Encoding Matters More Than the Solver

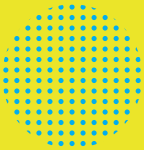
Penalty Coefficients in QUBO Formulations for Quantum Annealing

Roman Kužel[1] Mátyás Koniorczyk[2] Janez Povh[1]

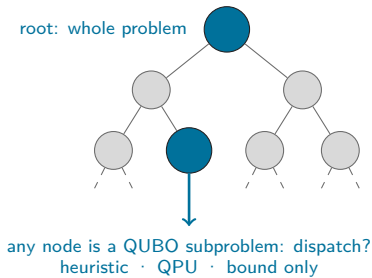
[1]Rudolfovo – Science and Technology Centre Novo Mesto, Slovenia
[2]HUN-REN Wigner Research Centre for Physics, Budapest, Hungary
roman.kuzel@rudolfovo.eu

Slovenski forum o kvantnem računalništvu, 2026

RUDOLFOVO
Znanstveno in tehnološko središče Novo mesto



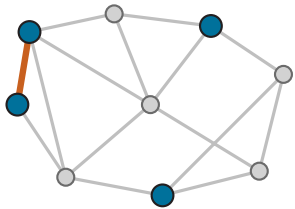
An exact solver that may call a quantum annealer



- **BiqBin** certifies optima by branch-and-bound on EuroHPC Vega.
- A node may go to a heuristic or a QPU for a better feasible solution (incumbent); the certificate never depends on that call.
- Annealers take QUBO only: every constraint must first become a **penalty**.

Before dispatching a node: what does its encoding demand, and what can the target deliver?

Model problem: the sparsest k-subgraph



$k = 4$ vertices (teal) spanning 1 edge (orange)

Choose exactly k vertices spanning the **fewest edges**: $\min \frac{1}{2} x^\top A x$ s.t. $e^\top x = k$. NP-hard; 270 instances with known optima.

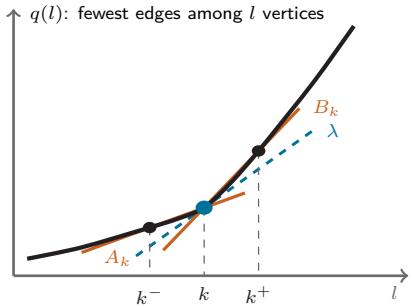
PENALTY FORM OF $e^\top x = k$

$$\min \ \tfrac{1}{2}\, x^\top Q x - \lambda\,(e^\top x - k) + \eta\,(e^\top x - k)^2$$

λ: linear (Lagrange multiplier) η: quadratic (the usual "big M")

100 % vs 2 % solved to optimality: same simulated-annealing solver, same 270 instances, 10 reads; nothing changed but λ and η.

What we proved for a cardinality constraint



λ tilts the curve so that k is its lowest point; η adds curvature only if no tilt works.

1. **Exact condition.** QUBO $\equiv$ constrained problem iff λ sits between the chord slopes A_k, B_k and $\eta > (A_k - B_k)/2$.
2. **No η needed** iff $A_k \leq B_k$: then $\eta = 0$ and λ alone enforces equality from both sides.
3. **Optimal η is NP-hard**, so we need estimates.

Sufficient penalties from one greedy pass, no SDP

Any feasible k-subgraph with $\tilde{m}$ edges (greedy max-degree removal) gives $\tilde{A}_k = \min\{\tilde{m},\, k-1\}$, $\tilde{B}_k = 0$:

T1 · symmetric split

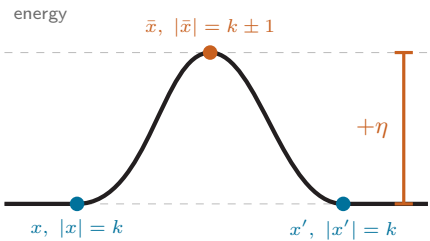
$$\lambda = \eta = \tfrac{1}{2}\min\{\tilde{m},\ k-1\}$$

T2 · pure quadratic (traditional)

$$\lambda = 0, \quad \eta = \min\{\tilde{m},\ k-1\}$$

- Both are provably sufficient: correct encodings for an exact solver. By construction $\eta_{T2} = 2\eta_{T1}$.
- Against the general bounds: $\eta_{T1} \le \sigma_{\mathrm{EXPEDIS}}/2 \le \eta_{\mathrm{Lasserre}}/4$; on our instances 5–180× smaller.

Mechanism 1 · penalty barriers in the landscape



Feasible solutions one **swap** apart connect only through an infeasible state, where the penalty adds exactly η, whatever the objective does.

- **Single-flip dynamics (SA):** escape $\propto e^{-\eta/T}$; the solver stays in its starting basin.
- **Branch-and-bound:** the inflated range loosens the bounds; many more nodes.

Mechanism 2 · η consumes the hardware range

RU DOL FO VO

Programming precision

$$\Delta J = J_{\text{dev}}/(2\eta + 1)$$

η sits in every coupler and sets the rescaling; below the precision floor the graph is invisible to the QPU.

Chain strength is a second penalty

$$J_{\text{chain}} > |h_i| + \Sigma_j |J_{ij}|$$

grows with n, η, λ (Choi 2008). Too weak: chains break. Strong: the chain sets the scale.

λ is different: it touches only the local fields h_i. **For a dispatcher:** both quantities follow from (λ, η) and the device properties, before any call.

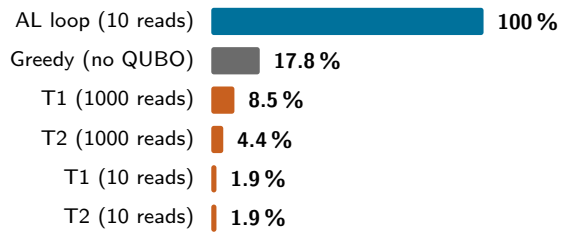
A practical augmented Lagrangian loop

ALGORITHM (SOLVER-AGNOSTIC)

1. λ_0 from a greedy pass, η_0 small
2. Inner solve:
   $x \leftarrow \arg\min \ \frac{1}{2}x^\top Ax - \lambda(e^\top x - k) + \eta(e^\top x - k)^2$
3. If $e^\top x \neq k$: $\lambda \leftarrow \lambda + 2\eta\,(k - e^\top x)$
4. $\eta \leftarrow \rho\eta$; repeat until $e^\top x = k$

- **Warm start for free:** the degree of the last vertex removed by greedy max-degree removal $\approx$ the slope the theorem asks for.
- Inner solver: anything that takes a QUBO.
- 2–4 iterations on average, with every solver.

Classical solvers: noise-free, same collapse



Terminal η, AL loop

≈ 0.1

T1 / T2 prescribe 26 / 52

MQLib hyper-heuristic and **SBM**, untuned: both 100 % under the AL loop. Under T1, MQLib 97 % optimal; SBM feasible on only 15 % of instances.

QPU pilot: the embedding is a second penalty

AL loop · QPU + descent
83 %
optimal (raw QPU: 0.4 %)

T1/T2 · QPU + descent
2 %
optimal

AL loop · random starts
92 %
same descent, same 500 reads

T1/T2 · random starts
10 %
same collapse

- **Pilot protocol**, D-Wave Advantage2: one minorminer embedding per size, chain strength 15 by trial, one anneal time, no gauge averaging. We know what we would change.
- **The encoding effect stands:** the classical control reproduces it. Not a quantum effect.
- **Chains set the scale:** the objective was programmed at $\approx 1.7\,\%$ of the coupler range, comparable to the device's control errors; at $n = 40$ no chain broke, yet 1 of 45 raw runs was optimal.

Exact solver: seconds of preprocessing save hours

Small · $\eta \in \{0, 1\}$ from five quick SA runs

270 / 270

solved; median 1 node for $n \leq 100$

T1 at $n = 160$ (η up to 60 instead of 0–1)

158×

median nodes; worst instance 59 s → 68 023 s

T2 at $n = 160$ (η twice T1)

15 / 45

timeouts; worst case $1 \rightarrow 57\,581$ nodes

BiqBin on EuroHPC Vega. No known optimum used: greedy $\tilde{m}$ for T1/T2; for Small, five SA runs at $k-2, \ldots, k+2$ and the theorem give (λ, η). The certificate is the theorem's; the encoding is paid for in nodes only.

A provably sufficient penalty instead of the problem-adapted one: two to three orders of magnitude.

The same analysis, per node, is what a dispatcher needs before it calls anything.

What a dispatch strategy for the exact solver needs

RU DOL FO VO

Encoding matters more than the solver, so the encoding of a node is decided before the solver for it is.

1 · PENALTIES · DONE

(λ, η) from one greedy pass per node. Deeper nodes: smaller graphs, smaller k, typically $\eta \approx 0$.

2 · REPRESENTABILITY · SPECIFIED

From (λ, η) and device properties: chain strength, auto-scaling, signal vs precision, clique limit. The same quantities set each solver's parameters (temperature schedule, step size, anneal time); no-go is the case where no setting works.

3 · MARGINAL BENEFIT · OPEN

What an incumbent from the device is worth at that node, against its access time. Needs a benchmark with the encoding as an axis.

Thank you

Encoding matters more than the solver: report your penalties, compare encodings, and keep a classical control in every run.

CONTACT

Roman Kužel
roman.kuzel@rudolfovo.eu

Centre for Applied Quantum Computing (CEQ), Rudolfovo, Novo Mesto

BiqBin exact solver: open source, pluggable bounds and heuristics

PAPER AND ACKNOWLEDGEMENTS

R. Kužel, M. Koniorczyk, J. Povh, *Linear and quadratic penalties for cardinality constraints in unconstrained binary optimization*, in preparation for J. Optim. Theory Appl.

EuroHPC Vega (IZUM) for HPC time; JSC/JUNIQ for D-Wave Leap access (Advantage2) time

RUDOLFOVO

Znanstveno in tehnološko središče Novo mesto

