Lower Bounds in Massively Parallel Computation

Peter Davies

AMG workshop, 28/10/22

Durham University

Main Aim of this Talk

Low-Space MPC Model

(Massively Parallel Computation)

s space

s space

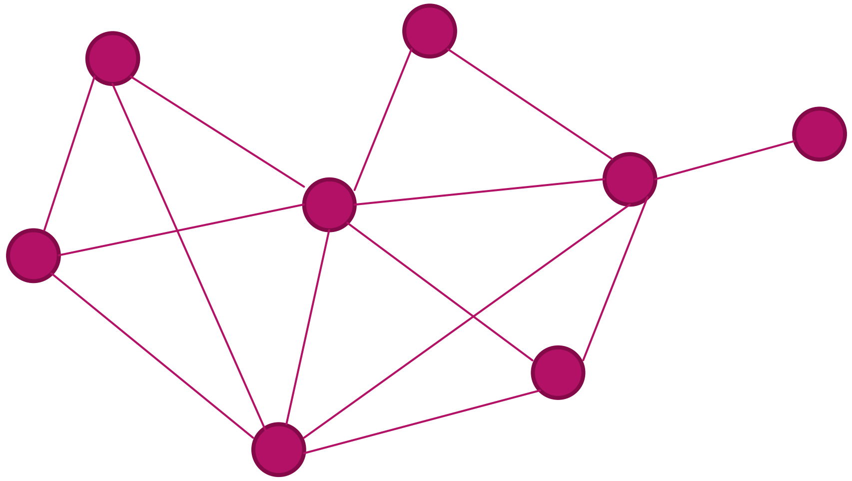

s space

s space

s space

s space

n nodes, **Δ** max degree

s space

s space

Low space regime: $s = n^{\varepsilon}$, $0 < \varepsilon < 1$

Low-Space MPC Model

(Massively Parallel Computation)

s space

s space

Number of machines:

Lower bounds are against any poly(**n**) number of machines / total space

s space

s space

s space

s space

n nodes, **Δ** max degree

s space

s space

Low space regime: $s = n^{\varepsilon}$, $0 < \varepsilon < 1$

MPC Communication

Graph arbitrarily partitioned among machines

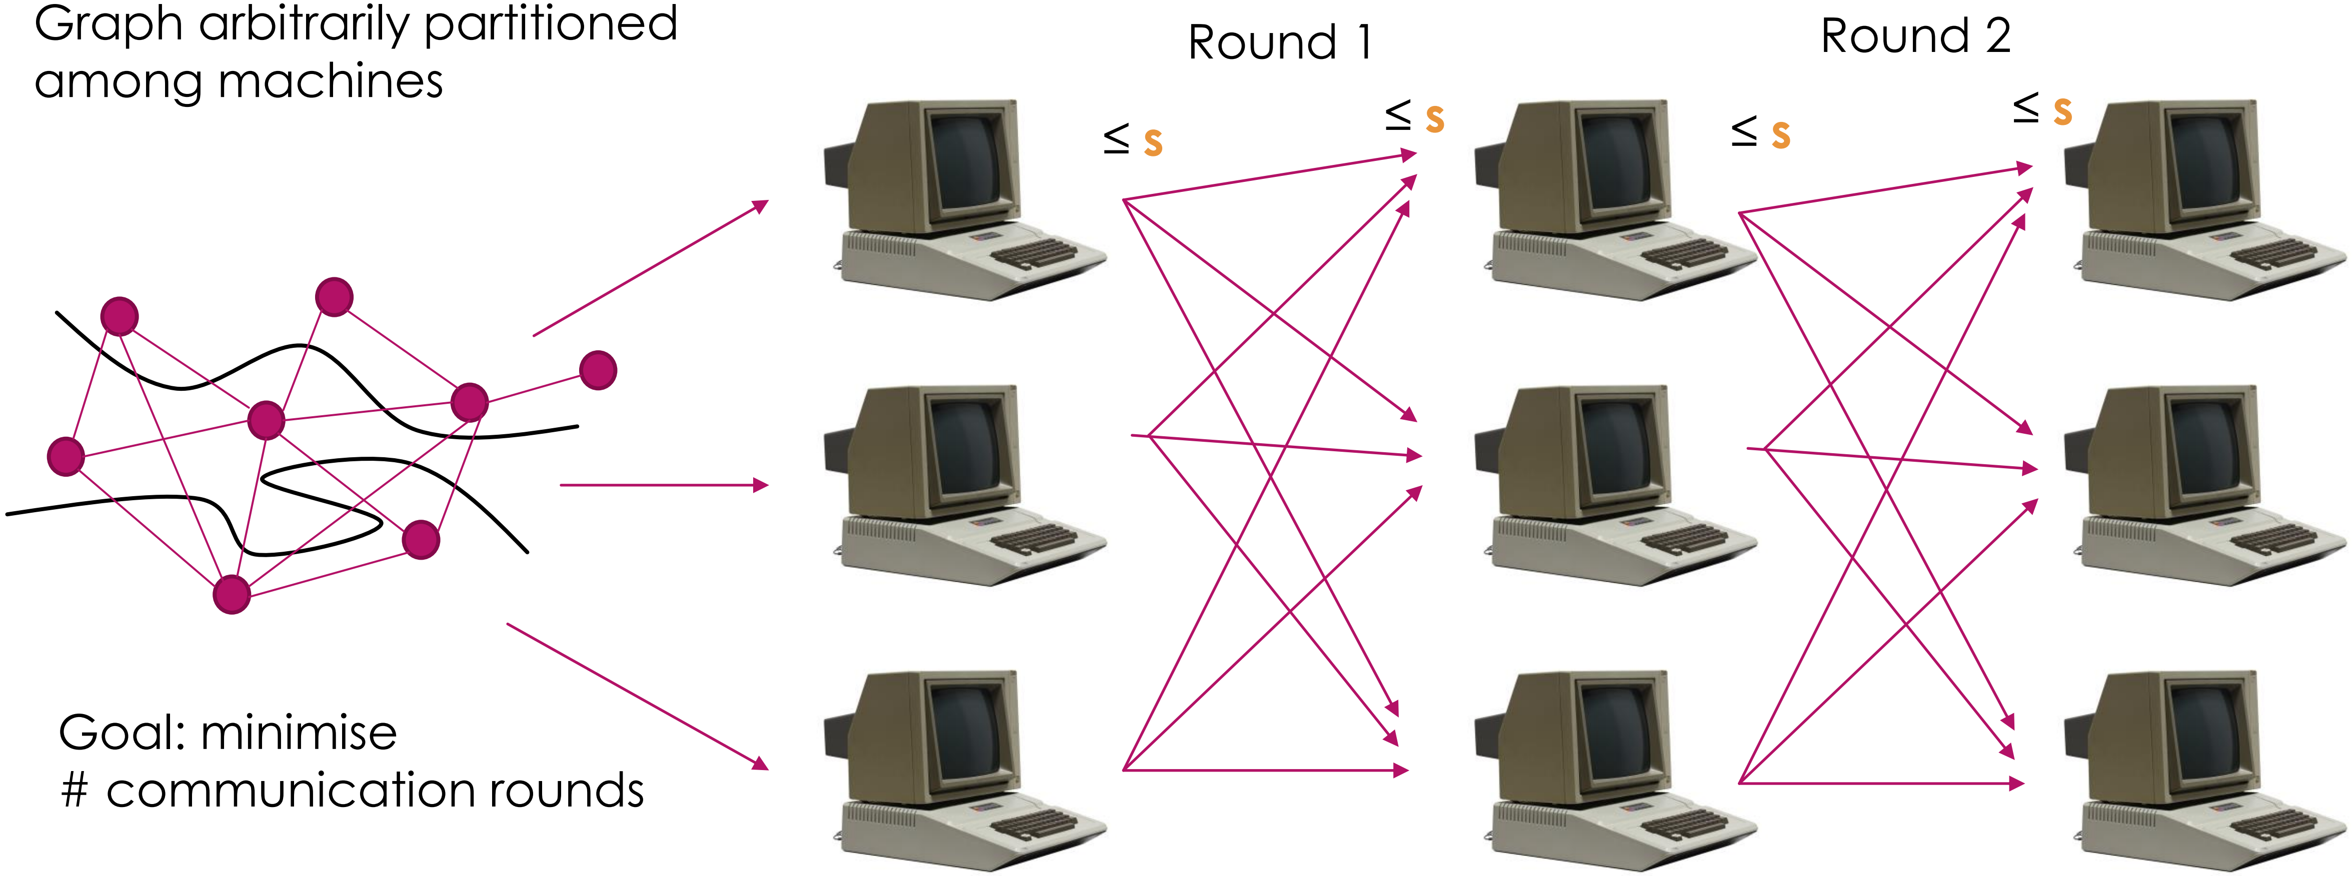

Goal: minimise
communication rounds

Known Lower Bounds

Unconditional Lower Bounds

If some problem in P cannot be solved in $O(\log_s n) = O(1)$ MPC rounds, then $NC^1 \neq P$.

Roughgarden, Vassilivistki, Wang, JACM 2018

1 vs 2 Cycle Conjecture

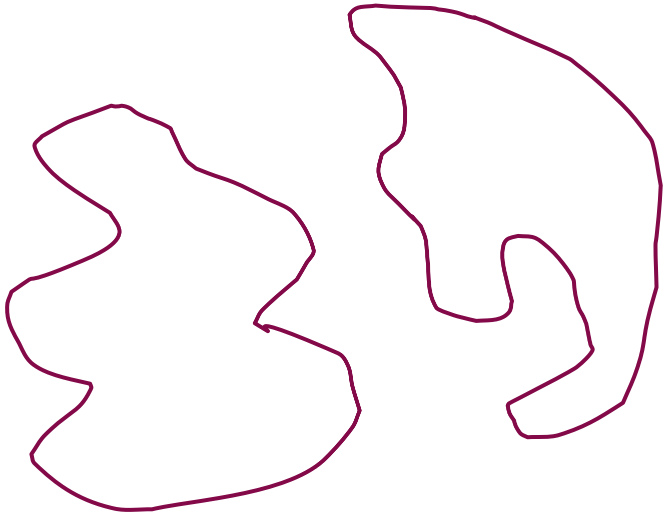

vs

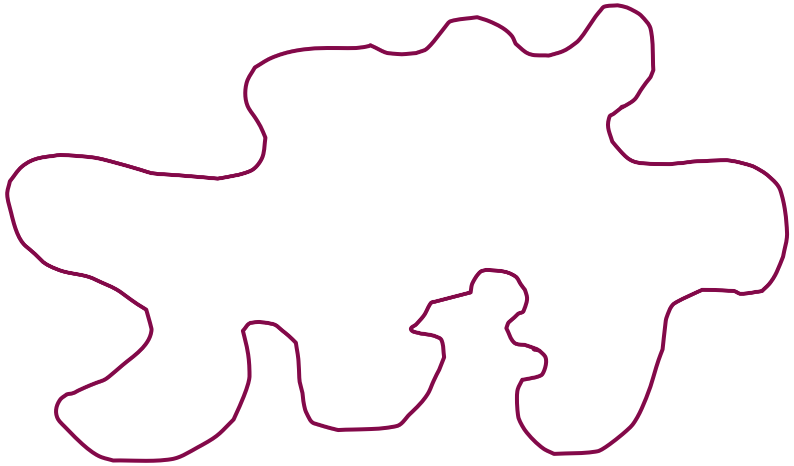

Two **n**/2-node cycles

One **n**-node cycle

1 vs 2 cycle conjecture: $\Omega(\log \mathbf{n})$ rounds required when $\mathbf{s} \leq \mathbf{n}^{0.99}$

even with any $\mathbf{n}^{O(1)}$ machines

Direct Conditional Lower Bounds

Some connectivity-based problems can be reduced to 1 vs 2 cycles

Example:

Finding connected components takes $\Omega(\log D)$ rounds*.

Behnezhad, Dhulipala, Esfandiari, Łącki & Mirrokni FOCS '19, Coy & Czumaj STOC '22

*conditional on the 1 vs 2 cycle conjecture

Conditional Lower Bounds From LOCAL

LOCAL Model

- Nodes of input graph are the processors
- Synchronous rounds of communication
- Can send arbitrary messages to neighbours each round
- Goal: minimise #communication rounds

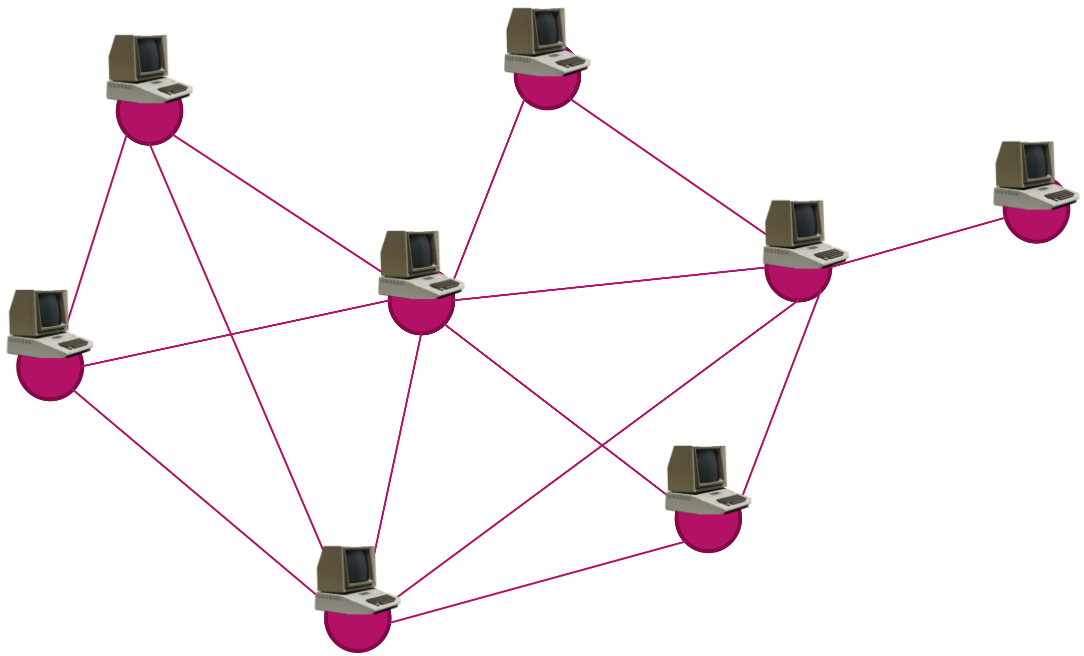

Many (unconditional) lower bounds are known

n nodes, **m** edges, **Δ** max degree

Conditional Lower Bounds from LOCAL

T-round LOCAL lower bound

↓

Ghaffari, Kuhn, Uitto (GKU)
Czumaj, Davies, Parter (CDP)

Ω(log T)-round MPC lower bound*

*conditional on the 1 vs 2 cycle conjecture, holds against component-stable algorithms

Component-Stability

Informally: algorithm treats connected components separately

GKU: a node's output depends only on its connected component (topology and IDs) and randomness…

…and the value of **n** and **Δ** (CDP)

Are MPC Algorithms Component-Stable?

CDP:

Many deterministic component-unstable algorithms surpass deterministic lower bounds (**MIS, maximal matching, Lovász Local Lemma, …**)

One deterministic component-unstable algorithm surpasses randomized lower bound (**$\Omega(n/\Delta)$-size independent set**)

All of these exploit the component-unstable **method of conditional expectations**

No randomized lower bound for locally-checkable problem has been broken

Checklist For Using Conditional MPC Lower Bounds

Lower Bound Statement

Theorem (CDP ‘21)

Let P be an O(1)-replicable graph problem that has a T(N, **Δ**)-round lower bound in the randomized LOCAL model with shared randomness, for constrained function T, on graphs with input estimate N and maximum degree **Δ**, from some normal family G. Then there is no o(log T(**n**, **Δ**))- round component-stable low-space MPC algorithm solving P w.h.p. on legal **n**-node, **Δ**-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

Lower Bound Statement

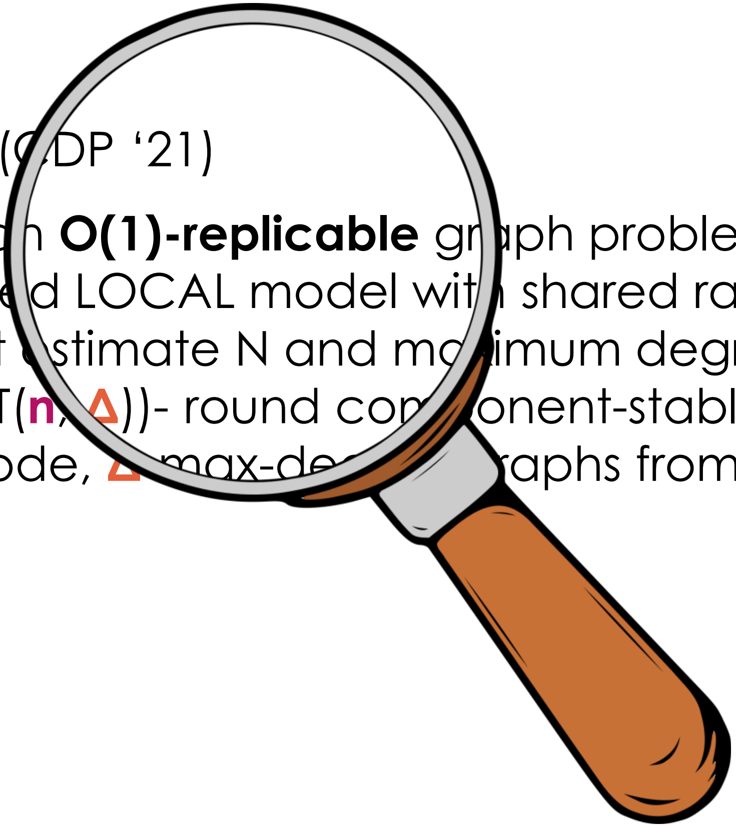

Theorem (CDP '21)

Let P be an **O(1)-replicable** graph problem that has a T(N, Δ)-round lower bound in the randomized LOCAL model with shared randomness, for constrained function T, on graphs with input estimate N and maximum degree Δ, from some normal family G. Then there is no o(log T(**n**, Δ))- round component-stable low-space MPC algorithm solving P w.h.p. on legal **n**-node, Δ max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

Replicability

O(1)-replicable graph problem:

An invalid solution on a graph G must also be invalid on multiple disjoint copies of G

plus a few isolated nodes

CDP'21
- Any locally-checkable problem is O(1)-replicable graph problem
- Some natural approximation problems are too

Checklist

- Problem is O(1)-replicable

Lower Bound Statement

Theorem (CDP '21)

Let P be an O(1)-replicable graph problem that has a T(N, **Δ**)-round lower bound in the randomized LOCAL model with shared randomness, for constrained function T, on graphs with input estimate N and maximum degree **Δ**, from some normal family G. Then there is no o(log T(**n**, **Δ**))- round component-stable low-space MPC algorithm solving P w.h.p. on legal **n**-node, **Δ**-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

Lower Bound Statement

Theorem (CDP '21)

Let P be an O(1)-replicable graph problem that has a T(N, Δ)-round lower bound in the randomized LOCAL model with **shared randomness**, for constrained function T, on graphs with input estimate N and maximum degree Δ, from some normal family G. Then there is no o(log T(**n**, Δ))- round component-stable low-space MPC algorithm solving P w.h.p. on legal **n**-node, Δ-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

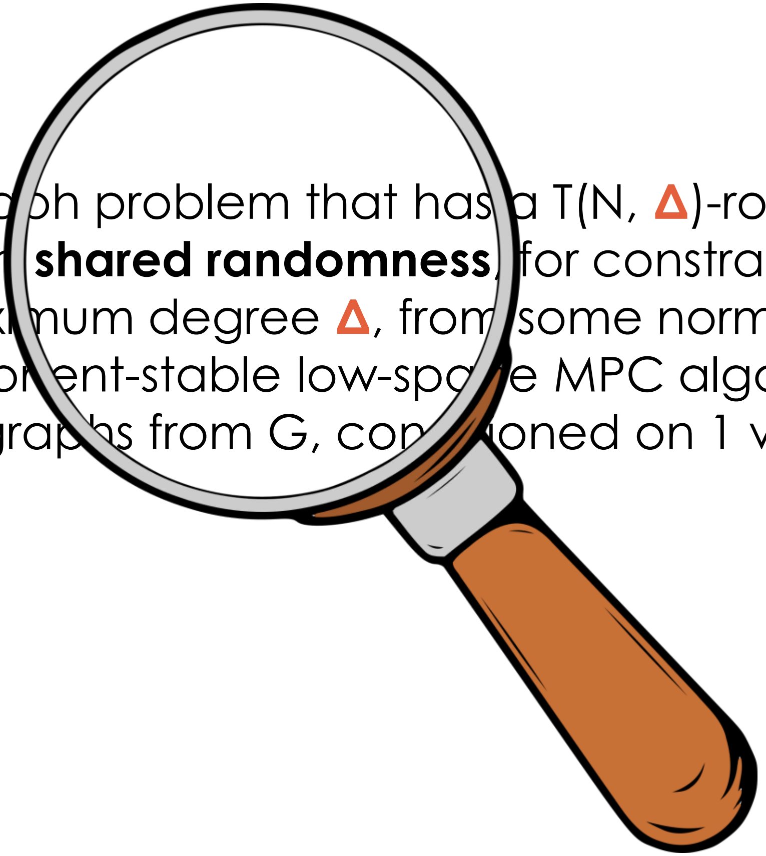

Shared Randomness

LOCAL lower bounds are generally against only local randomness

Ghaffari, Kuhn, and Uitto re-proved some against global randomness

Any other randomized bounds also need to be re-proven against global randomness to transport to MPC

So far this has not caused major problems

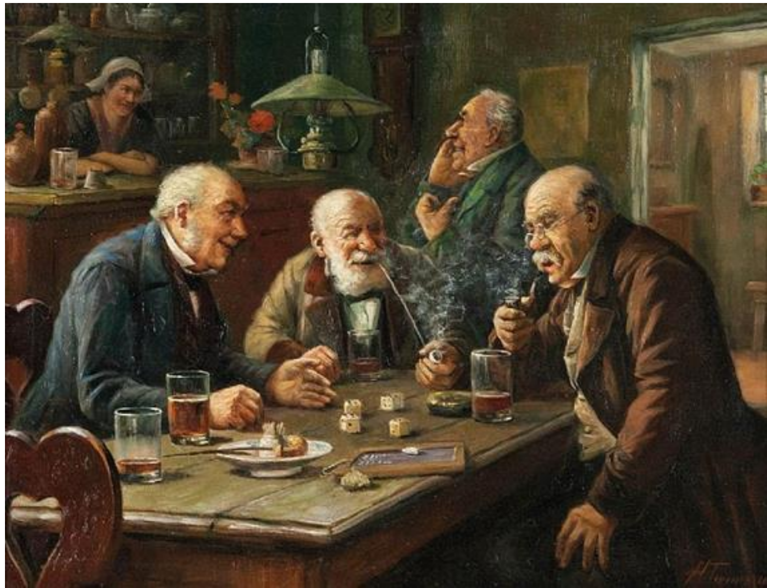

Checklist

- Problem is O(1)-replicable
- LOCAL lower bound holds under shared randomness

Lower Bound Statement

Theorem (CDP '21)

Let P be an O(1)-replicable graph problem that has a T(N, **Δ**)-round lower bound in the randomized LOCAL model with shared randomness, for constrained function T, on graphs with input estimate N and maximum degree **Δ**, from some normal family G. Then there is no o(log T(**n**, **Δ**))- round component-stable low-space MPC algorithm solving P w.h.p. on legal **n**-node, **Δ**-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

Lower Bound Statement

Theorem (CDP '21)

Let P be an O(1)-replicable graph problem that has a $T(N, \Delta)$-round lower bound in the randomized LOCAL model with shared randomness, for **constrained function** T, on graphs with input estimate N and maximum degree Δ, from some normal family G. Then there is no $o(\log T(n, \Delta))$- round component-stable low-space MPC algorithm solving P w.h.p. on legal n-node, Δ-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

Constrained Functions

$T(N, \Delta)$ is constrained if it is $\log^{O(1)} N$ for any Δ, 

Framework gives MPC lower bounds of *at most* $\Omega(\log \log n)$

Checklist

- Problem is O(1)-replicable
- LOCAL lower bound holds under shared randomness
- Target MPC bound is at most $\Omega(\log \log n)$

Lower Bound Statement

Theorem (CDP '21)

Let P be an O(1)-replicable graph problem that has a T(N, **Δ**)-round lower bound in the randomized LOCAL model with shared randomness, for constrained function T, on graphs with input estimate N and maximum degree **Δ**, from some normal family G. Then there is no o(log T(**n**, **Δ**))- round component-stable low-space MPC algorithm solving P w.h.p. on legal **n**-node, **Δ**-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

Lower Bound Statement

Theorem (CDP '21)

Let P be an O(1)-replicable graph problem that has a $T(N, \Delta)$-round lower bound in the randomized LOCAL model with shared randomness, for constrained function T, on graphs with input estimate N and maximum degree Δ, from some **normal family G**. Then there is no $o(\log T(n, \Delta))$- round component-stable low-space MPC algorithm solving P w.h.p. on legal n-node, Δ-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

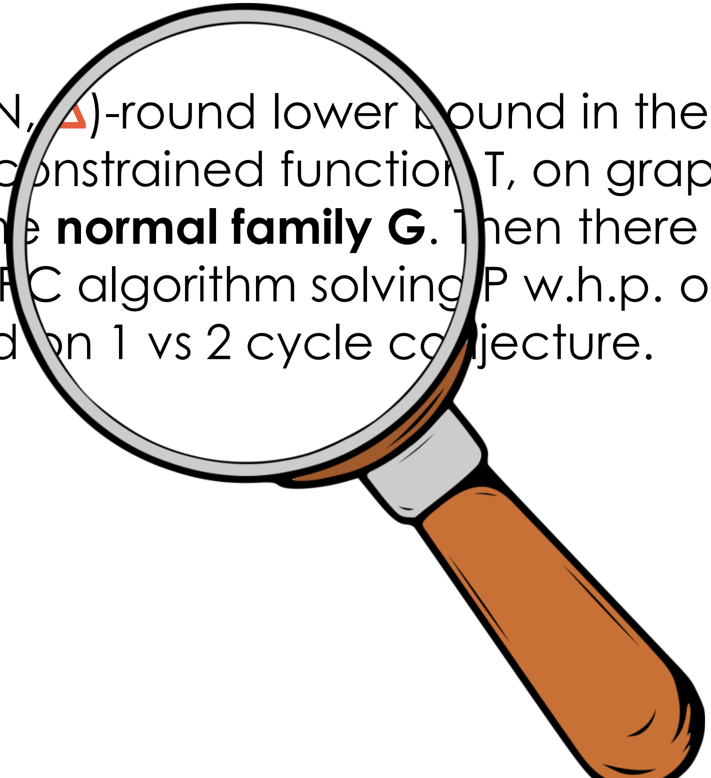

Normal Graph Families

Graph family G is normal if it is

- hereditary (closed under induced subgraph), and
- 

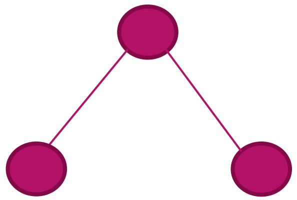

Example: family of trees is not normal, family of forests is.

Checklist

- Problem is O(1)-replicable
- LOCAL lower bound holds under shared randomness
- Target MPC bound is at most $\Omega(\log \log n)$
- Graph class is normal (hereditary)

Lower Bound Statement

Theorem (CDP '21)

Let P be an O(1)-replicable graph problem that has a T(N, **Δ**)-round lower bound in the randomized LOCAL model with shared randomness, for constrained function T, on graphs with input estimate N and maximum degree **Δ**, from some normal family G. Then there is no o(log T(**n**, **Δ**))- round component-stable low-space MPC algorithm solving P w.h.p. on legal **n**-node, **Δ**-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

Lower Bound Statement

Theorem (CDP '21)

Let P be an O(1)-replicable graph problem that has a T(N, Δ)-round lower bound in the randomized LOCAL model with shared randomness, for constrained function T, on graphs with **input estimate N** and maximum degree Δ, from some normal family G. Then there is no o(log T(**n**, Δ))- round component-stable low-space MPC algorithm solving P w.h.p. on legal **n**-node, Δ-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

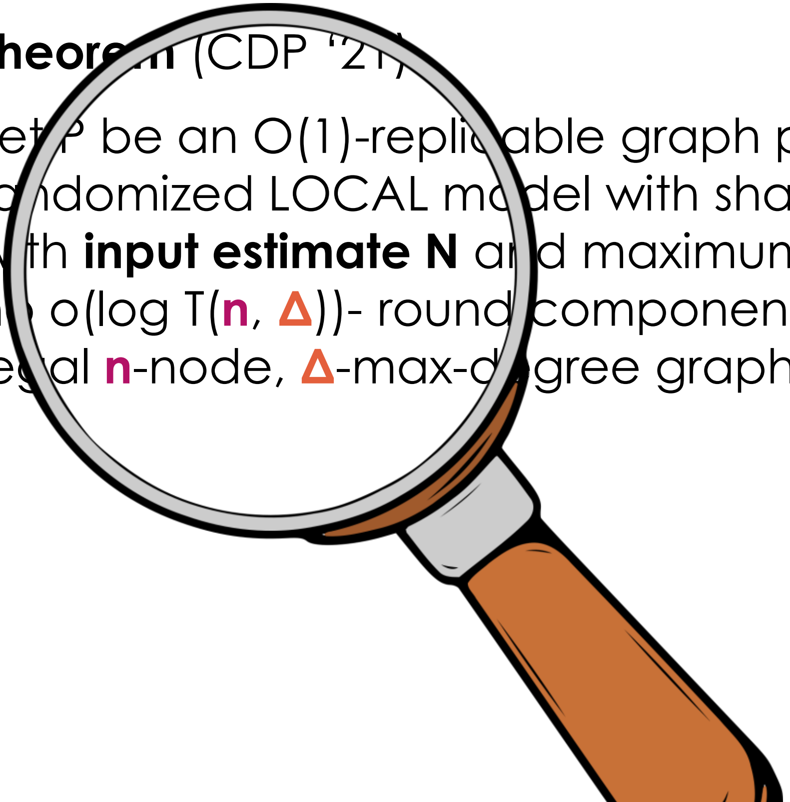

N vs n

In LOCAL, only polynomial estimate N of number of nodes assumed

In MPC, exact value **n** can be easily calculated

This is one of the reasons we must allow knowledge of **n** in component-stability

MPC lower bounds from the theorem are **stronger** in this sense

Lower Bound Statement

Theorem (CDP ‘21)

Let P be an O(1)-replicable graph problem that has a T(N, **Δ**)-round lower bound in the randomized LOCAL model with shared randomness, for constrained function T, on graphs with input estimate N and maximum degree **Δ**, from some normal family G. Then there is no o(log T(**n**, **Δ**))- round component-stable low-space MPC algorithm solving P w.h.p. on legal **n**-node, **Δ**-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

Lower Bound Statement

Theorem (CDP '21)

Let P be an O(1)-replicable graph problem that has a T(N, **Δ**)-round lower bound in the randomized LOCAL model with shared randomness, for constrained function T, on graphs with input estimate N and maximum degree **Δ**, from some normal family G. Then there is no o(log T(**n**, **Δ**))- round component-stable low-space MPC algorithm solving P w.h.p. on **legal** **n**-node, **Δ**-max-degree graphs from G, conditioned on 1 vs 2 cycle conjecture.

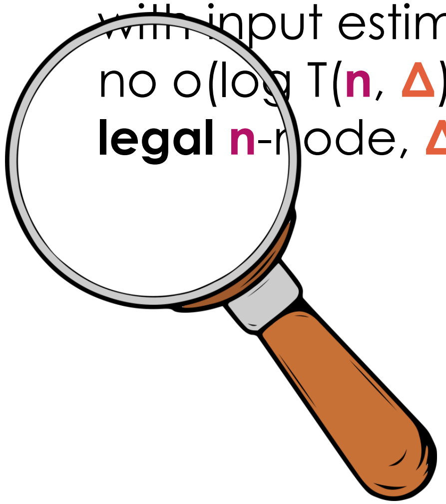

Legal Graphs

Checklist

- Problem is O(1)-replicable
- LOCAL lower bound holds under shared randomness
- Target MPC bound is at most $\Omega(\log \log n)$
- Graph class is normal (hereditary)

Can We Remove The Technicalities?

- Problem is O(1)-replicable — Maybe (or weaken)
- LOCAL lower bound holds under shared randomness — Unlikely
- Target MPC bound is at most $\Omega(\log \log n)$ — Maybe (or weaken)
- Graph class is normal (hereditary) — Unlikely

\+ Component stability — Maybe (but seems difficult)