit `g74@121` (expected shed 151.8 MW). As evidence shifts weight onto the storm regime, the optimum relocates *into* the razor-balanced, now-islanded Area 3 — the within-storm expected shed rises to 170.8 MW at a *smaller* 355 MW unit. Crucially, the per-regime objectives are cached on the union outage state, so in our implementation, once the finite regime basis has been evaluated (a one-time 39,314-solve warm-up), each evidence-conditioned re-derivation costs *zero* incremental physical solves. Re-running Sundar's published extensive-form workflow under new evidence, by contrast, re-samples and re-solves; that is a property of the workflow, not a theorem about all stochastic programs.

A decision boundary, not a hair-trigger. The relocation is confined to a severe *and* well-evidenced storm (Table II): the within-storm optimum enters Area 3 only for marginal tie rate $m \gtrsim 0.25$, and the evidence-conditioned optimum flips only near $P(\text{storm} \mid E) \approx 0.95$, because the within-storm margin is thin. A storm that degrades a *single* tie leaves the optimum unchanged (`g74@121`): the surviving ∼500 MW tie carries the ∼330 MW post-interdiction deficit. Only the correlated double-tie event islands the area. That physics-shaped boundary is a credibility feature.

Honesty. The probabilistic winner `g67@318` is itself a member of Sundar's published A_3/A_4; the non-circularity of the finding rests on its being *sub-maximal in capacity*, Area-3-local, and absent from the calm optimum — a topological relocation, not the recovery of a known pair. The *deterministic* double-tie limit (expected shed 350.3 MW, an upper bound on the probabilistic 170.8 MW) is exactly degenerate across five in-area units

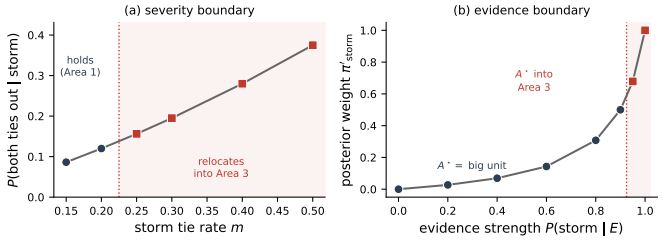


Fig. 3. The two-axis decision boundary (points read from `decision_boundary.json` and `evidence_sweep.json`). (a) The within-storm optimum relocates into Area 3 only above tie rate $m \approx 0.25$. (b) The evidence-conditioned optimum flips only at high posterior weight on the storm regime. Squares: optimum in Area 3; circles: optimum held at the largest unit.

TABLE III
EXPERIMENT B — AREA-1 ENUMERABILITY TO RESIDUAL MASS ϵ.

ϵ	states enumerated	cumulative mass
1e-02	9,714	0.9900
1e-03	111,873	0.9990
1e-04	829,001	0.9999

and additionally reaches units Sundar never selects at low budget (`g68@321`, `g72@323`).

IV. Experiment B: certified where the tail is enumerable

On an Area-1 subsystem the physical component-reliability model is enumerable (Table III). Running the exhaustive-leaf certified search with guard band ϵ_{sep}=2e−06 MW yields an *unconditional* certificate at k=1 and an honest UNCERTIFIED at k=2 on the flat objective (Table IV). The certificate is for the *certified-enumeration* objective, and it is issued only where the top intervals separate — neither Sundar method targets this object. Gate verdict: **GO**.

The k=1 certified margin (1.32 MW, i.e. $J^-(A^\star) - J^+(\text{runner-up})$) is six orders of magnitude above the cross-backend LP-objective agreement floor ($\sim 10^{-6}$ MW), so the verdict is solver-independent even though the certified run itself is executed on HiGHS.

V. Discussion: what this proves

Three results compose into one claim. Experiment A is a *reduction*: the generalized operator contains Sundar's objective exactly, so the comparison is a containment, not a re-derivation. Experiment C is the *generalization*: evidence conditioned on a physical uncertainty model moves the optimal interdiction, and moves it *topologically* — into an islanded area rather than up a capacity ranking — at zero incremental physical cost. Experiment B shows the inner value can be *certified*, not merely sampled, wherever the tail enumerates. Together they substantiate "equivalent to Sundar, and then some" in a precise sense: TERRA generalizes the fixed-scenario *workflow* at the decision-pipeline level, and Sundar's objective is the special case in which the state distribution is pinned to the empirical measure.

TABLE IV
EXPERIMENT B — CERTIFIED INTERDICTION (AREA-1). CERTIFICATION REQUIRES THE WINNER'S *LOWER* BOUND TO CLEAR THE RUNNER-UP'S *UPPER* BOUND; MARGIN $= J^-(A^\star) - J^+(\text{RUNNER-UP})$. k=2 IS UNCERTIFIED BECAUSE THE TOP INTERVALS *OVERLAP* (NEGATIVE MARGIN) — THE TWO LOWER BOUNDS DIFFERING IS NOT SUFFICIENT.

k	$A^\star$	$J^-(A^\star)$ (MW)	J^+(r.u.) (MW)	margin (MW)	certificate
1	74	427.4	426.1	+1.32	**certified**
2	20, 74	746.6	765.7	-19.04	UNCERTIFIED

The most satisfying part of the arc is its continuity with the Phase-1 negative result. Phase 1 showed that under *uniform* uncertainty the aggregate reserve margin dominates, the objective is knapsack-like, and ranking beats search — so the original search story failed. Phase 2 keeps the identical physics and asks a different question: does *structured evidence* expose latent structure that uniform weighting suppresses? It does. A nonuniform posterior over a physical common-mode regime activates the same super-additive corridor interaction Phase 1 measured but found inert, and the optimum reorders topologically. This is not "our first idea failed so we changed problems until TERRA won"; it is "the negative result told us *what* was suppressing structure, and we changed the uncertainty representation — not the physics — to test whether inference exposes it." The "then some" is that activation, not new physics.

VI. Threats to validity and scope

Enumerability. The certified variant is demonstrated where the inner expectation is the finite empirical measure (Experiments A, C) or an enumerable physical model (Experiment B, Area-1). On the full 96-generator fleet at realistic outage rates the certified enumeration is intractable ($\sim 10^{10}$ states); TERRA there falls back to sampling — the same tool Sundar uses — so the certification claim is explicitly bounded to the enumerable regime.

Rate dependence. The storm finding depends on modelled corridor rates (m, ρ) and a 10% storm prior, owned as constructs in `conventions.md`. The severity sweep (Table II, Fig. 3a) makes "no relocation at physically defensible rates" an admissible, reported outcome, not a failure to be tuned away. The relocation is confined to a corner of the $(m, P(\text{storm} \mid E))$ plane; we present that narrowness as a physical decision boundary, not as a headline generality.

Degeneracy and circularity. The probabilistic winner is a Sundar unit; the finding is stated as topological relocation into the area, and the non-Sundar units appear only in the deterministic double-tie limit. We avoid the stronger and false reading "TERRA finds a unit Sundar never selects."

Instance and physics. A single RTS-GMLC operating point and the DC min-shed model; no AC physics, unit commitment, or multi-period dynamics. The fungibility

that pins the argmax under uniform weighting is a property of this transmission-rich instance, and the topological effect requires the razor-balanced Area-3 corridor; other instances may differ in both directions.

VII. Reproduction

`PYTHONPATH=src python icn_run.py` writes the Experiment A/C artifacts; `icn_s6_gate.py` writes the Experiment B gate. Each driver asserts its load-bearing findings at runtime. The reusable operators (`certified_objective`, `certified_interdiction`, `re_interdict`, `corridor_config_probs`) carry unit tests under `tests/plugins/reliability/`. Run date: 2026-08-26.