Proofs of Statements

EC.1. Proofs from Section 4.1.

Proof of Thm. 6. Note that, by virtue of how the comparison condition is structured, in each application of the quantum search algorithm there is always at least one "good" element. Thus, the proof and the running time analysis given by Durr and Hoyer (1996) for the case $\epsilon = 0$ (i.e., exact minimum finding) also work for $\epsilon > 0$, using the same argument given by (Durr and Hoyer 1996, Sect. 3) for the case in which the values of $f(x)$ are not distinct. □

EC.2. Proofs from Section 5.1.

Proof of Prop. 1. We can rely on the construction of Grover and Rudolph (2002), that consists in a sequence of controlled rotations. The construction can be understood on a binary tree with m leaves, representing the elements of b, and where each inner node has value equal to the sum of the squares of the children nodes. Each inner node requires a controlled rotation with an angle determined by the square root of the ratio between the two children nodes. Assuming that m is a power of 2 for simplicity, the tree has $\log m$ levels and the construction when b is dense ($d = m$) requires $\sum_{j=0}^{\log m} 2^j = O(m)$ controlled rotations. Notice that at each inner node, a controlled rotation is necessary only if both children have nonzero value. If only one child has nonzero value the operation requires at most a controlled X, and if both children have value zero no operation takes place. If $d < m$, there can be at most d nodes at each level that require a controlled rotation, and in fact the deepest level of the binary tree such that all nodes contain nonzero value is $\lfloor \log d \rfloor$. We need $O(d)$ controlled rotations up to this level of the tree, and for each of these nodes we may need at most $O(\log m)$ susequent operations, yielding the total gate complexity $\tilde{O}(d)$. □

EC.3. Proofs from Section 5.2.

Proof of Prop. 3. After line 3, using Prop. 2 we are in the state:

$$\frac{1}{2}(1+\alpha_k)|0\rangle \otimes |k\rangle + \frac{1}{2}(1-\alpha_k)|1\rangle \otimes |k\rangle + \frac{1}{2}(|0\rangle - |1\rangle) \otimes \sum_{\substack{j=0 \\ j\neq k}}^{2^n-1} \alpha_j |j\rangle.$$

We now apply amplitude estimation to the state $|0\rangle \otimes |k\rangle$ to determine the magnitude of $\frac{1}{2}(1+\alpha_k)$. The exact phase that must be estimated by the phase estimation portion of the algorithm is the number θ such that:

$$\sin \pi\theta = \frac{1}{2}(1+\alpha_k).$$

Suppose $\alpha_k \geq -\epsilon$. Then $\sin \pi\theta \geq \frac{1}{2}(1-\epsilon)$, implying, by monotonicity of $\sin^{-1}$ over its domain:

$$\theta > \frac{\sin^{-1}\left(\frac{1}{2}(1-\epsilon)\right)}{\pi} \geq \frac{\frac{\pi}{6} + \frac{2}{\sqrt{3}}\left(\frac{1}{2}(1-\epsilon) - \frac{1}{2}\right)}{\pi} \geq \frac{\frac{\pi}{6} - \frac{\epsilon}{\sqrt{3}}}{\pi} \geq \frac{1}{6} - \frac{\epsilon}{\sqrt{3}\pi},$$

where for the second inequality we have used the Taylor expansion of $\sin^{-1}(x)$ at $x = \frac{1}{2}$:

$$\sin^{-1}(x) \approx \frac{\pi}{6} + \frac{2}{\sqrt{3}}(x - \frac{1}{2}) + \frac{2\sqrt{3}}{9}(x - \frac{1}{2})^2.$$

Rather than θ, we obtain an approximation $\tilde{\theta}$ up to a certain precision. By Thm. 3, using $\left\lceil \log \frac{\sqrt{3}\pi}{\epsilon} \right\rceil + 2$ qubits of precision, then

$$|\theta - \tilde{\theta}| \leq \frac{\epsilon}{\sqrt{3}\pi} \tag{EC.1}$$

with probability at least 3/4. Then we must have $\tilde{\theta} \geq \frac{1}{6} - \frac{2\epsilon}{\sqrt{3}\pi}$ with probability at least 3/4. In this case, the algorithm returns 1 (recall that if the first bit of the expansion is 1, i.e., $0.a > \frac{1}{2}$, we must take the complement $1 - 0.a$ because we do not know the sign of eigenvalue on which phase estimation is applied, see (Brassard et al. 2002) for details).

Now suppose the algorithm returns 1, implying $\tilde{\theta} \geq \frac{1}{6} - \frac{2\epsilon}{\sqrt{3}\pi}$. By (EC.1) we must have $\theta \geq \tilde{\theta} - \frac{\epsilon}{\sqrt{3}\pi} \geq \frac{1}{6} - \frac{3\epsilon}{\sqrt{3}\pi}$. Thus,

$$\begin{aligned}\alpha_k &= 2\sin \pi\theta - 1 \geq 2\sin(\frac{\pi}{6} - \sqrt{3}\epsilon) - 1 \geq 2\left(\frac{1}{2}\cos(-\sqrt{3}\epsilon) + \frac{\sqrt{3}}{2}\sin(-\sqrt{3}\epsilon)\right) - 1 \\ &\geq 2\left(\frac{1}{2}(1 - \frac{2\sqrt{3}}{\pi}\epsilon) - \frac{3\epsilon}{\pi}\right) - 1 \geq -\frac{\epsilon}{\pi}(3 + 2\sqrt{3}) \geq -2\epsilon.\end{aligned}$$

The remaining part of the proposition's statement follows immediately from the first part.

Regarding the gate complexity, amplitude estimation with $O(\log \frac{1}{\epsilon})$ digits of precision requires $O(\frac{1}{\epsilon})$ calls to U, controlled-U, or controlled-$U^\dagger$, and the reflection circuits of the Grover iterator (which can be implemented with $O(q)$ basic gates and auxiliary qubits). The inverse quantum Fourier transform on $O(\log \frac{1}{\epsilon})$ qubits can be implemented with $O(\log^2 \frac{1}{\epsilon})$ basic gates; finally, the controlled unitary $|0\rangle\langle 0| \otimes U' + |1\rangle\langle 1| \otimes I^{\otimes q}$ to construct $|0\rangle \otimes |k\rangle$, as well as the final bitwise comparison, can be implemented with $O(q)$ gates. □

Proof of Prop. 4. The proof is similar to Prop. 3, and we use the same symbols and terminology. We apply amplitude estimation to the state $|0\rangle \otimes |k\rangle$ to determine the magnitude of $\frac{1}{2}(1 + \alpha_k)$. Suppose $\alpha_k \leq -\epsilon$. Then $\sin \pi\theta \leq \frac{1}{2}(1-\epsilon)$, implying:

$$\begin{aligned}\theta &\leq \frac{\sin^{-1}\left(\frac{1}{2}(1-\epsilon)\right)}{\pi} \leq \frac{\frac{\pi}{6} + \frac{2}{\sqrt{3}}\left(\frac{1}{2}(1-\epsilon) - \frac{1}{2}\right) + \frac{8}{9\sqrt{3}}\left(\frac{1}{2}(1-\epsilon) - \frac{1}{2}\right)^2}{\pi} \\ &\leq \frac{\frac{\pi}{6} - \frac{\epsilon}{\sqrt{3}} + \frac{4\epsilon^2}{9\sqrt{3}}}{\pi} \leq \frac{1}{6} - \frac{7\epsilon}{9\sqrt{3}\pi},\end{aligned}$$

Algorithm 11 Sign estimation routine: SIGNESTNFP(U, k, ϵ).

1: **Input:** state preparation unitary U on q qubits (and its controlled version) with $U|0^q\rangle = \sum_{j=0}^{2^q-1} \alpha_j |j\rangle$ and α_j real up to a global phase factor, index $k \in \{0, \dots, 2^q - 1\}$, precision ϵ.

2: **Output:** 0 if $\alpha_k \leq -\epsilon$, with probability at least 3/4.

3: Let V map $|0^q\rangle$ to $|k\rangle$. Compute $|\psi\rangle = \text{INTERFERE}(U, V)$.

4: Apply amplitude estimation to the state $|0\rangle \otimes |k\rangle$ with $\left\lceil \log \frac{9\sqrt{3}\pi}{\epsilon} \right\rceil + 2$ qubits of accuracy; let $|a\rangle$ be the bitstring produced by the phase estimation portion of amplitude estimation.

5: If $0.a < \frac{1}{2}$, return 1 if $0.a > \frac{1}{6} - \frac{2\epsilon}{3\sqrt{3}\pi}$, 0 otherwise; if $0.a \geq \frac{1}{2}$, return 1 if $1 - 0.a > \frac{1}{6} - \frac{2\epsilon}{3\sqrt{3}\pi}$, 0 otherwise.

where we used $\epsilon \leq \frac{1}{2}$, and for the second inequality we have used the Taylor expansion of $\sin^{-1}(x)$ at $x = \frac{1}{2}$:

$$\sin^{-1}(x) \approx \frac{\pi}{6} + \frac{2}{\sqrt{3}}(x - \frac{1}{2}) + \frac{2\sqrt{3}}{9}(x - \frac{1}{2})^2 + \frac{8\sqrt{3}}{27}(x - \frac{1}{2})^3,$$

coupled with the fact that the third-order term is negative at $\frac{1}{2}(1-\epsilon)$. Rather than θ, we obtain an approximation $\tilde{\theta}$ up to a certain precision. By Thm. 3, using $\left\lceil \log \frac{9\sqrt{3}\pi}{\epsilon} \right\rceil + 2$ qubits of precision, then

$$|\theta - \tilde{\theta}| \leq \frac{\epsilon}{9\sqrt{3}\pi} \tag{EC.2}$$

with probability at least 3/4. Then we must have $\tilde{\theta} \leq \frac{1}{6} - \frac{2\epsilon}{3\sqrt{3}\pi}$ with probability at least 3/4. In this case, the algorithm returns 0.

Now suppose the algorithm returns 0, implying $\tilde{\theta} \leq \frac{1}{6} - \frac{2\epsilon}{3\sqrt{3}\pi}$. By (EC.2) we must have so that $\theta \leq \tilde{\theta} + \frac{\epsilon}{9\sqrt{3}\pi} \leq \frac{1}{6} + \frac{5\epsilon}{9\sqrt{3}\pi}$. Thus,

$$\begin{aligned}\alpha_k &= 2\sin\pi\theta - 1 \leq 2\sin(\frac{\pi}{6} + \frac{5\epsilon}{9\sqrt{3}}) - 1 \leq 2\left(\frac{1}{2}\cos\frac{5\epsilon}{9\sqrt{3}} + \frac{\sqrt{3}}{2}\sin\frac{5\epsilon}{9\sqrt{3}}\right) - 1 \\ &\leq \frac{5}{9\sqrt{3}}\epsilon \leq \frac{\epsilon}{3}.\end{aligned}$$

The remaining part of the proposition's statement follows immediately from the first part, and the running time analysis is the same as in Prop. 3. □

EC.4. Proofs from Section 5.3.

Proof of Prop. 5. Let us analyze REDCOST(A_B, A_k, c, ϵ). The QLSA encodes the solution to:

$$\begin{pmatrix} A_B & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} A_k \\ c_k \end{pmatrix} \tag{EC.3}$$

in a state $|\psi\rangle$ that is guaranteed to be $\frac{\epsilon}{10\sqrt{2}}$-close to the exact normalized solution $|(A_B^{-1}A_k, c_k)\rangle$. Call this state $|(\tilde{x}, \tilde{y})\rangle$, where $\tilde{x}, \tilde{y}$ correspond to the (approximate) solution of (EC.3).

On line 4 we apply the unitary $U_c^\dagger$, obtaining $U_c^\dagger|(\tilde{x},\tilde{y})\rangle$. We are now interested in tracking the value of the coefficient of the basis state $|0^{\lceil \log m+1\rceil}\rangle$, which is the input to the SignEstNFN routine on line 4 of CanEnter. This coefficient is equal to:

$$\langle 0^{\lceil \log m+1\rceil}|U_c^\dagger|(\tilde{x},\tilde{y})\rangle\rangle = \langle(-c_B,1)|(\tilde{x},\tilde{y})\rangle,$$

because by definition $\langle 0^{\lceil \log m+1\rceil}|U_c^\dagger = (U_c|0^{\lceil \log m+1\rceil}\rangle)^\dagger = \langle(-c_B,1)|$. Furthermore,

$$\langle(-c_B,1)|(\tilde{x},\tilde{y})\rangle = \frac{1}{\|(-c_B,1)\|}\left(\tilde{y} - \sum_{j=1}^m c_{B(j)}\tilde{x}_j\right).$$

Again by definition, $\tilde{y}$ is approximately equal to $c_k/\|(A_B^{-1}A_k,c_k)\|$ whereas $\sum_{j=1}^m c_{B(j)}\tilde{x}_j$ is approximately equal to $c_B^\top A_B^{-1}A_k/\|(A_B^{-1}A_k,c_k)\|$. The total error is $\frac{\epsilon}{10\sqrt{2}}$, hence, recalling that $\|(-c_B,1)\| = \sqrt{2}$, we have:

$$\left|\frac{1}{\sqrt{2}}\left(\tilde{y} - \sum_{j=1}^m c_{B(j)}\tilde{x}_j\right) - \frac{1}{\sqrt{2}}\frac{\bar{c}_k}{\|(A_B^{-1}A_k,c_k)\|}\right| \leq \frac{\epsilon}{10\sqrt{2}}.$$

This concludes the proof. □

Proof of Prop. 6. CanEnter relies on Algorithms 3 (SignEstNFN) and 4 (RedCost).

On line 4 of CanEnter, we apply SignEstNFN to the unitary that implements RedCost, with $|0^{\lceil \log m+1\rceil}\rangle$ as the target basis state and precision $\frac{11\epsilon}{10\sqrt{2}}$. Suppose this returns 0. By Prop. 3 and Prop. 5, we have:

$$-\frac{11\epsilon}{10\sqrt{2}} > \frac{1}{\sqrt{2}}\left(\tilde{y} - \sum_{j=1}^m c_{B(j)}\tilde{x}_j\right) \geq \frac{1}{\sqrt{2}}\frac{\bar{c}_k}{\|(A_B^{-1}A_k,c_k)\|} - \frac{\epsilon}{10\sqrt{2}}.$$

This implies:

$$\frac{\bar{c}_k}{\|(A_B^{-1}A_k,c_k)\|} < -\epsilon.$$

The above discussion guarantee that if the return value of the CanEnter is 1, then $\bar{c}_k < -\epsilon\|(A_B^{-1}A_k,c_k)\|$, with probability at least 3/4, as desired. The gate complexity is easily calculated: SignEst makes $O(\frac{1}{\epsilon})$ calls to RedCost and requires an additional $O(\log m + \log^2\frac{1}{\epsilon})$ gates. RedCost requires one call to the QLSA and additional $O(m)$ gates for U_c. Thus, the total gate complexity is $O(\frac{1}{\epsilon}T_{\text{LS}}(A_B,A_k,\frac{\epsilon}{2}) + m + \log^2\frac{1}{\epsilon})$. □

Proof of Thm. 9. FindColumn relies on three subroutines: U_{rhs}, CanEnter, and quantum search, as described in Thm. 4. U_{rhs} can be constructed by repeatedly applying the procedure of Prop. 1 controlled on the register containing the column index k; the total gate complexity is $\tilde{O}(d_c n)$. By Prop. 6, if the routine CanEnter returns 1 then the reduced cost of column A_k is $< -\epsilon\|(A_B^{-1}A_k,c_k)\|$ with respect to the rescaled data, with bounded probability. We can boost this probability with repeated applications and a majority vote to make it as close to 1 as desired.

Notice that CanEnter is applied to the rescaled data. In terms of the original data, the condition on the reduced cost becomes $\bar{c}_k < -\epsilon\|(A_B^{-1}A_k, \frac{c_k}{\|c_B\|})\|$, as claimed in the theorem statement (the rescaling of A_k and A_B^{-1} cancel out).

Finally, we apply the quantum search algorithm (Thm. 4) using CanEnter as the target function. The expected number of iterations before success is $O(\sqrt{n})$; the gate complexity is therefore $\tilde{O}(\frac{\sqrt{n}}{\epsilon}(T_{\text{LS}}(A_B, A_k, \frac{\epsilon}{2}) + m))$. By Thm. 7, $T_{\text{LS}}(A_B, A_k, \frac{\epsilon}{2})$ is $\tilde{O}(\kappa d d_c n + \kappa d^2 m)$, because P_{A_B} has gate complexity $\tilde{O}(dm)$ and P_b is the routine U_{rhs}, which has gate complexity $\tilde{O}(d_c n)$. Thus, we obtain a total gate complexity of $\tilde{O}(\frac{\kappa d\sqrt{n}}{\epsilon}(d_c n + dm))$, as claimed. □

Proof of Thm. 10. The main ingredient of the proof is to analyze the subroutine SteepestEdgeCompare, Alg. 7, which is used to run the quantum minimum finding algorithm of Thm. 6.

At Steps 3-4 we compute norm estimates using Thm. 8. (In a practical implementation, the estimates at Step 4 could be derived from those at Step 3, with properly adjusted error tolerances.) This requires time $\tilde{O}(\frac{1}{\epsilon}T_{\text{LS}}(A_B, A_j, O(\epsilon)))$.

Recall that $c_{\max} := \max_{\ell \in N} c_\ell$. At Steps 5-6 we define two unitaries whose first building block is RedCost. By Prop. 5, RedCost encodes an approximation of $\frac{\bar{c}_j}{\|(A_B^{-1}A_j, c_j)\|}, \frac{\bar{c}_k}{\|(A_B^{-1}A_k, c_k)\|}$ with additive error at most $\frac{\epsilon}{10}$ in the amplitude of the all-zero basis states; let α_0, β_0 be these amplitudes. We then multiply α_0 by $\frac{\tilde{d}_j}{\tilde{e}_j c_{\max}}$ and β_0 by $\frac{\tilde{d}_k}{\tilde{e}_k c_{\max}}$. Note that these terms are ≤ 1 thanks to the factor $c_{\max}$ at the denominator, and multiplying the all-zero basis state by such a coefficient is simply an application of a unitary with that coefficient in the top-left corner. It is known that such a unitary can be efficiently constructed with high precision with cost $\tilde{O}(1)$; for an explicit construction, see, e.g., (Gilyén et al. 2019, Lemma 48).

We then apply Interfere and SignEstNFN, to estimate the sign of $\frac{1}{2}(\beta_0 - \alpha_0)$ with error $\frac{\epsilon}{8c_{\max}}$. Suppose $\frac{\bar{c}_k}{\|A_B^{-1}A_k\|} < (1-\epsilon)\frac{\bar{c}_j}{\|A_B^{-1}A_j\|} - \epsilon$. We have:

$$\beta_0 \leq \left(\frac{\bar{c}_k}{\|(A_B^{-1}A_k, c_k)\|} + \frac{\epsilon}{10}\right)\frac{\|(A_B^{-1}A_k, c_k)\|}{c_{\max}\|A_B^{-1}A_k\|}(1+\frac{\epsilon}{2}) < (1-\frac{\epsilon}{2})\frac{\bar{c}_j}{c_{\max}\|A_B^{-1}A_j\|} - \frac{\epsilon}{c_{\max}} + \frac{\epsilon}{5c_{\max}} \leq \alpha_0 - \frac{\epsilon}{4c_{\max}}.$$

By Prop. 3, this implies that SignEstNFN returns 0. Thus, applying the quantum minimum finding algorithm of Thm. 6 returns an approximate minimizer in $O(\sqrt{n})$ iterations, i.e., a column k such that:

$$\frac{\bar{c}_k}{\|A_B^{-1}A_k\|} \leq (1+\epsilon)\frac{\bar{c}_h}{\|A_B^{-1}A_j\|} + \epsilon,$$

where h is the arg min as defined in the theorem statement.

The running time is obtained by multiplying the gate complexity of the SteepestEdgeCompare subroutine by the number of iterations $O(\sqrt{n})$. In each call to SteepestEdgeCompare, we perform a constant number of norm estimations using Thm. 8, we apply RedCost and SignEstNFN with precision $O(\frac{\epsilon}{c_{\max}})$. By Prop. 3 and Prop. 5, the running time is $\tilde{O}(\frac{\kappa c_{\max} d\sqrt{n}}{\epsilon}(d_c n + dm))$.

□

EC.5. Proofs from Section 5.5.

Proof of Prop. 7. The algorithm works as follows. We apply the unitary U_{rhs} that maps $|0^{\lceil \log n \rceil}\rangle \otimes |0^{\lceil \log m \rceil}\rangle \rightarrow \sum_{k \in N} \frac{\|A_k\|}{\|A_N\|_F} (|k\rangle \otimes |A_k\rangle)$; we call the first register, that loops over $k \in N$, the "column register". This unitary can be constructed with $\tilde{O}(d_c n)$ gates, using Prop. 1. (The time to classically compute the norms can be ignored, as this only needs to be done once, and its time complexity is less than the total complexity of the algorithm.) We then apply the QLSA of (Childs et al. 2017, Thm. 3), using U_{rhs} as the oracle P_b, P_{A_B} as the oracle for the constraint matrix, and precision $\frac{\epsilon}{2n}$. As a result, conditioned on the column register being $|k\rangle$, we obtain $|A_B^{-1} A_k\rangle$ in the output register of the QLSA algorithm. Following Childs et al. (2017), there exists an auxiliary register that has value $|0^r\rangle$ with amplitude $\frac{1}{\alpha}\|\tilde{A}_B^{-1}|A_k\rangle\| = \frac{1}{\alpha\|A_k\|}\|\tilde{A}_B^{-1}A_k\|$, where α is a known number with $\alpha = O(\kappa\sqrt{\log(n\kappa/\epsilon)})$, and $\tilde{A}_B^{-1}$ is an operator that is $\frac{\epsilon}{2n}$-close to A_B^{-1} in the operator norm. Since this is true for all k, the probability of obtaining $|0^r\rangle$ in the auxiliary register is:

$$\sum_{k \in N} \frac{\|A_k\|^2}{\|A_N\|_F^2} \frac{1}{\alpha^2\|A_k\|^2} \|\tilde{A}_B^{-1} A_k\|^2 = \frac{\|\tilde{A}_B^{-1} A_N\|_F^2}{\alpha^2 \|A_N\|_F^2}. \tag{EC.4}$$

Using Thm. 5 to estimate this probability, amplitude estimation with precision $\frac{\epsilon}{4\pi\alpha^2}$ yields an estimate $\tilde{a}$ of (EC.4) with error $\pm(\frac{\epsilon}{4\alpha^2} + \frac{\epsilon^2}{16\alpha^4})$, see (Brassard et al. 2002, Thm. 12). Our estimate for $\|A_B^{-1}A_N\|_F^2$ is $\rho := \tilde{a}\alpha^2\|A_N\|_F^2$. We have:

$$\begin{aligned}\frac{\rho}{\|A_B^{-1}A_N\|_F^2} &\leq \frac{\alpha^2\|A_N\|_F^2}{\|A_B^{-1}A_N\|_F^2}\left(\frac{\|\tilde{A}_B^{-1}A_N\|_F^2}{\alpha^2\|A_N\|_F^2} + \frac{\epsilon}{4\alpha^2} + \frac{\epsilon^2}{16\alpha^4}\right) \leq 1 + \frac{\epsilon}{2} + \frac{\alpha^2\|A_N\|_F^2}{\|A_B^{-1}A_N\|_F^2}\left(\frac{\epsilon}{4\alpha^2} + \frac{\epsilon^2}{16\alpha^4}\right) \\ &\leq 1 + \frac{\epsilon}{2} + \alpha^2\left(\frac{\epsilon}{4\alpha^2} + \frac{\epsilon^2}{16\alpha^4}\right) \leq 1 + \epsilon,\end{aligned}$$

where we have used the fact that $\frac{\|A_N\|_F^2}{\|A_B^{-1}A_N\|_F^2} \leq 1$ because the smallest singular value of A_B^{-1} is 1, and the fact that we can assume $\epsilon \leq \alpha^2$ so that $\frac{\epsilon^2}{16\alpha^2} \leq \frac{\epsilon}{16}$. A similar calculation yields the lower bound, yielding the desired approximation. Regarding the running time, amplitude estimation with precision $\frac{\epsilon}{4\pi\alpha^2}$ requires $\frac{4\pi\alpha^2}{\epsilon} = \tilde{O}(\frac{\kappa^2}{\epsilon})$ applications of the QLSA. The running time for the QLSA is $\tilde{O}(\kappa d_c n + \kappa d^2 m)$, because P_{A_B} has gate complexity $\tilde{O}(dm)$ and P_b is the routine U_{rhs}, which has gate complexity $\tilde{O}(d_c n)$. Thus, we obtain a total running time of $\tilde{O}(\frac{\kappa^2}{\epsilon}(\kappa d_c n + \kappa d^2 m))$. □

EC.6. Proofs from Section 5.6.

Proof of Prop. 8. Let $|\tilde{x}\rangle$ be the state produced by the QLSA, where $\tilde{x}$ is an approximate solution to the linear system; by Thm. 7, we have $\|\tilde{x} - \frac{A_B^{-1}A_k}{\|A_B^{-1}A_k\|}\| \leq \frac{\delta}{10}$.

Suppose the algorithm returns 1; this implies that the index ℓ is not found. Then if all algorithms are successful, the routine SignEstNFN$^+(U_{\text{LS}}, \ell, \frac{9\delta}{10})$ must have returned 0 for all ℓ. by Prop. 3, $\tilde{x}_\ell < \frac{9\delta}{10}$, thus:

$$\frac{(A_B^{-1}A_k)_\ell}{\|A_B^{-1}A_k\|} \leq \tilde{x}_\ell + \frac{\delta}{10} < \frac{9\delta}{10} + \frac{\delta}{10} = \delta,$$

which implies $(A_B^{-1}A_k)_\ell < \delta\|A_B^{-1}A_k\|$. Thus, if IsUnbounded returns 1 and all algorithms are successful, it must be that $(A_B^{-1}A_k)_\ell < \delta\|A_B^{-1}A_k\|$ for all ℓ, which is the condition in the statement of the proposition; in this case, the LP is indeed unbounded.

We now analyze the running time. To determine with constant probability if the sought index ℓ exists, i.e., if $g(\ell)=1$ for some value of ℓ, we apply amplitude estimation with $O(\sqrt{m})$ applications of the Grover operator, see Thm. 5. Each application takes time $O(\frac{1}{\delta}(\kappa d^2 m))$, which is the complexity of running the QLSA and the sign estimation. The probability of success for the QLSA is bounded, and we have a way of determining success, therefore we can boost the probability of correctness arbitrarily high with enough repetitions of the algorithm. □

Proof of Thm. 11. Recall the formula for the ratio test, eq. (1):

$$r^* := \min_{j=1,\dots,m:u_j>0} \frac{x_{B(j)}}{u_j},$$

where $x_B = A_B^{-1}b, u = A_B^{-1}A_k$. Note that since $x_B \geq 0$, this can be rewritten as follows:

$$r^* := \max\{r : x_B - ru \geq 0\}. \tag{EC.5}$$

Thus, an exact solution of the ratio test is attained by any index j such that $(A_B^{-1}b - r^* A_B^{-1}A_k)_j = (x_B - r^*u)_j = 0$.

In Alg. 9, on line 3 we define the unitary U_r. By Prop. 3, with the tolerances set in the algorithm, if:

$$\frac{(A_B^{-1}(b - rA_k))_j}{\|A_B^{-1}(b - rA_k)\|} < -\frac{\delta}{4\|A_B^{-1}(b - rA_k)\|},$$

then SignEstNFN returns 0. Removing the normalization factor and taking into account the error of the QLSA, which we set to be at most $\frac{\delta}{4\|A_B^{-1}(b-rA_k)\|}$, we have that if $(A_B^{-1}(b - rA_k))_j < -\frac{\delta}{2}$ then SignEstNFN returns 0 with high probability (that can be boosted, as usual, by doing multiple repetitions, taking the majority vote, and using the Chernoff bound). Similarly, by adjusting the tolerance and using the other side of Prop. 3, we can define a unitary that returns 1 if $(A_B^{-1}(b - rA_k))_j \geq -\frac{\delta}{2}$.

We now analyze the binary search at line 4. We can compute an upper bound on the maximum value of r by starting with $r = 1$ and doubling every time until we find a value such that $A_B^{-1}b - rA_B^{-1}A_k \ngeq -\frac{\delta}{2}\mathbf{1}_m$; this gives an estimate of the upper range for r that is off by at most a factor of two. Note that the maximum value for r is $O(\max\{\|A_B^{-1}b\|, \|A_B^{-1}A_k\|\})$. We are seeking a value $\tilde{r}$ such that $A_B^{-1}b - \tilde{r}A_B^{-1}A_k \ngeq -\frac{\delta}{2}\mathbf{1}_m$ but $A_B^{-1}b - (\tilde{r} - \frac{\delta}{2\kappa\|A_k\|})A_B^{-1}A_k \geq -\frac{\delta}{2}\mathbf{1}_m$; this can be done by binary search in time $O(\log \frac{\kappa \max\{A_B^{-1}b, A_B^{-1}A_k\}\|A_k\|}{\delta}) = \tilde{O}(1)$. For the return value $\tilde{r}$, for every (basic) index j, we have $(A_B^{-1}(b - (\tilde{r} - \frac{\delta}{2\kappa\|A_k\|})A_k))_j \geq -\frac{\delta}{2}$, which implies $(A_B^{-1}(b - \tilde{r}A_k))_j \geq -\frac{\delta}{2} - \frac{\delta}{2\kappa\|A_k\|}(A_B^{-1}A_k)_j \geq -\delta$. This implies that the new basic solution obtained by moving by $\tilde{r}$ along the nonbasic edge k

is δ-feasible, i.e., it violates nonnegativity by at most δ. Additionally, we know that there exists j such that $(A_B^{-1}(b-\tilde{r}A_k))_j < -\frac{\delta}{2}$, and such an index j is precisely the value returned by FindRow; thus, the algorithm returns a row index for the pivot that yields the next basic (approximately) feasible solution.

Regarding the running time, the binary search requires $\tilde{O}(1)$ iterations, as stated; at each binary search iteration, we need $\tilde{O}(\sqrt{m})$ amplitude estimation iterations (Thm. 5) at line 3, that multiplies the cost of running SignEstNFN on the solution of the QLSA, which is $\tilde{O}(\frac{\|A_B^{-1}A_k\|}{\delta}T_{\text{LS}}(A_B, A_k, \frac{\delta}{8\|A_B^{-1}(b-rA_k)\|})) = \tilde{O}(\frac{\|A_B^{-1}A_k\|}{\delta}\kappa d^2 m)$. The total running time of the algorithm is therefore $\tilde{O}(\frac{1}{\delta}\eta d^2\kappa^2 m^{1.5})$. □

EC.7. Proofs from Section 5.7.

Proof of Prop. 9. Let $|\tilde{x}\rangle$ be the state produced by the QLSA; by Thm. 7, we have $\|\tilde{x} - \frac{A_B^{-1}b}{\|A_B^{-1}b\|}\| \leq \frac{\delta}{10}$.

Suppose $A_B^{-1}b \ngeq -\delta\mathbf{1}_m$, i.e., the basic solution is infeasible. Then, for some index ℓ, we must have $\frac{(A_B^{-1}b)_\ell}{\|A_B^{-1}b\|} < -\frac{\delta}{\|A_B^{-1}b\|}$. This implies:

$$\tilde{x}_\ell \leq \frac{(A_B^{-1}b)_\ell}{\|A_B^{-1}b\|} + \frac{\delta}{10\|A_B^{-1}b\|} < -\frac{\delta}{\|A_B^{-1}b\|} + \frac{\delta}{10\|A_B^{-1}b\|} = -\frac{9\delta}{10\|A_B^{-1}b\|}.$$

By Prop. 4, if the routine SignEstNFN$(U_{\text{LS}}, \ell, \frac{9\delta}{20\|A_B^{-1}b\|})$ is successful it returns zero, so that the function $g(\ell)$ evaluates to 1. This implies that if all subroutines are successful, IsFeasible returns 0, as desired. The running time analysis is the same as in Prop. 8. □

EC.8. Proofs from Section 6

Proof of Prop. 10. By (Gilyén et al. 2019, Lemma 50), we can construct a $(\mu(A_B), O(\log n), \epsilon/(\kappa^2 \log^3 \frac{\kappa}{\epsilon}))$ block encoding for A_B using suitable data structures, in time $\tilde{O}(1)$. Using the techniques in (Chakraborty et al. 2018, Sect. 4.3), we can then compute a normalized version of $A_B^{-1}|x\rangle$ in time $\tilde{O}(\mu(A_B)\kappa)$.

The QRAM data structures describing A_B and A_N can be prepared in the claimed time following Kerenidis and Prakash (2018) and Prop. 1. After each iteration of the simplex method, we simply need to reindex the structures in memory, i.e., swap one nonbasic column with a basic column, which takes at most $\tilde{O}(m)$ operations since each column has size m. □