Supplementary Information

Structure-Aware Exact Priority Optimization for Real-Time Cause–Effect Chains under Correlated Execution Demands

Haojie Cai[1,*]; Hao Cai[2]; Wenlong Hu[3]; Zhiwen Li[1]

[1]Nanjing University, Nanjing, Jiangsu, China
[2]Henan Polytechnic University, Jiaozuo, Henan, China
[3]Huazhong University of Science and Technology, Wuhan, Hubei, China

Corresponding author: chatjie@163.com
ORCID: 0009-0002-4265-4662

1 Scope and reproducibility

These materials accompany the main article. The main paper gives the general scheduling semantics, projection theorem, and the complete proof of the primary nine-job single-edge reversal. Here we specify two additional outcome-selected variants, reproducible atom certificates for the graph counterexamples, and the per-condition timing data. All task models are synthetic. The two additional variants share a parent task set; they are not two independent population samples.

The script `certificates/verify_paper.py` in the research archive uses only Python's standard library. It checks the primary event identities, the endpoint certificates for the path and the augmented graph, both graph counterexamples, and all 24 additional certificates, each at one cap value. The accompanying research artifact provides frozen code, numerical inputs, saved results and verification scripts. Its README distinguishes certificate checks, recorded-run audits and timing regeneration.

2 Additional strict-reversal variants

The chain is $A \to B \to A \to B$, with $D = 13$, all relative job deadlines 80, and three task types A, B, H. In both variants, the original graph is the path through the nine demand indices in table order. The augmented graph adds edge $(3, 8)$. The two variants use the same scheduling semantics as the paper. At common cap $q = 1/4$, exact optimal risks change from $7/8$ to $3/4$ in variant U, and from $3/4$ to $5/8$ in variant V, with the unique winner changing from BAH to ABH in both.

Table 1: Variant U: job parameters.

Index	Job	Release	Low	High
0	B0	6	1	2
1	B1	8	1	1
2	H0	5	1	3
3	A0	6	2	3
4	A1	10	1	2
5	A2	4	1	2
6	H1	2	1	2
7	H2	8	1	1
8	B2	10	2	3

The exact piecewise risk functions in Tables 3 and 4 are established by feasible endpoint laws and common dual inequalities over the stated parameter intervals. The research artifact includes the variant job models and executable pointwise certificates; the included verifier checks selected parameter values rather than enumerating a separate certificate for every point in the intervals.

Table 2: Variant V: job parameters.

Index	Job	Release	Low	High
0	B0	6	1	2
1	B1	8	1	1
2	H0	5	1	3
3	A0	4	2	3
4	A1	10	1	1
5	A2	4	1	2
6	H1	2	1	2
7	H2	8	1	1
8	B2	10	2	3

Table 3: Variant U: exact risks. P is the path, C the path with added edge.

Priority	R_P	R_C
BAH	$(1 + \min(1, 3q))/2$	same
ABH	$\min(1, 1/2 + 5q/2)$	$\min(1, 1/2 + q)$
AHB, HAB	$\min(1, 1/2 + 5q/2)$	same
BHA, HBA	1	1

Table 4: Variant V: exact risks.

Priority	R_P	R_C
BAH	$\min(3/4, 1/2 + q)$	same
ABH	$\min(1, 1/2 + 3q/2)$	$1/2 + q/2$
AHB, HAB	$\min(1, 1/2 + 5q/2)$	same
BHA, HBA	1	1

3 Probability-graph counterexamples

The following certificates use integer atom indices $a = \sum_i 2^i x_i$. Each row of the primal law gives an atom and its mass. The dual coefficients specify the affine majorant in the main paper. The first entry of α is the intercept, followed by one coefficient per demand bit. The entries of β follow the edge order listed below each certificate and in its JSON record. All arithmetic can be checked in fractions by enumerating the Boolean assignments.

3.1 Steiner-tree upper bound is not exact

For the four-terminal OR event in the main paper's Steiner example, the exact risk is 29/40, below the tree upper bound of 3/4. The primal atoms and masses are:

Atom	Mass
0	11/40
1	1/40
2	1/40
4	1/20
8	1/40
56	1/40
72	1/40
82	1/40
97	1/40
129	1/40
130	1/40
180	1/40
255	17/40

Dual coefficients:

$$\alpha = (0, 1/2, 1/2, 1/4, 1/2, -1/4, -1/4, 0, -1/4),$$

$$\beta = (1/4, 0, 1/4, 1/4, 0, 1/4, 1/4, 1/4, 1/4, 1/4, 1/4, 0).$$

Edge ordering: $(0,5)$, $(0,6)$, $(0,7)$, $(1,4)$, $(1,6)$, $(1,7)$, $(2,4)$, $(2,5)$, $(2,7)$, $(3,4)$, $(3,5)$, $(3,6)$.

3.2 Degree-three shortest-path reduction is not exact

For even parity on the three retained bits, the original graph with the degree-three center has exact risk 7/8. The primal atoms and masses are:

Atom	Mass
0	5/16
3	1/16
5	1/16
6	1/16
11	1/8
13	1/8
14	1/8
15	1/8

Dual coefficients:

$$\alpha = (1, -1/2, -1/2, -1/2, 1/2),$$

$$\beta = (1/2, 1/2, 1/2).$$

Edge ordering: $(0,3)$, $(1,3)$, $(2,3)$.

4 Prospective benchmark data

Table 5 reports the condition-level results. The method columns give median end-to-end runtimes in milliseconds over three runs. Risk is the exact optimal failure probability, and m is the number of task types. The path has $q = 1/5$ and the star $q = 1/4$. An asterisk identifies an SMPSC+W measurement retimed separately with an increased state-transition budget. The original budget failures remain in the raw records. P = full pipeline, B = short chain plus background, R = reentrant. S = synchronous compiler with witness transfer, F = shared family, A = active witness, E = projection only, T = complete structural algorithm.

Table 5: Prospective benchmark: condition-level risks and median runtimes.

m	Kind	Graph	Risk	S	F	A	E	T
3	B	path	0	7.65	3.65	3.38	2.92	1.29
3	B	star	0	6.86	3.20	3.33	3.19	1.34
3	P	path	1/2	4.89	5.85	5.78	5.38	4.05
3	P	star	1/2	5.00	5.57	5.53	5.19	3.77
3	R	path	1/2	5.95	4.70	4.29	4.23	3.76
3	R	star	1/2	6.22	4.12	4.30	3.93	3.83
4	B	path	0	29.47	6.91	5.95	5.41	3.80
4	B	star	0	29.34	5.78	5.74	5.55	3.97
4	P	path	7/10	4.18	5.37	5.96	5.73	5.75
4	P	star	17/24	3.92	5.70	5.77	5.93	6.05
4	R	path	3/5	16.22	15.14	19.52	19.49	19.69
4	R	star	5/8	16.35	16.13	20.77	20.71	20.78
5	B	path	0	282.33	18.55	15.56	15.09	13.56
5	B	star	0	273.02	15.70	15.42	15.36	12.92
5	P	path	1/2	15.50	17.32	16.33	16.22	13.81
5	P	star	1/2	14.73	13.00	14.00	13.29	11.38
5	R	path	1/2	160.11	26.57	18.33	17.82	17.08
5	R	star	1/2	158.16	19.80	17.22	16.49	16.92
6	B	path	0	2370.13*	61.00	54.65	54.27	52.12
6	B	star	0	2370.30*	57.64	56.78	56.04	53.25
6	P	path	7/10	59.31	63.18	81.42	87.94	100.19
6	P	star	5/8	58.03	59.87	73.42	73.37	75.52
6	R	path	0	1179.15*	59.07	52.11	51.72	50.08
6	R	star	0	1090.43*	52.49	57.35	52.78	53.85

The full input design and saved run summaries are included as JSON, and the condition table is included as CSV.

5 Follow-up matched factorial ablation

Table 6 reports the condition-level results. The method columns report medians of five end-to-end timings per condition in milliseconds; Risk is the exact optimal failure probability and m is the number of task types. N = neither projection nor closed forms; E = projection only; C = closed forms only; T = both. For each condition, all four arms used the same revised implementation and returned the same exact optimal risk. These times are not combined with the earlier three-repeat timing batch. The archive contains the 480 raw run records, rational atom certificates, and scripts that independently re-evaluate the scheduling events and recompute every median and geometric mean.

Table 6: Matched factorial ablation: condition-level risks and median runtimes.

m	Kind	Graph	Risk	N	E	C	T
3	B	path	0	3.51	2.88	1.36	1.38
3	B	star	0	3.34	3.01	1.32	1.29
3	P	path	1/2	5.85	6.26	3.98	4.10
3	P	star	1/2	5.80	5.46	3.80	4.01
3	R	path	1/2	4.78	4.07	4.61	3.89
3	R	star	1/2	4.52	4.17	4.44	3.96
4	B	path	0	6.50	5.69	4.21	3.85
4	B	star	0	6.19	5.48	3.98	4.02

m	Kind	Graph	Risk	N	E	C	T
4	P	path	7/10	5.90	6.05	6.11	6.35
4	P	star	17/24	6.13	6.53	6.61	6.82
4	R	path	3/5	20.45	20.04	19.97	19.69
4	R	star	5/8	21.89	22.82	23.09	23.83
5	B	path	0	18.58	18.42	15.66	14.90
5	B	star	0	17.90	16.65	15.22	14.34
5	P	path	1/2	18.31	17.18	14.53	14.54
5	P	star	1/2	14.15	14.79	12.46	13.11
5	R	path	1/2	20.95	19.12	21.66	19.47
5	R	star	1/2	18.83	17.52	19.21	18.90
6	B	path	0	64.78	68.77	66.88	59.65
6	B	star	0	62.26	59.36	54.52	53.80
6	P	path	7/10	86.57	96.13	88.86	93.32
6	P	star	5/8	76.94	77.69	78.81	78.74
6	R	path	0	53.28	55.03	52.97	48.74
6	R	star	0	55.39	52.47	52.99	50.95