



Universiteit
Leiden



Advancing
Quantum
Algorithms

Which resource gives quantum computers their power?

SUMMIT Quantum Limits, Complexity session
Dokkum, 9 September 2026

Jordi Tura i Brugués · Applied Quantum Algorithms, Leiden University



Google Research



The Complexity theme

The programme is organised around four kinds of limit. I am representing the complexity one.

MASS

Is there a mass limit to quantum physics? Can gravity create entanglement?

COMPLEXITY

What are the quantum limits of complexity?

TIME

Can we realise a system that stays quantum indefinitely?

SPACE

Limits to multipartite nonlocality, and to the quantisation of space.

THE FOUR QUESTIONS IN THE COMPLEXITY THEME

1

For what quantum algorithms do we have efficient classical simulations, and which quantum computations produce an advantage?

Acts 1 to 3

2

Can quantum algorithms help us understand interacting quantum particles?

quantum PCP, fermions

3

Does quantum physics allow more transparent and understandable AI?

explainable quantum AI

4

Does quantum spatial nonlocality improve quantum algorithms?

Acts 4 and 5

The hour is about the first question, and it ends up in the fourth one as well.

The machines

Some results from the last few years. The platforms do different things, so I am not comparing them to each other.

2021 A three-node quantum network

Entanglement between separated solid-state nodes. Delft.

2022 Universal control of a silicon processor

Spin qubits, made on a semiconductor line. Delft.

2024 Below the surface code threshold

The logical error rate falls as the code grows. Google.

2025 Tweezer arrays of thousands of atoms

Held, and coherent. Caltech.

2025 Logical magic state distillation

On a neutral-atom processor. Harvard, MIT, QuEra.

2025 Trapped ions, all to all

Dozens of logical qubits encoded. Quantinuum.

2026 More couplers, longer circuits

Superconducting. IBM.

We are building all of this on the promise that these machines are powerful. Powerful why?

In Search of the Quantum Resource

Five stops. Approx.

1

Superposition

a biased coin is not a qubit

2

Entanglement

necessary, not sufficient, not even a good proxy

3

Magic

today's favourite, and what came before it

4

Can we even tell if the resource is there?

detecting any of this in the lab is hard

5

What survives

one theorem, and an answer of sorts


The naive answer

“

A quantum computer is more powerful because it doesn't only have the 0 and the 1, it also has the one in the middle.

A self-proclaimed quantum evangelist, corporate keynote, RISK 2018


“the one in the middle”



a coin with bias p
nothing quantum here

a qubit

$$\alpha |0\rangle + \beta |1\rangle$$


$$+ = 0$$

β can be negative (complex)
the sign is the whole point

The polite version of the same mistake: “it tries all answers in parallel”.

Measurement gives you one answer. The work is in making the wrong ones cancel.

Three candidates

Each was, for a while, the resource. What happened each time?

1

Superposition

why it was exciting

2ⁿ amplitudes at once

what we learned

so does a probability vector

2

Entanglement

why it was exciting

*correlations with no classical
analogue*

what we learned

GHZ: maximal, and trivial to simulate

3

Magic

why it was exciting

whatever Cliffords cannot make

what we learned

product magic states do nothing on
their own

To keep in mind: “which resource” assumes there is one o_O


$$\langle aQa^L \rangle$$

1. Superposition

What superposition actually buys

An exponential head-start, in memory. Where exactly?

n classical bits, randomised

$$p = (p_{0\dots 0}, \dots, p_{1\dots 1})$$

2^n entries

$$p_x \geq 0, \quad \sum_x p_x = 1$$

stochastic matrices

vs

n qubits

$$\psi = (\psi_{0\dots 0}, \dots, \psi_{1\dots 1})$$

2^n entries

$$\psi_x \in \mathbb{C}, \quad \sum_x |\psi_x|^2 = 1$$

unitaries


Nobody stores this vector. You sample it: n bits of memory, one trajectory at a time.

Sample this one the same way? The sign kills it: cancellations blow up the variance.

The head-start in memory is real, and the sign is why. It cannot be read out: n qubits return n bits.

Interference does the work

Two coin flips versus two Hadamards.



Every quantum algorithm is a way of arranging cancellations.

Coherence, formalised

Baumgratz, Cramer, Plenio, PRL 113, 140401 (2014).

free states

$$\Delta(\rho) = \sum_i |i\rangle\langle i| \rho |i\rangle\langle i|$$

diagonal in a fixed basis

free operations

$$\Lambda(\mathcal{I}) \subseteq \mathcal{I}$$

cannot create off-diagonal terms

monotone

$$C_{\ell_1}(\rho) = \sum_{i \neq j} |\rho_{ij}|$$

zero on free states, non-increasing under free operations

The case for it

No coherence, no interference: a classical sampler.

Necessary, and quantified. Does more mean more power?


Coherence is the orange part.



Same mould for every candidate: free states, free operations, a number.

Coherence. Geometrical interpretation

The free set is a line through the Bloch ball. Deciding is trivial.



incoherent: $\rho = \frac{1}{2}(\mathbb{I} + r_z Z)$, $C_{\ell_1}(\rho) = 2|\rho_{01}| = \sqrt{r_x^2 + r_y^2}$

The free set is a segment through the centre of the ball. Convex.

Coherence is the distance to it. Free operations move you towards it.

Deciding whether it is there: read the off-diagonal elements. Trivial.

Rotate the basis and the segment rotates. The resource is where you draw the line.

A convex free set. Membership is trivial.

What coherence is good for

Where it does real work

Metrology: phase sensitivity. The resource there is asymmetry, coherence relative to a symmetry.

Thermodynamics: coherence between energy levels is constrained on its own, beyond free energy.

Discrimination games: robustness of coherence is exactly the advantage you gain.

Any nonzero coherence converts into entanglement with an incoherent ancilla.

T2. Decoherence destroys exactly this quantity, and every experiment fights it.

Where it does not

$$C_{\ell_1}(|+\rangle) = 1 \text{ in } Z, \quad 0 \text{ in } X$$

Basis dependent.

$$C_{\ell_1}(|+\rangle^{\otimes n}) = 2^n - 1$$

The maximum, on a product state.

2n numbers. Zero power.

Coherence buys real things. As a measure of computational power?

One qubit at a time gets you nowhere

$$|\psi\rangle = \bigotimes_{k=1}^n (\alpha_k |0\rangle + \beta_k |1\rangle)$$

$2n$ numbers. A laptop does a million qubits before lunch.

All 2^n amplitudes are there. None of the head-start is being spent.

Superposition and interference are all there. Nothing is hard.

Superposition, so far

necessary	yes
-----------	-----

sufficient	no
------------	----

does more mean more power	no
---------------------------	----

You need gates on two qubits at once. That is where the next hope came from.

A background network diagram consisting of a complex web of thin grey lines connecting various black and grey dots of different sizes. The dots are scattered across the slide, with some acting as central hubs connected to multiple other dots. The overall pattern is abstract and geometric, resembling a molecular structure or a data network.
$$\langle aQa^L \rangle$$

2. Entanglement

Entanglement, for pure states

Write the coefficients as a matrix.

$$|\psi\rangle_{AB} = \sum_{ij} C_{ij} |i\rangle_A |j\rangle_B$$

Product state: C has rank 1, so it factorises.

$$|\psi\rangle = \sum_{k=1}^r \sqrt{\lambda_k} |a_k\rangle |b_k\rangle, \quad r = \text{rank } C$$

$$S(\rho_A) = - \sum_k \lambda_k \log \lambda_k$$

Schmidt rank r , and the entropy of what Bob holds.

An SVD decides it. For pure states this is cheap.

For pure states, entanglement is the rank of a matrix. Mixed states?

Two examples

$$|0\rangle \otimes \frac{1}{\sqrt{2}} (|0\rangle + |1\rangle) \rightarrow C = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}$$

rank 1. Product.

$$\frac{1}{\sqrt{2}} (|00\rangle + |11\rangle) \rightarrow C = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

rank 2, $\lambda = (\frac{1}{2}, \frac{1}{2})$, $S = 1$.

Maximally entangled, for two qubits.

What entanglement actually is

Mixed states are not optional.

$$\rho_{AB} \text{ separable} \iff \rho_{AB} = \sum_k p_k \rho_k^A \otimes \rho_k^B$$

a telephone line



Separable: anything Alice and Bob can prepare with local operations and a telephone line.

Entangled: everything else. Labs have noise, subsystems are mixed.

The decomposition is not unique, so there is no rank to read off.

Separable set is convex. Hardness of membership problem?

Bertlmann's socks

$$\rho = \frac{1}{2}(|00\rangle\langle 00| + |11\rangle\langle 11|)$$

A coin copied twice. See one sock, and you know the other.

Correlated, and separable. Entanglement is not correlation.

Why we believed it

Jozsa and Linden, Proc. R. Soc. A 459, 2011 (2003). Vidal, PRL 91, 147902 (2003).

Theorem (informal)


Bounded entanglement across every cut, at every step, implies a polynomial-time classical simulation.

$$|\psi\rangle = \sum_{\alpha=1}^{\chi} \lambda_{\alpha} |a_{\alpha}\rangle |b_{\alpha}\rangle, \quad \chi \leq \text{poly}(n) \Rightarrow \text{easy}$$

Entanglement is necessary. The theorem says nothing about sufficient.

The simulator that entanglement defines

Tensor networks: Jozsa-Linden as software.



physical legs: the n qubits. bonds: χ numbers each, χ = Schmidt rank across that cut.

$$|\psi\rangle = \sum_{\{\alpha\}} A_{\alpha_1}^{(1)} A_{\alpha_1 \alpha_2}^{(2)} \cdots A_{\alpha_{n-1}}^{(n)} |s_1 \cdots s_n\rangle, \quad \text{cost} \sim \text{poly}(n, \chi)$$

What it buys, and the wall

1D area law: constant χ , DMRG runs the field.

Operators too: the QFT is an MPO.

2D: efficient to write down, #P-hard to contract.

Every free theory comes with a simulator. This is the one entanglement defines.

First surprise: Gottesman and Knill

H, S, CNOT and Pauli measurements: simulable in polynomial time.

$$|\text{GHZ}_n\rangle = \frac{1}{\sqrt{2}}(|0 \cdots 0\rangle + |1 \cdots 1\rangle)$$

$$|C_G\rangle = \prod_{(i,j) \in E} CZ_{ij} |+\rangle^{\otimes n}$$

The most entangled states people know how to name.

A laptop tracks them with n^2 bits.

Entanglement is maximal. Nothing is hard.


Why it is easy

Track n Paulis that stabilise the state, not 2^n amplitudes.

Cliffords map Paulis to Paulis.

The QFT looks maximally entangling

Every qubit talks to every other qubit.



The belief

All-to-all. Early results: maximal operator entanglement.

This must be where Shor gets its power. Right?

$$R_k = \begin{pmatrix} 1 & 0 \\ 0 & e^{2\pi i/2^k} \end{pmatrix}, \quad \binom{n}{2} \text{ two-qubit gates}$$



Second surprise: the QFT barely entangles

Chen, Stoudenmire, White, PRX Quantum 4, 040318 (2023).

$$F_n = R_n Q_n \quad (R_n : \text{bit reversal}, Q_n : \text{the rest})$$

angle between qubits $l < m$: $\frac{\pi}{2^{m-l}} \rightarrow 0$ exponentially

Schmidt coefficients of Q_n decay exponentially

Bond dimension 8 already gives double precision, for any n .

All the operator entanglement is in the bit reversal. Relabelling qubits is free.

QFT on a low- χ MPS: linear time in n .

So where is Shor's power?

In the modular exponentiation, which entangles the two registers heavily.

The QFT reads the period out by interference.

SEEN BEFORE

2006, 2007

Aharonov, Landau and Makowsky; Yoran and Short; Browne. Simulable on product inputs and outputs.

2014

Woolfe, Hill and Hollenberg. The QFT as an MPO, numerically, without the reason.

2023 adds the operator itself, and the reason. 153 scites, and this question was asked under it.

Distant qubits barely talk. The circuit is all-to-all on paper and quasi-local in fact.

Factoring 15, and factoring 21

Two ways of asking how much of the algorithm was actually run.

Compiled with the answer in hand

“Pretending to factor large numbers on a quantum computer”

2001 was the landmark. What came after it was compiled with the answer already known.

Any product of two distinct odd primes, in constant time, with two coherent qubits.

The cost is set by the length of the period.

Smolin, Smith, Vargo, Nature 499, 163 (2013).

The count

15 21

6 CNOT, 2 Toffoli

21 ?

entangling gates

Factoring 15 takes 21 entangling gates. How many does factoring 21 take?



Factoring 15, and factoring 21

Two ways of asking how much of the algorithm was actually run.

Compiled with the answer in hand

“Pretending to factor large numbers on a quantum computer”

2001 was the landmark. What came after it was compiled with the answer already known.

Any product of two distinct odd primes, in constant time, with two coherent qubits.

The cost is set by the length of the period.

Smolin, Smith, Vargo, Nature 499, 163 (2013).

The count

15	21	6 CNOT, 2 Toffoli
----	----	-------------------

21	2405	191 CNOT, 369 Toffoli
----	------	-----------------------

entangling gates. About 115 times more.

Mod 15, six of the eight multiplications are by 1, and multiplying by 4 is two conditional swaps.

Gidney (2025).

Factoring 15 takes 21 entangling gates. Factoring 21 takes 2405. How much entanglement does a demonstration actually need?

For every $\varepsilon > 0$ there is a universal model of quantum computation in which every continuous entanglement measure, across every bipartition, at every step, is at most ε .

Van den Nest, PRL 110, 060504 (2013)


Fuel that works in homeopathic doses is not fuel.

How: dilute the computation



The whole computation rides on one flag qubit.

$$|\Psi_\varepsilon\rangle = \sqrt{1-\varepsilon} |0\rangle_{\text{flag}} |0\cdots 0\rangle + \sqrt{\varepsilon} |1\rangle_{\text{flag}} |\psi\rangle$$

the computation $|\psi\rangle$



the diluted state $|\Psi_\varepsilon\rangle$



Measure the flag, repeat $O(1/\varepsilon)$ times. Every continuous measure is $O(\varepsilon \log 1/\varepsilon)$.

What it means

A quantity of entanglement is not what you are paying with.

1

What dies

Speed-up proportional to entanglement.

The word 'amount'.

2


What survives

Jozsa-Linden. Rank is not continuous.

Necessary, yes. Meterable, no.

Same lesson as coherence, one level up: the number does not track the power.

The entanglement-simulability plane





The rank version, in the wild

Random-circuit sampling is a duel between hardware and tensor networks.

CZ : rank 2 iSWAP, fSim : rank 4

Every gate across a cut multiplies the simulation cost by its rank.

Sycamore's gates were chosen against the simulator.

And the simulators answered

2019: 200 s versus '10 000 years'.

2022: tensor networks on GPUs, hours.

2024: noise sets where the classical side breaks.

The currency is rank per gate per cut. A count, not an amount.

Caveat: random circuits are random Boolean functions one level up. Almost all hard, none useful.



Entanglement, so far

Where Act 2 leaves us.

Entanglement, so far

necessary	yes, in rank
-----------	--------------

sufficient	no
------------	----

does more mean more power	no
---------------------------	----

Two candidates, and the same lesson each time.

Change what is free and the resource moves.

Cliffords are free. What is left over has a name.

And hold this thought: deciding whether a state is entangled at all is NP-hard. Act 4.

A background network diagram consisting of a complex web of thin grey lines connecting various nodes. Some nodes are represented by small grey dots, while others are larger black dots. The overall structure is abstract and geometric, resembling a molecular or computational network.

$\langle aQa^L \rangle$

3. Magic

What Cliffords cannot do

Bravyi and Kitaev, PRA 71, 022316 (2005).

Free: the stabilizer world

$$\text{STAB}_n = \{ C |0\rangle^{\otimes n} : C \in \langle H, S, \text{CNOT} \rangle \}$$

Entanglement, teleportation, error correction.

All simulable. So none of it is the power.

Not free: magic

$$|T\rangle = \frac{1}{\sqrt{2}} (|0\rangle + e^{i\pi/4} |1\rangle)$$


Cliffords plus $|T\rangle$ states: universal.

Whatever Cliffords cannot make.

Third resource theory of the talk, same mould.

Magic, defined three ways

The way I think about it: is ρ inside the stabilizer polytope?



$$\text{STAB}_n = \text{conv}\{|s\rangle\langle s|\}, \quad \rho \in \text{STAB}_n ?$$

robustness

$$R(\rho) = \min \left\{ \sum_i |q_i| : \rho = \sum_i q_i |s_i\rangle\langle s_i| \right\}$$

negative weight needed to reach ρ from the polytope. A monotone. An LP with exponentially many vertices.

fidelity, extent

$$F(\psi) = \max_s |\langle s|\psi\rangle|^2, \quad \xi(\psi) = \min \|c\|_1^2 \text{ over } |\psi\rangle = \sum_i c_i |s_i\rangle$$

the nearest vertex. What the best simulators actually pay.

stabilizer entropy

$$M_\alpha(\psi) = \frac{1}{1-\alpha} \log_2 \left(\frac{1}{2^n} \sum_P \langle \psi | P | \psi \rangle^{2\alpha} \right)$$

flatness of the Pauli spectrum. Computable, measurable. Does not see the polytope.

Three objects. For mixed states only the polytope sees the convex hull: an exponential LP.



Why everybody talks about magic now

It became measurable, and it became expensive.

1

Measurable

Stabilizer entropy: zero on stabilizer states, computable, measurable by Pauli sampling.

A number people can plot.

2

Expensive

Cliffords cheap. Every T gate eats a distilled magic state.

Magic is the currency of a fault-tolerant machine.

Currency is not power. We have been here twice.

Magic: the same lesson, part 1

Maximal magic without power. Power that does not scale with magic.

Product magic states

$$M_2(|T\rangle^{\otimes n}) = n \log_2 \frac{4}{3}$$

Extensive magic, product state, $2n$ numbers.

Magic without entanglement does nothing.

The best known simulator


$$\text{cost} \leq \text{poly}(n) \cdot 2^{0.23t}, \quad t = \# T \text{ gates}$$

An upper bound on simulation cost. No matching lower bound is known, and proving one would separate BQP from BPP.

What the cost tracks is the T-count of the circuit, not the magic of the state.

Magic: the same lesson, part 2

The same state, through two lenses.



Change what you call free and the magic is gone.

Third time: coherence and a basis, entanglement and Cliffords, magic and matchgates.

A free theory with no simulator

Integrable models: exact spectra, and no efficient dynamics.

What integrability gives you

Extensively many conserved charges. The system does not thermalise.

Bethe ansatz: exact eigenvalues and eigenstates.

Solvable in every sense a physicist means.

Free fermions and matchgates are the easy corner of this.

What it does not give you

The spectrum comes in closed form. Correlation functions do not follow from it: they need determinant formulas and are expensive in general.

Generic eigenstates at finite energy density carry volume-law entanglement, so no matrix product state describes them either.

Dynamics is a separate problem: after a quench, entanglement grows linearly in time, so tensor-network evolution dies at long times.


No efficient classical simulation in general.

Exactly solvable, and computationally hard. Which of the two is the free theory?

The power of one qubit

Knill and Laflamme, PRL 81, 5672 (1998).

One clean qubit, n maximally mixed ones, one controlled unitary, one measurement.



$$\langle X \rangle = \frac{\text{Re Tr } U}{2^n}$$

The normalised trace is the average of $\langle x | U | x \rangle$ over all basis states. For time evolution it is the Fourier transform of the density of states.

Why it mattered

Sampling it classically would collapse the polynomial hierarchy (Morimae, Fujii, Fitzsimons 2014).

The register is maximally mixed the whole time.

Negativity: zero across the clean-qubit cut, bounded elsewhere (Datta, Flammia, Caves 2005).

A conditionally hard computation with essentially no entanglement. So what is doing the work?

A number for the gap: discord

Ollivier and Zurek, PRL 88, 017901 (2001). Henderson and Vedral, J. Phys. A 34, 6899 (2001).

$$I(A:B) = S(A) + S(B) - S(AB)$$

$$J(A|B) = S(A) - \min_{\{\Pi_j\}} \sum_j p_j S(\rho_{A|j})$$

$$D(A|B) = I(A:B) - J(A|B) \geq 0$$

$$D = 0 \iff \rho = \sum_j p_j \rho_A^{(j)} \otimes |j\rangle\langle j|_B$$

Two mutual informations that agree classically. The gap is what a measurement on B must disturb.

Datta, Shaji, Caves 2008: DQC1 has no entanglement and nonzero discord.

Why some people loved it

A real gap: quantumness of mixed states, in an NMR decade.


A headline algorithm with a number attached.

Closed form for two qubits. A curve for every noisy-dynamics paper.

Several thousand citations. Some of them non-silly.

What we learned

Ferraro, Aolita, Cavalcanti, Cucchietti, Acín, PRA 81, 052318 (2010).



orange: generic states, all with discord

The gap it named is real and still open. The number was not the answer to it.

Three cracks

Almost every state has it. A number that is nonzero almost everywhere separates nothing.

It can grow under local operations on the unmeasured side.

The DQC1 link never became a theorem.



Honourable mentions

Other corners of the Hilbert space, and the ones I have forgotten.

Steering

Between entanglement and nonlocality, with its own free set: local hidden state models.

Wiseman, Jones, Doherty (2007)

Non-Gaussianity

The continuous-variable version: Gaussian states and operations are simulable.

Bartlett, Sanders, Braunstein, Nemoto (2002)

Imaginariness

A resource theory of the complex phase. Too early to say.

Wu et al., PRL 126, 090401 (2021)

Nonlocality without entanglement

A product basis that LOCC cannot tell apart. No Bell violation anywhere: this one is about the operations.

Bennett et al., PRA 59, 1070 (1999)

Contextuality

In odd dimensions it is magic under another name. Backup slide if you want it.

Howard, Wallman, Veitch, Emerson (2014)

Anyons

Braiding Ising anyons gives Clifford. Fibonacci is universal. The same lesson again.

Nayak et al., RMP 80, 1083 (2008)

Athermality

Thermodynamics has its own free theory and its own monotones.

Brandão et al., PRL 111, 250404 (2013)

Each of these is somebody's whole field, and I have surely forgotten some. Tell me afterwards.



A resource is the shadow a simulator casts.

Fix what is free, and something looks like the resource. Change it, and the resource changes.
Power is a property of a computation against the best classical algorithm. Unconditional separations: open.

Magic, so far

Where Act 3 leaves us.

Magic, so far

necessary	yes, given Cliffords
sufficient	no
does more mean more power	no

free

product states, local gates
a basis
Cliffords
matchgates
positive Wigner

looks like the resource

entanglement
coherence
magic
non-Gaussianity
negativity

Whatever we make of resources, can we at least measure one?


A background network diagram consisting of a complex web of thin grey lines connecting various black and grey dots of different sizes. The dots are distributed across the entire slide, with a higher density in the lower half. The lines form a series of interconnected triangles and polygons, creating a mesh-like structure.

$\langle aQa^L \rangle$

4. Can we even tell?

A convex body you cannot describe

Separability is NP-hard (Gurvits 2003). Witnesses: Terhal, Phys. Lett. A 271, 319 (2000).



$$\sigma = \sum_k p_k \rho_k^A \otimes \rho_k^B \quad (\text{Act 2})$$


$$\text{Tr}(W\sigma) \geq 0 \quad \forall \sigma \text{ separable}, \quad \text{Tr}(W\rho) < 0$$

Nobody checks membership. You find a hyperplane and measure one observable.

You never test membership. You measure one observable instead.

The shape of the body

Where every criterion you have heard of lives.



$$\sigma = \sum_k p_k \rho_k^A \otimes \rho_k^B \quad (\text{separable, from before})$$

$$\text{PPT: } \rho^{TB} \succeq 0 \quad (\text{necessary; sufficient for } 2 \times 2, 2 \times 3)$$

Inside the ball: separable, no witness needed. Noisy states land here.

PPT: the first computable outer shape. Exact for 2×2 and 2×3 .

Detection is a question of which shape you can afford, and how far it sits from the body.

Every criterion is a shape in this gap.

The year 23 after Gurvits

The whole quantum information community is convinced the separability problem is hopeless. The whole community?

“The whole of quantum information has given up on the separability problem and is occupied by witnesses and relaxations. The whole of it? No. One small village of indomitable researchers still holds out.”

Gurvits, 2003: separability is NP-hard. The field, reasonably, moved to witnesses and hierarchies.

A small village kept working on the exact body.

The village

Lewenstein, Sanpera, the Horodeckis, Nechita, and friends.

The craft: fix what is free, and characterise, exactly, what does not come for free.

Two things the village found, in families where the body can be named.

Make it as simple as you like

Diagonal symmetric states: separability becomes a classical cone.

$$\rho_{DS} = \sum_{i \leq j} p_{ij} |D_{ij}\rangle\langle D_{ij}|, \quad M(\rho) = (p_{ij})$$

$$\rho_{DS} \text{ separable} \iff M \in \mathcal{CP}_d, \quad \rho_{DS} \text{ PPT} \iff M \in \mathcal{DNN}_d$$


$$\mathcal{CP}_d = \mathcal{DNN}_d \iff d \leq 4$$

PPT is enough for local dimension at most four; PPT-entangled DS states exist from five on.

Witnesses here are copositive matrices.

Strip it down to one matrix of populations. Membership is still NP-hard.

Tura, Aloy, Quesada, Lewenstein, Sanpera, Quantum 2, 45 (2018). Diananda (1962); Maxfield and Minc (1962).




Round on purpose: neither cone is polyhedral.

Absolute separability

Separable under every global unitary. A question about spectra, and one of the newer threads.

absolutely maximally entangled (AME): $\rho_S = \frac{\mathbb{I}}{d^{|S|}} \quad \forall |S| \leq \lfloor n/2 \rfloor$

absolutely separable: $U\rho U^\dagger$ separable $\forall U$



Global unitaries are free, so it is a property of the spectrum alone: a region in the simplex of eigenvalues.

Inside: the Gurvits-Barnum ball. Around it: absolutely PPT, a finite list of linear matrix inequalities (Hildebrand 2007).

Open: is absolutely PPT the same as absolutely separable? Known for $2 \times n$.

AME: combinatorics. This: convex geometry. Same question, opposite corner.

A newer thread, and one I would hand to a student.

The recipe: fix what is free, characterise what is not, and know when the computable shape is the body.

How you actually compute a witness

Doherty, Parrilo, Spedalieri, PRL 88, 187904 (2002); PRA 69, 022308 (2004).

ρ_{AB} separable $\Rightarrow \exists \rho_{AB_1 \dots B_k} \succeq 0$ symmetric in $B_1 \dots B_k$

$\text{Tr}_{B_2 \dots B_k} \rho_{AB_1 \dots B_k} = \rho_{AB}$ (level k : one SDP)

Each level is one semidefinite programme. Fail, and the dual is a witness.

The hardness does not vanish. It moves into the level you need.

In progress with Koenraad and PhD students in the group: the volume of the set of genuinely multipartite entangled states.

Same recipe for Bell inequalities, ground states, spectral gaps, and, soon, magic.

The pattern


A hard convex body.

Easy bodies around it, shrinking.

Each one hands you hyperplanes.

Bell nonlocality, geometrically

The most operational definition of all, and a polytope.



$$p(ab|xy) = \sum_{\lambda} q(\lambda) p(a|x, \lambda) p(b|y, \lambda)$$

Two labs, no communication, only statistics. Nothing assumed about the devices.

Local: a mixture of deterministic strategies. So the free set is a polytope and every Bell inequality is a facet.

Deciding: a linear programme with exponentially many vertices. The classical bound is NP-hard in general.

The most operational definition in the talk, and the same wall.



Device-independent, at scale

A long line of work. Our part of it is small.


2015	Loophole-free Bell test, Delft	nothing assumed about the devices <i>Hensen et al., Nature 526, 682</i>
2014	Bell inequalities for many bodies	symmetry: 2^n terms down to a handful <i>Tura et al., Science 344, 1256</i>
2025	73 superconducting qubits	Bell-correlation depth 24, from an energy measurement <i>Wang et al., PRX 15, 021024</i>
2026	Classical bounds by tensor networks	tropical contraction <i>Hu and Tura, PRL 136, 100202</i>

Bell nonlocality returns in a different role, as the strictest certificate we have.

Our contribution: making such witnesses cheap enough to compute and to measure.

Draw it again, for magic

The stabilizer polytope.



$$\text{STAB}_n = \text{conv}\{ |s\rangle\langle s| : |s\rangle \text{ stabilizer} \}$$

The same picture as Act 4, with a polytope in place of the convex body.

You cannot measure every Pauli at once. Which witnesses can you afford?

Same open problem, many groups on it.

Magic witnesses in mixed states: [arXiv:2504.18098](https://arxiv.org/abs/2504.18098). Bell inequalities as magic witnesses: [arXiv:2503.18734](https://arxiv.org/abs/2503.18734).

Which witnesses can you afford?

One modest step. Liu, Gasull, Hu, Zhang, Baccari, Tura, arXiv:2607.26154 (2026).

The problem

Limited Pauli measurements: a reduced polytope, the shadow of the full one on what you can see.

Which Paulis commute decides that shadow.

Two obstructions: what commutes, and the signs.

What we found

The signs drop out exactly when there are no active dependencies.

Perfect frustration graph: efficiently solvable, bounded by the clique number.

Trade-off: measurability against witness power.

For one class of measurement sets: exactly when it is tractable, and what it costs. The general case is open.

Natural definitions, hard questions

Every one of these is a sensible way to draw a line. Only one is easy to check.

resource	operational definition	free set, geometrically	is it there?
coherence	diagonal in a chosen basis	a simplex (a segment, for a qubit)	trivial: read the off-diagonals
discord	classical on one side	not convex, measure zero	efficient, and uninformative
entanglement	Alice, Bob, a telephone line	convex hull of product states	NP-hard (Gurvits 2003)
nonlocality	two labs, no communication, statistics	a polytope; Bell inequalities are facets	LP with exponentially many vertices
magic	preparable with Cliffords	a polytope with $\sim 2^{(n^2/2)}$ vertices	an exponential LP

The more natural the definition, the harder it is to say whether the resource is there. Hence hyperplanes.



Even deciding is hard

Pseudoentanglement (Aaronson et al. 2022). Pseudomagic (Gu, Oliviero, Leone 2024).

Pseudo-resources

Two ensembles, very different resource content, indistinguishable in polynomial time.

The amount is not even decidable.

Full circle

Whatever makes the machine hard to simulate makes its resources hard to certify.

Certification is itself a quantum limit.

Witnesses choose where to hit the wall.



5. What survives

The one theorem with no assumptions

Bravyi, Gosset, König, Science 362, 308 (2018). Noise-robust, and in 3D: Nat. Phys. 16, 1040 (2020).

The statement

Constant-depth quantum beats constant-depth classical. No assumptions.

The problem is the 2D hidden linear function. Bounded fan-in classical circuits need depth of order $\log n$.

The separation is small, and it is the only unconditional one.

The engine

A 2D grid, a GHZ-type state, a nonlocal game. A light cone cannot fake it.

The proof runs on nonlocality.

Nonlocality returns a second time: the only unconditional advantage we have runs on a resource the community had set aside.

Where this leaves the machines

The resource picture becomes an engineering budget.

1

Certification

first steps on real hardware, small sizes

2

Interplay

entanglement and magic together;
neither alone does anything

3

Fault tolerance

magic is what a fault-tolerant machine
spends

One question per theme: can any of this be certified in your systems?



Which resource? None of them on its own.

All of them, each relative to a simulator.
Perhaps “which” was the wrong word.



Advancing
Quantum
Algorithms

We went looking for the Quantum Resource.

We found a history of favourites, a convex body, and one theorem.

The resource itself we did not find. I doubt it is there to be found.

Thank you.



Applied Quantum Algorithms, Leiden, 2025

Liu, Gasull, Hu, Zhang, Baccari, Tura, arXiv:2607.26154 · Hu, Tura, PRL 136, 100202 (2026) · Wang et al., PRX 15, 021024 (2025)
jtura.cat



**Advancing
Quantum
Algorithms**

Backup

Where magic comes from

Howard, Wallman, Veitch, Emerson, Nature 510, 351 (2014). Veitch et al., NJP 14, 113011 (2012).

$$W_{\rho}(u) \geq 0 \quad \forall u \quad \Longleftrightarrow \quad \rho \in \text{conv}(\text{STAB})$$

$$\sum_u |W_{\rho}(u)| = 1 + 2 \text{neg}(\rho)$$

Odd dimension: negativity, contextuality and magic are one thing.

Positive Wigner function: simulable by sampling.

A map of the free theories

Backup. Each region has a simulator. Schematic, not to scale, and the intersections are indicative.

