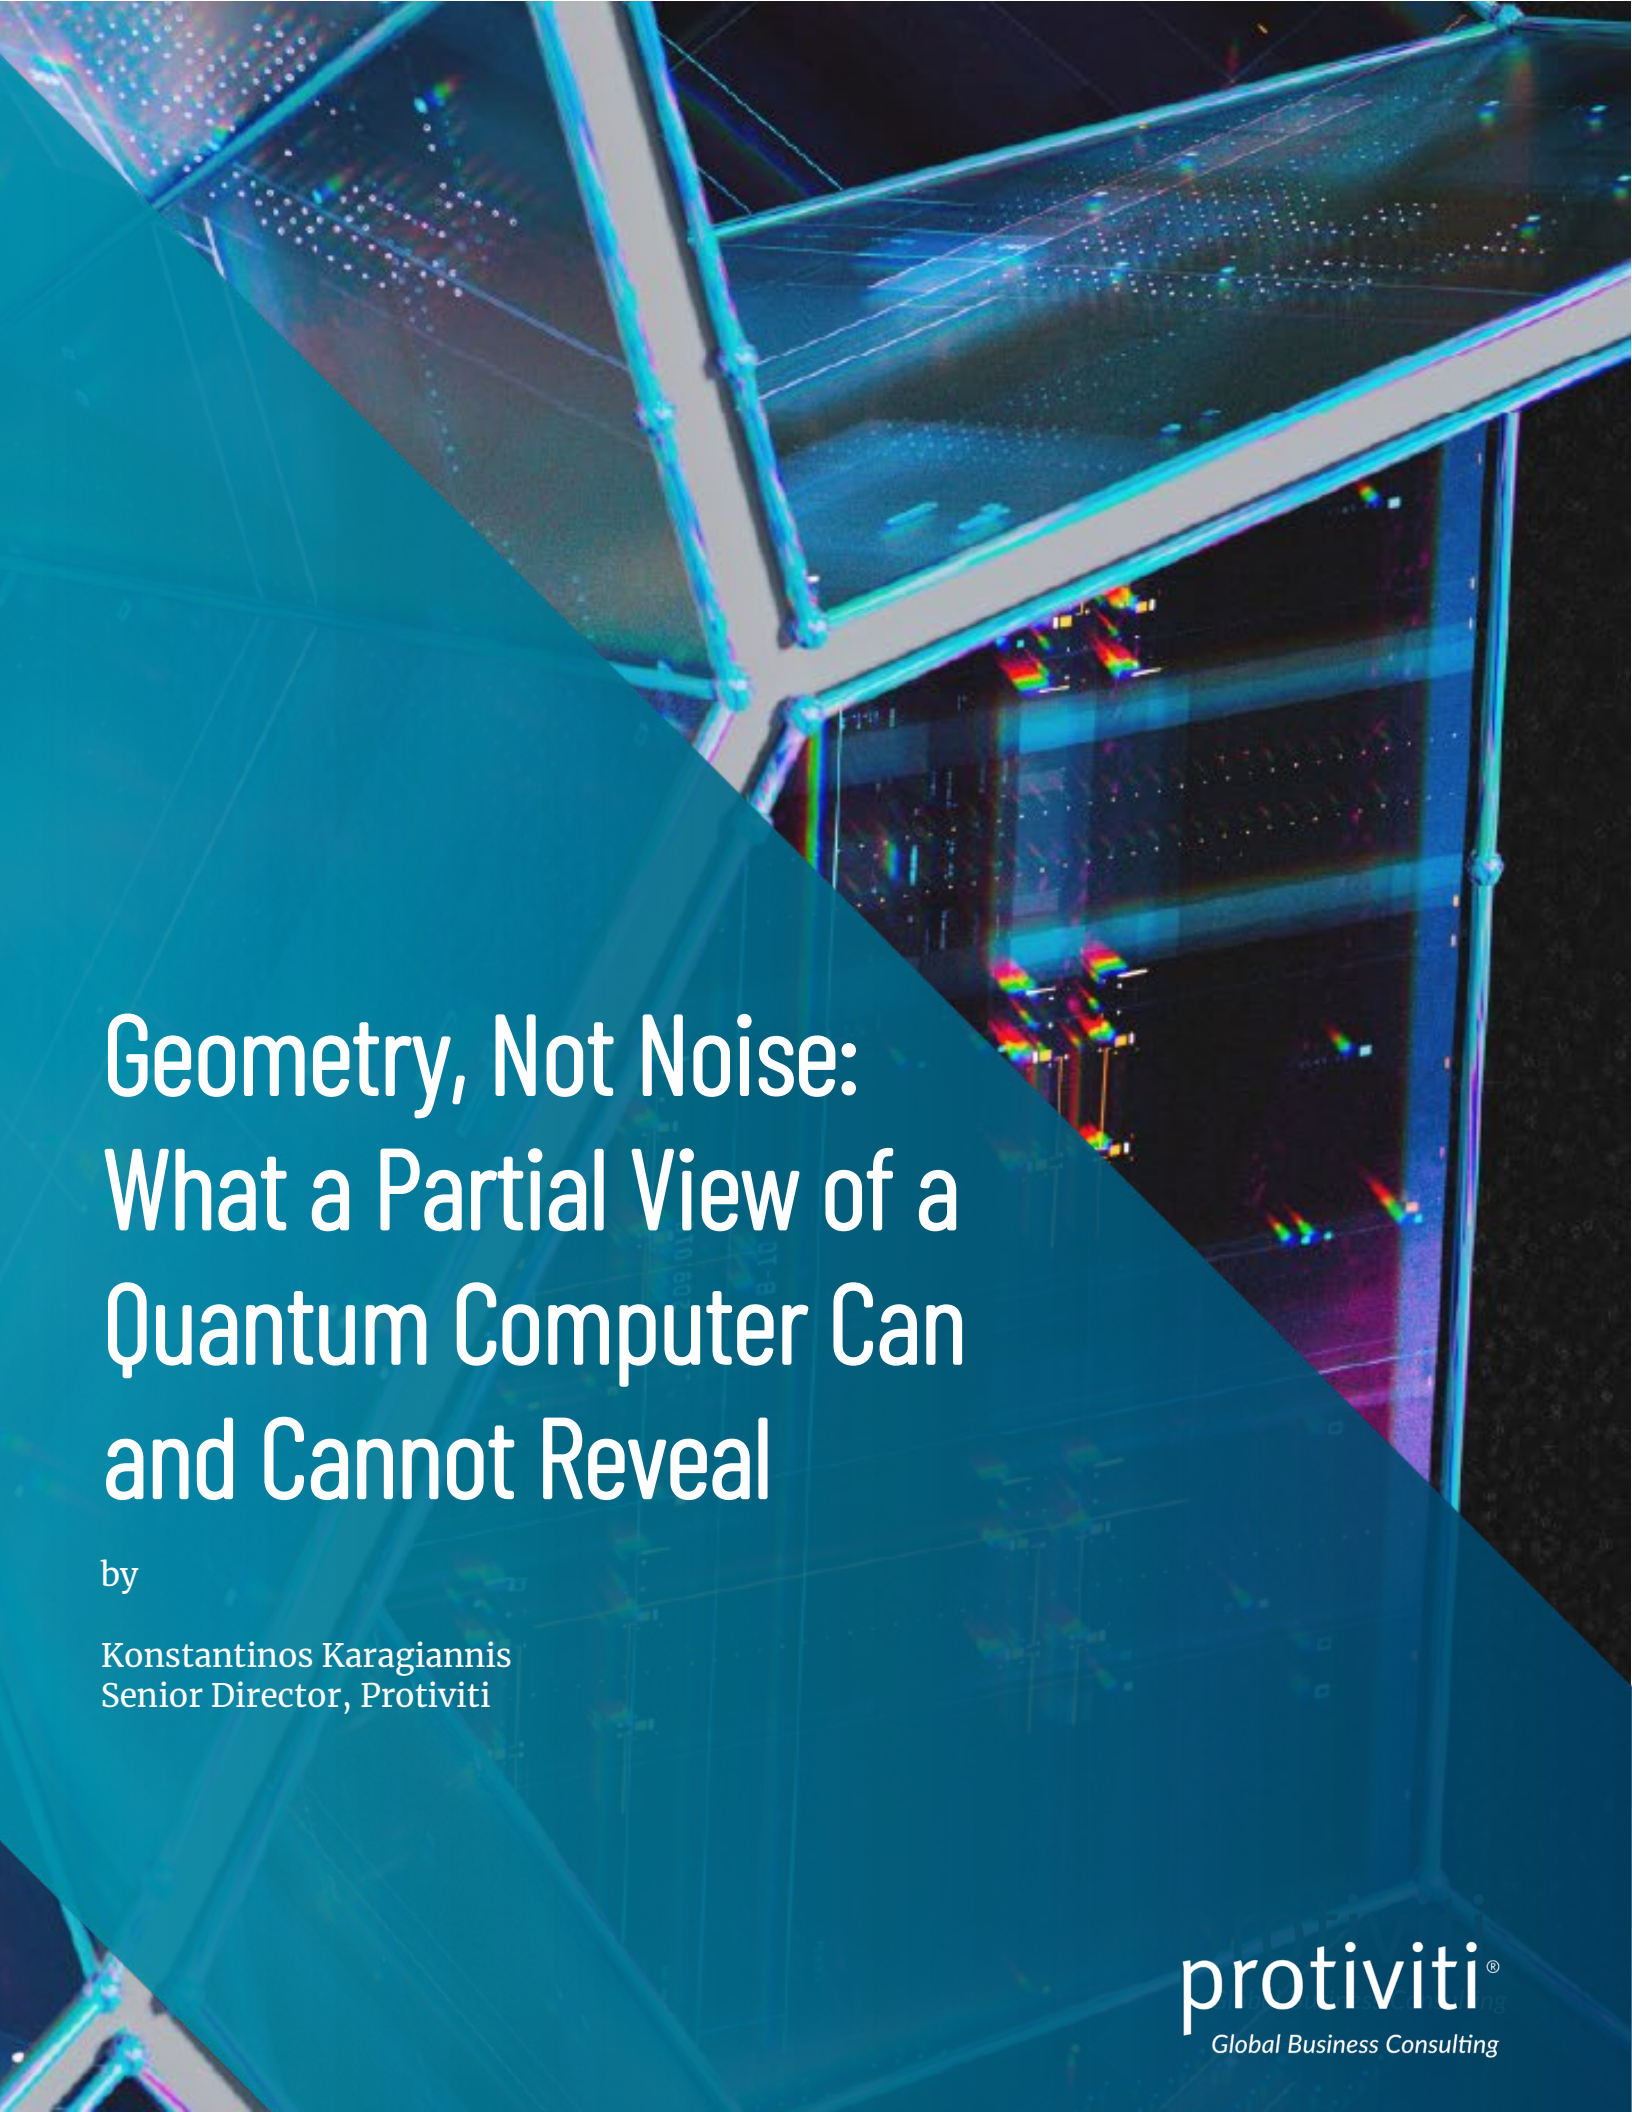


Geometry, Not Noise: What a Partial View of a Quantum Computer Can and Cannot Reveal

by

Konstantinos Karagiannis
Senior Director, Protiviti

protiviti®
Global Business Consulting

Split a password among a thousand guards, and kidnapping a hundred of them should get you nothing. But what if you could kidnap even fewer guards in a specific arrangement and learn more about the password? Weird, right?

Quantum error correction protects information in a similar way, spreading each logical bit across many physical qubits so that most small, local pieces reveal nothing on their own. New Protiviti research shows that whether a partial view leaks depends not on how much of the machine an observer can see, but on the shape of what they see.

Fault-tolerant quantum computers will not keep a logical qubit in one place. They will hand pieces of it to separate decoder modules and route its readout over shared lines. The research asks how much one of those pieces reveals, and the answer cuts both ways. An observer who sees every column of a surface code but one learns exactly nothing about the logical bit. An observer who sees a single row stretching all the way across the code can learn almost all of it. With a side channel that reads each qubit correctly 99% of the time, one such row of a distance-seven code lets an observer guess the logical result correctly about 93% of the time. The same channel missing a single column does no better than a coin flip. Quantum error correction, in other words, is not encryption.

New Protiviti research shows that whether a partial view leaks depends not on how much of the machine an observer can see, but on the shape of what they see.

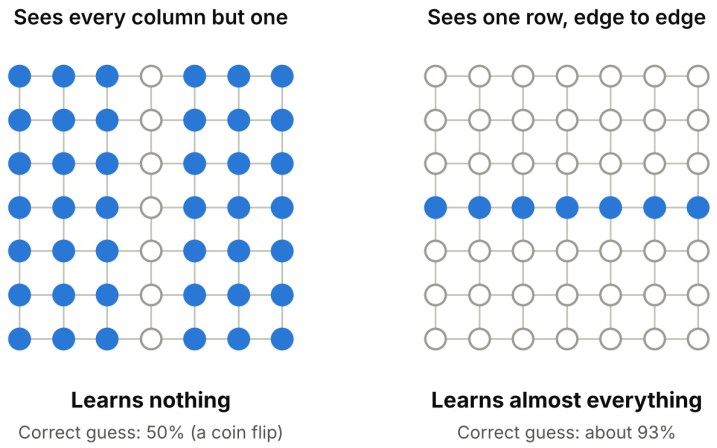


Figure 1. Shape, not size. An observer missing a single column learns nothing; one holding a single row that runs edge to edge can learn almost everything.

The paper, *How wide must an eavesdropper be? Leakage of a toric-code logical bit to a strip and the Nishimori point*, by Konstantinos Karagiannis, Senior Director of Quantum Computing Services at Protiviti, is available on arXiv. It works the question out mathematically, numerically and on IBM quantum hardware, and turns the results into design rules for anyone who has to share part of a quantum computer's readout with a party they do not fully trust.

Why a partial view matters

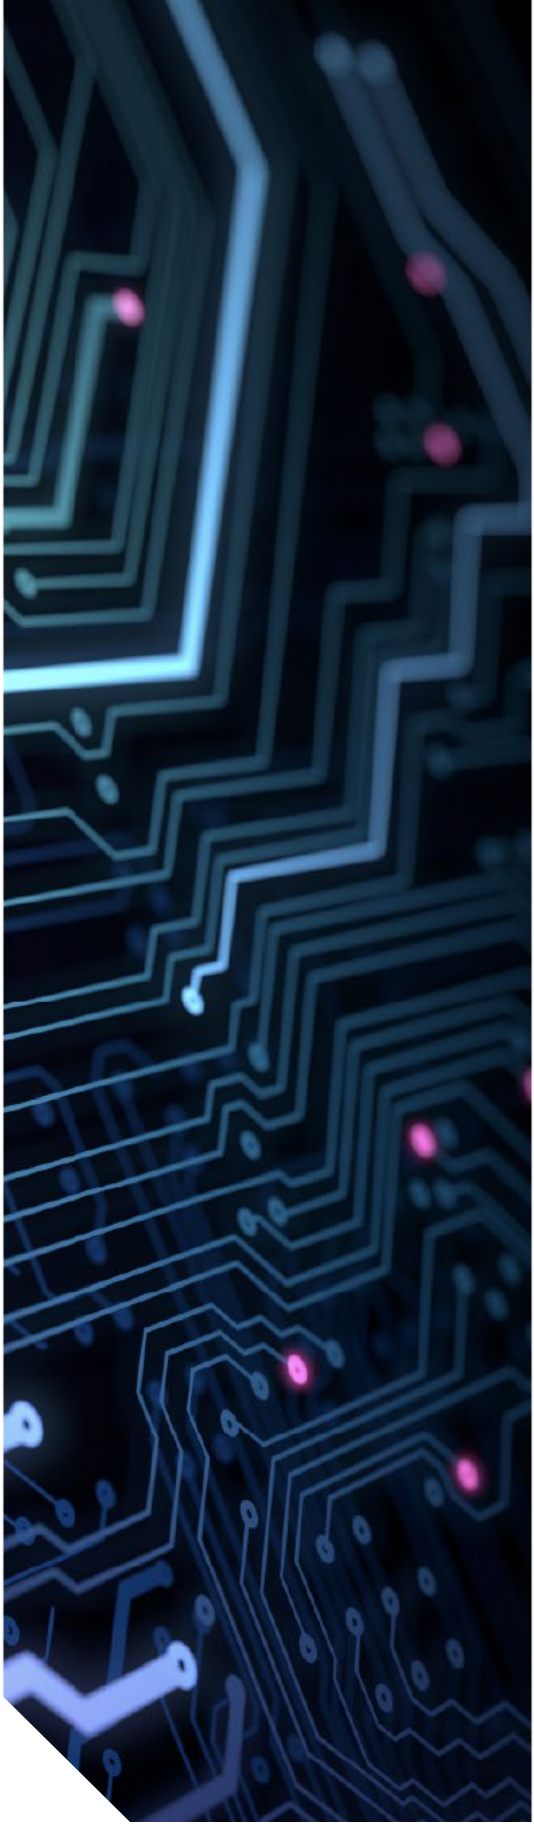
In the surface code, a leading approach to fault tolerance, one logical qubit lives across a two-dimensional grid of physical qubits. A long-established property of such codes is that a sufficiently small patch of the grid reveals nothing about the logical value.

Real machines complicate that picture. Decoding at scale is done in parallel, with regions of the grid assigned to separate classical decoder modules. Superconducting devices share readout lines among several qubits, and crosstalk on those lines can expose one user's outcomes to another. On a shared cloud machine, that amounts to an analog wiretap: a customer whose qubits sit on the same line as yours could, in principle, measure their own qubits and listen for yours. Either way, a party that should not learn a logical result can end up holding the readout of part of the code.

We call this security concern **Regional Exposure**: how much logical information an observer can recover from the part of a system they can see.

One exact rule: a region learns what it can decode

The paper's central result is a novel theorem called the regional leakage–decoding identity. It says that an observer's best possible advantage in telling the two logical states apart from a region is fixed exactly by how well an optimal error-correction decoder would perform if it were confined to that region. In plain terms, what a region can learn is what a decoder limited to that region can recover.



Previous work established when a region of a quantum code can access a logical bit in the absence of noise. This research answers a harder question: when noise is present, how much can that region actually learn? The regional leakage–decoding identity provides an exact answer, connecting the distinguishability of quantum states to the optimal decoding performance available within the observer's region. It bridges two questions that are usually treated separately: how well a quantum code protects information from errors, and how well it protects information from someone who can see only part of the system.

Three things make this useful. It is an equality, not a bound, so it measures exposure exactly. It holds for a region of any shape. And it holds for any Pauli noise, including correlated noise. It also turns a question about quantum information into a decoding problem, which the field already knows how to compute.

One consequence is immediate. If a region does not contain a complete logical operator (a complete logical path that wraps around the code or connects its relevant boundaries), it reveals exactly zero, at any noise level. If it does contain one, its exposure depends on how well it works as a small error-correcting code in its own right. Counting the physical qubits an observer holds tells you neither.

If a region does not contain a complete logical operator (a complete logical path that wraps around the code or connects its relevant boundaries), it reveals exactly zero, at any noise level.

Three results that run against intuition

1. **Missing one column is enough.** An observer whose region misses a single entire column of the code learns nothing about the corresponding logical bit, even with access to every other qubit. The result holds for the toric code and for the rotated planar surface code used on today’s hardware.
2. **Noise is a weak defense, and better hardware weakens it further.** Once a region spans the code, noise is the only thing protecting the logical bit. At a physical flip rate of one in a thousand, a single row of qubits still gives an observer half the maximum advantage at a code distance of about 350 (roughly a quarter of a million physical qubits for one logical qubit), far beyond any planned device. Against a three-row strip, the same protection takes a distance on the order of 40,000, which means

billions of qubits. Lower error rates make this exposure worse, not better.

3. **Width is a poor guide to risk.** One of the more surprising results is that, under uniform, independent noise, adding a second row provides exactly zero additional information to an observer. This is not an approximation or a numerical observation: it is an exact mathematical result, valid at every code distance and noise rate. The second row becomes useful only when real hardware introduces non-uniform error rates that an attacker can exploit. Yet below the error-correction threshold, the width an eavesdropper needs for a near-complete read grows only logarithmically with the code distance. Making the code much bigger buys little protection against an observer who gets slightly wider. The exposure disappears only above the threshold where error correction itself fails, and there the logical information is lost to everyone. The memory is perfectly secure only because it is perfectly useless.

There is also a deeper connection to statistical physics. The point at which regional leakage changes its behavior is not an unrelated security threshold: through the leakage–decoding identity, it is tied directly to the Nishimori point of the random-bond Ising model, a well-studied transition in statistical mechanics. This connects the question of who can learn a quantum computer's logical information to the same mathematical framework used to understand phase transitions in disordered systems. In other words, the same critical point that governs whether errors can be decoded also governs whether a spanning observer can recover the logical information.

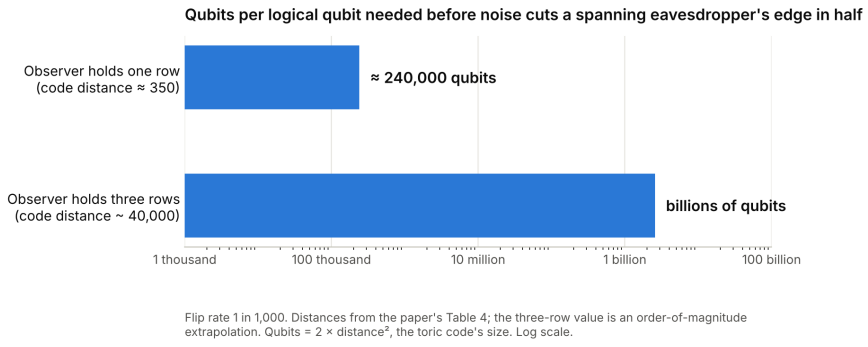


Figure 2. Once a region spans the code, noise is a weak defense. A third row can change the security picture dramatically, even when a second row adds nothing under uniform noise.



What real hardware adds

We prepared encoded surface-code states at code distances three and five on two IBM Heron processors, in five runs over three days, and emulated an observer restricted to a strip of the measured data.

The predicted jump in exposure at a width of three rows appeared on every run. The more interesting result concerns the second row, which theory says is worthless under uniform noise. On real devices qubit quality varies, and at distance five an attacker trained on the data profited from the second row on every run, in effect routing around the worst qubits column by column, much as a driver takes a side street around a stretch of potholes. A model with independent but unequal per-qubit error rates reproduced 90–99% of that gain, so the effect comes from non-uniform noise and correlations add very little.

Simulations built from each run's calibration underestimated the effective flip rate in all ten cases, by 6% to 42%.

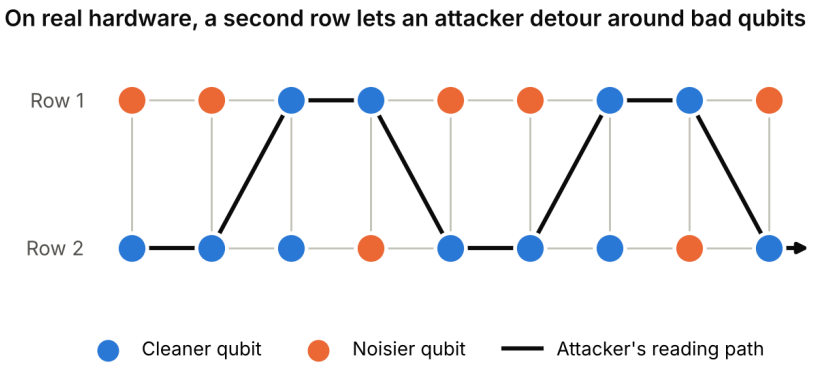


Figure 3. Illustration of the routing effect seen on IBM Heron processors: where one row is noisy, a trained attacker leans on the other.

Device calibration data proved a poor guide. Simulations built from each run's calibration underestimated the effective flip rate in all ten cases, by 6% to 42%. An eavesdropper exploits the noise present during the computation; a defender relying on calibration figures is reading yesterday's map.

These experiments use small codes, and the observer is emulated in software from the full readout. They show how regional exposure behaves under real device noise. They do not demonstrate a compromise of an IBM readout channel or a working attack on a deployed quantum service.

Design rules for untrusted readout

- **Partition around logical operators, not qubit counts.** Each decoder window, including its buffer, and each group of qubits sharing a readout line should miss at least one full column of the code, and one full row for the conjugate measurement basis. Compact blocks smaller than the code in both directions meet this test. A readout line running along a full row or column does not. Think of it as a moat: one full column, and one full row, that no untrusted region crosses.
- **Protect the moments when logical outcomes are exposed.** Routine syndrome data reveal nothing about the logical value on their own. The exposure arises when a logical measurement is read out and during lattice-surgery operations. The final parity that defines a logical result should be assembled only on trusted hardware.
- **Assess combinations, not just components.** Observers who pool their data act as one observer of the combined region, and two regions that are each safe can together span the code. Repeated exposure of the same logical value also increases what a spanning region learns. The geometric guarantee survives both: a combined region that still misses a column learns nothing, however many times it looks.
- **Measure the noise you actually have.** Where a partition cannot avoid spanning the code, assess it with measured per-qubit error rates. Averages and calibration figures can hide the differences an attacker would exploit.



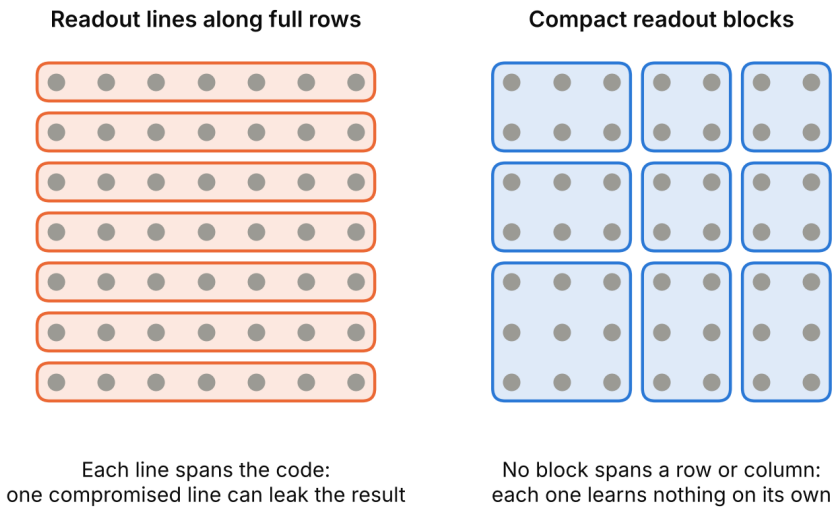


Figure 4. Partition around the code's geometry. Readout groupings that never span a full row or column keep each group blind to the logical result.

What this means for quantum cybersecurity

This suggests a different way to think about quantum cybersecurity. Instead of asking only how many qubits a component can access, designers can evaluate the logical information that its particular view of the system makes recoverable. That distinction could become important when designing distributed decoders, multi-tenant quantum processors, and hardware with shared measurement infrastructure.

Quantum error correction is not confidentiality. A code built to protect information from errors does not stop an observer with the right subset of measurement outcomes from reading it, and on that question geometry matters more than noise.

Machines large enough for this to matter in production are still ahead of us, but their decoder layouts, readout wiring and multi-tenant service models are being designed now. We propose Regional Exposure Assessment as part of that work: evaluating, for each party and each component that touches measurement data, how much logical information its view could recover. The regional leakage–decoding identity gives that assessment an exact mathematical basis. Securing a quantum computation means controlling who can observe what comes out of it, as well as protecting the quantum state itself.

Quantum error correction is not confidentiality.

About the author



Konstantinos Karagiannis

Senior Director, Protiviti

Konstantinos is the Senior Director of Quantum Computing Services at Protiviti, where he helps companies get ready for quantum opportunities and threats. He has been involved in the quantum computing industry since 2012, and in InfoSec since the 90s. Konstantinos is a frequent speaker at RSA, Black Hat, DEF CON, and dozens of conferences worldwide, and was named one of the Top 25 Most Influential People in Quantum Security by Techstrong. He also hosts The Post-Quantum World, which is the longest continuously running quantum computing podcast. He can be reached at konstantinos.karagiannis@protiviti.com.

Protiviti (www.protiviti.com) is a global consulting firm that helps clients transform and protect their businesses, and respond to planned and unexpected events. Through a network of more than 90 offices in over 25 countries, Protiviti and its independent and locally owned member firms deliver deep expertise and tailored capabilities across technology, artificial intelligence, data, operations, finance, legal, compliance, HR, marketing, digital, risk, and internal audit—enabling organizations to accelerate innovation, navigate risks and safeguard what matters most.

Named to the Fortune 100 Best Companies to Work For® list since 2015, Protiviti Inc. has served more than 80 percent of Fortune 100 and nearly 80 percent of Fortune 500 companies. The firm also works with government agencies and smaller, growing companies, including those looking to go public. Protiviti Inc. is a wholly owned subsidiary of Robert Half (NYSE: RHI).

© 2026 Protiviti Inc. An Equal Opportunity Employer M/F/Disability/Veterans. 07/26
Protiviti is not licensed or registered as a public accounting firm and does not issue opinions on financial statements or offer attestation services.

protiviti®