
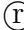
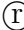



Almost Universally Optimal Distributed Laplacian Solvers via Low-Congestion Shortcuts*

Ioannis Anagnostides¹  Christoph Lenzen²  Bernhard Haeupler³  Goran Zuzic⁴ 
Themis Gouleakis⁵

¹Carnegie Mellon University

²CISPA Helmholtz Center for Information Security

³ETH Zürich & Carnegie Mellon University

⁴ETH Zürich

⁵National University of Singapore

Abstract





In this paper, we refine the (almost) *existentially optimal* distributed Laplacian solver recently developed by Forster, Goranci, Liu, Peng, Sun, and Ye (FOCS ‘21) into an (almost) *universally optimal* distributed Laplacian solver.

Specifically, when the topology is known, we show that any Laplacian system on an n -node graph with *shortcut quality* $\text{SQ}(G)$ can be solved within $n^{o(1)}\text{SQ}(G)\log(1/\varepsilon)$ rounds, where ε is the required accuracy. This almost matches our lower bound which guarantees that any correct algorithm on G requires $\tilde{\Omega}(\text{SQ}(G))$ rounds, even for a crude solution with $\varepsilon \leq 1/2$. Even in the unknown-topology case (i.e., standard CONGEST), the same bounds also hold in most networks of interest. Furthermore, conditional on conjectured improvements in state-of-the-art constructions of low-congestion shortcuts, the CONGEST results will match the known-topology ones.

Moreover, following a recent line of work in distributed algorithms, we consider a hybrid communication model which enhances CONGEST with limited global power in the form of the node-capacitated clique (NCC) model. In this model, we show the existence of a Laplacian solver with round complexity $n^{o(1)}\log(1/\varepsilon)$.

The unifying thread of these results, and our main technical contribution, is the study of novel *congested* generalization of the standard *part-wise aggregation* problem. We develop near-optimal algorithms for this primitive in the Supported-CONGEST model, almost-optimal algorithms in (standard) CONGEST, as well as a very simple algorithm for bounded-treewidth graphs with slightly worse bounds. This primitive can be readily used to accelerate the FOCS’21 Laplacian solver. We believe this primitive will find further independent applications.

⁰Funding acknowledgments are deferred to the next page.

*The author ordering was randomized using <https://www.aeaweb.org/journals/policies/random-author-order/generator>. It is requested that citations of this work list the authors separated by `\textcircled{r}` instead of commas: Anagnostides  Lenzen  Haeupler  Zuzic  Gouleakis.

Contents

1	Introduction	1
1.1	Overview of our Contributions and Techniques	2
1.1.1	The Congested Part-Wise Aggregation Problem	2
1.1.2	Almost Universally Optimal Laplacian Solvers	4
1.2	Further Related Work	5
2	Preliminaries	6
3	The Congested Part-Wise Aggregation Problem	10
3.1	Solving Congested Instances in the CONGEST Model	10
3.1.1	The Layered Graph	11
3.1.2	Treewidth-Bounded Graphs	13
3.1.3	General Graphs	14
3.2	The NCC Model	17
4	Almost Universally Optimal Laplacian Solvers	18
5	Conclusions	19
A	The Laplacian Solver	24
A.1	Low-Congestion Minors	24
A.2	The Laplacian Building Blocks	25
A.2.1	Ultra-Sparsification	25
A.2.2	Sparsified Cholesky	26
A.2.3	Minor Schur Complement	27
A.2.4	Schur Complement Chain	27
A.3	Putting Everything Together	28
B	Omitted Proofs	29
B.1	Proofs from Section 2	29
B.2	Proofs from Section 3	30
B.3	Useful Routines	31
B.4	Ultra-Sparsification: Proof of Lemma A.4	32
B.5	Sparsified Cholesky: Proof of Lemma A.5	34
B.6	Minor Schur Complement: Proof of Lemma A.6	37
B.7	Proof of Theorem A.9	39
B.8	Proof of Proposition 1.1	40
C	Congested Part-Wise Aggregation in the NCC Model	42

⁰Funding acknowledgments. Bernhard Haeupler: Supported in part by NSF grants CCF-1814603, CCF-1910588, NSF CAREER award CCF-1750808, a Sloan Research Fellowship, funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (ERC grant agreement 949272), and the Swiss National Foundation (project grant 200021-184735). Goran Zuzic: Supported in part by the Swiss National Foundation (project grant 200021-184735). Themis Gouleakis: Supported in part by an NRF Fellowship for AI (R-252-000-A33-133). Part of the work was done while visiting the Simons Institute for Theory of Computing.

1 Introduction

The *Laplacian paradigm* has emerged as one of the cornerstones of modern algorithmic graph theory. Integrating techniques from combinatorial optimization with powerful machinery from numerical linear algebra, it was originally pioneered in [ST14] who established the first nearly-linear time solvers for a (linear) Laplacian system. Thereafter, there has been a considerable amount of interest in providing simpler and more efficient solvers [KMP14; Kel+13; KS16]. Indeed, this framework has led to some state of the art algorithms for a wide range of fundamental graph-theoretic problems; e.g., see [AMV21; Mad16; Coh+17; Bra+20; Kel+14; Pen16; AMV20], and references therein. In the distributed setting, a major breakthrough was very recently made in [For+20]. In particular, the authors developed a distributed algorithm that solves any Laplacian system on an n -node graph after $n^{o(1)}(\sqrt{n} + D) \log(1/\varepsilon)$ rounds of the standard CONGEST model, where D represents the hop-diameter of the underlying network and $\varepsilon > 0$ is the error of the solver. Moreover, they showed that their algorithm is *existentially optimal*, up to the $n^{o(1)}$ factor, establishing a lower bound of $\tilde{\Omega}(\sqrt{n} + D)$ rounds via a reduction from the $s - t$ connectivity problem [Das+11].

This *existential* lower bound in the CONGEST model of distributed computing should hardly come as any surprise. Indeed, it is well-known by now that a remarkably wide range of *global* optimization problems, including minimum spanning tree (MST), minimum cut (Min-Cut), maximum flow, and single-source shortest paths (SSSP), require $\tilde{\Omega}(\sqrt{n} + D)$ rounds¹ [PR99; Elk04; Das+11]. The same limitation generally applies to any non-trivial approximation and even under randomization. Nonetheless, these lower bounds are constructed on some pathological graph instances which arguably do not occur in practice. This begs the question: *Can we obtain more refined performance guarantees based on the underlying topology of the communication network?* The framework of *low-congestion shortcuts*, introduced by [GH16], demonstrated that bypassing the notorious $\Omega(\sqrt{n})$ lower bound is possible: MST and Min-Cut on *planar graphs* can be solved in $\tilde{O}(D)$ rounds. This is crucial, given that in many graphs of practical significance the diameter is remarkably small; e.g., $D = \text{polylog}(n)$ (as is folklore, this holds for most social networks), implying *exponential improvements* over generic algorithms used for general graphs. In the context of the distributed Laplacian paradigm, we raise the following question:

Is there a faster distributed Laplacian solver under “non-worst-case” families of graphs in the CONGEST model?

The only known technique in distributed computing for designing algorithms that go below the \sqrt{n} -bound is the low-congestion shortcut framework of Ghaffari and Haeupler [GH16], and its large ecosystem of tools built around it [HIZ16a; HIZ16b; HWZ21; GH20; Zuz+22; GZ22; HRG22]. However, the “ ρ -congested minor” primitive introduced and extensively used in the novel distributed Laplacian solver [For+20] is out-of-reach from the current set of tools available in the low-congestion shortcut framework. We address this issue by introducing an analogous primitive called *ρ -congested part-wise aggregation*, which greatly simplifies the interface used by [For+20]. We then extend the low-congestion shortcut framework with new techniques that enables it to near-optimally solve this primitive: we provide both an algorithm that utilizes the very recent hop-constrained expander decompositions for shortcut construction [HRG22] to solve the primitive in general graphs with a linear dependence on ρ , as well as a very simple algorithm with a quadratic ρ -dependence for bounded-treewidth graphs. Finally, we settle our original question in the positive by establishing

¹As usual, we use the notation $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$ to suppress polylogarithmic factors on n .

that our new primitive can be readily used to accelerate the distributed Laplacian solver for non-worst-case topologies.

Specifically, we show our new techniques are sufficient to lift the existentially optimal algorithm [For+20] to a *universally optimal* algorithm—modulo $n^{o(1)}$ factor inherent in the prior approach—for distributedly solving a Laplacian system, meaning that, *for any topology*, our algorithm is essentially as fast as possible. In other words, for any graph, our algorithm almost matches the best possible (correct) algorithm for that graph. This result is unconditional in essentially all settings of interest (see [Theorem 1.2](#) for details), but relies on conjectured improvements of current state-of-the-art constructions of low-congestion shortcuts to achieve unqualified universal optimality—like all other results in the area.

Furthermore, another concrete way of bypassing the $\tilde{\Omega}(\sqrt{n} + D)$ lower bound, besides investigating non-worst-case families of graphs, is by enhancing the local communication network with a limited amount of *global power*. Indeed, research concerning *hybrid* networks was recently initiated in the realm of distributed algorithms [Aug+20], although networks combining different communication modes have already found numerous applications in real-life computing systems; as such, hybrid networks have been intensely studied in other areas of distributed computing (see [CGC16; Wan+10; KS18], and references therein). In this paper, we will enhance the standard CONGEST model with the recently introduced *node-capacitated clique* (henceforth NCC) [Aug+19]. The latter model enables all-to-all communication, but with severe capacity restrictions for every node. The integration of these models will be referred to as the HYBRID model for the rest of this work. This leads to the following central question:

Is there a faster distributed Laplacian solver in the HYBRID model?

Our paper essentially settles this question by showing the same ρ -congested part-wise aggregation primitive can be efficiently solved in $\tilde{O}(\rho)$ rounds of NCC, implying an almost optimal $n^{o(1)}$ -round distributed algorithm for solving Laplacian systems in the HYBRID model. A conceptual contribution of our approach is that we treat both CONGEST, Supported-CONGEST, and HYBRID in a *unified way* through the lens of the low-congestion shortcut framework, by designing our algorithm using high-level primitives and leaving the model-specific translations to the framework itself. We note that a similar unified view of PRAM (i.e., parallel) and CONGEST (i.e., distributed) graph algorithms through the same lens has led to very recent breakthroughs on long-standing open problems for both of these settings [Li22].

1.1 Overview of our Contributions and Techniques

The unifying thread and the main technical ingredient of our (almost) universally optimal distributed Laplacian solvers is a new fundamental communication primitive which we refer to as the *congested part-wise aggregation problem*. Specifically, we develop near-optimal algorithms for solving this problem in the (Supported-)CONGEST and the NCC model ([Section 3](#)), and then we utilize this primitive to develop almost universally optimal Laplacian solvers in [Section 4](#).

1.1.1 The Congested Part-Wise Aggregation Problem

To introduce the congested part-wise aggregation problem, let us first give some basic background. The aforementioned Ghaffari-Haeupler framework of low-congestion shortcuts revolves around the

so-called *part-wise aggregation problem* posed as follows: “The graph is partitioned into *disjoint* and individually-connected parts, and we need to compute some simple aggregate function for each part, e.g., the minimum of the values held by the nodes in a given part” [GH16] (see Definition 2.1 for a formal definition). Importantly, it has been shown that this primitive can be solved efficiently in *structured* topologies, and that many problems (including the MST, shortest path, min-cut, etc.) reduce to a small number of calls to a part-wise aggregation oracle, leading to universally optimal algorithms. Unfortunately, it is not clear how to reduce solving a Laplacian system to (a small number of) part-wise aggregation calls and in this paper, we primarily address this issue.

Our first technical contribution is to extend the framework of low-congestion shortcuts by studying a more general primitive: one that incorporates *congestion* (of the input parts) into the underlying *part-wise aggregation* instance. More precisely, unlike the standard part-wise aggregation problem, we allow each node to participate in up to $\rho \in \mathbb{Z}_{\geq 1}$ aggregation parts (see Definition 3.1). We later show that efficient solutions to this primitive leads to efficient distributed Laplacian solvers.

We first remark that a natural strategy for solving congested part-wise aggregation instances does not work: congested instances *cannot*, in general, be directly reduced to a “small” collection of 1-congested instances, thereby necessitating a more refined approach. To this end, our approach is based on “lifting” the underlying communication network \overline{G} into its ρ -layered version $\widehat{G}_{O(\rho)}$: every edge is replaced with a matching and every node with a ρ -clique. The importance of this transformation is that, as we show in Lemma 3.3, the ρ -congested part-wise aggregation problem can be reduced to a 1-congested instance on the ρ -layered graph (Section 3.1.1). This is first established under the assumption that individual parts correspond to simple paths, and then we extend our results to general parts by following Haeupler, Wajc, and Zuzic [HWZ21]. In light of this reduction, we next focus on solving the 1-congested part-wise aggregation instance on the layered graph.

As a warm-up, we treat graphs with bounded *treewidth* $\text{tw}(G)$ (Definition 2.8). It is known from [HIZ16b] that on a graph G with treewidth $\text{tw}(G)$, a 1-congested part-wise aggregation instance can be solved in $\tilde{O}(\text{tw}(G)D)$ rounds of CONGEST. Keeping this in mind, we first show that the treewidth of the ρ -layered graph \widehat{G}_ρ can only increase by a factor of ρ compared to the original graph (Lemma 3.8). Hence, we can solve 1-congested instances in $\widehat{G}_{O(\rho)}$ in $\tilde{O}(\rho \text{tw}(\overline{G})D)$ rounds (when the underlying network is $\widehat{G}_{O(\rho)}$), which in turn allows us to solve ρ -congested instances on \overline{G} in $\tilde{O}(\rho^2 \text{tw}(G)D)$ time in G (another ρ factor is necessary to simulate $\widehat{G}_{O(\rho)}$ in \overline{G}). This positive result poses a natural question: can we achieve similar results on graphs with bounded *minor density* $\delta(G)$ (Definition 2.6)? However, the answer to this question is negative: *minor density* can blow up even for a 2-layered planar graph (see Observation 3.10), making such a result impossible.

Then, we look at arbitrary graphs G : it is known that 1-congested part-wise aggregation instances can be solved in a number of rounds that is controlled by $\text{SQ}(G)$, where $\text{SQ}(G)$ is the *shortcut quality* of G (a certain graph parameter we formalize in Definition 2.4). Specifically, it can be solved in $\tilde{O}(\text{SQ}(G))$ rounds when the topology is known in advance² [HWZ21] and $\text{poly}(\text{SQ}(G)) \cdot n^{o(1)}$ in general CONGEST [HRG22]. The shortcut quality parameter is significant because it was shown that many distributed problems (including the MST, shortest path, min-cut, and—Laplacian solving, as we show later) require $\tilde{\Omega}(\text{SQ}(G))$ rounds in CONGEST to be solved on G [HWZ21]. Therefore,

²This model is also known as the *supported* CONGEST. That is, CONGEST under the assumption that the topology is known; see Section 2 for a formal description of the model. Our techniques also apply in the full generality of CONGEST, as we explain in the sequel.

algorithms that have an upper bound close to $\text{SQ}(G)$ are *universally optimal*.

With the end goal of solving the 1-congested part-wise aggregations on layered graphs \widehat{G}_ρ in time controlled by $\text{SQ}(G)$, our main result established that *the shortcut quality of the ρ -layered graph does not increase* (modulo polylogarithmic factors) as compared to the original graph (**Theorem 3.11**). This has a plethora of important consequences: (1) when $\text{SQ}(G) \leq n^{o(1)}$, we can unconditionally solve ρ -congested part-wise aggregation instances in $\rho \cdot n^{o(1)}$ CONGEST rounds and (2) when the topology of G is known, there exists a distributed algorithm which solves any ρ -congested part-wise aggregation problem in $\rho \cdot \widetilde{O}(\text{SQ}(G))$ rounds. As a consequence of our general result, the shortcut quality of any 2-layered planar graph is $\widetilde{O}(D)$ since it is known that the shortcut quality of a planar graph is $\widetilde{O}(D)$ [GH16]. This constitutes perhaps the most natural example of a graph whose minor density is very far from the shortcut quality; the only other example documented in the literature so far is that of *expander graphs*.

Our proof proceeds by employing alternative characterizations of the shortcut quality in terms of certain communication tasks. Specifically, shortcut quality can be shown to be equal (modulo polylogarithmic factors) to the following two-player max-min game: the first (max) player chooses k sources and k sinks in the graph such that we can find k node-disjoint paths matching the sources with the sinks; then the second (min) player finds the smallest so-called *quality* Q such that there exist k paths matching the sources with the sinks with the path lengths being at most Q and each edge of the underlying graph supporting at most Q of second player's paths. This characterization allows us to compare the shortcut quality of \widehat{G}_ρ with \overline{G} as follows: take the worst-case (first player's) set of sources and sinks in \widehat{G}_ρ . Project them to \overline{G} and note they have node congestion ρ (due to the construction of \widehat{G}_ρ). Then, we show we can decompose (i.e., partition) these set of sources and sinks into $\widetilde{O}(\rho)$ pairs of sub-sources and sub-sinks that are node-disjointly connectable in G . However, each such set enjoys paths of quality $\text{SQ}(G)$, hence embedding each such pair in a separate layer of \widehat{G}_ρ shows that the shortcut quality of $\text{SQ}(\widehat{G}_\rho)$ is at most $\widetilde{O}(\text{SQ}(\overline{G}))$. Although this general approach improves over our result for treewidth-bounded graphs we previously described, our approach for the latter class of graphs is substantially simpler and more suited for potential practical applications.

1.1.2 Almost Universally Optimal Laplacian Solvers

First, we note that any distributed Laplacian solver that always correctly outputs an answer on a fixed graph G must take at least $\widetilde{\Omega}(\text{SQ}(G))$ rounds, giving us a lower bound to compare ourselves with. Our refined lower bound uses the hardness result recently shown by [HWZ21] for the spanning connected subgraph problem, applicable for *any* (i.e., non-worst-case) graph G . Specifically, we show that a Laplacian solver can be leveraged to solve the spanning connected subgraph problem, thereby substantially strengthening the lower bound in [For+20].

Proposition 1.1. *Consider a graph \overline{G} with shortcut quality $\text{SQ}(\overline{G})$. Then, solving a Laplacian system on \overline{G} with $\varepsilon \leq \frac{1}{2}$ requires $\widetilde{\Omega}(\text{SQ}(\overline{G}))$ rounds in both CONGEST and Supported-CONGEST models.*

On the upper-bound side, we utilize the congested part-wise aggregation primitive to improve and refine the Laplacian solver of [For+20], leading to a substantial improvement in the round complexity under *structured* network topologies.

Theorem 1.2. *Consider any n -node graph G with shortcut quality $\text{SQ}(G)$ and hop-diameter D .*

There exists a distributed Laplacian solver with error $\varepsilon > 0$ with the following guarantees:

- In the Supported-CONGEST model, it requires $n^{o(1)} \text{SQ}(G) \log(1/\varepsilon)$ rounds.
- In the CONGEST model, it requires $n^{o(1)} \text{poly}(\text{SQ}(G)) \log(1/\varepsilon)$ rounds.
- In the CONGEST model on graphs with minor density δ , it requires $n^{o(1)} \delta D \log(1/\varepsilon)$ rounds.

We note that the above algorithm is almost (up to inherent $n^{o(1)}$ factors) universally optimal for most settings of interest. Since it is (almost) matching the $\text{SQ}(G)$ -lower-bound, it is unconditionally universally optimal when the topology is known in advance (i.e., Supported-CONGEST). Furthermore, in standard CONGEST, we give almost universally optimal $Dn^{o(1)} \log(1/\varepsilon)$ -round algorithms for topologies that include planar graphs, $n^{o(1)}$ -genus graphs, $n^{o(1)}$ -treewidth graphs, excluded-minor graphs, since all of them are graphs with minor density $\delta(G) = n^{o(1)}$. Furthermore, for the realistic case of $D \leq n^{o(1)}$, it holds for most networks of interest that $\text{SQ}(G) \leq n^{o(1)}$ (e.g., expanders, hop-constrained expanders, as well as all classes mentioned earlier), for which we get $n^{o(1)} \log(1/\varepsilon)$ -round solvers. Finally, the conjectured improvements of the state-of-the-art of almost-optimal low-congestion shortcut constructions would immediately lift our results to be unconditionally universally optimal in CONGEST. However, the issue is orthogonal and out-of-scope of this paper.

Furthermore, in HYBRID we obtain an almost optimal complexity in *general graphs*:

Theorem 1.3. *Consider any n -node graph. There exists a distributed Laplacian solver in the HYBRID model with round complexity $n^{o(1)} \log(1/\varepsilon)$, where $\varepsilon > 0$ is the error of the solver.*

This implies a remarkably fast subroutine for solving a Laplacian system in HYBRID under arbitrary topologies. As a result, we corroborate the observation that a very limited amount of global power can lead to substantially faster algorithms for certain optimization problems, supplementing a recent line of work [CLP21b; Aug+20; KS20; FHS20; CLP21a; Göt+21; KS22; Coy+22]. Furthermore, our framework based on the congested part-wise aggregation problem allows for a unifying treatment of both (Supported-)CONGEST and HYBRID, and we consider this to be an important conceptual contribution of our work. Indeed, as we previously explained, both of our accelerated Laplacian solvers rely on faster algorithms for solving the congested part-wise aggregation problem. In particular, for (Supported-)CONGEST we have already described our approach in detail, while in the HYBRID model we employ certain communication primitives developed in [Aug+19] for dealing with congestion in part-wise aggregations. A byproduct of our results is that the framework of low-congestion shortcuts interacts particularly well with the HYBRID model, as was also observed in [AG21].

1.2 Further Related Work

Our main reference point is the recent Laplacian solver of Forster, Goranci, Liu, Peng, Sun, and Ye [For+20] with *existentially* almost-optimal complexity of $n^{o(1)}(\sqrt{n} + D) \log(1/\varepsilon)$ rounds, where $\varepsilon > 0$ represents the error of the solver. Specifically, they devised several new ideas and techniques to circumvent certain issues which mostly relate to the bandwidth restrictions of the CONGEST model; these building blocks, as well as the resulting Laplacian solver are revisited in our work to refine the performance of the solver. We are not aware of any previous research addressing this problem in the distributed context. On the other hand, the Laplacian paradigm has attracted a

considerable amount of interest in the community of parallel algorithms. Most notably, we refer to [PS14; Ble+14]. These approaches in the PRAM model of parallel computing fail—at least without non-trivial modifications—to lead to a almost-optimal solver in the distributed context [For+20].

In addition to being a problem of independent interest, solving Laplacian systems often leads to a plethora of very fast algorithms (albeit typically polynomially-away from being optimal) for other problems such as (exact) maximum flow [Mad16], min-cost flow [AMV21], shortest paths with negative weights [Coh+17], etc. The recent distributed Laplacian solver [For+20] also contributed fast analogues of these algorithms in the distributed model. A natural question to ask is whether we can also use our techniques to make these algorithms work for more structured graphs. However, these algorithms rely on directed or exact shortest path computations, which currently represent a major barrier for shortcut-based approaches. Moreover, the same set of problems represent a barrier even for existentially-optimal approaches as the current state-of-the-art is a factor of $D^{1/4}$ away from achieving unqualified existential optimality [CM21].

Research concerning hybrid communication networks in distributed algorithms was recently initiated by [Aug+20]. Specifically, they investigated the power of a model which integrates the standard LOCAL model [Lin92] with the recently introduced node-capacitated clique (NCC) [Aug+19], focusing mostly on distance computation tasks. Several of their results were subsequently improved and strengthened in subsequent works [KS20; CLP21a] under the same model of computation. In our work we consider a substantially weaker model, imposing a severe limitation on the communication over the “local edges”. This particular variant has been already studied in some recent works for a variety of fundamental problems [FHS20; Göt+21].

The NCC model, which captures the global network in all hybrid models studied thus far, was introduced in [Aug+19] partly to address the unrealistic power of the *congested clique* (CLIQUE) [Lot+03]. In the latter model each node can communicate *concurrently and independently* with *all* other nodes by $O(\log n)$ -bit messages. In contrast, the NCC model allows communication with $O(\log n)$ (arbitrary) nodes per round. As a result, in the HYBRID model and under a sparse local network, only $\tilde{\Theta}(n)$ bits can be exchanged overall per round, whereas CLIQUE allows for the exchange of up to $\tilde{\Theta}(n^2)$ (distinct) bits. As evidence for the power of CLIQUE we note that even slightly super-constant lower bounds would give new lower bounds in circuit complexity, as implied by a simulation argument in [DKO14].

2 Preliminaries

General notation We denote with $[k] := \{1, 2, \dots, k\}$. Graphs throughout this paper are undirected. The nodes and the edges of a given graph G are denoted as $V(G)$ and $E(G)$, respectively. We also use $n := |V(G)|$ for brevity. The graphs are often weighted, in which case we assume (as is standard) that for all $e \in E(G)$, $w(e) \in \{1, 2, \dots, \text{poly}(n)\}$. We will denote the hop-diameter of a graph G with $D(G)$ (the hop-diameter ignores weights). Moreover, we use $A \uplus B$ to denote the multiset union, i.e., each element is repeated according to its multiplicity; this operation corresponds to disjoint unions when $A \cap B = \emptyset$.

Communication models The communication network consists of a set of \bar{n} entities with $[\bar{n}] := \{1, 2, \dots, \bar{n}\}$ being the set of their IDs, and a local communication *topology* given by a graph \overline{G} .³ We define $D := D(\overline{G})$ to be the (hop-)diameter of the underlying network. At the beginning, each node knows its own unique $O(\log \bar{n})$ -bit identifier as well as the weights of the incident edges. Communication occurs in *synchronous rounds*, and in every round nodes have unlimited computational power to process the information they possess. We will consider models with both *local* and *global* communication modes.

The *local* communication mode will be modeled with the *CONGEST model* [Pel00] and *Supported-CONGEST model* [SS13], for which in each round every node can exchange an $O(\log \bar{n})$ -bit message with each of its neighbors in \overline{G} via the *local* edges. In the (standard) CONGEST model, each node $v \in V(\overline{G})$ initially only knows the identifiers of each node in v 's own neighborhood, but has no further knowledge about the topology of the graph. On the other hand, in the Supported-CONGEST model, all nodes know the entire topology of \overline{G} upfront, but not the input.

The *global* communication mode will be modeled using *NCC* [Aug+19], for which in each round every node can exchange $O(\log \bar{n})$ -bit messages with $O(\log \bar{n})$ arbitrary nodes via *global* edges. If the capacity of some channel is exceeded, i.e., too many messages are sent to the same node, it will only receive an *arbitrary* (potentially adversarially selected) subset of the information based on the capacity of the network; the rest of the messages are dropped. In this context, we will let HYBRID be the integration of CONGEST and NCC (i.e., nodes have both a *local* and a *global* communication mode at their disposal).

The performance of a distributed algorithm will be measured in terms of its *round complexity*—the number of rounds required so that every node knows its part of the output. For randomized algorithms it will suffice to reach the desired state with high probability.⁴ We will assume throughout this work that nodes have access to a common source of randomness; this comes without any essential loss of generality in our setting [Gha15]. When talking about a distributed algorithm for a specific problem (e.g., Laplacian solving, part-wise aggregation, etc.) we assume the input is appropriately *distributedly stored* (i.e., each node will know its own part) and, upon termination, it will be required that the output is appropriately distributedly stored. The appropriate way to distributedly store the input and output will be explained in the problem definition.

Low-Congestion Shortcuts A recurring scenario in distributed algorithms for global problems (e.g. MST) boils down to solving the following part-wise aggregation problem:

Definition 2.1 (Part-Wise Aggregation Problem). Consider an n -node graph G whose node set $V(G)$ is *partitioned* into k (disjoint) parts $P_1 \uplus \dots \uplus P_k \subseteq V(G)$ such that each induced subgraph $G[P_i]$ is *connected*. In the *part-wise aggregation* problem, each node $v \in V$ is given its part-ID (if any) and an $O(\log n)$ -bit value $\mathbf{x}(v)$ as input. The goal is that, for every part P_i , all nodes in P_i learn the part-wise aggregate $\bigoplus_{w \in P_i} \mathbf{x}(w)$, where \bigoplus is an arbitrary pre-defined *aggregation function*.

Throughout this paper, we will assume that the aggregation function \bigoplus is commutative and

³To avoid any possible confusion we point out that, for consistency with the nomenclature of Forster, Goranci, Liu, Peng, Sun, and Ye [For+20], we henceforth reserve \overline{G} to denote the underlying *communication network*, while G is used in statements regarding arbitrary graphs.

⁴We say that an event holds with high probability if it occurs with probability at least $1 - 1/n^c$ for a (freely choosable) constant $c > 0$.

associative (e.g. min, sum, logical-AND), although this is not strictly needed (e.g., see [GZ22]). To give a concrete example, in the context of Boruvka’s algorithm for the MST problem, determining the minimum-weight outgoing edge for each part is an instance of a part-wise aggregation problem with $\oplus := \min$. To solve such problems, Ghaffari and Haeupler [GH16] introduced a natural combinatorial graph structure which they refer to as *low-congestion shortcuts*.

Definition 2.2 (Low-Congestion Shortcuts). Consider a graph G whose node set $V(G)$ is *partitioned* into k (disjoint) parts $P_1 \uplus \dots \uplus P_k \subseteq V(G)$ such that each induced subgraph $G[P_i]$ is *connected*. A collection of subgraphs H_1, \dots, H_k is a *shortcut* of G with *congestion* c and *dilation* d if the following properties hold: (i) the (hop) diameter of each subgraph $G[P_i] \cup H_i$ is at most d , and (ii) every edge is included in at most c many of the subgraphs H_i . The quantity $Q = c + d$ will be referred to as the *quality* of the shortcut.

Importantly, a shortcut of quality Q allows us to solve the part-wise aggregation problem in $\tilde{O}(Q)$ rounds of CONGEST, as formalized below. For self-sufficiency, we include the proof in [Appendix B.1](#).

Proposition 2.3. *Suppose that P_1, \dots, P_k is any part-wise aggregation instance in a communication network \bar{G} . Given a shortcut of quality Q , we can solve with high probability the part-wise aggregation problem in $\tilde{O}(Q)$ CONGEST rounds.*

Shortcut Quality and Construction of Shortcuts Shortcut quality, introduced below, is a fundamental graph parameter that has been proven to characterize the complexity of many important problems in distributed computing.

Definition 2.4. Given a graph $G = (V, E)$, we define the *shortcut quality* $\text{SQ}(G)$ of G as the optimal (smallest) shortcut quality of the worst-case partition of V into disjoint and connected parts $P_1 \uplus P_2 \uplus \dots \uplus P_k \subseteq V$.

For fundamental problems such as MST, SSSP, and Min-Cut any correct algorithm requires $\tilde{\Omega}(\text{SQ}(\bar{G}))$ rounds on any network \bar{G} , even if we allow randomized solutions and (non-trivial) approximation factors. In fact, this limitation holds even when the network topology \bar{G} is known to all nodes in advance [HWZ21]. We remark that $\tilde{\Omega}(D(\bar{G})) \leq \text{SQ}(\bar{G}) \leq O(D(\bar{G}) + \sqrt{\bar{n}})$, and the upper bound is known to be tight in certain (pathological) worst-case graph instances. This explains the notorious (existential) $\tilde{\Omega}(D + \sqrt{\bar{n}})$ lower bound pervasive in distributed computing [Das+11].

Moreover, *assuming* fast distributed algorithms for constructing shortcuts of quality competitive with $\text{SQ}(\bar{G})$, all of the aforementioned problems can be solved in $\tilde{O}(\text{SQ}(\bar{G}))$ rounds [GH16; Zuz+22; GZ22]. However, the key issue here is the algorithmic construction of the shortcuts upon which the above papers rely. While there has been a lot of recent progress in this regard, current algorithms are quite complicated and have sub-optimal guarantees. We recall below these state-of-the-art $\text{SQ}(\bar{G})$ -competitive construction results.

Theorem 2.5. *There exists a distributed algorithm that, given any part-wise aggregation instance on any \bar{n} -node graph \bar{G} , computes with high probability a shortcut with the following guarantees:*

- *In CONGEST, the shortcut has quality $\text{poly}(\text{SQ}(\bar{G})) \cdot \bar{n}^{o(1)}$ and the algorithm terminates in $\text{poly}(\text{SQ}(\bar{G})) \cdot \bar{n}^{o(1)}$ rounds [HRG22].*

- In Supported-CONGEST, the shortcut has quality $\tilde{O}(\text{SQ}(\overline{G}))$ and the algorithm terminates in $\tilde{O}(\text{SQ}(\overline{G}))$ rounds [HWZ21].

Universal Optimality A distributed algorithm is said to be α -*universally optimal* if, on every network graph \overline{G} , it is α -competitive with the fastest correct algorithm on \overline{G} [HWZ21]. Even the existence of such algorithms is not at all clear as it would seem possible that vastly different algorithms are required to leverage the structure of different networks. Nevertheless, a remarkable consequence of [Theorem 2.5](#) is that in Supported-CONGEST we can design $\tilde{O}(1)$ -universally optimal algorithms for many fundamental optimization problems. Moreover, efficient shortcut construction is the only obstacle towards achieving these results in the full generality of CONGEST, which is an issue orthogonal and out of scope for this paper. Still, the aforementioned results are sufficient to design $\overline{n}^{o(1)}$ -universally optimal algorithms on graphs that have shortcut quality $\text{SQ}(\overline{G}) = \overline{n}^{o(1)}$, as it is arguably the case in most networks of practical interest.

Graphs Excluding Dense Minors It turns out that the crucial issue of efficient shortcut construction can be resolved with a near-optimal, simple, and even deterministic algorithm for the rich class of graphs with *bounded minor density*. Formally, let us first recall the following definition.⁵

Definition 2.6 (Minor Density). The *minor density* $\delta(G)$ of a graph G is defined as

$$\delta(G) = \max \left\{ \frac{|E'|}{|V'|} : H = (V', E') \text{ is a minor of } G \right\}.$$

It should be noted that $\delta(G) = \tilde{\Theta}(r(G))$, where $r(G)$ is the *complete-graph minor size*, i.e., $r(G) = \max\{r : K_r \text{ is a minor of } G\}$ [Tho84; Tho01]. Furthermore, any family of graphs closed under taking minors (such as planar graphs) has a constant minor density. For such graphs, Ghaffari and Haeupler [GH20] established efficient shortcut construction:

Theorem 2.7 ([GH20]). *Any graph G with hop-diameter D and minor density $\delta(G)$ admits shortcuts of quality $\tilde{O}(\delta D)$, which can be constructed with high probability in $\tilde{O}(\delta D)$ rounds of CONGEST.*

The (linear) dependency on the minor density is existentially optimal [GH20, Lemma 3.2]. It should be noted that, in the context of [Theorem 2.7](#), there is also a deterministic distributed algorithm with a slightly worse guarantee [GH20]. Some of our results apply for communication networks with *bounded treewidth*, so let us recall the following definition.

Definition 2.8 (Tree Decomposition and Treewidth). A *tree decomposition* of a graph G is a tree T with tree-nodes X_1, \dots, X_k , where each X_i is a subset of $V(G)$ satisfying the following properties:

1. $V = \bigcup_{i=1}^k X_i$;
2. For any node $u \in V(G)$, the tree-nodes containing u form a connected subtree of T ;
3. For every edge $\{u, v\} \in E(G)$, there exists a tree-node X_i which contains both u and v .

The *width* w of the tree decomposition is defined as $w := \max_{i \in [k]} |X_i| - 1$. Moreover, the *treewidth* $\text{tw}(G)$ of G is defined as the minimum of the width among all possible tree decompositions of G .

⁵See the first part of [Definition A.2](#) for a formal description of a minor.

Bounded-treewidth graphs inherit all of the nice properties guaranteed by [Theorem 2.7](#), as implied by the following well-known fact.

Fact 2.9. *For any graph G , $\delta(G) \leq \text{tw}(G)$.*

3 The Congested Part-Wise Aggregation Problem

This section is concerned with a *congested* generalization of the standard part-wise aggregation problem ([Definition 2.1](#)), formally introduced below.

Definition 3.1 (Congested Part-Wise Aggregation Problem). Consider an n -node graph G with a collection of k subsets of nodes $P_1, \dots, P_k \subseteq V(G)$ called *parts* such that each induced subgraph $G[P_i]$ is *connected* and each node $v \in V(G)$ is contained in at most $\rho \in \mathbb{Z}_{\geq 1}$ many parts, i.e., $\forall v \in V(G) \quad |\{i : P_i \ni v\}| \leq \rho$. In the ρ -congested part-wise aggregation problem, each node v is given the following as input: for each part $P_i \ni v$ node v knows the part-ID i and an $O(\log n)$ -bit part-specific value $\mathbf{x}_i(v)$. The goal is that, for each part P_i , all nodes in P_i learn the part-wise aggregate $\bigoplus_{w \in P_i} \mathbf{x}_i(w)$, where \bigoplus is an arbitrary pre-defined *aggregation function*.

This congested generalization of the standard part-wise aggregation problem that we study in this section turns out to be a central ingredient in our refined Laplacian solver; this is further explained in [Section 4](#). The remainder of this section is organized as follows. In [Section 3.1](#) we establish near-optimal algorithms for solving congested part-wise aggregations in CONGEST, which is also the main focus of this section. We conclude by pointing out the construction for NCC in [Section 3.2](#).

3.1 Solving Congested Instances in the CONGEST Model

The first natural strategy for solving the ρ -congested part-wise aggregation problem of [Definition 3.1](#) is through a reduction to $\text{poly}(\rho)$ 1-congested instances. However, this approach immediately fails even if we allow $\rho = 2$. Indeed, there exist congested part-wise aggregation instances for which every two (distinct) parts share a common node, even when $\rho = 2$, leading to the following observation.

Observation 3.2. *For an infinite family of values \bar{n} , there exists an \bar{n} -node planar graph \bar{G} and a 2-congested part-wise aggregation instance \mathcal{I} with $k = \Theta(\sqrt{\bar{n}})$ parts such that reducing \mathcal{I} to the union of k' 1-congested part-wise aggregation instances on \bar{G} requires $k' = \Omega(\sqrt{\bar{n}})$.*

Such a pattern is illustrated in [Figure 1](#). As a result, directly employing a 1-congested part-wise aggregation oracle is of little use since it would introduce an overhead depending on the number of parts. In light of this, we develop a more refined approach that leverages what we refer to as the *layered graph*. This concept is introduced in [Section 3.1.1](#), where we show that the congested part-wise aggregation problem can be reduced to the 1-congested part-wise aggregation problem in the layered graph. Then, we give an algorithm for the ρ -congested part-wise aggregation problem in treewidth-bounded graphs through a simple approach in [Section 3.1.2](#), yielding an $\tilde{O}(\rho^2 \text{tw}(\bar{G})D)$ -round algorithm. Finally, we show that the shortcut quality SQ of the ρ -layered graph does not increase (modulo polylogarithmic factors) as compared to the original graph ([Theorem 3.11](#)). This implies a solution for ρ -congested part-wise aggregations in general graphs with a runtime with the optimal, linear, dependence on ρ , albeit at the cost of a more involved argument ([Section 3.1.3](#), specifically [Corollary 3.12](#)).

3.1.1 The Layered Graph

Here we introduce the *layered graph* \hat{G}_ρ associated with the underlying graph \bar{G} . Then, we reduce the problem of ρ -congested part-wise aggregation on \bar{G} to a 1-congested instance on $\hat{G}_{O(\rho)}$.

The Layered Graph Consider an underlying network \bar{G} and some $\rho \in \mathbb{Z}_{\geq 1}$, corresponding to the congestion parameter in [Definition 3.1](#). The *layered graph* \hat{G}_ρ is constructed in the following way. First, we let \hat{G}_ρ be a disjoint union of ρ copies of \bar{G} (called *layers*), namely $\bar{G}_1, \bar{G}_2, \dots, \bar{G}_\rho$. Each node $v \in V(\bar{G})$ is associated with its copies $v_1, v_2, \dots, v_\rho \in V(\hat{G}_\rho)$. We also add an edge between each two copies that originate from the same node (i.e., we add a clique to \hat{G}_ρ on the set of copies associated with the same node $v \in V(\bar{G})$); this construction is illustrated in [Figure 2](#). The layered graph induces a natural *projection* operation $\pi : V(\hat{G}_\rho) \rightarrow V(\bar{G})$ which maps a copy v_i to its original node $v = \pi(v_i)$. Furthermore, we often talk about simulating \hat{G}_ρ in \bar{G} , by which we mean that each node v simulates—learns all the inputs and can generate all outputs—for its copies v_1, \dots, v_ρ . Throughout this paper, we will assume that $\rho = \text{poly}(\bar{n})$ so that any $O(\log n)$ -bit message on \hat{G}_ρ can be sent within $O(1)$ rounds in \bar{G} ; this also keeps the \tilde{O} -notation well-defined.

The main goal of this section is to establish that the ρ -congested part-wise aggregation problem on \bar{G} can be reduced to a 1-congested instance on $\hat{G}_{O(\rho)}$, as formalized below.

Lemma 3.3 (Unrestricted Congested Part-Wise Aggregation). *Let \bar{G} be an \bar{n} -node graph and let $\mathbb{Z}_{\geq 1} \ni \rho \leq \text{poly}(\bar{n})$. Suppose that any (1-congested) part-wise aggregation on $\hat{G}_{O(\rho)}$ can be solved with a τ -round CONGEST algorithm on $\hat{G}_{O(\rho)}$. Then, there exists an $\tilde{O}(\rho \cdot \tau)$ -round CONGEST algorithm on \bar{G} that solves any ρ -congested part-wise aggregation instance on \bar{G} .*

The remainder of this section is dedicated to the proof of this result. We first point out that any CONGEST algorithm on \hat{G}_ρ can be simulated with only a ρ multiplicative overhead in the round complexity (see [Appendix B.2](#)).

Lemma 3.4 (Simulating \hat{G}_ρ in \bar{G}). *For any \bar{G} and any $\mathbb{Z}_{\geq 1} \ni \rho \leq \text{poly}(\bar{n})$, we can simulate any τ -round CONGEST algorithm on \hat{G}_ρ with a $(\rho \cdot \tau)$ -round CONGEST algorithm on \bar{G} .*




Figure 1: A 2-congested part-wise aggregation problem on a 6×6 grid (the instance immediately extends to a $\sqrt{n} \times \sqrt{n}$ topology). Different colors highlight different parts of the instance.




Figure 2: An example of a transformation from \bar{G} to the layered graph \hat{G}_ρ with $\rho = 3$. We have highlighted with different colors different layers of the graph.

Furthermore, we will use a folklore result showing how to color a (multi)graph of maximum degree Δ in $O(\Delta)$ colors in $O(\log n)$ rounds of CONGEST. By multigraph here we simply mean that there can be multiple parallel edges between the same pair of nodes, and every such edge can carry an independent message per round. To keep the paper self-contained we provide a short sketch of the proof in [Appendix B.2](#).

Fact 3.5 (Folklore, [Joh99]). *Given a (multi)graph G with n nodes and maximum degree $\Delta \leq \text{poly}(n)$, there exists a randomized CONGEST algorithm that colors the edges of G with $O(\Delta)$ colors and completes in $O(\log n)$ rounds, with high probability. The coloring is proper, i.e., two edges that share an endpoint are assigned a different color.*

Now we are ready to prove a version of our main reduction ([Lemma 3.3](#)), but with the slightly twist that we restrict each part of the ρ -congested part-wise aggregation problem to be a simple path. This restriction will be removed later.

Lemma 3.6 (Path-Restricted Congested Part-Wise Aggregation). *Let \bar{G} be a \bar{n} -node graph and let $\mathbb{Z}_{\geq 1} \ni \rho \leq \text{poly}(\bar{n})$. Suppose that there exists a τ -round CONGEST algorithm solving the (1-congested) part-wise aggregation on $\hat{G}_{O(\rho)}$. Then, there exists an $\tilde{O}(\rho \cdot \tau)$ -round CONGEST algorithm on \bar{G} that solves any ρ -congested part-wise aggregation instance on \bar{G} when each part is restricted to be a simple path⁶ (nodes are not repeated in simple paths).*

Proof. Let $\mathcal{P} = \{P_1, P_2, \dots, P_k\}$ be subsets of nodes in \bar{G} comprising the parts of some ρ -congested part-wise aggregation on \bar{G} . We will construct paths $\mathcal{P}' = \{P'_1, P'_2, \dots, P'_k\}$ in $\hat{G}_{O(\rho)}$ in a way that solving a part-wise aggregation on \mathcal{P}' corresponds to solving a ρ -congested part-wise aggregation on \mathcal{P} .

Let E_i be the set of edges of \bar{G} comprising the simple path traversing all the nodes in P_i , and consider the graph $G' := (V(\bar{G}), \uplus_{i=1}^k E_i)$. First, we observe that the degree of any node in $v \in V(G') = V(\bar{G})$ is at most 2ρ since at most ρ many parts contain v and each part contributes at most 2 to the degree (since P_i is a simple path). Furthermore, we can simulate any ψ -round CONGEST algorithm on G' with a $(\psi \cdot \rho)$ -round CONGEST algorithm on \bar{G} as each edge $e \in E(\bar{G})$ appears at most ρ times in $E(G')$ due to the part-wise aggregation instance being at most ρ -congested. Therefore, using [Fact 3.5](#) we can distributedly color the edges of G' into at most $O(\rho)$ colors in $O(\log n)$ CONGEST rounds on G' , which translates to $\tilde{O}(\rho)$ CONGEST rounds on \bar{G} . Suppose that the algorithm assigns a color $\mathbf{c}(e) \in \{1, \dots, O(\rho)\}$ to each edge $e \in \uplus_i E_i$.

We now construct $P'_i \subseteq \hat{G}_{O(\rho)}$ as follows: consider each edge $\{u, v\} \in E_i$ and add both $u_{\mathbf{c}(\{u,v\})}, v_{\mathbf{c}(\{u,v\})} \in V(\hat{G}_{O(\rho)})$ to P'_i (i.e., the $\mathbf{c}(\{u, v\})$ -th copy of both u and v). By construction, P'_i induces a connected subgraph and the projection P'_i to \bar{G} is exactly P_i . Next, we invoke the (1-congested) part-wise aggregation τ -round algorithm for $\{P'_1, \dots, P'_k\}$ on $\hat{G}_{O(\rho)}$, which can be converted to an $\tilde{O}(\tau \cdot \rho)$ -round algorithm on \bar{G} ([Lemma 3.4](#)). Thus, we obtain an $\tilde{O}(\tau \cdot \rho)$ -round CONGEST algorithm on \bar{G} which solves any path-restricted ρ -congested part-wise aggregation problem. \square

Finally, our reduction in [Lemma 3.3](#) follows by reformulating [[HWZ21](#), Lemma 7.2], as we argue in [Appendix B.2](#).

⁶I.e., there exists a simple path traversing all the nodes of the part, and each node knows the corresponding incident edges of that path.

3.1.2 Treewidth-Bounded Graphs

Here we leverage the reduction we established in [Lemma 3.3](#) to obtain a simple algorithm for solving the congested part-wise aggregation problem in treewidth-bounded graphs. The crucial observation is that the treewidth of the layered graph can only grow by a factor of ρ compared to the treewidth of the underlying graph, as we show in [Lemma 3.8](#).

Claim 3.7. $D(\widehat{G}_\rho) \leq D(\overline{G}) + 1$.

Lemma 3.8. *If the treewidth of \overline{G} is $\text{tw}(\overline{G})$, then $\text{tw}(\widehat{G}_\rho) \leq \rho \text{tw}(\overline{G}) + \rho - 1$.*

Proof. Consider a tree decomposition (in the sense of [Definition 2.8](#)) of \overline{G} into tree-nodes $\{X_j\}_{j=1}^k$ such that the width of the decomposition satisfies $w = \text{tw}(\overline{G})$. We will show that there exists a tree decomposition on the graph \widehat{G}_ρ with width at most $\rho(w + 1) - 1$, which in turn will imply that $\text{tw}(\widehat{G}_\rho) \leq \rho(w + 1) - 1 = \rho(\text{tw}(\overline{G}) + 1) - 1$. Indeed, consider the following sets:

$$\widehat{X}_j := \{u_i : u \in X_j, i \in [\rho]\},$$

for all $j \in [k]$. In words, each node $V(\overline{G}) \ni u \in X_j$ is replaced by all of its copies u_i in \widehat{X}_j . Observe that, by construction, $|\widehat{X}_j| = \rho|X_j|$. Thus, it suffices to show that the collection of sets $\{\widehat{X}_j\}_{j=1}^k$ forms a legitimate tree decomposition. First, since $V(\overline{G}) \subseteq \bigcup_j X_j$, it follows that $V(\widehat{G}_\rho) \subseteq \bigcup \widehat{X}_j$. Moreover, consider any two sets $\widehat{X}_j, \widehat{X}_\ell$, both containing a node $u_i \in V(\widehat{G}_\rho)$ for some $i \in [\rho]$. Then, we know that all the tree-nodes in the (unique) path between X_j and X_ℓ based on the original tree decomposition include u since X_j and X_ℓ both include u and $\{X_j\}$ is a tree decomposition of \overline{G} . In turn, this implies that all the tree-nodes in the path between \widehat{X}_j and \widehat{X}_ℓ also contain u_i . Thus, the tree-nodes containing u_i form a connected subtree. Finally, we know that for every edge $\{u, v\} \in E(\overline{G})$ there exists a subset X_j such that $u, v \in X_j$. Hence, we can infer that for every edge in $E(\widehat{G}_\rho)$ there is a tree-node \widehat{X}_j which includes both incident endpoints. As a result, we have constructed a tree decomposition in \widehat{G}_ρ with width $\max_{j \in [k]} |\widehat{X}_j| - 1 \leq \rho(w + 1) - 1$. \square

Corollary 3.9. *Let \overline{G} be an \bar{n} -node communication network of diameter at most D and treewidth $\text{tw}(\overline{G})$. Then, we can solve with high probability any ρ -congested part-wise aggregation problem in \overline{G} within $\tilde{O}(\rho^2 \cdot \text{tw}(\overline{G}) \cdot D)$ rounds of CONGEST.*

Proof. First, we know from [Lemma 3.8](#) that $\text{tw}(\widehat{G}_\rho) = O(\rho \text{tw}(\overline{G}))$, in turn implying that the minor density of \widehat{G}_ρ can be bounded as $\delta(\widehat{G}_\rho) \leq \text{tw}(\widehat{G}_\rho) = O(\rho \text{tw}(\overline{G}))$ ([Fact 2.9](#)). Thus, [Theorem 2.7](#) implies that \widehat{G}_ρ admits shortcuts of quality $\tilde{O}(\rho \text{tw}(\overline{G}) D(\overline{G}))$, which can be additionally constructed in $\tilde{O}(\rho \text{tw}(\overline{G}) D(\overline{G}))$ rounds of communication on \widehat{G}_ρ . Finally, we have shown in [Lemma 3.3](#) that this is sufficient to solve any ρ -congested part-wise aggregation problem on \overline{G} in $\tilde{O}(\rho^2 \cdot \text{tw}(\overline{G}) \cdot D(\overline{G}))$ rounds of CONGEST, concluding the proof. \square

Minor Density in the Layered Graph In light of [Lemma 3.8](#), a natural question is whether an analogous bound holds with respect to the minor density of the underlying graph; i.e., whether $\delta(\widehat{G}_\rho) = \text{poly}(\rho)\delta(\overline{G})$. Such a result would be strictly stronger as it would apply to the broader class of graphs with bounded minor density, and would essentially lift all the results in [\[GH20\]](#), such as [Theorem 2.7](#), to the node-congestion setting in a black-box manner. Unfortunately, this is not possible.

Indeed, consider a $\sqrt{n} \times \sqrt{n}$ grid G —where \sqrt{n} is assumed to be an integer—such that every node in the graph is 2-congested. Then, it is clear that $\delta(G) = \tilde{O}(1)$ (since planar graphs have excluded minors). On the other hand, we claim that $\delta(\hat{G}_\rho) = \Omega(\sqrt{n})$. To see this, denote by (i, j) the node positioned in the i -th row and j -th column with respect to the original graph, and by (i', j') the node positioned in the i -th row and j -th column of the "duplicate" layer, for $i, j \in [\sqrt{n}]$. Moreover, let $C_j = \{(i, j) : i \in [\sqrt{n}]\}$ be the nodes comprising the j -th column of the original graph and $R_i = \{(i', j') : j' \in [\sqrt{n}]\}$ be the nodes comprising the i -th row of the duplicate layer. Then, it follows that the minor graph induced by the connected components $R_1, \dots, R_{\sqrt{n}}, C_1, \dots, C_{\sqrt{n}}$ contains the complete bipartite graph $K_{\sqrt{n}, \sqrt{n}}$ as a subgraph (Figure 3). As a result, this implies that the minor density of \hat{G}_ρ is $\Omega(\sqrt{n})$.

Observation 3.10. *There exists an n -node graph G with minor density $\delta(G) = \tilde{O}(1)$, but its 2-layered version \hat{G}_2 has minor density $\delta(\hat{G}_2) = \Omega(\sqrt{n})$.*




Figure 3: The layered graph \hat{G}_ρ corresponding to a 3×3 grid with every node having congestion $\rho = 2$ (leftmost image), and a minor of \hat{G}_ρ induced by the connected components $\{C_1, C_2, C_3, R_1, R_2, R_3\}$ (rightmost image).

3.1.3 General Graphs

We conclude with our main result of Section 3.1: a near-optimal distributed algorithm for solving the ρ -congested part-wise aggregation problem in general graphs. In light of our reduction in Lemma 3.3, the technical crux is to control the degradation in the shortcut quality incurred by the transformation into the layered graph. Surprisingly, we show that the shortcut quality of \hat{G}_ρ does not increase by more than a polylogarithmic factor even when the number of layers is polynomial:

Theorem 3.11. *For any \bar{n} -node graph \bar{G} and any $\mathbb{Z}_{\geq 1} \ni \rho \leq \text{poly}(\bar{n})$, we have that $\text{SQ}(\hat{G}_\rho) = \tilde{O}(\text{SQ}(\bar{G}))$.*

This theorem improves over our previous result for treewidth-bounded graphs (Lemma 3.8) since the latter guarantee inevitably induces a linear factor of ρ in the shortcut quality of \hat{G}_ρ . While this will not affect the asymptotic performance of the Laplacian solver, this improvement might prove to be important for future applications. Assuming that we have shown Theorem 3.11, we can then utilize the efficient shortcut constructions given in Theorem 2.5 to solve ρ -congested part-wise aggregations on any graph.

Corollary 3.12. *There exists a randomized distributed algorithm that, for any \bar{n} -node graph \bar{G} and*

$\rho \in \mathbb{Z}_{\geq 1} \leq \text{poly}(\bar{n})$, solves with high probability any ρ -congested part-wise aggregation instance on \bar{G} with the following guarantees:

- In the CONGEST, the algorithm terminates in at most $\rho \cdot \text{poly}(\text{SQ}(\bar{G})) \cdot \bar{n}^{o(1)}$ rounds.
- In the CONGEST model on graphs with minor density δ , it requires $\tilde{O}(\rho \cdot \delta \cdot D)$ rounds.
- In the Supported-CONGEST, the algorithm terminates in $\tilde{O}(\rho \cdot \text{SQ}(\bar{G}))$ rounds.

Proof. We construct (virtually) the graph $\hat{G}_{O(\rho)}$. By [Theorem 3.11](#), we know that $\text{SQ}(\hat{G}_{O(\rho)}) = \tilde{O}(\text{SQ}(\bar{G}))$. Thus, we can construct shortcuts on $\hat{G}_{O(\rho)}$ in $\text{poly}(\text{SQ}(\bar{G})) \cdot \bar{n}^{o(1)}$, $\tilde{O}(\delta D)$, and $\tilde{O}(\text{SQ}(\bar{G}))$ rounds in general CONGEST ([Theorem 2.5](#)), Supported-CONGEST ([Theorem 2.5](#)), and CONGEST with minor density δ ([Theorem 2.7](#)), respectively. Therefore, we can solve 1-congested part-wise aggregation instances using those shortcuts in the required times using [Proposition 2.3](#). Since solving a 1-congested part-wise aggregation on $\hat{G}_{O(\rho)}$ suffices to solve ρ -congested part-wise aggregations on \bar{G} with only an $\tilde{O}(\rho)$ slowdown ([Lemma 3.3](#)), the proof is complete. \square

The rest of this subsection is dedicated to the proof of [Theorem 3.11](#). To argue about the shortcut quality of the layered graph, we need to develop several generalized notions of node connectivity. Pair node and any-to-any connectivity are essentially the multi- and single-commodity versions of node connectivity, respectively.

Pair Node Connectivity Given a (multi)set of source-sink pairs $\mathcal{P} = \{(s_i, t_i)\}_{i=1}^k$ in G , we say that \mathcal{P} has pair node connectivity ρ if there exist paths P_1, \dots, P_k , with s_i and t_i being the endpoints of each P_i , such that every node $v \in V(G)$ is contained in at most ρ many paths, i.e., for all v we have $|\{i : v \in P_i\}| \leq \rho$. If \mathcal{P} has pair node connectivity 1 we say that they are pair *node-disjointly connectable*.

Any-to-Any Node Connectivity Suppose that we are given multisets of k sources $S = \{s_1, \dots, s_k\}$ and k sinks $T = \{t_1, \dots, t_k\}$. We say that (S, T) have any-to-any node connectivity ρ if there is a permutation $\pi : \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ such that the pairs $\{(s_i, t_{\pi(i)})\}_{i=1}^k$ have pair node connectivity ρ . If (S, T) have any-to-any node connectivity 1 we say they are any-to-any *node-disjointly connectable*.

The following decomposition lemma states that two sets with any-to-any node connectivity ρ can be decomposed into $\tilde{O}(\rho)$ many pairs of subsets that are any-to-any node-disjointly connectable.

Lemma 3.13. *Given a graph G , suppose we are given any two multisets of nodes $S \subseteq V(G)$ and $T \subseteq V(G)$ of size $k := |S| = |T|$ that have any-to-any node connectivity ρ . Then, we can partition $S = S_1 \uplus S_2 \uplus \dots \uplus S_{O(\rho \log k)}$ and $T = T_1 \uplus T_2 \uplus \dots \uplus T_{O(\rho \log k)}$ such that $|S_i| = |T_i|$ and (S_i, T_i) are any-to-any node-disjointly connectable.*

Proof. Suppose that each edge in G has infinite capacity while each node in G has unit capacity. Then, let us connect a super-source s to each node $x \in S$ with a unit-capacity edge, and a super-sink t to each node $x \in T$ with a unit capacity edge. By assumption, we know that there exists a flow f over $E(G)$ which sends k units of flow from s to t with edge congestion 1 and node congestion at most ρ . Therefore, the flow f/ρ sending k/ρ units of flow from s to t is a feasible solution of the

maximum flow linear program with node constraints (i.e., it satisfies both edge and node capacity constraints). Since that linear program is integral (i.e., has an integrality gap of 1), there exists an integral flow f' which sends at least k/ρ units of flow and satisfies both node and edge capacity restrictions. In other words, there exist at least k/ρ node disjoint paths (with the exception of the endpoints) between s and t . Let $S_1 \subseteq S$ ($T_1 \subseteq T$) be the set of nodes on these paths immediately following the super-source (just before the super-sink, respectively). Clearly, by construction, (S_1, T_1) are any-to-any node-disjointly connectable. Finally, we define $S' := S \setminus S_1, T' := T \setminus T_1$ and proceed iteratively as above (producing S_2, T_2 instead of S_1, T_1). In each step, the size of S' and T' decreases by at least a multiplicative factor of $1 - 1/\rho$. Hence, $O(\rho \log k)$ steps suffice so that $S' = T' = \emptyset$. \square

Next, we introduce two communication tasks that will be useful for characterizing the shortcut quality.

Multiple-Unicast Problem Suppose that we are given k source-sink pairs $\mathcal{P} = \{(s_i, t_i)\}_{i=1}^k$. The goal is to find the smallest possible *completion time* τ such that there are k paths P_1, \dots, P_k for which (1) the endpoints of each P_i are exactly s_i and t_i ; (2) the dilation is τ , i.e., each path P_i has at most τ hops; and (3) the congestion is τ , i.e., each edge $e \in E(G)$ is contained in at most τ many paths.

Any-to-Any-Cast Problem Suppose we are given k sources $S = \{s_1, \dots, s_k\}$ and k sinks $T = \{t_1, \dots, t_k\}$. The goal is to find the smallest *completion time* τ such that there exists a permutation $\pi : \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ for which the multiple-unicast problem on $\{(s_i, t_{\pi(i)})\}_{i=1}^k$ has a completion time of at most τ .

Finally, we now recall (a reinterpretation of) a result characterizing shortcut quality from [HWZ20; HWZ21]. Shortcut quality was originally defined as the smallest completion-time of the worst-case generalized (with respect to parts) multiple-unicast (i.e., multi-commodity) problem over an a *pair* node-disjointly connectable instance (Definition 2.4). Using recent network coding gap results we can equivalently express shortcut quality as the smallest completion-time of the worst-case any-to-any-cast (i.e., single-commodity) problem over sources and sinks that are *any-to-any* node-disjointly connectable. The formal statement follows.

Theorem 3.14 ([HWZ20; HWZ21]). *Consider any graph G and let τ be the worst-case completion time of any-to-any-cast problems taken over all any-to-any node-disjointly connectable sets $(S \subseteq V(G), T \subseteq V(G))$. Then, $\tau = \tilde{\Theta}(\text{SQ}(G))$.*

Proof. It was proven in [HWZ21, Lemma 2.8 in the Full Version] that $\text{SQ}(G)$ is, up to $\tilde{\Theta}(1)$ factors, equal to the completion time C of some multiple-unicast instance with respect to some source-sink pairs $\mathcal{P} := \{(s_i, t_i)\}_{i=1}^k$ that are pair node-disjointly connectable. We note that, since sources and sinks are disjoint, it follows that $k = \text{poly}(n)$ and $O(\log n) = O(\log k)$. Furthermore, [HWZ20] proved that there exists a sub-instance $\mathcal{P}' = \{(s'_i, t'_i)\}_{i=1}^{k'} \subseteq \mathcal{P}$ such that $\text{SQ}(G)$ is (up to $\tilde{\Theta}(1)$ factors) equal to the completion time τ of the any-to-any-cast problem with respect to $(\{s'_i\}_{i=1}^{k'}, \{t'_i\}_{i=1}^{k'})$. One side of the claim is clear: for any sub-instance $\mathcal{P}' \subseteq \mathcal{P}$ we have that $\tau \leq C$. The other direction is harder and we sketch its proof here using the terminology in [HWZ20]. By definition and strong duality, $\text{Cut}_{\mathcal{P}}(2C) = \text{ConcurrentFlow}_{\mathcal{P}}(2C) \leq 1$. Furthermore, $\text{Cut}_{\mathcal{P}}(C/10) = \text{Cut}_{\mathcal{P}}(2C)/20 \leq 1/10$. Hence, by [HWZ21, Lemma 2.6] there is a sub-instance $\mathcal{P}' \subseteq \mathcal{P}$ with a

moving cut of distance $\tau := \tilde{\Omega}(C)$ and capacity less than $|\mathcal{P}'|$. Therefore, this proves that the completion time of any-to-any-cast problem on \mathcal{P}' is at least τ . With this in mind, we have that $\tilde{\Omega}(\text{SQ}(G)) = \tilde{\Omega}(C) = \tau \leq C = \tilde{\Theta}(\text{SQ}(G))$.

Finally, since $\mathcal{P} = \{(s_i, t_i)\}_{i=1}^k$ was pair node-disjointly connectable, it follows from definition that the sub-instance $(\{s'_i\}_{i=1}^{k'}, \{t'_i\}_{i=1}^{k'})$ is any-to-any node-disjointly connectable. Therefore, $(\{s'_i\}_{i=1}^{k'}, \{t'_i\}_{i=1}^{k'})$ satisfies the constraints of this result and has completion-time $\tau = \tilde{\Theta}(\text{SQ}(G))$, as required. It is also clear that, by shortcut quality, any any-to-any node-disjointly connectable instance has completion time at most $\text{SQ}(G)$ using the node-disjoint paths that witness the any-to-any node-disjointness as parts of the shortcut, making $(\{s'_i\}_{i=1}^{k'}, \{t'_i\}_{i=1}^{k'})$ the worst-case such instance (modulo polylogarithmic factors). \square

We now combine all of the previous ingredients to prove the main result of this section.

Proof of Theorem 3.11. Let $S \subseteq V(\widehat{G}_\rho)$ and $T \subseteq V(\widehat{G}_\rho)$ be any-to-any node-disjointly connectable sets such that the completion time of any-to-any-cast between S and T is $\tilde{\Theta}(\text{SQ}(\widehat{G}_\rho))$ (Theorem 3.14). Let $k := |S| = |T|$, and suppose that $S' := \biguplus_{s \in S} \{\pi(s)\} \subseteq V(\overline{G})$ and $T' := \biguplus_{t \in T} \{\pi(t)\} \subseteq V(\overline{G})$ are the multisets induced by projecting S and T to \overline{G} , respectively. By construction of \widehat{G}_ρ , S' and T' have any-to-any node connectivity ρ ; to see this, consider the witness paths disjointly connecting them in \widehat{G}_ρ and project them to \overline{G} . Therefore, we can partition $S' = S'_1 \uplus \dots \uplus S'_{O(\rho \log k)}$ and $T' = T'_1 \uplus \dots \uplus T'_{O(\rho \log k)}$ such that $|S'_i| = |T'_i|$ and (S'_i, T'_i) are any-to-any node-disjointly connectable in \overline{G} (Lemma 3.13).

By definition of shortcut quality, for each $i \in \{1, \dots, O(\rho \log k)\}$ there exists a set of paths $(P_j^i)_{j=1}^{|S'_i|}$ in \overline{G} between S'_i and T'_i of quality (i.e., both congestion and dilation) at most $\text{SQ}(\overline{G})$. Then, we inject the first $O(\log k)$ collections of paths $(P_j^1)_j, (P_j^2)_j, \dots, (P_j^{O(\log k)})_j$ to the first layer \overline{G}_1 of \widehat{G}_ρ ; the second $O(\log k)$ collections to the second layer \overline{G}^2 , and so on, until we finally inject the last $O(\log k)$ collections to the last layer \overline{G}_ρ . Note that only the paths on the same layer interact, so both the congestion and dilation after injecting all paths into \widehat{G}_ρ is $O(\text{SQ}(\overline{G}) \log k)$. Hence, the same applies for the shortcut quality. Finally, to solve the any-to-any-cast problem on S and T one might need to add an between-layer edge at the beginning and at the end since each injected path is restricted to some adversarially chosen layer. However, this only increases the congestion and dilation by $O(1)$. Hence, the completion time of any-to-any-cast between S and T is $\tilde{O}(\text{SQ}(\overline{G}))$, implying that $\text{SQ}(\widehat{G}_\rho) = \tilde{O}(\text{SQ}(\overline{G}))$. \square

3.2 The NCC Model

We next turn our attention to the NCC model. We observe that the ρ -congested part-wise aggregation problem admits a solution in $\text{poly}(\rho, \log \bar{n})$ rounds of NCC. This is established after appropriately translating the communication primitives established for NCC in [Aug+19]; the details are provided in Appendix C.

Lemma 3.15. *Let \overline{G} be an \bar{n} -node communication network. Then, we can solve with high probability any ρ -congested part-wise aggregation problem on \overline{G} after $O(\rho + \log \bar{n})$ rounds of NCC.*

4 Almost Universally Optimal Laplacian Solvers

In this section we relate the congested part-wise aggregation problem we studied in the previous section with the Laplacian solver of Forster, Goranci, Liu, Peng, Sun, and Ye [For+20]. To present a unifying analysis for both CONGEST and HYBRID, as well as for future applications and extensions, we analyze the distributed Laplacian solver under the following hypothesis.

Assumption 4.1. *Consider a model of computation which incorporates CONGEST. We assume that we can solve with high probability any ρ -congested part-wise aggregation problem in $Q(\rho) = O(\rho^c Q(1))$ rounds, for some universal constant $c \geq 1$.*

One of our crucial observations is that the performance of the Laplacian solver of Forster, Goranci, Liu, Peng, Sun, and Ye [For+20] can be parameterized in terms of the complexity of the congested part-wise aggregation problem. Indeed, we revisit and refine the main building blocks of their solver in [Appendix A](#), leading to the following result.

Theorem 4.2 (Full Version in [Theorem A.9](#)). *Consider a weighted \bar{n} -node graph \bar{G} for which [Assumption 4.1](#) holds for some $Q(\rho) = O(\rho^c Q)$, where c is a universal constant and $Q = Q(\bar{G})$ is some parameter. Then, we can solve any Laplacian system after $\bar{n}^{o(1)} Q \log(1/\varepsilon)$ rounds.*

Combining this theorem with [Corollary 3.12](#) and [Lemma 3.15](#) yields the following immediate consequences.

Theorem 1.2. *Consider any n -node graph G with shortcut quality $\text{SQ}(G)$ and hop-diameter D . There exists a distributed Laplacian solver with error $\varepsilon > 0$ with the following guarantees:*

- *In the Supported-CONGEST model, it requires $n^{o(1)} \text{SQ}(G) \log(1/\varepsilon)$ rounds.*
- *In the CONGEST model, it requires $n^{o(1)} \text{poly}(\text{SQ}(G)) \log(1/\varepsilon)$ rounds.*
- *In the CONGEST model on graphs with minor density δ , it requires $n^{o(1)} \delta D \log(1/\varepsilon)$ rounds.*

Theorem 1.3. *Consider any n -node graph. There exists a distributed Laplacian solver in the HYBRID model with round complexity $n^{o(1)} \log(1/\varepsilon)$, where $\varepsilon > 0$ is the error of the solver.*

Lower Bound in Supported-CONGEST Finally, we complement our positive results with a almost-matching lower bound on any graph \bar{G} , applicable even under the Supported-CONGEST model, thereby establishing universal optimality up to an $\bar{n}^{o(1)}$ factor. Our reduction leverages the refined hardness result established in [HWZ21] for the *spanning connected subgraph* problem [Das+11]. In this problem a subgraph \bar{H} of \bar{G} is specified with nodes knowing all of the incident edges belonging to \bar{H} . The goal is to let every node learn whether \bar{H} is connected and spans the entire network.

Theorem 4.3 ([HWZ21]). *Let \mathcal{A} be any algorithm which is always correct with probability⁷ at least $\frac{2}{3}$ for the spanning connected subgraph problem, and $T(\bar{G}) = \max_{\mathcal{I}} T_{\mathcal{A}}(\mathcal{I}; \bar{G})$ be the worst-case round-complexity of \mathcal{A} under \bar{G} . Then,*

$$T(\bar{G}) = \tilde{\Omega}(\text{SHORTCUTQUALITY}(\bar{G})).$$

⁷Note that Haeupler, Wajc, and Zuzic [HWZ21] only proved this for always-correct algorithms with probability 1, but the extension we claim here follows readily from their argument.

In this context, we show that a Laplacian solver can be leveraged to solve the spanning connected subgraph problem, leading to the following lower bound.

Proposition 1.1. *Consider a graph \overline{G} with shortcut quality $\text{SQ}(\overline{G})$. Then, solving a Laplacian system on \overline{G} with $\varepsilon \leq \frac{1}{2}$ requires $\tilde{\Omega}(\text{SQ}(\overline{G}))$ rounds in both CONGEST and Supported-CONGEST models.*

This substantially strengthens the *existential lower bound* in [For+20], and deviates from their argument which is based on a reduction from the $s - t$ connectivity problem. The proof is deferred to [Appendix B.8](#).

5 Conclusions

We established almost universally optimal Laplacian solvers for both the (Supported-)CONGEST and the HYBRID model. One of our main technical contributions was to introduce and study a congested generalization of the standard part-wise aggregation problem, which we believe may find further applications beyond the Laplacian paradigm in the future. For example, one candidate problem would be to refine the distributed algorithm for max-flow due to Ghaffari, Karrenbauer, Kuhn, Lenzen, and Patt-Shamir [Gha+15]. We also hope that our accelerated Laplacian solvers will be used as a basic primitive for obtaining improved distributed algorithms for other fundamental optimization problems as well. Indeed, Forster, Goranci, Liu, Peng, Sun, and Ye [For+20] showed that the Laplacian paradigm can offer *sublinear* and *exact* distributed algorithms for problems such as max-flow, an objective which previously appeared elusive.

References

- [Ach03] Dimitris Achlioptas. “Database-friendly random projections: Johnson-Lindenstrauss with binary coins”. In: *J. Comput. Syst. Sci.* 66.4 (2003), pp. 671–687.
- [AG21] Ioannis Anagnostides and Themis Gouleakis. “Deterministic Distributed Algorithms and Lower Bounds in the Hybrid Model”. In: *35th International Symposium on Distributed Computing, DISC 2021*. Ed. by Seth Gilbert. Vol. 209. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 5:1–5:19.
- [Alo+95] Noga Alon, Richard M. Karp, David Peleg, and Douglas West. “A Graph-Theoretic Game and Its Application to the k -Server Problem”. In: *SIAM J. Comput.* 24.1 (Feb. 1995), pp. 78–100.
- [AMV20] Kyriakos Axiotis, Aleksander Madry, and Adrian Vladu. “Circulation Control for Faster Minimum Cost Flow in Unit-Capacity Graphs”. In: *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*. IEEE, 2020, pp. 93–104.
- [AMV21] Kyriakos Axiotis, Aleksander Madry, and Adrian Vladu. “Faster Sparse Minimum Cost Flow by Electrical Flow Localization”. In: *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021*. IEEE, 2021, pp. 528–539.
- [Aug+19] John Augustine, Mohsen Ghaffari, Robert Gmyr, Kristian Hinnenthal, Christian Scheideler, Fabian Kuhn, and Jason Li. “Distributed Computation in Node-Capacitated Networks”. In: *The 31st ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA 2019*. Ed. by Christian Scheideler and Petra Berenbrink. ACM, 2019, pp. 69–79.

- [Aug+20] John Augustine, Kristian Hinnenthal, Fabian Kuhn, Christian Scheideler, and Philipp Schneider. “Shortest Paths in a Hybrid Network Model”. In: *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020*. Ed. by Shuchi Chawla. SIAM, 2020, pp. 1280–1299.
- [Ble+14] Guy E. Blelloch, Anupam Gupta, Ioannis Koutis, Gary L. Miller, Richard Peng, and Kanat Tangwongsan. “Nearly-Linear Work Parallel SDD Solvers, Low-Diameter Decomposition, and Low-Stretch Subgraphs”. In: *Theory Comput. Syst.* 55.3 (2014), pp. 521–554.
- [Bra+20] Jan van den Brand, Yin Tat Lee, Danupon Nanongkai, Richard Peng, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. “Bipartite Matching in Nearly-linear Time on Moderately Dense Graphs”. In: *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*. IEEE, 2020, pp. 919–930.
- [BS07] Surender Baswana and Sandeep Sen. “A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs”. In: *Random Struct. Algorithms* 30.4 (2007), pp. 532–563.
- [CGC16] Tao Chen, Xiaofeng Gao, and Guihai Chen. “The features, hardware, and architectures of data center networks: A survey”. In: *Journal of Parallel and Distributed Computing* 96 (2016), pp. 45–74.
- [CLP21a] Keren Censor-Hillel, Dean Leitersdorf, and Volodymyr Polosukhin. “Distance Computations in the Hybrid Network Model via Oracle Simulations”. In: *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021*. Vol. 187. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 21:1–21:19.
- [CLP21b] Keren Censor-Hillel, Dean Leitersdorf, and Volodymyr Polosukhin. “On Sparsity Awareness in Distributed Computations”. In: *SPAA ’21: 33rd ACM Symposium on Parallelism in Algorithms and Architectures*. Ed. by Kunal Agrawal and Yossi Azar. ACM, 2021, pp. 151–161.
- [CM21] Shiri Chechik and Doron Mukhtar. “Single-source shortest paths in the CONGEST model with improved bounds”. In: *Distributed Computing* (2021), pp. 1–18.
- [Coh+17] Michael B. Cohen, Aleksander Madry, Piotr Sankowski, and Adrian Vladu. “Negative-Weight Shortest Paths and Unit Capacity Minimum Cost Flow in $\tilde{O}(m^{10/7} \log W)$ Time (Extended Abstract)”. In: *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*. Ed. by Philip N. Klein. SIAM, 2017, pp. 752–771.
- [Coy+22] Sam Coy, Artur Czumaj, Michael Feldmann, Kristian Hinnenthal, Fabian Kuhn, Christian Scheideler, Philipp Schneider, and Martijn Struijs. *Near-Shortest Path Routing in Hybrid Communication Networks*. 2022. arXiv: [2202.08008](https://arxiv.org/abs/2202.08008) [cs.DC].
- [Das+11] 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”. In: *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*. STOC ’11. San Jose, California, USA: Association for Computing Machinery, 2011, pp. 363–372.
- [DKO14] Andrew Drucker, Fabian Kuhn, and Rotem Oshman. “On the power of the congested clique model”. In: *ACM Symposium on Principles of Distributed Computing, PODC ’14*. Ed. by Magnús M. Halldórsson and Shlomi Dolev. ACM, 2014, pp. 367–376.

- [Dur+19] David Durfee, Yu Gao, Gramoz Goranci, and Richard Peng. “Fully dynamic spectral vertex sparsifiers and applications”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019*. ACM, 2019, pp. 914–925.
- [Elk04] Michael Elkin. “Unconditional Lower Bounds on the Time-Approximation Tradeoffs for the Distributed Minimum Spanning Tree Problem”. In: *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*. STOC ’04. Chicago, IL, USA: Association for Computing Machinery, 2004, pp. 331–340.
- [FHS20] Michael Feldmann, Kristian Hinnenthal, and Christian Scheideler. *Fast Hybrid Network Algorithms for Shortest Paths in Sparse Graphs*. 2020. arXiv: [2007.01191](https://arxiv.org/abs/2007.01191) [cs.DC].
- [For+20] Sebastian Forster, Gramoz Goranci, Yang P. Liu, Richard Peng, Xiaorui Sun, and Mingquan Ye. “Minor Sparsifiers and the Distributed Laplacian Paradigm”. In: *CoRR* abs/2012.15675 (2020). arXiv: [2012.15675](https://arxiv.org/abs/2012.15675).
- [GH16] Mohsen Ghaffari and Bernhard Haeupler. “Distributed Algorithms for Planar Networks II: Low-Congestion Shortcuts, MST, and Min-Cut”. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*. Ed. by Robert Krauthgamer. SIAM, 2016, pp. 202–219.
- [GH20] Mohsen Ghaffari and Bernhard Haeupler. “Low-Congestion Shortcuts for Graphs Excluding Dense Minors”. In: *CoRR* abs/2008.03091 (2020). arXiv: [2008.03091](https://arxiv.org/abs/2008.03091).
- [Gha+15] Mohsen Ghaffari, Andreas Karrenbauer, Fabian Kuhn, Christoph Lenzen, and Boaz Patt-Shamir. “Near-Optimal Distributed Maximum Flow: Extended Abstract”. In: New York, NY, USA: Association for Computing Machinery, 2015, pp. 81–90.
- [Gha15] Mohsen Ghaffari. “Near-Optimal Scheduling of Distributed Algorithms”. In: *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*. New York, NY, USA: Association for Computing Machinery, 2015, pp. 3–12.
- [Göt+21] Thorsten Götte, Kristian Hinnenthal, Christian Scheideler, and Julian Werthmann. “Time-Optimal Construction of Overlay Networks”. In: *PODC ’21: ACM Symposium on Principles of Distributed Computing*. ACM, 2021, pp. 457–468.
- [GZ22] Mohsen Ghaffari and Goran Zuzic. “Universally-Optimal Distributed Exact Min-Cut”. In: 2022.
- [HIZ16a] Bernhard Haeupler, Taisuke Izumi, and Goran Zuzic. “Low-Congestion Shortcuts without Embedding”. In: *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016*. Ed. by George Giakkoupis. ACM, 2016, pp. 451–460.
- [HIZ16b] Bernhard Haeupler, Taisuke Izumi, and Goran Zuzic. “Near-Optimal Low-Congestion Shortcuts on Bounded Parameter Graphs”. In: *Distributed Computing - 30th International Symposium, DISC 2016*. Ed. by Cyril Gavoille and David Ilcinkas. Vol. 9888. Lecture Notes in Computer Science. Springer, 2016, pp. 158–172.
- [HRG22] Bernhard Haeupler, Harald Räcke, and Mohsen Ghaffari. “Hop-Constrained Expander Decompositions, Oblivious Routing, and Distributed Universal Optimality”. In: 2022.
- [HWZ20] Bernhard Haeupler, David Wajc, and Goran Zuzic. “Network coding gaps for completion times of multiple unicasts”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 494–505.
- [HWZ21] Bernhard Haeupler, David Wajc, and Goran Zuzic. “Universally-optimal distributed algorithms for known topologies”. In: *STOC ’21: 53rd Annual ACM SIGACT Symposium*

- on *Theory of Computing, Virtual Event, Italy, June 21-25, 2021*. Ed. by Samir Khuller and Virginia Vassilevska Williams. ACM, 2021, pp. 1166–1179.
- [Ind06] Piotr Indyk. “Stable distributions, pseudorandom generators, embeddings, and data stream computation”. In: *J. ACM* 53.3 (2006), pp. 307–323.
- [Joh99] Öjvind Johansson. “Simple distributed $\delta + 1$ -coloring of graphs”. In: *Information Processing Letters* 70.5 (1999), pp. 229–232.
- [Kel+13] Jonathan A. Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. “A simple, combinatorial algorithm for solving SDD systems in nearly-linear time”. In: *Symposium on Theory of Computing Conference, STOC’13, 2013*. Ed. by Dan Boneh, Tim Roughgarden, and Joan Feigenbaum. ACM, 2013, pp. 911–920.
- [Kel+14] Jonathan A. Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. “An Almost-Linear-Time Algorithm for Approximate Max Flow in Undirected Graphs, and its Multicommodity Generalizations”. In: *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014*. Ed. by Chandra Chekuri. SIAM, 2014, pp. 217–226.
- [KMP10] Ioannis Koutis, Gary L. Miller, and Richard Peng. “Approaching Optimality for Solving SDD Linear Systems”. In: *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. 2010, pp. 235–244.
- [KMP14] Ioannis Koutis, Gary L. Miller, and Richard Peng. “Approaching Optimality for Solving SDD Linear Systems”. In: *SIAM J. Comput.* 43.1 (2014), pp. 337–354.
- [Kou14] Ioannis Koutis. “Simple parallel and distributed algorithms for spectral graph sparsification”. In: *26th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA ’14*. Ed. by Guy E. Blelloch and Peter Sanders. ACM, 2014, pp. 61–66.
- [KS16] Rasmus Kyng and Sushant Sachdeva. “Approximate Gaussian Elimination for Laplacians - Fast, Sparse, and Simple”. In: *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016*. Ed. by Irit Dinur. IEEE Computer Society, 2016, pp. 573–582.
- [KS18] Udit Narayana Kar and Debarshi Kumar Sanyal. “An overview of device-to-device communication in cellular networks”. In: *ICT Express* 4.4 (2018), pp. 203–208.
- [KS20] Fabian Kuhn and Philipp Schneider. “Computing Shortest Paths and Diameter in the Hybrid Network Model”. In: *Proceedings of the 39th Symposium on Principles of Distributed Computing*. PODC ’20. Association for Computing Machinery, 2020, pp. 109–118.
- [KS22] Fabian Kuhn and Philipp Schneider. *Routing Schemes and Distance Oracles in the Hybrid Model*. 2022. arXiv: [2202.06624](https://arxiv.org/abs/2202.06624) [cs.DC].
- [Kyn+16] Rasmus Kyng, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Daniel A. Spielman. “Sparsified Cholesky and Multigrid Solvers for Connection Laplacians”. In: *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*. STOC ’16. Cambridge, MA, USA: Association for Computing Machinery, 2016, pp. 842–850.
- [Li22] Václav Rozhoň  Christoph Grunau  Bernhard Haeupler  Goran Zuzic  Jason Li. “Undirected $(1+\epsilon)$ -Shortest Paths via Minor-Aggregates: Near-Optimal Deterministic Parallel & Distributed Algorithms”. In: 2022.
- [Lin92] Nathan Linial. “Locality in Distributed Graph Algorithms”. In: *SIAM J. Comput.* 21.1 (1992), pp. 193–201.

- [Lot+03] Zvi Lotker, Elan Pavlov, Boaz Patt-Shamir, and David Peleg. “MST construction in $O(\log \log n)$ communication rounds”. In: *SPAA 2003: Proceedings of the Fifteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures, 2003*. Ed. by Arnold L. Rosenberg and Friedhelm Meyer auf der Heide. ACM, 2003, pp. 94–100.
- [LS18] Huan Li and Aaron Schild. “Spectral Subspace Sparsification”. In: *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018*. Ed. by Mikkel Thorup. IEEE Computer Society, 2018, pp. 385–396.
- [Mad16] Aleksander Madry. “Computing Maximum Flow with Augmenting Electrical Flows”. In: *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016*. Ed. by Irit Dinur. IEEE Computer Society, 2016, pp. 593–602.
- [Pel00] David Peleg. *Distributed computing: a locality-sensitive approach*. SIAM, 2000.
- [Pen16] Richard Peng. “Approximate Undirected Maximum Flows in $O(m \text{polylog}(n))$ Time”. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*. Ed. by Robert Krauthgamer. SIAM, 2016, pp. 1862–1867.
- [PR99] D. Peleg and V. Rubinfeld. “A near-tight lower bound on the time complexity of distributed MST construction”. In: *40th Annual Symposium on Foundations of Computer Science*. 1999, pp. 253–261.
- [PS14] Richard Peng and Daniel A. Spielman. “An efficient parallel solver for SDD linear systems”. In: *Symposium on Theory of Computing, STOC 2014*. Ed. by David B. Shmoys. ACM, 2014, pp. 333–342.
- [SRS18] Aaron Schild, Satish Rao, and Nikhil Srivastava. “Localization of Electrical Flows”. In: *SODA ’18*. USA: Society for Industrial and Applied Mathematics, 2018, pp. 1577–1584.
- [SS08] Daniel A. Spielman and Nikhil Srivastava. “Graph sparsification by effective resistances”. In: *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, 2008*. Ed. by Cynthia Dwork. ACM, 2008, pp. 563–568.
- [SS13] Stefan Schmid and Jukka Suomela. “Exploiting locality in distributed SDN control”. In: *HotSDN 2013 - Proceedings of the 2013 ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (2013)*, pp. 121–126.
- [ST14] Daniel A. Spielman and Shang-Hua Teng. “Nearly Linear Time Algorithms for Preconditioning and Solving Symmetric, Diagonally Dominant Linear Systems”. In: *SIAM J. Matrix Anal. Appl.* 35.3 (2014), pp. 835–885.
- [Tho01] Andrew Thomason. “The Extremal Function for Complete Minors”. In: *Journal of Combinatorial Theory, Series B* 81.2 (2001), pp. 318–338.
- [Tho84] Andrew Thomason. “An extremal function for contractions of graphs”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 95.2 (1984), pp. 261–265.
- [Vis12] Nisheeth Vishnoi. “ $Lx=b$. Laplacian solvers and their algorithmic applications”. In: *Foundations and Trends in Theoretical Computer Science* 8 (Jan. 2012).
- [Wan+10] Guohui Wang, David G. Andersen, Michael Kaminsky, Konstantina Papagiannaki, T.S. Eugene Ng, Michael Kozuch, and Michael Ryan. “C-Through: Part-Time Optics in Data Centers”. In: *Proceedings of the ACM SIGCOMM 2010 Conference*. SIGCOMM ’10. Association for Computing Machinery, 2010, pp. 327–338.
- [Zuz+22] Goran Zuzic, Gramoz Goranci, Mingquan Ye, Bernhard Haeupler, and Xiaorui Sun. “Universally-Optimal Distributed Shortest Paths and Transshipment via Graph-Based ℓ_1 -Oblivious Routing”. In: *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2022, pp. 2549–2579.

A The Laplacian Solver

In this section we describe the basic building blocks of the distributed Laplacian solver of [For+20]. Our goal will be to cast their guarantees within our more general framework, leading to the proof of [Theorem 4.2](#). First, let us introduce some notation related to Laplacian systems.

The Laplacian Matrix Consider a weighted undirected graph $G = (V, E, \mathbf{w} > 0)$. The *Laplacian* of the graph G is defined as

$$\mathcal{L}(G)_{u,v} = \begin{cases} \sum_{\{u,z\} \in E} \mathbf{w}(u,z) & \text{If } u = v, \\ -\mathbf{w}(u,v) & \text{otherwise.} \end{cases}$$

The Laplacian matrix of a graph is (i) *symmetric* ($\mathcal{L}(G)^T = \mathcal{L}(G)$); (ii) *positive semi-definite* ($\mathbf{x}^T \mathcal{L}(G) \mathbf{x} \geq 0$ for any \mathbf{x}); and (iii) *weakly diagonally dominant* ($\mathcal{L}(G)_{u,u} \geq \sum_{v \neq u} |\mathcal{L}(G)_{u,v}|$).

Definition A.1 (Schur Complement). For a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ and a partition of $[n]$ into \mathcal{T} and S , permute the rows and columns of \mathbf{A} such that

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{[S,S]} & \mathbf{A}_{[S,\mathcal{T}]} \\ \mathbf{A}_{[\mathcal{T},S]} & \mathbf{A}_{[\mathcal{T},\mathcal{T}]} \end{bmatrix}.$$

Then, the *Schur complement* of \mathbf{A} onto \mathcal{T} is defined as $\mathbf{SC}(\mathbf{A}, \mathcal{T}) := \mathbf{A}_{[\mathcal{T},\mathcal{T}]} - \mathbf{A}_{[\mathcal{T},S]} \mathbf{A}_{[S,S]}^\dagger \mathbf{A}_{[S,\mathcal{T}]}$, where \mathbf{M}^\dagger denotes the *Moore-Penrose pseudo-inverse* of matrix \mathbf{M} . For a graph G and a subset $\mathcal{T} \subseteq V(G)$, we will write $\mathbf{SC}(G, \mathcal{T}) := \mathbf{SC}(\mathcal{L}(G), \mathcal{T})$.

Notation Consider two positive semi-definite matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$. For a vector $\mathbf{x} \in \mathbb{R}^n$ we define $\|\mathbf{x}\|_{\mathbf{A}} := \sqrt{\mathbf{x}^T \mathbf{A} \mathbf{x}}$ (Mahalanobis norm). We will write $\mathbf{A} \approx_\varepsilon \mathbf{B}$ if $\exp(-\varepsilon) \mathbf{A} \preceq \mathbf{B} \preceq \exp(\varepsilon) \mathbf{A}$, where $\mathbf{A} \preceq \mathbf{B}$ if and only if the matrix $\mathbf{B} - \mathbf{A}$ is positive semi-definite. For an edge $e = \{u, v\}$, we will let $\mathbf{b}(e) := \mathbb{1}_u - \mathbb{1}_v$, where $\mathbb{1}_u \in \mathbb{R}^n$ represents the characteristic vector of node u . For a graph G with *resistances* $\mathbf{r}(e)$, we define the *leverage scores* as $\text{lev}_G(e) := \mathbf{r}(e)^{-1} \mathbf{b}^T(e) \mathcal{L}(G)^\dagger \mathbf{b}(e)$. Note that $0 \leq \text{lev}_G(e) \leq 1$.

A.1 Low-Congestion Minors

Here we introduce the concept of a *low-congestion minor*, a central component in the distributed Laplacian solver of [For+20].

Definition A.2 ([For+20]). A graph G is a *minor* of \overline{G} if the following properties hold:

1. For every node $u^G \in V(G)$ there exists:
 - (i) A subset of nodes of \overline{G} , which is termed as a *super-node*, $S^{G \rightarrow \overline{G}}(u^G)$, with a leader node $\ell(u^G) \in S^{G \rightarrow \overline{G}}(u^G)$;
 - (ii) A connected subgraph of \overline{G} on $S^{G \rightarrow \overline{G}}(u^G)$, for which we maintain a spanning tree $T^{G \rightarrow \overline{G}}(u^G)$.

2. There exists a mapping of the edges of G onto edges of \overline{G} , or self-loops, such that for any $\{u^G, v^G\} \in E(G)$, the mapped edge $\{u, v\}$ satisfies $u \in S^{G \rightarrow \overline{G}}(u^G)$ and $v \in S^{G \rightarrow \overline{G}}(v^G)$.

Moreover, we say that this minor G has *congestion* ρ , or G is a ρ -minor, if:

1. Every node $u \in \overline{G}$ is contained in at most ρ super-nodes $S^{G \rightarrow \overline{G}}(u^G)$, for some $u^G \in V(G)$;
2. Every edge of \overline{G} appears as the image of an edge of G or in one of the trees connecting super-nodes (i.e., $T^{G \rightarrow \overline{G}}(u^G)$ for some u^G) at most ρ times.

Finally, we say that G is ρ -minor distributed over \overline{G} if every $u \in V(\overline{G})$ stores:

1. All $u^G \in V(G)$ for which $u \in S^{G \rightarrow \overline{G}}(u^G)$;
2. For every edge e incident to u , (i) all the nodes u^G for which $e \in T^{G \rightarrow \overline{G}}(u^G)$, and (ii) all edges e^G that map to it.

We remark that the basis of [Definition A.2](#) was the earlier concept of a *distributed cluster graph* of Ghaffari, Karrenbauer, Kuhn, Lenzen, and Patt-Shamir [[Gha+15](#)]. The important connection is that the congested part-wise aggregation problem we introduced is the central ingredient that allows performing certain “local” operations on a graph ρ -minor distributed into the underlying communication network. The following lemma is a direct consequence of [Definition A.2](#).

Lemma A.3. *Let $G = (V, E)$ be an n -node graph ρ -minor distributed into an \bar{n} -node communication network $\overline{G} = (\overline{V}, \overline{E})$ for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Then, we can perform with high probability the following operations in the NCC model, simultaneously for all $u^G \in V(G)$, within $O(Q(\rho))$ rounds:*

1. Every leader $\ell(u^G)$ sends an $O(\log \bar{n})$ -bit message to all the nodes in $S^{G \rightarrow \overline{G}}(u^G)$;
2. All the nodes in $S^{G \rightarrow \overline{G}}(u^G)$ compute an aggregation function on $O(\log \bar{n})$ -bit inputs.

A.2 The Laplacian Building Blocks

To keep the exposition reasonably self-contained, here we review the basic ingredients of the distributed Laplacian solver developed in [[For+20](#)]. Our main goal is to extend the guarantees established in [[For+20](#)] under [Assumption 4.1](#). Then, we will combine these pieces in [Appendix A.3](#) to complete the construction.

A.2.1 Ultra-Sparsification

As is standard in the Laplacian paradigm, we will require a preconditioner in the form of an *ultra-sparsifier*. In particular, the following lemma is established in [Appendix B.4](#), and it is a refinement of [[For+20](#), Lemma 4.9]:

Lemma A.4 (Ultra-Sparsification). *Consider an n -node m -edge graph G which is ρ -minor distributed into an \bar{n} -node communication network \overline{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Then, $\text{ULTRASPARSIFY}(G, k)$ takes as input a parameter k and returns after $n^{o(1)}Q(\rho)$ rounds a graph H such that*

1. H is a subgraph of G ;

2. H has $n - 1 + m2^{O(\sqrt{\log n \log \log n})}/k$ edges;
3. $\mathcal{L}(G) \preceq \mathcal{L}(H) \preceq k\mathcal{L}(G)$.

Moreover, the algorithm returns $\widehat{G}, \mathbf{Z}_1, \mathbf{Z}_2, C$ such that

1. \widehat{G} 1-minor distributes into H such that $\widehat{G} = \mathbf{SC}(H, C)$, with $|C| = m2^{O(\sqrt{\log n \log \log n})}/k$;
2. The operators \mathbf{Z}_1 and \mathbf{Z}_2 can be evaluated in $O(Q(\rho) \log n)$ rounds, and are such that

$$\mathcal{L}(H)^\dagger = \mathbf{Z}_1^T \begin{bmatrix} \mathbf{Z}_2 & 0 \\ 0 & \mathcal{L}(\widehat{G})^\dagger \end{bmatrix} \mathbf{Z}_1.$$

Let us briefly review the pieces required for this lemma. First, we need the distributed implementation of the low-stretch spanning tree algorithm of Alon, Karp, Peleg, and West [Alo+95] which is due to Ghaffari, Karrenbauer, Kuhn, Lenzen, and Patt-Shamir [Gha+15]. Then, this spanning tree is augmented with off-tree edges based on the sampling procedure of Koutis, Miller, and Peng [KMP10], leading to a graph with a spectral approximation guarantee with respect to the original graph. Finally, the parallel elimination procedure of Blalock, Gupta, Koutis, Miller, Peng, and Tangwongsan [Ble+14] is used to perform a series of contractions, leading to a subset with size analogous to the number of off-tree edges. We revisit these steps in detail in [Appendix B.4](#).

A.2.2 Sparsified Cholesky

The next building block is the sparsified Cholesky algorithm of Kyng, Lee, Peng, Sachdeva, and Spielman [Kyn+16], which manages to effectively eliminate in every iteration a non-negligible fraction of the nodes. In the distributed context, we state the following lemma which is a refinement of [For+20, Lemma 4.10].

Lemma A.5 (Sparsified Cholesky). *Let G be an n -node graph ρ -minor distributed into a communication network \overline{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Then, for a given parameter d and error ε , the algorithm `ELIMINATE`(G, d, ε) runs in $O(Q(\rho)(\log^c n/\varepsilon^c)^d)$ rounds, where c represents some universal constant, and returns a subset $\mathcal{T} \subset V(G)$ and access to operators \mathbf{Z}_1 and \mathbf{Z}_2 such that*

1. $|\mathcal{T}| \leq (49/50)^d |V(G)|$;
2. The operators $\mathbf{Z}_1, \mathbf{Z}_1^T, \mathbf{Z}_2$ can be applied to vectors in $O(Q(\rho)(\log^c n/\varepsilon^c)^d)$ rounds;
- 3.

$$(1 - \varepsilon)^d \mathcal{L}(G)^\dagger \preceq \mathbf{Z}_1^T \begin{bmatrix} \mathbf{Z}_2 & 0 \\ 0 & \mathbf{SC}(G, \mathcal{T})^\dagger \end{bmatrix} \mathbf{Z}_1 \preceq (1 + \varepsilon)^d \mathcal{L}(G)^\dagger.$$

This lemma is established based on a distributed implementation of the *sparsified Cholesky* algorithm of Kyng, Lee, Peng, Sachdeva, and Spielman [Kyn+16]. In particular, the Cholesky decomposition essentially reduces solving a Laplacian to inverting (i) any sub-matrix of the Laplacian induced on a set S , and (ii) the Schur complement on $V \setminus S$. Thus, Kyng, Lee, Peng, Sachdeva, and Spielman [Kyn+16] initially develop a procedure for identifying an "almost independent" subset of nodes F (more precisely, a *strongly* diagonally dominant subset) for which inverting the Laplacian restricted on F can be done efficiently through preconditioning (e.g. via the *Jacobi method*), while F also contains

at least a constant fraction of the nodes. Next, a combinatorial view of the Schur complement based on a certain family of random walks (see [Dur+19]) is employed to construct a spectral sparsifier of the Schur complement on $\mathcal{T} = V \setminus F$. This process is then repeated for d iterations, leading to Lemma A.5. Several technical challenges that arise are discussed in Appendix B.5. Next, the main idea is to recurse on the set of terminals \mathcal{T} . However, in our context this requires maintaining the invariant that the underlying subgraph is cast as a minor (with a reasonable congestion) of \overline{G} . This is ensured in the following subsection.

A.2.3 Minor Schur Complement

This subsection introduces a subroutine that will be invoked after the ELIMINATE algorithm to return a low-congestion minor based on the set of terminals \mathcal{T} returned by ELIMINATE; while doing so, the algorithm will incur a small overhead in the spectral guarantee, and a limited growth in the number of nodes with respect to \mathcal{T} . This increase will be eventually negligible due to the selection of parameter d in ELIMINATE. In this context, the following lemma is a refinement of [For+20, Theorem 3].

Lemma A.6. *Let G be an n -node graph ρ -minor distributed into an \bar{n} -node communication network for which Assumption 4.1 holds for some $Q = Q(\rho)$. Then, for an error parameter $0 < \varepsilon < 0.1$ and a subset \mathcal{T} of nodes, the algorithm APPROXSC returns with high probability a graph H as a ρ -minor distribution into \overline{G} such that*

1. $\mathcal{T} \subseteq V(H)$;
2. H has $O(|\mathcal{T}| \log^2 n / \varepsilon^2)$ edges;
3. $\mathbf{SC}(H, \mathcal{T}) \approx_\varepsilon \mathbf{SC}(G, \mathcal{T})$.

This algorithm requires $O(\log^{10} n / \varepsilon^3)$ calls to a distributed Laplacian solver to accuracy $1/\text{poly}(n)$ on graphs that 2ρ -minor distribute into \overline{G} , and an overhead of $O(Q(\rho) \log^{10} \bar{n} / \varepsilon^3)$ rounds.

This result builds upon the work of Li and Schild [LS18], who (roughly speaking) established that randomly contracting an edge with probability equal to its leverage score (and otherwise deleting) would suffice. In the distributed context, Forster, Goranci, Liu, Peng, Sun, and Ye [For+20] devise a parallelized implementation of this scheme based on the *localization of electrical flows* [SRS18]. More precisely, they manage to identify a non-negligible subset of edges—which they refer to as *steady edges*—with small mutual (electrical) “correlation”, allowing for independent (and hence highly parallelized) contractions/deletions within this set. This approach employs the recursive and sketching-based method of random projections due to Spielman and Srivastava [SS08], similarly to [LS18], to estimate quantities such as leverage scores and electrical correlation. These steps are carefully reviewed in Appendix B.6.

A.2.4 Schur Complement Chain

Finally, let us introduce the concept of a Schur complement chain, and explain how it can be employed to produce a Laplacian solver.

Definition A.7. For an n -node graph G , $\{(G_i, \mathbf{Z}_{i,1}, \mathbf{Z}_{i,2}, \mathcal{T}_i)\}_{i=1}^t$ is a (γ, ε) -Schur complement chain if the following conditions hold:

1. $G_1 = G$;
2. $\mathcal{T}_i \subset V(G_{i+1}) \subset V(G_i)$ and $\mathbf{SC}(G_i, \mathcal{T}_i) \approx_\varepsilon \mathbf{SC}(G_{i+1}, \mathcal{T}_i)$;
3. $|V(G_{i+1})| \leq |V(G_i)|/\gamma$ for $i < t$, and $|V(G_t)| \leq \gamma$.
- 4.

$$(1 - \varepsilon)\mathcal{L}(G_i)^\dagger \preceq \mathbf{Z}_{i,1}^T \begin{bmatrix} \mathbf{Z}_{i,2} & 0 \\ 0 & \mathbf{SC}(G_i, \mathcal{T}_i)^\dagger \end{bmatrix} \mathbf{Z}_{i,1} \preceq (1 + \varepsilon)\mathcal{L}(G_i)^\dagger.$$

In the sequel, a Schur complement chain will be developed through [Lemmas A.4](#) to [A.6](#). Next, the following lemma implies a solution to the Laplacian system based on a suitable Schur complement chain.

Lemma A.8 ([\[For+20\]](#)). *Consider an \bar{n} -node communication network for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$, and let $\{(G_i, \mathbf{Z}_{i,1}, \mathbf{Z}_{i,2}, \mathcal{T}_i)\}_{i=1}^t$ be a (γ, ε) -Schur complement chain for an n -node graph G for some $\gamma \geq 2$ and $\varepsilon \leq 1/(C \log n)$, for a sufficiently large constant C , such that for all i :*

1. G_i ρ -minor distributes into \bar{G} ;
2. The linear operators $\mathbf{Z}_{i,1}$ and $\mathbf{Z}_{i,2}$ can be evaluated in at most $\bar{n}^{o(1)}Q(\rho)$ rounds.

Then, for any given vector \mathbf{b} , there is an algorithm which computes a vector \mathbf{x} in $\bar{n}^{o(1)}Q(\rho)$ rounds such that

$$\|\mathbf{x} - \mathcal{L}(G)^\dagger \mathbf{b}\|_{\mathcal{L}(G)} \leq \varepsilon \log n \|\mathbf{b}\|_{\mathcal{L}(G)^\dagger}.$$

A.3 Putting Everything Together

In this subsection we combine the building blocks we previously developed to establish [Theorem 4.2](#). The distributed Laplacian solver of Forster, Goranci, Liu, Peng, Sun, and Ye [\[For+20\]](#) is given in [Algorithm 1](#). We also include below the formal version of [Theorem 4.2](#).

Algorithm 1: Distributed Laplacian Solver [\[For+20\]](#): SOLVER(G, ε)

Input: An undirected weighted graph G ;

$G' := \text{SPECTRALSPARSIFY}(G)$;

$(G_1, \mathbf{Z}_{1,1}, \mathbf{Z}_{1,2}, \mathcal{T}_1, G_2) := \text{ULTRASPARSIFY}(G', k)$; // [Lemma A.4](#)

$\{(G_i, \mathbf{Z}_{i,1}, \mathbf{Z}_{i,2}, \mathcal{T}_i)\}_{i=2}^t := \text{BUILDCHAIN}(G_2, d, \varepsilon, k)$;

Solve $\mathcal{L}(G)\mathbf{x} = \mathbf{b}$ via Chebyshev preconditioning ; // [Lemma A.8](#)

Procedure BUILDCHAIN(G, d, ε, k) ;

if $|V(G)| \leq k$ **then**

 | **return** \emptyset ;

end

$(\mathbf{Z}_1, \mathbf{Z}_2, C) := \text{ELIMINATE}(G, d, \varepsilon)$; // [Lemma A.5](#)

$H := \text{APPROXSC}(G, C, \varepsilon)$; // [Lemma A.6](#)

return $(G, \mathbf{Z}_1, \mathbf{Z}_2, C) \cup \text{BUILDCHAIN}(H, d, \varepsilon, k)$;

Theorem A.9 (Full-Version of [Theorem 4.2](#)). *Consider a weighted \bar{n} -node graph \bar{G} for which [Assumption 4.1](#) holds for some $Q(\rho) = O(\rho^c \mathcal{Q}(\bar{G}))$, where c is a universal constant and $\mathcal{Q} = \mathcal{Q}(\bar{G})$*

is some parameter. Then, for any vector $\mathbf{b} \in \mathbb{R}^{\bar{n}}$ stored on its nodes and a sufficiently small error parameter $\varepsilon > 0$, $\text{SOLVER}(\bar{G}, \varepsilon)$ returns after $\bar{n}^{o(1)} \mathcal{Q} \log(1/\varepsilon)$ rounds a vector \mathbf{x} distributed on its nodes such that

$$\|\mathbf{x} - \mathcal{L}(G)^\dagger \mathbf{b}\|_{\mathcal{L}(G)} \leq \varepsilon \|\mathbf{b}\|_{\mathcal{L}(G)}.$$

The proof of this theorem is included in [Appendix B.7](#). We note that a guarantee with respect to the $\mathcal{L}(G)^\dagger$ -norm—as in [Lemma A.8](#)—can be translated to a guarantee in the $\mathcal{L}(G)$ -norm. This incurs only a logarithmic multiplicative overhead since it is assumed that the weights are polynomially bounded and the dependence on $1/\varepsilon$ is logarithmic [[Vis12](#), pp. 19–20]. Thus, the overhead is subsumed by the factor $\bar{n}^{o(1)}$.

B Omitted Proofs

In this section we include all of the proofs deferred from the main body and [Appendix A](#). We commence from [Section 2](#).

B.1 Proofs from [Section 2](#)

Proposition 2.3. *Suppose that P_1, \dots, P_k is any part-wise aggregation instance in a communication network \bar{G} . Given a shortcut of quality Q , we can solve with high probability the part-wise aggregation problem in $\tilde{O}(Q)$ CONGEST rounds.*

Proof sketch. Consider only one part P_i in isolation over the network $G[P_i] + H_i$. First, we claim that there exists a simple deterministic algorithm that computes the AND-aggregate (where each node $v \in P_i$ has a input bit $\mathbf{x}(v)$) in $O(d)$ rounds, where each edge is used to send at most $O(1)$ messages. Concretely, any node whose input is 0 will forward its input to all neighbors and deactivate itself. Any node which hears about the existence of an input-0 will forward this to all of its neighbors and deactivate itself. After $O(d)$ rounds, either all nodes have heard about the existence of a 0 or they can conclude all inputs are 1.

We continue considering only one part P_i in isolation. The next step is to elect a leader of P_i by finding the node with the smallest ID in P_i ; then, (1) iterate from the most significant bit of the ID to the least significant bit of the ID; (2) compute the AND-aggregate of the current bit of all the nodes' IDs; (3) if the AND-aggregate is 0, all nodes whose current bit of the ID is 1 will drop out.

Putting these together we have a way of computing the aggregate of a part P_i in isolation in $\tilde{O}(d)$ rounds with each edge carrying $\tilde{O}(1)$ messages: First, we elect a leader of P_i . Then, the leader initiates the computation of a spanning BFS tree of $\bar{G}[P_i] + H_i$ by broadcasting from itself to all other nodes, and each node forwards the message to all neighbors; the neighbor from which it hears the message first is the parent in the tree. Finally, by performing a convergecast over the BFS tree, one can easily compute the aggregate in $O(d)$ rounds for a single part P_i .

Finally, we have to run the algorithms on all the parts $\{P_i\}_i$ simultaneously. However, this might incur congestion issues on some edges: algorithms associated with multiple parts want to send a message through the same edge in the same round. To prevent this, we randomly delay the start of each algorithm by selecting the delay uniformly at random between 0 and $\tilde{O}(c)$. This guarantees that the total number of messages (across all parts and all rounds) that want to cross a given

edge is $\tilde{O}(c)$. Hence, randomly delaying all algorithms makes the *expected* number of messages crossing a given edge in a fixed round is $\Theta(1)$. By Chernoff bounds, this number is bounded by $\tilde{O}(1)$ with high probability. Therefore, by simulating each round of the algorithm using $\tilde{O}(1)$ rounds of communication (where each round of communication carries at most a single message across an edge), we can schedule the algorithms on all parts simultaneously [Gha15]. In turn, this allows us to complete all of the aggregates in $\tilde{O}(d + c) = \tilde{O}(Q)$ rounds. \square

B.2 Proofs from Section 3

Lemma 3.3 (Unrestricted Congested Part-Wise Aggregation). *Let \bar{G} be an \bar{n} -node graph and let $\mathbb{Z}_{\geq 1} \ni \rho \leq \text{poly}(\bar{n})$. Suppose that any (1-congested) part-wise aggregation on $\hat{G}_{O(\rho)}$ can be solved with a τ -round CONGEST algorithm on $\hat{G}_{O(\rho)}$. Then, there exists an $\tilde{O}(\rho \cdot \tau)$ -round CONGEST algorithm on \bar{G} that solves any ρ -congested part-wise aggregation instance on \bar{G} .*

Proof. Armed with Lemma 3.6, the claim essentially follows by leveraging Haeupler, Wajc, and Zuzic [HWZ21, Lemma 7.2 in the Full Version]. More precisely, we will have to slightly reformulate their result. Haeupler, Wajc, and Zuzic [HWZ21] show how, for a given part P_i , one can solve the part-wise aggregate problem on P_i by reducing it to a sequence of $\tilde{O}(1)$ -many (1-congested) part-wise aggregations between disjoint parts that are restricted to be simple paths $\mathcal{P}'_i = \{P'_{i,j}\}_{j=1}^{\tilde{O}(1)}$ (where nodes know the paths' edges they participate in). This is sufficient to prove our result: suppose we run that reduction on all parts P_i simultaneously. A single call to the part-wise aggregation on all of parts P_i combined, asks to find a ρ -congested part-wise aggregation in which the parts $\bigsqcup_i \mathcal{P}'_i$ are all simple paths. This is ρ -congested since at most ρ parts P_i use any node v , and within each such P_i , ever oracle call uses the node v at most once (since they are disjoint). \square

Let us briefly comment on the validity of our interpretation of [HWZ21, Lemma 7.2]. Their statement has a few easily reconciled differences compared to our previous usage. Most notably, they compute shortcuts for a set of parts, assuming an oracle for doing so, which is a harder problem than simply computing part-wise aggregates. However, it can be easily verified that the shortcuts are used only to facilitate solving part-wise aggregations. Hence, the proof can easily be translated to require an oracle computing only part-wise aggregations.

Lemma 3.4 (Simulating \hat{G}_ρ in \bar{G}). *For any \bar{G} and any $\mathbb{Z}_{\geq 1} \ni \rho \leq \text{poly}(\bar{n})$, we can simulate any τ -round CONGEST algorithm on \hat{G}_ρ with a $(\rho \cdot \tau)$ -round CONGEST algorithm on \bar{G} .*

Proof. Let us consider one round of communication in \hat{G}_ρ . Each node v will simulate (learn all messages coming into) its copies $v_1, \dots, v_\rho \in V(\hat{G}_\rho)$. Therefore, in each round node $v \in V(\bar{G})$ needs to learn all messages sent to v 's copies $v_1, \dots, v_\rho \in V(\hat{G}_\rho)$ from their neighbors in \hat{G}_ρ . Note that, by definition, v already knows the messages sent between any two copies v_i and v_j . Hence, in a single round v can learn all messages sent to any fixed v_i . As a result, ρ rounds of communication in \bar{G} suffice to simulate a single round in \hat{G}_ρ . \square

Fact 3.5 (Folklore, [Joh99]). *Given a (multi)graph G with n nodes and maximum degree $\Delta \leq \text{poly}(n)$, there exists a randomized CONGEST algorithm that colors the edges of G with $O(\Delta)$ colors and completes in $O(\log n)$ rounds, with high probability. The coloring is proper, i.e., two edges that share an endpoint are assigned a different color.*

Proof sketch. A simple edge-coloring algorithm presented in [Joh99] works by choosing a color uniformly at random from the set $\{1, \dots, O(\Delta)\}$ for each edge. Each edge will, with constant probability, choose a color not used by its neighbors. Then, this color stays fixed and the edge drops out. Hence, after $O(\log n)$ iterations the edges will be properly colored. Implementation-wise, we can assume there is an additional node in the middle of each edge which represents that edge (this only makes the problem harder). Each edge randomly chooses and sends its color to its endpoints which, in turn, inform on whether there is a conflict. Then, the edges send back to its endpoints whether it dropped out. This iteration is then repeated until we reach a proper coloring. \square

Claim 3.7. $D(\widehat{G}_\rho) \leq D(\overline{G}) + 1$.

Proof. First, consider any two nodes $u_i, v_j \in V(\widehat{G}_\rho)$ such that $\pi(u_i) \neq \pi(v_j)$, with $i, j \in [\rho]$. By construction of the layered graph \overline{G}_i , there exists a path of length at most $D(\overline{G})$ in the i -th layer of \widehat{G}_ρ between u_i to v_i . Thus, it follows that the (hop) distance between u_i and v_j is at most $D(\overline{G})$ given that v_j and v_i , with $i \neq j$, are adjacent—the copies form a clique in the layered graph. This also implies that the distance between any two nodes u_i and u_j , with $\pi(u_i) = \pi(u_j)$, is 1, concluding the proof. \square

B.3 Useful Routines

Before diving into the proofs of the Laplacian building blocks it will be useful to present several operations that can be performed efficiently under [Assumption 4.1](#). We stress that the proofs related to the Laplacian solver closely follow the approach in [For+20]. Our goal here is to translate them into our more general setting.

Corollary B.1 (Matrix-Vector Products). *Consider a matrix \mathbf{A} with non-zeroes supported on the edges of an n -node graph G which is ρ -minor distributed over a communication network \overline{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$, with values stored in the endpoints of the corresponding edges, and a vector $\mathbf{x} \in \mathbb{R}^n$ stored on the nodes $\ell(u^G)$ for $u^G \in V(G)$. Then, we can compute the vector $\mathbf{A}\mathbf{x} \in \mathbb{R}^n$ stored on the leader nodes $\ell(u^G)$ for all $u^G \in V(G)$ after $O(Q(\rho))$ rounds with high probability.*

The proof of this corollary follows the one by Forster, Goranci, Liu, Peng, Sun, and Ye [For+20, Corollary 4.4], but nonetheless we state it here for completeness.

Proof of Corollary B.1. The first step is to use [Assumption 4.1](#) to disseminate the coordinates of vector \mathbf{x} to the corresponding super-nodes after $Q(\rho)$ rounds; that is, for every $u^G \in V(G)$ the leader $\ell(u^G)$ passes to $S^{G \rightarrow \overline{G}}(u^G)$ the corresponding coordinate. Then, every node performs locally all the multiplications for its corresponding indices, and after ρ rounds the node can deliver this information to the corresponding super-node. Observe that this is possible because \mathbf{A} is supported on edges of G , and [Definition A.2](#) imposes an edge-congestion bound. Finally, we invoke again [Assumption 4.1](#) to sum all of the values of each super-node to the leader node, which gives the desired output requirement. \square

Another important corollary of [Assumption 4.1](#) is that we can simulate the spectral sparsification algorithm of Koutis (henceforth SPECTRALSPARSIFY) on G [Kou14]:

Corollary B.2 (Spectral Sparsification). *Consider an n -node graph G that ρ -minor distributes into an \bar{n} -node communication network \bar{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Then, for any $0 < \varepsilon < 0.1$ we can implement the SPECTRALSPARSIFY algorithm of Koutis for G after $O(Q(\rho) \log^7 n / \varepsilon^2)$ rounds, which returns with high probability a graph \tilde{G} distributed as a ρ -minor into \bar{G} such that*

- $\mathcal{L}(G) \approx_\varepsilon \mathcal{L}(\tilde{G})$ (Spectral approximation);
- \tilde{G} is a reweighted subgraph of G with $O(n \log^6 n / \varepsilon^2)$ edges in expectation.

The proof of this corollary is fairly simple (see [[For+20](#), Corollary 4.4]), but we give a sketch for completeness.

Proof of [Corollary B.2](#). The SPECTRALSPARSIFY algorithm of Koutis iteratively uses the spanner scheme of Baswana and Sen [[BS07](#)]. The latter algorithm gradually grows clusters. In particular, in each round clusters are sampled at random—a “leader” node determines whether the cluster is included in the sample, and then forwards the information to the rest of the cluster. Then, nodes compare the weights of their incident edges to decide whether they will join some cluster, and which incident edges will be added to the spanner. As a result, all the operations of the Baswana-Sun algorithm can be performed via the routine of [Assumption 4.1](#), and the claim follows. \square

Composition of Minors We also state the extensions of [[For+20](#), Lemma 4.6] and [[For+20](#), Corollary 4.7], which are related to the composition of ρ -minors.

Lemma B.3 (Composing Minors). *Consider a graph G_2 which is ρ_2 -minor distributed into a communication network \bar{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$, and a graph G_1 which is ρ_1 -minor distributed into G_2 . Then, we can compute with high probability and after $\tilde{O}(Q(\rho_1 \rho_2))$ rounds a $(\rho_1 \times \rho_2)$ -minor distribution of G_1 into \bar{G} .*

Corollary B.4 (Parallel Contraction). *Consider a graph G which is ρ -minor distributed into a communication network \bar{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. If F represents a subset of the edges of graph G , we can obtain with high probability a ρ -minor distribution of G/F into \bar{G} in $\tilde{O}(Q(\rho))$ rounds.*

Recall that the notation G/F implies the graph obtained from G after contracting all the edges in the set $F \subseteq E(G)$.

B.4 Ultra-Sparsification: Proof of [Lemma A.4](#)

The first ingredient required for [Lemma A.4](#) is a distributed version of the celebrated Alon-Karp-Peleg-West (AKPW) low-stretch spanning tree construction [[Alo+95](#)], which is due to Ghaffari, Karrenbauer, Kuhn, Lenzen, and Patt-Shamir [[Gha+15](#)]. We commence by stating their definition of a *distributed N -node cluster graph*, which incidentally was the basis for [Definition A.2](#).

Definition B.5 (Distributed Cluster Graph, [[Gha+15](#)]). A distributed N -node cluster graph is a 5-tuple $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{L}, \mathcal{T}, \psi)$ satisfying the following properties:

1. $\mathcal{V} = \{S_1, \dots, S_N\}$ forms a partition of the node set into N clusters;
2. \mathcal{E} represents a multi-set of (weighted) edges;

3. \mathcal{L} is the set of leaders such that every cluster S_i has *exactly* one leader $\ell_i \in \mathcal{L}$. The ID of the leader node will also serve as the ID of the cluster, while it is assumed that nodes know the ID of their leader, as well as the size of their cluster;
4. $\mathcal{T} = \{T_1, \dots, T_N\}$ is a set of cluster trees such that each cluster tree $T_i = (S_i, E_i)$ is a (rooted) spanning tree of the induced subgraph $G[S_i]$ of G , with root the leader of the cluster $\ell_i \in S_i$ (observe that this implies that the subgraph induced by each cluster S_i is connected);
5. $\psi : \mathcal{E} \mapsto E$ is a bijective function that maps every edge $\{S_i, S_j\} \in \mathcal{E}$ to some edge $\{u_i, u_j\} \in E$ connecting the corresponding clusters; i.e., it holds that $u_i \in S_i$ and $u_j \in S_j$. It is assumed that the two nodes u_i and u_j know that the edge $\{u_i, u_j\}$ is used to connect their respective clusters, as well as its weight.

Having introduced the concept of a distributed cluster graph, we state the following lemma, which is a direct corollary of the communication primitives we previously described.

Lemma B.6. *Let $G = (V, E)$ be an n -node graph ρ -minor distributed into an \bar{n} -node communication network $\bar{G} = (\bar{V}, \bar{E})$ for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. If $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{L}, \mathcal{T}, \psi)$ is a distributed cluster graph for G , the following operations can be performed in $\tilde{O}(Q(\rho))$ rounds:*

1. *The leader ℓ_i of each cluster S_i broadcasts an $O(\log \bar{n})$ -bit message to every node in S_i ;*
2. *Computing aggregation functions on $O(\log \bar{n})$ -bit inputs simultaneously for all clusters, assuming the tree T_i is known.*

Proof. The definition of a distributed N -node cluster graph ([Definition B.5](#)) implies that \mathcal{G} is 1-minor distributed over G , and in turn ρ -minor distributed into \bar{G} . Note that the induced distributed mapping can be obtained using $\tilde{O}(Q(\rho))$ rounds of communication by virtue of [Lemma B.3](#). Thus, [Assumption 4.1](#) leads to the desired claim. \square

As a result, it follows that the SPLITGRAPH algorithm in [[Gha+15](#)] can be simulated on a graph G which is ρ -minor distributed into \bar{G} after $n^{o(1)}Q(\rho)$ communication rounds—under [Assumption 4.1](#). In particular, this observation directly gives a distributed construction of a low-stretch spanning tree:

Lemma B.7 ([[For+20](#); [Gha+15](#)]). *Consider an n -node m -edge graph G which is ρ -minor distributed into an \bar{n} -node communication network \bar{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Then, we can construct a spanning tree T of G after $n^{o(1)}Q(\rho)$ rounds such that the nodes know upper bounds on the corresponding stretches that sum to at most $m2^{O(\sqrt{\log n \log \log n})}$.*

Importantly, it turns out that the guarantee of [Lemma B.7](#) suffices to sample edges by stretch, as implied by the following lemma.

Lemma B.8 ([[KMP10](#)]). *Consider an n -node graph G and a tree T such that the nodes know upper bounds on the corresponding stretches that sum up to α . Then, for any parameter k there is a sampling procedure, implementable locally, that gives a graph H which satisfies with high probability the following:*

1. $\mathcal{L}(G) \preceq \mathcal{L}(H) \preceq k\mathcal{L}(G)$;
2. H contains the edges of T and $O(\alpha \log n/k)$ additional edges.

The final step for establishing [Lemma A.4](#) uses the parallel elimination procedure of Blelloch, Gupta, Koutis, Miller, Peng, and Tangwongsan [Ble+14], which requires a logarithmic number of rounds under the PRAM model of computation. Thus, we can show the following lemma:

Lemma B.9. *Consider an n -node graph H which is ρ -minor distributed into an \bar{n} -node communication network \bar{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Moreover, let T be a spanning tree of H and W be the set of off-tree edges of H with respect to T . Then, there is an algorithm which runs in $O(Q(\rho) \log n)$ rounds and returns a graph \hat{G} , 1-embeddable into H , satisfying the following:*

1. \hat{G} contains $O(|W|)$ nodes and edges;
2. There are operators \mathbf{Z}_1 and \mathbf{Z}_2 , which can be evaluated in $O(Q(\rho) \log n)$ rounds, such that

$$\mathcal{L}(H)^\dagger = \mathbf{Z}_1^T \begin{bmatrix} \mathbf{Z}_2 & 0 \\ 0 & \mathcal{L}(\hat{G})^\dagger \end{bmatrix} \mathbf{Z}_1.$$

With these pieces in place, [Lemma A.4](#) follows directly from [Lemma B.7](#), [Lemma B.8](#), and [Lemma B.9](#).

B.5 Sparsified Cholesky: Proof of [Lemma A.5](#)

The proof of [Lemma A.5](#) mainly relies on a distributed implementation of the *Schur Complement Chain* (SCC) construction of Kyng, Lee, Peng, Sachdeva, and Spielman [Kyn+16]. In particular, the first step is to formalize a notion of almost-independence:

Definition B.10. A matrix \mathbf{M} is α -diagonally dominant (henceforth α -DD) if

$$\mathbf{M}_{i,i} \geq (1 + \alpha) \sum_{j \neq i} |\mathbf{M}_{i,j}|, \quad \forall i.$$

Moreover, an index set F is α -DD if $\mathbf{M}_{[F,F]}$ is α -DD.

An important observation is that computing the inverse $\mathbf{M}^{-1}[F, F]$ for an α -DD set can be efficiently performed using a preconditioned gradient descent method. In this context, Kyng, Lee, Peng, Sachdeva, and Spielman [Kyn+16] give a simple sampling algorithm for finding “large” α -DD sets given a Laplacian matrix. More precisely, their algorithm initially selects a random subset of nodes, and then it filters out these which do not meet the condition of [Definition B.10](#). This leads to the following result:

Lemma B.11. *Let G be an n -node graph ρ -minor distributed into a communication network \bar{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Then, if \mathcal{L} is the Laplacian matrix of G and $\alpha \geq 0$ some parameter, there is an algorithm which computes an α -DD subset F of \mathcal{L} of size at least $n/(8(1 + \alpha))$ in $O(Q(\rho) \log n)$ rounds with high probability.*

Indeed, the algorithm of Kyng, Lee, Peng, Sachdeva, and Spielman [Kyn+16] determines an α -DD subset of size $n/(8(1 + \alpha))$, while the round-complexity guarantee follows similarly to the proof in [For+20, Lemma 6.7]. Here we should note that the global aggregation steps required in the distributed implementation of [For+20, Lemma 6.7] can be trivially performed in $O(Q(1))$ rounds.

The next step is to construct an operator that approximates $\mathcal{L}_{[F,F]}^{-1}$, where F is an α -DD set, and can be efficiently applied to vectors. This is ensured by the following lemma:

Lemma B.12 ([For+20]). *Let G be a graph ρ -minor distributed into an \bar{n} -node communication network \bar{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Moreover, let \mathcal{L} be the Laplacian matrix associated with G , and F be a subset of $V(G)$ such that $\mathcal{L}_{[F,F]}$ is α -DD for some $\alpha \geq 4$. Then, for any vector \mathbf{b} stored on the leaders of the super-nodes, there is an algorithm which returns in $O(Q(\rho) \log(1/\varepsilon))$ rounds the vector $\mathbf{Z}\mathbf{b}$ stored on the same nodes, where \mathbf{Z} is a linear operator such that*

$$\mathcal{L}_{[F,F]} \preceq \mathbf{Z}^{-1} \preceq \mathcal{L}_{[F,F]} + \varepsilon \cdot \mathbf{SC}(\mathcal{L}, F),$$

for any sufficiently small $\varepsilon > 0$.

Again, this lemma follows from the guarantee in [Kyn+16] regarding the *Jacobi procedure*, as well as by directly adapting the distributed implementation in [For+20] using [Corollary B.1](#).

Approximating the Schur Complement. Moreover, α -DD sets will be useful in the approximation of the Schur complement induced by the complementary subset of nodes. First, let us recall a combinatorial view of the Schur complement as a Laplacian matrix with weights estimated by certain random walks:

Lemma B.13 ([Dur+19]). *Let G be an n -node weighted graph and a subset of nodes \mathcal{T} . Moreover, consider parameters $0 < \varepsilon < 1$ and $\mu = O(\log n/\varepsilon^2)$. If H is an initially empty graph, repeat for every edge $\{u, v\} \in E(G)$ and for μ iterations the following procedure:*

1. *Simulate a random walk starting from u until it first hits \mathcal{T} at some node t_1 ;*
2. *Simulate a random walk starting from v until it first hits \mathcal{T} at some node t_2 ;*
3. *Combine these two walks to get a walk $t_1 = u_0, \dots, u_\ell = t_2$, where ℓ is the length of the combined walk.*
4. *Add the edge $\{t_1, t_2\}$ to H with weight*

$$\frac{1}{\mu \sum_{i=0}^{\ell-1} 1/w(u_i, u_{i+1})}.$$

Then, the resulting graph H satisfies $\mathcal{L}(H) \approx_\varepsilon \mathbf{SC}(G, \mathcal{T})$ with high probability.

It should be noted that the random walks in the lemma are implied in the usual sense, wherein a step from a node is taken with probability proportional to the edge-weights of the incident edges. In the sequel, we will compute an α -DD set F via [Lemma B.11](#), and then the goal will be to approximate the Schur complement on the set $\mathcal{T} = V \setminus F$. Importantly, given that F is α -DD, we can guarantee that the random walks required in [Lemma B.13](#) will be short in expectation. Nonetheless, a challenge that arises in the distributed context—and in particular under the CONGEST model—is that the expected congestion of an edge may be prohibitively large. This issue will be resolved by incorporating new nodes to the terminals whenever they exceed some threshold of congestion. At the same time, however, we also have to limit the node-congestion since G is minor distributed into \bar{G} , and we can only deal with limited congestion. This will be addressed by invoking the spectral sparsification algorithm, ensuring that the average degree, and subsequently the congestion, remains limited.

Before we proceed with the algorithm that approximates the Schur complement, we note that we can implement the random walks of [Lemma B.13](#) in $\tilde{O}(Q(\rho))$ rounds under [Assumption 4.1](#), as implied by the approach in [\[For+20\]](#).

Lemma B.14 ([\[For+20\]](#)). *Let G be an n -node graph ρ -minor distributed into an \bar{n} -node communication network \bar{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Moreover, let F be an α -DD set, $\mathcal{T} = V \setminus F$ the set of terminals, $\varepsilon \in (0, 1)$ some error parameter, and $\gamma \geq 1$ the congestion parameter. Then, the algorithm `RANDOMWALKSCHUR` runs in $O(\alpha^{-1}\gamma Q(\rho) \log^2 n/\varepsilon^2)$ rounds, and returns a graph H along with its $(\alpha^{-1}\gamma \log n\rho)$ -minor distribution into \bar{G} such that*

$$\mathcal{L}(H) \approx_\varepsilon \mathbf{SC}(G, \hat{\mathcal{T}}),$$

with high probability, where $\hat{\mathcal{T}} \supseteq \mathcal{T}$ has size at most $n - |F| + O(\alpha^{-1}m\varepsilon^{-2} \log^2 n/\gamma)$.

Proof. Let us briefly describe the `RANDOMWALKSCHUR` algorithm. First, we compute the expected congestion of the family of random walks W predicted by [Lemma B.13](#) with respect to the set of terminals \mathcal{T} . This is done by propagating the congestion to neighbors for $O(\alpha^{-1} \log n)$ steps. Then, we create a new set $\hat{\mathcal{T}}$ which includes \mathcal{T} as a subset, as well as all the nodes which exceeded the congestion threshold of γ based on the estimation procedure of the previous step. Note that the congestion of a node with respect to W is simply the number of times this particular node participates in some random walk of W . By construction, it follows that the size of $\hat{\mathcal{T}}$ is $n - |F|$ along with all the nodes that exceeded the congestion threshold of γ . However, since F is an α -DD set it follows that the length of a random walk is $O(\alpha^{-1} \log n)$ with high probability, while for every edge we simulate $\mu = O(\log n/\varepsilon^2)$ random walks (this is related to the concentration of the corresponding random variables, as implied by [Lemma B.13](#)), in turn implying that the total congestion generated by these random walks is $O(\alpha^{-1}m\varepsilon^{-2} \log^2 n)$. As a result, only $O(\alpha^{-1}m\varepsilon^{-2} \log^2 n/\gamma)$ nodes can have congestion more than γ , verifying the assertion regarding the size of $\hat{\mathcal{T}}$. Next, the algorithm implements the random walks of [Lemma B.13](#), but with respect to the augmented set of terminals $\hat{\mathcal{T}}$. A Chernoff bound argument assures us that all nodes in $V \setminus \hat{\mathcal{T}}$ will have congestion $O(\gamma)$ with high probability.

In terms of the distributed implementation, estimating the congestion can be implemented in $O(\alpha^{-1}Q(\rho) \log^2 n/\varepsilon^2)$ rounds; this follows since every walk has length $O(\alpha^{-1} \log n)$ with high probability, and we execute $\mu = O(\log n/\varepsilon^2)$ iterations for every edge. Also note that a single step in the procedure estimating the congestion can be implemented in $O(Q(\rho))$ rounds. Next, the generation of the random walks with respect to the augmented set $\hat{\mathcal{T}}$ can be performed in $O(\alpha^{-1}\gamma Q(\rho) \log^2 n/\varepsilon^2)$ rounds with high probability; this uses the aforementioned guarantee for the congestion. The final step is to minor-distribute the graph H with weights as dictated by [Lemma B.13](#). This is done by assigning to the terminals the leaders of all intermediate (non-terminal) nodes. The congestion guarantee ensures that the resulting mapping is an $O(\alpha^{-1}\gamma \log n\rho)$ -minor distribution into \bar{G} . \square

Proof of [Lemma A.5](#). The `ELIMINATE` algorithm proceeds in d rounds, initializing $\mathbf{M}^{(0)}$ to be an ε -spectral sparsifier of $\mathcal{L}(G)$ (recall [Corollary B.2](#)). In every round $i \geq 1$, (i) we compute an α -DD set F_i with $\alpha := 4$; (ii) we employ [Lemma B.12](#) to have access to an operator that approximates $\mathbf{M}_{[F, F]}^{(i-1)}$; and (iii) we compute an ε -spectral sparsifier $\mathbf{M}^{(i)}$ of the Schur complement $\mathbf{SC}(\mathbf{M}^{(i-1)}, \hat{\mathcal{T}}_i)$ approximated via [Lemma B.14](#); here, $\hat{\mathcal{T}}_i = \hat{\mathcal{T}}_{i-1} - F_i + U_i$, where U_i represents the set of extra nodes

added to ensure low congestion. In particular, [Lemma B.14](#) is invoked with congestion parameter $\gamma := 1000C\alpha^{-1}\log^8 n/\varepsilon^4$, where C is a sufficiently large constant. The sparsification algorithm of Koutis ([Corollary B.2](#)) tells us that the number of edges will be $m = (n \log^6 n/\varepsilon^2)$, in turn implying that the number of nodes drops by at least a multiplicative factor of $49/50$.

In terms of the distributed implementation, notice that due to the selection of the parameters the approximation of the Schur complement ([Lemma B.14](#)) can be performed in $O(Q(\rho) \log^{10} n/\varepsilon^6)$ rounds. Next, the spectral sparsification step can be implemented in $O(Q(\rho') \log^7 n/\varepsilon^2)$, where $\rho' = \alpha^{-1}\gamma \log n \rho = O(\log^9 n/\varepsilon^4)\rho$. Thus, by virtue of [Assumption 4.1](#) we can infer that $Q(\rho') = O(\log^{c'} n/\varepsilon^{c'})Q(\rho)$, where c' is some universal constant. Thus, after d iterations the cost of these operations is bounded by $O(Q(\rho)(\log^c n/\varepsilon^c)^d)$, where c is some universal constant. Finally, the error guarantee follows directly from [Lemmas B.12 to B.14](#), after a direct argument bounding the accumulation of the error. \square

B.6 Minor Schur Complement: Proof of [Lemma A.6](#)

We commence this subsection by introducing the notion of *steady edges*, which are in a sense edges which are mutually "uncorrelated":

Definition B.15 ([\[For+20\]](#)). A stochastic subset of edges $Z \subseteq E$ is called (α, δ) -steady with respect to an m -edge graph H if

1. $\mathbb{E}_Z \left[\sum_{e \in Z} \mathbf{r}(e)^{-1} \mathbf{b}(e) \mathbf{b}(e)^T \right] \preceq \alpha \mathcal{L}(H)$;
2. For all $e \in Z$ we have $\sum_{e \neq f \in Z} \frac{|\mathbf{b}(e)^T \mathcal{L}(H)^\dagger \mathbf{b}(f)|}{\sqrt{\mathbf{r}(e)} \sqrt{\mathbf{r}(f)}} \leq \delta$;
3. For all $e \in Z$ it holds that

$$\mathbf{r}(e)^{-1} \mathbf{b}(e)^T \mathcal{L}(H)^\dagger \begin{bmatrix} \mathbf{SC}(H, \mathcal{T}) & 0 \\ 0 & 0 \end{bmatrix} \mathcal{L}(H)^\dagger \mathbf{b}(e) \leq \frac{32|\mathcal{T}|}{m}.$$

In words, the first constraint ensures that no edge will be selected in the steady set with too high of a probability; the second corresponds to the localization constraint, circumscribing the (mutual) correlation of edges within the set; and the final constraint imposes a bound on the variance, and will be used in the martingale analysis (to apply Freedman's inequality). It should be stressed that the existence of such objects is highly non-trivial, and follows from the localization of electrical flows recently shown by Schild, Rao, and Srivastava [\[SRS18\]](#). In the distributed setting, the following result will be established:

Lemma B.16 ([\[For+20\]](#)). *Let G be an n -node m -edge graph ρ -minor distributed into \overline{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. For a constant $\delta \in (0, 1)$ and a subset of terminals $\mathcal{T} \subseteq V(G)$, there exists an algorithm which has access to a distributed Laplacian solver, and returns with high probability a set of at least $\delta m / (2000C \log^2 m)$ edges in expectation which is $(\delta / (1000C \log^2 m), \delta)$ -steady, where C is a sufficiently large constant. This algorithm requires $O(\log^2 n)$ calls to a distributed Laplacian solver to $1/\text{poly}(n)$ accuracy on graphs that 2ρ -minor distribute into \overline{G} , and $O(Q(\rho) \log^2 n)$ communication rounds.*

The first step towards establishing this lemma is to approximate the correlation of edges within some arbitrary set:

Lemma B.17 ([For+20]). *Let G be an n -node graph with resistances \mathbf{r} , ρ -minor distributed into a communication network \overline{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Then, there is an algorithm, with access to a distributed Laplacian solver, which for any subset $W \subseteq E(G)$ and any edge $e \in W$ returns with high probability the quantity*

$$\sum_{e \neq f \in W} \frac{|\mathbf{b}(e)^T \mathcal{L}(G)^\dagger \mathbf{b}(e)|}{\sqrt{\mathbf{r}(e)} \sqrt{\mathbf{r}(f)}}$$

to within a factor of 2. This algorithm requires $O(\log^2 n)$ calls to a distributed Laplacian solver on graphs that ρ -minor distribute into \overline{G} to accuracy $1/\text{poly}(n)$, and an additional $O(Q(\rho) \log^2 n)$ communication rounds.

The proof of this lemma follows directly from [For+20, Lemma 5.13], and leverages the ℓ_1 -sketch of Indyk [Ind06]. Similarly, a sketch can be employed to estimate the effect of each edge on the Schur complement:

Lemma B.18 ([For+20]). *Let G be an n -node with resistances \mathbf{r}_e , ρ -minor distributed into a communication network \overline{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Then, for a subset $\mathcal{T} \subseteq V(G)$, there exists an algorithm which returns with high probability an estimate of*

$$\mathbf{r}(e)^{-1} \mathbf{b}(e)^T \mathcal{L}(G)^\dagger \begin{bmatrix} \mathbf{SC}(G, \mathcal{T}) & 0 \\ 0 & 0 \end{bmatrix} \mathcal{L}(G)^\dagger \mathbf{b}(e)$$

to within a factor of 2. This algorithm requires $O(\log n)$ calls to a distributed Laplacian solver to accuracy $1/\text{poly}(n)$ on graphs that 2ρ -minor distribute into \overline{G} , and $O(Q(\rho) \log n)$ communication rounds.

As a result, [Lemma B.16](#) is established based on the algorithm FINDSTEADY in [For+20], with the round complexity guarantee following directly from [Lemma B.17](#) and [Lemma B.18](#).

The next ingredient is a pre-processing step which ensures that all the edges have leverage scores bounded away from 0 and 1.

Lemma B.19 ([LS18]). *Let G be an n -node graph ρ -minor distributed into a communication network \overline{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. If 1.1-approximate leverage scores $\widetilde{\text{lev}}_G(e)$ for the edges in G are known, then there exists a process which returns after $\widetilde{O}(Q(\rho))$ rounds a graph H such that*

1. H is electrically equivalent to G ;
2. H is 2ρ -minor distributed into \overline{G} ;
3. All the leverage scores of edges in H are between $[3/16, 13/16]$.

Moreover, there exists a procedure which takes as input G and returns in $O(Q(\rho))$ rounds a graph resulting from collapsing paths and parallel edges, and removing non-terminal leaves, along with a ρ -minor distribution into \overline{G} .

The distributed implementation of this lemma is fairly simple, and relies on [Lemma B.3](#). We will also use the following lemma, which is based on the random projection scheme of Spielman and Srivastava [SS08]:

Lemma B.20 ([For+20]). *Let G be an n -node graph ρ -minor distributed into a communication network \overline{G} for which [Assumption 4.1](#) holds for some $Q = Q(\rho)$. Then, there is an algorithm with access to a distributed Laplacian solver which for all edges $e \in E(G)$ approximates the leverage score $\text{lev}_G(e)$ to within a factor of $1 + \delta$ with high probability. This algorithm requires $O(\log n/\delta^2)$ calls to a distributed Laplacian solver on graphs which ρ -minor distribute into \overline{G} to accuracy $1/\text{poly}(n)$, as well as $O(Q(\rho) \log n/\delta^2)$ communication rounds.*

The proof of this lemma follows directly from [For+20, Lemma 5.4], and uses Achliopta’s variant of the Johnson-Lindenstrauss lemma [Ach03]. With these pieces at hand, we are ready to describe the algorithm for computing a minor Schur complement. At each iteration we first determine a set of steady edges via [Lemma B.16](#). Then, we estimate the leverage scores via the random projection scheme of [Lemma B.20](#), and each edge in the steady set is contracted (independently) with probability given by its (approximate) leverage scores; otherwise, the edge is deleted (for this we will use [Corollary B.4](#)). We also employ [Lemma B.19](#) in every iteration to ensure that leverage scores are bounded away from 0 and 1. This process is repeated as long as the number of edges exceeds a threshold, leading to the algorithm APPROXSC in [For+20]. The next theorem was shown in [For+20] using matrix martingale analysis:

Lemma B.21 ([For+20]). *The algorithm APPROXSC takes as input a graph G with a set of terminals \mathcal{T} and an error parameter ε , and returns with high probability a graph H satisfying $|E(H)| = O(|\mathcal{T}| \log^2 n/\varepsilon^2)$ and $\mathbf{SC}(H, \mathcal{T}) \approx_\varepsilon \mathbf{SC}(G, \mathcal{T})$.*

Proof of [Lemma A.6](#). First, the algorithm only performs deletions and contractions, implying that it indeed returns a minor. Moreover, the correctness follows directly from [Lemma B.21](#). To bound the requirements of the algorithm note that APPROXSC executes $O(\log m/\alpha)$ iterations, where $\alpha := \delta/(1000C \log^2 m) = O(\varepsilon/\log^4 m)$, with high probability. In each iteration the dominant cost in terms of calls to a distributed Laplacian solver follows from the subroutine approximating leverage scores, which requires $O(\log n/\delta^2) = O(\log^5 n/\varepsilon^2)$. Thus, we may conclude that APPROXSC requires $O(\log^{10} n/\varepsilon^3)$ calls to a distributed Laplacian solver. The bound in terms of the round complexity follows similarly. \square

B.7 Proof of [Theorem A.9](#)

Theorem A.9 (Full-Version of [Theorem 4.2](#)). *Consider a weighted \overline{n} -node graph \overline{G} for which [Assumption 4.1](#) holds for some $Q(\rho) = O(\rho^c \mathcal{Q}(\overline{G}))$, where c is a universal constant and $\mathcal{Q} = \mathcal{Q}(\overline{G})$ is some parameter. Then, for any vector $\mathbf{b} \in \mathbb{R}^{\overline{n}}$ stored on its nodes and a sufficiently small error parameter $\varepsilon > 0$, $\text{SOLVER}(\overline{G}, \varepsilon)$ returns after $\overline{n}^{o(1)} \mathcal{Q} \log(1/\varepsilon)$ rounds a vector \mathbf{x} distributed on its nodes such that*

$$\|\mathbf{x} - \mathcal{L}(G)^\dagger \mathbf{b}\|_{\mathcal{L}(G)} \leq \varepsilon \|\mathbf{b}\|_{\mathcal{L}(G)}.$$

Proof. The correctness of the algorithm follows directly from [Lemmas A.4](#) to [A.6](#) and [A.8](#), so let us focus on the round complexity. By the guarantee of [Lemma A.4](#) we know that the ULTRASPARSIFY routine returns a graph G_2 such that $|V(G_2)| = |V(G_1)| 2^{O(\sqrt{\log n \log \log n})}/k$; this follows since we have sparsified the graph in the first step. Thus, for $k = 2^{(\log \overline{n})^{2/3}}$ we can infer that $|V(G_2)| \leq |V(G_1)|/k^{1-o(1)}$. Next, with regards to the Schur complement chain, [Lemmas A.5](#) and [A.6](#) imply that $|V(G_{i+1})| \leq |V(G_i)| O(0.98^d \log^2 n/\varepsilon^2)$. Hence, setting $d = 2^{(\log \log \overline{n})^2}$ and $\varepsilon = 1/(\log \overline{n})^2$ gives us that $|V(G_{i+1})| \leq |V(G_i)| 2^{-\Theta((\log \log \overline{n})^2)}$.

As a result, BUILDCHAIN returns a $(2^{\Theta((\log \log \bar{n})^2)}, \varepsilon)$ -Schur complement chain, which in turn implies that this chain has length $O(\log \bar{n}/(\log \log \bar{n})^2)$. Thus, [Lemma A.8](#) implies that we can use this chain to produce a solution in $\rho \bar{n}^{o(1)} Q(\rho)$ rounds, where ρ represents the maximum congestion of a graph along the chain; it will be establish that $\rho = \bar{n}^{o(1)}$.

Let $f(n, \rho)$ represent the number of rounds required by SOLVER on a graph with n nodes which ρ -minor distributes into \bar{G} , and $g(n, \rho)$ the number of rounds required by BUILDCHAIN with input an n -node graph which ρ -minor distributes into \bar{G} . Then, if we ignore lower order terms, it follows that

$$f(n, \rho) = \bar{n}^{o(1)} Q(\rho) + g(n/k^{1-o(1)}, \rho),$$

where we used that $|V(G_2)| \leq |V(G_1)|/k^{1-o(1)}$. Moreover, we have that

$$\begin{aligned} g(n, \rho) &= O\left((\log^c n/\varepsilon^c)^{(\log \log \bar{n})^2} Q(\rho)\right) + f(n, 2\rho) O(\log^{10} n/\varepsilon^3) + g(n/2^{\Theta((\log \log \bar{n})^2)}, \rho) \\ &= \bar{n}^{o(1)} Q(\rho) + \text{polylog}(\bar{n}) f(n, 2\rho) + g(n/2^{\Theta((\log \log \bar{n})^2)}, \rho), \end{aligned}$$

where we used that $|V(G_{i+1})| \leq |V(G_i)| 2^{-\Theta((\log \log \bar{n})^2)}$, and we ignored lower order terms. As a result, the overall increase in congestion is $2^{O(\log \bar{n}/(\log \log \bar{n})^2)} = \bar{n}^{o(1)}$. That is, all the graphs constructed $(\bar{n}^{o(1)})$ -minor distribute into \bar{G} . Finally, the theorem follows since by [Assumption 4.1](#) the dependence of $Q(\rho)$ on ρ is polynomial. \square

B.8 Proof of [Proposition 1.1](#)

Proposition 1.1. *Consider a graph \bar{G} with shortcut quality $\text{SQ}(\bar{G})$. Then, solving a Laplacian system on \bar{G} with $\varepsilon \leq \frac{1}{2}$ requires $\tilde{\Omega}(\text{SQ}(\bar{G}))$ rounds in both CONGEST and Supported-CONGEST models.*

Proof. First of all, as pointed out in [[For+20](#), Theorem 2], it suffices to establish the lower bound for a high-precision solver, i.e. for a sufficiently small $\varepsilon = 1/\text{poly}(\bar{n})$. Indeed, a low-accuracy solver ($\varepsilon \leq \frac{1}{2}$) can always be “boosted” with only an $O(\log \bar{n})$ overhead in the overall complexity.

In this context, let \bar{H} be the input to the spanning connected subgraph problem. We construct a resistor network H' so that $\mathbf{r}(e) = 1$ if $e \in E(\bar{H})$, and $\mathbf{r}(e) = \bar{n}^4$ for every edge $e \notin E(\bar{H})$. Moreover, let us select arbitrarily a node $v \in V(\bar{G})$. The key idea of the proof is to consider as input to the Laplacian solver a vector $\mathbf{b} \in \mathbb{R}^{\bar{n}}$ such that $\mathbf{b}(u) = -1$ for all $u \in V(\bar{G}) \setminus \{v\}$, while $\mathbf{b}(v) = \bar{n} - 1$.

To analyze the output of that Laplacian system, we first analyze the simpler Laplacian system with input a vector $\chi_{v,u} \in \mathbb{R}^{\bar{n}}$ for which the coordinate corresponding to node v is 1; the coordinate corresponding to node u is -1 ; and any other coordinate is set to 0. We recall the following well-known facts.

Fact B.22. *Let $\phi = \mathcal{L}(H')^\dagger \chi_{v,u}$. Then, for any node $w \in V(\bar{G})$ it holds that $\phi(v) \geq \phi(w) \geq \phi(u)$.*

Fact B.23. *Let $\phi = \mathcal{L}(H')^\dagger \chi_{v,u}$. Then, the $v - u$ effective resistance is such that $\text{res}_{H'}(v, u) = \phi(v) - \phi(u)$.*

As argued in [For+20], the output of the Laplacian with input $\chi_{v,u}$ and a sufficiently small error $\varepsilon = 1/\text{poly}(\bar{n})$ can be used to determine whether v and u are connected. Indeed, the following arguments have been extracted from their lower bound.

Claim B.24. *If u and v are connected in \overline{H} it follows that $\text{res}_{H'}(v, u) \leq \bar{n} - 1$.*

Proof. It is well-known that the effective resistances satisfy the triangle inequality. Moreover, given that v and u are connected in \overline{H} , it follows that there exists a path of length at most $\bar{n} - 1$ in H' so that every edge has resistance 1 (by construction of the resistor network H'). As a result, the triangle inequality implies that $\text{res}_{H'}(v, u) \leq \bar{n} - 1$. \square

Claim B.25. *If v and u are not connected in \overline{H} it follows that $\text{res}_{H'}(v, u) \geq \bar{n}^2$.*

Proof. Suppose that e_1, \dots, e_k are the edges leaving the connected component of v in \overline{H} , for some $k \leq \bar{n}^2$. Then, the Nash-Williams inequality implies that

$$\text{res}_{H'}(v, u) \geq \frac{1}{\sum_{i=1}^k \frac{1}{r(e_i)}} \geq \bar{n}^2,$$

by construction of the resistor network. \square

The next step of the proof is to incorporate in the analysis the error of the solver. To this end, let ϕ' be an ε -approximate solution to the linear system $\mathcal{L}(H')\phi = \chi_{v,u}$ in the sense that

$$\|\phi' - \mathcal{L}(H')^\dagger \chi_{v,u}\|_{\mathcal{L}(H')} \leq \varepsilon \|\chi_{v,u}\|_{\mathcal{L}(H')^\dagger} = \varepsilon \sqrt{\text{res}_{H'}(v, u)}.$$

Moreover, since the Laplacian matrix has integer resistances up to range $\text{poly}(\bar{n})$, it follows that for any \mathbf{x} , $\|\mathbf{x}\|_\infty \leq \text{poly}(\bar{n})\|\mathbf{x}\|_{\mathcal{L}}$. Thus, by setting $\varepsilon = 1/\text{poly}(\bar{n})$ to be sufficiently small, we have that

$$\text{res}_{H'}(v, u) - \frac{1}{\bar{n}} \leq \phi'(v) - \phi'(u) \leq \text{res}_{H'}(v, u) + \frac{1}{\bar{n}}.$$

Now we will use these bounds to argue about the initial Laplacian system with input vector \mathbf{b} . By linearity, a solution of the Laplacian system with input \mathbf{b} can be expressed as the sum of solutions of Laplacians with input $\chi_{v,u}$ over all $u \in V(\overline{G}) \setminus \{v\}$. Next, we let $\phi = \mathcal{L}(H')^\dagger \mathbf{b}$, and ϕ' be the output of the Laplacian solver for a sufficiently small $\varepsilon = 1/\text{poly}(\bar{n})$. Our analysis distinguishes between the following cases.

Case I Suppose that \overline{H} is connected. In turn, this implies that v is connected with any node $u \in V(\overline{G})$. As a result, it follows from [Fact B.22](#), [Fact B.23](#) and [Claim B.24](#) that for any node u ,

$$\phi'(v) - \phi'(u) \leq (\bar{n} - 1)^2 + 1. \tag{1}$$

Case II In the contrary case, there must be node u such that v and u are disconnected on \overline{H} . By [Claim B.25](#) and [Fact B.22](#) this yields that

$$\phi'(v) - \phi'(u) \geq \bar{n}^2 - 1. \quad (2)$$

Thus, (1) and (2) imply that the output ϕ' of the Laplacian solver contains enough information to determine whether \overline{H} is connected or not since $\bar{n}^2 - 1 > (\bar{n} - 1)^2 + 1$ for any $\bar{n} \geq 2$.

To leverage this in the CONGEST model we proceed as follows. First, node v sends to every other node in the graph its own part of the output from the Laplacian solver. This step can be clearly completed after $D(\overline{G})$ rounds. Then, each node u inspects whether the value $\phi'(v) - \phi'(u)$ is larger than $\bar{n}^2 - 1$. In that case, node u can transmit this information to the entire network; this step is easily seen to be implementable in $D(\overline{G})$. As a result, assuming that $\text{SQ}(\overline{G}) \geq 3D(\overline{G})$, the proof follows immediately from [Theorem 4.3](#). But the contrary case is also immediate since on *any* topology solving a Laplacian system trivially requires $\Omega(D(\overline{G}))$ rounds. This completes the proof. \square

C Congested Part-Wise Aggregation in the NCC Model

The purpose of this section is to establish [Lemma 3.15](#) by appropriately leveraging the machinery developed by Augustine, Ghaffari, Gmyr, Hinnenthal, Scheideler, Kuhn, and Li [[Aug+19](#)]. To this end, let us first describe one of their key communication primitives.

The Aggregation Problem In the *aggregation problem*, as defined by Augustine, Ghaffari, Gmyr, Hinnenthal, Scheideler, Kuhn, and Li [[Aug+19](#)], we are given a distributive function and a set of *aggregation parts* $\{P_1, \dots, P_k\}$, with $P_i \subseteq V(\overline{G})$ for all i . Every aggregation part is associated with some target node $t_i \in P_i$.⁸ Assuming that every node holds *exactly one* input value for each aggregation part of which it is a member, the goal is to let all the target nodes learn the aggregate values with respect to the associated aggregation parts. This setting allows a node to be part of multiple groups, and in particular, we let ℓ be the *local load*: the number of groups a given node may be included in—or an upper bound thereof. In addition, if $L = \sum_{i=1}^k |P_i|$ represents the *global load* of the aggregation problem, Augustine, Ghaffari, Gmyr, Hinnenthal, Scheideler, Kuhn, and Li [[Aug+19](#), Theorem 2.3] established the following result.

Lemma C.1 ([\[Aug+19\]](#)). *There exists an aggregation algorithm which solves with high probability the aggregation problem in $O(L/\bar{n} + \ell/\log \bar{n} + \log \bar{n})$ rounds of NCC.*

In the context of the ρ -congested part-wise aggregation problem ([Definition 3.1](#)), it is clear that $\ell \leq \rho$ and $L \leq \rho\bar{n}$. Thus, we are now ready to establish [Lemma 3.15](#), the statement of which is recalled below.

Lemma 3.15. *Let \overline{G} be an \bar{n} -node communication network. Then, we can solve with high probability any ρ -congested part-wise aggregation problem on \overline{G} after $O(\rho + \log \bar{n})$ rounds of NCC.*

Proof. We first employ the communication protocol of [Lemma C.1](#) so that after $O(\rho + \log \bar{n})$ rounds of NCC each target node learns with high probability the aggregate values with respect to the

⁸In [[Aug+19](#)] the target node does not have to belong to the corresponding aggregation part, but this additional flexibility will not be required for our purposes.

associated aggregation parts. Next, we can essentially reverse in time the previous communication pattern, but this time using the aggregate values as determined by the target nodes. As a result, every node will know with high probability the aggregate value for each of its aggregation parts after $O(\rho + \log \bar{n})$ rounds of NCC. \square