Certified Region-Decomposed Contingency Inference: Local-to-Global TERRA on Power and Thermal Networks

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

***Abstract*—Probability-ordered explicit enumeration of the joint contingency space certifies its own coverage but does not scale: a 124-component test system needs more than 5×10^6 states to leave 0.019 of the probability mass un-enumerated. We ask whether that boundary belongs to enumeration or to structured inference. The network is partitioned into regions, each region is enumerated separately, and the regions are joined through physics-derived separators—the region's Ward equivalent and two governor statistics—so that every joint quantity is a *certified bracket* that includes the un-enumerated mass (TERRA-watt is the graph-resolved physics engine). Two regions reproduce exhaustive global enumeration to round-off and reach residual 0.020 with 22,672 local region evaluations and 1.0×10^6 scalar interface pairs (4.4$\times$ fewer region evaluations than a 10^5-state global budget that leaves 0.111); beyond two regions, where no enumeration exists, single-bus chains of six regions (10^{24} enumerated joint states) and the three RTS-GMLC areas as separate regions get certified per-branch and per-region overload probabilities. The program's own gate—does locality compress the separator?—is answered negatively on RTS-GMLC three-area, IEEE 118 and PEGASE 1354, and the mechanisms that replace it (certified boxing, lattice trees with a sheared coordinate, environment messages for any interconnection shape, a shared-box joint bracket, tie outages as a mixture) are measured against exact references. The decomposition wins on coverage; it is a wall-clock win only in the two-region single-bus case, and the joint probability of any overload at multi-bus interfaces among three or more regions remains a stated bound. The same code runs unchanged on a data-hall thermal network through a 38-line adapter, exhaustively equal to brute force.**

***Index Terms*—Contingency analysis, probabilistic risk, network decomposition, Ward equivalent, junction tree, certified bounds, interval arithmetic, reliability test system.**

I. Introduction

The companion study [1] located a boundary: probability-ordered explicit enumeration of the joint contingency space (the enumeration branch of reliability assessment [2], [3]) reproduces Monte Carlo estimates and certifies coverage of individually rare states, but the number of states needed for a certified residual ε grows exponentially with the component count (from 11,336 at 62 components to more than 5×10^6 at 124 for $\varepsilon = 10^{-2}$). The question this paper answers is whether that boundary belongs to *global explicit enumeration* or to structured probabilistic inference as such. The hypothesis, stated as an R&D direction for the TERRA engine, was that the inference partition should be aligned with the physical network partition: regions enumerated locally, coupled through separator variables that the physics defines, with messages passed between them—the pattern of local-to-global solvers for nonlinear network physics [4] and of exact inference on junction trees [5], [6]. The closest prior art is multi-area reliability evaluation by decomposition [7], [8], which partitions the state space of tie-line capacities; here the partition is of the network and the coupling is the full DC physics through Ward equivalents [9].

Two success criteria were fixed in advance: (i) local-to-global must *reproduce* global TERRA, and (ii) it must *improve scaling*. Two rules were fixed as well: no independence assumption between regions beyond what the physics gives, and every approximation term reported as a first-class number next to the result it widens—the difference $\mathrm{P}_{\text{partitioned}} - \mathrm{P}_{\text{global}}$ is never hidden. The program was gated: before any engineering, the compressibility of the separator on realistic grids was to be measured, and the plan pivots if the measurement is unfavorable. It was, and this paper is largely the account of the pivot.

Five research questions structure the paper: RQ1, does local-to-global reproduce global TERRA; RQ2, what coverage does it reach against a global budget; RQ3, does the separator compress on realistic grids (the gate); RQ4, what certified mechanisms work at multi-bus interfaces and beyond two regions; RQ5, does the probability layer transfer to a second physical domain unchanged.

Contributions

1) A separator for DC-network regions—the Ward equivalent [10], [11] plus two governor statistics—and exact two-region messages that reproduce exhaustive global enumeration to round-off (Section II, RQ1).
2) A measurement, on RTS-GMLC three-area, IEEE 118 and PEGASE 1354, that separator compressibility comes from fleet degeneracy and not from locality (RQ3), with the mechanism named.
3) Certified mechanisms for interfaces where value-grouped messages do not pay: boxing at $k \geq 2$,

Companion to the contingency-risk study TERRA-CS-PF-08 [1]; same test system, same reliability data. TERRA = TERRA-watt, the graph-resolved physics engine the studies run on; RTS-GMLC = the 2019 update of the IEEE Reliability Test System. Every number in this paper is read from committed, test-pinned result artifacts of the program's case studies (Appendix); wall-clock timings were measured on one shared workstation in Python/NumPy and vary by machine.

lattice trees for single-bus chains with a sheared coordinate, environment messages for any interconnection shape, a shared-box joint bracket, and ties as reliability components—each bracketing brute force or an exact reference in tests (RQ4).
4) A consolidated benchmark that states the crossover per metric: coverage is where the decomposition wins; wall-clock only for two single-bus regions (RQ2).
5) Domain independence demonstrated rather than argued: the probability layer runs unchanged on a data-hall thermal network (RQ5).

II. Method

A. Regions, Ward equivalents and the separator

The DC network [12] (buses, branches with susceptance $\beta_{ij} = 1/x_{ij}$ and rating, generators with dispatch and headroom, loads) is partitioned by bus into regions $r = 1, \dots, R$; branches with ends in two regions are *ties* and are fixtures (never outaged) unless stated (Section II-G). A region's outage state s_r (a subset of its own generators and internal branches) determines, by Kron/Ward reduction of the region's interior onto its boundary buses ∂r (dimension k_r), the reduced susceptance $B_{\mathrm{eq},r}(s_r)$ ($k_r \times k_r$; itself a singular Kron-reduced Laplacian) and the equivalent boundary injection. Generator outages are re-dispatched by the companion study's governor rule: with L_r the region's lost dispatch and H_r the headroom of its *surviving* thermal units, $\Sigma L = \sum_r L_r$ and $\Sigma H = \sum_r H_r$ over the whole system, the lost dispatch is shared by the surviving units in proportion to headroom, $\alpha = \min(\Sigma L, \Sigma H)/\Sigma H$ ($\alpha = 0$ when $\Sigma H = 0$), and any shortfall $D = \max(\Sigma L - \Sigma H, 0)$ is absorbed at the slack bus, so the interface injections always balance. Islanded joint states are classified before any flow is evaluated (below). Because α and D are global scalars, the equivalent injection of a region is affine in them, $p_{\mathrm{eq},r} = p_r^0 + \alpha\, p_r^h + D\, p_r^D$, and the same holds for every internal branch flow, $f_r = f_r^0 + \alpha f_r^h + D f_r^D + M_{\theta,r}\,\theta_{\partial r}$. The *separator* of a region state is therefore

$$S_r(s_r) = \big(\mathrm{isl}_r,\ \text{boundary components},\ L_r,\ H_r,\ p_r^0,\ p_r^h,\ p_r^D,\ B_{\mathrm{eq},r}\big),$$

with the MW quantities rounded to 10^{-6} MW and B_{eq} to 10^{-9} pu for hashing: everything the rest of the network needs to know about the state. Islanding is integer: a region state is locally islanded if an interior component loses its boundary; a joint state is islanded if the interface supergraph (boundary components as nodes, live ties as edges) is disconnected. Islanded joint states carry no overloads by convention, matching the companion study.

B. Exact two-region messages

For two regions the interface Laplacian $Y = T + \mathrm{blockdiag}(B_{\mathrm{eq},A}, B_{\mathrm{eq},B})$ (T the tie Laplacian) is singular; one interface node is pinned as the angle reference and dropped, and the reduced system is nonsingular whenever the interface is connected, giving the boundary angles $\theta = Y^{-1}(p_{\mathrm{eq},A} \oplus p_{\mathrm{eq},B})$, hence every flow and every Tier-1 outcome (Tier 1 is the companion study's pre-corrective outcome layer: overloaded branches, severity, islanding, deficit). Region A's outcomes depend on its own state exactly and on region B only through S_B, so region B's enumerated states are grouped by separator value: $|S_B|$ distinct values instead of $|\mathrm{enum}_B|$ states. Region A's message is built by evaluating each of its enumerated states against every value of S_B ($|\mathrm{enum}_A| \cdot |S_B|$ interface solves, each $k \times k$) and summing into a table $\mu_A[a, b]$ indexed by both separator values; symmetrically for B; the join is $\sum_{a,b} \mu_A[a, b]\, \mu_B[b, a]$ over $2|S_A| \cdot |S_B|$ table entries—the *interface pairs* reported throughout. Each region is enumerated in probability order to a residual ε_r (its un-enumerated mass); the joint residual is $1 - \prod_r (1 - \varepsilon_r)$, the mass of joint states in which at least one region is un-enumerated; and every certified interval is $[\mathrm{lower}, \mathrm{lower} + \mathrm{residual} \cdot \mathrm{cap}]$ with cap the largest value the quantity can take on an un-enumerated state (1 for a probability; the sum of the ratings for severity), exactly as in the companion study. Combinatorics are relocated, not eliminated: interface pairs are reported next to physics evaluations everywhere.

C. Certified boxing

When $|S_A| \cdot |S_B|$ is too large, the separator support is covered by boxes and each pair of boxes is evaluated once with interval arithmetic [13], [14]: the governor scalars range over intervals from the (L, H) box ranges; at $k \geq 2$ a box is a *topology class* (internal outage subset, which fixes B_{eq}, M_θ, p^D and the islanding structure—exact) times an injection box in (L, H, p^0, p^h); the interface Laplacian is exact for the class pair, so boundary angles and tie flows have exact interval hulls by sign-splitting the linear map. Each state's internal flows are kept in the affine form $f = c_0 + \alpha c_1 + D c_2$ with the box's own states exact and only the neighbour hulled, which captures the α–θ correlation that pure interval arithmetic loses. Every branch is definitely overloaded, definitely not, or uncertain; the mass of the box pairs whose flag is uncertain, $\sum_{(a,b)\ \text{uncertain}} m_A(a)\, m_B(b)$, is the *box slack*, reported, and refinement splits the boxes carrying it. The slack is the mass within one box width of a threshold, so it is at least first-order in the box width (on the $k = 1$ replica at $\varepsilon_r = 10^{-2}$ the first-pass slack is 2.8×10^{-3} / 3.8×10^{-4} / 6.3×10^{-5} at 200 / 50 / 20 MW).

D. Single-bus chains: lattice trees

With one boundary bus per region ($k = 1$) a region's export is its own imbalance and its internal flows depend on the rest only through α and D, i.e. through the *sums* ΣL, ΣH. Per-region histograms of (L, H) on a common lattice are convolved (by fast Fourier transform, FFT [15]) into the total lattice; each total box fixes the α/D ranges, so per-region definitely-ok / possibly-ok histograms are convolved for $\mathrm{P}(\text{no overload})$, leave-one-out convolutions give per-branch and severity brackets, prefix/suffix convolutions bracket the tie flows, and joint islanding is exact ($\mathrm{P}(\text{some region islanded}) = \mathrm{covered} - \prod_r(\mathrm{covered}_r - \mathrm{isl}_r)$, with $\mathrm{covered}_r = 1 - \varepsilon_r$ and $\mathrm{covered} = \prod_r \mathrm{covered}_r$ the mass of joint states in which every

region is enumerated). The cost is $R\times$(lattice convolution) per box, independent of $\prod_r |\text{enum}_r|$. The width is a lattice effect—a total cell spans $R\,\delta_L$ in L and $R\,\delta_H$ in H, an α window $\approx R(\delta_L + \alpha\delta_H)/\Sigma H$—so the width is first-order in the cell size. Because L is a subset sum over the lost units and H is the fleet's total headroom minus the subset sum over the same units, H is nearly affine in L when headroom is roughly proportional to dispatch; taking the second axis as $W = H + c\,L$ with a fitted c (shear) collapses the lattice towards one dimension and removes the δ_H part of the window.

E. Environment messages for any interconnection

For a target region i with boundary nodes I and the rest O, the Schur complement of the interface Laplacian gives

$$\theta_I = (B_{\text{eq},i} + B_{\text{env}})^{-1}(p_i + p_{\text{env}}),$$
$$B_{\text{env}} = T_{II} - T_{IO}Y_{OO}^{-1}T_{OI}, \qquad p_{\text{env}} = -T_{IO}Y_{OO}^{-1}p_O,$$

where the off-diagonal blocks are ties only (the reference node is dropped from whichever side owns it, so Y_{OO} is nonsingular whenever the interface is connected), so the *environment message* $(B_{\text{env}}, p_{\text{env}})$ depends on the other regions' classes and states and never on the target's, and $p_{\text{env}} = \sum_{j\neq i} G_j p_j$ is linear in their injections (affine in α, D). Nothing requires a tree: chains, trees and cycles are the same computation. The others' state tuples are binned per class tuple into environment cells with tight hulls (min/max over members) on $(L_{\text{env}}, H_{\text{env}}, p^0_{\text{env}}, p^h_{\text{env}})$; the target's exact states are evaluated against every cell with the affine core; tie flows $f_t = \beta_t(\theta_a - \theta_b)$ are affine in θ_I and in a far-side potential the cells carry. Per-region and per-branch probabilities and severity are certified brackets and islanding is exact through class tuples; the joint $\mathrm{P}(\text{any overload})$, lacking a shared coordinate, is bounded by the union bound $[\max_r \text{lower}_r,\ \min(\sum_r \text{upper}_r, \text{covered})]$ over region and tie events, whose extra width is reported as the *union gap*.

F. A shared coordinate for the joint event

Boxing every region's states (class $\times$ injection cell) and keying each target's environment cells by the neighbours' *box tuple* makes every region's definitely-ok / possibly-ok flag valid for every joint state of a joint box tuple $b = (b_1, \ldots, b_R)$. Writing b_{-r} for the tuple without r and $m_r^{\text{def}}(b_r \mid b_{-r})$ (m_r^{pos}) for the mass of region r's enumerated states in box b_r that are definitely (possibly) free of overload given the environment tuple b_{-r}, the intersection factorises, $\mathrm{P}(\text{all ok}) \in [\sum_b \prod_r m_r^{\text{def}}(b_r \mid b_{-r}), \sum_b \prod_r m_r^{\text{pos}}(b_r \mid b_{-r})]$—the product is legitimate because the regional priors are independent (no component belongs to two regions; common causes are conditioned on explicitly, as in the companion study)—and $\mathrm{P}(\text{any}) = \text{covered} - \mathrm{P}(\text{all ok})$ is a bracket. The price is up to n_{boxes}^{R-1} tuples per target (only occupied tuples are formed).

G. Ties as reliability components

Ties enter only through the interface geometry, so tie outages are a mixture: every subset τ of outaged ties (ties are few) is evaluated with the surviving ties in service—Laplacian, tie flows and interface connectivity on the live ties, so a lost interface islands exactly—and the certified intervals are mixed with $\mathrm{P}(\tau)$. A mixture of brackets is a bracket.

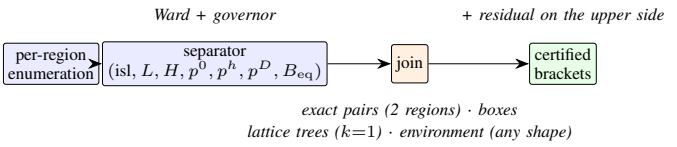


Fig. 1. The pipeline. Regions are enumerated separately in probability order; each state's separator is exact physics; the join is exact for two regions and certified (bracketed, slack reported) otherwise; every result carries the unenumerated mass.

H. Domain independence

The probability layer consumes, through a protocol, only region-local physics (separator + internal outcome operators per state) and the interface geometry. Any linear conservation network enters through an adapter onto the DC network object (38 lines of Python): a steady-state thermal network of data halls (racks with heat loads, cooling units with capacity and duty, paths with conductance and heat-flow capacity) maps exactly with $x = 1/G$, W↔MW, cooling unit↔generator, heat load↔load, path capacity↔rating, reference plenum↔slack, taking the cooling potential $-T$ as the angle so that flow magnitudes and limits are sign-independent.

III. Experimental design

Systems. (a) RTS-GMLC Area 1 [16] (24 buses, 62 reliability components) replicated $\times 2 \ldots \times 6$ and joined slack-to-slack by one tie per interface ($k = 1$; $x = 0.06$ pu, 500 MW) or by two parallel circuits between the same bus pair ($k = 2$); (a′) for the exhaustive check, the $\times 2$ replica with the reliability model restricted to 4 branches and 4 generators per region (16 components, 65,536 joint states) on the full 48-bus network; (b) the RTS-GMLC three-area system with the areas as regions (ties AB1–3, CA-1, CB-1: a cycle, $k = 4/4/2$; 188 components); (c) IEEE 118 and PEGASE 1354 for the compressibility gate; (d) a data-hall thermal chain (halls of racks around a return plenum, cooling units on coil nodes, wall paths between halls). *Reliability data.* The published RTS-GMLC branch and generator forced-outage parameters used by the companion study; IEEE 118 and PEGASE 1354 have no published reliability data and use RTS-96 class rates by assumption; ties are fixtures except in Section II-G. *Dispatch* (it matters: the same replica appears under three dispatches, named in every table row). The companion study's peak DC optimal-power-flow snapshot—a merit-order dispatch, "peak" in the tables—for the $\times 2$ replica in the two-region exact and boxing $k = 2$ rows and for the $\times 3$ chain in the shear check; the same snapshot with twin units perturbed apart ("de-twinned") for boxing $k = 1$, so that the separator support is not trivially degenerate; pro-rata to $P_{\max}$ at 0.6 load for the chains, the environment rows, the three-area system (also 0.9 for one boxing row) and the gate. *Partitions* for the gate: spectral bisection [17], spectral + Kernighan–Lin refinement

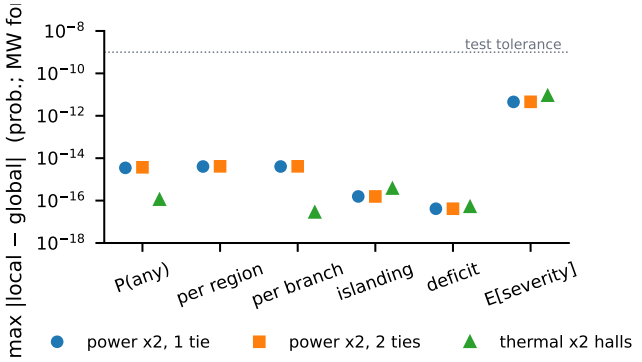


Fig. 2. Exactness (RQ1). Maximum absolute difference between local-to-global and the exhaustive reference (global TERRA on the 16-component $\times 2$ subset with one and two ties, brute force on the thermal two-hall system) per outcome type; every point is round-off, two to eight orders below the 10^{-9} test tolerance (severity in MW, the rest probabilities).

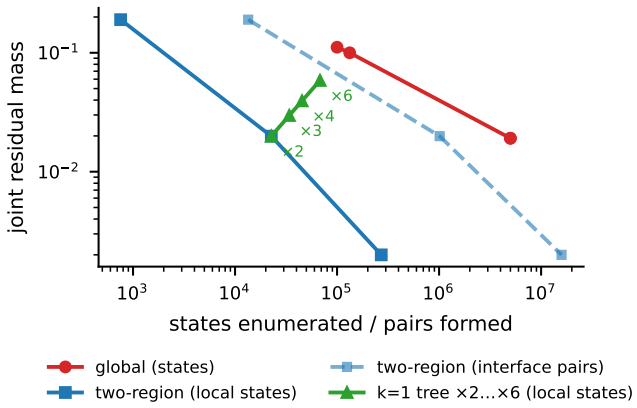


Fig. 3. Coverage (RQ2): un-enumerated mass against the states enumerated or pairs formed. Global enumeration (red) vs the two-region path (blue, local states and interface pairs) and the $k = 1$ tree on chains of $R = 2, 3, 4, 6$ regions (green; plotted against local states—its cost is lattice convolutions, see the width–cost plot of Section IV-F).

[18] (plain and cut-weighted) and random bisections per grid, plus the area partition on the three-area system. *References.* Brute force where the joint space is enumerable; the exhaustive two-region path on the *same* per-region enumerations where two regions apply (for a three-region system the two others are merged into one region of their product enumeration); a monolithic Monte Carlo of 20 000 draws [19] as statistical cross-check everywhere, reported with its 95 % Wilson interval [20] (half-width ≈ 0.0065 at $p = 0.68$, $\approx 2 \times 10^{-4}$ at $p = 2 \times 10^{-4}$). *Protocol.* One row per (path, system, R, k, ε_r): joint residual, the combinatorics that replace global states, the certified $\mathrm{P}(\text{any overload})$ bracket before and after the residual, reference, wall-clock.

IV. Results

A. RQ1: reproducing global TERRA

On the exhaustive 16-component two-region system (65,536 joint states) local-to-global equals global enumeration (a monolithic power-transfer / line-outage distribution-factor resolve per state [21]) on every branch, region, severity and deficit number to a maximum absolute difference of 4.5×10^{-12} (one tie) and 4.6×10^{-12} (two ties); the thermal two-hall system (4,096 states) equals brute force to 8.9×10^{-12} (Fig. 2). Under truncation on the full $\times 2$ replica the certified intervals overlap the global ones and the Monte Carlo 95 % interval [0.6724, 0.6853] at every ε_r (at $\varepsilon_r = 10^{-3}$: [0.6806, 0.6806]).

B. RQ2: coverage

Fig. 3 and Table II give the crossover. Global enumeration on the $\times 2$ replica leaves 0.111 of the mass un-enumerated at 10^5 states (24 s) and 0.019 at 5,000,000 states. The two-region path reaches 0.020 with 22,672 local states and 1.0×10^6 interface pairs in 19 s. Counting one region evaluation (Ward reduction plus internal operators) per enumerated local state, that is 22,672 region evaluations (21,916 new beyond the $\varepsilon_r = 10^{-1}$ tier) against 10^5 global state solves—4.4× fewer, with the 1.0×10^6 scalar interface pairs counted separately—and 221× fewer states than the 5×10^6 global run. It is also the one row that is faster than the global budget while covering more (the 5×10^6-state global run was not timed). At $\varepsilon_r = 10^{-3}$ the path reaches 0.002 with 271,032 local states (1.6×10^7 interface pairs, 646 s), which no global budget reached. Beyond two regions no enumeration exists to compare against: the $k = 1$ tree covers six regions (10^{24} enumerated joint states, the product of the per-region enumerations; the joint contingency space is 2^{372}) with residual 0.059, and environment messages cover the three RTS-GMLC areas separately with residual 0.030. The two-region cut of the same system that was run (Area 3 against Areas 1+2, whose 131-component complement was capped at 60,000 enumerated states because its own $\varepsilon_r = 10^{-2}$ enumeration would need more than 5×10^6) left 0.13. Coverage is the metric on which the decomposition wins outright.

C. RQ3: the gate

The plan's gate asked whether realistic partitions give a separator compression $C_r = |\text{enum}_r|/|S_r|$ of at least 10 at $k_r \leq 3$ on two of the grids. Fig. 4 and Table I. No candidate partition of IEEE 118 or PEGASE 1354 produces a region with $k_r \leq 3$ at all (their smallest boundaries are 8 and 13), so the threshold is unmeasurable there rather than failed. On the RTS grids the $k_r \leq 3$ regions are three small cuts of Area 1 (81–98 enumerated states at $\varepsilon_r = 10^{-2}$, two of them the same region under two partitioners; $C = 1.6$–1.8) and Area 3 of the three-area system (13 736 states, $C = 5.7$)—one substantive datum. Because C grows with $|\text{enum}_r|$, the small-region rows and the PEGASE rows ($\varepsilon_r = 10^{-1}$, 5 000-state cap) are low by construction and support the trend only. Over every partition of every grid C is never large (1–6) and shows no rise at small k; the inference-optimal partition of the three-area system forms 5.1×10^7 interface pairs at $\varepsilon_r = 10^{-2}$ in the gate run (6.4×10^7 in the boxing run of Table II, under its complement cap); the verdict is *unfavorable*. The mechanism: $|S| \approx$ (generator statistics) $\times$ (boundary-relevant topology classes); topology compresses, generator

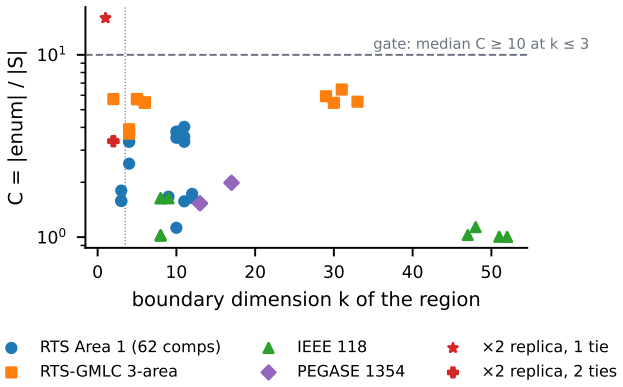


Fig. 4. The gate (RQ3): separator compression C_r against boundary dimension k_r over candidate partitions of realistic grids at $\varepsilon_r = 10^{-2}$ (10^{-1} on PEGASE). No grid clears the threshold; the $\times 2$ replica with one tie is the twin-unit outlier.

TABLE I
GATE PER GRID: MEDIAN C_r OVER REGIONS WITH $k_r \leq 3$ AND THE NUMBER OF SUCH REGIONS (RTS-GMLC AREA 1: TWO DISTINCT REGIONS, ONE OF THEM UNDER TWO PARTITIONERS).

grid	median C ($k \leq 3$)	regions with $k \leq 3$
IEEE 118	—	0
PEGASE 1354	—	0
RTS-GMLC Area 1	1.6	3
RTS-GMLC 3-area	5.7	1

statistics collapse only through twin units at identical MW—the compression the two-region feasibility study saw ($16\times$) is fleet degeneracy, not locality, and coarser separator binning is inert. The fitted separator law ranks candidate partitions like the measured cost on the three-area system (the area partition wins), but the boundary dimension alone ranks identically and the law's magnitudes are off by an order of magnitude at the median ($|\log_{10}(\text{predicted}/\text{measured})| = 1.05$ over 47 rows, 13 censored at the enumeration cap): rank by $\sum_r k_r^2$; do not use the law for magnitudes.

D. RQ4: mechanisms at multi-bus interfaces

Boxing. At $k = 1$, boxes replace 2.7×10^8 exact pairs by 4.1×10^5 at $\varepsilon_r = 10^{-3}$ with slack 7.4×10^{-6}—where the exact pair path was not run (hours of scalar pairs); at $\varepsilon_r = 10^{-2}$ the boxed ladder costs 375 s against 63 s exact. At $k = 2$ on the two-tie replica the boxed bracket has slack 1.0×10^{-5} at 1.3×10^6 pairs against 2.3×10^7 exact pairs; each pass costs about what the exact path costs, and the four-pass ladder 1120 s against 239 s: not a win (the second tie is a parallel circuit between the same bus pair; the 500 MW ties never bind, so the two-tie replica has the same P(any) to four decimals as the one-tie replica and exercises $k = 2$ in the mechanism, not a different reliability picture). On the RTS three-area system, Area 3 ($k = 2$) against Areas 1+2 ($k = 2$) has slack 4.4×10^{-6} at 0.6 load and brackets the exact value 0.1841 as [0.1793, 0.1888] at 0.9 load, but the complement's enumeration residual 0.13 dominates the certified width—the case for trees and environment messages. *Trees.* Fig. 5: on chains $R = 2, 3, 4, 6$ the plain lattice's slack floors near 2.0×10^{-2} and the sheared lattice (c =0.589; W constant under the pro-rata dispatch, the lattice one-dimensional) reaches 2.4×10^{-3} at $R = 6$ in 1243 s per ladder against 5090 s; at $R = 2$ the bracket [0.12065, 0.12273] contains the exact 0.12269. Under the peak (merit-order) dispatch on the $\times 3$ chain the fitted shear is zero and the slack stalls at 0.18: the shear is a property of the dispatch. *Environment messages.* Fig. 6: on the $\times 3$ two-tie replica at $\varepsilon_r = 10^{-1}$ the per-region brackets close to width zero at 2 MW cells and every ladder row brackets the exact two-region path (0.0443) on the same enumerations, but P(any) is the union bound [0.0405, 0.1216]—the regions overload together; at $\varepsilon_r = 10^{-2}$ only 50 MW cells were affordable (per-region slack 2.2×10^{-2}, union gap 2.3×10^{-1}, 1.6 h per row: explicit binning of 10^8 environment tuples). On the RTS three areas at $\varepsilon_r = 10^{-2}$, likewise at 50 MW cells, the per-area brackets are [0.0000, 0.0301], [0.0000, 0.0458], [0.0000, 0.0006] (per-area Monte Carlo 0.00015, 0.00000, 0.00000; P(any) 0.0002), i.e. per-area upper bounds that are box slack, and the union bound P(any) $\leq$ 0.079 adds a union gap 7.9×10^{-2} of the same size: informative only as an upper bound. The same areas at $\varepsilon_r = 10^{-1}$ with 2 MW cells resolve the covered mass to width $\sim 10^{-8}$ (Area 1 $[1.1 \times 10^{-6}, 1.1 \times 10^{-6}]$) but carry a joint residual 0.27, and the Monte Carlo Area-1 estimate $[5.1 \times 10^{-5}, 4.4 \times 10^{-4}]$ lies above the covered-mass bracket: almost all of Area 1's overload mass sits in the residual, so the tight bracket certifies little. Cell resolution and residual trade against each other on this system. *Joint bracket.* Fig. 7: the shared box coordinate turns the union bound into [0.0440, 0.0492] at 101 boxes per region and 104 s ($\varepsilon_r = 10^{-1}$; exact 0.0443), $16\times$ tighter; at $\varepsilon_r = 10^{-2}$ the affordable coarse row [0.046, 0.447] is wider than the union bound [0.084, 0.317] (both are valid, so the certified interval is their intersection [0.084, 0.317]—the union bound itself); the boxes fine enough to resolve the flags cost n^2_{boxes} tuples. *Ties.* On the $\times 3$ two-tie replica at ε_r =0.1 (covered mass 1−0.27) with q =0.01 per tie, the 16-subset mixture leaves P(any) unchanged at 500 MW ratings and raises islanding from 0 to 1.5×10^{-4} over the covered mass (both circuits of an interface must fail: $2q^2 = 2.0 \times 10^{-4}$ over the full mass), exact.

E. RQ5: the second domain

Nothing in the probability layer changed to run the thermal network. The independent thermal solve equals the mapped network's flows; two halls are exhaustively equal to brute force (RQ1); three halls with two wall paths per interface (28 components) give certified per-hall brackets (hall 1: [0.0111, 0.0111] vs Monte Carlo 0.0123) and an any-wall union bound [0.00053, 0.00212] against Monte Carlo 0.0021—the walls do overload when a cooling unit is lost, and it was this case that exposed that the environment path's first version did not evaluate tie flows (severity bracket violated), which was then fixed and re-run.

F. The consolidated protocol

Table II condenses the 32-row benchmark; Fig. 8 plots width against cost per mechanism.

TABLE II

THE PROTOCOL, CONDENSED: JOINT RESIDUAL, COMBINATORICS THAT REPLACE GLOBAL STATES, CERTIFIED P(ANY OVERLOAD) BEFORE THE RESIDUAL, REFERENCE (EXACT WHERE ONE EXISTS, ELSE MONTE CARLO), WALL-CLOCK OF THE LADDER / PASS TOTAL. INTERFACE PAIRS COUNT BOTH MESSAGE TABLES ($2|S_A||S_B|$); k IS THE LARGEST BOUNDARY DIMENSION; † MARKS A MONTE CARLO REFERENCE (20 000 DRAWS, WILSON 95 % HALF-WIDTH ≈ 0.0065 AT $p \approx 0.7$, ≈ 0.004 AT $p \approx 0.1$, $\approx 2 \times 10^{-4}$ AT $p \approx 2 \times 10^{-4}$), ALL OTHERS EXACT OVER THE COVERED MASS. FULL TABLE IN THE BENCHMARK ARTIFACT.

path	system, dispatch	R	k	ε_r	residual	combinatorics	P(any) before residual	reference	wall s
global enum., 10^5 states	×2, peak	1	1	—	0.111	states 100000	[0.5765, 0.5765]	0.6789†	24
two-region exact	×2, peak	2	1	0.01	0.020	pairs 1013888	[0.6638, 0.6638]	0.6789†	19
two-region exact	×2, peak	2	1	0.001	0.002	pairs 15859712	[0.6806, 0.6806]	0.6789†	646
boxing k=1	×2, de-twinned peak	2	1	0.01	0.020	box pairs 109836	[0.9786, 0.9786]	0.9786	375
boxing k=1	×2, de-twinned peak	2	1	0.001	0.002	box pairs 405544	[0.9963, 0.9963]	—	2693
boxing k=2	×2 two-tie, peak	2	2	0.01	0.020	box pairs 1297321	[0.6638, 0.6638]	0.6638	1120
boxing k=2	A3 vs A1+2, pro-rata 0.6	2	2	0.01	0.132	box pairs 910000	$[2.6 \times 10^{-5}, 3.0 \times 10^{-5}]$	2.8×10^{-5}	4423
boxing k=2	A3 vs A1+2, pro-rata 0.9	2	2	0.01	0.132	box pairs 1041180	[0.1793, 0.1888]	0.1841	3937
tree, sheared	×3 chain, pro-rata 0.6	3	1	0.01	0.030	total boxes 1835	[0.1000, 0.1023]	0.1129†	222
tree, sheared	×6 chain, pro-rata 0.6	6	1	0.01	0.059	total boxes 2806	[0.0452, 0.0476]	0.0558†	1243
tree, plain	×6 chain, pro-rata 0.6	6	1	0.01	0.059	total boxes 2918	[0.0378, 0.0575]	0.0558†	5090
env. messages	×3 two-tie, pro-rata 0.6	3	2	0.01	0.030	env cells 5592/16137/10785	[0.0835, 0.3173]	0.1104†	5621
env. messages	3 areas (cycle), pro-rata 0.6	3	4	0.01	0.030	env cells 9012/12609/10547	$[4.1 \times 10^{-6}, 0.0787]$	2.0×10^{-4}†	8499
joint bracket	×3 two-tie, pro-rata 0.6	3	2	0.1	0.271	occupied tuples 10201	[0.0440, 0.0492]	0.0443	104
joint bracket	×3 two-tie, pro-rata 0.6	3	2	0.01	0.030	occupied tuples 2991	[0.0463, 0.4473]	0.1104†	1645

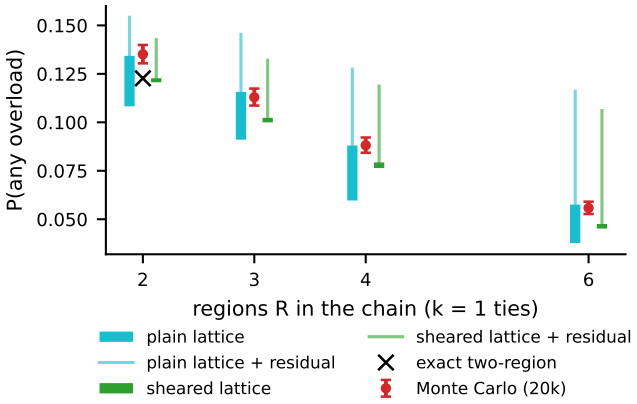


Fig. 5. Chains of $R = 2, 3, 4, 6$ single-bus regions (RQ4): certified P(any overload) before (thick) and after (thin) the joint residual, plain vs sheared lattice, with the Monte Carlo point and the exact two-region value at $R = 2$.

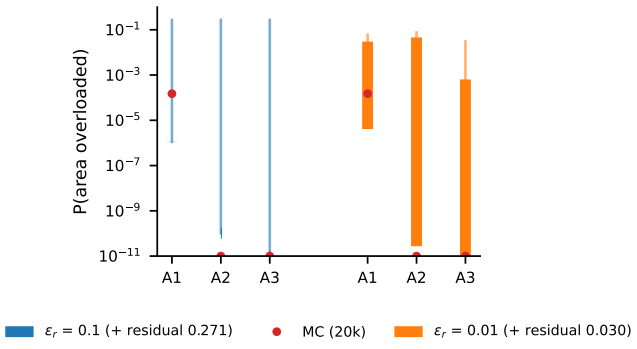


Fig. 6. RTS-GMLC three areas as regions (RQ2/RQ4): per-area certified P(overload) (thick: before residual; thin: with residual) at two truncations, with the Monte Carlo point. Lower bounds that are exactly zero are drawn at the axis floor.

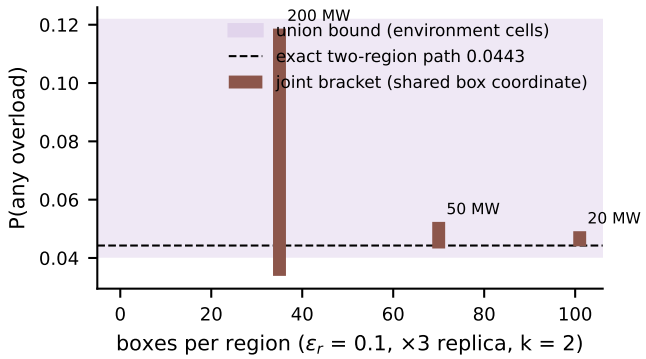


Fig. 7. Union bound vs the shared-box joint bracket (RQ4) on the ×3 two-tie replica at $\varepsilon_r = 10^{-1}$: the bracket narrows with the box count and contains the exact value; the union bound does not narrow.

V. DISCUSSION

The decomposition does what the hypothesis asked on the first criterion and on the coverage half of the second: regions enumerated separately and joined with certified brackets reach residuals that no global budget reaches, and every per-branch and per-region number comes with its un-enumerated mass. On wall-clock the honest statement is narrower: the two-region exact path at $k = 1$ reaches residual 0.020 in 19 s where a 10^5-state global enumeration leaves 0.111 in 24 s (the 5×10^6-state global run was not timed here), and that is the one case that is also faster; the exact pair path scales as $|S|^2$, boxing at $k \geq 2$ is on par with or slower than it, trees are seconds to minutes per pass, and environment rows are hours at $\varepsilon_r = 10^{-2}$ on three regions because the environment tuples are enumerated explicitly. Every boxing and lattice mechanism has a slack that is first-order in its discretisation (a bound; the measured ladders are at least that fast), and the one coordinate trick that removes a dimension of it (the shear) is a property of pro-rata-like dispatches. The joint probability of any overload at multi-bus interfaces among three or more regions is a union bound unless n_{boxes}^{R-1} is affordable, and the two events regions tend to share—twin-unit overloads driven by a common α—

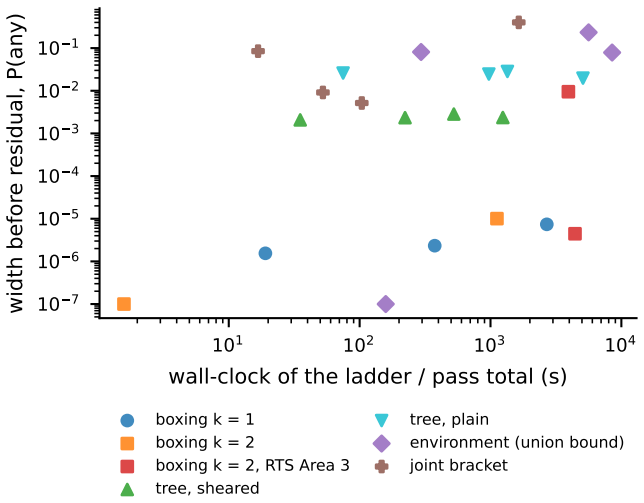


Fig. 8. Certified width before the residual against wall-clock, one marker per benchmark row and mechanism (power systems).

are exactly where the union bound is loose. Two defects were caught by the tests along the way and are recorded (in the program log, not in an artifact) rather than hidden: joint islanding is a joint property (a first tree version violated a lower bound by 3×10^{-5}), and tie flows must be evaluated in the environment path (the thermal walls exposed it). We regard the bracket discipline that caught them as the result that transfers.

VI. Limitations

DC / steady-state linear physics only; ties as fixtures except in the mixture; the gate is unmeasured (not failed) at $k_r \leq 3$ on IEEE 118 and PEGASE 1354, whose candidate partitions have no such region, and PEGASE was measured at $\varepsilon_r = 10^{-1}$ under a 5 000-state cap; IEEE 118 / PEGASE reliability rates are assumed (RTS-96 classes); environment tuples enumerated explicitly (three regions at $\varepsilon_r = 10^{-2}$; the lattice convolution of the tree path is the route beyond); thermal outcomes are path heat-flow limits, not node temperatures; single-machine Python/NumPy wall-clocks (Rust parity is out of scope by decision); the shear was compared across dispatches only on the tree path; multi-region runs at $k \geq 2$ reach $R = 3$ (environment tuples explicit) and $k = 1$ chains $R = 6$.

VII. Conclusions

TERRA local-to-global gives certified (bracketed, residual-stated) contingency-risk numbers for an interconnection whose regions are enumerated separately—per-branch and per-region overload probability, expected severity, islanding and deficit—by mechanisms that assume nothing about the number of regions, the interface dimension or the interconnection shape, demonstrated to $R = 6$ at $k = 1$ and $R = 3$ at $k \geq 2$ including a cycle. The joint probability of any overload is exact for two regions at any k; on single-bus chains it is a lattice bracket (joint islanding exact); beyond that it is a union bound or a shared-box bracket, with the gap reported. Coverage is where the decomposition wins; wall-clock only for two single-bus regions; the separator does not compress with locality. Future work, in order: a slack-driven splitting of the joint box tuples; lattice-convolved environments beyond three regions; conditioning on observed component states (local evidence) on the boxed, tree and environment paths; the companion study's Tier-2 corrective-dispatch bounds; and node-potential limits in the thermal domain.

Appendix

Every number is read from a committed `results.json` of the program's case studies (`local_to_global_poc`, `separator_compressibility`, `certified_boxing`, `certified_boxing_k`, `region_tree_chain`, `environment_messages`, `local_to_global_scaling`, `foundation_hvac/thermal_local_to_global`); each has a byte-pinned rendered report and artifact tests, and this paper's `paper.tex` is regenerated from them by its own test. The mechanisms are tested against brute force on exhaustive small systems (3-region two-tie chains, 12–16 components) for every metric including per-branch, ties, islanding, severity and deficit; the exact two-region path is tested against exhaustive global enumeration; the thermal adapter against an independent thermal solve. The consolidated benchmark is itself a derived artifact regenerated byte-identically from the source studies. Interface pairs, boxes, cells and tuples are reported in every row: the combinatorics are relocated, not eliminated.

References

[1] TERRA case study TERRA-CS-PF-08, "Structured probabilistic inference for power-system contingency risk: A comparison with Monte Carlo simulation," Silara EMS Studio, repository path `docs/terra_graph_engine/case_studies/foundation_power_flow/probabilistic_contingency_risk`, Tech. Rep., 2026, internal report; companion study on the same test system and reliability data.

[2] R. Billinton and R. N. Allan, *Reliability Evaluation of Power Systems*, 2nd ed. Plenum Press, 1996.

[3] A. M. Rei and M. T. Schilling, "Reliability assessment of the Brazilian power system using enumeration and Monte Carlo," *IEEE Transactions on Power Systems*, vol. 23, no. 3, pp. 1480–1487, 2008.

[4] S. Srinivasan and K. Sundar, "The solution of potential-driven, steady-state nonlinear network flow equations via graph partitioning," 2025, arXiv:2512.22124.

[5] S. L. Lauritzen and D. J. Spiegelhalter, "Local computations with probabilities on graphical structures and their application to expert systems," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 50, no. 2, pp. 157–224, 1988.

[6] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.

[7] D. P. Clancy, G. Gross, and F. F. Wu, "Probabilistic flows for reliability evaluation of multiarea power system interconnections," *International Journal of Electrical Power & Energy Systems*, vol. 5, no. 2, pp. 101–114, 1983.

[8] C. Singh and A. Lago-Gonzalez, "Improved algorithms for multi-area reliability evaluation using the decomposition-simulation approach," *IEEE Transactions on Power Systems*, vol. 4, no. 1, pp. 321–328, 1989.

[9] A. Monticelli, S. Deckmann, A. Garcia, and B. Stott, "Real-time external equivalents for static security analysis," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-98, no. 2, pp. 498–508, 1979.

[10] J. B. Ward, "Equivalent circuits for power-flow studies," *Transactions of the American Institute of Electrical Engineers*, vol. 68, no. 1, pp. 373–382, 1949.

[11] F. Dörfler and F. Bullo, "Kron reduction of graphs with applications to electrical networks," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 60, no. 1, pp. 150–163, 2013.
[12] A. J. Wood, B. F. Wollenberg, and G. B. Sheblé, *Power Generation, Operation, and Control*, 3rd ed. Wiley, 2013.
[13] R. E. Moore, R. B. Kearfott, and M. J. Cloud, *Introduction to Interval Analysis*. Philadelphia: SIAM, 2009.
[14] Z. Wang and F. L. Alvarado, "Interval arithmetic in power flow analysis," *IEEE Transactions on Power Systems*, vol. 7, no. 3, pp. 1341–1349, 1992.
[15] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965.
[16] C. Barrows *et al.*, "The IEEE reliability test system: A proposed 2019 update," *IEEE Transactions on Power Systems*, vol. 35, no. 1, pp. 119–127, 2020.
[17] M. Fiedler, "Algebraic connectivity of graphs," *Czechoslovak Mathematical Journal*, vol. 23, no. 2, pp. 298–305, 1973.
[18] B. W. Kernighan and S. Lin, "An efficient heuristic procedure for partitioning graphs," *The Bell System Technical Journal*, vol. 49, no. 2, pp. 291–307, 1970.
[19] R. Billinton and W. Li, *Reliability Assessment of Electric Power Systems Using Monte Carlo Methods*. Plenum Press, 1994.
[20] E. B. Wilson, "Probable inference, the law of succession, and statistical inference," *Journal of the American Statistical Association*, vol. 22, no. 158, pp. 209–212, 1927.
[21] T. Guler, G. Gross, and M. Liu, "Generalized line outage distribution factors," *IEEE Transactions on Power Systems*, vol. 22, no. 2, pp. 879–881, 2007.