Maximum Flow With a Logarithmic Potential

August 2026

Abstract

We give a short deterministic $m^{1+o(1)}$-time maximum-flow algorithm using the min-ratio cut data structure in the form stated by Li and Wice. One capped logarithmic potential replaces their residual outer loop. The proof is the static specialization of the dual-potential framework of van den Brand et al.; the cap keeps every maintained number polynomial while introducing only a negligible dual infeasibility.

1 Introduction

Almost-linear maximum flow was obtained by Chen et al. [CKL+22]. The dual-potential framework was developed by van den Brand et al. [VDBCK+24]. Li's balancing viewpoint [Li26] and the companion note of Li and Wice [LW26] isolate the same inner min-ratio cut data structure. The latter note invokes it repeatedly on residual graphs. Here one logarithmic potential trajectory replaces that outer loop.

Theorem 1. *Let G be a directed graph with m arcs, integral polynomially bounded capacities, and distinct vertices s, t. An exact maximum s–t flow can be found deterministically in $m^{1+o(1)}$ time.*

We prove a threshold transshipment statement below and reduce it to maximum flow in Section 3. For an arc $e = (u, v)$ let $(By)_e = y_u - y_v$, so $B^\top f$ is outflow minus inflow. Parallel arcs are allowed.

2 Main proof

Consider

$$\mathrm{OPT} = \min\{c^\top f : B^\top f = d,\ f \geq 0\} = \max\{d^\top y : c - By \geq 0\}. \tag{1}$$

The directed graph has p arcs and is strongly connected, $c \in \mathbb{Z}^E_{\geq 1}$, $d \in \mathbb{Z}^V$, and $\mathbf{1}^\top d = 0$. Fix an integer threshold $F \geq 1$. Let $P \geq 2$ be a polynomial bound on p, the number of vertices, F, $\|c\|_\infty$, $\|d\|_1$, and $\|f^\star\|_\infty$ for some optimal integral solution $f^\star$. We will decide whether $\mathrm{OPT} < F$ and produce a corresponding flow.

Set

$$A = 10, \qquad \delta = P^{-10}, \qquad \tau = P^{-30},$$

and define the capped logarithm

$$\psi_\tau(s) = \begin{cases} -\log(s + \tau), & s \geq 0, \\ -\log \tau - s/\tau, & s < 0. \end{cases}$$

For

$$r = F - d^\top y, \qquad s = c - By, \qquad w_e = \frac{1}{\tau + \max\{s_e, 0\}},$$

use the potential

$$\Phi(y) = Ap \log r + \sum_e \psi_\tau(s_e), \qquad g = \nabla\Phi(y) = -\frac{Ap}{r} d + B^\top w. \tag{2}$$

We start at $y = 0$. The logarithm requires only $r > 0$; individual slacks may be slightly negative.

Lemma 2 (Scalar estimates)**.** *Let $w(s) = (\tau + \max\{s, 0\})^{-1}$. If $w(s)|z| \le \eta < 1$, then*

$$\frac{w(s)}{1+\eta} \le w(s - z) \le \frac{w(s)}{1-\eta}$$

and

$$\psi_\tau(s - z) \le \psi_\tau(s) + w(s)z + \frac{(w(s)z)^2}{2(1-\eta)}.$$

Proof. For $t \in [0, 1]$, the denominator defining $w(s - tz)$ differs from that defining $w(s)$ by at most $t|z|$; this proves the first display, including when the interval crosses zero. Since

$$\psi_\tau(s - z) - \psi_\tau(s) = z \int_0^1 w(s - tz)\, dt,$$

putting $\theta = w(s)|z|$ in the first bound gives

$$|w(s - tz) - w(s)| \le w(s) \frac{t\theta}{1 - t\theta}.$$

Thus the part beyond the linear term is at most $(w(s)z)^2 \int_0^1 t/(1 - t\theta)\, dt \le (w(s)z)^2/(2(1-\eta))$, proving the second display. □

For $a \perp \mathbf{1}$ and positive edge weights u, write

$$\rho(a, u) = \min_{z \notin \mathrm{span}\{\mathbf{1}\}} \frac{a^\top z}{\|\operatorname{diag}(u)Bz\|_1}.$$

Thresholding the coordinates of z shows that the same minimum is attained by a cut direction $z = \mathbf{1}_C$: both numerator and denominator are the integrals of $a(C_\theta)$ and $u(\partial C_\theta)$ over level sets.

Lemma 3 (Direction or flow)**.** *While $r > 0$, the following statements hold.*

1. *If a dual-feasible $y^\star$ satisfies $d^\top y^\star \ge F$, then $\rho(g, w) \le -1$.*

2. *If $\rho(g, w) > -\gamma$ for $\gamma = 1/20$, then there is $f \ge 0$ with $B^\top f = d$ and $c^\top f < F$. When the ratio of the largest to smallest w_e is polynomial, such an f is found deterministically in almost-linear time.*

Proof. For the first claim put $h = y^\star - y$ and $s^\star = c - By^\star \geq 0$. Then $d^\top h \geq r$ and $Bh = s - s^\star$. With $q_e = w_e(s_e - s_e^\star)$, every positive q_e is strictly less than 1, because then $s_e > s_e^\star \geq 0$ and

$$q_e = \frac{s_e - s_e^\star}{s_e + \tau} < 1.$$

If $Q_+ = \sum_e [q_e]_+$ and $Q_- = \sum_e [-q_e]_+$, then $Q_+ \leq p$ and

$$g^\top h \leq -Ap + Q_+ - Q_- \leq -(Q_+ + Q_-) = -\|\operatorname{diag}(w)Bh\|_1.$$

For the second claim, norm duality gives

$$-\rho(g,w) = \min_{B^\top x = g} \|\operatorname{diag}(w)^{-1}x\|_\infty. \tag{3}$$

This already supplies an x with $|x_e| < \gamma w_e$. Under the stated polynomial-ratio hypothesis, applying the cut inequality to S and $V \setminus S$ gives $|g(S)| < \gamma w(\partial S)$ for every nontrivial cut. Thus the deterministic approximate undirected-flow routine of [CGL+20, Corollary 6.10], with capacities γw and accuracy 1 and with signs reversed to our incidence convention, returns $B^\top x = g$ with $|x_e| \leq 2\gamma w_e$; its cut outcome is impossible. In either case fix such an x with $|x_e| \leq 2\gamma w_e$ and define

$$f = \frac{r}{Ap}(w - x).$$

Then $f \geq 0$ and, by (2), $B^\top f = d$. If $s_e < 0$, then $s_e(w_e - x_e) \leq 0$; if $s_e \geq 0$, then $s_e(w_e - x_e) \leq 1 + 2\gamma$. Consequently

$$c^\top f = d^\top y + s^\top f \leq F - r + \frac{1+2\gamma}{A} r < F.$$

$\square$

Min-ratio cut black box. We use the following local restatement of Definition 12 and Theorem 13 of Li–Wice [LW26].

Definition 4 (κ-approximate min-ratio cut structure; Definition 12 of [LW26])**.** *Let $H = (V, E, \bar{w})$ be an undirected multigraph whose active edge weights lie in $[1/U, U]$. Given an initial gradient $\bar{g} \perp \mathbf{1}$ with $\bar{g}_v \in [-U, U]$, an arbitrary initial potential $y \in \mathbb{R}^V$, and a detection threshold $\varepsilon_{\text{det}} > 0$, the structure supports*

$$\textsc{InsertEdge}(e), \quad \textsc{DeleteEdge}(e), \quad \textsc{UpdateGradient}(u, v, \xi), \quad \textsc{Potential}(v).$$

An insertion specifies $\bar{w}_e$, UpdateGradient *adds ξ to $\bar{g}_u$ and subtracts it from $\bar{g}_v$, and* Potential *returns y_v.*

At all times it maintains and exposes scalars $a \leq 0, b > 0$ and an implicit cut $C \subseteq V$ such that

$$a = \bar{g}(C), \qquad \bar{w}(\partial C) \leq b, \qquad \frac{a}{b} \leq \frac{1}{\kappa} \min_{\substack{S \subseteq V \\ \bar{w}(\partial S) \neq 0}} \frac{\bar{g}(S)}{\bar{w}(\partial S)}.$$

It also supports ToggleCut(η) *for $0 < \eta \leq 1/b$. This operation implicitly replaces y by $y + \eta \mathbf{1}_C$ and returns a set E' containing every edge $e = \{u, v\}$ for which*

$$\left| \bar{w}_e\big((y_u - y_v) - (y_u^{\text{last}} - y_v^{\text{last}})\big) \right| \geq \varepsilon_{\text{det}},$$

where y^{last} is the potential when e was inserted or last returned. The set E' may contain additional edges.

Theorem 5 (Theorem 13 of [LW26])**.** *If* $\log U = \widetilde{O}(1)$*, there is a deterministic structure from Definition 4 with* $\kappa = |V|^{o(1)}$*. Every operation takes amortized* $|V|^{o(1)} \log U$ *time. After* N *calls to* ToggleCut*, the total number of returned edge occurrences is at most* $|V|^{o(1)} N/\varepsilon_{\text{det}}$*.*

Our underlying undirected graph is connected, so the level-set identity above identifies the cut minimum in Definition 4 with $\rho(\bar{g}, \bar{w})$.

One ideal step. Let κ be the factor in Theorem 5, retain $\gamma = 1/20$, and put $\lambda = \gamma/\kappa$. First suppose that the maintained pair is the exact current pair (g, w), and let (a, b) and C be the tuple and implicit cut of Definition 4. If $a/b > -\lambda$, then $\rho(g, w) > -\gamma$ and Lemma 3 extracts a flow. Otherwise take

$$\eta = \frac{\lambda^2}{100|a|}, \qquad \Delta = \eta \mathbf{1}_C, \qquad y \leftarrow y + \Delta. \tag{4}$$

Since $|a|/b \geq \lambda$, Definition 4 permits this update, and

$$g^\top \Delta = -\frac{\lambda^2}{100}, \qquad D := \|\operatorname{diag}(w) B\Delta\|_1 \leq \frac{\lambda}{100}. \tag{5}$$

Lemma 6 (Progress)**.** *The update* (4) *preserves* $r > 0$ *and decreases* Φ *by at least* $\lambda^2/200$*.*

Proof. From (2) and (5),

$$\frac{Ap}{r} |d^\top \Delta| \leq |g^\top \Delta| + D < 1,$$

so the new residual is positive. Concavity of log and Lemma 2, applied edgewise with $|w_e(B\Delta)_e| \leq D < 1/2$, give

$$\Phi(y + \Delta) - \Phi(y) \leq g^\top \Delta + D^2 \leq -\frac{\lambda^2}{200}.$$

□

The cap in ψ_τ is used only in the next lemma. A literal logarithm would not give the polynomial weight range needed by the black box.

Lemma 7 (Range and certificate)**.** *Suppose* $\Phi(y) \leq \Phi(0)$ *and* $r \geq \delta/2$*. After fixing one vertex potential to zero,*

$$\max_v y_v - \min_v y_v < 2|V|P, \qquad \frac{1}{4P^2} \leq w_e \leq P^{30},$$

and

$$H := \frac{1}{\tau} \sum_e [-s_e]_+ \leq 130p \log P. \tag{6}$$

If also $r \leq \delta$*, then* $\mathrm{OPT} \geq F$*.*

Proof. Let $R = \max_v y_v - \min_v y_v$. If $R \geq 2|V|P$, a directed path from a maximum-potential vertex to a minimum-potential vertex contains an arc with $(By)_e \geq R/|V|$, hence $s_e \leq -R/(2|V|)$. Its contribution to ψ_τ is at least $R/(2|V|\tau)$, whereas every other contribution is at least $-\log(P+R+1)$. Since $\Phi(0) \leq Ap \log P$ and $r \geq \delta/2$,

$$\frac{R}{2|V|\tau} \leq Ap \log(2P/\delta) + p \log(P + R + 1). \tag{7}$$

For $R \geq 2|V|P$, the left side minus the last term is increasing: its derivative is at least $(2|V|\tau)^{-1} - p/(P+R+1) > 0$. At $R = 2|V|P$ it is already larger than $Ap\log(2P/\delta)$, because $\tau = P^{-30}$ and $P \geq p, |V|, 2$. This contradicts (7).

Thus $R < 2|V|P \leq 2P^2$. With one fixed vertex set to zero this also gives $|y_v| \leq R$ for every v, and every positive slack is at most $3P^2$. Every nonnegative slack contributes at least $-4\log P$, while every negative slack contributes at least $[-s_e]_+/\tau$. Therefore

$$H - 4p\log P \leq \sum_e \psi_\tau(s_e) \leq \Phi(0) - Ap\log r \leq Ap\log(2P/\delta),$$

which gives (6). The stated weight range follows immediately. Moreover,

$$|g_v| \leq \frac{Ap}{r}|d_v| + \sum_{e \text{ incident to } v} w_e \leq 20P^{12} + P^{31} \leq P^{40},$$

so the gradient coordinates have polynomial range as well.

Finally, integrality of the transshipment polytope makes OPT integral. For the optimal integral $f^\star$ fixed above,

$$\mathrm{OPT} = d^\top y + s^\top f^\star \geq F - r - P\sum_e [-s_e]_+ \geq F - \delta - P\tau H > F - 1.$$

Indeed, the total loss is at most $P^{-10} + 130P^{-28}\log P < 1$ for $P \geq 2$; hence $\mathrm{OPT} \geq F$. □

Implementation. Set $\varepsilon_{\mathrm{det}} = \lambda/1000$. Initially $y = 0$, so $r = F > \delta$ and $\Phi(y) = \Phi(0)$. Inductively, every block begins with $\Phi(y) \leq \Phi(0)$. First stop if $r \leq \delta$. Otherwise $r > \delta$, so Lemma 7 gives

$$\frac{1}{4P^2} \leq w_e \leq P^{30}, \qquad |g_v| \leq P^{40}.$$

Set $U = P^{40}$, query all vertex potentials, set

$$\widetilde{r} = r, \qquad \widetilde{w} = w, \qquad \widetilde{g} = -\frac{Ap}{\widetilde{r}}d + B^\top \widetilde{w},$$

and initialize the structure of Theorem 5 at the current potential y. The displayed bounds, $\widetilde{g} \perp \mathbf{1}$, and $\log U = O(\log P) = \widetilde{O}(1)$ verify all of its initial hypotheses.

Hold $\widetilde{r}$ fixed for a block of at most p toggles. At each query use the maintained tuple (a, b). If $a/b > -\lambda$, stop and extract a flow as below. Otherwise apply (4) with $(\widetilde{g}, \widetilde{w})$ in place of (g, w). The resulting ToggleCut returns an edge set. Process every returned edge, including false positives: DeleteEdge, remove its old endpoint-gradient contribution, use two Potential queries to compute its current weight, then InsertEdge and add the new endpoint contribution. The cut itself is never enumerated.

We prove by induction within the block that the exact potential never increases and $r \geq \delta/2$. The detector and Lemma 2 imply, before every toggle,

$$\frac{\widetilde{w}_e}{1+\varepsilon_{\mathrm{det}}} \leq w_e \leq \frac{\widetilde{w}_e}{1-\varepsilon_{\mathrm{det}}}. \tag{8}$$

For a non-stopping tuple, the scaled step satisfies

$$\widetilde{g}^\top \Delta = -\frac{\lambda^2}{100}, \qquad \|\operatorname{diag}(\widetilde{w})B\Delta\|_1 \le \frac{\lambda}{100}.$$

The identity defining $\widetilde{g}$ therefore gives

$$\frac{Ap}{\widetilde{r}}|d^\top \Delta| \le |\widetilde{g}^\top \Delta| + \|\operatorname{diag}(\widetilde{w})B\Delta\|_1 \le \frac{\lambda}{50}.$$

For every prefix of at most p toggles,

$$\frac{|r-\widetilde{r}|}{\widetilde{r}} \le \frac{\lambda}{50A} \le \frac{\lambda}{500}. \tag{9}$$

Together with (8), this yields

$$|(g-\widetilde{g})^\top \Delta| \le \frac{\lambda^2}{5000}, \qquad \|\operatorname{diag}(w)B\Delta\|_1 \le \frac{\lambda}{90}.$$

Hence the exact potential satisfies

$$\Phi(y+\Delta) - \Phi(y) \le -\frac{\lambda^2}{100} + \frac{\lambda^2}{5000} + \left(\frac{\lambda}{90}\right)^2 < -\frac{\lambda^2}{200}.$$

Also, since the block starts with $\widetilde{r} > \delta$, (9) gives $r > \delta/2$ throughout the block. This closes the induction and makes Lemma 7 applicable at every iterate. In particular every refreshed weight lies in $[1/U, U]$. Moreover, the fixed $\widetilde{r} > \delta$ and the stored weight range give $|\widetilde{g}_v| \le P^{40}$, while $\widetilde{g} \perp \mathbf{1}$; thus the data-structure hypotheses remain valid.

The stale ratio test is correct. If $\mathrm{OPT} \ge F$, let $y^\star$ be a dual optimum and set $h = y^\star - y$ and $s^\star = c - By^\star \ge 0$. For $\widetilde{q}_e = \widetilde{w}_e(s_e - s^\star_e)$ and $\widetilde{Q}_\pm = \sum_e [\pm \widetilde{q}_e]_+$, (8) and (9) give

$$\widetilde{Q}_+ \le (1+\varepsilon_{\mathrm{det}})p, \qquad \widetilde{g}^\top h \le -Ap\frac{r}{\widetilde{r}} + \widetilde{Q}_+ - \widetilde{Q}_- \le -(\widetilde{Q}_+ + \widetilde{Q}_-).$$

Thus $\rho(\widetilde{g}, \widetilde{w}) \le -1$, and Definition 4 gives $a/b \le -1/\kappa < -\lambda$. Consequently this case cannot stop at the ratio test.

Conversely, if $a/b > -\lambda$, then Definition 4 implies $\rho(\widetilde{g}, \widetilde{w}) > -\gamma$. Applying the deterministic routine used in Lemma 3 to $(\widetilde{g}, \widetilde{w})$ returns

$$B^\top x = \widetilde{g}, \qquad |x_e| \le 2\gamma \widetilde{w}_e.$$

The flow

$$f = \frac{\widetilde{r}}{Ap}(\widetilde{w} - x)$$

is nonnegative and satisfies $B^\top f = d$. For $s_e < 0$, its contribution to $s^\top(\widetilde{w} - x)$ is nonpositive because $\widetilde{w}_e - x_e \ge (1-2\gamma)\widetilde{w}_e > 0$. For $s_e \ge 0$, the contribution is at most $(1+2\gamma)(1+\varepsilon_{\mathrm{det}})$. Therefore

$$c^\top f \le F - r + \frac{\widetilde{r}}{A}(1+2\gamma)(1+\varepsilon_{\mathrm{det}}) \le F - r\left(1 - \frac{(1+2\gamma)(1+\varepsilon_{\mathrm{det}})}{A(1-\lambda/500)}\right) < F.$$

This congestion-flow computation is separate from Theorem 5.

Check r again after the block. At the first rebuild with $r \le \delta$, the preceding block began with $r_{\text{old}} = \widetilde{r} > \delta$, so

$$r \ge \left(1 - \frac{\lambda}{50A}\right) r_{\text{old}} > \frac{\delta}{2}.$$

Lemma 7 then certifies $\mathrm{OPT} \ge F$.

At every active iterate,

$$\Phi(y) \ge Ap\log(\delta/2) - 4p\log P = -O(p\log P),$$

whereas $\Phi(0) \le Ap\log P$. The exact decrease above therefore permits at most

$$T = O(p\lambda^{-2}\log P) = p^{1+o(1)} \tag{10}$$

toggles. There are $1 + \lceil T/p \rceil = p^{o(1)}$ rebuilds, each using $O(p)$ initialization operations. By Theorem 5, the total number of returned edge occurrences is at most

$$p^{o(1)}T/\varepsilon_{\text{det}} = p^{1+o(1)}.$$

Since every data-structure operation costs $p^{o(1)}\log U$, the T toggles, all rebuilds, and all refreshes take $p^{1+o(1)}$ total time. The one congestion-flow computation at a ratio stop has the same bound.

Theorem 8 (Threshold transshipment)**.** *Under the assumptions preceding (1), there is a deterministic $p^{1+o(1)}$-time algorithm which returns either a feasible f with $c^\top f < F$, or the correct assertion* $\mathrm{OPT} \ge F$*.*

Proof. Run the block algorithm above. A ratio stop returns the displayed feasible flow of cost below F, while a residual stop is certified by Lemma 7. If $\mathrm{OPT} \ge F$, the stale witness above forbids a ratio stop; if $\mathrm{OPT} < F$, the residual certificate forbids a residual stop. If neither stop occurred within the number of toggles in (10), the exact potential would be below its proved lower bound. Thus the algorithm terminates with one of the two correct outcomes, and the preceding operation count proves the stated running time. ☐

3 Cleanup

Let the input capacities be u_e and put $M = \sum_e u_e$. Discard zero-capacity arcs. If $M = 0$ or either terminal is isolated, return the zero flow; otherwise discard isolated nonterminals, after which the construction below has size $O(m)$. For an integer $k \in [0, M]$, form a transshipment graph H_k as follows. For each original arc $e = (u, v)$ add a vertex x_e and arcs $(u, x_e), (v, x_e)$, both of cost 1, and set

$$d_k(z) = k(\mathbf{1}_s - \mathbf{1}_t)_z + \sum_{e=(v,z)} u_e, \qquad d_k(x_e) = -u_e.$$

Add a root q, and for every other vertex z add both (q, z) and (z, q) with cost $M + 2$ and set $d_k(q) = 0$. The graph is strongly connected and always has an integral feasible transshipment: route every nonroot imbalance directly to or from q. This solution has cost $O(M^2)$, so positivity of all costs gives an optimal integral solution of norm $O(M^2)$. Thus one polynomial P satisfies all assumptions of Theorem 8 uniformly over k.

An s–t flow of value k corresponds to the root-free transshipment

$$h_{ux_e} = f_e, \qquad h_{vx_e} = u_e - f_e,$$

whose cost is M, and the converse follows from the demand equation at x_e. By integrality, if value k is infeasible then every optimum uses at least one unit of a root arc and has cost at least $M+2$. Thus Theorem 8 with $F = M+1$ decides feasibility of k. Binary search finds the maximum value in $m^{1+o(1)}$ total time.

It remains only to make the returned flow integral. For completeness, given any feasible fractional h, subtract $\lfloor h \rfloor$. The remaining demands are integral and every remaining coordinate lies in $[0,1)$. While its undirected support contains a cycle, push a signed circulation around that cycle, choosing the one of its two directions with nonpositive cost, until an edge becomes 0 or 1. Fold an edge reaching 1 into the integral part and update the residual integer demands. When the fractional support is a forest, a leaf equation forces its incident value to be integral; induction empties the forest. Hence this rounding preserves demands, nonnegativity, and does not increase cost. A link-cut tree maintains a fractional-support spanning forest, signed cost on a fundamental cycle, path addition, and the first 0/1 event in $O(\log p)$ amortized time per removed edge, so the rounding is near-linear.

On the final feasible target the threshold routine returns cost $< M+1$. After rounding, the integral cost is at most M, so no root arc is used; mapping back yields an exact maximum flow. This proves Theorem 1.

References

[CGL+20] Julia Chuzhoy, Yu Gao, Jason Li, Danupon Nanongkai, Richard Peng, and Thatchaphol Saranurak. A deterministic algorithm for balanced cut with applications to dynamic connectivity, flows, and beyond. *arXiv:1910.08025*, 2020.

[CKL+22] Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *FOCS*, 2022.

[Li26] Jason Li. Balancing weights, directed sparsification, and augmenting paths. *arXiv:2604.14633*, 2026.

[LW26] Jason Li and Alex Wice. Maximum flow without the outer IPM. *arXiv:2608.17384*, 2026.

[VDBCK+24] Jan van den Brand, Li Chen, Rasmus Kyng, Yang P. Liu, Simon Meierhans, Maximilian Probst Gutenberg, and Sushant Sachdeva. Almost-linear time algorithms for decremental graphs: min-cost flow and more via duality. In *FOCS*, 2024.