Distributed Minimum Weight Cycle Approximation

Gopinath Mishra

Yi-Jun Chang Yanyu Chen

Hung Thuan Nguyen Mingyang Yang

Dipan Dey

Yonggang Jiang

https://arxiv.org/abs/2603.25368

Minimum weight cycle

Also known as the girth

$\mathrm{OPT} =$ length of the minimum weight cycle (MWC).

Sum of all edge weights in the cycle

Minimum weight cycle

Also known as the girth

OPT = length of the minimum weight cycle (MWC).

Sum of all edge weights in the cycle

Cycle C is ℓ-approximation of MWC ⟷ Length of $C \leq \ell \cdot \mathrm{OPT}$

△

Goal: Compute this in the CONGEST model of distributed computing.

Prior work

- Minimum weight cycle (MWC) and replacement paths (RP) were recently studied by Manoharan and Ramachandran in the CONGEST model.

Manoharan and Ramachandran [PODC 2024]

Manoharan and Ramachandran [SIROCCO 2024]

Prior work

- Minimum weight cycle (MWC) and replacement paths (RP) were recently studied by Manoharan and Ramachandran in the CONGEST model.

Manoharan and Ramachandran [PODC 2024]

Manoharan and Ramachandran [SIROCCO 2024]

- MWC and RP have some similarities.

In CONGEST, they are particularly interesting because their optimal round complexities do not fall into the common complexity classes: $\tilde{\Theta}(D)$, $\tilde{\Theta}(\sqrt{n} + D)$, $\tilde{\Theta}(n)$.

Prior work

- Minimum weight cycle (MWC) and replacement paths (RP) were recently studied by Manoharan and Ramachandran in the CONGEST model.

Manoharan and Ramachandran [PODC 2024]

Manoharan and Ramachandran [SIROCCO 2024]

- MWC and RP have some similarities.

In CONGEST, they are particularly interesting because their optimal round complexities do not fall into the common complexity classes: $\tilde{\Theta}(D)$, $\tilde{\Theta}(\sqrt{n} + D)$, $\tilde{\Theta}(n)$.

- $\tilde{\Theta}\left(n^{2/3} + D\right)$ is the tight round complexity bound for unweighted directed RP.

Chang, Chen, Dey, Mishra, Nguyen and Sanchez [PODC 2025]

Variants

- There are four variants of the MWC problem:

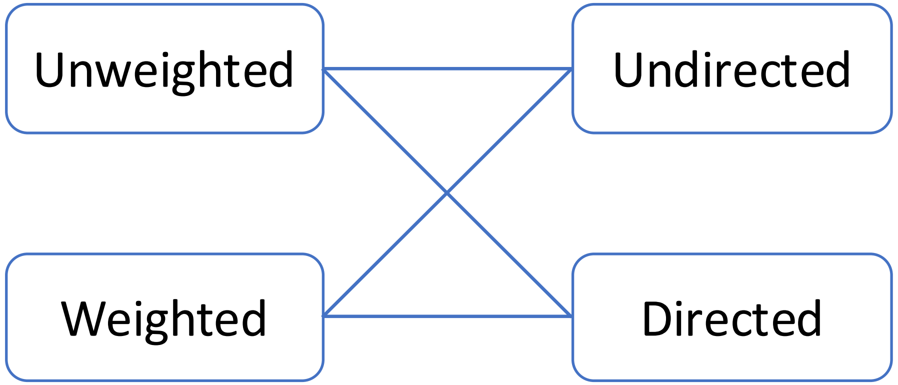

Variants

- There are four variants of the MWC problem:

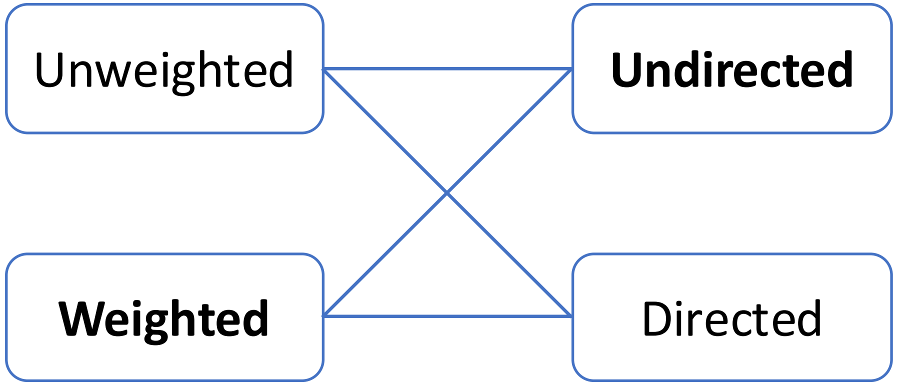

- Our focus: Weighted + Undirected.

We briefly discuss the remaining cases near the end of the talk.

Previous bounds

- There are four variants of the MWC problem:

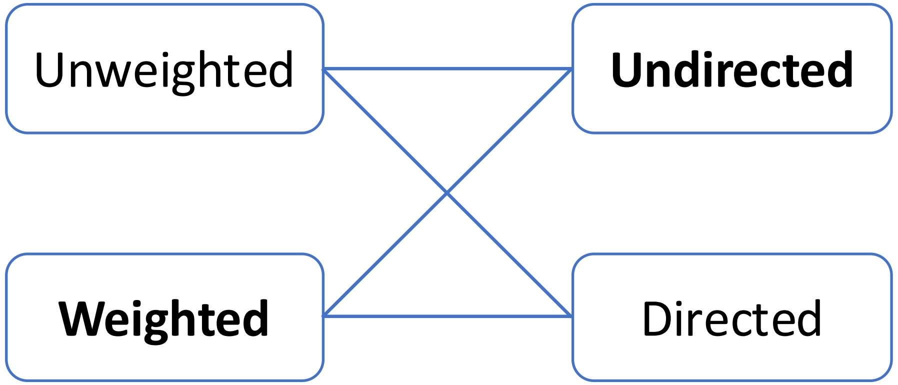

- Our focus: Weighted + Undirected.

Prior work:

- $(2+\epsilon)$-approximation: $\tilde{O}\left(n^{2/3}+D\right)$
- $(2-\epsilon)$-approximation: $\Omega(n)$
- Any approximation: $\Omega(\sqrt{n})$

Manoharan and Ramachandran [PODC 2024]

Main questions

- There are four variants of the MWC problem:

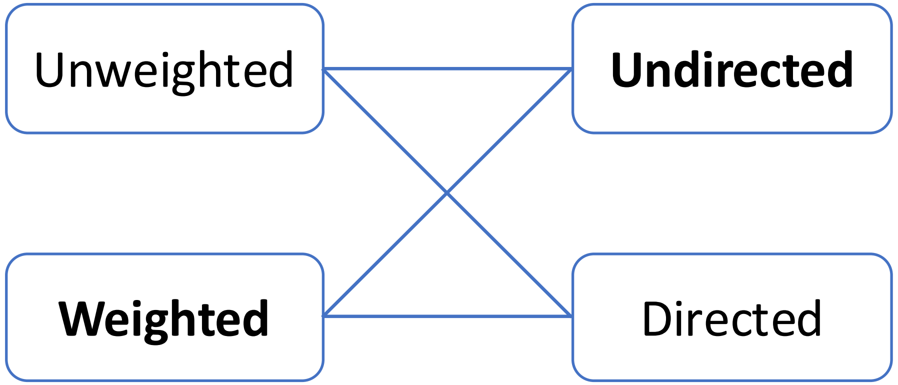

Main questions:

- Can we close this gap?
- What is the tight round-approximation tradeoff?

- Our focus: Weighted + Undirected.

Prior work:

- $(2+\epsilon)$-approximation: $\tilde{O}\left(n^{2/3}+D\right)$
- $(2-\epsilon)$-approximation: $\Omega(n)$
- Any approximation: $\Omega(\sqrt{n})$

Manoharan and Ramachandran [PODC 2024]

Improving the approximation factor by $-\Theta(1/\log n)$ only increases the round complexity by a constant factor.

Our results

New upper bound. $(k+1)$-approximation: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **real** $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$.

Prior work:

- $(2+\epsilon)$-approximation: $\tilde{O}\left(n^{2/3} + D\right)$
- $(2-\epsilon)$-approximation: $\Omega(n)$
- Any approximation: $\Omega(\sqrt{n})$

Manoharan and Ramachandran [PODC 2024]

Improving the approximation factor by $-\Theta(1/\log n)$ only increases the round complexity by a constant factor.

Our results

New upper bound. $(k+1)$-approximation: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **real** $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$.

Improving the round complexity at the cost of a worse approximation ratio.

Prior work:

- $(2+\epsilon)$-approximation: $\tilde{O}\left(n^{2/3} + D\right)$
- $(2-\epsilon)$-approximation: $\Omega(n)$
- Any approximation: $\Omega(\sqrt{n})$

Manoharan and Ramachandran [PODC 2024]

Improving the approximation factor by $-\Theta(1/\log n)$ only increases the round complexity by a constant factor.

Our results

New upper bound. $(k+1)$-approximation: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **real** $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$.

New lower bound. $(k+1-\epsilon)$-approximation: $\tilde{\Omega}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **integer** $k \geq 1$.

Conditioning on the Erdos girth conjecture.

Higher lower bound

Prior work:

- $(2+\epsilon)$-approximation: $\tilde{O}\left(n^{2/3} + D\right)$
- $(2-\epsilon)$-approximation: $\Omega(n)$
- Any approximation: $\Omega(\sqrt{n})$

Manoharan and Ramachandran [PODC 2024]

Improving the approximation factor by $-\Theta(1/\log n)$ only increases the round complexity by a constant factor.

Our results

New upper bound. $(k+1)$-approximation: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **real** $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$.

They **match** at $k = 2, 3, 4, 5, \ldots$

New lower bound. $(k+1-\epsilon)$-approximation: $\tilde{\Omega}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **integer** $k \geq 1$.

Conditioning on the Erdos girth conjecture.

The guiding tradeoff

New upper bound. $(k+1)$-approximation: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **real** $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$.

They **match** at $k = 2, 3, 4, 5, \ldots$

New lower bound. $(k+1-\epsilon)$-approximation: $\tilde{\Omega}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **integer** $k \geq 1$.

Both upper and lower bounds are due to this congestion-dilation tradeoff.

Congestion: the amount of information needed to cross some edge.

Dilation: the number of hops between the sender and the receiver of some message.

The guiding tradeoff:

- Congestion: $\alpha^{1+1/k}$
- Dilation: n/α

$$\min_{\alpha} \left\{\alpha^{1+1/k} + n/\alpha\right\} = \Theta\left(n^{(k+1)/(2k+1)}\right)$$

Lower bound

The existence is guaranteed by the Erdos girth conjecture.

- H = any α-vertex graph with girth $\geq 2k + 2$ and $\Omega\left(\alpha^{1+1/k}\right)$ edges.

Lower bound

The existence is guaranteed by the Erdos girth conjecture.

↓

- H = any α-vertex graph with girth $\geq 2k+2$ and $\Omega(\alpha^{1+1/k})$ edges.
- Create two subgraphs G_1 and G_2 of H.

Both subgraphs encode an $|E(H)| \in \Omega(\alpha^{1+1/k})$-bit string.

- $1 \leftrightarrow \text{existence of an edge}$
- $0 \leftrightarrow \text{non-existence of an edge}$

Lower bound

The existence is guaranteed by the Erdos girth conjecture.

- H = any α-vertex graph with girth $\geq 2k+2$ and $\Omega(\alpha^{1+1/k})$ edges.
- Create two subgraphs G_1 and G_2 of H.

Both subgraphs encode an $|E(H)| \in \Omega(\alpha^{1+1/k})$-bit string.

- $1 \leftrightarrow$ existence of an edge
- $0 \leftrightarrow$ non−existence of an edge

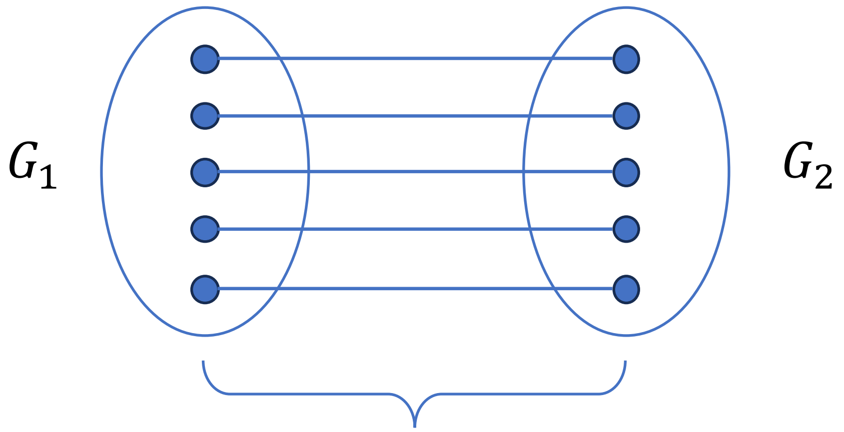

For each vertex in G_1 and its corresponding vertex in G_2, add a zero-weight path of $\Theta(n/\alpha)$ hops connecting them.

Lower bound

The existence is guaranteed by the Erdos girth conjecture.

- H = any α-vertex graph with girth $\geq 2k+2$ and $\Omega(\alpha^{1+1/k})$ edges.
- Create two subgraphs G_1 and G_2 of H.

Both subgraphs encode an $|E(H)| \in \Omega(\alpha^{1+1/k})$-bit string.

- $1 \leftrightarrow$ existence of an edge
- $0 \leftrightarrow$ non−existence of an edge

Add an infinite-weight overlay tree over this graph to shrink the diameter to $O(\log n)$.

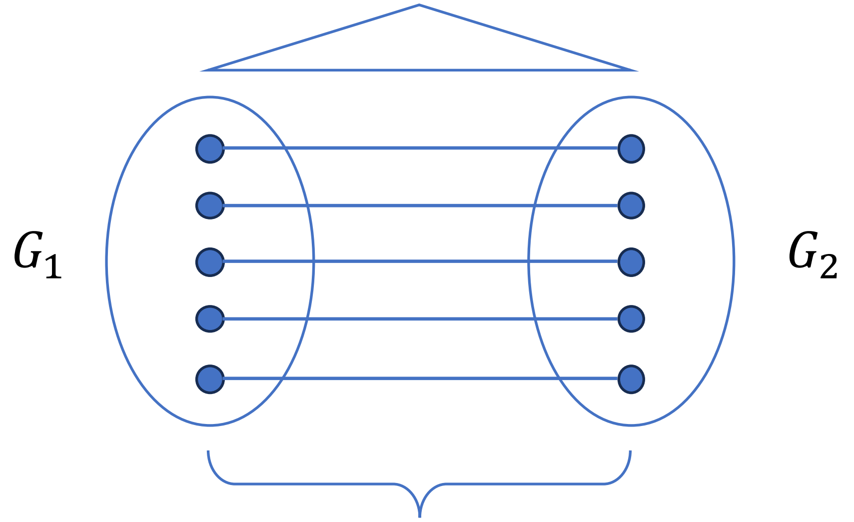

For each vertex in G_1 and its corresponding vertex in G_2, add a zero-weight path of $\Theta(n/\alpha)$ hops connecting them.

Lower bound

The existence is guaranteed by the Erdos girth conjecture.

↓

- H = any α-vertex graph with girth $\geq 2k + 2$ and $\Omega\left(\alpha^{1+1/k}\right)$ edges.
- Create two subgraphs G_1 and G_2 of H.

Both subgraphs encode an $|E(H)| \in \Omega\left(\alpha^{1+1/k}\right)$-bit string.
- $1 \leftrightarrow \text{existence of an edge}$
- $0 \leftrightarrow \text{non-existence of an edge}$

Claim: $(k + 1 - \epsilon)$-approximate MWC requires solving set-disjointness over the two strings.

↑

Distinguishing between $\text{OPT} \leq 2$ and $\text{OPT} \geq 2\text{k} + 2$

Add an infinite-weight overlay tree over this graph to shrink the diameter to $O(\log n)$.

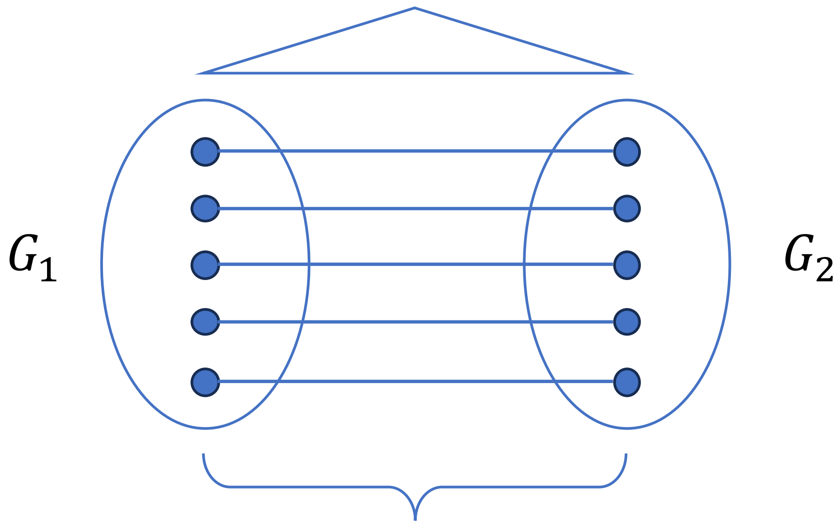

For each vertex in G_1 and its corresponding vertex in G_2, add a zero-weight path of $\Theta(n/\alpha)$ hops connecting them.

$\exists e \in E(H)$ in both G_1 and G_2 → $\mathrm{OPT} \leq 2$

Lower bound

The existence is guaranteed by the Erdos girth conjecture.

- H = any α-vertex graph with girth $\geq 2k+2$ and $\Omega(\alpha^{1+1/k})$ edges.
- Create two subgraphs G_1 and G_2 of H.

Both subgraphs encode an $|E(H)| \in \Omega(\alpha^{1+1/k})$-bit string.

- $1 \leftrightarrow$ existence of an edge
- $0 \leftrightarrow$ non−existence of an edge

Claim: $(k+1-\epsilon)$-approximate MWC requires solving set-disjointness over the two strings.

Distinguishing between $\mathrm{OPT} \leq 2$ and $\mathrm{OPT} \geq 2k+2$

Add an infinite-weight overlay tree over this graph to shrink the diameter to $O(\log n)$.

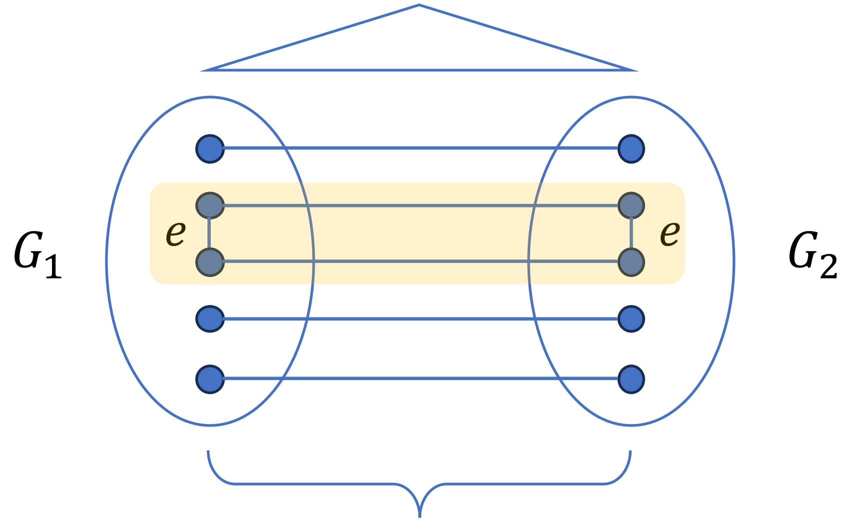

For each vertex in G_1 and its corresponding vertex in G_2, add a zero-weight path of $\Theta(n/\alpha)$ hops connecting them.

Lower bound

$\exists e \in E(H)$ in both G_1 and G_2 → $\mathrm{OPT} \leq 2$

$\nexists e \in E(H)$ in both G_1 and G_2 → $\mathrm{OPT} \geq 2k+2$

The existence is guaranteed by the Erdos girth conjecture.

- H = any α-vertex graph with girth $\geq 2k+2$ and $\Omega(\alpha^{1+1/k})$ edges.
- Create two subgraphs G_1 and G_2 of H.

Both subgraphs encode an $|E(H)| \in \Omega(\alpha^{1+1/k})$-bit string.

- $1 \leftrightarrow$ existence of an edge
- $0 \leftrightarrow$ non−existence of an edge

Claim: $(k+1-\epsilon)$-approximate MWC requires solving set-disjointness over the two strings.

Distinguishing between $\mathrm{OPT} \leq 2$ and $\mathrm{OPT} \geq 2k+2$

Add an infinite-weight overlay tree over this graph to shrink the diameter to $O(\log n)$.

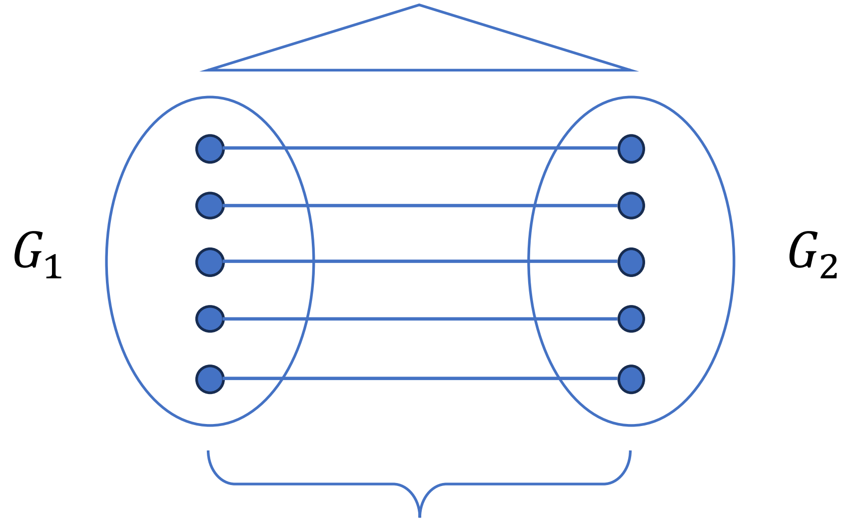

For each vertex in G_1 and its corresponding vertex in G_2, add a zero-weight path of $\Theta(n/\alpha)$ hops connecting them.

Lower bound

Two ways to solve set-disjointness:

- Use the overlay tree – congestion: $\Omega\left(\alpha^{1+\frac{1}{k}}/\log n\right)$.
- Go along the long paths – dilation: $\Omega(n/\alpha)$.

The existence is guaranteed by the Erdos girth conjecture.

- H = any α-vertex graph with girth $\geq 2k+2$ and $\Omega\left(\alpha^{1+1/k}\right)$ edges.
- Create two subgraphs G_1 and G_2 of H.

Both subgraphs encode an $|E(H)| \in \Omega\left(\alpha^{1+1/k}\right)$-bit string.

- $1 \leftrightarrow$ existence of an edge
- $0 \leftrightarrow$ non−existence of an edge

Claim: $(k+1-\epsilon)$-approximate MWC requires solving set-disjointness over the two strings.

Distinguishing between $\mathrm{OPT} \leq 2$ and $\mathrm{OPT} \geq 2k+2$

Add an infinite-weight overlay tree over this graph to shrink the diameter to $O(\log n)$.

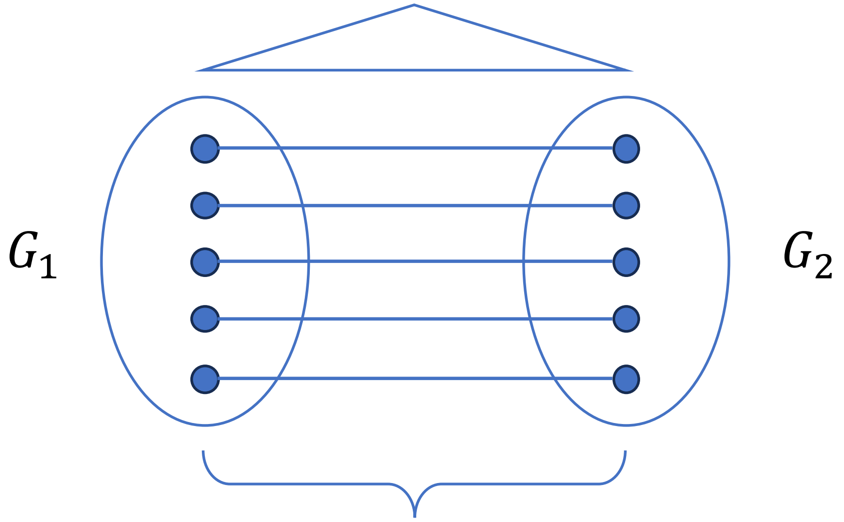

For each vertex in G_1 and its corresponding vertex in G_2, add a zero-weight path of $\Theta(n/\alpha)$ hops connecting them.

Lower bound

Two ways to solve set-disjointness:

- Use the overlay tree – congestion: $\Omega\left(\alpha^{1+\frac{1}{k}}/\log n\right)$.
- Go along the long paths – dilation: $\Omega(n/\alpha)$.

$\widetilde{\Omega}\left(n^{(k+1)/(2k+1)} + D\right)$ lower bound

Claim: $(k+1-\epsilon)$-approximate MWC requires solving set-disjointness over the two strings.

Distinguishing between $\mathrm{OPT} \leq 2$ and $\mathrm{OPT} \geq 2k+2$

The guiding tradeoff:

- Congestion: $\alpha^{1+1/k}$
- Dilation: n/α

$$\min_{\alpha} \left\{\alpha^{1+1/k} + n/\alpha\right\} = \Theta\left(n^{(k+1)/(2k+1)}\right)$$

Upper bound

- **Key idea:** Using low-diameter decompositions.

Partitioning the vertex set into low-diameter clusters.

Upper bound

- **Key idea:** Using low-diameter decompositions.

Partitioning the vertex set into low-diameter clusters.

- **Intuition:**
  - Suppose $\mathrm{OPT} = 2d$.
  - Suppose there is a cluster of radius at most $(k+1)d$ containing the MWC.

Distance from the cluster center to any point is at most $(k+1)d$.

Not just vertices but also any intermediate point on an edge.

Upper bound

- **Key idea:** Using low-diameter decompositions.

Partitioning the vertex set into low-diameter clusters.

- **Intuition:**
  - Suppose $\mathrm{OPT} = 2d$.
  - Suppose there is a cluster of radius at most $(k+1)d$ containing the MWC.

Consider the SSSP tree T of the cluster rooted at the cluster center s.

Upper bound

- **Key idea:** Using low-diameter decompositions.

  Partitioning the vertex set into low-diameter clusters.

- **Intuition:**
  - Suppose $\mathrm{OPT} = 2d$.
  - Suppose there is a cluster of radius at most $(k+1)d$ containing the MWC.

Consider the SSSP tree T of the cluster rooted at the cluster center s.

There is a non-tree-edge $e = \{u, v\}$ in the cluster.

Upper bound

- **Key idea:** Using low-diameter decompositions.

  Partitioning the vertex set into low-diameter clusters.

- **Intuition:**
  - Suppose $\mathrm{OPT} = 2d$.
  - Suppose there is a cluster of radius at most $(k+1)d$ containing the MWC.

Consider the SSSP tree T of the cluster rooted at the cluster center s.

There is a non-tree-edge $e = \{u, v\}$ in the cluster.

Can find a $(k+1)$-approximation of MWC just by considering cycles with only one non-tree-edge.

Weight of cycle $\leq 2 \cdot (k+1)d \leq (k+1) \cdot \mathrm{OPT}$

Upper bound

Next: Construct such a clustering.

- **Key idea:** Using low-diameter decompositions.

  Partitioning the vertex set into low-diameter clusters.

- **Intuition:**
  - Suppose $\mathrm{OPT} = 2d$.
  - Suppose there is a cluster of radius at most $(k+1)d$ containing the MWC.

Consider the SSSP tree T of the cluster rooted at the cluster center s.

There is a non-tree-edge $e = \{u, v\}$ in the cluster.

Can find a $(k+1)$-approximation of MWC just by considering cycles with only one non-tree-edge.

Weight of cycle $\leq 2 \cdot (k+1)d \leq (k+1) \cdot \mathrm{OPT}$

MPX clustering

Miller, Peng, and Xu [SPAA 2013]

Weighted MPX clustering:

- Each vertex v grows an SSSP tree at time $-\delta_v$.
  - δ_v follows the exponential distribution with mean $\frac{kd}{\ln n}$.
- Each vertex joins the cluster of the first SSSP tree that hits it.
  - The propagation time of SSSP tree along an edge equals its weight.

MPX clustering

Miller, Peng, and Xu [SPAA 2013]

Weighted MPX clustering:

- Each vertex v grows an SSSP tree at time $-\delta_v$.
  - δ_v follows the exponential distribution with mean $\frac{kd}{\ln n}$.
- Each vertex joins the cluster of the first SSSP tree that hits it.
  - The propagation time of SSSP tree along an edge equals its weight.

Claim: With probability $\Omega\left(n^{-1/k}\right)$, the clustering has the desired properties.

↓

Repeating for $O\left(n^{1/k}\right)$ times suffice for $(k+1)$-approximation of MWC.

Implementation

Known tradeoff for $(1+\epsilon)$-approximate undirected weighted SSSP:

- Congestion: α
- Dilation: n/α

Becker, Forster, Karrenbauer, Lenzen [DISC 2017]

Claim: With probability $\Omega\left(n^{-1/k}\right)$, the clustering has the desired properties.

↓

Repeating for $O\left(n^{1/k}\right)$ times suffice for $(k+1)$-approximation of MWC.

Implementation

$\tilde{O}\left(n^{(k+1)/2k} + D\right)$

Known tradeoff for $(1+\epsilon)$-approximate undirected weighted SSSP:

- Congestion: α
- Dilation: n/α

Becker, Forster, Karrenbauer, Lenzen [DISC 2017]

Tradeoff for $(k+1)$-approximation of MWC:

- Congestion: $\alpha \cdot n^{1/k}$
- Dilation: n/α

Claim: With probability $\Omega\left(n^{-1/k}\right)$, the clustering has the desired properties.

Repeating for $O\left(n^{1/k}\right)$ times suffice for $(k+1)$-approximation of MWC.

Caveat

Known tradeoff for $(1+\epsilon)$-approximate undirected weighted SSSP:

- Congestion: α
- Dilation: n/α

Becker, Forster, Karrenbauer, Lenzen [DISC 2017]

We need a the approximate SSSP to be tree-like.

Rozhoň, Haeupler, Martinsson, Grunau, Zuzic [STOC 2023]

Implementation

We are not done yet!

How to tighten the upper bound?

$\tilde{O}(n^{(k+1)/2k} + D)$

Tradeoff for $(k+1)$-approximation of MWC:
- Congestion: $\alpha \cdot n^{1/k}$
- Dilation: n/α

$\tilde{O}(n^{(k+1)/(2k+1)} + D)$

The guiding tradeoff:
- Congestion: $\alpha^{1+1/k}$
- Dilation: n/α

Implementation

We are not done yet!

How to tighten the upper bound?

Solution: Handle long-hop and short-hop cycles separately, using n/α as the threshold.

$\tilde{O}\left(n^{(k+1)/2k} + D\right)$

Tradeoff for $(k+1)$-approximation of MWC:

- Congestion: $\alpha \cdot n^{1/k}$
- Dilation: n/α

$\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$

The guiding tradeoff:

- Congestion: $\alpha^{1+1/k}$
- Dilation: n/α

Short-hop cycles

For cycles of at most n/α hops, using rounding + BFS, the congestion is reduced to $n^{1/k}$.

$\leq n^{(k+1)/(2k+1)}$ when $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$

$\tilde{O}\left(n^{(k+1)/2k} + D\right)$

Tradeoff for $(k+1)$-approximation of MWC:

- Congestion: $\alpha \cdot n^{1/k}$
- Dilation: n/α

$\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$

The guiding tradeoff:

- Congestion: $\alpha^{1+1/k}$
- Dilation: n/α

Long-hop cycles

What about long-hop cycles?

It suffices to improve the number of SSSP computation from $O(n^{1/k})$ to $O(\alpha^{1/k})$.

Solution sketch: Modify MPX to initiate an SSSP tree only from the $\tilde{O}(\alpha)$ randomly sampled vertices. Why this is fine? Intuitively, for long-hop cycles, with high probability there will be $\tilde{\Omega}(1)$ sampled vertices on the cycle.

A lot of technical details omitted…

$\tilde{O}(n^{(k+1)/2k} + D)$

↑

Tradeoff for $(k+1)$-approximation of MWC:

- Congestion: $\alpha \cdot n^{1/k}$
- Dilation: n/α

$\tilde{O}(n^{(k+1)/(2k+1)} + D)$

↑

The guiding tradeoff:

- Congestion: $\alpha^{1+1/k}$
- Dilation: n/α

Improving the approximation factor by $-\Theta(1/\log n)$ only increases the round complexity by a constant factor.

Summary

New upper bound. $(k+1)$-approximation: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **real** $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$.

They **match** at $k = 2, 3, 4, 5, \ldots$

New lower bound. $(k+1-\epsilon)$-approximation: $\tilde{\Omega}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **integer** $k \geq 1$.

Conditioning on the Erdos girth conjecture.

Improving the approximation factor by $-\Theta(1/\log n)$ only increases the round complexity by a constant factor.

Summary

New upper bound. $(k+1)$-approximation: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **real** $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$.

They **match** at $k = 2, 3, 4, 5, \ldots$

New lower bound. $(k+1-\epsilon)$-approximation: $\tilde{\Omega}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **integer** $k \geq 1$.

Conditioning on the Erdos girth conjecture.

Take-away message: MPX is the right tool to tackle the MWC problem in the CONGEST model.

In the sense that it yields tight bounds.

Improving the approximation factor by $-\Theta(1/\log n)$ only increases the round complexity by a constant factor.

Summary

New upper bound. $(k+1)$-approximation: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **real** $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$.

They **match** at $k = 2, 3, 4, 5, \ldots$

New lower bound. $(k+1-\epsilon)$-approximation: $\tilde{\Omega}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **integer** $k \geq 1$.

Conditioning on the Erdos girth conjecture.

Take-away message: MPX is the right tool to tackle the MWC problem in the CONGEST model.

and beyond

Some implications: (via known results for $(1+\epsilon)$-approximate undirected weighted SSSP)

- $\tilde{O}\left(mn^{1/k}\right)$-work $\tilde{O}(1)$-depth parallel algorithm.
- $\tilde{O}\left(n^{1/k}\right)$-round broadcast congested clique. ← **Optimal!**

Possible extension to directed graphs?

- There are four variants of the MWC problem:

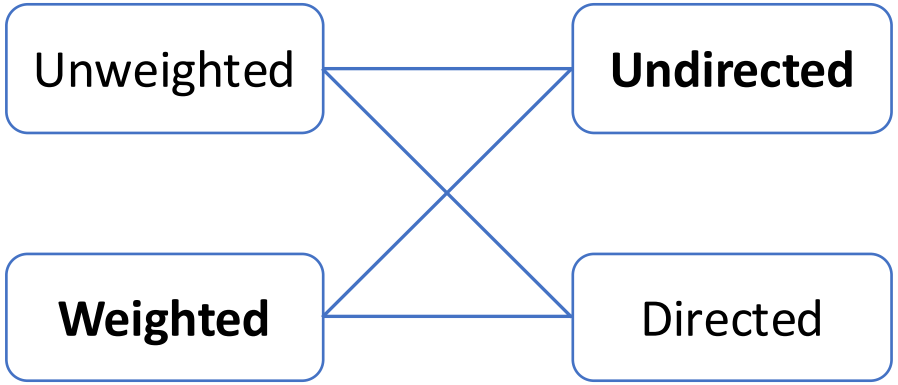

What about directed graphs?

- Our focus: Weighted + Undirected.

$$\tilde{\Theta}\left(n^{(k+1)/(2k+1)} + D\right)$$

Possible extension to directed graphs?

- There are four variants of the MWC problem:

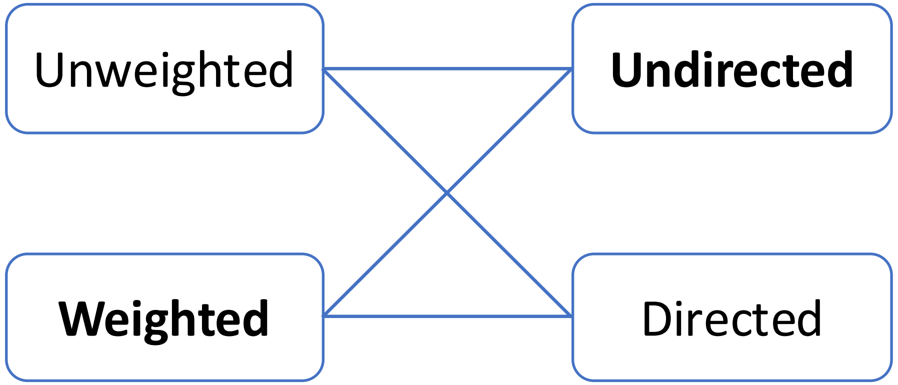

Conjecture: Our approach also extends to directed graphs, and $\tilde{\Theta}(n^{(k+1)/(2k+1)} + D)$ is also the tight bound.

What about directed graphs?

- Our focus: Weighted + Undirected.

$\tilde{\Theta}(n^{(k+1)/(2k+1)} + D)$

Possible extension to directed graphs?

Our **lower bound** extends to directed graphs immediately.

▽

Conjecture: Our approach also extends to directed graphs, and $\widetilde{\Theta}\left(n^{(k+1)/(2k+1)} + D\right)$ is also the tight bound.

△

Upper bound: Replace MPX with some directed version of the clustering?

Improving the approximation factor by $-\Theta(1/\log n)$ only increases the round complexity by a constant factor.

Independent work

New upper bound. $(k+1)$-approximation: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **real** $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$.

They **match** at $k = 2, 3, 4, 5, \ldots$

New lower bound. $(k+1-\epsilon)$-approximation: $\tilde{\Omega}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **integer** $k \geq 1$.

Conditioning on the Erdos girth conjecture.

Independent work: "Girth Approximations in the CONGEST Model" by Shiri Chechik, Gur Lifshitz, Doron Mukhtar [PODC 2026]

- Same lower bound.
- Upper bound: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$
    - Higher approximation ratio: $2k - 1 + o(1)$
    - Only for integers k.

Improving the approximation factor by $-\Theta(1/\log n)$ only increases the round complexity by a constant factor.

Independent work

New upper bound. $(k+1)$-approximation: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **real** $k \geq \frac{1+\sqrt{5}}{2} \approx 1.62$.

They **match** at $k = 2, 3, 4, 5, \ldots$

New lower bound. $(k+1-\epsilon)$-approximation: $\tilde{\Omega}\left(n^{(k+1)/(2k+1)} + D\right)$, for any **integer** $k \geq 1$.

Conditioning on the Erdos girth conjecture.

Independent work: “Girth Approximations in the CONGEST Model” by Shiri Chechik, Gur Lifshitz, Doron Mukhtar [PODC 2026]

- Unweighted undirected: $\tilde{O}\left(n^{1/f} + D\right)$ rounds.
  - f-approximation
- Directed: $\tilde{O}\left(n^{2/3} + D\right)$ rounds.
  - 2-approximation (unweighted)
  - $(2+\epsilon)$-approximation (weighted)

- Same lower bound.
- Upper bound: $\tilde{O}\left(n^{(k+1)/(2k+1)} + D\right)$
  - Higher approximation ratio: $2k - 1 + o(1)$
  - Only for integers k.

Open problems

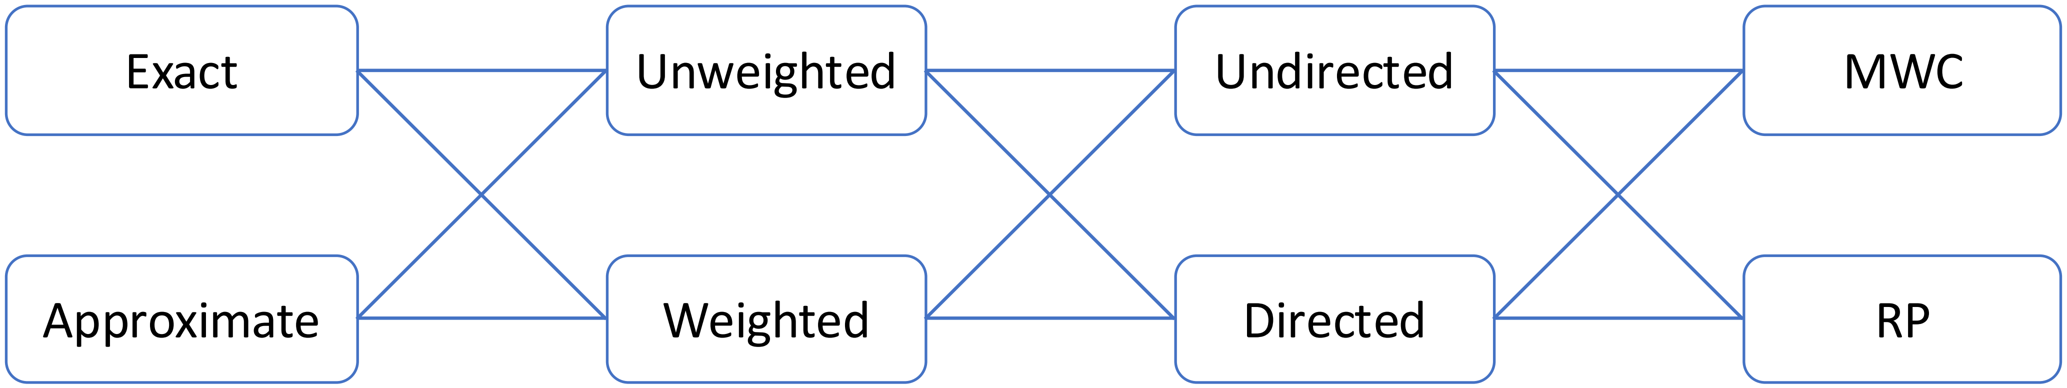

Determine their optimal round complexities.

- Many of these problems do not fall into the common complexity classes: $\tilde{\Theta}(D)$, $\tilde{\Theta}(\sqrt{n} + D)$, $\tilde{\Theta}(n)$.

Can we get universally-optimal algorithms?

Open problems

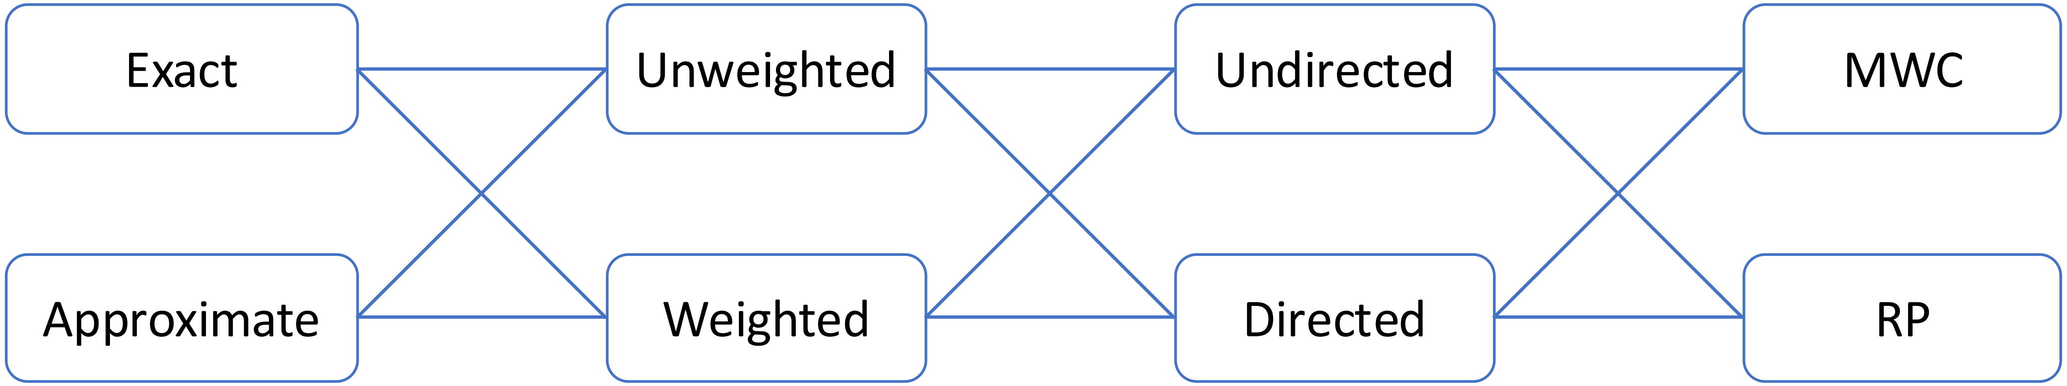

Determine their optimal round complexities.

- Many of these problems do not fall into the common complexity classes: $\tilde{\Theta}(D)$, $\tilde{\Theta}(\sqrt{n} + D)$, $\tilde{\Theta}(n)$.

Can we get universally-optimal algorithms?

Thank you!