

Computing Shortest Paths and Diameter in the Hybrid Network Model

Fabian Kuhn
kuhn@cs.uni-freiburg.de
University of Freiburg

Philipp Schneider
philipp.schneider@cs.uni-freiburg.de
University of Freiburg

ABSTRACT

The HYBRID model, introduced in [Augustine et al., SODA '20], provides a theoretical foundation for networks that allow multiple communication modes. The model follows the principles of synchronous message passing, whereas nodes are allowed to use two fundamentally different communication modes. First, a local mode where nodes may exchange arbitrary information per round over edges of a local communication graph G (akin to the LOCAL model). Second, a global mode where every node may exchange $O(\log n)$ messages of size $O(\log n)$ bits per round with arbitrary nodes in the network. The HYBRID model intends to reflect the conditions of many real hybrid networks, where high-bandwidth but inherently local communication is combined with highly flexible global communication with restricted bandwidth.

We continue to explore the power and limitations of the HYBRID model by investigating the complexity of computing shortest paths and diameter of the local communication graph G . We show that the all pair shortest paths problem can be solved exactly in $\tilde{O}(\sqrt{n})$ rounds, which improves on the previous $\tilde{O}(n^{2/3})$ round algorithm and closes the gap to the known $\tilde{\Omega}(\sqrt{n})$ lower bound (up to polylog n factors). Furthermore, we give constant approximations for the k -source shortest paths problem (k -SSP) with runtime $\tilde{O}(\sqrt{k})$, provided that k is sufficiently large. As k -SSP has a lower bound of $\tilde{\Omega}(\sqrt{k})$ even for large approximation ratios, our k -SSP algorithms are almost tight for large enough k . In the case of a single source we give an exact $\tilde{O}(n^{2/5})$ -round algorithm, improving on the known $\tilde{O}(\sqrt{D})$ -round algorithm for graphs with large diameter D .

For the diameter problem we provide algorithms with complexities $\tilde{O}(n^{1/3}/\varepsilon)$ and $\tilde{O}(n^{0.397}/\varepsilon)$ and approximation factors $(3/2+\varepsilon)$ and $(1+\varepsilon)$, respectively. On the negative side, we demonstrate that the classical 2-party set-disjointness framework can be adapted for the HYBRID model to prove a lower bound of $\tilde{\Omega}(n^{1/3})$ rounds for computing the diameter exactly. For the *weighted* diameter problem the same holds for computing $(2-\varepsilon)$ -approximations for any $\varepsilon > 0$.

1 INTRODUCTION

In networking contexts the communication nodes often have access to different modes of communication with fundamentally different properties [3, 13]. For instance, in a network of mobile wireless devices the nodes are able to communicate through the cellular networking infrastructure, however close devices can also communicate locally with their wireless interfaces [21]. While point-to-point communication between any two devices through the cellular network is more flexible and faster for devices that are at great distance, short ranged communication among close devices is typically cheaper and tends to allow higher bandwidth. In a different scenario, a company or organization might combine their own local area network with communication over the Internet [33, 36]. For

data centers there are suggestions to augment the standard wired electrical communication infrastructure with high-speed optical or wireless connections [10, 14]. Solving shortest paths and diameter problems in the local communication infrastructure of such *hybrid networks* has direct applications, e.g., for learning the topology of the local network which can be used for efficient IP-routing.

The Hybrid Network Model. From an algorithmic point of view, it is interesting to ask whether the combination of multiple communication modes can lead to an asymptotic advantage over using one of the communication modes by itself. The recent paper by Augustine et al. [3] has introduced the *hybrid network model* (in the following called HYBRID) as a formal distributed communication model to study this question. The basic assumption underlying the HYBRID model is that the nodes of a network can use two modes of communication: (1) a relatively low bandwidth but flexible, global communication mode, where any two nodes can directly communicate with each other and (2) a high bandwidth local communication mode, where only neighboring nodes in the local network can directly communicate among each other.

Formally, the network consists of a set V of n nodes with IDs $[n]$ ¹ and a local communication topology given by a graph $G = (V, E)$. Communication happens in synchronous rounds. In each round nodes can do arbitrary computations based on the information they possess. Local communication is modeled with the standard LOCAL model [28, 30]: in each round, every node can exchange an arbitrary message with each of its neighbors in G . The global communication mode uses the recently introduced *node-capacitated clique* (NCC) model [2]: in each round, every node can exchange $O(\log n)$ -bit messages with up to $O(\log n)$ arbitrary nodes.

As discussed in [3], one can parameterize hybrid networks by the maximum message size λ for the local mode and the number of bits γ each node can exchange using the global mode. Then the standard LOCAL and CONGEST, the CLIQUE and the NCC model are just marginal cases of that parametrization.² Considering the whole spectrum of *non-marginal* hybrid network parameterizations, the combination LOCAL + NCC pushes both communication modes to one extreme end of the spectrum, as it leaves local communication completely unrestricted and allows only polylog n bits of global communication per node. Note that restricting the global communication regime even further, e.g., to just $O(\log n)$ bits per node and round would not change our main narrative. Any HYBRID model algorithm could still be simulated in such a slightly more restrictive model at the cost of a polylog n multiplicative factor. Permitting each node to exchange $O(\log n)$ many messages of size $O(\log n)$ bits per round simplifies the formulation of randomized algorithms.

¹We define $[k] := \{1, \dots, k\}$ for $k \in \mathbb{N}$.

²LOCAL: $\lambda = \infty, \gamma = 0$, CONGEST: $\lambda = O(\log n), \gamma = 0$, CLIQUE (+“Lenzen-Routing” [24]): $\lambda = 0, \gamma = n \log n$, NCC: $\lambda = 0, \gamma = O(\log^2 n)$.

Except for lower bounds, we neglect $\text{polylog } n$ factors anyway by using the $\tilde{O}(\cdot)$ -notation.³

Distance Computations in the HYBRID Model. As in [3], we investigate the complexity of shortest paths problems in the local communication graph G and of (approximately) computing the diameter of G for the HYBRID communication model. When only using the LOCAL model, computing the diameter or shortest paths for any number of sources clearly takes $\Theta(D)$ rounds, where D is the diameter of G . This is also true if distances only need to be computed approximately (for reasonably good approximations). Therefore, all of our algorithmic results can be interpreted as the minimum of D and the given round complexity. Note that there are graphs, for which D is linear in n .

If only the NCC model is used to communicate, the (approximate) APSP problem clearly requires $\tilde{\Omega}(n)$ rounds because in each round, each node can only receive $O(\log^2 n)$ bits of information and learning the distances to all other nodes requires at least $\Omega(n)$ bits of information. Even for the (approximate) SSSP problem nothing better than an $\tilde{O}(n)$ bound is known for general graphs in the NCC model. As the work of [3] demonstrates, in the HYBRID model, i.e., the combination of LOCAL and NCC, significantly faster algorithms are possible than in each individual model.

1.1 Contributions

The key algorithmic tool used in the APSP algorithms of [3] is an algorithm to solve a multi-message broadcast problem (called the *token dissemination problem*) in the HYBRID model. In this work, we consider a more general problem, which we call the *token routing problem*. We are given a set of sources $S \subseteq V$ from which we must route a total of K point-to-point messages of size $O(\log n)$ bits each, where each node is source and destination of at most k messages. We show that if the sources and destinations are sufficiently well spaced in the local network, then all messages can be routed in time $\tilde{O}(K/n + \sqrt{k} + \sqrt{|S|})$. By contrast, simply broadcasting all messages using the algorithm of [3] takes $\tilde{\Omega}(\sqrt{k|S|})$. We refer to Section 2 for more details.

Leveraging the token routing protocol we obtain our first result in Section 3, where we speed up a key step in the exact APSP algorithm of [3] that constituted a runtime bottleneck in the original algorithm. With this modification we improve on the previous runtime of $\tilde{O}(n^{2/3})$ and close the gap to the known lower bound $\tilde{\Omega}(\sqrt{n})$ (up to $\text{polylog}(n)$ factors).

THEOREM 1.1. *There is an algorithm that solves APSP exactly in the HYBRID model in $\tilde{O}(\sqrt{n})$ rounds w.h.p.*

As another direct application of the token routing algorithm, we show in Section 4 that a single round of the much more powerful CLIQUE model⁴ can be simulated on (a suitably chosen) subset of n^x ($x \leq 1$) nodes of G in $\tilde{O}(n^{x/2} + n^{2x-1})$ rounds in the HYBRID model. We first apply this result to the k -source shortest paths problem (k -SSP) problem, where each node needs to compute (approximate) distances to an arbitrary subset of k nodes of G .

We provide a framework algorithm for k -SSP that intuitively works as follows. It takes a CLIQUE-algorithm \mathcal{A} that approximates k -SSP as input and uses it to solve k -SSP on a small skeleton graph \mathcal{S} of G .⁵ The solution on \mathcal{S} is then extended to the whole graph G using the local network. Combining this framework algorithm with the fast CLIQUE algorithms of [7, 8], we obtain an array of results (with potential for further improvement if faster or more precise algorithms for the CLIQUE model are found).

THEOREM 1.2. *The k -SSP problem can be solved in the HYBRID model w.h.p. for $\varepsilon > 0$ and the following parameters.*

- In $\tilde{O}(n^{1/3}/\varepsilon)$ rounds for $k = n^{1/3}$ sources and approximation factor $(3+\varepsilon)$ on weighted and $(1+\varepsilon)$ on unweighted graphs.
- In $\tilde{O}(n^{1/3}/\varepsilon + \sqrt{k})$ rounds and approximation factor $(7+\varepsilon)$ on weighted and $(2+\varepsilon)$ on unweighted graphs.
- In $\tilde{O}(n^{0.397} + \sqrt{k})$ rounds and approximation factor $(3+o(1))$ on weighted graphs and runtime $\tilde{O}(n^{0.397}/\varepsilon + \sqrt{k})$ and approximation factor $(1+\varepsilon)$ on unweighted graphs.

Furthermore, we apply the above framework algorithm to the fast exact SSSP algorithm of [7], which yields an $\tilde{O}(n^{2/5})$ -round, exact SSSP algorithm for the HYBRID model. Note that this improves on the $\tilde{O}(\sqrt{\text{SPD}})$ shortest path algorithm of [3] for graphs with large shortest path diameter SPD .

THEOREM 1.3. *There is an algorithm that solves SSSP exactly in the HYBRID model in $\tilde{O}(n^{2/5})$ rounds w.h.p.*

Apart from studying shortest paths problems, we also consider a more basic distance problem, namely the problem of computing the diameter of the given network graph G . In Section 5, we reuse the observation that we can simulate the CLIQUE model on small skeleton graphs of G , in order to simulate fast diameter algorithms from [7, 8] for the CLIQUE model. We show that the diameter of the skeleton graph yields decent diameter approximations for the whole graph G . We obtain the following upper bounds.

THEOREM 1.4. *The diameter can be computed w.h.p. in the HYBRID model for $\varepsilon > 0$ and the following parameters. In $\tilde{O}(n^{1/3}/\varepsilon)$ rounds with approximation factor $(3/2+\varepsilon)$ and in $\tilde{O}(n^{0.397}/\varepsilon)$ rounds with approximation factor $(1+\varepsilon)$.*

In Sections 6 and 7 we provide algorithmic lower bounds for the HYBRID model. First we slightly modify the $\tilde{\Omega}(\sqrt{n})$ lower bound for the APSP problem given in [3] to obtain a $\tilde{\Omega}(\sqrt{k})$ lower bound for the k -SSP problem even for large approximation factors. This implies that our corresponding upper bounds are almost tight if the number of sources is sufficiently large ($k \geq n^{2/3}$ suffices).

THEOREM 1.5. *A randomized algorithm that computes an α approximation of the k -source shortest paths problem in the HYBRID model takes $\tilde{\Omega}(\sqrt{k})$ rounds, even on unweighted graphs, for any $\alpha \leq \alpha'$ and some fixed $\alpha' \in \Theta(n/\sqrt{k})$.*

Furthermore, we show that the 2-party set disjointness framework can be adapted for the HYBRID model. Specifically, we adapt the worst case graphs of [11, 17] (constructed for the CONGEST model) to create an according information bottleneck in the HYBRID model. We obtain the following lower bounds.

⁵A skeleton graph \mathcal{S} is a graph on a random subset of nodes of G that conserves the distances of G . See Appendix C for details and properties.

³The $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$ notations hide factors that are polylogarithmic in n .

⁴The congested clique model (or CLIQUE model) refers to the synchronous message passing model, where communication among nodes works as follows: in each round each node may send a $O(\log n)$ bit message to every other node.

THEOREM 1.6. *Computing the diameter in the HYBRID model takes $\Omega((n/\log^2 n)^{1/3})$ rounds even for randomized algorithms. For weighted diameter the same holds for $(2-\epsilon)$ -approximations for any $\epsilon > 0$ (weights up to $n^{1/3}$ suffice).*

Note that for the *weighted* diameter problem there is an upper bound of $\tilde{O}(n^{1/3})$ rounds for a $(2+o(1))$ -approximation. This upper bound is implied by the $(1+o(1))$ -approximative HYBRID-algorithm for SSSP with same complexity given in [3].⁶

1.2 Related Work

Hybrid Communication Networks. Apart from the work of Augustine et al. [3], there is little work on distributed algorithms in hybrid network models. In [3] an information dissemination routine to broadcast k tokens in $\tilde{O}(\sqrt{k})$ time is employed to solve various shortest paths problems. They solve APSP exactly in time $\tilde{O}(n^{2/3})$ and show that approximations have complexity $\Theta(\sqrt{n})$. For SSSP they give an exact solution in time $\tilde{O}(\sqrt{SPD})$ for graphs with shortest path diameter SPD . In a similar approach to this paper, the token dissemination protocol is used to simulate the *broadcast congested clique* (BCC) model. This results in a $(1+o(1))$ -approximation with complexity $\tilde{O}(n^{1/3})$ when combined with the fast SSSP algorithm by [5] for that model. Finally, they give an approximation scheme for SSSP with constant but large factor $(1/\epsilon)^{O(1/\epsilon)}$ in $\tilde{O}(n^\epsilon)$ rounds.

Global Communication Networks. Notably, the work of Gmyr et al. [13] studies the minimum spanning tree (MST) problem in a similar setting as [3]. However, their algorithm uses only the global communication edges (with an explicit polylogarithmic degree overlay network). They show that if the local network has bounded degree, an MST can be computed in polylogarithmic time. Following up on [13], Augustine et al. [2] introduce the more abstract NCC model and show that the MST problem can be solved in polylogarithmic time on general graphs. The NCC model can be seen as a (much weaker) generalization of the *congested clique* (CLIQUE) model (however in general, there is an $\Omega(n/\log n)$ -factor slowdown when going from the CLIQUE to the NCC model). Shortest paths problems have been intensively studied in the CLIQUE model (see, e.g., [7–9, 17, 23]). Notably, [8] gives polylogarithmic-time, constant approximation algorithms for APSP and k -SSP for the CLIQUE model. As discussed above, many of our results are obtained by simulating existing CLIQUE algorithms in the HYBRID model on small, random subgraphs.

Local Communication Networks. In the CONGEST model, there is a long series of work on exactly and approximately solving the SSSP problem (see, e.g., [5, 15, 25, 29, 34]) and the APSP problem (see, e.g., [1, 6, 18, 19, 26, 27, 29]). For SSSP, we know how to compute an $(1+o(1))$ -approximation in time $\tilde{O}(D+\sqrt{n})$ due to [5], which is known to be optimal up to polylogarithmic factors even for a much coarser approximation [34]. For the APSP problem, it was shown that $\tilde{\Omega}(n)$ rounds are needed, even for a polynomial approximation by [29]. Recently, it was shown that this bound can be matched (up to polylogarithmic factors) even for computing APSP exactly [6].

Graph Diameter. Work on distributed algorithms for the diameter of a graph was initiated by Frischknecht et al. [11]. They show that in the CONGEST model, one needs $\tilde{\Omega}(n)$ rounds for an exact solution and $\tilde{O}(D+\sqrt{n})$ rounds for a $(\frac{3}{2}-\epsilon)$ -approximation even in unweighted graphs. Both bounds have later been achieved: in [31], an elegant $O(n)$ -time algorithm for unweighted graphs was given and in [16], an $\tilde{O}(D+\sqrt{n})$ -time algorithm for computing a $\frac{3}{2}$ -approximation of the diameter in unweighted graphs was presented. For weighted graphs, computing a $(2-\epsilon)$ -approximation takes $\tilde{\Omega}(n)$ rounds [17]. The diameter problem has also been considered in the CLIQUE model by [8]. They show that a $(\frac{3}{2}+\epsilon)$ -approximation can be computed in polylogarithmic time.

1.3 Preliminaries

General Definitions. The scope of this paper is solving graph problems in the local communication graph $G = (V, E)$. We consider edges to be *undirected*. Edges are assigned *weights* $w : E \rightarrow [W]$, where W is at most polynomial in n , thus the weight of an edge (and a simple path) fits into a message. A graph is considered *unweighted* if $W = 1$. Let $w(P) = \sum_{e \in P} w(e)$ denote the length of a path $P \subseteq E$. Then the *distance* between two nodes $u, v \in V$ is defined as

$$d(u, v) := \min_{u-v\text{-path } P} w(P).$$

A path with smallest length between two nodes is called a *shortest path*. Let $|P|$ be the number of edges (or *hops*) of a path P . The *hop-distance* between two nodes u and v is defined as

$$\text{hop}(u, v) := \min_{u-v\text{-path } P} |P|.$$

The *diameter* of G is defined as

$$D(G) := \max_{u, v \in V} \text{hop}_G(u, v).$$

Let the *h -limited distance* from u to v

$$d_h(u, v) := \min_{u-v\text{-path } P, |P| \leq h} w(P).$$

If there is no u - v path P with $|P| \leq h$ we define $d_h(u, v) := \infty$.

Problem Definitions. In the general k -sources *shortest paths problem* (k -SSP) we are a given set of k sources. Every $v \in V$ has to learn $d(v, s)$ for all sources s . In the (α, β) -approximate version of the problem for $\alpha \geq 1, \beta \geq 0$, every $u \in V$ has to learn values $\tilde{d}(v, s)$ such that $d(v, s) \leq \tilde{d}(v, s) \leq \alpha \cdot d(v, s) + \beta$ for all sources s . For $\beta = 0$ we speak of the “ α -approximation” problem.

The *all-pairs shortest paths problem* (APSP) equals the case $k = n$ and the single-source shortest paths problem (SSSP) equals the case $k = 1$. In the *diameter problem* all nodes have to learn $D(G)$. We define (α, β) -approximations of the diameter problem analogously.

The *token routing problem* describes the following information routing problem. We are given a set of sender nodes $S \subseteq V$ and a set of receiver nodes $R \subseteq V$. Each sender $s \in S$ must send at most k_S and each receiver $r \in R$ must receive at most k_R tokens (of size $O(\log n)$ bits each), respectively. Each token has a dedicated receiver node $r \in R$, and each receiver $r \in R$ knows the senders it must receive a token from and how many tokens it needs to receive from each sender. The problem is solved when all nodes in R know all tokens they are the receiver of.

⁶Due to the fact that the eccentricity $e(v) := \max_{u \in V} d(u, v)$ of a node can be computed with a SSSP algorithm and $D/2 \leq e(v) \leq D$.

2 TOKEN ROUTING

In the first technical section we solve the token routing problem defined above if sender nodes S and receiver nodes R are “well distributed” which is the case (w.h.p.) if senders and receivers are sampled from V uniformly at random. We start by defining *helper sets* for the sender nodes S and receiver nodes R , respectively:

DEFINITION 2.1. *Let $W \subseteq V$ and let $k \in O(n/|W|)$ be the maximum number of tokens per node. A family $\{H_w \subseteq V \mid w \in W\}$ of helper sets fulfills the following properties for all $w \in W$ and some integer $\mu \in \Theta(\min(\sqrt{k}, n/|W|))$. (1) Each H_w has size at least μ . (2) The hop-distance is $\text{hop}(w, x) \in \tilde{O}(\mu)$ between w and any $x \in H_w$. (3) Each node is member of at most $\tilde{O}(1)$ sets H_w .*

Figuratively speaking, a family of helper sets assigns each node in W a sufficiently large set of helper nodes nearby, so that no node from V has to help too many nodes in W . Suppose the sets H_s for $s \in S$ and H_r for $r \in R$ form two families of helper sets for senders S and receivers R , respectively. Intuitively, each sender $s \in S$ (or receiver $r \in R$) can increase its bandwidth over the global network by a factor of μ , by relying on its μ helpers from H_s (or H_r).

This works as follows (assume for the sake of this explanation that $|S| \in \Theta(|R|)$ and $k := k_S \in \Theta(k_R)$, so that $\mu := \mu_S \in \Theta(\mu_R)$). In a first step we use the local network to distribute the tokens of each sender to its respective helpers. Due to property (2) this can be done in $\tilde{O}(\mu)$ rounds (each sender simply floods its tokens to depth $\tilde{O}(\mu)$ in the local network and each helper picks the tokens assigned to it by its sender).

Then each helper $v \in H_s$ takes care of sending some part of the tokens of $s \in S$ to some helper $u \in H_r$ of the according receiver $r \in R$ of that message. We ensure that this happens in a balanced manner, so that each helper has to exchange at most $\tilde{O}(k/\mu)$ tokens, due to property (1). The main challenge in this step is to pair up helpers of senders with helpers of receivers in a balanced way without helper sets knowing each other. We achieve this by relaying tokens via pseudo-random intermediate nodes (i.e., determined by a publicly known, randomly seeded hash-function).

More precisely each helper $v \in H_s$ forwards the messages it is responsible for, to a set of intermediate nodes. Afterwards each helper $u \in H_r$ requests its share of all messages to r from the appropriate intermediate nodes. Here we leverage that a receiver $r \in R$ knows the set of messages it must receive from each sender, so that its helpers $u \in H_r$ can compute which intermediates they must fetch their share of messages from.

After all tokens are sent in that manner, the helper sets H_r will collectively own all tokens for receiver $r \in R$. Finally, r collects its tokens from H_r in time $\tilde{O}(\mu)$ via the local network (again the helpers can simply flood their tokens to depth μ and the receivers pick the tokens they are recipient of). The whole process takes $\tilde{O}(\mu + k/\mu) = \tilde{O}(K/n + \sqrt{k})$ rounds, where $K \in \Theta(|S| \cdot k)$.

Note that the helper sets specified in Definition 2.1 can only exist if the sets of senders S and receivers R are sufficiently separated from each other in the network. We will see that we can find appropriate helper sets if $S \subseteq V$ and $R \subseteq V$ are sampled uniformly at random, which is sufficient for our purposes. Ultimately, we will show the following theorem.

THEOREM 2.2. *Let $S, R \subseteq V$ be sampled from V with probabilities $p_S = 1/n^\epsilon$, $p_R = 1/n^\delta$ for const. $\epsilon, \delta \in [0, 1)$, respectively. Let k_S and k_R be the number of tokens to be sent or received by any node in S and R , respectively. Let $K := |S| \cdot k_S + |R| \cdot k_R$ be the total workload. Then the token routing problem can be solved in $\tilde{O}(K/n + \sqrt{k_S} + \sqrt{k_R})$ rounds.*

2.1 Computing Helper Sets

Our first task is to compute helper sets satisfying the conditions of Definition 2.1. We propose a simple protocol that leverages fast symmetry breaking algorithms that are known for the local network (even for the more restrictive CONGEST model). Specifically, we first compute a (α, β) -ruling set. Such a set has the following properties:

DEFINITION 2.3 (C.F., [4]). *An (α, β) -ruling set for $G = (V, E)$ is a subset $\mathcal{R} \subseteq V$, such that for every $v \in V$ there is a node $r \in \mathcal{R}$ with $\text{hop}(v, r) \leq \beta$ and for any $r_1, r_2 \in \mathcal{R}$, $r_1 \neq r_2$ we have $\text{hop}(r_1, r_2) \geq \alpha$.*

There is a plethora of work on how to compute ruling sets in the distributed setting. A particular result that suits our needs in terms of parameter range and runtime is given by the following Lemma 2.4. This lemma is a Corollary of Theorem 1.1 in [22] (which also works in the stronger CONGEST model). The classic paper [4] provides the same result for LOCAL, which would also be sufficient.

LEMMA 2.4 (C.F., [4, 22]). *A $(2\mu + 1, 2\mu \lceil \log n \rceil)$ -ruling set can be computed (deterministically) in the local network in $O(\mu \log n)$ rounds.*

Before we start with the details, we give some high level view how Algorithm 1 (given below) uses a ruling set to compute helper sets. Based on an (α, β) -ruling set $\mathcal{R} \subseteq V$, we can compute a *clustering* of G . Initially, each “ruler” $r \in \mathcal{R}$ represents a separate cluster C_r . Then each node $v \in V$ determines the ruler $r \in \mathcal{R}$ closest to it (break ties arbitrarily), and joins its cluster C_r . Now we have the following properties. Each cluster contains at least $(\alpha - 1)/2$ nodes (due to the min. hop-distance of α between rulers). Furthermore, the diameter of each cluster is at most β .

Presume that each node joins some set W with probability $p \in (0, 1)$. That means, W is randomly sampled set of nodes that persists in the network in a distributed sense, i.e., each node is aware whether it is a member of W or not. The goal is to compute helper sets for W based on the aforementioned clustering. First we compute such a clustering based on a $(2\mu + 1, 2\mu \lceil \log n \rceil)$ -ruling set as described above. Note that μ is the integer parameter used in Definition 2.1 (but neglect its exact value for now). Let C be one of the clusters. The number of nodes from W within that cluster is roughly $|C \cap W| \in \tilde{O}(p \cdot |C|)$ w.h.p. Then each node v in cluster C joins the helper set H_w of each $w \in W \cap C$ with probability $q \in \tilde{\Theta}(\mu/|C|)$.

Now we have the following properties: (1) each H_w has size $|H_w| \in \tilde{\Theta}(q|C|) = \tilde{\Theta}(\mu)$ w.h.p. In fact, we can adjust the constant factors in q such that $|H_w| \geq \mu$ w.h.p. (2) the hop-distance between any pair of nodes in the same cluster C is at most $2\mu \lceil \log n \rceil \in \tilde{O}(\mu)$ and therefore the same is true for any helper set $H_w \subseteq C$. (3) each $v \in C$ joins roughly $\tilde{O}(q|C \cap W|) = \tilde{O}(qp|C|) = \tilde{O}(\mu p) = \tilde{O}(\mu|W|/n)$ helper sets. Then for $\mu \in \tilde{O}(n/|W|)$ (as required by Definition 2.1) we have that each v joins at most $\tilde{O}(1)$ sets H_w . Note that these were the three requirements for helper sets of Definition 2.1.

Algorithm 1 Compute-Helpers(μ) $\triangleright \mu$ as in Lemma 2.5

compute a $(2\mu+1, 2\mu\lceil\log n\rceil)$ -ruling set \mathcal{R} with algorithm [22]
for $2\mu\lceil\log n\rceil$ rounds **do** \triangleright learn \mathcal{R} , W within $\mu\lceil\log n\rceil$ hops
 v forwards information on \mathcal{R} , W via incident local edges
 $\mathcal{R}_v \leftarrow$ rulers that v discovered
 $r \leftarrow \arg \min_{r' \in \mathcal{R}_v} \text{hop}(v, r')$ \triangleright closest ruler
 v joins cluster C_r \triangleright establish clustering
for $4\mu\lceil\log n\rceil$ rounds **do** $\triangleright v$ learns all members of C_r
 v forwards information on C_w with $w \in \mathcal{R}$ via local edges
 $q \leftarrow \min(2\mu/|C_r|, 1)$
for each $w \in W \cap C_r$ **do**
 v joins H_w with probability q

LEMMA 2.5. Let $W \subseteq V$ be a random set obtained with sampling probability $p = \frac{1}{n^\epsilon}$ for constant $\epsilon \in (0, 1)$. Let $k \in \mathbb{N}$. Running Algorithm 1 with $\mu := \lfloor \min(\sqrt{k}, 1/p) \rfloor$ establishes a family of helper sets $\{H_w \subseteq V \mid w \in W\}$ with respect to Definition 2.1 w.h.p. Algorithm 1 takes $O(\mu \log n)$ rounds.

PROOF. The round complexity is the sum of the runtime for computing of the ruling set and the local explorations (by local exploration to some depth d we mean that each node forwards all new learned information for d rounds). Each of these steps takes $O(\mu \log n)$ (c.f. Lemma 2.4). Next we focus on the correctness.

First we show that $W \in \Theta(pn)$ w.h.p. From this it follows that $\mu \in \Theta(\min(\sqrt{k}, n/|W|))$ (as we require in Definition 2.1). Second we prove that $H_w \geq \mu$ w.h.p., which is condition (1) of Definition 2.1. Third, we show that each node is member of at most $\tilde{O}(1)$ helper sets, corresponding to condition (3). Condition (2) can be concluded immediately from the algorithm, as a node $v \in C_r$ will only join H_w of some w in within the same cluster C_r , which has diameter at most $4\mu\lceil\log n\rceil$ hops by construction.

The random number $|W| \sim \text{Bin}(n, p)$ is a sum of n independent Bernoulli variables and the expected value is $\mathbb{E}(|W|) = pn$. In the following we apply standard Chernoff bounds (c.f. Lemma A.1⁷) and use that by asymptotic growth properties, for any constant $c > 0$ there is a $n_0 \in \mathbb{N}$, s.t. $n^\delta > c \log n$ for any const. $\delta \in (0, 1)$ and all $n \geq n_0$. Then the tail distributions can be bounded as follows

$$\mathbb{P}(|W| < \frac{pn}{2}) \leq e^{-pn/8} = e^{-n^{1-\epsilon}/8} \stackrel{n \geq n_0}{\leq} e^{-c \log n} = \frac{1}{n^c},$$

$$\mathbb{P}(|W| > 2pn) \leq e^{-pn/3} = e^{-n^{1-\epsilon}/3} \stackrel{n \geq n_0}{\leq} e^{-c \log n} = \frac{1}{n^c}.$$

As outlined before, we have that $|C_r| \geq \mu$ due to the minimum distance of $2\mu+1$ between any two rulers from \mathcal{R} and since nodes join their closest ruler and G is connected. Let $w \in W \cap C_r$. If $q = 1$ in the algorithm, all $v \in C_r$ join H_w , hence $|H_w| \geq \mu$. Otherwise we have $q = 2\mu/|C_r|$. Now $|H_w| \sim \text{Bin}(|C_r|, q)$ is a sum of independent random variables over $\{0, 1\}$ and $\mathbb{E}(|H_w|) = 2\mu \leq 2/p$. Again we can bound the tail distribution of $|H_w|$ (using Lemma A.1):

$$\begin{aligned} \mathbb{P}(|H_w| < \mathbb{E}(|H_w|)/2) &\leq \mathbb{P}(|H_w| < \frac{1}{p}) \leq e^{-1/8p} \\ &= e^{-n^\epsilon/8} \stackrel{n \geq n_0}{\leq} e^{-c \log n} = \frac{1}{n^c} \end{aligned}$$

and therefore we have $|H_w| \geq \mathbb{E}(|H_w|)/2 = \mu$ w.h.p.

⁷For the purpose of self-containedness we include the Chernoff- and union bounds we are using in Appendix A

It remains to be shown that no node joins too many helper sets. Let X_v be the random number of helper sets H_w for $w \in W \cap C_r$ that v joins. Then $X_v \sim \text{Bin}(|W \cap C_r|, q)$ is again a sum of independent Bernoulli variables. The expectation is $\mathbb{E}(X_v) = pq|C_r| = 2p\mu \leq 2$. We bound the probability that X_v is at most an additive term $3c \log n$ higher than its expectation (with Lemma A.1).

$$\begin{aligned} \mathbb{P}(X_v > \mathbb{E}(X_v) + 3c \log n) &= \mathbb{P}(X_v > (1 + \frac{3}{2}c \log n) \mathbb{E}(X_v)) \\ &\leq e^{-c \log n} = \frac{1}{n^c}. \end{aligned}$$

Finally, the total number of events we showed to occur w.h.p. above is polynomial in n . Using Lemma A.3 in Appendix A (“union bound”) we can guarantee that *all* of the above events occur w.h.p. \square

2.2 Token Routing Protocol

With the sub-protocol to compute helper sets in place, the main challenge is to pair up helpers H_s of senders $s \in S$ to the according helpers H'_r of receivers $r \in R$. This means that for each token that has to go from $s \in S$ to some $r \in R$ we must pair up two helpers from H_s and H'_r that do the actual transmission. The challenge is to ensure that the pairing has a balanced degree on each side, i.e., the number of tokens each helper has to send or receive is roughly balanced. The difficulty comes from the fact that we can not rely on global information about helper sets, since broadcasting this would incur too much contention on the global network. Instead, we route messages via intermediate nodes determined through a randomly seeded hash function.

In order to keep our algorithms reasonably concise, we make a few assumptions. First, we require that each token has a *label* of the form (s, r, i) where s is the sender r is the receiver and $i \in [x]$ enumerates all tokens from s to r , where $x \leq k_S$ is the total number of tokens from s to r . Second, each receiver initially knows the set of *labels* of tokens it must receive. Since the receivers are required to know how many tokens they must receive from each sender, each receiver can compute the set of labels from each sender.

The last assumption is that each node knows a randomly seeded hash function $h : V \times V \times \mathbb{N} \rightarrow V$ (from a k -wise independent family, c.f. Definition D.1), which maps each token label (s, r, i) to an intermediate node with ID $h(s, r, i)$. We will show that a random seed of size $\tilde{O}(1)$ bits suffices for our purposes, which can be broadcast in $\tilde{O}(1)$ rounds (c.f., Lemma B.1). Algorithm 2 gives an overview of the involved subroutines.

Algorithm 2 Token-Routing(ps, pr)

$\mu_S \leftarrow \lfloor \min(\sqrt{k_S}, 1/ps) \rfloor$; $\mu_R \leftarrow \lfloor \min(\sqrt{k_R}, 1/pr) \rfloor$
Compute-Helpers(μ_S); Compute-Helpers(μ_R)
Routing-Preparation(μ_S, μ_R) \triangleright assign tasks to helpers
Routing-Scheme \triangleright forward tokens via intermediate nodes

The following subroutine Routing-Preparation summarizes the necessary steps to prepare the final protocol Routing-Scheme. Each sender and receiver determines its respective helper set. Then every sender balances its tokens it must send and every receiver balances the labels of tokens it must receive among their helpers in time $\tilde{O}(\mu_S + \mu_R)$.

Algorithm 3 Routing-Preparation(μ_S, μ_R)

for $2(\mu_S + \mu_R)\lceil \log n \rceil$ rounds **do** \triangleright detect helpers via flooding
 v forwards information on H_w, H'_w via incident edges
if $v \in S$ (or $v \in R$) **then**
 compute *balanced* assignment of tasks to helpers H_v (H'_v)
for $2(\mu_S + \mu_R)\lceil \log n \rceil$ rounds **do** \triangleright flood token(-labels) to helpers
 v picks token(-labels) assigned to it
 v forwards all other token(-labels) to local neighbors

FACT 2.6. Assume that helper sets $H_s, s \in S, H'_r, r \in R$ have already been established for parameters μ_S, μ_R . Then by executing Algorithm 3, all tokens of any sender $s \in S$ and all token-labels of any receiver $r \in R$ are distributed among their respective helpers via the local network, such that each helper $w \in H_s$ has at most $\lceil \frac{k_S}{\mu_S} \rceil$ tokens from s and each $w \in H'_r$ has at most $\lceil \frac{k_R}{\mu_R} \rceil$ token-labels from r . The assignment of token(-labels) to helpers can be computed by each sender and receiver after learning its helper set (in the first loop). Algorithm 3 takes $O((\mu_S + \mu_R) \log n)$ rounds, due to Lemma 2.5.⁸

The final sub-procedure Routing-Scheme uses the hash function $h : V \times V \times \mathbb{N} \rightarrow V$ to route each token $t_{(s,r,i)}$ via the node $h(s, r, i)$. Essentially, each helper $w \in H_s$ of a sender $s \in S$ sends its tokens to the respective intermediate nodes given by the hash-value of its label, $O(\log n)$ tokens at a time until all tokens are at intermediate nodes. Subsequently, we do the same in reverse on the receiver side. Each helper $w \in H'_r$ of a receiver $r \in R$, sends requests for its assigned token labels ($O(\log n)$ per round) to the respective intermediate nodes. The intermediate nodes forward all requested tokens in the subsequent round.

Algorithm 4 Routing-Scheme $\triangleright h : V \times V \times \mathbb{N} \rightarrow V$ globally known

$T \leftarrow$ tokens v must send \triangleright tasks of v as sender-helper
while $T \neq \emptyset$ **do**
 $\text{out} \leftarrow$ pick $O(\log n)$ tokens from T ;
 send each $t_{(s,r,i)} \in \text{out}$ to node $h(s, r, i)$ via global network
 $T \leftarrow T \setminus \text{out}$
 $I \leftarrow$ tokens v received as intermediate node
 $T' \leftarrow$ labels of tokens v must receive \triangleright tasks as receiver-helper
while $T' \neq \emptyset$ or $I \neq \emptyset$ **do**
 $\text{req} \leftarrow$ pick $O(\log n)$ token labels from T' ;
 for each $(s, r, i) \in \text{req}$ send (s, r, i) and $ID(v)$ to node $h(s, r, i)$
 $T' \leftarrow T' \setminus \text{req}$
 $\text{rcv} \leftarrow$ all requests received
 for each $(s, r, i) \in \text{rcv}$ received from some u , send $t_{(s,r,i)}$ to u
 $I \leftarrow I \setminus \{t_{(s,r,i)} \mid (s, r, i) \in \text{rcv}\}$
for $2\mu_R\lceil \log n \rceil$ rounds **do** \triangleright collect tokens from helpers
 v forwards tokens to its neighbors via local edges

LEMMA 2.7. During the execution of Algorithm 4 no node receives more than $O(\log n)$ messages per round. After Algorithm 4 has terminated, every receiver knows all tokens that it was target of. A random hash function h that guarantees the above, requires a random seed of $O(\log^2 n)$ bits. Algorithm 4 takes $\tilde{O}(\frac{k_S}{\mu_S} + \frac{k_R}{\mu_R} + \mu_R)$ rounds.

⁸We will only briefly summarize claims that we believe to be evident or easily verifiable by the reader, and refer to those claims as “Fact”.

PROOF. We start with the runtime. By Fact 2.6, each node has at most $\lceil \frac{k_S}{\mu_S} \rceil$ tokens to send and at most $\lceil \frac{k_R}{\mu_R} \rceil$ tokens to receive. Every node receives $O(\log n)$ messages per round (as we argue below), thus every node can answer all requests for tokens they receive in their role as an intermediate node directly in the subsequent round. Therefore the process of funneling tokens from sender-helpers to receiver-helpers via intermediates takes $\tilde{O}(\frac{k_S}{\mu_S} + \frac{k_R}{\mu_R})$ rounds. It takes another $\tilde{O}(\mu_R)$ rounds until receivers have collected all tokens from their helpers in the final loop.

The algorithm is correct if each receiver eventually obtains all the tokens it must receive. That is no message must be dropped during the runtime by exceeding the receive bound of $O(\log n)$ messages per node and round. In Lemma D.4 we show that this is guaranteed by selecting targets for tokens (or token requests) is a uniform and k -wise independent random process, for some $k \in \Theta(\log n)$.

Therefore, we source our hash function $h : V \times V \times \mathbb{N} \rightarrow V$ from a k -wise independent family \mathcal{H} as defined in Definition D.1. That such a family exists and that selecting a hash function from it requires $O(\log^2 n)$ random bits is implied by Lemma D.2. This guarantees that selecting targets by applying the hash function $h(s, r, i)$ on token labels (s, r, i) is k -wise independent if the labels are distinct (which they are). \square

We prove Theorem 2.2 by showing that Algorithm 2 solves token routing in $\tilde{O}(K/n + \sqrt{k_S} + \sqrt{k_R})$ rounds.

PROOF OF THEOREM 2.2. The correctness of Algorithm 2 was shown in Lemma 2.7 and what remains is to sum up the runtime of the sub-procedures. The first task is the computation of helper sets in Algorithm 1 which takes $\tilde{O}(\mu_S + \mu_R)$ rounds due to Lemma 2.5. Assigning tokens and labels to helpers in Algorithm 3 also takes $\tilde{O}(\mu_S + \mu_R)$ rounds, as outlined in Fact 2.6. The final routing process among helpers via the global network and subsequent collection of tokens via the local network can be done in $\tilde{O}(\frac{k_S}{\mu_S} + \frac{k_R}{\mu_R} + \mu_R)$ rounds, as shown by Lemma 2.7.

The total runtime is $\tilde{O}(\frac{k_S}{\mu_S} + \frac{k_R}{\mu_R} + \mu_S + \mu_R)$. In Lemma 2.5 we showed that for the parameters μ_S, μ_R computed in Algorithm 2 we have $\mu_S \in \Theta(\min(\sqrt{k_S}, n/|S|))$ and $\mu_R \in \Theta(\min(\sqrt{k_R}, n/|R|))$. Therefore, the total runtime is $\tilde{O}(\frac{|S| \cdot k_S}{n} + \frac{|R| \cdot k_R}{n} + \sqrt{k_S} + \sqrt{k_R})$. \square

3 ALL PAIRS SHORTEST PATHS

Our first application of the token routing protocol is the APSP problem. Specifically, we show that APSP can be solved *exactly* in the HYBRID model on weighted graphs in $\tilde{O}(\sqrt{n})$ rounds. This is an improvement over the $\tilde{O}(n^{2/3})$ complexity algorithm of [3]. In fact, [3] gives a $\tilde{\Omega}(\sqrt{n})$ lower bound for APSP algorithms in the HYBRID model, even for α -approximations for some $\alpha \in \Theta(\sqrt{n})$. Consequently, our algorithm settles the APSP problem in the HYBRID model in terms of runtime (up to polylog n factors) and approximation factors (i.e. admitting approximation factors up to $\tilde{O}(\sqrt{n})$ does not allow faster algorithms).

The algorithm presented in this paper reuses most of the routines given by [3] with a major rework of just one subroutine that constitutes a runtime bottleneck in the original algorithm. This section shows that the bottleneck can be resolved using the token routing protocol of Section 2. In order to keep this section reasonably

concise, we refer to [3] for full proofs and detailed pseudo-codes of the base algorithm. In order to preserve some measure of self-containedness, we summarize the exact APSP algorithm of [3] in the following few paragraphs (supplemented by some key Lemmas given in Appendix C), therefore knowledge of [3] is not strictly required to understand this section.

As in [3], we leverage the concept of a skeleton graph $\mathcal{S} = (V_{\mathcal{S}}, E_{\mathcal{S}})$ (first introduced by [37] in the context of the PRAM model). The idea is to sample a subset $V_{\mathcal{S}} \subseteq V$ randomly and compute shortest paths among sampled nodes $V_{\mathcal{S}}$ that are in relatively close hop distance of each other. These shortest paths form the edge set $E_{\mathcal{S}}$ of \mathcal{S} . A key insight is that if we sample each node into $V_{\mathcal{S}}$ with probability $\frac{1}{x}$ for some $x \in [n]$, then there will be a sampled node on a shortest path $P_{u,v}$ between any pair $u, v \in V$ roughly every $\tilde{O}(x)$ hops (c.f., Lemma C.1).

If we set the hop-length of paths in $E_{\mathcal{S}}$ to some value $h \in \tilde{O}(x)$ (the correct value of h is given in Lemma C.1), then this implies that for any pair $u, v \in V$ with sufficient hop distance $\text{hop}(u, v) \geq h$ a shortest u - v -path $P_{u,v}$ must run over the skeleton (c.f., Lemma C.2). Then the length $w(P_{u,v})$ of $P_{u,v}$ can be computed by adding the length of the sub-paths $w(P_{u,v}) = w(P_{u,s_1}) + w(P_{s_1,s_2}) + w(P_{s_2,v})$, where s_1 and s_2 are the first and the last skeleton node on $P_{u,v}$, respectively.

The computation of the distances among skeleton nodes $s_1, s_2 \in V_{\mathcal{S}}$ relies on the *global* knowledge of the skeleton \mathcal{S} . To achieve that, each skeleton node determines its neighbors in \mathcal{S} using the local network. This can be done in $\tilde{O}(x)$ rounds as edges $E_{\mathcal{S}}$ correspond to paths in G with at most $h \in \tilde{O}(x)$ hops. Since \mathcal{S} will have much fewer nodes than G w.h.p. (namely $\tilde{O}(\frac{n}{x})$) the size of $E_{\mathcal{S}}$ is manageable and we can make $E_{\mathcal{S}}$ globally known using the *token dissemination protocol* of [3] (replicated in Lemma B.1). Specifically, the $\tilde{O}((n/x)^2)$ edges of $E_{\mathcal{S}}$ can be made public knowledge in $\tilde{O}(\frac{n}{x})$ rounds.

Subsequently, the length of shortest paths P_{s_1,s_2} for any $s_1, s_2 \in V_{\mathcal{S}}$ can be computed locally by all nodes in the network. Moreover, any node $u \in V$ can compute its distance to any $s_2 \in V_{\mathcal{S}}$ by computing the minimum of $w(P_{u,s_1}) + w(P_{s_1,s_2})$, where $s_1 \in V_{\mathcal{S}}$ are the skeleton nodes in u 's vicinity (that is within $h \in \tilde{O}(x)$ hops). The latter can be determined fast in the local network. Now that any node is able to compute its own distance to all skeleton nodes, the part that remains is to compute distances to *any* node in V that has hop distance at least h . This is arguably the hardest part. In order that $u \in V$ can learn its distance to some $v \in V$, it must learn over which skeleton node $s_2 \in V_{\mathcal{S}}$ in v 's vicinity a shortest path $P_{u,v}$ runs, as well as the length of the corresponding sub-path $P_{s_2,v}$.

In [3] this is done by another run of *token dissemination* where each token corresponds to a distance label $d_h(s, v)$ for $(s, v) \in V_{\mathcal{S}} \times V$. Having all these distance labels available at every node in the network makes computing APSP simple, as any node $u \in V$ can now compute $w(P_{u,s_1}) + w(P_{s_1,s_2}) + w(P_{s_2,v})$ for any $v \in V$ and any $s_1, s_2 \in V_{\mathcal{S}}$ within h hops of u and v respectively (the minimum over $s_1, s_2 \in V_{\mathcal{S}}$ yields the true distance $d(u, v)$). However, this approach requires that $|V| \cdot |V_{\mathcal{S}}| \in \tilde{O}(\frac{n^2}{x}) \gg n$ distance labels are broadcast in the worst case, taking $\tilde{O}(n/\sqrt{x})$ rounds. This distorts the trade-off between local and global computation and pushes the

runtime up to $\tilde{O}(n^{2/3})$ (the total runtime including the last step amounts to $\tilde{O}(x + n/\sqrt{x})$ which is optimized for $x = n^{2/3}$).

Note that the above procedure actually broadcasts more information than necessary, which one can see from the fact that afterwards, *every* node can compute the distance between *every* pair $u, v \in V$ (as long as u and v have hop distance at least h). However, in the distributed version of the APSP problem, a given node is only required to learn its *own* distance to the other nodes. This is where the token routing protocol from Section 2 can help.

PROOF OF THEOREM 1.1. Suppose we do the same steps as described above, except for the last, where distance labels $d_h(s, v)$ for $(s, v) \in V_{\mathcal{S}} \times V$ are broadcast in the *whole* network with the token dissemination protocol. Up to this point, every node $v \in V$ knows its distance $d(v, s)$ to every skeleton node $s \in V_{\mathcal{S}}$ and the first $s' \in V_{\mathcal{S}}$ on a shortest path $P_{v,s}$. We call s' a *connector* for v to s (note that $\text{hop}(v, s') \leq h$ and $s' = s$ is possible). For the proof of correctness and runtime up to this point we refer to our explanations above and to [3] for a detailed analysis.

Now each node $v \in V$ creates a token $t_{(v,s)}$ with sender v and receiver s containing $\langle d_h(v, s'), ID(v), ID(s') \rangle$, i.e. the distance from v to the connector s' for the path from v to s . This creates an instance of the token routing problem, where senders correspond to V and the receivers correspond to $V_{\mathcal{S}}$. Each sender has exactly one token for each receiver, that is $k_S = |V_{\mathcal{S}}| \in \tilde{O}(\frac{n}{x})$. Conversely, each receiver expects $k_R = n$ tokens, precisely one from each node.

Solving this instance of token routing gives each skeleton node $s \in V_{\mathcal{S}}$ enough information to compute its distance to every node $v \in V$ by locally computing $d(s, s') + d_h(s', v)$ where $s' \in V_{\mathcal{S}}$ is the connector given in the token $t_{(v,s)}$, which also contains $d_h(v, s')$. In another $\tilde{O}(x)$ rounds of local communication, each skeleton nodes can distribute distance labels $\langle d(s, v), ID(s), ID(v) \rangle$ to all nodes in their $\tilde{O}(x)$ neighborhood via the unrestricted local network. Equipped with the information from these labels, each node $u \in V$ can compute its distance to each $v \in V$ by taking the minimum of $d_h(u, s) + d(s, v)$ for each $s \in V_{\mathcal{S}}$ in the h -hop neighborhood of u .

The steps of the original algorithm (excluding the last step) take $\tilde{O}(x)$ rounds for the local communication up to hop-distance h and $\tilde{O}(\frac{n}{x})$ for making the skeleton public knowledge. By Theorem 2.2, solving the above instance of token-routing takes

$$\tilde{O}\left(\frac{n \cdot (n/x) + (n/x) \cdot n}{n} + \sqrt{|V|} + \sqrt{|V_{\mathcal{S}}|}\right) = \tilde{O}\left(\frac{n}{x} + \sqrt{n}\right)$$

rounds. This results to a total runtime of $\tilde{O}(x + \frac{n}{x} + \sqrt{n}) = \tilde{O}(\sqrt{n})$ by choosing $x = \sqrt{n}$. \square

4 k-SOURCE SHORTEST PATHS

Next, we apply the token routing algorithm to simulate the *Congested Clique* model (or CLIQUE model) on a subset of nodes in a HYBRID network. This allows us to simulate powerful algorithms from [7, 8] for the CLIQUE model and transfer their results for shortest paths (this section) and diameter problems (subsequent section) into the domain of the HYBRID model. In order to be efficient, we provide a general framework that allows to translate CLIQUE algorithms that are parametrized by runtime complexity, number of sources and multiplicative and additive approximation

parameters into HYBRID algorithms. A subset of results we obtain that way are tight in terms of runtime (bar polylog n factors), which we demonstrate in Section 6.

We start by defining the CLIQUE model, which is a member of the family of synchronous message passing models (like the HYBRID model). Nodes possess unlimited computing power, unique IDs and exchange messages in synchronous rounds. In a CLIQUE each node is allowed to send one (different) $O(\log n)$ bit message to every node each round.⁹ The following corollary of Theorem 2.2 demonstrates the connection between CLIQUE and HYBRID networks that we aim to exploit in this section.

COROLLARY 4.1. *Let $S \subseteq V$ be a subset of nodes from a HYBRID network, obtained by sampling each node with probability $\frac{1}{n^{1-x}}$ for $x \in (0, 1]$. One round of the CLIQUE model can be simulated on S in $\tilde{O}(n^{2x-1} + n^{x/2})$ rounds w.h.p.*

PROOF. Each node $v \in S$ is a sender and a receiver of exactly $|S|$ messages, one to and one from each node. We make S public knowledge with a run of token dissemination in $\tilde{O}(\sqrt{|S|})$ (c.f. Lemma B.1), so each node in S knows from whom it receives a message. Then one round of CLIQUE on S corresponds to an instance of the token routing problem. Hence the result follows immediately from Theorem 2.2. In the simulation of the CLIQUE model, each $v \in S$ must send/receive at most one message to/from every other node in S , i.e. at most $|S|$ messages in total. This takes $\tilde{O}(|S|^2/n + \sqrt{|S|})$. Using Chernoff bounds (Lemma A.1) we have $|S| \in \Theta(n^x)$ w.h.p. \square

Based on this simulation result we can provide the framework to transform shortest path algorithms in the CLIQUE model into algorithms for the HYBRID model. The following theorem gives the details on how runtime, number of sources and approximation parameters of CLIQUE algorithms can be translated accordingly. Further into this section, we give the detailed simulation algorithm and then prove individual claims in separate lemmas, which, when combined, culminate in the proof of Theorem 4.2. At the end of this section, we apply the theorem to a series of known CLIQUE algorithms which generates according results for the HYBRID model. We remark that future advances in CLIQUE algorithms for shortest paths may give further improvements for the HYBRID model.

THEOREM 4.2. *Let \mathcal{A} be an (α, β) -approximation CLIQUE algorithm that computes weighted shortest paths for n^γ , $\gamma \in [0, 1]$ sources in time $T_{\mathcal{A}} \in \tilde{O}(\eta n^\delta)$ for constants $\delta \geq 0, \eta \geq 1$. Let $x := \frac{2}{3+2\delta}$.*

Then there is a HYBRID algorithm \mathcal{B} that solves the $n^{x\gamma}$ source shortest paths problem with runtime $T_{\mathcal{B}} \in \tilde{O}(\eta n^{1-x})$ and the following additional properties. On weighted graphs, \mathcal{B} computes a $(2\alpha+1+\beta/T_{\mathcal{B}})$ -approximation. On unweighted graphs, \mathcal{B} computes an $(\alpha+\frac{2}{\eta}+\frac{\beta}{T_{\mathcal{B}}})$ -approximation. For $\gamma = 1$ (\mathcal{A} solves APSP), we can do the above for an arbitrary number of $k \in [n]$ sources but this adds $\tilde{O}(\sqrt{k})$ rounds to the runtime $T_{\mathcal{B}}$. For $\gamma = 0$ (SSSP) the approximation factor improves to $(\alpha+\beta/T_{\mathcal{B}})$ for weighted and unweighted graphs.

⁹Using the constant time routing scheme from the CLIQUE model by Lenzen [24] one can even assume that in the CLIQUE model, each node can send n messages to arbitrary targets, i.e. allowing multiple messages to the same node.

4.1 Simulation Algorithm

The intuition behind the algorithm is as follows. Initially, we compute a skeleton graph $\mathcal{S} = (V_{\mathcal{S}}, E_{\mathcal{S}})$ in the HYBRID network by first sampling nodes to $V_{\mathcal{S}}$ with probability $\frac{1}{n^{1-x}}$. After the sampling, we use the local network to determine the skeleton edges $E_{\mathcal{S}}$, which represent paths in G with hop distance at most $h \in \tilde{O}(n^{1-x})$. The parameter x determines the size of the skeleton (i.e., $|V_{\mathcal{S}}| \in \tilde{O}(n^x)$ w.h.p.) and optimizes the trade-off between the time required for the simulation of \mathcal{A} on \mathcal{S} and the time spent for communication on the local network (e.g., to compute $E_{\mathcal{S}}$). This establishes a fundamental property of skeleton graphs, namely that on a shortest path between two endpoints more than h hops apart there must be a skeleton node at least every h hops. Moreover, a shortest path must run over the skeleton via two skeleton “entry-nodes”, each within h hops to one of the endpoints (we outlined this earlier in Section 3 and formalize it in Lemmas C.1, C.2).

Note that at this point each skeleton node knows only its incident edges in \mathcal{S} . In contrast to the APSP algorithm in Section 3, we do *not* make \mathcal{S} public knowledge, since this would take too long with our available tools. Instead, the local knowledge of skeleton edges by incident skeleton nodes serves as input for the corresponding graph problem on \mathcal{S} , which we solve by simulating an appropriate CLIQUE algorithm \mathcal{A} in \mathcal{S} as outlined in Corollary 4.1. Afterwards, the distances to the set of sources (or estimations thereof) are known to all skeleton nodes. But since every node in V has a skeleton node within h hops on a shortest path to a source (as formalized by Lemma C.1), they can simply gather the required information to compute their distance to all sources via the local network.

There is one catch however, as sources in a graph problem given on the whole network G will probably not coincide with the randomly sampled skeleton nodes $V_{\mathcal{S}}$. Generally speaking, we can not simply add all sources from G to the skeleton since they might be too regionally concentrated in G , thereby obstructing the creation of helper sets, where each node must only help $\tilde{O}(1)$ senders or receivers, which would also inhibit the (fast) simulation of the CLIQUE. Note however, that if we are dealing with only a single source in G , adding it to $V_{\mathcal{S}}$ is actually not a problem as each helper will be assigned at most one additional skeleton in the simulation routine (c.f., Corollary 4.1).

In general, we circumvent the mentioned problem by letting each source $v \in V$ that has not been sampled to $V_{\mathcal{S}}$ tag its closest skeleton $r_v \in \mathcal{S}$ as its *representative*, i.e. as substitute source in $V_{\mathcal{S}}$. Subsequently, we make representatives and distances to their original source publicly known (a comparatively small amount of information). After computing distances to the representatives as described above, each node can estimate its distance to the original source by adding its distance to the according representative to the (publicly known) distance from representative to the original source. The drawback is that the estimations will increase the additive error by up to $2h$ (two times the maximum hop-length of a skeleton edge) in the unweighted case and the multiplicative error by $(1+\alpha)$ in the weighted case. We show that additive errors of $\Theta(h)$ are in fact not that troublesome. On one hand, we can use the local network to compute distances up to some arbitrarily large multiple of h *exactly*. On the other hand, the multiplicative impact of additive errors on paths longer than that is relatively small.

Algorithm 5 SP-Simulation($\mathcal{A}, \gamma, \delta, \eta$) \triangleright param. of Theorem 4.2

```

 $x \leftarrow \frac{2}{3+2\delta}$   $\triangleright x$  optimizes overall runtime
Compute-skeleton( $\gamma, x$ )  $\triangleright$  construct  $\mathcal{S} = (V_S, E_S)$ 
Compute-Representatives  $\triangleright$  pick representatives in  $V_S$ 
Clique-Simulation( $\mathcal{A}, x$ )  $\triangleright$  run  $\mathcal{A}$  on  $\mathcal{S}$ 
for  $\eta h$  rounds do  $\triangleright$  flood dist. to sources locally
     $v$  forwards distances to sources via local edges
     $v$  computes own distance to all sources with Equation (1)

```

FACT 4.3. From the simulation of \mathcal{A} and the subsequent local distribution of the results in Algorithm 5, each node $v \in V$ knows an (α, β) -estimation $\tilde{d}(u, r_s)$ between each skeleton $u \in V_S$ and each representative r_s of some source s . Additionally each node knows the distance $d_h(r_s, s)$ between each source s and its representative r_s from the respective token dissemination in the sub-procedure given by Algorithm 6. If two nodes $u, v \in V$ are within h hops of each other they know their ηh -hop limited distance $d_{\eta h}(u, v)$ from the local exploration. With the available information every node $v \in V$ can compute an estimation $\tilde{d}(v, s)$ to each source $s \in V$ as follows

$$\tilde{d}(v, s) := \min(d_{\eta h}(v, s), \min_{u \in V_S} d_h(v, u) + \tilde{d}(u, r_s) + d_h(r_s, s)). \quad (1)$$

We outsource several tasks from the main algorithm to sub-procedures with the intent of reusing some of them in the next section for the diameter problem. The task of computing the skeleton graph is moved to Algorithm 6. The task of establishing the relations among sources and representatives and making these public knowledge is done by Algorithm 7. The simulation routine as outlined in Corollary 4.1 is given by Algorithm 8.

Algorithm 6 Compute-Skeleton(γ, x) $\triangleright h \in \tilde{O}(n^{1-x})$ see Lem. C.1

```

 $v$  joins  $V_S$  with probability  $\frac{1}{n^{1-x}}$   $\triangleright |V_S| \in \tilde{\Theta}(n^x)$  w.h.p.
if  $v$  is a source and  $\gamma = 0$  then  $v$  joins  $V_S$ 
for  $h \in \tilde{O}(n^{1-x})$  rounds do  $\triangleright$  locally determine  $E_S$ 
     $v$  forwards all graph information  $\leftarrow$  via local network

```

FACT 4.4. The size of V_S is $\tilde{\Theta}(n^x)$ w.h.p. Let $h := \xi n^{1-x} \ln n$, where ξ is the parameter given in Lemma C.1. Algorithm 6 establishes a weighted graph $\mathcal{S} = (V_S, E_S)$ among the sampled nodes V_S in $\tilde{O}(n^{1-x})$ rounds, whereas we define $E_S := \{\{u, v\} \mid u, v \in V_S, \text{hop}(u, v) \leq h\}$. The weight of $\{u, v\} \in E_S$ is defined as $d_h(u, v)$. After the subroutine, all skeleton nodes know their neighbors in \mathcal{S} and the distances of the incident edges in E_S . Further properties of \mathcal{S} are outlined in Lemma C.2.

Algorithm 7 Compute-Representatives

```

if  $v$  is a source then
     $v$  tags closest node in  $V_S$  as rep.  $r_v$   $\triangleright$  set  $v = r_v$  if  $v \in V_S$ 
     $v$  creates token  $\langle d_h(v, r_v), ID(v), ID(r_v) \rangle$ 
     $\triangleright d_h(v, r_v) = 0$  and  $ID(v) = ID(r_v)$  indicates  $v \in V_S$ 
     $v$  participates in token dissemination protocol
if  $v$  received token  $\langle \cdot, \cdot, ID(v) \rangle$  then
     $v$  becomes source in  $\mathcal{S}$   $\triangleright v$  was tagged by some source in  $G$ 

```

FACT 4.5. After Algorithm 7 terminates, every source s of G has a representative $r_s \in V_S$ ($s = r_s$ if $s \in V_S$) and r_s and the distances $d_h(s, r_s)$ are known by the whole network for all sources s of G .

Algorithm 8 Clique-Simulation(\mathcal{A}, x) \triangleright CLIQUE alg. \mathcal{A}

```

for  $i \leftarrow 1$  to  $T_{\mathcal{A}}$  do  $\triangleright$  simulation of  $\mathcal{A}$ 
    if  $v \in V_S$  then  $\triangleright$  nodes  $V_S$  take part in simulation
         $v$  performs computations based on  $\mathcal{A}$  round  $i$ 
         $v$  creates tokens  $t_{(v,u)}^i$  to nodes  $u \in V_S$  for round  $i$ 
    Token-Routing( $\frac{1}{n^{1-x}}, \frac{1}{n^{1-x}}$ )  $\triangleright V \setminus V_S$  take part as helpers

```

The full proof of Theorem 4.2 is given in the following subsection. In a nutshell, the proof of the runtime of Algorithm 5 consists of summing up the runtimes of the simulation of \mathcal{A} (which is $\tilde{O}(\eta n^{x\delta+x/2})$) and the local exploration to compute small distances exactly (which is $\tilde{O}(\eta n^{1-x})$). The trade-off is optimized for the parameter $x := \frac{2}{3+2\delta}$ given in the proof.

Although lengthy, proving the claimed approximation ratios is not that technically involved. It boils down to proving the error inflicted by computing shortest paths to representatives of sources instead of directly to the sources. For this we rely extensively on the properties of skeleton graphs (given in Lemmas C.1 and C.2) and the triangle inequality. Subsequently, we leverage the exact knowledge of distances of paths with few hops to turn additive errors into small multiplicative errors.

4.2 Proof of Theorem 4.2

The most straight forward part of the proof of Theorem 4.2 is showing the runtime of Algorithm 5. Note that in the formulation of this and subsequent lemmas we will assume Algorithm 5 receives a CLIQUE algorithm \mathcal{A} as input, which is parametrized according to the specifications given in Theorem 4.2.

LEMMA 4.6. Given n^{xy} sources, Algorithm 5 takes $\tilde{O}(\eta n^{1-x})$ rounds.

PROOF. During the execution of Algorithm 5 we conduct two explorations in the local network, the longer of which takes $\eta h \in \tilde{O}(\eta n^{1-x})$ rounds (h is the parameter from Lemma C.1). Moreover, we simulate the CLIQUE algorithm \mathcal{A} on the skeleton \mathcal{S} . A round of simulation takes $\tilde{O}(n^{2x-1} + n^{x/2})$ rounds in accordance with Corollary 4.1.

Note that for the chosen $x = \frac{2}{3+2\delta}$ one exponent dominates the runtime of the simulation, as $2x-1 = \frac{1-2\delta}{3+2\delta} \leq \frac{1}{3+2\delta} = \frac{x}{2}$. Therefore the runtime simplifies to $\tilde{O}(n^{x/2})$. Since the number of nodes on which we simulate \mathcal{A} is only $|V_S| = \tilde{O}(n^x)$, running \mathcal{A} takes $\tilde{O}(\eta n^{x\delta})$ simulated rounds, that is $\tilde{O}(\eta n^{x\delta+x/2})$ actual rounds. For the value we chose for x the exponent equals

$$x\delta + \frac{x}{2} = \frac{2\delta}{3+2\delta} + \frac{1}{3+2\delta} = 1 - \frac{2}{3+2\delta} = 1 - x,$$

thus running \mathcal{A} takes $\tilde{O}(\eta n^{1-x})$ rounds in total.

Finally, we note that since each source in G picks at most one representative in V_S that acts as a source in \mathcal{S} , the number of sources $\Theta(|V_S|^y) = \Theta(n^{xy})$ in \mathcal{S} is roughly in line with the allowed number to execute \mathcal{A} on \mathcal{S} . Specifically, we have at most a constant factor more sources on \mathcal{S} than \mathcal{A} allows w.h.p. (due to the randomized

sampling of V_S). This does not matter much as we can repeat \mathcal{A} a constant number of times with a constant fraction of sources. \square

Let us now turn to the correctness of the distance approximation computed by Algorithm 5.

LEMMA 4.7. *On unweighted graphs with n^{γ} sources and $\gamma > 0$, Alg. \mathcal{B} (represented by Alg. 5) computes an $(\alpha + \frac{2}{\eta} + \frac{\beta}{T_{\mathcal{B}}})$ -approximation w.h.p., where $T_{\mathcal{B}} \in \tilde{O}(\eta n^{1-x})$. We get a $(1+2\alpha + \frac{\beta}{T_{\mathcal{B}}})$ -approximation on weighted graphs, w.h.p.*

PROOF. Let $v, s \in V$ be a node and a source, respectively. In case $\text{hop}(u, v) \leq \eta h$, the value $\tilde{d}(v, s)$ computed in Equation (1) is equal to $d_{\eta h}(v, s)$ (since s is a source and due to Lemma C.1), i.e. the algorithm computes the exact distance. From now on we assume $\text{hop}(u, v) > \eta h$.

Next we give some definitions and state a few facts that we implicitly use for the subsequent calculations. Let $s \in V$ be a source, let $r_s \in V_S$ be the representative of s and let Q be a shortest v - r_s -path. According to Lemma C.1 there is a skeleton node $w \in V_S$ on Q within h hops of v (possibly $w = r_s$). Note that v learns the distance $d_h(v, w)$ and in fact $d_h(v, z)$ for any $z \in V_S$ within h hops during the exploration via the local network in Algorithm 5.

By definition, \mathcal{A} computes an approximation $\tilde{d}(u, r_s)$ with

$$d(u, r_s) \leq \tilde{d}(u, r_s) \leq \alpha \cdot d(u, r_s) + \beta$$

from any skeleton node $u \in V_S$ to r_s and after the algorithm has finished, the values $\tilde{d}(u, r_s)$ are known by v (c.f. Fact 4.3). Finally, since s must have at least one skeleton node within h hops distance, its representative $r_s \in V_S$ (the skeleton closest to v) must clearly be within h hops of s as well. The distance label $d_h(r_s, s)$ is made public knowledge, i.e., also known by v (c.f. Fact 4.5). We have

$$\begin{aligned} w(Q) &= d(v, r_s) = d_h(v, w) + d(w, r_s) \\ &= \min_{u \in V_S} d_h(v, u) + d(u, r_s) \\ &\leq \min_{u \in V_S} d_h(v, u) + \tilde{d}(u, r_s). \end{aligned}$$

Let P be an actual shortest path from v to s . We know $|P| \geq \eta h$. We show that we do not underestimate $w(P)$

$$\begin{aligned} w(P) &= d(v, s) \leq d(v, r_s) + d_h(r_s, s) = w(Q) + d_h(r_s, s) \\ &\leq \min_{u \in V_S} d_h(v, u) + \tilde{d}(u, r_s) + d_h(r_s, s) = \tilde{d}(v, s). \end{aligned}$$

Note that the right hand side of the above inequality equals the right hand side of Equation (1), given that $\text{hop}(u, v) > \eta h$. Next we upper bound the estimation $\tilde{d}(v, s)$ given in Equation (1) with a constant multiple of $w(P)$ as follows

$$\begin{aligned} \tilde{d}(v, s) &= \min_{u \in V_S} d_h(v, u) + \tilde{d}(u, r_s) + d_h(r_s, s) \\ &\leq d_h(v, w) + \tilde{d}(w, r_s) + d_h(r_s, s) \\ &\leq d_h(v, w) + \alpha \cdot d(w, r_s) + \beta + d_h(r_s, s). \end{aligned} \quad (2)$$

The first inequality is due to the fact that $d_h(v, w) + \tilde{d}(w, r_s)$ is one of the options occurring in the minimum function. Assume G is unweighted. Then from the above inequality and by exploiting $d(v, s) = \text{hop}(v, s) > \eta h$ we obtain

$$\begin{aligned} \tilde{d}(v, s) &\leq d_h(v, w) + \alpha \cdot d(w, r_s) + \beta + d_h(r_s, s) \leq \alpha \cdot d(v, s) + \beta + 2h \\ &\leq \alpha \cdot d(v, s) + \frac{\beta}{\eta h} d(v, s) + \frac{2}{\eta} d(v, s) = (\alpha + \frac{\beta}{\eta h} + \frac{2}{\eta}) d(v, s) \end{aligned}$$

Now assume that G is weighted. We continue from Inequality (2), where we replace w with the first skeleton node $w' \in V_S$ on a shortest path from v to s (which also upper bounds the minimum). Note that $\text{hop}(v, w') \leq h$ (by Lemma C.1). More explanations about the individual steps are given after the following equations.

$$\begin{aligned} \tilde{d}(v, s) &\leq d_h(v, w') + \alpha \cdot d(w', r_s) + \beta + d_h(r_s, s) \\ &\leq d_h(v, w') + \alpha \cdot (d(w', s) + d(s, r_s)) + \beta + d_h(r_s, s) \\ &\leq \alpha \cdot (d_h(v, w') + d(w', s)) + (1 + \alpha) d_h(r_s, s) + \beta \quad (\alpha \geq 1) \\ &= \alpha \cdot d(v, s) + (1 + \alpha) d_h(r_s, s) + \beta \quad (\text{Def. of } w') \\ &\leq \alpha \cdot d(v, s) + (1 + \alpha) d(v, s) + \beta \quad (d_h(r_s, s) \leq d(v, s)) \\ &= (1 + 2\alpha) \cdot d(v, s) + \beta \\ &\leq (1 + 2\alpha + \beta/\eta h) \cdot d(v, s) \quad (d(v, s) \geq \eta h) \end{aligned}$$

In the second last inequality we exploit that $r_s \in V_S$ minimizes $d_h(s, r_s)$ and therefore $d_h(s, r_s) \leq d(v, s)$ w.h.p., as we show in the following. For a contradiction, presume $d_h(s, r_s) > d(v, s)$. Since $\text{hop}(v, s) \geq h$ there is a skeleton node r' on a shortest v - s path P within h hops of s w.h.p. due to Lemma C.1. But $d_h(s, r') \leq d(v, s) < d_h(s, r_s)$, thus r' would be the representative of s .

In the last inequality we use $d(v, s) \geq \eta h$ for the weighted case. Since the minimum weight of an edge is one, this follows because P has more than ηh edges.

Finally, we note that we can do the exploration of the graph via the local network in parallel during the whole runtime $T_{\mathcal{B}} \in \tilde{O}(\eta n^{1-x})$ of Algorithm 5, which gives exact $T_{\mathcal{B}}$ -hop distances. Thus, we can actually replace occurrences of ηh in the above approximations with $T_{\mathcal{B}}$ (the terms are equal up to polylog n factors anyway). \square

Next we show that the number of sources does not have to be restricted if \mathcal{A} is an APSP algorithm. The downside of having many sources is an additional $\tilde{O}(\sqrt{k})$ rounds required to make representatives public knowledge.

LEMMA 4.8. *If \mathcal{A} solves APSP in CLIQUE we achieve the same parameters as in Lemma 4.7 for an arbitrary number of $k \in [n]$ sources, w.h.p. This adds $\tilde{O}(\sqrt{k})$ rounds to the runtime $T_{\mathcal{B}}$ of Algorithm 5.*

PROOF. From the point of correctness, the only reason that is prohibiting the use of an arbitrary number of sources is that the number of representatives tagged as sources in \mathcal{S} by sources in G might be significantly more than $|V_S|^\gamma$, i.e. the number of sources allowed by \mathcal{A} . If \mathcal{A} is an APSP algorithm we have no such restriction and may allow an arbitrary number of sources in G .

With respect to runtime, each source in G must announce its representative in \mathcal{S} and the distance to it, in order that correctness is preserved. In case we have k sources this takes $\tilde{O}(\sqrt{k})$ time using token dissemination (c.f., Lemma B.1), whereas the rest of the algorithm and its runtime remains unaffected. Therefore, the runtime only increases by $\tilde{O}(\sqrt{k})$ rounds. \square

Finally, we show that the approximation factor improves significantly for SSSP. This is due to the observation that we can always summon the single source into the skeleton without straining helpers in the subsequent token routing protocol too much (each node has to help at most one additional node).

LEMMA 4.9. *If \mathcal{A} solves SSSP in the CLIQUE, then \mathcal{B} gives an $(\alpha + \beta/T_{\mathcal{B}})$ -approximation on weighted graphs w.h.p.*

PROOF. Algorithm 6 always adds s to the skeleton if it is a single source. It is not hard to see that the token routing protocol still works if we add a single node to the randomly sampled skeleton \mathcal{S} . Therefore, runtime of Algorithm 5 remains unaffected. However, the approximation factor improves since the distance to the source does not have to take into account the detour via a representative.

Specifically, all nodes learn the distance approximations $\tilde{d}(u, s)$ computed by \mathcal{A} for all $u \in V_{\mathcal{S}}$. Additionally each node $v \in V$ knows the distance to the first skeleton $w \in V_{\mathcal{S}}$ on a shortest path P from v to s w.h.p. (since w must be within h hops due to Lemma C.1). Hence in any case all nodes learn at least an (α, β) -approximation by taking the minimum of $d_h(v, u) + d(u, s)$ for all $u \in V_{\mathcal{S}}$ (c.f. Equation (1) where $r_s = s$).

If the shortest v - s -path has fewer than ηh hops, we can give an exact solution (c.f. Fact 4.3 and Equation (1)). Otherwise, we know that $d(v, s) \geq \eta h$ (and even $d(v, s) \geq T_{\mathcal{B}} \in \tilde{O}(\eta h)$ since each node learns G up to depth $T_{\mathcal{B}}$ during the whole runtime of \mathcal{B}). Thus, we can turn the additive error into a small multiplicative error

$$\begin{aligned} \tilde{d}(v, s) &\leq \alpha \cdot d(v, s) + \beta \leq \left(\alpha + \frac{\beta}{d(v, s)}\right) \cdot d(v, s) \\ &\leq \left(\alpha + \frac{\beta}{\eta h}\right) \cdot d(v, s) \leq \left(\alpha + \frac{\beta}{T_{\mathcal{B}}}\right) \cdot d(v, s). \quad \square \end{aligned}$$

4.3 Implied Results for k-SSP and SSSP

In the following subsection we apply Theorem 4.2 on known results for the CLIQUE model [7, 8], which will generate a series of results for the HYBRID model for different numbers of sources, approximation parameters and time complexity.

The first CLIQUE algorithm is from [7] (Theorem 1.2), which allows for a variable number n^{γ} of sources and achieves a runtime of $\tilde{O}(1/\varepsilon)$ for $\gamma = 1/2$ and approximation factor $(1 + \varepsilon)$. Note that the corollary below could be stated more generally for arbitrary $\gamma \geq 1/2$, which we opted not to do for conciseness.

COROLLARY 4.10. *There is an algorithm that approximates the $n^{1/3}$ -source shortest path problem in the HYBRID model with running time $\tilde{O}(n^{1/3}/\varepsilon)$ and approximation factor $(1 + \varepsilon)$ on unweighted graphs and $(3 + \varepsilon)$ on weighted graphs, w.h.p.*

PROOF. Let \mathcal{A} be the algorithm from [7] Theorem 1.2. We fix $\gamma = 1/2$. Then \mathcal{A} has the parameters $\delta = 0$, $\eta = 1/\varepsilon$ and $\alpha = (1 + \varepsilon)$. This implies $x = 2/3$. From Theorem 4.2 we obtain a runtime of $\tilde{O}(n^{1/3}/\varepsilon)$ on graphs with $n^{1/3}$ sources. On weighted graphs the approximation factor is $2\alpha + 1 = 3 + 2\varepsilon \stackrel{\varepsilon' := 2\varepsilon}{=} 3 + \varepsilon'$. On unweighted graphs it is $\alpha + \frac{2}{\eta} = \alpha + 2\varepsilon = 1 + 3\varepsilon \stackrel{\varepsilon' := 3\varepsilon}{=} 1 + \varepsilon'$. \square

The next CLIQUE algorithm stems from [7] (Theorem 1.1). It solves APSP in $\tilde{O}(1/\varepsilon)$ rounds with approximation parameters $(\alpha, \beta) = (2 + \varepsilon, (1 + \varepsilon)w_{uv})$ for a path from u to v , whereas w_{uv} is the heaviest edge on that path.

COROLLARY 4.11. *There is an algorithm that approximates shortest paths for k sources in the HYBRID model with running time $\tilde{O}(\frac{n^{1/3}}{\varepsilon} + \sqrt{k})$ and approximation factor $(7 + \varepsilon)$ on weighted graphs and $(2 + \varepsilon)$ on unweighted graphs, w.h.p.*

PROOF. The defining parameters of the algorithm from [7] Theorem 1.1 are $\gamma = 1$, $\delta = 0$ and $\eta = 1/\varepsilon$. This implies $x = 2/3$ and a runtime of $\tilde{O}(n^{1/3}/\varepsilon)$ of the resulting HYBRID algorithm by Theorem 4.2. The multiplicative error is $\alpha = (2 + \varepsilon)$ and the additive error computed on the skeleton is $\beta = (1 + \varepsilon)w_{uv}$, where $w_{u,v}$ is the heaviest edge on a shortest path from u to v .

We first consider the case where G is weighted. We observe that $(1 + \varepsilon)w_{uv} \leq (1 + \varepsilon)d(u, v)$, thus we can get rid of the additive error of \mathcal{A} by simply adding $(1 + \varepsilon)$ to the multiplicative error instead. That is, we can consider \mathcal{A} as $(3 + 2\varepsilon)$ -approximation algorithm. Then, by Theorem 4.2, we have a multiplicative error $2(3 + 2\varepsilon) + 1 = 7 + 4\varepsilon \stackrel{\varepsilon' := 4\varepsilon}{=} 7 + \varepsilon'$.

Let us consider the unweighted case. We can also bound the additive error by $(1 + \varepsilon)W_{\mathcal{S}}$ where $W_{\mathcal{S}}$ is the largest edge weight of the skeleton \mathcal{S} . Since G is unweighted we have $W_{\mathcal{S}} \leq h$. Then, by Theorem 4.2 we obtain a multiplicative error of

$$\alpha + \frac{2}{\eta} + \frac{(1 + \varepsilon)W_{\mathcal{S}}}{T_{\mathcal{B}}} \stackrel{T_{\mathcal{B}} \geq \eta h}{\leq} \alpha + 2\varepsilon + \frac{(1 + \varepsilon)h}{\eta h} \leq 2 + 3\varepsilon + \frac{1 + \varepsilon}{1/\varepsilon} \stackrel{\varepsilon \leq 1}{\leq} 2 + 5\varepsilon \stackrel{\varepsilon' := 5\varepsilon}{=} 2 + \varepsilon'. \quad \square$$

We turn to a result of [8], which offers a more accurate APSP algorithm for the CLIQUE. They solve APSP with multiplicative error $(1 + o(1))$ in $\tilde{O}(n^{\rho + o(1)})$ rounds, where ρ is the so called *distributed matrix multiplication coefficient* [8]. The value $\rho \leq 1 - \frac{2}{\omega}$ is algorithmically bounded by the *matrix multiplication coefficient* ω , for which $2 \leq \omega < 2.3728639$ is known [12], thus $\rho < 0.15715$. It is conjectured by some that $\omega = 2$, raising the possibility of further improvements in the bound of ρ . To keep the result brief, we do not express the runtime as a function of ρ and only insert its known upper bound. The algorithm of [8] allows edge weights up to $2^{(n^{o(1)})}$ which is n^c if we choose $(\log \log n + \log c)/\log n$ for $o(1)$. This and the previous two corollaries constitute the proof of Theorem 1.2.

COROLLARY 4.12. *There is a HYBRID algorithm for the k -SSP problem with runtime $\tilde{O}(n^{0.397} + \sqrt{k})$ and approx. factor $(3 + o(1))$ on weighted graphs and runtime $\tilde{O}(n^{0.397}/\varepsilon + \sqrt{k})$ and approx. factor $(1 + \varepsilon)$ on unweighted graphs w.h.p.*

PROOF. We characterize the above algorithm with $\delta = \rho < 0.15715$, $\eta = 0$ and $\alpha = 1 + o(1) \leq 1 + \varepsilon$ (for sufficiently large n) and $\beta = 0$. We (somewhat artificially but w.l.o.g.) increase $\eta = 1/\varepsilon$ for the unweighted case (to do a more extensive local exploration). Then the runtime is $\tilde{O}(n^{0.397}/\varepsilon + \sqrt{k})$ in the unweighted and $\tilde{O}(n^{0.397} + \sqrt{k})$ in the weighted case. On unweighted graphs the approximation is $\alpha + \frac{2}{\eta} \leq 1 + 3\varepsilon \stackrel{\varepsilon' := 3\varepsilon}{=} 1 + \varepsilon'$. On weighted graphs we have $2\alpha + 1 = 3 + o(1)$. \square

The last corollary is derived from [7] Theorem 5.2, stating that SSSP can be solved *exactly* in the CLIQUE model in $\tilde{O}(n^{1/6})$ rounds. As we have shown in the proof of Theorem 4.2, we can conserve this approximation ratio (or rather lack thereof) in case of a single source, which gives us an exact algorithm for the HYBRID model.

COROLLARY 4.13 (THEOREM 1.3). *There is a HYBRID algorithm that solves SSSP exactly in $\tilde{O}(n^{2/5})$ rounds, w.h.p.*

PROOF. The SSSP algorithm from [7] can be parametrized by $\gamma = 0$, $\delta = 1/6$, $\eta = 1$ and $\alpha = \beta = 0$. \square

5 SIMULATION RESULTS FOR DIAMETER

In the distributed diameter problem all nodes must learn the parameter $D = D(G)$, which equals the maximum hop-distance $\text{hop}(u, v)$ for any pair $u, v \in V$. There are fast CLIQUE algorithms known for the *weighted* diameter problem and the *weighted* APSP problem (which obviously solves the weighted diameter problem). We aim to apply the same technique of simulating the CLIQUE on a skeleton graph of a HYBRID network to the diameter problem.

The core idea is relatively simple. If D is large, then the diameter of the skeleton will give a good approximation of the diameter of the whole graph. Therefore, computing a good approximation on the skeleton using a known CLIQUE algorithm will give a decent approximation for the HYBRID model. If, on the other hand, D is relatively small we can use the local network to compute D exactly. The following theorem formalizes the result in a more general way.

THEOREM 5.1. *Let \mathcal{A} be an (α, β) -approximation algorithm for the CLIQUE model for the weighted diameter with runtime $T_{\mathcal{A}} \in \tilde{O}(\eta n^\delta)$ for const. $\delta \geq 0, \eta \geq 1$. Let $x := \frac{2}{3+2\delta}$. Then there is an algorithm \mathcal{B} for the HYBRID model with runtime $T_{\mathcal{B}} \in \tilde{O}(\eta n^{1-x})$ that gives an $(\alpha + \frac{2}{\eta} + \frac{\beta}{T_{\mathcal{B}}})$ -approximation of the unweighted diameter D w.h.p.*

We give a brief summary of the algorithm and its proof before going into the details. As in Section 4, we first compute a skeleton graph $\mathcal{S} = (V_{\mathcal{S}}, E_{\mathcal{S}})$ of appropriate size $\Theta(n^x)$ (where $x := \frac{2}{3+2\delta}$ balances the trade-off between the runtime of the simulation and the runtime for the local exploration of G). Second, we simulate the CLIQUE algorithm \mathcal{A} on \mathcal{S} , which gives each node in $V_{\mathcal{S}}$ an approximation $\tilde{D}(\mathcal{S})$ of $D(\mathcal{S})$.

Third, we let nodes exchange information via the local model for $\eta h + 1$ rounds (where $h \in \tilde{O}(n^{1-x})$ is the maximum length of skeleton edges). This achieves two things. (I) it floods $\tilde{D}(\mathcal{S})$ to all nodes (recall that each node has a skeleton within h hops w.h.p.). (II) each node learns its largest hop distance $h_v := \max_{w \in N_v} \leq \eta h + 1$ to any node w in v 's $(\eta h + 1)$ -hop neighborhood N_v . Fourth, we apply a standard technique from [2] (given in Lemma B.2) to aggregate the maximum value out of n values $h_u, u \in V$ (one held by each node) and to announce the biggest value $\hat{h} = \max_{u \in V} h_u$ to all nodes in $\tilde{O}(1)$ rounds. Finally, all nodes can compute an approximation \tilde{D} of $D(G)$ by setting $\tilde{D} := \hat{h}$ if $\hat{h} \leq \eta h$ and $\tilde{D} := \tilde{D}(\mathcal{S}) + 2h$ else.

For the correctness, the important observation is that the diameter of the skeleton can be at most $2h$ (the maximum length of two skeleton edges) shorter than the diameter of G and we make up for that by adding $2h$ to the estimation $\tilde{D}(\mathcal{S})$. This gives us an additive error of at most $2h + \beta$, which does not impact our approximation much if we turn it into a multiplicative error, using that we know D exactly if $D \leq \eta h$. We move on to the application of Theorem 5.1.

5.1 Diameter Simulation - Algorithm and Proof

We use the same basic approach of simulating the CLIQUE model on a skeleton, thus the following algorithm is quite similar to the one of the previous section and reuses the according sub-procedures.

Algorithm 9 Diam-Simulation($\mathcal{A}, \delta, \eta$) $\triangleright \delta, \eta, \mathcal{A}$ as in Th. 5.1

```

 $x \leftarrow \frac{2}{3+2\delta}$   $\triangleright x$  optimizes overall runtime
Compute-Skeleton( $\perp, x$ )  $\triangleright$  compute  $\mathcal{S} = (V_{\mathcal{S}}, E_{\mathcal{S}}), |V_{\mathcal{S}}| \in \tilde{O}(n^x)$ 
Clique-Simulation( $\mathcal{A}, x$ )  $\triangleright$  run  $\mathcal{A}$  on  $\mathcal{S}$ 
for  $\eta h + 1$  rounds do  $\triangleright v$  learns approx.  $\tilde{D}(\mathcal{S})$  of  $D(\mathcal{S})$ 
     $v$  forwards diameter approximations via local edges
 $h_v \leftarrow \max_{u \in N_{\eta h+1}(v)} \text{hop}(v, u)$   $\triangleright$  largest hop-dist.  $v$  sees locally
 $v$  participates in aggregation for  $\hat{h} = \max_{u \in V} h_u$ 
 $v$  computes approx.  $\tilde{D}$  of  $D(G)$  with Equation (3)

```

FACT 5.2. *With the same arguments as in Lemma 4.6 the runtime of Algorithm 9 is $\tilde{O}(\eta n^{1-x})$. Algorithm 9 first computes a skeleton with the properties outlined in Fact 4.4. The simulation of a diameter algorithm \mathcal{A} makes an (α, β) -approximation $\tilde{D}(\mathcal{S})$ of $D(\mathcal{S})$ known to all skeleton nodes. The subsequent local exploration of graph G for $\eta h + 1$ rounds accomplishes two things. First, it spreads $\tilde{D}(\mathcal{S})$ to all nodes (recall that each node has a skeleton within h hops w.h.p.). Second, each node learns the largest hop distance $h_v \leq \eta h + 1$ to any node u in v 's $(\eta h + 1)$ -neighborhood. We apply the aggregation scheme of [2] (c.f., Lemma B.2) to aggregate the maximum value out of n values $h_u, u \in V$ (one held by each node) and announce the biggest value $\hat{h} = \max_{u \in V} h_u$ to all nodes in $\tilde{O}(1)$ rounds. Finally, each node computes an approximation of \tilde{D} of D as follows*

$$\tilde{D} := \begin{cases} \hat{h}, & \text{if } \hat{h} \leq \eta h \\ \tilde{D}(\mathcal{S}) + 2h, & \text{else} \end{cases} \quad (3)$$

The runtime is already accounted for in Fact 5.2. The following lemma completes the proof of Theorem 5.1.

LEMMA 5.3. *Algorithm 9 computes an $(\alpha + \frac{2}{\eta} + \frac{\beta}{T_{\mathcal{B}}})$ -approximation of the diameter D of an unweighted graph G w.h.p.*

PROOF. As summarized in Fact 5.2, each node already knows an (α, β) -approximation $\tilde{D}(\mathcal{S})$ of the diameter of \mathcal{S} and the maximum hop distance \hat{h} that any node sees in its $(\eta h + 1)$ -neighborhood. First, assume $D \leq \eta h$. By definition there is a pair of nodes $u, v \in V$ with $\text{hop}(u, v) = D$ and therefore $\tilde{D} = \hat{h} = D$ by Equation (3) and the definition of \hat{h} . We would like to point that if $\hat{h} \leq \eta h$, then all nodes can be sure that $D \leq \eta h$ and thus $\hat{h} = D$. Because if D would be larger, then $\hat{h} = \eta h + 1$ and the nodes use the diameter approximation $\tilde{D} = \tilde{D}(\mathcal{S}) + 2h$ instead, by Equation (3).

Consequently, we now have to look at the case $D > \eta h$. We analyze how good the approximation $\tilde{D} = \tilde{D}(\mathcal{S}) + 2h$ is for G . By definition we know $D(\mathcal{S}) \leq \tilde{D}(\mathcal{S}) \leq \alpha D(\mathcal{S}) + \beta$. Next, we relate the diameter of the skeleton \mathcal{S} to the diameter of G . Let P be a shortest path in G with $|P| = D$. By Lemma C.1 there are two skeleton nodes $u, v \in V_{\mathcal{S}}$ within the first h hops and the last h hops of P w.h.p. By Lemma C.2 we have $d_{\mathcal{S}}(u, v) = d(u, v) = \text{hop}(u, v)$. Consequently, $D(\mathcal{S}) \geq d(u, v) \geq |P| - 2h = D - 2h$. As the skeleton graph preserves distances (Lemma C.2) we have $D(\mathcal{S}) \leq D$. Thus

$$D \leq D(\mathcal{S}) + 2h \leq \tilde{D}(\mathcal{S}) + 2h = \tilde{D} \leq \alpha D(\mathcal{S}) + \beta + 2h \leq \alpha D + \beta + 2h.$$

Note that the computed approximation \tilde{D} in the middle is squeezed from above and below by expressions in D . Since we know that we have computed the diameter exactly if it was smaller than ηh , we

can apply the same trick as in the previous section and turn the additive approximation error into a smaller multiplicative one.

$$\tilde{D} \leq \alpha D + \beta + 2h \leq (\alpha + \frac{\beta}{D} + \frac{2h}{D})D \leq (\alpha + \frac{\beta}{\eta h} + \frac{2}{\eta})D$$

Finally we point out that we can do the local exploration for $T_B + 1$ rounds without changing the asymptotic runtime of the algorithm (since $T_B \in \tilde{\Theta}(\eta h)$ anyway). Hence we can get rid of the parameter h in our expression: $\tilde{D} \leq (\alpha + \frac{\beta}{T_B} + \frac{2}{\eta})D$. \square

5.2 Implied Results for Computing Diameter

First, we apply a dedicated diameter CLIQUE algorithm from [7]. It achieves a $(\frac{3}{2} + \epsilon, W)$ -approximation in $\tilde{O}(\frac{1}{\epsilon})$ rounds, where W is the maximum edge weight. (Note that this and the subsequent corollary prove Theorem 1.4.)

COROLLARY 5.4. *There is a HYBRID algorithm that computes a $(\frac{3}{2} + \epsilon)$ -approximation of $D(G)$ in $\tilde{O}(\frac{n^{1/3}}{\epsilon})$ rounds w.h.p.*

PROOF. The algorithm of [7] can be characterized with $\delta = 0, \eta = 1/\epsilon$ and $(\alpha, \beta) = (\frac{3}{2} + \epsilon, W)$. The runtime follows immediately from Theorem 5.1. Since G is unweighted, the maximum edge weight in the skeleton S is h . Thus, the approx. factor is

$$\alpha + \frac{2}{\eta} + \frac{\beta}{T_B} \leq \frac{3}{2} + \epsilon + 2\epsilon + \frac{W}{\eta h} \leq \frac{3}{2} + 3\epsilon + \frac{h}{h/\epsilon} = \frac{3}{2} + 4\epsilon \stackrel{\epsilon' := 4\epsilon}{=} \frac{3}{2} + \epsilon'. \quad \square$$

Another noteworthy result can be deduced from the APSP algorithm of [8] that we used before for shortest paths. It has runtime $\tilde{O}(n^{\rho+o(1)}) \subseteq \tilde{O}(n^{0.15715})$ (for n large enough) and a multiplicative error $(1+o(1))$ (recall that ρ is the distributed matrix multiplication coefficient with $\rho \leq 1 - \frac{2}{\omega} < 0.15715$). It allows for edge weights up to $2^{(n^{o(1)})}$ which is n if we choose $\log \log n / \log n$ for the $o(1)$. To keep the result concise, we omit expressing it as a function of ρ and insert the known upper bound instead.

COROLLARY 5.5. *There is a HYBRID algorithm that computes a $(1 + \epsilon)$ -approximation of $D(G)$ in $\tilde{O}(\frac{n^{0.397}}{\epsilon})$ rounds w.h.p.*

PROOF. Clearly, solving APSP on the skeleton also solves the diameter problem, hence Theorem 5.1 is applicable. We characterize the above algorithm with parameters $\delta = 0.15715, \eta = 1/\epsilon$ and $\alpha = 1 + o(1) \leq 1 + \epsilon$ (for sufficiently large n) and $\beta = 0$. Then the runtime is $\tilde{O}(n^{0.397/\epsilon})$ and the approximation parameter is $\alpha + \frac{2}{\eta} + \frac{\beta}{T_B} \leq 1 + 3\epsilon \stackrel{\epsilon' := 3\epsilon}{=} 1 + \epsilon'$. \square

6 LOWER BOUNDS FOR k -SSP

Since the lower bound of $\tilde{\Omega}(\sqrt{k})$ for k -source shortest paths is only a slight modification of the lower bound of $\tilde{\Omega}(\sqrt{n})$ for APSP given in [3], we would like to keep this subsection brief by referring to [3] for technical details. In order to keep this paper reasonably self-contained, we give the definition of the according worst case graph and the core idea behind the lower bound. Thus, knowledge of [3] is not strictly required to obtain an understanding of the proof.

The worst case graph that we use (depicted in Figure 1) is unweighted. It consists of a path of length $\Omega(n)$ with a dedicated node b at one end. Additionally, there are two sets of sources S_1, S_2 attached to the path, one to a node v_1 at distance $L \in \tilde{\Omega}(\sqrt{k})$ from b the other to a node v_2 at the opposite end of the path from b .

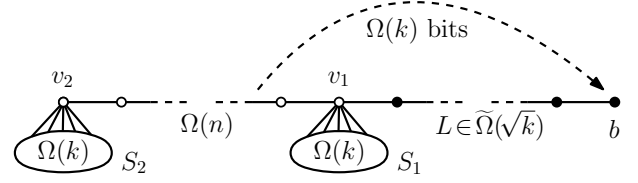


Figure 1: Worst case graph for k -SSP.

Attached means that each of the sources in the sets S_1, S_2 has an edge to the respective node v_1, v_2 .

The first insight is the following. Given a set S of sources of size $|S| = k$ then we randomly assign $|S|/2$ nodes to each set S_1 and S_2 . This creates a random state which b must necessarily learn in order to learn the distance to each node. Considered as a random variable, the state of S carries a Shannon entropy of $\tilde{\Omega}(k)$ bits that must somehow be transmitted to b since in order that b can solve the k -SSP problem. Note that the difference to [3] is $k = n$.

The second insight is that given some graph G that consists of a subgraph G' and a path with L hops attached to G' and a random variable X with entropy of $\Omega(L^2)$ bits that is initially only known to nodes in G' , then it takes $\tilde{\Omega}(L)$ rounds until the node b at the end of the path learns all bits. The intuition behind this is as follows.

Assume we are given at most $L-1$ rounds. Then any information that reaches b from G' could not have relied only local edges, since the path has L hops. Therefore, any information that ever reaches b must at some point be transmitted to a node on the path via a global edge. However, the combined capacity of all nodes on the path for receiving information is $O(L \log^2 n)$ bits. Hence this takes at least $\tilde{\Omega}(L)$ rounds. In [3] this is formalized in the technical Lemma 4.4.

Combining the two insights above already leads to a lower bound of $\tilde{\Omega}(L) = \tilde{\Omega}(\sqrt{k})$ for solving the k -source shortest path problem exactly. The third insight concerns the approximation ratio. Assume b does not know for a single node u whether $u \in S_1$ or $u \in S_2$. Then b must assume $u \in S_2$ in order to give a valid approximation (recall that approximations must be lower bounded by the true distance). But if in reality $u \in S_1$, then b 's distance approximation is off by a factor $\alpha' \in \Theta(n/\sqrt{k})$. This means that, given slightly less time than required above, b can not even give an α' -approximation. Hence, computing an α -approximation for any $\alpha \leq \alpha'$ takes $\tilde{\Omega}(\sqrt{k})$ rounds.

From the above insights we conclude Theorem 1.5 (which can be seen as corollary of its counterpart in [3]).

7 LOWER BOUND FOR DIAMETER

The second lower bound is for the diameter problem. We demonstrate that classical proofs of lower bounds for the CONGEST model based on 2-party set disjointness problems, can be adapted for lower bounds in the HYBRID model. The 2-party set disjointness problem asks whether for two given bit sequences $a, b \in \{0, 1\}^K$, there is an index $i \in [K]$, such that $a_i = b_i = 1$. If two parties (Alice and Bob) are given one of the bit sequences a, b each, then they must exchange $\Omega(K)$ bits to solve this problem [20, 32].

To obtain lower bounds in models of distributed computing one has to create a problem instance that solves the 2-party set disjointness problem by solving the problem instance (i.e., provide an appropriate reduction). Additionally, the problem instance must

have an information bottleneck between the part of the network that represents Alice and the part that represents Bob. As communication in HYBRID is conceptually different from communication in CONGEST, we need to slightly alter the existing worst case graphs and provide additional technical lemmas to demonstrate a corresponding information bottleneck for the HYBRID model. The technical tools to adapt existing proofs using this framework to the HYBRID model may also be of use for repeating the procedure for other graph problems based on the same proof strategy. The worst case graph in our proof is an adaptation of the one by [17] for the CONGEST model and we use a similar simulation argument as in [34]. Specifically, we show that $\tilde{\Omega}(n^{1/3})$ rounds are required to compute the unweighted diameter exactly and to compute a $(2-\varepsilon)$ -approximation of the weighted diameter. The lower bounds hold for randomized protocols, even if both players have access to public randomness.

More specifically, we show how to construct a family of graphs, which for some parameter k encodes set disjointness instances over a universe of size k^2 and such that for a given set disjointness input $a, b \in \{0, 1\}^{k^2}$, the graph has a slightly higher diameter if and only if a and b are disjoint. We will show that if we give a to Alice and b to Bob, Alice and Bob can together efficiently simulate a given HYBRID model algorithm for computing the diameter of the graph defined by a and b . Because by the end of the simulation Alice and Bob know the diameter of the graph (with sufficiently large probability), they have solved set disjointness and they must therefore have exchanged $\Omega(k^2)$ bits. We will see that the parts of the graph simulated by Alice and Bob respectively, can exchange only $\tilde{O}(n)$ bits in the simulated HYBRID model protocol, resulting in a lower bound of $\tilde{\Omega}(k^2/n)$ for the diameter problem (we are able to choose $k \in \Theta(n^{2/3})$).

We use the following family of graphs $\Gamma_{k,\ell,W}^{a,b} = (V, E)$ to show a bound of $\tilde{\Omega}(n^{1/3})$ rounds to compute the diameter in the HYBRID model. It is based on the construction of [17] and an example is depicted in Figure 2. We define $\Gamma_{k,\ell,W}^{a,b}$ as follows. We have four sets of nodes V_1, V_2, U_1, U_2 each of size k . These node sets form *cliques* with edges of weight W . Then the node sets V_i and U_i ($i = 1, 2$) are perfectly matched with ℓ -hop-paths of edges of weight 1. We further have two nodes \hat{v} and \hat{u} , connecting to all nodes in $V_1 \cup V_2$ and in $U_1 \cup U_2$, respectively, with edges of weight W . The nodes \hat{v}, \hat{u} are connected by a path of length and hop-length ℓ .

Let $a, b \in \{0, 1\}^{k^2}$ be two input sequences of the set disjointness problem. We represent a and b in $\Gamma_{k,\ell,W}^{a,b}$ with edges (c.f., Figure 2, red lines). More specifically, we first map each bit a_i in the sequence a to a pair in $p_i \in V_1 \times V_2$ and each bit b_j in b to a pair in $q_j \in U_1 \times U_2$. Since $|V_1 \times V_2| = |U_1 \times U_2| = k^2$, we have a pair for each bit.

Additionally, we assign bits to pairs such that the order of bits in the sequence is consistent with the matching between V_1, U_1 and V_2, U_2 . If for some $i \in [k^2]$ the bit a_i is mapped to $p_i = (v_1, v_2)$ and b_i to $q_i = (u_1, u_2)$ then v_1, u_1 and v_2, u_2 are matched with paths of length ℓ . Finally, we add an actual edge between a pair $p_i = (v_1, v_2)$ or $q_j = (u_1, u_2)$ if and only if the corresponding bit a_i or b_j in the sequence is 0 (c.f., Figure 2, red lines).

The main property of the graph $\Gamma_{k,\ell,W}^{a,b}$ is that it has low diameter if and only if a and b are disjoint. This is shown by the following

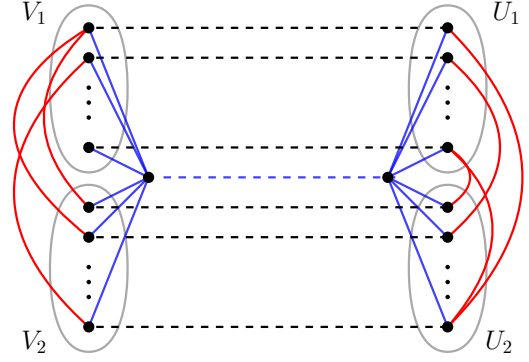


Figure 2: Graph $\Gamma_{k,\ell,W}^{a,b}$ with $|V_i| = |U_i| = k$. The V_i, U_i form *internal cliques* (gray oval shapes) with edges of weight W . Inputs $a, b \in \{0, 1\}^{k^2}$ are encoded via red edges. Dashed lines are paths with ℓ hops. Bold lines are edges of weight W . The blue network guarantees $D \leq 2W + \ell$. It is $D = W + \ell$, if and only if a, b are disjoint.

lemma. We first consider the weighted case, where we are interested in large weights $W > \ell$.

LEMMA 7.1. *Suppose $W > \ell$. Then sequences $a, b \in \{0, 1\}^{k^2}$ are disjoint, if and only if the diameter of $\Gamma_{k,\ell,W}^{a,b}$ is at most $W + 2\ell$. Else, the diameter is at least $2W + \ell$.*

PROOF. Considering the graph *without* the edges added through a, b (everything except red edges in Figure 2), the distance between any pair of nodes is at most $2W + \ell$. We believe that the respective cases are readily verifiable by the reader (the most noteworthy case is for two nodes on different paths, represented by the dashed lines in Figure 2).

Next, we show that the distance between any pair of nodes is at most $W + 2\ell$ if a, b are disjoint. Regardless of a, b , this is true for any node on the path from \hat{u} to \hat{v} (including those two) and any other node (c.f., blue path in Figure 2). The distance between such a pair is in fact at most $W + \ell$.

This is also true, independently from a, b , for any pair of nodes only from either the *top half* of the graph (defined as $V_1 \cup U_1$ and nodes on the matching paths of length ℓ running between V_1 and U_1) or the *bottom half* of the graph (defined analogously, c.f. Figure 2). More specifically, the distance between such a pair is at most $W + \ell$ due to the fact that each of the sets V_1, V_2, U_1, U_2 forms an internal clique with edges of weight W .

If a, b are disjoint, then (and only then) we can also bound the distance by $W + 2\ell$ between two nodes from opposing parts of the graph (one node in the top, the other in the bottom half). Given that a, b are disjoint, for each $i \in [k^2]$ we have $a_i = 0$ or $b_i = 0$ (or both) in the input sequences a, b . That means that for any two pairs $p_i = (v_1, v_2)$ and $q_i = (u_1, u_2)$ we have $\{v_1, v_2\} \in E$ or $\{u_1, u_2\} \in E$ by construction.

Let w^\uparrow, w^\downarrow be two nodes in the top and the bottom half (as defined above) respectively. Let $v_1 \in V_1, u_1 \in U_2$ be the two nodes with a matching path that contains w^\uparrow (w^\uparrow may be equal to v_1 or u_1), and let $v_2 \in V_2, u_2 \in U_2$ be the analogous nodes for w^\downarrow . Since

$\{v_1, v_2\} \in E$ or $\{u_1, u_2\} \in E$ the distance between w^\uparrow and w^\downarrow can be at most $W + 2\ell$.

So far we have shown that all possible combinations of two nodes have distance at most $W + 2\ell$ if a, b are *disjoint*. It remains to be shown that this is *not* the case if a, b are *not* disjoint. Then we have nodes $v_1 \in V_1, u_1 \in U_1, v_2 \in V_2, u_2 \in U_2$ such that $\{v_1, v_2\} \notin E$ and $\{u_1, u_2\} \notin E$. It is not hard to see that in order to get from v_1 to u_2 one must use at least two edges of weight W and at least one path of total weight ℓ . Hence the diameter of $\Gamma_{k,\ell,W}^{a,b}$ is at least $2W + \ell$. \square

In the unweighted case $W = 1$ we have to proceed even more carefully with our case distinction, specifically for nodes on different paths. However, we can reuse parts of the proof of the previous lemma. We show the following.

LEMMA 7.2. *Suppose $W = 1$. Then $a, b \in \{0, 1\}^{k^2}$ are disjoint, if and only if the diameter of $\Gamma_{k,\ell,1}^{a,b}$ is $\ell + 1$. Else, the diameter is $\ell + 2$.*

PROOF. As in the previous proof the maximum distance among any pair of nodes is at most $\ell + 2W = \ell + 2$. For nodes on the path from \hat{v} to \hat{u} (blue path) we have shown a distance of at most $\ell + W = \ell + 1$. The same (distance at most $\ell + W = \ell + 1$) is true for pairs of nodes both either in the top half or the bottom half (as defined in the previous proof).

A somewhat different approach is needed to show a distance of at most $\ell + 1$ between nodes w^\uparrow and w^\downarrow in the top and the bottom half respectively, if a, b are disjoint. Like before let $v_1 \in V_1, u_1 \in U_2$ be the two nodes with a matching path that contains w^\uparrow and let L^\uparrow be the according path. Let $v_2 \in V_2, u_2 \in U_2$ and L^\downarrow be the analogous objects for w^\downarrow . Let $v_1 = x_1, \dots, x_\ell = u_1$ and $v_2 = y_1, \dots, y_\ell = u_2$ be the nodes on the paths L^\uparrow and L^\downarrow ordered from left to right. Let $i, j \in [\ell]$ such that $x_i = w^\uparrow$ and $y_j = w^\downarrow$. Presume $i + j \leq \ell - 1$ then

$$\begin{aligned} d(w^\uparrow, w^\downarrow) &= d(x_i, y_j) \leq d(x_i, x_1) + d(x_1, y_1) + d(y_1, y_j) \\ &= i + 2 + j \leq \ell - 1. \end{aligned}$$

Presume $i + j \geq \ell + 1$. This implies $(\ell - i) + (\ell - j) \leq \ell - 1$. Then almost analogously

$$\begin{aligned} d(w^\uparrow, w^\downarrow) &= d(x_i, y_j) \leq d(x_i, x_\ell) + d(x_\ell, y_\ell) + d(y_\ell, y_j) \\ &= (\ell - i) + 2 + (\ell - j) \leq \ell - 1. \end{aligned}$$

Finally presume $i + j = \ell$. In this case we actually require the set disjointness. This means we have $\{x_1, y_1\} \in E$ or $\{x_\ell, y_\ell\} \in E$, making one of the distances $d(x_1, y_1)$ or $d(x_\ell, y_\ell)$ equal to $W = 1$. Assume $\{x_1, y_1\} \in E$. Then

$$\begin{aligned} d(w^\uparrow, w^\downarrow) &= d(x_i, y_j) \leq d(x_i, x_1) + d(x_1, y_1) + d(y_1, y_j) \\ &= i + 1 + j = \ell + 1. \end{aligned}$$

The case $\{x_\ell, y_\ell\} \in E$ is analogous. This covers all possible pairs of nodes, i.e., the distance between any such pair is at most $\ell + 1$ if a, b are *disjoint*. As in the previous proof, it remains to argue that there is a pair of nodes with distance strictly bigger than $\ell + 1$, if the sets are *not disjoint*. We have already shown the existence of a pair with distance $2W + \ell = \ell + 2$ in the previous proof and are therefore done. \square

We move on to the reduction of the 2-party set disjointness problem to the diameter graph problem using the graph $\Gamma_{k,\ell,W}^{a,b}$ and its properties. Assume we are given a randomized HYBRID model

algorithm \mathcal{A} that computes the diameter D of the graph. We assume that at the end of Algorithm \mathcal{A} , at least one of the nodes of the graph must know the output of Algorithm \mathcal{A} (i.e., the diameter or an approximation of the diameter).

Assume that we are given a set disjointness instance, where Alice is given set a and Bob is given set b . Then the two parties together simulate $\Gamma_{k,\ell,W}^{a,b}$ and run \mathcal{A} on the simulated graph. During the simulation the players exchange messages that go between nodes simulated by opposing parties so they can continue the correct simulation.

LEMMA 7.3. *Presume Alice knows $a \in \{0, 1\}^{k^2}$ and Bob knows $b \in \{0, 1\}^{k^2}$. Let \mathcal{A} be a randomized HYBRID model algorithm. Assume further that Alice and Bob have access to public randomness. Then the two parties can together correctly simulate up to $\lfloor \ell/2 \rfloor$ rounds of Algorithm \mathcal{A} , such that (1) the only communication between Alice and Bob stems from messages via the global network and (2) each node is simulated by at least one party throughout the whole simulation of \mathcal{A} .*

PROOF. We organize the nodes of $\Gamma_{k,\ell,W}^{a,b}$ into columns. The first column consists of $V_1 \cup V_2 \cup \{\hat{v}\}$. The last (ℓ^{th}) column consists of $U_1 \cup U_2 \cup \{\hat{u}\}$. The columns in between are made up of the inner nodes on the paths that connect the first and the last column (c.f., Figure 2). Formally, column i consists of all nodes that have hop-distance i to their closest node in the first column.

First some nomenclature. If we say that Alice (or Bob) simulates some round r of a node x , we assume that Alice knows the correct state of node x at the beginning of round r (i.e., after finishing round $r - 1$) and Alice can thus compute the messages sent by node x in round r . This does *not* necessarily require that Alice also learns all the messages x receives in round r (unless Alice also simulates round $r + 1$ of x). This is due to the fact that x can only react in round $r + 1$ or later rounds to the messages x receives in round r anyway.

Based on her knowledge of a , Alice is able to construct the subgraph of $\Gamma_{k,\ell,W}^{a,b}$ induced by the first $\ell - 1$ columns. By symmetry, Bob can construct the subgraph induced by all columns except the first. Then the simulation works as follows. In the first round Alice simulates columns $1, \dots, \ell - 1$ and Bob simulates columns $2, \dots, \ell$. With each subsequent simulation round the simulated node set shrinks by one column from the opposite side. Specifically, in round i , Alice only continues simulating columns $1, \dots, \ell - i$ and Bob $i + 1, \dots, \ell$. Since we only run \mathcal{A} for at most $\lfloor \ell/2 \rfloor$ rounds, condition (2) is satisfied.

Note that the set of nodes simulated by Alice and Bob overlaps. In order to ensure a correct simulation of the protocol, a node x that is simulated by both Alice and Bob has to be simulated consistently by the two parties. In particular, since \mathcal{A} can be a randomized algorithm, Alice and Bob have to use the same randomness in the simulation of node x . Note that they can consistently simulate the node because they have access to public randomness. In fact, if we think of the public randomness as part of the input given to the players, the whole simulation becomes deterministic.

In order to make sure that the simulation of the protocol is correct, we need to guarantee that Alice and Bob always know the correct state of all the nodes they simulate. We prove that this is true by induction on the number of rounds. For the first round,

clearly both Alice and Bob know the initial states of all the nodes they simulate and they can therefore perform the correct internal first round computations and compute the first round messages of all their nodes. For any round $r \geq 2$, we need to show that a correct simulation up to round $r - 1$ implies that Alice and Bob can learn the state at the end of round $r - 1$ (i.e., at the beginning of round r) for all the nodes they simulate in round r .

Consider some node x that sends a message to node y in round $r - 1$. W.l.o.g., assume that Alice simulates node x in round $r - 1$ and Alice thus knows the message x sends to y in round $r - 1$. If Bob also simulates node x in round $r - 1$, he also knows the message from x to y and Alice and Bob do not need to exchange any information about this message. Let us therefore assume that Bob does not simulate node x in round $r - 1$, but that Bob simulates node y in round r . In this case, Bob needs to know the message from x to y in round $r - 1$ to correctly simulate y in round r and Alice has to send the content of the message to Bob.

In order to show that Alice never needs to send information to Bob about a message that is sent over a local edge, we need to show that the following constellation can not occur. There can not be two neighbors x and y in G such that x is simulated by Alice in round $r - 1$, but x is *not* simulated by Bob in round $r - 1$, and y is simulated by Bob in round r .

In round $r - 1$ Alice simulates the nodes in columns $1, \dots, \ell - (r - 1)$ and Bob simulates the nodes in columns r, \dots, ℓ . Node x therefore needs to be in one of the columns $1, \dots, r - 1$. However, in the next round r , Bob only simulates columns $r + 1, \dots, \ell$ and therefore y needs to be in one of those columns. Therefore, there cannot be an edge between x and y in G . For a correct simulation of \mathcal{A} , Alice and Bob therefore only need to exchange information about messages sent over the global edges and thus also claim (1) of the lemma is proven. \square

Now we have the required tools to prove Theorem 1.6. Specifically, we show that computing the diameter in the HYBRID model takes $\Omega((n/\log^2 n)^{1/3})$ even for Monte Carlo algorithms with success probability $2/3$. Additionally, we want to show that the same is true for computing a $(2-\epsilon)$ -approximation for any $\epsilon > 0$ by using only weights up to $n^{1/3}$.

PROOF. Assume \mathcal{A} is a (randomized) diameter algorithm that computes the diameter of a graph with probability at least $2/3$ and takes at most $\lfloor \frac{\ell}{2} \rfloor - 1$ rounds for some $\ell \in [n]$ yet to be determined. Note that this implies that after correctly simulating \mathcal{A} for $\lfloor \frac{\ell}{2} \rfloor$ rounds every node knows the diameter of the graph.

Consider the two-player set disjointness problem and assume that Alice and Bob are given two strings $a, b \in \{0, 1\}^{k^2}$ as inputs. Then Lemma 7.3 tells us that Alice and Bob can together correctly simulate $\lfloor \frac{\ell}{2} \rfloor$ rounds Algorithm \mathcal{A} on graph $\Gamma_{k, \ell, W}^{a, b}$. This implies that Alice and Bob know the correct value of the diameter (in the unweighted case) or a $(2-\epsilon)$ -approximation of the diameter (in the weighted case) with probability at least $2/3$. By Lemmas 7.1 and Lemma 7.2 this implies that they have solved the given set disjointness instance with probability $2/3$ (more on the weighted case further below).

Due to Lemma 7.3, the only information that Alice and Bob ever exchange during the simulation are the messages sent via local

messages in the simulated HYBRID network. But according to the lower bound for set disjointness [20, 32], Alice and Bob must have exchanged $\Omega(k^2)$ bits during the whole process, since they solve the set disjointness problem with sufficiently good probability.

Specifically, this means that during the simulation $\Omega(k^2/\log n)$ messages must have been exchanged between nodes via the global network, since a message can hold at most $O(\log n)$ bits. However, all nodes combined can send (or receive) at most $O(n \log n)$ messages per round. Hence, it takes at least $\Omega(k^2/(n \log^2 n))$ rounds until \mathcal{A} is finished.

Summing up, \mathcal{A} must take at least $\Omega(\min\{\lfloor \frac{\ell}{2} \rfloor - 1, k^2/(n \log^2 n)\})$ rounds. If n is the total number of nodes of $\Gamma_{k, \ell, W}^{a, b}$, then we know from the construction that $k \cdot \ell \in \Theta(n)$. This is fulfilled by choosing $\ell \in \Theta((n/\log^2 n)^{1/3})$ and $k \in \Theta((n \log n)^{2/3})$. Then the runtime of \mathcal{A} is $\Omega((n/\log^2 n)^{1/3})$.

In the weighted case, assume that the approximation ratio of Algorithm \mathcal{A} is α . In order to guarantee that Alice and Bob can distinguish disjoint from non-disjoint instances, α has to satisfy $\alpha(W+2\ell) < (2W+\ell)$ (cf. Lemma 7.1). This means $\alpha < \frac{2W+\ell}{W+2\ell}$. We have $\frac{2W+\ell}{W+2\ell} = 2 - o(1)$, for any $W \in \omega(\ell) \subseteq \Omega(n^{1/3})$. Thus $\alpha = 2-\epsilon$ fulfills the requirement for large enough n . \square

REFERENCES

- [1] Udit Agarwal, Vijaya Ramachandran, Valerie King, and Matteo Pontecorvi. 2018. A Deterministic Distributed Algorithm for Exact Weighted All-Pairs Shortest Paths in $\tilde{O}(n^{3/2})$ Rounds. In *Proc. of the 2018 ACM Symposium on Principles of Distributed Computing (PODC)*. 199–205.
- [2] John Augustine, Mohsen Ghaffari, Robert Gmyr, Kristian Hinnenthal, Christian Scheideler, Fabian Kuhn, and Jason Li. 2019. Distributed Computation in Node-Capacitated Networks. In *Proc. 31st ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*. 69–79.
- [3] John Augustine, Kristian Hinnenthal, Fabian Kuhn, Christian Scheideler, and Philipp Schneider. 2020. Shortest Paths in a Hybrid Network Model. In *Proc. 31st ACM-SIAM Symp. on Discrete Algorithms (SODA)*. 1280–1299.
- [4] B. Awerbuch, A. V. Goldberg, M. Luby, and S. A. Plotkin. 1989. Network Decomposition and Locality in Distributed Computation. In *Proc. 30th Symp. on Found. of Computer Science (FOCS)*. 364–369.
- [5] Ruben Becker, Andreas Karrenbauer, Sebastian Krinninger, and Christoph Lenzen. 2017. Near-Optimal Approximate Shortest Paths and Transshipment in Distributed and Streaming Models. In *Proc. 31st Int. Symp. on Distributed Computing (DISC)*. 7:1–7:16.
- [6] Aaron Bernstein and Danupon Nanongkai. 2019. Distributed Exact Weighted All-Pairs Shortest Paths in Near-Linear Time. In *Proc. 51st ACM Symp. on Theory of Computing (STOC)*. 334–342.
- [7] Keren Censor-Hillel, Michal Dory, Janne H. Korhonen, and Dean Leitersdorf. 2019. Fast Approximate Shortest Paths in the Congested Clique. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*. 74–83.
- [8] Keren Censor-Hillel, Petteri Kaski, Janne H. Korhonen, Christoph Lenzen, Ami Paz, and Jukka Suomela. 2019. Algebraic methods in the congested clique. *Distributed Computing* 32, 6 (2019), 461–478.
- [9] Keren Censor-Hillel, Dean Leitersdorf, and Elia Turner. 2018. Sparse Matrix Multiplication and Triangle Listing in the Congested Clique Model. In *Proc. 22nd Int. Conf. on Principles of Distributed Systems (OPODIS)*. 4:1–4:17.
- [10] Nathan Farrington, George Porter, Sivasankar Radhakrishnan, Hamid Hajabdolali Bazzaz, Vikram Subramanya, Yeshaiah Fainman, George Papen, and Amin Vahdat. 2010. Helios: a hybrid electrical/optical switch architecture for modular data centers. In *ACM SIGCOMM 2010 Conference*. 339–350.
- [11] Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. 2012. Networks cannot compute their diameter in sublinear time. In *Proc. 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1150–1162.
- [12] François Le Gall. 2014. Powers of tensors and fast matrix multiplication. In *Symposium on Symbolic and Algebraic Computation (ISSAC '14)*.
- [13] Robert Gmyr, Kristian Hinnenthal, Christian Scheideler, and Christian Sohler. 2017. Distributed Monitoring of Network Properties: The Power of Hybrid Networks. In *Proc. 44th Int. Coll. on Automata, Languages, and Programming (ICALP)*. 137:1–137:15.
- [14] Daniel Halperin, Srikanth Kandula, Jitendra Padhye, Paramvir Bahl, and David Wetherall. 2011. Augmenting data center networks with multi-gigabit wireless links. In *Proc. of the ACM SIGCOMM 2011 Conference*. 38–49.
- [15] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. 2016. A deterministic almost-tight distributed algorithm for approximating single-source shortest paths. In *48th ACM Symposium on Theory of Computing (STOC)*. 489–498.
- [16] Stephan Holzer, David Peleg, Liam Roditty, and Roger Wattenhofer. 2014. Brief Announcement: Distributed $3/2$ -Approximation of the Diameter. In *Proc. 28th Int. Symp. on Distributed Computing (DISC)*. 562–564.
- [17] Stephan Holzer and Nathan Pinski. 2016. Approximation of Distances and Shortest Paths in the Broadcast Congest Clique. In *Proc. 19th Int. Conf. on Principles of Distributed Systems (OPODIS)*. 1–16.
- [18] Stephan Holzer and Roger Wattenhofer. 2012. Optimal distributed all pairs shortest paths and applications. In *2012 ACM Symposium on Principles of Distributed Computing (PODC)*. 355–364.
- [19] Chien-Chung Huang, Danupon Nanongkai, and Thatchaphol Saranurak. 2017. Distributed Exact Weighted All-Pairs Shortest Paths in $\tilde{O}(n^{5/4})$ Rounds. In *58th IEEE Symposium on Foundations of Computer Science (FOCS)*. 168–179.
- [20] B. Kalyanasundaram and G. Schnitger. 1992. The Probabilistic Communication Complexity of Set Intersection. *SIAM J. Discrete Math* 5, 4 (1992), 545–557.
- [21] Udit Narayana Kar and Debarshi Kumar Sanyal. 2018. An overview of device-to-device communication in cellular networks. *ICT Express* 4, 3 (2018), 203–208.
- [22] Fabian Kuhn, Yannic Maus, and Simon Weidner. 2018. Deterministic Distributed Ruling Sets of Line Graphs. In *Proc. 25th Int. Coll. on Structural Information and Communication Complexity (SIROCCO)*. 193–208.
- [23] François Le Gall. 2016. Further algebraic algorithms in the congested clique model and applications to graph-theoretic problems. In *Proc. 30th Int. Symp. on Distributed Computing (DISC)*. 57–70.
- [24] Christoph Lenzen. 2013. Optimal deterministic routing and sorting on the congested clique. In *Proc. 32nd Symp. on Principles of Distr. Comp. (PODC)*. 42–50.
- [25] Christoph Lenzen and Boaz Patt-Shamir. 2013. Fast routing table construction using small messages. In *Proc. 45th ACM Symposium on Theory of Computing (STOC)*. 381–390.
- [26] Christoph Lenzen and Boaz Patt-Shamir. 2015. Fast Partial Distance Estimation and Applications. In *2015 ACM Symposium on Principles of Distributed Computing (PODC)*. 153–162.
- [27] Christoph Lenzen and David Peleg. 2013. Efficient distributed source detection with limited bandwidth. In *2013 ACM Symposium on Principles of Distributed Computing (PODC)*. 375–382.
- [28] Nathan Linial. 1992. Locality in Distributed Graph Algorithms. *SIAM J. Comput.* 21, 1 (1992), 193–201.
- [29] Danupon Nanongkai. 2014. Distributed approximation algorithms for weighted shortest paths. In *Symp. on Theory of Computing (STOC)*. 565–573.
- [30] D. Peleg. 2000. *Distributed Computing: A Locality-Sensitive Approach*. SIAM.
- [31] David Peleg, Liam Roditty, and Elad Tal. 2012. Distributed Algorithms for Network Diameter and Girth. In *Proc. 39th Int. Coll. on Automata, Languages, and Programming (ICALP)*. 660–672.
- [32] A. A. Razborov. 1992. On the Distributional Complexity of Disjointness. *Theor. Comp. Sci.* 106 (1992), 385–390.
- [33] Michael Rossberg and Guenter Schaefer. 2011. A Survey on Automatic Configuration of Virtual Private Networks. *Computer Networks* 55, 8 (2011).
- [34] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. 2012. Distributed Verification and Hardness of Distributed Approximation. *SIAM J. Comput.* 41, 5 (jan 2012), 1235–1265.
- [35] Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. 1995. Chernoff-Hoeffding Bounds for Applications with Limited Independence. *Journal on Discrete Mathematics* 8, 2 (1995), 223–250.
- [36] Anis Tell, Wale Babalola, George Kalebiala, and Krishna Chinta. 2018. SD-WAN: A Modern Hybrid-WAN to Enable Digital Transformation for Businesses. *IDC White Paper* (April 2018).
- [37] Jeffrey D. Ullman and Mihalis Yannakakis. 1991. High-Probability Parallel Transitive-Closure Algorithms. *Journal on Computing* 20, 1 (1991), 100–125.
- [38] Salil P. Vadhan. 2012. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science* 7, 1–3 (2012), 1–336. <https://doi.org/10.1561/04000000010>

A SOME BASIC PROBABILISTIC CONCEPTS

LEMMA A.1 (CHERNOFF BOUND). *We use the following forms of Chernoff bounds in our proofs:*

$$\mathbb{P}(X > (1+\delta)\mu_H) \leq \exp\left(-\frac{\delta\mu_H}{3}\right),$$

with $X = \sum_{i=1}^n X_i$ for i.i.d. random variables $X_i \in \{0, 1\}$ and $\mathbb{E}(X) \leq \mu_H$ and $\delta \geq 1$. Similarly, for $\mathbb{E}(X) \geq \mu_L$ and $0 \leq \delta \leq 1$ we have

$$\mathbb{P}(X < (1-\delta)\mu_L) \leq \exp\left(-\frac{\delta^2\mu_L}{2}\right).$$

REMARK A.2. *Note that the first inequality even holds if we have k -wise independence among the random variables X_i for $k \geq \lceil \mu_H \delta \rceil$ (c.f., [35] Theorem 2, note that a substitution $\mu_H := (1+\varepsilon)\mathbb{E}(X)$ generalizes the result for any $\mu_H \geq \mathbb{E}(X)$).*

LEMMA A.3 (UNION BOUND). *Let E_1, \dots, E_k be events, each taking place w.h.p. If $k \leq p(n)$ for a polynomial p then $E := \bigcap_{i=1}^k E_i$ also takes place w.h.p.*

PROOF. Let $d := \deg(p)+1$. Then there is an $n_0 \geq 0$ such that $p(n) \leq n^d$ for all $n \geq n_0$. Let $n_1, \dots, n_k \in \mathbb{N}$ such that for all $i \in \{1, \dots, k\}$ we have $\mathbb{P}(\overline{E_i}) \leq \frac{1}{n_i^c}$ for some (yet unspecified) $c > 0$. With Boole's Inequality (union bound):

$$\mathbb{P}(\overline{E}) = \mathbb{P}\left(\bigcup_{i=1}^k \overline{E_i}\right) \leq \sum_{i=1}^k \mathbb{P}(\overline{E_i}) \leq \sum_{i=1}^k \frac{1}{n_i^c} \leq \frac{p(n)}{n^c} \leq \frac{1}{n^{c-d}}$$

for all $n \geq n'_0 := \max(n_0, \dots, n_k)$. Let $c' > 0$ be arbitrary. We choose $c \geq c'+d$. Then $\mathbb{P}(\overline{E}) \leq \frac{1}{n^{c'}}$ for all $n \geq n'_0$. \square

REMARK A.4. *If a constant number of events is involved we use the above lemma without explicitly mentioning it. It is possible to use the lemma in a nested fashion as long as the number of applications is polynomial in n .*

B DISSEMINATION PROTOCOLS IN HYBRID

In this section we give a quick overview about the known protocols to disseminate information in HYBRID networks that we will use in this paper. The first one is the so-called token dissemination protocol of [3], which solves the problem to broadcast k tokens from arbitrary sources, whereas each source has at most ℓ tokens.

LEMMA B.1. (Theorem 2.1 in [3]) *Assume there are k tokens of size $O(\log n)$ bits in the network and each node has at most ℓ tokens. There is an algorithm that solves the token dissemination problem in the HYBRID model on connected graphs in $\tilde{O}(\sqrt{k+\ell})$ rounds, w.h.p.*

Another protocol that we use frequently, solves the following aggregation problem. Given a subset of nodes $A \subseteq V$ that have some information $x_v, v \in A$, all nodes must learn the result $f(x_v \mid v \in A)$, where f is an aggregate distributive function.¹⁰ An example is computing a sum or the maximum value. The algorithm by [2] solves this problem, relying only on the global network.

LEMMA B.2. (Theorem 2.1 in [2]) *The aggregation problem can be solved in $O(\log n)$ in the NCC model.*

¹⁰ An aggregate function f maps a multiset S of inputs to some $f(S)$. For an aggregate distributive function f there is an aggregate function g such that for any partition of $S = S_1 \oplus \dots \oplus S_k$ it is $f(S) = g(f(S_1), \dots, f(S_k))$

C TECHNIQUES FOR SKELETON GRAPHS

By sampling nodes from G with prob. $\frac{1}{x}$ there will be a sampled node on some shortest u - v -path P at least every $\tilde{O}(x)$ hops for any $u, v \in V$ w.h.p. We formalize that in the following lemma. This was first observed by [37].

LEMMA C.1 (FROM [3]). *Let $M \subseteq V$ be a subset of nodes of $G = (V, E)$ by sampling each node with probability at least $\frac{1}{x}$. Then there is a constant $\xi > 0$, such that for any $u, v \in V$ with $\text{hop}(u, v) \geq \xi x \ln n$, there is at least one shortest path P from u to v , such that any sub-path Q of P with at least $\xi x \ln n$ nodes contains a node in M w.h.p.*

PROOF. Let $u, v \in V$ with $\text{hop}(u, v) \geq \xi x \ln n$. Fix a shortest u - v -path $P_{u,v}$ and let Q be a sub-path of $P_{u,v}$ with at least $\xi x \ln n$ nodes. Let $X_{u,v}$ be the random number of marked nodes on Q . Then we have $\mathbb{E}(X_{u,v}) \geq \frac{|Q|}{x} \geq \xi \ln n$. Let $c > 0$ be arbitrary. We use a Chernoff bound:

$$\mathbb{P}(X_{u,v} < \frac{\xi \ln n}{2}) \leq \exp\left(-\frac{\xi \ln n}{8}\right) \stackrel{\xi \geq 8c}{\leq} \frac{1}{n^c}.$$

Thus we have $X_{u,v} \geq 1$ w.h.p. for constant $\xi \geq \max(8c, 2/\ln n)$. Therefore the claim holds w.h.p. for the pair u, v . We claim that w.h.p. the event $X_{u,v} \geq 1$ occurs for all pairs $u, v \in V$ and for all sub-paths Q of $P_{u,v}$ longer than $\xi x \ln n$ hops, for at least one shortest path $P_{u,v}$ from u to v . There are at most n^2 many pairs $u, v \in V$. Moreover we can select at most n sub-paths Q of P that do not fully contain any other selected sub-path. Hence the claim follows with the union bound given in Lemma A.3. \square

The above lemma implies additional properties of skeletons that we use (sometimes implicitly) in this paper.

LEMMA C.2. *Let $\mathcal{S} = (V_{\mathcal{S}}, E_{\mathcal{S}})$ be a skeleton graph of a connected graph G with n nodes by sampling each node of G to $V_{\mathcal{S}}$ with prob. at least $\frac{1}{x}$. The edges of \mathcal{S} are $E_{\mathcal{S}} = \{\{u, v\} \mid u, v \in V_{\mathcal{S}}, \text{hop}(u, v) \leq h\}$ (whereas $h := \xi x \ln n$ with parameter ξ as in Lemma C.1) and edge weights $d_h(u, v)$ for $\{u, v\} \in E_{\mathcal{S}}$. Then \mathcal{S} is connected w.h.p. Furthermore, for any $u, v \in \mathcal{S}$ we have $d_G(u, v) = d_{\mathcal{S}}(u, v)$ w.h.p.*

PROOF. Let $u, v \in \mathcal{S}$ and let P be a shortest u - v -path in G . By Lemma C.1 there is a skeleton node at least every h hops on P w.h.p. Hence, consecutive skeleton nodes must be connected by edges in $E_{\mathcal{S}}$ by definition. These form a u - v -path Q in \mathcal{S} . Therefore u, v are connected in G .

Let $u = v_1, \dots, v_k = v$ be the set of consecutive skeleton nodes on P . Then $\text{hop}(v_i, v_{i+1}) \leq h$ and $\{v_i, v_{i+1}\} \in E_{\mathcal{S}}$ due to Lemma C.1. Since P is a shortest path, all its sub-paths must be shortest paths as well. Hence we have $d_h(v_i, v_{i+1}) = d(v_i, v_{i+1})$. Therefore $d_{\mathcal{S}}(u, v) \leq \sum_{i=1}^{k-1} d(v_i, v_{i+1}) = d_G(u, v)$. Since edges of \mathcal{S} have at least the weight of the path in G they represent, we also have $d_{\mathcal{S}}(u, v) \geq d_G(u, v)$. \square

D k -WISE INDEPENDENT HASH FUNCTIONS

We will use families of k -wise independent hash functions, which are defined as follows.

DEFINITION D.1. *For finite sets A, B , let $\mathcal{H} := \{h : A \rightarrow B\}$ be a family of hash functions. Then \mathcal{H} is called k -wise independent if for a random function $h \in \mathcal{H}$ and for any k distinct keys $a_1, \dots, a_k \in A$*

we have that $h(a_1), \dots, h(a_k) \in B$ are independent and uniformly distributed random variables in B .

From literature we know that such a family of hash functions exists in the following form (c.f., [38]).

LEMMA D.2. For $A := \{0, 1\}^a$ and $B := \{0, 1\}^b$, there is a family of k -wise independent hash functions $\mathcal{H} := \{h : A \rightarrow B\}$ such that selecting a function from \mathcal{H} requires $k \cdot \max(a, b)$ random bits and computing $h(x)$ for any $x \in A$ can be done in $\text{poly}(a, b, k)$ time.

REMARK D.3. We can use a random member $h \in \mathcal{H}$ of a family as described in Definition D.1 to limit the number of messages any node receives in a given round. If all nodes send at most $O(\log n)$ messages to targets that are determined using that hash function h with a distinct key for each message, then any node receives at most $O(\log n)$ messages per round w.h.p. The details are given by the following lemma.

LEMMA D.4. Presume some algorithm takes at most $p(n)$ rounds for some polynomial p . Moreover, presume that each round, every node sends at most $\sigma \in \Theta(\log n)$ messages via global edges to k -wise

independent, uniformly random targets in V . There is a $k \in \Theta(\log n)$ and a $\rho \in \Theta(\log n)$ such that for sufficiently large n , in every round, every node in V receives at most ρ messages per round w.h.p.

PROOF. Let $X_{v,r}$ be the (random) number of messages v receives in round r . Node v is targeted by some token for some node from V with probability $\frac{1}{n}$ (for simplicity we assume that v can send the token to itself, in reality it can just keep it). Since each node sends at most σ tokens per round we have $\mathbb{E}(X_{v,r}) \leq \sigma \cdot n \cdot \frac{1}{n} = \sigma$. By definition we have $\sigma \geq \xi \ln n$ for some constant ξ and large enough n . Let $c > 0$ be arbitrary. We choose $\rho \geq (1 + \frac{3c}{\xi})\sigma$ and $k \geq \lceil \frac{3c}{\xi}\sigma \rceil$. We employ the Chernoff bound given in Lemma A.1, which holds for k -wise independent random variables as outlined in Remark A.2.

$$\mathbb{P}(X_{v,r} > \rho) \leq \mathbb{P}(X_{v,r} > (1 + \frac{3c}{\xi})\sigma) \leq \exp\left(-\frac{3\xi c \ln n}{3\xi}\right) = \frac{1}{n^c}.$$

In accordance with the union bound given in Lemma A.3 the event $\bigcap_{\substack{v \in V \\ r \leq p(n)}} (X_{v,r} \leq \rho)$ takes place w.h.p. \square