Quantum Lattice Enumeration in Limited Depth

Nina Bindel[1] Xavier Bonnetain[2] Marcel Tiepelt[3] **Fernando Virdia**[4]

[1] SandboxAQ, Palo Alto, CA, USA

[2] Université de Lorraine, CNRS, Inria, Nancy, France

[3] KASTEL, Karlsruhe Institute of Technology, Karlsruhe, Germany

[4] NOVA LINCS, Universidade NOVA de Lisboa, Lisbon, Portugal

- Aim: assess the concrete threat posed by a specific quantum algorithm on newly standardised cryptography.
- This work was published at Crypto 2024 as [BBTV24].
- Morally a follow up to MSR internship work [JNRV20].
- Our results are partial: many known unknowns captured as conjectures, and backed by small-scale experiments.

- In 1994 Peter Shor develops an period-finding *quantum* algorithm in polynomial time [Sho97].
- This algorithm's results in quantum polynomial time attacks on RSA and discrete logarithm.
- Recently, significant investments from industry into developing quantum computing technology [MQT18, MN18, AAB+19, Gib19, WFG21].
- Increased urgency to develop alternative public-key cryptography primitives conjectured to resit quantum-computing attacks ("post-quantum" cryptography).

How we think quantum algorithms

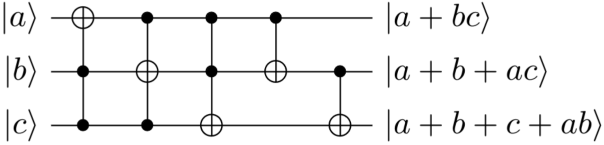

- Width 3, depth 5 and gate count 5.
- The wires are qubits, the nodes are gate evaluations.
- The cost can be expressed in terms of different metrics, e.g. by counting wires, components, depth, area. . .

[JS19] suggests that one can compare the # of quantum gates with CPU cycles.

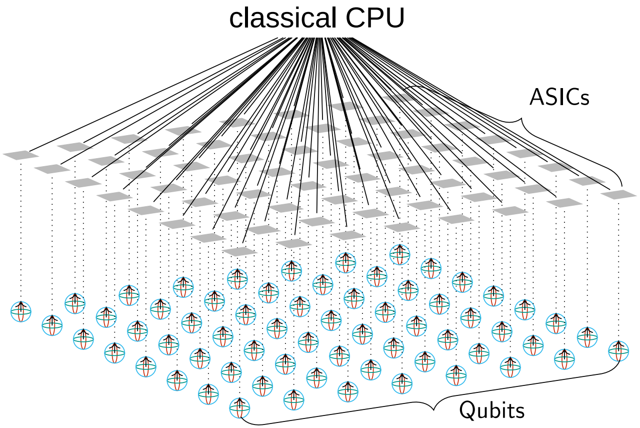

⇒ We consider number of gates as an estimate for the cost of a circuit.

[0] Image courtesy of Sam Jaques.

- In 2016, NIST publishes a call for proposals for post-quantum signature schemes and key encapsulation mechanisms [Nat16].
- They propose a model for thinking about concrete post-quantum security:
  - A candidate scheme should be as hard to break "as the AES block cipher".
  - Quantum computers that can perform a limited number max-depth (MD) of serial gate evaluations: qubits are hard to error-correct.

Proposed values for max-depth (MD):

- $MD = 2^{40} \approx$ "gates that presently envisioned quantum computing architectures are expected to serially perform in a year".
- $MD = 2^{64} \approx$ "gates that current classical computing architectures can perform serially in a decade".
- $MD = 2^{96} \approx$ "gates that atomic scale qubits with speed of light propagation times could perform in a millennium".

The max-depth constraint can significantly impact quantum attack perfromance.

- Attackers may be limited in the size of the instances they can solve before decoherence.
- Multiple quantum circuits may have to be run in parallel to solve larger instances.

Example: Quantum exaustive key-search on AES

- AES-256: naively, Grover's requires depth/gates $\approx \sqrt{2^{|\text{key}|}} = 2^{128} > MD$.
- Grover search almost certainly fails if stopped early:
  $\implies$ We need to account for Grover's parallelisation.
- Grover search parallelises badly [Zal99], causing the concrete quantum advantage to strongly reduce [JNRV20]: AES-256 ($MD = 2^{96}$) $\Rightarrow 2^{192}$ gates)

- In 2023 NIST posts the first draft standards for comments.
- Four candidates are selected to become new standards.
- $3/4$ depend on computational hardness conjectures about algebraic lattices.

Natural questions

- What are the best quantum attacks on lattice problems?
- What is their cost against the standards?

Case-study: Kyber (ML-KEM).

- Depends on the hardness of distinguishing a specific distribution of integer matrices "modulo q" from uniformly random.
- Classically, the two best attack approaches require performing "lattice reduction".
  - Given an public key, build a matrix $\boldsymbol{B} \in \mathbb{Z}_q^{m\times m}$. Want to find a "short" non-zero vector in the integer span of the columns of $\boldsymbol{B}$ (the "lattice with basis $\boldsymbol{B}$").
  - To do so, call a "block reduction" algorithm on $\boldsymbol{B}$ (*e.g.* BKZ [SE91], Slide reduction [GN08], Progressive BKZ [AWHT16], Self-Dual BKZ [MW16]...).
  - Block reduction constructs a polynomially long sequence of related, smaller-rank matrices $(\boldsymbol{B}_i \in \mathbb{Z}_q^{m\times n})_i$, and looks for a short-enough non-zero vector in the integer span of each $\boldsymbol{B}_i$.
  - Finding a short vector in such lattices is considered hard, and is an instance of the "short vector problem" (SVP). An SVP solver is used for each $\boldsymbol{B}_i$.

- Block reduction is a classical algorithm, its cost is dominated by that of solving SVP.
- The best quantum attacks of Kyber involve applying quantum speed-ups to SVP solvers.
- There are many approaches for building an SVP solver.
- At least two of these, *sieving* and *enumeration*, can be "compiled" into quantum algorithms using black-box methods [LMv13, KMPM19, ANS18, BCSS23].
- The resulting asymptotic quantum speedups are understood, but there's not a lot of work on their concrete cost [AGPS20] (and now **[BBTV24]**).

Our work: new conjectured lower bounds on the concrete cost of quantum enumeration with extreme cylinder pruning (incl. a new quantum enumeration algorithm).

- Quantum enumeration algorithms were first demonstrated by Aono *et al.* [ANS18]; asymptotically, they provide a $\approx$ quadratic speedup.
- Our work looks at the "max-depth" setting [Nat16, Pre18].
- Our results suggest that quantum speedups in this setting **may** not apply (just as for Grover [JNRV20]).

Lattice enumeration

- Say we are looking for a short vector $v \neq 0$ in a lattice L with basis $(\boldsymbol{b_1}, \ldots, \boldsymbol{b_{n-1}}, \boldsymbol{b_n})$.
- Suppose we know an upper bound R on $\|v\|$.
- In enumeration, we explore all (or most) vectors in L of norm $\leq R$, optionally stopping when we find one.
- Conceptually, enumeration consists of depth-first search on a tree T containing short vectors as leaves.
- As used in lattice reduction, in dimension n, this requires $\text{poly}(n)$ memory, and $\mathbb{E}[\#T] = 2^{\frac{1}{8} n \log n + o(n)}$ time on average.

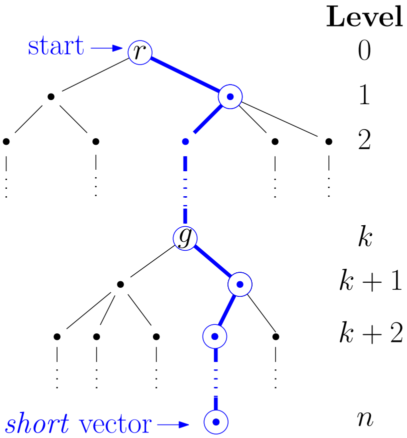

A look at the enumeration tree T

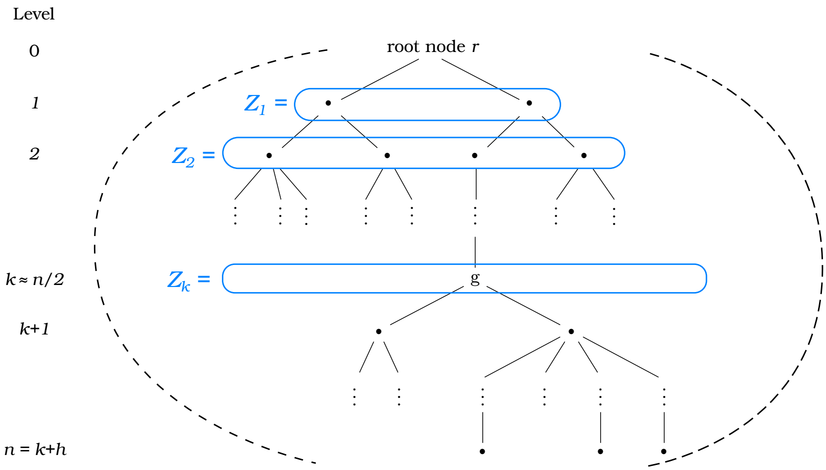

- Nodes located on different levels Z_k.
- "Middle" levels super-exponentially large [GNR10]:
  $$\#T \approx \#Z_{n/2}$$
- The tree size can be somewhat reduced by "pruning" unlikely paths early.

Montanaro's quantum tree search

- In 2018, Montanaro introduces two quantum tree-search algorithms, DetectMV and FindMV [Mon18].
- Given a tree T and a predicate P, DetectMV returns whether $\exists$ leaf $\in T$ such that $P(\text{leaf}) = \text{true}$ in $\tilde{O}(\sqrt{\mathcal{T} \cdot n})$ evaluations of P, where $\mathcal{T}$ = upper bound of $\#T$.
- By performing decision on every level, *DetectMV* $\mapsto$ *FindMV*, which returns such a leaf.
- For trees with $O(1)$ marked leaf and $\#T \approx \mathcal{T}$:

  Classical avg. case runtime $O(\#T) \mapsto$ quantum avg. case depth $\tilde{O}(\sqrt{\#T \cdot n})$.

Montanaro's quantum tree search

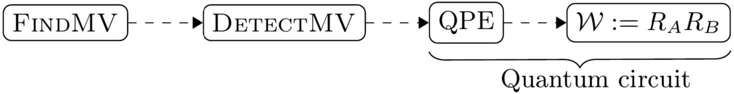

- DetectMV = repeating multiple Quantum Phase Estimations (QPE) of an operator W that checks the predicate P; evaluating QPE(W) is the *quantum part*.
- QPE(W) = serially evaluate $\tilde{O}(\sqrt{\#T \cdot n})$ times the operator W.
- Our objective: estimate/lower-bound the expected gate-cost of FindMV(T), while keeping the depth of QPE(W) within max-depht MD.

A back of the envelop estimation/lower bound of the depth of QPE(W)

- Lower-bound the size of W by assuming $\text{Depth}(W) = \text{Gates}(W) = 1$.
- Using the LWE estimator we find the required block size β to break Kyber.
  - β is the depth n of tree.
  - From n we obtain $\#T$ by using lower bounds for the cost of enumeration with cylinder pruning [ANSS18].
- Finally, we check if the resulting circuit depth of QPE(W) is $\leq MD$.

$$\mathop{\mathbb{E}}_{\substack{\text{random}\\\text{tree } T}}[\mathrm{Depth}(\mathrm{QPE}(W))] \approx \mathbb{E}[\sqrt{\#T \cdot \beta}] \approx \sqrt{\mathbb{E}[\#T] \cdot \beta} \approx \begin{cases} 2^{90.3} & \text{for Kyber-512,} \\ 2^{166.2} & \text{for Kyber-768,} \\ 2^{263.7} & \text{for Kyber-1024.} \end{cases}$$

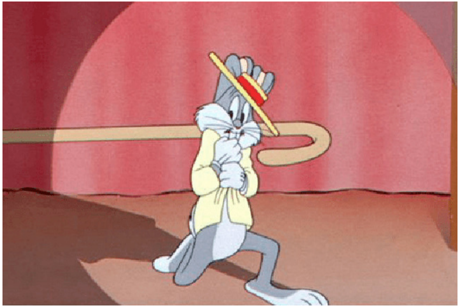

- Wait, don't drag me out of the room.
- I do know Jensen's inequality! $\mathbb{E}[\sqrt{\#T}] \leq \sqrt{\mathbb{E}[\#T]}$.
- We plausibly don't fit within 2^{96} depth.

We need smaller trees to enumerate.

Classic trick from parallel enumeration

- Precompute nodes up to level $k > 1$, run FindMV on the subtrees.
- We can estimate the size of subtrees using similar techniques to those used for the full tree.

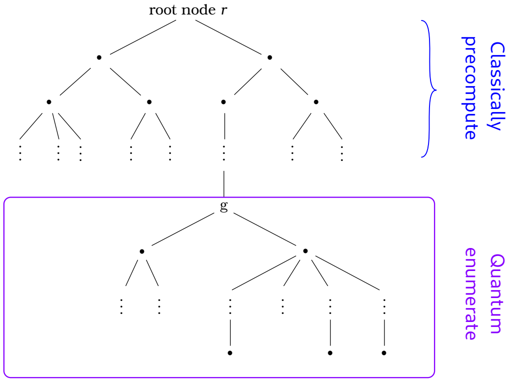

Up to what level k should we precompute?

- $k \approx 1$: QPE(W) covering most of the tree would have to fit within max-depth: likely not our case.
- $k \approx n/2$: we run $\approx \left|Z_{n/2}\right|$ quantum enumeration calls: cost $\approx$ classical enumeration.
- $k \approx n$: we precompute most of the classical tree, no speedup.

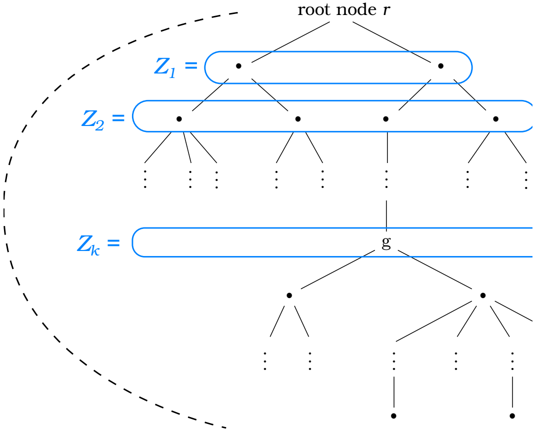

- Quantum enumeration on level $k \ll n/2$ is likely impossible.
- On level $k \geq n/2$ it is pointless.
- Our best chance is $k \lessapprox n/2$, somehow reducing the number of calls to be $\ll |Z_{n/2}|$.

Bundle trees rooted in Z_k into bunches

- Precompute sets of 2^y elements in Z_k.
- Collect them under a 'virtual' node v.
- Run FindMV over the tree $T(v)$ with root v.

Disclaimer

- Bundling requires 2^y QRACM.
- QRACM may be quite costly to access [JR23].
- Yet, many quantum-classical speedups assume it.

Having identified a more general combined classical-quantum enumeration strategy, we would like to estimate its cost.

- Want a formula for the average cost of the attack, in terms of quantum gates and circuit depth.
- If not possible, we'd settle for lower bounds and hope they are very high.
- We now look at the depth of QPE(W), the gate count follow similarly.

First conjecture

Let $T(v)$ be a tree of height h. Since Depth(QPE(W)) $\in \tilde{O}(\sqrt{\# T(v) \cdot h})$, our first conjectured lower bound is

$$\text{Depth}(\text{QPE}(W)) \geq \sqrt{\# T(v) \cdot h}.$$

- Given a specific attack target, the value of h will be determined by k as part of the attack strategy.
- Therefore $\underset{\substack{\text{random}\\\text{tree } T}}{\mathbb{E}}[\mathrm{Depth}(\mathrm{QPE}(W))] \geq \underset{\substack{\text{random}\\\text{tree } T}}{\mathbb{E}}\left[\sqrt{\#T(v)}\right] \cdot \sqrt{h}$.
- There is no theory about estimating $\mathbb{E}\left[\sqrt{\#T(v)}\right]$ in the lattice literature (Aono *et al.* [ANS18] already mention this issue).
- Jensen's gap only gives us an upper bound: $\mathbb{E}\left[\sqrt{\#T(v)}\right] \leq \sqrt{\mathbb{E}\left[\#T(v)\right]}$.

Definition: Multiplicative Jensen's gap

Let X be a random variable. We say X has multiplicative Jensen's gap 2^z if

$$\sqrt{\mathbb{E}[X]} = 2^z\, \mathbb{E}[\sqrt{X}].$$

Ideally, we'd like an upper bound to z. We will estimate "around it".

- The Jensen's gap gives us $\mathbb{E}\left[\mathrm{Depth}(\mathrm{QPE}(W))\right] \geq 2^{-z}\sqrt{\mathbb{E}\left[\# T(v)\right]} \cdot \sqrt{h}$.
- We now need $\mathbb{E}[\# T(v)]$.
- Standard lattice theory gives us this for the full enumeration tree T, and for cylinder-pruned trees.
- However, we are looking at sub-trees rooted on level k.

Second + third conjectures combined

Let $T(g)$ be a sub-tree with root $g \in Z_k$. Then

$$\# T(g) \approx \mathbb{E}\left[\frac{\sum_{i>0} |Z_{k+i}|}{|Z_k|}\right] \gtrapprox \sum_{i>0} \frac{\mathbb{E}[|Z_{k+i}|]}{\mathbb{E}[|Z_k|]}.$$

All combined, we arrive at our conjectures lower bounds for the $\mathbb{E}[\text{cost}]$ of the attack.

Quantum depth

$$\mathbb{E}\left[\mathrm{Depth}(\mathrm{QPE}(W))\right] \geq \frac{1}{2^z}\sqrt{\mathbb{E}\left[\# T(v)\right] \cdot (n-k+1)} \cdot \mathrm{Depth}(W), \text{ for } g \in Z_k.$$

Quantum gate-cost

$$\mathbb{E}[\mathrm{Gates}(\mathrm{FindMV})] \geq \frac{\mathbb{E}[|Z_k|]}{2^y} \cdot \frac{1}{2^z}\sqrt{\mathbb{E}\left[\# T(v)\right] \cdot (n-k+1)} \cdot \mathrm{Gates}(W), \text{ for } g \in Z_k.$$

We can now try to compute some estimates.

- We assume either $\mathrm{Depth}(W) = \mathrm{Gates}(W) = 1$ (in the "query-model") or an estimated lower bound based on best-known quantum arithmetic circuits (in the "circuit-model", similar to independent work [BvHJ+23]).
- We decide how to lower bound $\#T(g) \gtrapprox \sum_{i>0} \frac{\mathbb{E}[|Z_{k+i}|]}{\mathbb{E}[|Z_k|]}$: for the numerator should we use our best known estimates, or absolute lower bounds [ANSS18]?
- We use the LWE-estimator to find the enumeration dimension $n = \beta$.
- We estimate costs for every $k \leq n$, $y \leq 64$, $z \leq 64$.
- We report smallest z such that our lower bound of *classical* $+$ *quantum gate-cost* $\leq$ Grover search on AES.

more likely to be feasible → less likely to be feasible						
	Kyber-512		Kyber-768		Kyber-1024	
	GCost of quantum walk operator $\mathcal{W}$					
MaxDepth	1	*minimal*	1	*minimal*	1	*minimal*
2^{40}	$z \geq 0$	$z \geq 0$	$z \geq 2$	$z \geq 17$	$z \geq 50$	$z > 64$
2^{64}	$z \geq 0$	$z \geq 0$	$z \geq 1$	$z \geq 17$	$z \geq 49$	$z > 64$
2^{96}	$z \geq 0$	$z \geq 0$	$z \geq 1$	$z \geq 19$	$z \geq 51$	$z > 64$

Figure: Smallest Jensen's gap for which lower bound on attack cost $\leq$ Grover-on-AES' cost. Using [ANSS18]'s lower bounds for subtree sizes: it requires maximally improving current cylinder pruning technique.

more likely to be feasible → less likely to be feasible

	Kyber-512		Kyber-768		Kyber-1024	
	GCost of quantum walk operator $\mathcal{W}$					
MaxDepth	1	*minimal*	1	*minimal*	1	*minimal*
2^{40}	$z \geq 20$	$z \geq 36$	$z \geq 61$	$z > 64$	$z > 64$	$z > 64$
2^{64}	$z \geq 20$	$z \geq 36$	$z \geq 61$	$z > 64$	$z > 64$	$z > 64$
2^{96}	$z \geq 15$	$z \geq 40$	$z \geq 61$	$z > 64$	$z > 64$	$z > 64$

Figure: Smallest Jensen's gap for which lower bound on attack cost $\leq$ Grover-on-AES' cost. Using current understanding of cylinder pruning to estimate subtree sizes.

Take aways

- Likely we can exclude quantum enumeration on Kyber-768 and -1024.
- In the "circuit-model" for W, attacks on Kyber-512 also looks unlikely.
  - And we are being quite strict in various parts of the computation.
- There's "good hope" that quantum enumeration does not pose a threat.

Clarification

Yet, we can't fully exclude it without a clear understanding of the Jensen gap.

Can we say anything about this gap?

Open problems: Jensen's gap

- The overall classical+quantum cost changes smoothly as a funciton of z $\implies$ rough estimates of z may already help.
- Experimental evidence up to $\beta = 70$ says $z \approx 1$.
- Alternatively, we can prove lower bounds on $\mathbb{E}[\sqrt{\#T}]$:

$$\mathbb{E}[\sqrt{\#T}] \geq \max\left\{\sqrt{\mathbb{E}[\#T]} - \sqrt[4]{\mathbb{V}[\#T]},\quad 2^{-\frac{1}{2\ln 2}\sqrt[4]{\mathbb{V}[\#T]}} \cdot \sqrt{\mathbb{E}[\#T]}\right\}.$$

But both depend on $\mathbb{V}[\#T]$, which is also not known.

Open problems: other directions

- We've only covered cylinder pruning. What about discrete pruning? Or ad-hoc pruning for quantum enumeration?
- Currently, searching for attack costs is an optimisation problem. Can we find a closed formula? This would allow running it as part of "estimator" scripts.
- There quite a few other places where our analysis is not be tight, meaning actual costs are likely higher.

Conclusions

- Asymptotically quadratic quantum speedups on enumeration look unlikely against lattice-based cryptography under max-depth constraints.
- Technically hard to fully exclude the viability of quantum enumeration.
- More needs to be learnt about the distribution of enumeration trees, to reduce conjectures and learn the Jensen's gap for enumeration tree sizes.

Thank you

Paper @ https://eprint.iacr.org/2023/1423
Slides @ https://fundamental.domains

Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, and William *et al.* Courtney.
Quantum supremacy using a programmable superconducting processor.
Nature, 574(7779):505–510, Oct 2019.

Martin R. Albrecht, Vlad Gheorghiu, Eamonn W. Postlethwaite, and John M. Schanck.
Estimating quantum speedups for lattice sieves.
In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 583–613. Springer, Heidelberg, December 2020.

Yoshinori Aono, Phong Q. Nguyen, and Yixin Shen.
Quantum lattice enumeration and tweaking discrete pruning.
In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part I*, volume 11272 of *LNCS*, pages 405–434. Springer, Heidelberg, December 2018.

Yoshinori Aono, Phong Q. Nguyen, Takenobu Seito, and Junji Shikata.
Lower bounds on lattice enumeration with extreme pruning.
In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 608–637. Springer, Heidelberg, August 2018.

Yoshinori Aono, Yuntao Wang, Takuya Hayashi, and Tsuyoshi Takagi.
Improved progressive BKZ algorithms and their precise cost estimation by sharp simulator.
In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 789–819. Springer, Heidelberg, May 2016.

Nina Bindel, Xavier Bonnetain, Marcel Tiepelt, and Fernando Virdia.
Quantum lattice enumeration in limited depth.
Springer-Verlag, 2024.

Xavier Bonnetain, André Chailloux, André Schrottenloher, and Yixin Shen.
Finding many collisions via reusable quantum walks - application to lattice sieving.
In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 221–251. Springer, 2023.

Shi Bai, Maya-Iggy van Hoof, Floyd B. Johnson, Tanja Lange, and Tran Ngo.
Concrete analysis of quantum lattice enumeration.
In *Advances in Cryptology - ASIACRYPT 2023*, Lecture Notes in Computer Science. Springer-Verlag, 2023.

Elizabeth Gibney.
Quantum gold rush: the private funding pouring into quantum start-ups.
Nature, 574(7776):22–24, October 2019.

Nicolas Gama and Phong Q. Nguyen.
Finding short lattice vectors within Mordell's inequality.
In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 207–216. ACM Press, May 2008.

Nicolas Gama, Phong Q. Nguyen, and Oded Regev.
Lattice enumeration using extreme pruning.

In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 257–278. Springer, Heidelberg, May / June 2010.

Samuel Jaques, Michael Naehrig, Martin Roetteler, and Fernando Virdia.
Implementing grover oracles for quantum key search on AES and LowMC.
In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 280–310. Springer, Heidelberg, May 2020.

Samuel Jaques and Arthur G. Rattew.
Qram: A survey and critique, 2023.

Samuel Jaques and John M. Schanck.
Quantum cryptanalysis in the RAM model: Claw-finding attacks on SIKE.
In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 32–61. Springer, Heidelberg, August 2019.

Elena Kirshanova, Erik Mårtensson, Eamonn W. Postlethwaite, and Subhayan Roy Moulik.
Quantum algorithms for the approximate k-list problem and their application to lattice sieving.
In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 521–551. Springer, Heidelberg, December 2019.

Thijs Laarhoven, Michele Mosca, and Joop van de Pol.
Solving the shortest vector problem in lattices faster using quantum search.
In Philippe Gaborit, editor, *Post-Quantum Cryptography - 5th International Workshop, PQCrypto 2013*, pages 83–101. Springer, Heidelberg, June 2013.

Samuel K. Moore and Amy Nordrum.
Intel's new path to quantum computing.
IEEE Spectrum, 2018.

Ashley Montanaro.
Quantum-walk speedup of backtracking algorithms.
Theory Comput., 14(1):1–24, 2018.

Microsoft Quantum Team.
Developing a topological qubit.
Cloud Perspectives Blog, 2018.

Daniele Micciancio and Michael Walter.
Practical, predictable lattice basis reduction.
In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 820–849. Springer, Heidelberg, May 2016.

National Institute of Standards and Technology.
Submission requirements and evaluation criteria for the Post-Quantum Cryptography standardization process.
http://csrc.nist.gov/groups/ST/post-quantum-crypto/documents/call-for-proposals-final-dec-2016.pdf, December 2016.

John Preskill.
Quantum Computing in the NISQ era and beyond.
Quantum, 2:79, August 2018.

Claus-Peter Schnorr and M Euchner.
Lattice basis reduction: Improved practical algorithms and solving subset sum problems.
In *FCT*, 1991.

Peter W. Shor.
Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer.
SIAM J. Comput., 26(5):1484–1509, October 1997.

Karl Wehden, Ismael Faro, and Jay Gambetta.
IBM's roadmap for building an open quantum software ecosystem.
IBM Research Blog, 2021.

Christof Zalka.
Grover's quantum searching algorithm is optimal.
Phys. Rev. A, 60:2746–2751, Oct 1999.