Limits of Distributed Quantum Advantage

Massimo Equi[2]

Joint work with:

A. Balliu[1], F. Casagrande[1], C. Coupette[2], A. Cruciani[2], F. D'Amore[1], B. Keller[2], H. Lievonen[2], A. Modanese[2], D. Olivetti[1], G. Schmid[3], J. Suomela[2]

[1]GSSI, Italy; [2]Aalto University, Finland; [3]University of Freiburg, Germany

Estonian-Latvian CS Theory Days 2026

Models and Problems

Models of distributed computing

- LOCAL, deterministic and randomized
- quantum-LOCAL
- Non-signaling distributions

Classes of problems

- Locally Checkable Labeling (LCL) problems [SODA 2026]
- Linear Optimization problems [DISC 2025]

LOCAL

- The network is a graph $G = (V, E)$
- Nodes have input and unique IDs
- IDs belong to $\{1, 2, \ldots, n^c\}$
- No limits on computation or message size

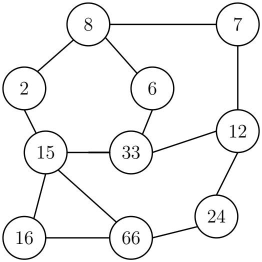

Definitions: LOCAL

LOCAL

- The network is a graph $G = (V, E)$
- Nodes have input and unique IDs
- IDs belong to $\{1, 2, \ldots, n^c\}$
- No limits on computation or message size

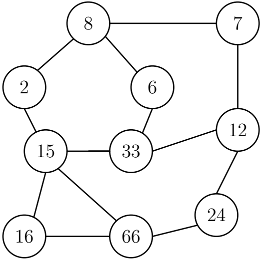

In **one communication round**, every node $v \in V$ performs the following:

- **send** messages to all neighbors
- **receive** messages from all neighbors
- **update** internal state, and possibly stop

Definitions: LOCAL

LOCAL

- The network is a graph $G = (V, E)$
- Nodes have input and unique IDs
- IDs belong to $\{1, 2, \ldots, n^c\}$
- No limits on computation or message size

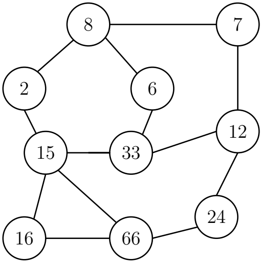

In **one communication round**, every node $v \in V$ performs the following:

- **send** messages to all neighbors
- **receive** messages from all neighbors
- **update** internal state, and possibly stop

Complexity:
Number of communication rounds as a function of n before every node stops

Definitions: rand-LOCAL and quantum-LOCAL

rand-LOCAL

- every node has a **random binary string** as input

quantum-LOCAL

- nodes hold **qubits**, can apply local unitary transformations, and can send and receive qubits

Definitions: Finitely Dependent Distributions

Finitely dependent distributions (T-dependent)

- for any two subsets of nodes X and Y, if $\mathcal{N}_T(X) \cap \mathcal{N}_T(Y) = \varnothing$, then their output is **independent**.

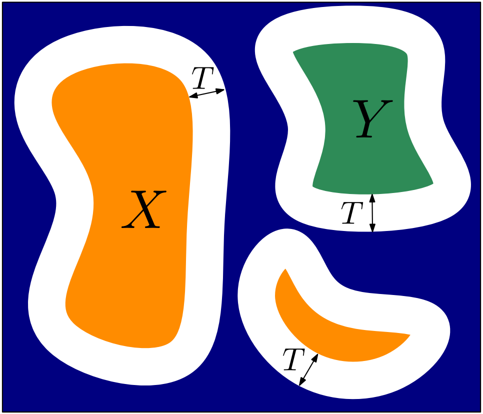

Definitions: Non-Signaling Distributions

Non-signaling distributions

- for any subset of nodes X, changes at distance $> T$ **do not affect** the output distribution of X.

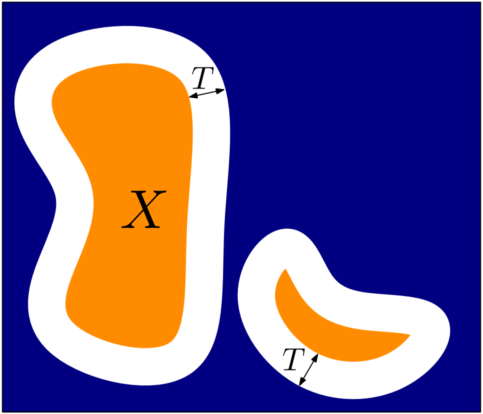

Definitions: Non-Signaling Distributions

Non-signaling distributions

- for any subset of nodes X, changes at distance $> T$ **do not affect** the output distribution of X.

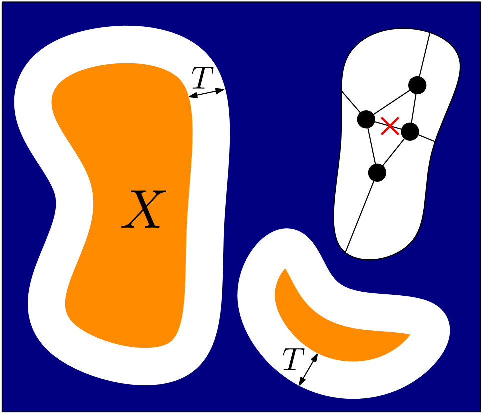

Distributed Landscape

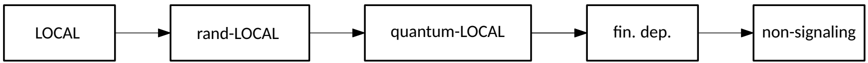

Distributed Landscape

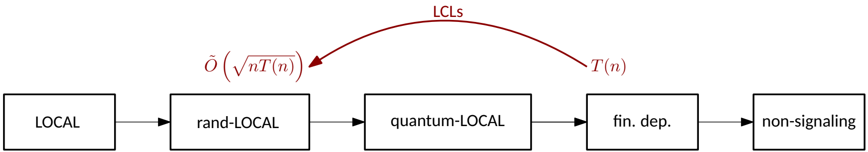

Today's results:

- For LCLs: $T(n)$ in fin. dep. implies $\tilde{O}\left(\sqrt{nT(n)}\right)$ in LOCAL

Distributed Landscape

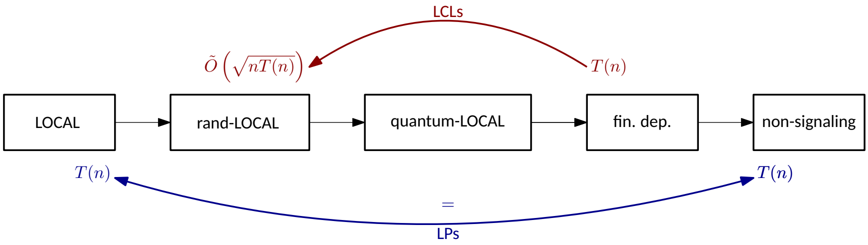

Today's results:

- For LCLs: $T(n)$ in fin. dep. implies $\tilde{O}\left(\sqrt{nT(n)}\right)$ in LOCAL
- For LPs: no separation

Limits for Locally Checkable Labelling Problems

Locally Checkable Labellings (LCLs)
[Naor and Stockmeyer - STOC '93]

- An LCL defines a labelling problem
- Allowed configurations defined as a **set of neighbourhoods**
- Checking radius r is constant

Checking radius r: if every node $v \in V$ sees a correct labelling in its neighbourhood $\mathcal{N}_r(v)$, then the labelling is globally correct.

Limits for Locally Checkable Labelling Problems

Locally Checkable Labellings (LCLs)
[Naor and Stockmeyer - STOC '93]

- An LCL defines a labelling problem
- Allowed configurations defined as a **set of neighbourhoods**
- Checking radius r is constant

Checking radius r: if every node $v \in V$ sees a correct labelling in its neighbourhood $\mathcal{N}_r(v)$, then the labelling is globally correct.

Why LCLs? In a sense, they are the NP problems of distributed computing

Limits for Locally Checkable Labelling Problems

Today we show the following:
[Balliu et al. - SODA '26]

if LCL problem $\mathcal{P}$ is solvable in $T(n)$ in quantum-LOCAL
then $\mathcal{P}$ is solvable in $O(\sqrt{nT(n) \operatorname{poly}\log n})$ in rand-LOCAL

Limits for Locally Checkable Labelling Problems

Today we show the following:
[Balliu et al. - SODA '26]

if LCL problem $\mathcal{P}$ is solvable in $T(n)$ in quantum-LOCAL
then $\mathcal{P}$ is solvable in $O(\sqrt{nT(n)\,\text{poly}\log n})$ in rand-LOCAL

We actually prove this for finitely dependent distributions

Limits for Locally Checkable Labelling Problems

Assume a T-dependent distribution exists for problem $\mathcal{P}$

We want to find a rand-LOCAL algorithm that solves $\mathcal{P}$

Limits for Locally Checkable Labelling Problems

Assume a T-dependent distribution exists for problem $\mathcal{P}$

We want to find a rand-LOCAL algorithm that solves $\mathcal{P}$

We cannot just sample from the distribution blindly

- nodes at distance less than T are dependent

Limits for Locally Checkable Labelling Problems

Assume a T-dependent distribution exists for problem $\mathcal{P}$

We want to find a rand-LOCAL algorithm that solves $\mathcal{P}$

We cannot just sample from the distribution blindly

- nodes at distance less than T are dependent

Our approach: we cluster the graph [Coiteux-Roy et al. - STOC '24]

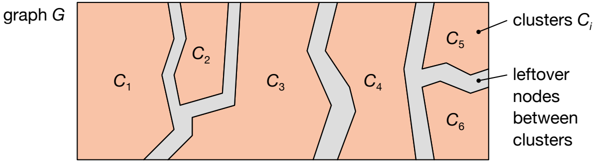

Complexity: $\tilde{O}(\sqrt{nT})$ Cluster diameter: $\tilde{O}(\sqrt{nT})$ Unclust. nodes: $O(\sqrt{n/T})$

Limits for Locally Checkable Labelling Problems

Complexity: $\tilde{O}(\sqrt{nT})$ Cluster diameter: $\tilde{O}(\sqrt{nT})$ Unclust. nodes: $O(\sqrt{n/T})$

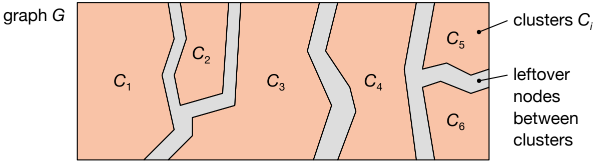

Limits for Locally Checkable Labelling Problems

Complexity: $\tilde{O}(\sqrt{nT})$ Cluster diameter: $\tilde{O}(\sqrt{nT})$ Unclust. nodes: $O(\sqrt{n/T})$

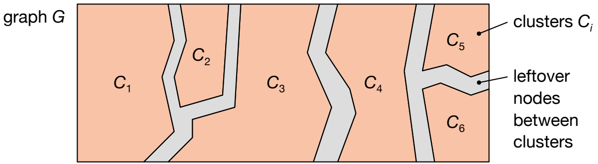

We organize the leftover nodes into partitions at distance $\Omega(T)$

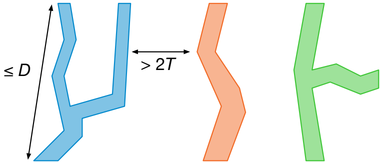

Maximum diameter for a partition: $D = O(T \cdot \sqrt{n/T}) = O(\sqrt{nT})$

Reason: Scout up to distance $2T$, repeat if unclustered node found

Limits for Locally Checkable Labelling Problems

Complexity: $\tilde{O}(\sqrt{nT})$ Cluster diameter: $\tilde{O}(\sqrt{nT})$ Unclust. nodes: $O(\sqrt{n/T})$

Maximum diameter for a partition: $D = O(T \cdot \sqrt{n/T}) = O(\sqrt{nT})$

Now we can sample in every partition independently in time $O(\sqrt{nT})$

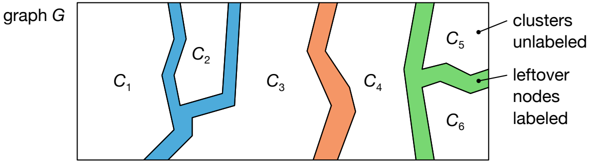

Limits for Locally Checkable Labelling Problems

Complexity: $\tilde{O}(\sqrt{nT})$ Cluster diameter: $\tilde{O}(\sqrt{nT})$ Unclust. nodes: $O(\sqrt{n/T})$

Maximum diameter for a partition: $D = O(T \cdot \sqrt{n/T}) = O(\sqrt{nT})$

Now we can sample in every partition independently in time $O(\sqrt{nT})$

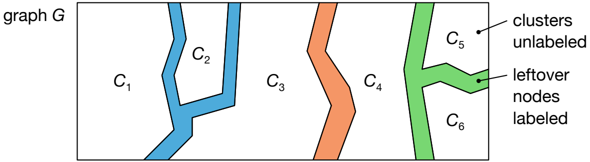

Then complete what's in the clusters by brute force in time $O(\sqrt{nT})$

Total time: $\tilde{O}(\sqrt{nT})$

- Clustering: $\tilde{O}(\sqrt{nT})$
- Partitioning and sampling: $O(\sqrt{nT})$
- Completing by brute force: $O(\sqrt{nT})$

Summary

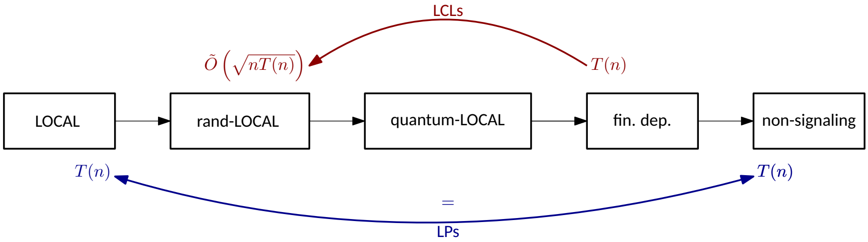
LCLs
$\tilde{O}\left(\sqrt{nT(n)}\right)$
$T(n)$
LOCAL
rand-LOCAL
quantum-LOCAL
fin. dep.
non-signaling
$T(n)$
$T(n)$
=
LPs

Summary

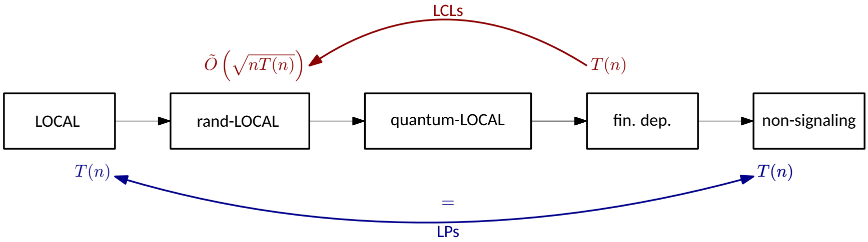

What we proved

- For LCLs: $T(n)$ in non-signaling implies $\tilde{O}(\sqrt{nT(n)})$ in LOCAL

Summary

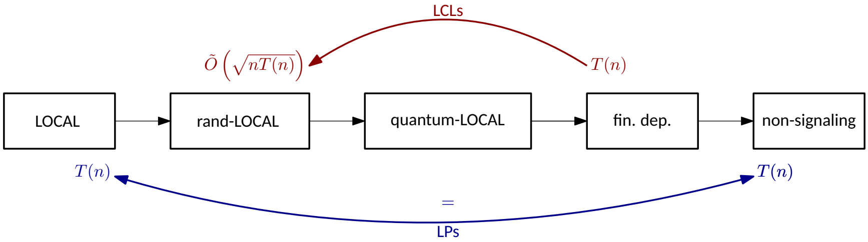

What we proved

- For LCLs: $T(n)$ in non-signaling implies $\tilde{O}(\sqrt{nT(n)})$ in LOCAL

Up next

- For LPs: $T(n)$ in non-signaling implies $T(n)$ in LOCAL

Limits for Linear Programs

Reminder:

- **Independent Set (IS)**: a subset of nodes $I \subseteq V$ such that $\forall v, u \in I : (v, u) \notin E$

Limits for Linear Programs

Reminder:

- **Independent Set (IS)**: a subset of nodes $I \subseteq V$ such that $\forall v, u \in I : (v, u) \notin E$
- **Maximal Independent Set (MIS)**: a subset of nodes $I \subseteq V$ such that $\forall u \in V \setminus I,\ \exists v \in I : (v, u) \in E$

Limits for Linear Programs

Reminder:

- **Independent Set (IS)**: a subset of nodes $I \subseteq V$ such that $\forall v, u \in I : (v, u) \notin E$
- **Maximal Independent Set (MIS)**: a subset of nodes $I \subseteq V$ such that $\forall u \in V \setminus I,\ \exists v \in I : (v, u) \in E$

Example of Distributed LP: **fractional maximum independent set**

$$
\begin{aligned}
\text{maximize} & \sum_{v \in V} x_v \\
\text{subject to} & \sum_{(v,u) \in E} x_v + x_u \leq 1 \quad \forall v \in V \\
& 0 \leq x_v \leq 1 \quad \forall v \in V
\end{aligned}
$$

Limits for Linear Programs

Reminder:

- **Independent Set (IS)**: a subset of nodes $I \subseteq V$ such that $\forall v, u \in I : (v, u) \notin E$
- **Maximal Independent Set (MIS)**: a subset of nodes $I \subseteq V$ such that $\forall u \in V \setminus I,\ \exists v \in I : (v, u) \in E$

Example of Distributed LP: **fractional maximum independent set**

$$
\begin{aligned}
\text{maximize} \sum_{v \in V} x_v \\
\text{subject to} \sum_{(v,u) \in E} x_v + x_u \leq 1 \quad \forall v \in V \\
0 \leq x_v \leq 1 \quad \forall v \in V
\end{aligned}
$$

When $x_e \in \{0, 1\} \quad \forall e \in E$, the solution is a **maximum independent set**, i.e. an MIS of maximum cardinality.

Limits for Linear Programs

We prove that for any distributed LP:

if there is a non-signaling distribution of α-approx. with locality $T(n)$
then there is a LOCAL algorithm producing α-approx. with locality $T(n)$

Limits for Linear Programs

We prove that for any distributed LP:

if there is a non-signaling distribution of α-approx. with locality $T(n)$
then there is a LOCAL algorithm producing α-approx. with locality $T(n)$

Example for fractional MIS. Let

- X_v be the random variable for the output of node v in non-signaling
- $X = \sum_{v \in V} X_v$ be the random variable for the output for entire graph

Limits for Linear Programs

We prove that for any distributed LP:

if there is a non-signaling distribution of α-approx. with locality $T(n)$
then there is a LOCAL algorithm producing α-approx. with locality $T(n)$

Example for fractional MIS. Let

- X_v be the random variable for the output of node v in non-signaling
- $X = \sum_{v \in V} X_v$ be the random variable for the output for entire graph

If X is an α-approx. for P then:

- In LOCAL, explore your radius-$T(n)$ neighborhood
- Set the LOCAL output for v to be $\mathbb{E}[X_v]$
- By linearity of $\mathbb{E}$,

$$\sum_{v \in V} \mathbb{E}[X_v] = \mathbb{E}\left[\sum_{v \in V} X_v\right] = \mathbb{E}[X]$$

- By monotonicity of $\mathbb{E}$, if X is an α-approx. then $\mathbb{E}[X]$ is an α-approx.

Summary

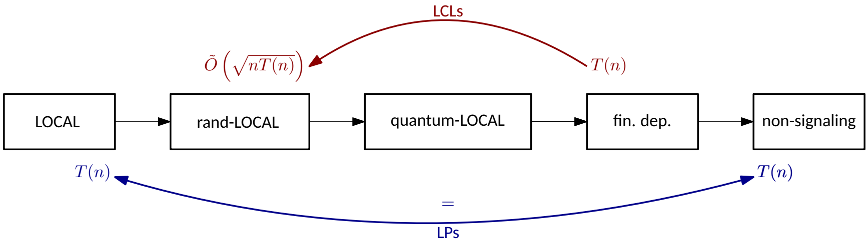

Distributed Quantum Advantage in Locally Checkable Labeling Problems
A. Balliu, F. Casagrande, F. D'Amore, M. Equi, B. Keller, H. Lievonen, D. Olivetti, G. Schmid, J. Suomela
SODA 2026

New Limits on Distributed Quantum Advantage: Dequantizing Linear Programs
A. Balliu, C. Coupette, A. Cruciani, F. D'Amore, M. Equi, H. Lievonen, A. Modanese, D. Olivetti, J. Suomela
DISC 2025

Summary

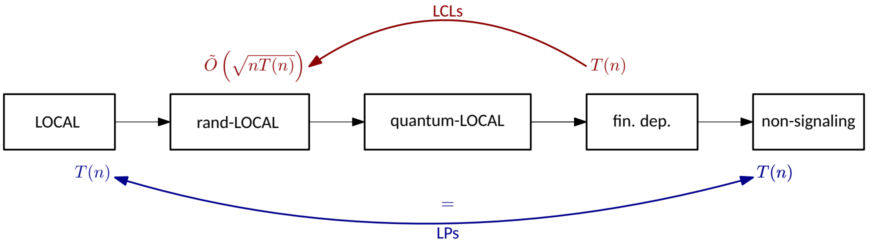

Distributed Quantum Advantage in Locally Checkable Labeling Problems
A. Balliu, F. Casagrande, F. D'Amore, M. Equi, B. Keller, H. Lievonen, D. Olivetti, G. Schmid, J. Suomela
SODA 2026

New Limits on Distributed Quantum Advantage: Dequantizing Linear Programs
A. Balliu, C. Coupette, A. Cruciani, F. D'Amore, M. Equi, H. Lievonen, A. Modanese, D. Olivetti, J. Suomela
DISC 2025

Open question: find an actual LCL showing this separation

Summary

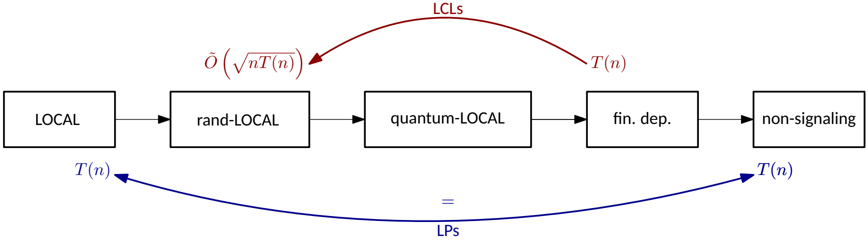

Distributed Quantum Advantage in Locally Checkable Labeling Problems
A. Balliu, F. Casagrande, F. D'Amore, M. Equi, B. Keller, H. Lievonen, D. Olivetti, G. Schmid, J. Suomela
SODA 2026

New Limits on Distributed Quantum Advantage: Dequantizing Linear Programs
A. Balliu, C. Coupette, A. Cruciani, F. D'Amore, M. Equi, H. Lievonen, A. Modanese, D. Olivetti, J. Suomela
DISC 2025

Open question: find an actual LCL showing this separation

Thank you!