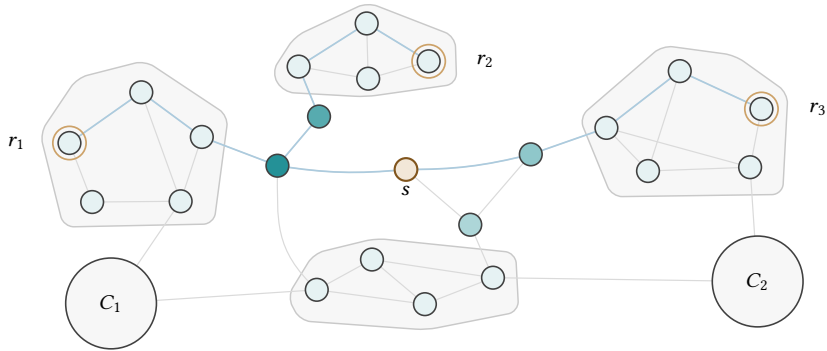
Parallel Small Vertex Connectivity

in Near-Linear Work and Polylogarithmic Depth



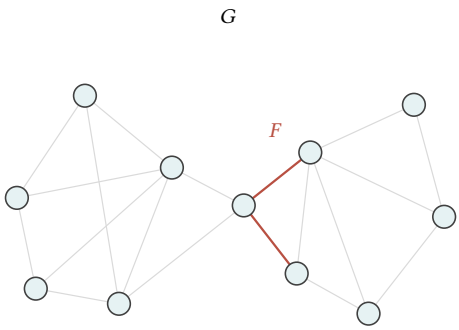
Yonggang Jiang
MPI-INF and Saarland University

Changki Yun
University of Maryland

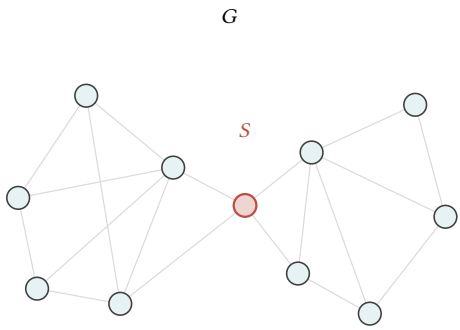
FOCS 2026

Edge and vertex connectivity

Edge cut F



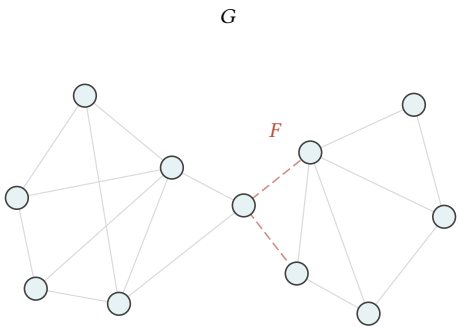
Vertex cut S

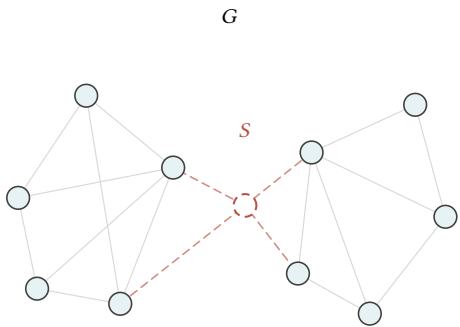


Edge and vertex connectivity

$$\text{Edge connectivity} := \min\{|F| : G - F \text{ is disconnected}\},$$

$$\text{Vertex connectivity} := \min\{|S| : G - S \text{ is disconnected}\}.$$





Connectivity problems

Edge connectivity

$\widetilde{O}(m)$ sequential time [Kar00]

Vertex connectivity

$m^{1+o(1)}$ sequential time [LNP+25, CKL+25]

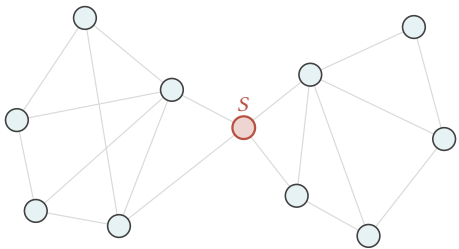
Edge connectivity also admits a parallel algorithm [GG21, LMN21, AB23].

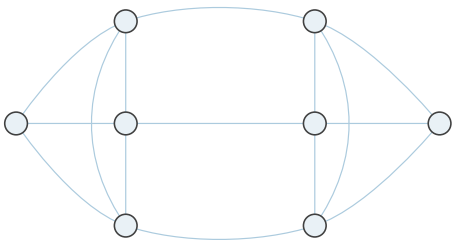
$\widetilde{O}(m)$ work, $\widetilde{O}(1)$ depth

Can vertex connectivity achieve the same parallel bounds?

Deciding k-vertex connectivity

Find a cut $|S| < k$ or find a witness of k-connectivity.





Parallel *small* vertex connectivity

Range of k	Work	Depth	Consequence / reference
Arbitrary k	$\widetilde{O}(n^\omega)$	$\widetilde{O}(1)$	[BJMY25]
$k = \text{polylog}(n)$ (dense graphs)	$\widetilde{O}(n^2k^4)$	$\widetilde{O}(k^2)$	[CKT93]
$k < 4$	$\widetilde{O}(m)$	$\widetilde{O}(1)$	[SV82, TV85, MR92]
Arbitrary k	$\widetilde{O}(m)$	$\widetilde{O}(1)$	Implies dense directed reachability of $\widetilde{O}(n^2)$ work, $\widetilde{O}(1)$ depth [BJMY25]
$k = n^{\Theta(1)}$	$\widetilde{O}(m)$	$\widetilde{O}(1)$	Implies "direct sum" of dense reachability of $\widetilde{O}(nk)$ work, $\widetilde{O}(1)$ depth [This work]
$k = \text{polylog}(n)$	$\widetilde{O}(m)$	$\widetilde{O}(1)$	Randomized PRAM, w.h.p. [This work]

Design Principle: Avoiding Directed Reachability

Sparsification [NI92, CKT93] allows us to assume $m \leq nk$.

- Sparse graphs are already hard.

Directed reachability enters through residual searches [FNY+20] and exact flow [LNP+25] in essential subroutines.

- Conditioned on $\widetilde{O}(m)$ work, the best known depth is $n^{1/2+o(1)}$ [JLS19].

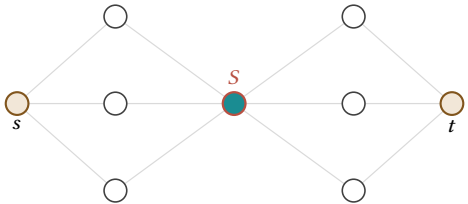
Warm-up: (s, t) k-connectivity

Vertex weights $w \geq 0$, $\quad w(s) = w(t) = 0$, $\quad \sum_v w(v) = 1$.

A cut $|S| < k$

Take $w \propto \mathbf{1}_S$.

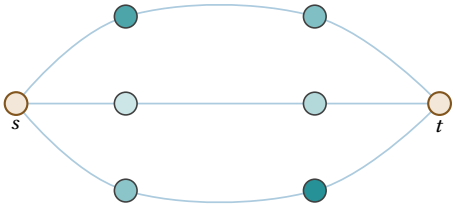
$\exists w : \quad \mathrm{dist}_w(s, t) \geq 1/(k-1)$



No cut $|S| < k$

k internally disjoint paths (Menger)

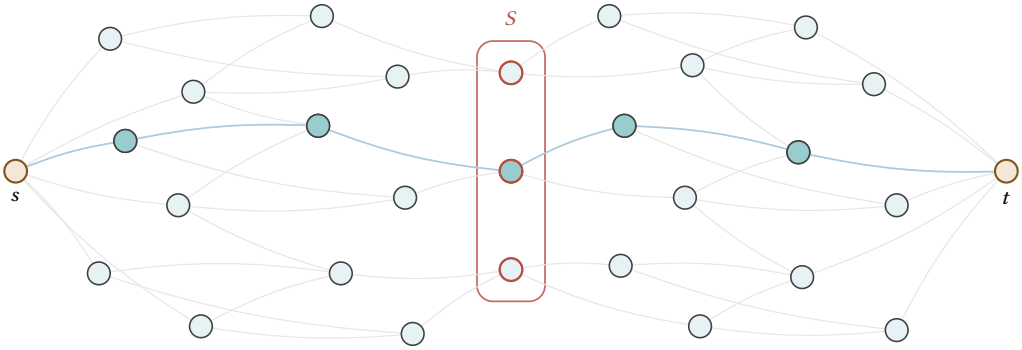
$\forall w : \quad \mathrm{dist}_w(s, t) \leq 1/k$



Finding the weights with MWU [GK07, BGS21]

Starting from uniform weights, increase weights along a shortest path P.

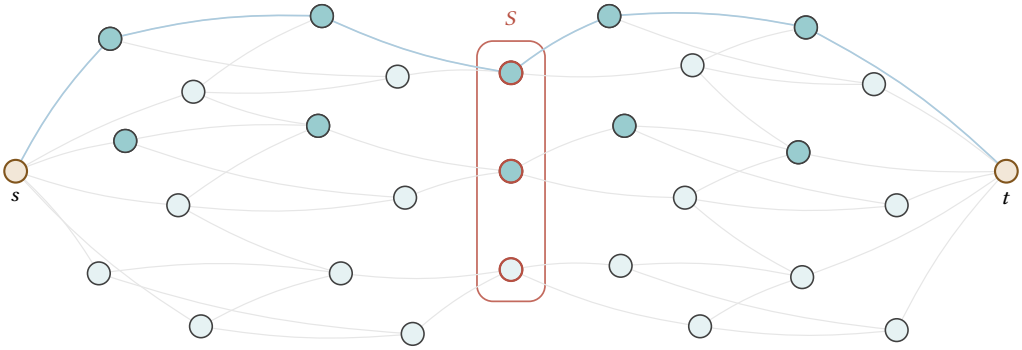
$w(v) \leftarrow (1 + 1/\text{poly}(k))w(v)$ for $v \in P \setminus \{s, t\}$.



Finding the weights with MWU [GK07, BGS21]

Starting from uniform weights, increase weights along a shortest path P.

$w(v) \leftarrow (1 + 1/\text{poly}(k))w(v)$ for $v \in P \setminus \{s, t\}$.

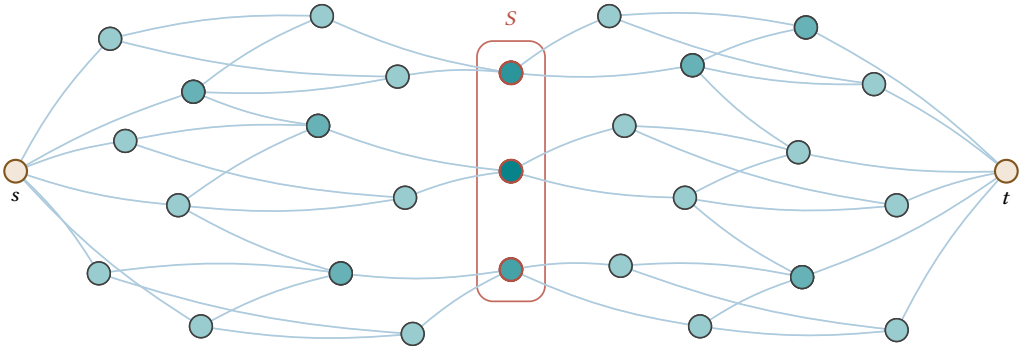


Finding the weights with MWU [GK07, BGS21]

Starting from uniform weights, increase weights along a shortest path P.

$w(v) \leftarrow (1 + 1/\text{poly}(k))w(v)$ for $v \in P \setminus \{s, t\}$.

If $|S| < k$, MWU reaches $\text{dist}_w(s, t) > 1/k$ in $T = \widetilde{O}(\text{poly}(k))$ rounds.



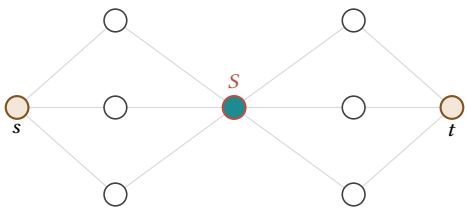
Caveat: can we find a shortest path?

Exact directed SSSP $\longrightarrow$ exact undirected SSSP [HJS26]

Use $(1 + 1/\Theta(k))$-approximate SSSP [RGH+22]. $\widetilde{O}(mk^2)$ work, $\widetilde{O}(k^2)$ depth

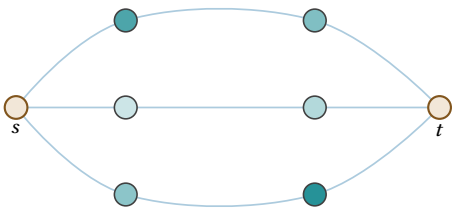
A cut $|S| < k$

$\exists w : \quad \mathrm{dist}'_w(s,t) \geq 1/(k - \mathbf{0.6})$



No cut $|S| < k$

$\forall w : \quad \mathrm{dist}'_w(s,t) < 1/(k - \mathbf{0.4})$

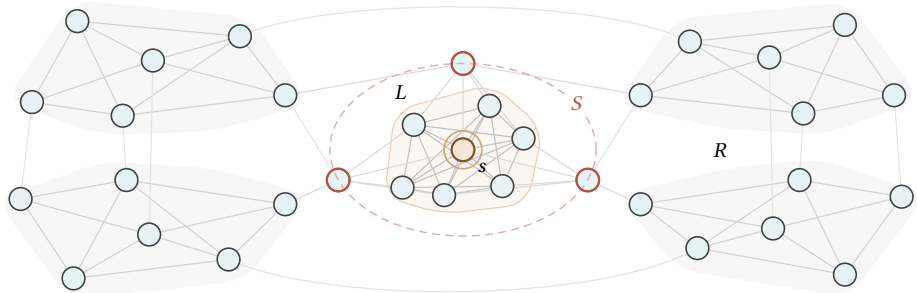


Global connectivity via local cuts

Suppose $|S| < k$ and $\mathrm{vol}_G(L) := \sum_{v \in L} \deg_G(v) \in [\mu, 2\mu)$. [FNY+20]

Try $\mu = 1, 2, 4, \ldots$ and sample $\widetilde{O}(m/\mu)$ vertices by degree to hit L.

From $s \in L$, it suffices to certify some cut of size $< k$.



Parallel local cuts

Theorem (this work, informal)

After shared preprocessing, randomized LocalCut(s, μ, k) uses

$$\widetilde{O}(\mu \operatorname{poly}(k)) \text{ work} \qquad \widetilde{O}(\operatorname{poly}(k)) \text{ depth.}$$

If a cut (L, S, R) satisfies

$$s \in L, \qquad \operatorname{vol}_G(L) \in [\mu, 2\mu), \qquad |S| < k,$$

it returns a fractional (s, t)-cut of size $< k$ for some t, with constant probability.

If G is k-connected, it always returns $\perp$.

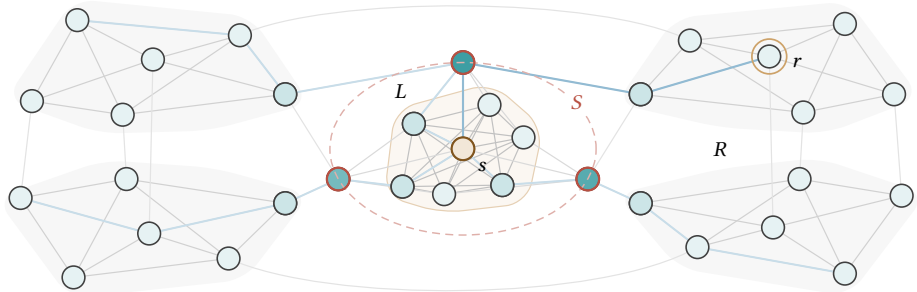
A separate rounding algorithm converts the fractional cut to an integral cut S. (Omitted)

MWU with changing endpoints

Let $r \in R$ change between rounds.

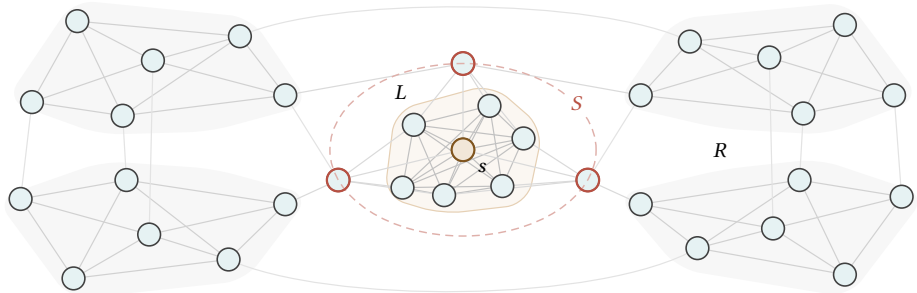
Each approximate shortest s–r path still crosses S.

MWU finds separating weights for one of these endpoints.



Caveat: how do we choose $r \in R$?

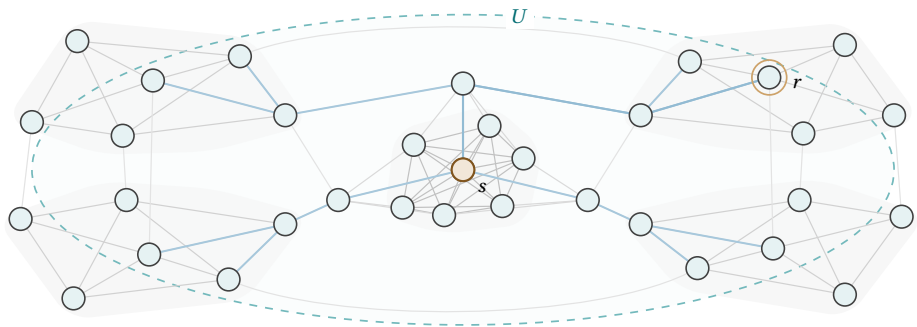
We know s, but the target cut (L, S, R) is unknown.



Random termination

If no cut is found, one may grow a single-source shortest path subtree U of $\mathrm{vol}_G(U) = \Theta(kT\mu)$, while all vertices of U is distant from s by at most $1/(k - 0.4)$.

Sample $r \in U$ by degree, and the (s, r)-path in U is MWU updated.

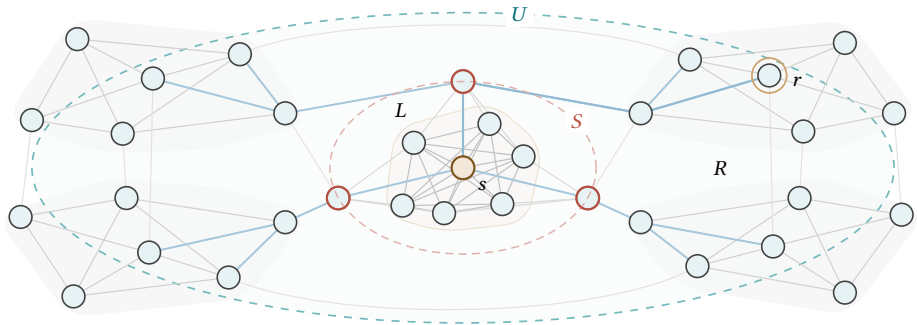


Random termination

For simplicity, assume $\Delta(G) = O(\mu)$.

$$\Pr[r \notin R \mid \text{history}] \leq \frac{\text{vol}_G(L \cup S)}{\text{vol}_G(U)} \leq \frac{O(k\mu)}{\Theta(kT\mu)} = O(1/T)$$

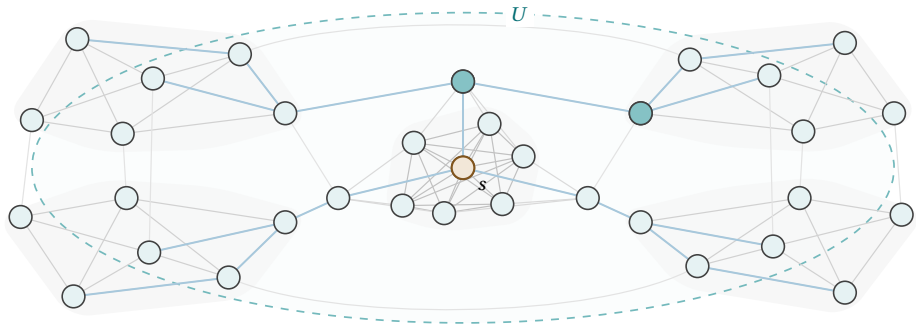
With $(1 - O(1/T))^T = \Omega(1)$ probability, each of T rounds finds $r \in R$.



Caveat: can each search stay local?

$\widetilde{O}(m/\mu)$ sources $\implies$ only $\widetilde{O}(\mu\,\mathrm{poly}(k))$ work per source.

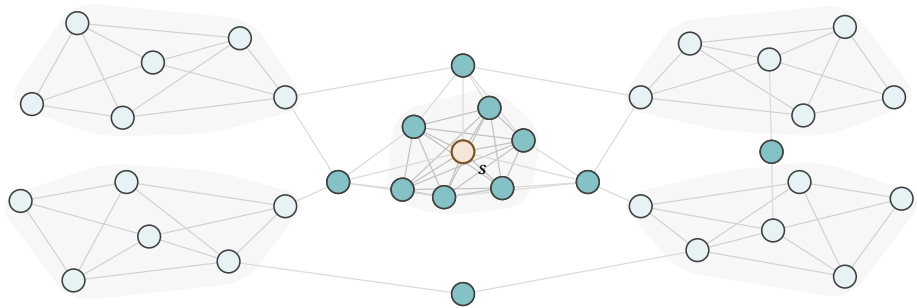
Even approximate SSSP on G already costs $\widetilde{O}(m\,\mathrm{poly}(k))$ work per round.



Only a small volume changes weight

Over all T rounds, updated vertices have total degree $\widetilde{O}(\mu \operatorname{poly}(k))$.

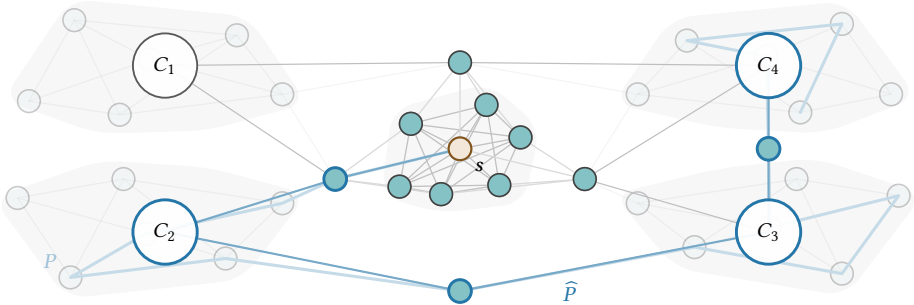
Keep these vertices explicit and contract the weight-1 components.



Contracting weight-1 components

$$\widehat{w}(C_i) = 1 \qquad \Longrightarrow \qquad 0 \;\leq\; w(P) - \widehat{w}(\widehat{P}) \;\leq\; n.$$

An additive loss of n may distort short paths and affect MWU.



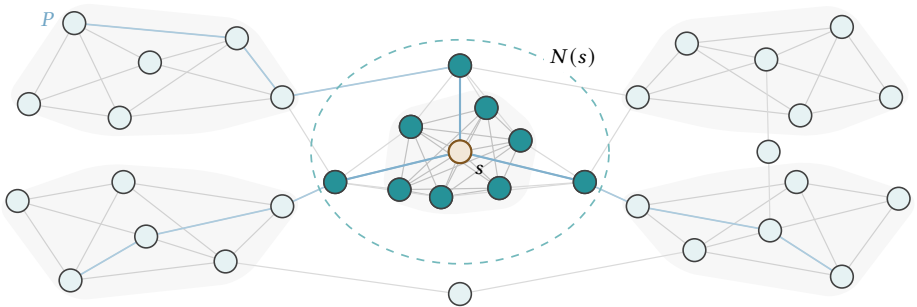
Making the contraction loss negligible

$w_{\text{init}}(v) \leftarrow n^3$ for $v \in N(s)$.

Only a constant blowup in # of MWU rounds.

Every path P leaving s satisfies $w(P) \geq n^3$.

$(1 - n^{-2})w(P) \leq \widehat{w}(\widehat{P}) \leq w(P)$.



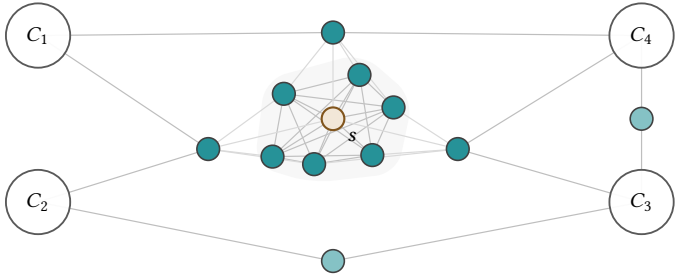
Single-source length sparsifier

Initially, $\mathrm{vol}_G(\{s\} \cup N(s)) \le \mathrm{vol}_G(L \cup S) = O(k\mu)$.

After updates, $|E(\widehat{G})| \le \sum_{v \text{ explicit}} \deg_G(v) = \widetilde{O}(\mu \,\mathrm{poly}(k))$.

- SSSP and other operations fit within this bound.

$\widehat{G}$

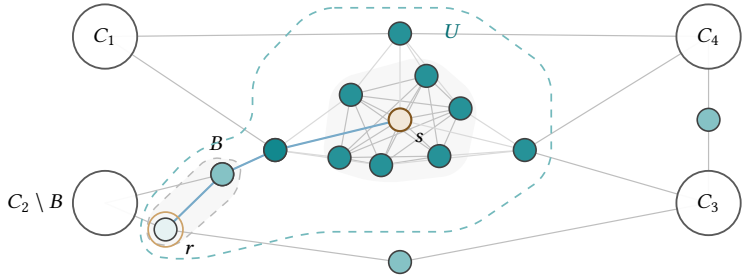


Maintaining weight-1 components

Need to unfold the compression to perform random termination.

The *sensitivity spanning forest* data structure supports all operations under batch removal of new explicit vertices.

- Based on *(non-amortized) connectivity* [KKM13] and *batch-parallel dynamic trees* [TDB19].

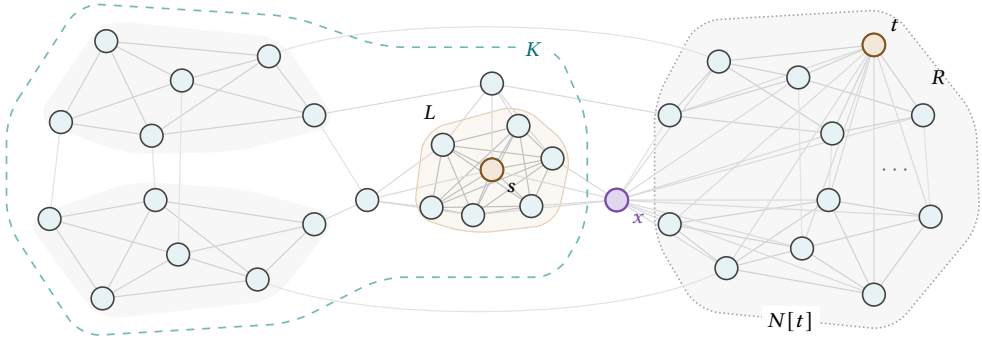


Removing the degree bound ($\Delta(G) \gg \mu$)

Sample t by degree. $\Pr[t \in R] = \Omega(1)$.

If $t \in R$, low-degree vertices form a connected component K of $G - N[t]$ containing L.

Extend paths through a boundary exit x to t for weight updates.



Our contributions

First algorithm for all $k \geq 3$ with

$$\widetilde{O}(m \operatorname{poly}(k)) \text{ work} \qquad \widetilde{O}(\operatorname{poly}(k)) \text{ depth}$$

For $k = n^{o(1)}$, this is $\widehat{O}(m) = mn^{o(1)}$ work and $\widehat{O}(1) = n^{o(1)}$ depth.

Our barrier through dense directed reachability (DDR)

For polynomial $k = n^{\alpha}$, $\widehat{O}(m)$ work and $\widehat{O}(1)$ depth would let us compute

$$\bigwedge_{i=1}^{n^{1-\alpha}} \mathrm{DDR}(G_i), \qquad |V(G_i)| = n^{\alpha}$$

$$\text{in } \widehat{O}\left(n^{1-\alpha} \cdot n^{2\alpha}\right) \text{ work and } n^{o(1)} \text{ depth.}$$

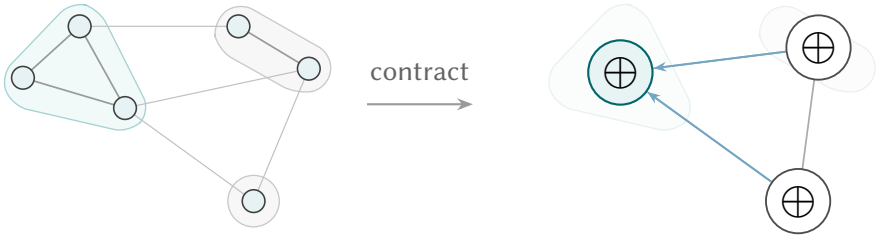
Just as surprising as the corresponding breakthrough for a single DDR instance.

One algorithm for PRAM and CONGEST?

The CONGEST result and parallel question of [JM23] motivated this work.

Our PRAM algorithm and their CONGEST algorithm are different.

Minor-Aggregation



Efficient simulations in PRAM and CONGEST [RGH+22].

k-vertex connectivity in $\text{poly}(k, \log n)$ MA rounds?