Faster near-exact weighted diameter and radius in quantum CONGEST

June 30, 2026

Abstract

We give randomized algorithms for weighted diameter and radius in the quantum CONGEST model. On a connected, undirected n-vertex network of unweighted diameter D and positive polynomially bounded integer edge weights, every vertex outputs a one-sided $(1 + o(1))$ overestimate, with high probability, in

$$\widetilde{O}\Big(\min\{n,\, n^{5/6} + \sqrt{n}\, D\}\Big)$$

rounds. Wu and Yao previously obtained $\widetilde{O}(\min\{n^{9/10}D^{3/10}, n\})$ rounds for the same problem. For every fixed $\delta \in (0, 1/2)$, the new bound is polynomially smaller throughout $D \le n^{1/2-\delta}$; when $D = \Theta(\sqrt{n})$, both bounds are $\widetilde{O}(n)$. The algorithm samples candidate vertices independently of the landmarks used for shortest-path reconstruction. A candidate that is not a landmark is attached as a virtual root to a landmark shortcut graph, and nested distributed quantum optimization selects a good candidate and a good trial.

Note. This paper was generated by an internal harness powered by GPT-5.5, with prompting by Aidan Bai. Cleanup and final preparation were performed by Codex, powered by GPT-5.

1 Introduction

Let $G = (V, E)$ be the communication network, let $w : E \to \mathbb{N}_{>0}$, and write $d_w(u, v)$ for weighted shortest-path distance. The weighted diameter and radius are

$$\Delta_w = \max_{u,v \in V} d_w(u, v), \qquad R_w = \min_{u \in V} \max_{v \in V} d_w(u, v).$$

We study the number of synchronous rounds needed to approximate these values when each edge carries $O(\log n)$ qubits per round. The parameter D denotes the diameter of the unweighted communication graph.

There is a nearly linear classical exact baseline. Bernstein and Nanongkai's randomized algorithm computes exact weighted all-pairs shortest paths in $\widetilde{O}(n)$ rounds with high probability; their verification and restart procedure gives a Las Vegas implementation with $\widetilde{O}(n)$ expected rounds [BN19]. Thus exact weighted diameter and radius also have a $\widetilde{O}(n)$ baseline. At the other end, for every fixed $\varepsilon \in (0, 1/2)$, Abboud, Censor-Hillel, and Khoury proved an

$\Omega(n/\log^3 n)$ classical lower bound, applying also to randomized algorithms with constant success probability, for $(3/2-\varepsilon)$-approximating diameter or radius, even on sparse unweighted networks [ACHK16]. Exact weighted single-source shortest paths can be computed, with high probability, by the randomized Broadcast CONGEST algorithm of Chechik and Mukhtar in $\widetilde{O}(\sqrt{n}\,D^{1/4}+D)$ rounds [CM20]; its polynomially bounded nonnegative-integer weight setting includes ours. One arbitrary source gives a factor-2 baseline for diameter and radius, but it does not give a near-exact answer. At a coarser guarantee, Ancona, Censor-Hillel, Dalirrooyfard, Efron, and Vassilevska Williams give $(2+\varepsilon)$-approximations to weighted undirected diameter, radius, and all eccentricities with high probability in $\widetilde{O}(\sqrt{n}+D)$ rounds for $\varepsilon = 1/\operatorname{polylog}(n)$ [ACHD+21]. This differs in approximation strength from the task here.

The quantum CONGEST model was introduced by Elkin, Klauck, Nanongkai, and Pandurangan [EKNP14]. For unweighted networks, Le Gall and Magniez obtained, with high probability, an exact-diameter algorithm using $\widetilde{O}(\sqrt{nD})$ rounds and developed the distributed quantum optimization framework behind it [GM18]. The framework of van Apeldoorn and de Vos gives a bounded-error exact-radius algorithm at the same scale and treats distributed quantum queries systematically [vAdV22]. These results do not directly extend to near-exact weighted distances: reconstructing a weighted shortest path requires bounded-hop distance estimates and a landmark overlay, rather than breadth-first distances alone. Stronger output tasks do not share the same speedup: Wang, Wu, and Yao prove $\Omega(n/\log n)$ quantum lower bounds for exact unweighted all-pairs shortest paths and $\Omega(n/\log^2 n)$ for all exact eccentricities, even at constant hop diameter [WWY21].

Wu and Yao studied precisely this weighted problem [WY22]. Combining distributed quantum optimization with Nanongkai's weighted shortest-path framework [Nan14], they obtained one-sided $(1+o(1))$ approximations with high probability in

$$\widetilde{O}\Big(\min\{n^{9/10}D^{3/10}, n\}\Big)$$

rounds. They also proved that, for every fixed $\varepsilon \in (0, 1/2]$, any quantum CONGEST algorithm that succeeds with probability at least $11/12$ and returns a $(3/2-\varepsilon)$-approximation requires $\Omega(n^{2/3}/\log^2 n)$ rounds, even on graphs with $D = \Theta(\log n)$. Thus their lower bound leaves a polynomial gap already at $D = \Theta(\log n)$.

Our result improves the Wu–Yao upper bound.

Theorem 1.1 (Main theorem)**.** *Let G be a connected, undirected n-vertex network with unweighted diameter D, and suppose that every edge has a positive integer weight at most $n^{O(1)}$. There are randomized quantum CONGEST algorithms for weighted diameter and radius. With high probability, the diameter algorithm makes every vertex output $\widehat{\Delta}$, and the radius algorithm makes every vertex output $\widehat{R}$, satisfying*

$$\Delta_w \le \widehat{\Delta} \le (1+o(1))\Delta_w, \qquad R_w \le \widehat{R} \le (1+o(1))R_w.$$

Each algorithm uses

$$\widetilde{O}\Big(\min\{n,\, n^{5/6}+\sqrt{n}\,D\}\Big)$$

rounds. Running the two algorithms consecutively gives both values within the same asymptotic bound.

For $D = n^\delta$ and every fixed $0 \le \delta < 1/2$, the exponent in Theorem 1.1 is strictly smaller than the exponent in the Wu–Yao bound. More generally, its displayed algebraic bound is $o(n)$ whenever $D = o(\sqrt{n})$, whereas the previous expression reaches its $\widetilde{O}(n)$ cap at $D = n^{1/3}$. There is no claimed asymptotic improvement at $D = \Theta(\sqrt{n})$, and the known $\widetilde{\Omega}(n^{2/3})$ lower bound still leaves a polynomial gap from our $\widetilde{O}(n^{5/6})$ bound for small D.

Proof idea. Wu and Yao use a random sample both as the domain over which an eccentricity is optimized and as the landmark set that hits long shortest paths. We separate these roles. In each of n trials, a candidate set C_i has expected size c, while an independent landmark set L_i has expected size r. Every fixed vertex appears in $\Theta(c)$ candidate sets with high probability, so an outer quantum threshold search needs only $\widetilde{O}(\sqrt{n/c})$ trial evaluations to encounter a diameter endpoint or a radius center. The landmark sample can remain smaller and is used only for distance reconstruction.

This separation means that a candidate $s \in C_i$ need not lie in L_i. Bounded-hop estimates from s to the landmarks define a star from a virtual source to the existing landmark shortcut graph. We show that Nanongkai's scaled bounded-hop routine can simulate these star edges by local initialization at their landmark endpoints; no candidate-specific shortcut graph is needed. The path-hitting property then reconstructs every s-to-v distance within a factor $(1+\varepsilon)^2$.

One trial has preprocessing and evaluation cost

$$\widetilde{O}\left(D + \frac{n}{r} + c + rk + \sqrt{c}\left(D + \frac{rD}{k} + r\right)\right), \tag{1}$$

where k is the shortcut parameter. For $D \le n^{1/3}$, the choices $r = n^{1/3}$, $c = n^{2/3}$, and $k = \Theta(n^{1/6}\sqrt{D})$ give $\widetilde{O}(n^{5/6})$ rounds after the outer search. For $n^{1/3} < D \le n^{1/2}$, the choices $r = n^{1/3}$, $c = n^{4/3}/D^2$, and $k = \Theta(r)$ give $\widetilde{O}(\sqrt{n}D)$. We also verify that the bounded-error inner optimization can be amplified and uncomputed to form a clean oracle for the outer quantum search.

The output is only the value of the diameter or radius, not a witnessing pair or center. The algorithms are randomized and succeed with high probability; the result is restricted to connected undirected graphs with positive polynomially bounded integer weights. No algorithmic or novelty claim beyond this input and output regime is intended.

Section 2 records the model and imported primitives. Section 3 proves the virtual-root reduction, Section 4 analyzes separated sampling and coherent trial evaluation, and Section 5 completes the outer search and parameter optimization.

2 Model and source results

Every vertex has a distinct $O(\log n)$-bit identifier, knows n, and knows the weights of its incident edges. Local computation and memory are unbounded. In one quantum CONGEST round, each endpoint of an edge may send $O(\log n)$ qubits to the other endpoint. Classical protocols are valid quantum protocols. All displayed outputs have $O(\log n)$ bits under the polynomial weight bound. An event holds *with high probability* if its failure probability is n^{-A} for an arbitrarily large fixed constant A, after constants in the algorithm are increased. The notation $\widetilde{O}$ suppresses polylogarithmic factors in n.

For an integer $h \geq 0$, let $d_w^h(u, v)$ be the minimum weight of a u–v walk using at most h edges, with value $+\infty$ if no such walk exists. Because weights are positive, a minimum walk is a simple path. The eccentricity of s is

$$e_w(s) = \max_{v \in V} d_w(s, v),$$

so $\Delta_w = \max_s e_w(s)$ and $R_w = \min_s e_w(s)$.

Fix one canonical shortest path P_{uv} for every ordered pair (u, v), breaking ties by the sequence of vertex identifiers. A set $L \subseteq V$ is *h-hitting* if every block of h consecutive vertices on every canonical path contains a vertex of L. Thus, on a canonical path, the first and last landmark are fewer than h edges from the respective endpoints, and two consecutive landmarks are at most h edges apart.

We use the following consequences of Nanongkai's weighted shortest-path framework. Wu and Yao give the quantum-compatible formulation and account for the pipelined broadcasts [WY22, Sec. 3 and App. A]; in the arXiv v2 numbering, the underlying bounded-hop and shortcut statements are Nanongkai's Theorems 3.2, 3.6, 3.10, and 4.5 and Lemmas 3.7 and 4.6 [Nan14].

Lemma 2.1 (Bounded-hop distances and shortcuts)**.** *Let $S \subseteq V$, let $h \geq 1$, and let $0 < \varepsilon \leq 1$. In*

$$\widetilde{O}(D + h/\varepsilon + |S|)$$

rounds, all vertices can obtain values $\widetilde{d}^h(s, v)$, for $s \in S$ and $v \in V$, such that, with high probability,

$$d_w(s, v) \leq \widetilde{d}^h(s, v) \leq (1 + \varepsilon) d_w^h(s, v). \tag{2}$$

Let H be the graph on $m = |S|$ vertices whose finite edge weights are these bounded-hop estimates between sources. For every integer $1 \leq k < m$, a k-shortcut graph Q can be constructed and embedded in $\widetilde{O}(D + mk)$ rounds. It satisfies

$$d_Q(x, y) = d_H(x, y) \quad \text{for all } x, y \in S, \qquad \mathrm{SPD}(Q) < \frac{4m}{k}, \tag{3}$$

where SPD *is the largest number of edges on a minimum-hop shortest path. Every shortcut edge represents an H-path and hence has weight at least the true distance between its endpoints in G.*

We will apply a threshold form of distributed quantum optimization. It is important that the threshold need not be known to the algorithm.

Lemma 2.2 (Distributed quantum optimization)**.** *Suppose a leader can prepare a search state on a finite domain X, and a distributed unitary evaluation oracle and its inverse use T rounds. If an unknown threshold M is met by total search probability at least ρ, a marked x can be found, with high probability, using $\widetilde{O}(T/\sqrt{\rho})$ rounds, in addition to one-time initialization. The same statement holds for a lower threshold. Exact maximum or minimum over a domain of size q uses $\widetilde{O}(\sqrt{q}\,T)$ rounds.*

Proof. The threshold statement is Wu and Yao's Lemma 3.1, which is the unknown-threshold form of the distributed optimization framework of Le Gall and Magniez [WY22, GM18]. Replacing values by their order-reversing encodings gives the lower-threshold and minimum forms. Error reduction contributes only logarithmic factors. □

Finally, exact weighted all-pairs shortest paths are available in $\widetilde{O}(n)$ rounds [BN19]. A BFS-tree convergecast of the local eccentricities and a broadcast then give exact diameter or radius. We use this as a fallback. We also elect a leader and compute a BFS tree in $O(D)$ rounds. If D_0 is the leader's unweighted eccentricity, then $D_0 \le D \le 2D_0$; hence the algorithms need no exact value of D as input. Using $2D_0$ in the parameter choices changes only constants, and below we write D for this estimate.

3 A virtual root for the landmark overlay

Let L be a landmark set of size m, and suppose that $\widetilde{d}^h(x,v)$ is available for every $x \in L \cup \{s\}$ and $v \in V$. Form the landmark graph H_L with edge weights $\widetilde{d}^h(x,y)$, and let Q_L be a k-shortcut graph for H_L. For each landmark x, define

$$\alpha_s(x) = \widetilde{d}^h(s,x).$$

An infinite value is simply omitted. The root-initialized overlay distance is

$$\delta_{L,s}(u) = \min_{x \in L}\big(\alpha_s(x) + d_{Q_L}(x,u)\big). \tag{4}$$

Lemma 3.1 (Virtual-root overlay)**.** *Let $1 \le k < m$. From the endpoint labels $\alpha_s(x)$ and the embedded graph Q_L, the network can compute and disseminate values $\widetilde{D}_{L,s}(u)$ satisfying*

$$\delta_{L,s}(u) \le \widetilde{D}_{L,s}(u) \le (1+\varepsilon)\delta_{L,s}(u) \qquad (u \in L) \tag{5}$$

in

$$\widetilde{O}\left(D + \frac{mD}{\varepsilon k} + m\right) \tag{6}$$

rounds, with high probability. The construction uses no communication edge incident to s other than the original network edges and does not rebuild the shortcut graph.

Proof. Introduce a formal vertex σ and, for every finite $\alpha_s(x)$, a star edge (σ, x) of that weight. If $\alpha_s(x) = 0$ (which occurs when $s = x \in L$), we regard x as an additional copy of the source rather than as the endpoint of a positive-weight edge. Call the resulting weighted graph A_s. Its restriction to L is Q_L, and $d_{A_s}(\sigma,u) = \delta_{L,s}(u)$.

By (3), after its first star edge a shortest σ–u path can be chosen to have fewer than $4m/k$ further edges. Consequently every such distance has a shortest realization with at most

$$h_Q = \left\lceil \frac{4m}{k} \right\rceil + 1$$

edges. Apply Nanongkai's scaled light-weight h_Q-hop SSSP routine to A_s. The only point not covered verbatim by a physical source in the overlay is the first star edge. At scale j, landmark x locally forms exactly the rounded weight that this edge would have in the routine,

$$\left\lceil \frac{2h_Q\alpha_s(x)}{\varepsilon 2^j} \right\rceil,$$

and schedules the initial label/message that it would receive from σ. Infinite labels schedule nothing, and a zero label is present at time zero. After this initialization every relaxation and every message is an ordinary operation on an embedded edge of Q_L. Thus the simulated execution is identical to the execution on A_s after the silent first hop.

The rounding proof for light-weight h_Q-hop SSSP gives the one-sided guarantee (5). The polynomial weight bound requires only $O(\log n)$ scales and $O(\log n)$-bit labels. Simulating the h_Q/ε overlay rounds through the BFS tree and then pipelining the m answers costs

$$\widetilde{O}(Dh_Q/\varepsilon + m) = \widetilde{O}\left(D + \frac{mD}{\varepsilon k} + m\right),$$

as claimed. This adapts the light-weight SSSP routine of Theorem 3.2 and the overlay simulation of Lemma 4.6 in Nanongkai's arXiv v2; Theorem 3.10 supplies the shortcut-diameter bound. The local injection above accounts explicitly for the one additional virtual hop. □

The overlay labels reconstruct distances to all network vertices.

Lemma 3.2 (Rooted reconstruction)**.** *Suppose L is h-hitting and that (2) holds for $L \cup \{s\}$. Let Q_L and $\widetilde{D}_{L,s}$ be as above, set $\widetilde{D}_{L,s}(s) = 0$, and define*

$$\widetilde{d}_{L,s}(v) = \min_{u \in L \cup \{s\}} \big(\widetilde{D}_{L,s}(u) + \widetilde{d}^{h}(u, v)\big). \tag{7}$$

Then, simultaneously for all $v \in V$,

$$d_w(s, v) \le \widetilde{d}_{L,s}(v) \le (1+\varepsilon)^2 d_w(s, v). \tag{8}$$

Proof. Every finite bounded-hop estimate is at least the true distance. Every edge of Q_L represents an H_L-path, so it too is at least the true distance between its endpoints. The virtual-root routine returns the weight of an approximate route in this graph. Each term minimized in (7) is therefore the weight of a sequence from s to v, with every segment no shorter than its true endpoint distance. The triangle inequality proves the lower bound.

Fix v and inspect the canonical shortest path P_{sv}. If it has at most h edges, choosing $u = s$ in (7) gives an upper bound of $(1+\varepsilon)d_w(s, v)$. Otherwise list the landmarks on P_{sv} in their path order and let x and y be the first and last. Since every block of h consecutive vertices is hit, the prefix s–x, the suffix y–v, and the segment between every two consecutive listed landmarks have at most h edges. Their bounded-hop estimates form an H_L-walk, and hence

$$\alpha_s(x) + d_{Q_L}(x, y) = \alpha_s(x) + d_{H_L}(x, y) \le (1+\varepsilon) d_w(s, y).$$

The virtual-root guarantee and the bounded-hop estimate on the final segment therefore imply

$$\begin{aligned} \widetilde{d}_{L,s}(v) &\le \widetilde{D}_{L,s}(y) + \widetilde{d}^{h}(y, v) \\ &\le (1+\varepsilon)^2 d_w(s, y) + (1+\varepsilon) d_w(y, v) \\ &\le (1+\varepsilon)^2 d_w(s, v). \end{aligned}$$

□

4 Separated sampling and trial evaluation

Fix sufficiently large absolute constants a and b. For each trial $i \in [n]$, include every vertex in L_i independently with probability r/n, and independently include every vertex in C_i with probability c/n. We use

$$h = \left\lceil \frac{bn \log n}{r} \right\rceil. \tag{9}$$

For the asymptotic parameter choices below, $a \log n \le r \le c \le n$ and $h \le n$ once n exceeds an absolute constant. Smaller instances use the exact fallback.

Lemma 4.1 (Simultaneous sampling event)**.** *With high probability, all of the following hold simultaneously:*

(i) $r/2 \le |L_i| \le 2r$ *and* $c/2 \le |C_i| \le 2c$ *for every trial;*

(ii) every L_i *is* h*-hitting; and*

(iii) every vertex belongs to between $c/2$ *and* $2c$ *candidate sets.*

Proof. The three relevant binomial expectations are r, c, and c, respectively. Chernoff bounds, followed by a union bound over the n trials and n vertices, prove (i) and (iii) when a is sufficiently large.

For (ii), a canonical path has at most n blocks of h consecutive vertices, so all n^2 canonical paths have at most n^3 such blocks. For a fixed block B and trial i,

$$\Pr[L_i \cap B = \varnothing] \le (1 - r/n)^h \le \exp(-rh/n) \le n^{-b}.$$

There are n trials, so the union bound is at most n^{4-b}. □

Condition on this event for the moment, and write $m_i = |L_i|$. We choose a nominal shortcut parameter $k \le r/16$. Hence, for all sufficiently large n, $1 \le k < m_i$ in every good trial. This constant slack is necessary: choosing $k = r$ would be invalid whenever the realized landmark sample has fewer than r vertices.

For trial i, the network performs the following reversible preprocessing. It computes the bounded-hop estimates from $L_i \cup C_i$, orders the candidate identifiers at the leader and pads the list to length $2c$, and constructs the k-shortcut graph on L_i. A dummy list entry has value 0 for a maximum and $+\infty$ for a minimum. Given a genuine candidate s, the virtual-root routine and (7) give

$$\widetilde{e}_i(s) = \max_{v \in V} \widetilde{d}_{L_i,s}(v).$$

The maximum is convergecast to the leader. Define

$$F_\Delta(i) = \max_{s \in C_i} \widetilde{e}_i(s), \qquad F_R(i) = \min_{s \in C_i} \widetilde{e}_i(s). \tag{10}$$

Trials failing the explicit size tests are assigned the same extreme dummy values. This makes the evaluation circuit total even outside the event of Theorem 4.1.

Lemma 4.2 (One trial)**.** *On the event of Theorem 4.1, every candidate satisfies*

$$e_w(s) \le \widetilde{e}_i(s) \le (1+\varepsilon)^2 e_w(s). \tag{11}$$

After preprocessing, a candidate value can be evaluated in

$$T_{\mathrm{cand}} = \widetilde{O}\left(D + \frac{rD}{k} + r\right) \tag{12}$$

rounds. Either trial value in (10) can be computed by an inner quantum optimizer, including preprocessing, in

$$T_{\mathrm{trial}} = \widetilde{O}\left(D + \frac{n}{r} + c + rk + \sqrt{c}\left(D + \frac{rD}{k} + r\right)\right) \tag{13}$$

rounds.

Proof. The estimate (11) follows by maximizing (8) over v. Since $m_i = \Theta(r)$, Theorem 3.1 costs the quantity in (12); the final convergecast is absorbed by its D term.

The hop-distance computation has $|L_i \cup C_i| = O(c)$ sources and, by (9), costs

$$\widetilde{O}(D + h/\varepsilon + c) = \widetilde{O}(D + n/r + c).$$

Ordering and broadcasting the candidate list costs $O(D + c)$. Constructing the shortcut graph costs $\widetilde{O}(D + rk)$. Exact quantum maximum or minimum over the padded list of $2c$ entries takes $\widetilde{O}(\sqrt{c}\,T_{\mathrm{cand}})$ rounds by Theorem 2.2. Adding these contributions proves (13). □

4.1 The nested search is a coherent oracle

The outer search queries a trial index in superposition, so a probabilistic statement about a measured inner optimizer is not enough. We record the standard amplification and uncomputation argument in the present setting.

At the start, fix all classical coins: the membership bits defining all L_i, C_i, and the delay and rounding coins for every bounded-hop and overlay instance that can be queried. Give each classical subroutine failure probability at most n^{-A-4} before taking a union bound. There are only polynomially many trials, candidates, and subroutines, so with high probability one event $\mathcal{G}$ makes every estimate used by every possible query obey its one-sided guarantee. Conditioned on $\mathcal{G}$, preprocessing and candidate evaluation are fixed reversible circuits controlled by the trial and candidate registers.

Let B_i consist of reversible trial preprocessing followed by $q = \Theta(\log n)$ coherent, measurement-free repetitions of the inner optimizer and a reversible majority computation on their value strings. All branches are padded to the same worst-case number of rounds. The repetitions use deferred measurement. For either mode and every i, amplification makes the probability that B_i writes the wrong value at most

$$\eta = n^{-K},$$

where K is an arbitrarily large fixed constant. On a fresh zeroed work register, use the circuit

$$B_i^\dagger \ \mathrm{XOR}_{\mathrm{value}\to\mathrm{answer}} \ B_i. \tag{14}$$

Its action differs in Euclidean norm by at most $2\sqrt{\eta}$ from the ideal XOR oracle for $F_\Delta(i)$ or $F_R(i)$. The same bound holds when i is in superposition: the controlled-i sectors are orthogonal.

Allocate a fresh zeroed work block for every forward or inverse oracle call of the outer optimizer. This avoids making any claim about (14) on dirty work states. If the ideal outer algorithm uses Q_{out} queries, a hybrid argument changes its final state by at most $2Q_{\mathrm{out}}\sqrt{\eta}$. Here $Q_{\mathrm{out}} = \widetilde{O}(\sqrt{n/c}) \le \widetilde{O}(\sqrt{n})$, so choosing K sufficiently large makes this error inverse-polynomial. The repetitions, reversible sorting, padding, and uncomputation add only polylogarithmic round factors. Thus (13) is a valid coherent evaluation cost for the outer search. This explicit composition is the only place where unbounded local memory is used materially.

5 Outer optimization and parameter choices

Let z_Δ be a vertex of eccentricity Δ_w. By Theorem 4.1, it belongs to at least $c/2$ candidate sets. In each corresponding trial,

$$F_\Delta(i) \ge \widetilde{e}_i(z_\Delta) \ge \Delta_w,$$

whereas every good trial satisfies

$$F_\Delta(i) \le (1+\varepsilon)^2 \Delta_w.$$

Thus an unknown upper-threshold search with $\rho = c/(2n)$ finds a value in this interval. For the radius, let z_R be a center. At least $c/2$ trials satisfy

$$F_R(i) \le \widetilde{e}_i(z_R) \le (1+\varepsilon)^2 R_w,$$

and every good trial has $F_R(i) \ge R_w$. The lower-threshold form gives the same factor $\widetilde{O}(\sqrt{n/c})$. Bad-size trial sentinels cannot meet the relevant threshold.

Combining this factor with (13) gives

$$T(r,c,k) = \widetilde{O}\left[\sqrt{\frac{n}{c}}\left(D + \frac{n}{r} + c + rk + \sqrt{c}\left(D + \frac{rD}{k} + r\right)\right)\right]. \tag{15}$$

We now choose the parameters. Integer parts and the constant factors needed for $k \le r/16$ are suppressed in the displays. They affect neither (15) nor the sample event.

If $D \le n^{1/3}$, take

$$r = n^{1/3}, \qquad c = n^{2/3}, \qquad k = \Theta(n^{1/6}\sqrt{D}), \quad k \le r/16. \tag{16}$$

The factor outside the parentheses in (15) is $n^{1/6}$. The preprocessing terms obey

$$D + n/r + c + rk = O(n^{2/3}),$$

and

$$\sqrt{c}\,(D + rD/k + r) = n^{1/3}\left(D + O(n^{1/6}\sqrt{D}) + n^{1/3}\right) = O(n^{2/3}).$$

Hence $T = \widetilde{O}(n^{5/6})$.

If $n^{1/3} < D \le n^{1/2}$, take

$$r = n^{1/3}, \qquad c = \frac{n^{4/3}}{D^2}, \qquad k = \Theta(r), \quad k \le r/16. \tag{17}$$

Now $r \le c \le n^{2/3}$ and the outer factor is $n^{-1/6}D$. Moreover,

$$D + n/r + c + rk = O(n^{2/3})$$

and

$$\sqrt{c}\,(D + rD/k + r) = \frac{n^{2/3}}{D}\big(O(D) + n^{1/3}\big) = O(n^{2/3}).$$

Therefore $T = \widetilde{O}(\sqrt{n}D)$. For $D > n^{1/2}$, the exact APSP fallback uses $\widetilde{O}(n)$ rounds. It is also available in every regime, which yields the minimum with n in Theorem 1.1.

Take $\varepsilon = 1/\log n$. On the intersection of the sampling event in Theorem 4.1 and the classical subroutine event $\mathcal{G}$ from Section 4.1, the preceding threshold argument returns values in

$$[\Delta_w, (1+\varepsilon)^2\Delta_w] \quad \text{and} \quad [R_w, (1+\varepsilon)^2 R_w].$$

The amplified quantum procedures fail with inverse-polynomial probability by Section 4.1. The leader broadcasts the selected value in $O(D)$ rounds, which is absorbed in every displayed bound. Since $(1+\varepsilon)^2 = 1 + o(1)$, the approximation and round bounds in Theorem 1.1 follow.

References

[ACHD+21] Bertie Ancona, Keren Censor-Hillel, Mina Dalirrooyfard, Yuval Efron, and Virginia Vassilevska Williams. Distributed distance approximation. In *24th International Conference on Principles of Distributed Systems (OPODIS 2020)*, volume 184 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 30:1–30:17. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021.

[ACHK16] Amir Abboud, Keren Censor-Hillel, and Seri Khoury. Near-linear lower bounds for distributed distance computations, even in sparse networks. In *Distributed Computing—30th International Symposium (DISC 2016)*, volume 9888 of *Lecture Notes in Computer Science*, pages 29–42. Springer, 2016.

[BN19] Aaron Bernstein and Danupon Nanongkai. Distributed exact weighted all-pairs shortest paths in near-linear time. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*, pages 334–342. ACM, 2019.

[CM20] Shiri Chechik and Doron Mukhtar. Single-source shortest paths in the CONGEST model with improved bound. In *PODC '20: Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 464–473. ACM, 2020.

[EKNP14] Michael Elkin, Hartmut Klauck, Danupon Nanongkai, and Gopal Pandurangan. Can quantum communication speed up distributed computation? In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing (PODC 2014)*, pages 166–175. ACM, 2014.

[GM18] François Le Gall and Frédéric Magniez. Sublinear-time quantum computation of the diameter in CONGEST networks. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing (PODC 2018)*, pages 337–346. ACM, 2018.

[Nan14] Danupon Nanongkai. Distributed approximation algorithms for weighted shortest paths. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC 2014)*, pages 565–573. ACM, 2014. Full version: https://arxiv.org/abs/1403.5171v2.

[vAdV22] Joran van Apeldoorn and Tijn de Vos. A framework for distributed quantum queries in the CONGEST model. In *PODC '22: ACM Symposium on Principles of Distributed Computing*, pages 109–119. ACM, 2022.

[WWY21] Changsheng Wang, Xudong Wu, and Penghui Yao. Complexity of eccentricities and all-pairs shortest paths in the quantum CONGEST model. *SPIN*, 11(3):2140007, 2021.

[WY22] Xudong Wu and Penghui Yao. Quantum complexity of weighted diameter and radius in CONGEST networks. In *PODC '22: ACM Symposium on Principles of Distributed Computing*, pages 120–130. ACM, 2022. Full version: https://arxiv.org/abs/2206.02767v2.