

Electrical Flows, Laplacian Matrices, and New Approaches to the Maximum Flow Problem


Aleksander Mądry



Maximum flow problem

Input: Directed graph G ,
integer **capacities** u_e ,
source s and **sink** t

value = net flow out of s



Here, value = 7

no overflow on arcs:
 $0 \leq f(e) \leq u(e)$

no leaks at all $v \neq s, t$

Task: Find a **feasible s-t flow** of **max value**

Breaking the $O(n^{3/2})$ barrier

Undirected graphs and **approx.** answers ($O(n^{3/2})$ barrier still holds here)

[M '10]: **Crude approx. of** max flow **value** in **close to linear** time

[CKMST '11]: **(1- ϵ)-approx.** to max flow in $\tilde{O}(n^{4/3}\epsilon^{-3})$ time

[LSR '13, S '13, KLOS '14]: **(1- ϵ)-approx.** in **close to linear** time

But: What about the **directed** and **exact** setting?

[M '13]: Exact $\tilde{O}(n^{10/7}) = \tilde{O}(n^{1.43})$ -time alg.

This week

(n = # of vertices, $\tilde{O}()$ hides polylog factors)

Breaking the $O(n^{3/2})$ barrier

Undirected graphs and **approx.** answers ($O(n^{3/2})$ barrier still holds here)

[M '10]: **Crude approx. of** max flow **value** in **close to linear** time

[CKMST '11]: **(1- ϵ)-approx.** to max flow in $\tilde{O}(n^{4/3}\epsilon^{-3})$ time

[LSR '13, S '13, KLOS '14]: **(1- ϵ)-approx.** in **close to linear** time

But: What about the **directed** and **exact** setting?

[M '13]: Exact $\tilde{O}(n^{10/7}) = \tilde{O}(n^{1.43})$ -time alg.

Today

This week

(n = # of vertices, $\tilde{O}()$ hides polylog factors)

New approach:

Bring linear-algebraic techniques into play

Idea: Probe the **global flow structure** of the graph by **solving linear systems**

How to relate **flow structure** to **linear algebra**?
(And why should it even help?)

Key object: Electrical flows

Electrical flows



Input: **Undirected** graph G ,
resistances r_e ,
source s and sink t

Principle of least energy

Electrical flow of value F :

The unique minimizer of the **energy**

$$E(\mathbf{f}) = \sum_e r_e f(e)^2$$

among all **s-t** flows \mathbf{f} of value F

Electrical flows = ℓ_2 -minimization

How to compute an electrical flow?

Bottom line:



$$\mathbf{L} \varphi = \chi_{s,t}$$

Electrical flow
computation

Solving a **Laplacian** system

Bad news: Solving a linear system can take $O(n^\omega) = O(n^{2.373})$

(Prohibitive!)

Key observation:

$BR^{-1}B^T$ is the **Laplacian** matrix \mathbf{L}
of the underlying graph

How to utilize it?

Result: Electrical flow is a **nearly-linear time** primitive

From electrical flows to **undirected** max flow

[CKMST '11]

Approx. undirected max flow via electrical flows


Assume: F^* known (via binary search)

- Treat edges as resistors of resistance **1**
- Compute electrical flow of value F^*
(This flow has **no leaks**, but **can overflow** some edges)
- To fix that: **Increase resistances** on the overflowing edges
Repeat (**hope**: it doesn't happen too often)

Surprisingly: This approach can be made work!

But: One needs to be careful how to fill in the blanks

We will do this now



Filling in the blanks

Recall: We are dealing with **undirected** graphs

From now on: All capacities are **1**, $m=O(n)$
and the value F^* of max flow is known

Electrical vs. maximum flows

Fix some resistances r and consider the elect. flow f_E of value F^*

We don't expect f_E to obey **all** capacity constraints
(i.e., we can have $|f_E(e)| \gg 1$ for some edge e)

Still, f_E obeys these constraints in a certain sense...

We have:

$$\sum_e r_e |f_E(e)| \leq \sum_e r_e$$

In other words: Capacity constraints are preserved on **average** (weighted wrt to r_e s)

Proof:



Electrical vs. maximum flows

This gives rise to a **very fast** algorithm for the following task:

‘Feasibility on average’:

Given weights \mathbf{w} compute a flow \mathbf{f} of value \mathbf{F}^* s.t.

$$\sum_e w_e |\mathbf{f}(e)| \leq \sum_e w_e$$

Key point: We already know how to make such a crude algorithm useful to us!

Multiplicative weights update method

[FS '97, PST '95, AHK '05]

‘Technique for turning weak algorithms
into strong ones’

In our setting:

Crude algorithm computing ‘feasible on average’ flows



(1- ϵ)-approx. max flow

[(1+ ϵ)-approx. feasibility everywhere]

How does this method work?

Underlying idea

Crude algorithm

Maintain weights w

(Initially, all weights $w_e=1$)

A

w^0

f^1

feasible on average

A

w^1

Update weights
(based on f^1)

A

f^2

Update weights
(based on f^2)



w^3



(Process continues for **N** rounds)

At the end: Return the average of all f 's
(This is still a flow of value **F^***)

Updating weights



⋮

w^{i-1}



f^i



w^i




⋮



Want this term to be between 1 and $1+\epsilon$

Update step: For each e
 $w_e^i \leftarrow w_e^{i-1} (1+\epsilon |f^i(e)| / \rho_i)$

Maximum congestion in f^i
 $\rho_i = \max_e |f^i(e)|$



Updating weights



⋮

w^{i-1}

Weights w^{i-1}



f^i



w^i



⋮

Update step: For each e
 $w_e^i \leftarrow w_e^{i-1}(1 + \varepsilon |f^i(e)| / \rho_i)$

Underlying dynamics:

Edge e suffers large overflow $\rightarrow w_e$ grows rapidly

Average overflow small $\rightarrow \sum_e w_e$ grows slowly






No edge suffers large overflow **too often**

\rightarrow **averaging out** yields (almost) no overflow

Updating weights

⋮
 w^{i-1}

Weights w^{i-1}



⋮

Update step: For each e
 $w_e^i \leftarrow w_e^{i-1} (1 + \epsilon |f^i(e)| / \rho_i)$

Width $\rho = \max_i \rho_i$

[AHK '05]: It suffices to repeat this step $N = \tilde{O}(\rho \epsilon^{-2})$ times to get a **(1- ϵ)-approx** to max flow



Think: ρ measures the **electrical vs. max** flow discrepancy

Note: Linear dependence on ρ is unavoidable


Updating weights

⋮
 w^{i-1}

Weights w^{i-1}



Update step: For each e
 $w_e^i \leftarrow w_e^{i-1} (1 + \epsilon |f^i(e)| / \rho_i)$



⋮

Width $\rho = \max_i \rho_i$

[AHK '05]: It suffices to repeat this step $N = \tilde{O}(\rho \epsilon^{-2})$ times to get a $(1-\epsilon)$ -approx to max flow

Details:



Bottom line:



Electrical flow primitive gives us the crude algorithm

We can use MWU framework
to fill in our blanks!

Our algorithm

- Treat edges as resistors of resistance $r_e=1$
 - Compute electrical flow \mathbf{f} of value \mathbf{F}^*
 - **Increase resistances** on overflowing edges
- Repeat

Our algorithm

- Treat edges as resistors of resistance $r_e=1$
- Compute electrical flow \mathbf{f} of value F^*
- **Increase resistances:** for each e ,

$$r_e^i \leftarrow r_e^{i-1}(1+\varepsilon|f^i(e)|/\rho_i)$$

Repeat **$N=\tilde{O}(\rho\varepsilon^{-2})$ times**

- **At the end:** Take an **average** of all the flows as the final answer

- Resistances r_e evolve as weights w_e
- Convergence condition: “execute N rounds”

Our algorithm

- Treat edges as resistors of resistance $r_e=1$
- Compute electrical flow \mathbf{f} of value F^*
- **Increase resistances:** for each e ,

$$r_e^i \leftarrow r_e^{i-1}(1+\varepsilon|f^i(e)|/\rho_i)$$

Repeat **$N=\tilde{O}(\rho\varepsilon^{-2})$ times**

- **At the end:** Take an **average** of all the flows as the final answer

Result: This algorithm gives us an **$(1-\varepsilon)$ -approx.** max flow in **$\tilde{O}(\rho\varepsilon^{-2})\cdot\tilde{O}(n) = \tilde{O}(n\rho\varepsilon^{-2})$** time

Crucial question: How large the worst-case overflow ρ can be?

Our question: Let f be an elect. flow of value F^* wrt resist. r_e
How large $\rho = \max_e |f(e)|$ can be?

In general: ρ can be very large
(**Think:** one edge having an extremely small resistance)

Fix: Regularize the resistances with a uniform distribution
$$r_e' \leftarrow r_e + \varepsilon \sum |r|_1 / m$$

Can show: ρ is bounded by $O(n^{1/2} \varepsilon^{-1})$ then

Proof:



This gives a **(1- ε)-approx.** $\tilde{O}(n^{3/2} \varepsilon^{-3})$ -time algorithm

Going beyond the $\tilde{O}(n^{3/2})$ Barrier


Speeding up our algorithm

Running time is dominated by $\approx \rho$ elect. flow computations

Can we improve our $O(n^{1/2} \epsilon^{-1})$ bound on ρ ?

Not really...

$\approx n^{1/2}$ paths with $\approx n^{1/2}$ vertices each



one edge

Speeding up our algorithm


Running time is dominated by $\approx \rho$ elect. flow computations

Can we improve our $O(n^{1/2} \epsilon^{-1})$ bound on ρ ?

Not really...

$\approx n^{1/2}$ paths with $\approx n^{1/2}$ vertices each

Max flow:



$F^* \approx n^{1/2}$

one edge


Speeding up our algorithm

Running time is dominated by $\approx \rho$ elect. flow computations


Can we improve our $O(n^{1/2} \epsilon^{-1})$ bound on ρ ?

Not really...

Electr. flow:




Speeding up our algorithm



Key observation: If we remove this bad edge...

→ The **max flow** does not change much

Speeding up our algorithm



Key observation: If we remove this bad edge...

- The **max flow** does not change much
- But the resulting **electrical flow** is much better behaved!

Can we turn this observation into an algorithmic idea?

Speeding up our algorithm

Idea: Let our electrical flow oracle **self-enforce** a smaller overflow $\rho' \ll \rho$

Modification of the oracle: If the computed electrical flow has some edge e flow more than ρ' :

- **Remove** this edge from the graph (permanently)
- **Recompute** the electrical flow

Note: If this oracle always **successfully** terminates, its effective overflow is ρ'

Speeding up our algorithm

Crucial question: What is the right setting of ρ' ?

- We want ρ' to be as small as possible
- But if it becomes too small the edge removal might be too aggressive and cut too many of them

Sweet spot: $\rho' \approx n^{1/3}$

- Key reason:** Removal of edges that flow a lot
- significantly increases the **energy** of the electr. flow
 - But perturbs the **max flow** only slightly

Speeding up our algorithm

Our potential: The energy $E_r(\mathbf{f})$ of the electrical flow \mathbf{f} wrt current resistances \mathbf{r}

Can show:

- $E_r(\mathbf{f})$ is not too small initially and cannot become too large
 ↑ (as long as we remove no more than $\approx \epsilon F^*$ edges)
- As the resistances only increase, $E_r(\mathbf{f})$ never decreases

This makes $E_r(\mathbf{f})$ a convenient potential

Need to show: Removal of an overflowing edge increases $E_r(\mathbf{f})$ significantly

Speeding up our algorithm

Need to show: Removal of an overflowing edge increases $E_r(\mathbf{f})$ significantly

Fact: If an edge e contributes a δ -fraction of energy then removing it increases $E_r(\mathbf{f})$ by a factor of $1+\Omega(\delta)$

Further: If an edge e flows at least ρ' in \mathbf{f} then its energy contribution is $\Delta=\Omega(\epsilon(\rho')^2/n)$



Speeding up our algorithm

Need to show: Removal of an overflowing edge increases $E_r(\mathbf{f})$ significantly

Fact: If an edge e contributes a δ -fraction of energy then removing it increases $E_r(\mathbf{f})$ by a factor of $1+\Omega(\delta)$

Further: If an edge e flows at least ρ' in \mathbf{f} then its energy contribution is $\Delta=\Omega(\varepsilon(\rho')^2/n)$

Putting it all together: We can have $\leq \tilde{O}(\Delta^{-1})=\tilde{O}(n/\varepsilon(\rho')^2)$ edge removals before $E_r(\mathbf{f})$ grows by too much

Taking $\rho' \approx n^{1/3}\varepsilon^{-1}$ makes $\tilde{O}(\Delta^{-1})=\tilde{O}(\varepsilon n^{1/3})$ be smaller than $\varepsilon F^* \leq \varepsilon \rho'$ as needed

This gives the $\tilde{O}(n^{4/3}\varepsilon^{-3})$ -time $(1-\varepsilon)$ -approx. algorithm

Thank you

Tomorrow: Computing an exact max flow