universität wien

BA: Does Preprocessing Help under Congestion?

Klaus-Tycho Foerster, Janne H. Korhonen (IST Austria), Joel Rybicki (IST Austria), Stefan Schmid

universität wien

Motivation

- Standard distributed computing models (eg *CONGEST* with O(log n) message sizes & IDs):
  - Network **topology** is **unknown**
  - Compute from **scratch**

- But in many networking applications:
  - Communication **topology** remains **unchanged**
  - Only the problem **input changes**
  - Can we leverage **preprocessing**?

Idea of the ***SUPPORTED*** model (Schmid and Suomela, 2013)

1. Perform any preprocessing on communication graph H
2. Solve problem for subgraph G⊆H in eg *CONGEST* model
   - Use preprocessing information
   - Communicate on H

universität wien

Brief Background

- ***Congested Clique***
  - Introduced at SPAA 2003: Lotker, Pavlov, Patt-Shamir, Peleg
    - Analogy: *SUPPORTED CONGEST* model if communication graph H is a clique

- ***SUPPORTED*** model
  - Introduced for *LOCAL* and *CONGEST* at HotSDN 2013: Schmid and Suomela
  - *CONGEST*: Applications to subgraph detection at OPODIS 2017: Korhonen and Rybicki
  - *LOCAL*: Approximation bounds and connections to *SLOCAL* at INFOCOM 2019: Foerster, Hirvonen, Suomela, Schmid

- **This** BA: How do *CONGEST* lower bounds transfer to the *SUPPORTED CONGEST* model?

universität wien

Many Communication-Complexity Bounds Transfer

- Common observation:
  - Many *CONGEST* lower bounds rely on small cuts
  - Topology information needs to be transferred over congested cut

- High-level idea:
  - If small cut is also present on communication graph, then preprocessing does not help
  - Topology information of input/problem graph still needs to get across congested cut

- Adapt proof from Abboud, Censor-Hillel, Khoury, Paz (arXiv 2019)
  - *Family of lower bound graphs* construction

universität wien

Transfer of Lower Bounds from *CONGEST* to *SUPPORTED CONGEST*

Lower bound	Problem
$\Omega\left(n^{1/2}/\log n\right)$	4-cycle [*Drucker, Kuhn, Oshman PODC'14*], 2k-cycle [*Korhonen, Rybicki OPODIS'17*], Girth ((2 − ε)-apx.) [*Frischknecht, Holzer, Wattenhofer SODA'12*]
$\Omega(n/\log n)$	(2k + 1)-cycle [*DKO PODC'14*], APSP, Diameter ((3/2 − ε)-apx.) [*FHW SODA'12*]
$\Omega\left(n/(\log n)^2\right)$	Diameter on sparse graphs [*Abboud, Censor-Hillel, Khoury DISC'16*]
$\Omega\left(n/(\log n)^3\right)$	On sparse graphs: Diameter and radius ((3/2 − ε)-apx.), eccentricities ((5/3 − ε)-apx.) [*ACHK DISC'16*]
$\Omega\left(n^{2-1/k}/(\mathrm{k}\log n)\right)$	Subgraph detection (for any k) [*Fischer, Gonen, Kuhn, Oshman SPAA'18*]
$\Omega\left(n^2/(\log n)^2\right)$	Min. vertex cover, max. independent set, chrom. number ((4/3 − ε)-apx.), weighted 8-cycle [*Censor-Hillel, Khoury, Paz DISC'17*]
$\Omega\left(n^2\right)$	Identical subgraphs (deterministic only) [*CHKP DISC'17*]

universität wien

Summary and Outlook

- We investigated the power of **preprocessing** in the *CONGEST* model
- Many CONGEST lower bounds hold even under arbitrary **preprocessing**
  - Is SUPPORTED CONGEST maybe the proper way to look at lower bounds?
- Is there a “proper” **separation** between *CONGEST* and *SUPPORTED CONGEST* for **general graphs**?
  - “Proper”: Without relying on identifiers and graph size?
    - Note: Easy on restricted graph classes, e.g., if H has small chromatic number

universität wien

References

[1] Amir Abboud, Keren Censor-Hillel, and Seri Khoury. Near-linear lower bounds for distributed distance computations, even in sparse networks. In *Proc. 30th International Symposium on Distributed Computing (DISC 2016)*, pages 29–42. 2016.

[2] Amir Abboud, Keren Censor-Hillel, Seri Khoury, and Ami Paz. Smaller cuts, higher lower bounds, 2019. arxiv:1901.01630 [cs.DC].

[3] Keren Censor-Hillel, Seri Khoury, and Ami Paz. Quadratic and near-quadratic lower bounds for the CONGEST model. In *Proc. 31st International Symposium on Distributed Computing (DISC 2017)*. 2017.

[4] Andrew Drucker, Fabian Kuhn, and Rotem Oshman. On the power of the congested clique model. In *Proc. Symposium on Principles of Distributed Computing (PODC 2014)*, pages 367–376. 2014.

[5] Orr Fischer, Tzlil Gonen, Fabian Kuhn, and Rotem Oshman. Possibilities and impossibilities for distributed subgraph detection. In *Proc. 30th Symposium on Parallelism in Algorithms and Architectures (SPAA 2018)*, pages 153–162. 2018.

[6] Klaus-Tycho Foerster, Juho Hirvonen, Stefan Schmid, and Jukka Suomela. On the power of preprocessing in decentralized network optimization. In *Proc. 39th IEEE International Conference on Computer Communications (INFOCOM 2019)*. 2019.

[7] Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proc. 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)*, pages 1150–1162. 2012.

[8] Janne H. Korhonen and Joel Rybicki. Deterministic subgraph detection in broadcast CONGEST. In *Proc. 21st International Conference on Principles of Distributed Systems (OPODIS 2017)*. 2017.

[9] Zvi Lotker, Elan Pavlov, Boaz Patt-Shamir, and David Peleg. MST construction in $O(\log \log n)$ communication rounds. In *Proc. 15th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2003)*, pages 94–100. 2003.

[10] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. Distributed verification and hardness of distributed approximation. *SIAM Journal on Computing*, 41 (5): pages 1235–1265, 2012.

[11] Stefan Schmid and Jukka Suomela. Exploiting locality in distributed SDN control. In *Proc. 2nd ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN 2013)*, pages 121–126. 2013.

universität wien

BA: Does Preprocessing Help under Congestion?

Klaus-Tycho Foerster, Janne H. Korhonen (IST Austria), Joel Rybicki (IST Austria), Stefan Schmid