On the Stability Constant for ε-light Vertex Subsets

Question 6: Obstructions and a Greedy Barrier Framework

Dietmar Wolz, Ingo Althöfer

February 11, 2026

Abstract

For a graph $G = (V, E)$ with Laplacian L, call $S \subseteq V$ *ε-light* if the induced Laplacian satisfies $L_S \preceq \varepsilon L$. We study whether there exists a universal constant $c > 0$ such that for every G and every $\varepsilon \in (0, 1)$ there exists an ε-light set with $|S| \geq c\,\varepsilon\,|V|$.

We prove the universal *upper bound* $c \leq \frac{1}{2}$ using perfect matchings. We also explain why "half-star" linearization techniques are too strong: they implicitly control the cut Laplacian and therefore fail on matchings even when feasible sets exist. Finally, we give a fully explicit greedy spectral barrier algorithm (Route D) that controls L_S directly. The barrier analysis is fully rigorous and reduces the remaining difficulty to a single structural hypothesis: a *mass τ-control* lemma (stated explicitly). Conditional on this hypothesis with parameters (θ, β), the algorithm constructs sets of size

$$|S| \;\geq\; \frac{\varepsilon}{\alpha + \varepsilon}\, n - O(1), \qquad \text{with} \;\; \alpha \geq \frac{2}{\beta(1-\theta)}.$$

In particular, any improvement of the τ-control/selection constants directly propagates to the final c.

Contents

1 Problem statement

Let $G = (V, E)$ be a (simple, unweighted) graph with $n := |V|$ and Laplacian $L = \sum_{e \in E} L_e$, where for $e = \{p, q\}$ we write

$$L_e := (e_p - e_q)(e_p - e_q)^\top.$$

For a subset $S \subseteq V$, define the induced Laplacian L_S as the Laplacian of the induced subgraph $G[S]$ embedded in $\mathbb{R}^V$ (i.e. entries outside S are zero):

$$L_S := \sum_{e \in E(S,S)} L_e.$$

We call S *ε-light* if

$$L_S \preceq \varepsilon L. \tag{1}$$

Question 6 asks for the best universal constant c such that for all G and all $\varepsilon \in (0,1)$ there exists an ε-light set S with $|S| \geq c\,\varepsilon\,n$.

2 Universal upper bound: perfect matchings force $c \leq \frac{1}{2}$

Proposition 2.1 (Perfect matching obstruction)**.** *Let G be a perfect matching on n vertices (so n is even and E consists of $n/2$ disjoint edges). Then for every $\varepsilon \in (0,1)$, any set $S \subseteq V$ satisfying* (1) *must be an independent set. Consequently, $|S| \leq n/2$.*

Proof. Write the edges as $e_i = \{u_i, v_i\}$, $i = 1, \dots, n/2$. If S contains both endpoints of some edge e_k, then $L_S \succeq L_{e_k}$. Since $L = L_{e_1} + \cdots + L_{e_{n/2}}$ and the L_{e_i} have disjoint supports, testing $L_{e_k} \preceq \varepsilon L$ against the vector $x := e_{u_k} - e_{v_k}$ gives

$$x^\top L_{e_k} x = x^\top L x.$$

A direct calculation with $x = (1, -1)$ on the $\{u_k, v_k\}$ coordinates yields $x^\top L_{e_k} x = 4$ and hence $x^\top L x = 4$ as well. Therefore $4 \leq \varepsilon \cdot 4$, so $\varepsilon \geq 1$, contradicting $\varepsilon \in (0,1)$. Thus S cannot contain both endpoints of any matching edge, i.e. S is independent, and hence $|S| \leq n/2$. ☐

Corollary 2.2 (Upper bound on the universal constant)**.** *Any universal constant c with the property "for all G and all $\varepsilon \in (0,1)$, there exists S with $|S| \geq c\varepsilon n$ and $L_S \preceq \varepsilon L$" must satisfy $c \leq \frac{1}{2}$.*

Proof. Apply Proposition 2.1 to a perfect matching and choose ε arbitrarily close to 1. Then $c\varepsilon n \leq n/2$ forces $c \leq 1/(2\varepsilon)$ for all such ε, hence $c \leq 1/2$. ☐

3 Why half-star linearization fails on matchings

A common "linearization" writes L as a sum of vertex-local PSD matrices and then tries to select many summands. Define for each vertex u the *half-star* matrix

$$B_u := \sum_{v \sim u} \frac{1}{2} L_{\{u,v\}}.$$

Then $L = \sum_{u \in V} B_u$.

Lemma 3.1 (Half-star decomposition)**.** *For every $S \subseteq V$,*

$$\sum_{u \in S} B_u \;=\; L_S \;+\; \frac{1}{2} L_{\partial S},$$

where $L_{\partial S} := \sum_{e \in E(S, V \setminus S)} L_e$ is the cut Laplacian.

Proof. Fix an edge $e = \{p, q\}$. If $p, q \in S$, then e appears once in B_p and once in B_q, each time with coefficient $1/2$, so its total contribution to $\sum_{u \in S} B_u$ equals L_e, matching its contribution to L_S. If exactly one endpoint lies in S, then e contributes $(1/2)L_e$ to $\sum_{u \in S} B_u$ and contributes to $L_{\partial S}$. If neither endpoint lies in S it contributes nothing. Summing over edges gives the identity. ☐

Remark 3.2 (Why this is too strong for small ε)**.** *On a perfect matching, take S to be any independent set of size $n/2$ (one endpoint per edge). Then $L_S = 0$, hence S is ε-light for every $\varepsilon \in (0, 1)$. However, $L_{\partial S} = L$ (all edges cross the cut), so $\sum_{u\in S} B_u = \frac{1}{2}L$. Thus the linearized constraint $\sum_{u\in S} B_u \preceq \varepsilon L$ would force $\varepsilon \geq 1/2$, even though feasible sets exist for all ε. Therefore any proof of the sharp constant must control L_S without controlling $L_{\partial S}$.*

4 Route D: a greedy barrier framework (conditional lower bound)

4.1 Normalization on $\mathrm{im}(L)$

Let $L^\dagger$ denote the Moore–Penrose pseudoinverse of L and write

$$A_S \coloneqq L^{\dagger/2} L_S L^{\dagger/2}, \qquad P \coloneqq L^{\dagger/2} L L^{\dagger/2}.$$

Then P is the orthogonal projection onto $\mathrm{im}(L)$ and A_S is PSD with $\mathrm{im}(A_S) \subseteq \mathrm{im}(L)$.

Lemma 4.1 (Equivalence of constraints)**.** *For $\varepsilon > 0$, the inequality $L_S \preceq \varepsilon L$ holds on $\mathbb{R}^V$ if and only if*

$$A_S \preceq \varepsilon P \quad \textit{on } \mathrm{im}(L).$$

Proof. If $L_S \preceq \varepsilon L$, left- and right-multiply by $L^{\dagger/2}$ to obtain $A_S \preceq \varepsilon\, L^{\dagger/2} L L^{\dagger/2} = \varepsilon P$. Conversely, assume $A_S \preceq \varepsilon P$ on $\mathrm{im}(L)$. For any $x \in \mathbb{R}^V$ write $y = L^{1/2}x$ (interpreting $L^{1/2}$ on $\mathrm{im}(L)$ and 0 on $\ker(L)$). Then $x^\top L_S x = y^\top A_S y$ and $x^\top L x = y^\top P y = \|Py\|^2$. Thus $y^\top A_S y \leq \varepsilon\, y^\top P y$ for all $y \in \mathrm{im}(L)$, hence $L_S \preceq \varepsilon L$. □

Lemma 4.2 (Support on $\mathrm{im}(L)$)**.** *For every $S \subseteq V$, $PA_S = A_S P = A_S$. Moreover, if $u \notin S$ and we define the update*

$$\Delta_u(S) \coloneqq A_{S\cup\{u\}} - A_S,$$

then $\Delta_u(S) \succeq 0$ and $P\Delta_u(S) = \Delta_u(S)P = \Delta_u(S)$.

Proof. All matrices have the form $L^{\dagger/2}(\cdot)L^{\dagger/2}$, hence their images lie in $\mathrm{im}(L)$ and are annihilated by $I - P$. Positivity follows because $L_{S\cup\{u\}} - L_S$ is a sum of edge Laplacians and thus PSD. □

4.2 A key identity: sum of updates equals the cut Laplacian

For $S \subseteq V$, let $L_{\partial S} \coloneqq \sum_{e\in E(S, V\setminus S)} L_e$.

Lemma 4.3 (Sum-of-updates identity)**.** *For any $S \subseteq V$,*

$$\sum_{u\notin S} \Delta_u(S) \;=\; L^{\dagger/2} L_{\partial S} L^{\dagger/2} \;\preceq\; P.$$

Proof. Adding $u \notin S$ introduces exactly the edges from u to S, hence $L_{S\cup\{u\}} - L_S = \sum_{v\in S\cap N(u)} L_{\{u,v\}}$. Summing over all $u \notin S$ counts each cut edge exactly once, giving $L_{\partial S}$. Finally $L = L_S + L_{\partial S} + L_{V\setminus S}$ implies $L_{\partial S} \preceq L$, and conjugating by $L^{\dagger/2}$ yields the PSD bound by P. □

4.3 Barrier potential and one-step bounds

Fix a barrier level $b > 0$ and define the (unsubtracted) potential

$$\Phi_b(A) \coloneqq \mathrm{tr}\left((bP - A)^{-1}\right),$$

where all inverses and traces are taken on $\mathrm{im}(L)$ (so on $\mathrm{im}(L)$ one may read $P = I$).

Given a barrier increment $\delta > 0$, set

$$\Psi_\delta \coloneqq \left((b + \delta)P - A\right)^{-1}.$$

For an update $\Delta \succeq 0$ supported on $\mathrm{im}(L)$ define the shifted scores

$$\tau_\delta(\Delta) \coloneqq \mathrm{tr}(\Psi_\delta \Delta), \qquad \eta_\delta(\Delta) \coloneqq \mathrm{tr}(\Psi_\delta^2 \Delta).$$

Lemma 4.4 (Average τ bound). *Let $U = V \setminus S$. For fixed $A = A_S$ and $\delta > 0$,*

$$\sum_{u\in U} \tau_{u,\delta} \;\le\; \Phi_{b+\delta}(A), \quad \textit{where} \;\; \tau_{u,\delta} := \tau_\delta(\Delta_u(S)).$$

Proof. By Lemma 4.3, $\sum_{u\in U} \Delta_u(S) \preceq P$. Since $\Psi_\delta \succeq 0$,

$$\sum_{u\in U} \tau_{u,\delta} = \operatorname{tr}\left(\Psi_\delta \sum_{u\in U} \Delta_u(S)\right) \le \operatorname{tr}(\Psi_\delta P) = \operatorname{tr}(\Psi_\delta) = \Phi_{b+\delta}(A).$$

□

Lemma 4.5 (τ–η comparison). *Assume $0 \preceq A \prec (b+\delta)P$ on $\operatorname{im}(L)$. Then for every PSD update Δ supported on $\operatorname{im}(L)$,*

$$\tau_\delta(\Delta) \;\le\; (b+\delta)\,\eta_\delta(\Delta).$$

Proof. On $\operatorname{im}(L)$, $P = I$ and the eigenvalues of $(b+\delta)I - A$ are $b+\delta-\lambda_i(A) \le b+\delta$, hence $\Psi_\delta = ((b+\delta)I - A)^{-1} \succeq \frac{1}{b+\delta} I$. Let $X = \Psi_\delta^{1/2}\Delta\Psi_\delta^{1/2} \succeq 0$. Then $\tau_\delta(\Delta) = \operatorname{tr}(X)$ and

$$\eta_\delta(\Delta) = \operatorname{tr}(\Psi_\delta X) \;\ge\; \lambda_{\min}(\Psi_\delta)\operatorname{tr}(X) \;\ge\; \frac{1}{b+\delta}\,\tau_\delta(\Delta).$$

□

Lemma 4.6 (One-step resolvent bound). *Assume $0 \preceq A \prec (b+\delta)P$ on $\operatorname{im}(L)$ and let $\Delta \succeq 0$ be supported on $\operatorname{im}(L)$. If $\tau_\delta(\Delta) < 1$, then*

$$\Phi_{b+\delta}(A+\Delta) \;\le\; \Phi_{b+\delta}(A) + \frac{\eta_\delta(\Delta)}{1-\tau_\delta(\Delta)}.$$

Proof. On $\operatorname{im}(L)$, set $\Psi_\delta = ((b+\delta)I - A)^{-1}$ and $X = \Psi_\delta^{1/2}\Delta\Psi_\delta^{1/2} \succeq 0$. Then

$$((b+\delta)I - (A+\Delta))^{-1} = \Psi_\delta^{1/2}(I-X)^{-1}\Psi_\delta^{1/2}.$$

If $\tau := \operatorname{tr}(X) < 1$, then $\lambda_{\max}(X) \le \operatorname{tr}(X) = \tau$. For scalar $x \in [0,\tau]$ one has $(1-x)^{-1} \le 1 + \frac{x}{1-\tau}$, hence by functional calculus $(I-X)^{-1} \preceq I + \frac{1}{1-\tau}X$. Taking traces and using cyclicity yields

$$\Phi_{b+\delta}(A+\Delta) \le \Phi_{b+\delta}(A) + \frac{1}{1-\tau}\operatorname{tr}(\Psi_\delta X) = \Phi_{b+\delta}(A) + \frac{\eta_\delta(\Delta)}{1-\tau_\delta(\Delta)}.$$

□

Lemma 4.7 (Barrier-drop lower bound). *Let $0 \preceq A \prec bP$ on $\operatorname{im}(L)$ and $\delta > 0$. With $\Psi_\delta = ((b+\delta)P - A)^{-1}$,*

$$\Phi_b(A) - \Phi_{b+\delta}(A) \;\ge\; \delta\,\operatorname{tr}(\Psi_\delta^2).$$

Proof. Diagonalize A on $\operatorname{im}(L)$ with eigenvalues $\lambda_i \in [0,b)$. Then

$$\Phi_b(A) - \Phi_{b+\delta}(A) = \sum_i \left(\frac{1}{b-\lambda_i} - \frac{1}{b+\delta-\lambda_i}\right) = \sum_i \frac{\delta}{(b-\lambda_i)(b+\delta-\lambda_i)} \ge \sum_i \frac{\delta}{(b+\delta-\lambda_i)^2} = \delta\,\operatorname{tr}(\Psi_\delta^2).$$

□

Proposition 4.8 (Shift–then–update step condition). *Let $\Delta \succeq 0$ be supported on $\operatorname{im}(L)$ and fix $\delta > 0$. If $\tau_\delta(\Delta) < 1$ and*

$$\frac{\eta_\delta(\Delta)}{1-\tau_\delta(\Delta)} \;\le\; \Phi_b(A) - \Phi_{b+\delta}(A),$$

then $\Phi_{b+\delta}(A+\Delta) \le \Phi_b(A)$.

Proof. By Lemma 4.6, $\Phi_{b+\delta}(A+\Delta) \le \Phi_{b+\delta}(A) + \frac{\eta_\delta(\Delta)}{1-\tau_\delta(\Delta)}$. Using the assumed inequality yields $\Phi_{b+\delta}(A+\Delta) \le \Phi_b(A)$. □

4.4 The remaining missing ingredient: mass τ-control

At a step k, let $S = S_k$, $A = A_k$, and let $U := V \setminus S$ with $|U| = n - k$. Fix $\delta = \delta_k$ and define $\tau_{u,\delta} = \tau_\delta(\Delta_u(S))$.

Lemma 4.9 (Mass τ-control hypothesis). *Fix parameters $\theta \in (0,1)$ and $\beta \in (0,1]$. At each step k, define the set of τ-good candidates*

$$T_k := \{u \in U : \ \tau_{u,\delta_k} \le \theta\}.$$

Hypothesis: $|T_k| \ge \beta|U|$ *for every k.*

Remark 4.10 (A sufficient condition via a global potential bound). *By Markov's inequality and Lemma 4.4,*

$$\big|\{u \in U : \tau_{u,\delta} > \theta\}\big| \le \frac{1}{\theta} \sum_{u \in U} \tau_{u,\delta} \le \frac{1}{\theta}\, \Phi_{b+\delta}(A).$$

Thus a sufficient route to Lemma 4.9 is to prove an a priori bound of the form $\Phi_{b+\delta}(A) \le (1-\beta)\theta|U|$ along the run. This is exactly where one typically introduces a baseline-subtracted potential in BSS-style analyses. No such unconditional bound is proved in this write-up; the lower bound below is therefore conditional.

4.5 Selection and constant propagation

Lemma 4.11 (Selection among τ-good vertices). *Assume Lemma 4.9 holds at a step with parameters (θ, β). Let $\delta > 0$ and $\Psi_\delta = ((b+\delta)P - A)^{-1}$. Then there exists $u \in T_k$ such that*

$$\eta_{u,\delta} \ \le \ \frac{1}{\beta|U|}\, \mathrm{tr}(\Psi_\delta^2).$$

Proof. All $\eta_{u,\delta} \ge 0$ and $T_k \subseteq U$, so

$$\sum_{u \in T_k} \eta_{u,\delta} \le \sum_{u \in U} \eta_{u,\delta} = \mathrm{tr}\left(\Psi_\delta^2 \sum_{u \in U} \Delta_u(S)\right) \le \mathrm{tr}(\Psi_\delta^2 P) = \mathrm{tr}(\Psi_\delta^2),$$

using Lemma 4.3. Averaging over $|T_k| \ge \beta|U|$ gives the claim. □

Theorem 4.12 (Conditional barrier step, with explicit constants). *Assume mass τ-control holds with parameters (θ, β) at a step. Let $|U| = n - k$ and choose a barrier increment*

$$\delta := \frac{\alpha}{|U|+1} \qquad \text{with} \qquad \alpha \ge \frac{2}{\beta(1-\theta)}.$$

Then there exists a candidate $u \in U$ such that the step condition of Proposition 4.8 holds, and the update can be performed while maintaining $\Phi_{b+\delta}(A + \Delta_u) \le \Phi_b(A)$.

Proof. By Lemma 4.11 there exists $u \in T_k$ with $\eta_{u,\delta} \le \mathrm{tr}(\Psi_\delta^2)/(\beta|U|)$, and by definition of T_k we have $\tau_{u,\delta} \le \theta < 1$. Hence

$$\frac{\eta_{u,\delta}}{1 - \tau_{u,\delta}} \le \frac{1}{1-\theta} \cdot \frac{1}{\beta|U|}\, \mathrm{tr}(\Psi_\delta^2).$$

By Lemma 4.7, $\Phi_b(A) - \Phi_{b+\delta}(A) \ge \delta\, \mathrm{tr}(\Psi_\delta^2)$. Therefore the step condition holds provided

$$\frac{1}{1-\theta} \cdot \frac{1}{\beta|U|} \le \delta = \frac{\alpha}{|U|+1}.$$

Since $|U|/(|U|+1) \ge 1/2$, it suffices that $\alpha \ge 2/(\beta(1-\theta))$. Then Proposition 4.8 applies. □

4.6 Algorithm and stopping time

Greedy barrier algorithm (Route D). Fix $\alpha \geq 1$ and set $\delta_k := \alpha/(n-k+1)$. Initialize $S_0 = \emptyset$, $A_0 = 0$, and $b_0 := \varepsilon/(n+1)$. For $k = 0, 1, 2, \dots$ while $b_k \leq \varepsilon$:

1. Set $\delta = \delta_k$ and $\Psi_\delta := ((b_k + \delta)P - A_k)^{-1}$.
2. Choose $u_k \in U_k := V \setminus S_k$ such that the step condition in Proposition 4.8 holds.
3. Update $S_{k+1} := S_k \cup \{u_k\}$, $A_{k+1} := A_k + \Delta_{u_k}(S_k)$, and $b_{k+1} := b_k + \delta$.

Return $S_{k_\star}$ where $k_\star := \max\{k : \ b_k \leq \varepsilon\}$.

Lemma 4.13 (Stopping time bound)**.** *Let $\alpha \geq 1$ and $\delta_k = \alpha/(n-k+1)$. Then the stopping time satisfies*

$$k_\star \ \geq \ n\big(1 - e^{-\varepsilon/\alpha}\big) - 3.$$

Consequently, using $1 - e^{-x} \geq \frac{x}{1+x}$ for $x \in (0,1)$,

$$k_\star \ \geq \ \frac{\varepsilon}{\alpha + \varepsilon} n - 3.$$

Proof. We have

$$b_k = b_0 + \sum_{j=0}^{k-1} \frac{\alpha}{n-j+1} = b_0 + \alpha \sum_{t=n-k+2}^{n+1} \frac{1}{t}.$$

Using the lower integral bound $\sum_{t=a}^{b} \frac{1}{t} \geq \ln\big(\frac{b+1}{a}\big)$, the failure of the stopping condition at $k_\star + 1$ implies

$$\varepsilon < b_{k_\star+1} \leq b_0 + \alpha \ln\left(\frac{n+2}{n-k_\star+1}\right).$$

Rearranging gives $n - k_\star + 1 < (n+2)\exp\big(-(\varepsilon - b_0)/\alpha\big)$. With $b_0 = \varepsilon/(n+1)$, the slack can be absorbed into the additive constant 3, yielding $k_\star \geq n(1 - e^{-\varepsilon/\alpha}) - 3$. The second inequality uses $1 - e^{-x} \geq \frac{x}{1+x}$. ☐

Theorem 4.14 (Conditional lower bound, parameterized by (θ, β))**.** *Assume the mass τ-control hypothesis (Lemma 4.9) holds along the entire run with parameters (θ, β). Run the greedy barrier algorithm with any*

$$\alpha \ \geq \ \frac{2}{\beta(1-\theta)}.$$

Then the returned set S is ε-light (i.e. $L_S \preceq \varepsilon L$) and satisfies

$$|S| \ \geq \ \frac{\varepsilon}{\alpha + \varepsilon} n - 3.$$

Proof. At each step, Theorem 4.12 provides a valid choice of u_k ensuring $\Phi_{b_{k+1}}(A_{k+1}) \leq \Phi_{b_k}(A_k) < \infty$ and hence $A_{k+1} \prec b_{k+1}P$ on $\mathrm{im}(L)$. When the algorithm stops at $k_\star$ we have $b_{k_\star} \leq \varepsilon$ and therefore $A_{k_\star} \preceq b_{k_\star}P \preceq \varepsilon P$, which is equivalent to $L_{S_{k_\star}} \preceq \varepsilon L$ by Lemma 4.1. The size bound is Lemma 4.13. ☐

5 Status summary

Remark 5.1 (What is proved unconditionally vs. conditionally)**.** *Unconditionally, this write-up proves the sharp obstruction $c \leq \frac{1}{2}$ (Section 2) and pinpoints why half-star linearization cannot yield the sharp constant (Section 3). The greedy barrier analysis in Section 4 is rigorous, but the* lower bound *is conditional: it depends on proving a mass τ-control statement (Lemma 4.9) or an equivalent global upper bound on the shifted potential (Remark 4.10). Once such a control is available with parameters (θ, β), the resulting constant follows explicitly from Theorem 4.14.*