Quantum Laboratory Notebook

Chapter 3 — Chasing Tsirelson: How Close Can a Public Quantum Processor Get to 2√2?

Author: Alessandro Pezzali
Project type: Independent personal research and educational project
Series: Quantum Laboratory Notebook — Chapter 3 (companion to the Bell–CHSH and GHZ notebooks)
Platform used: Quantum Inspire — Tuna-17 backend
Language: cQASM 3.0
Date of experiment: 22–23 July 2026
Session: Phase A pair-scan and tie-break, 22 July; fringe calibration, definitive CHSH, and readout calibration, 23 July (~20:27 – 20:59)

> **Notebook entry — 23 July 2026**
>
> In Chapter 1 I showed that two qubits on this machine can break the CHSH limit. That was a yes-or-no result. This time I wanted a number: how close to the quantum ceiling — the Tsirelson bound of 2√2 — can this freely accessible processor actually get, and where does the rest of the gap go? So I did it properly. I scanned the chip for its best pair of qubits, measured the true measurement angles instead of trusting the nominal ones, ran the definitive CHSH test, and finally characterized the detectors. The headline number came out *below* Chapter 1's, which stung — until I corrected for the readout and saw where the missing visibility had been hiding all along. This notebook records the chase, the surprises the chip taught me on the way, and an honest accounting of the gap.

---

1. Abstract

Chapter 1 of this series asked a yes-or-no question — can a real quantum processor break the classical limit of the CHSH inequality? — and answered it: on the Tuna-17 processor the CHSH quantity came out at S = 2.585, comfortably above the local-realistic bound of $\mathbf{S \le 2}$. This chapter asks a harder, quantitative one. Quantum mechanics does not merely allow S to exceed 2; it allows it to reach a precise maximum — the *Tsirelson bound*, $\mathbf{S = 2\sqrt{2} \approx 2.828}$ — and no higher. How close can a freely accessible cloud processor actually get to that ceiling, and where, exactly, does the remaining gap go?

The two questions are tied together by a single number. For the entangled Bell state the CHSH value is $\mathbf{S = 2\sqrt{2} \cdot v}$, where the *visibility* v — a number between 0 and 1 — measures how close the hardware comes to a perfect entangled pair. Chasing the Tsirelson bound is therefore the same as chasing $v \to 1$, and every imperfection in the machine shows up as a shortfall in v. I worked up to it in stages: a scan to

find the best pair of physical qubits on the chip, a calibration measurement to pin down the true measurement angles, and then the definitive CHSH run. That run returned

$$S_{\text{raw}} = 2.494 \pm 0.039,$$

which clears the classical bound by about 12.7 standard deviations but reaches only **88.2%** of the Tsirelson maximum — and, honestly, a little *below* the 91.4% of Chapter 1. On its face, the careful version did worse than the first attempt.

The resolution is the most interesting result in the chapter. When I measured the two qubits' readout errors separately and corrected the same data for them, S rose to

$$S_{\text{mitigated}} = 2.736 \pm 0.047,$$

or **96.7%** of Tsirelson. The gap to the quantum ceiling, in other words, lives mostly in the *detectors*, not in the entangled state: what Tuna-17 prepared was already within a few percent of a perfect Bell pair, and it was mainly the act of reading it out that pulled the measured S down. I report S_{raw} as the headline — it is what the machine actually returned — but the corrected value is what tells us where the chase really stands.

---

2. Objective

The goal is a direct continuation of the first two chapters, but with the emphasis moved from *whether* to *how close*:

> **Prepare the entangled Bell state on real hardware, measure the CHSH quantity S, and see how close to the Tsirelson bound of $2\sqrt{2}$ a freely accessible processor can get — then account for the gap that remains.**

Chapter 1 established that the violation happens at all. Here the interesting content is the *shortfall*: the difference between the S the machine returns and the 2.828 it is allowed in principle. That shortfall is not a single anonymous "noise" term. It has parts — the choice of physical qubits, the fidelity of the entangling gate, drift in the calibration, and the errors made when the qubits are finally read out — and the aim of this chapter is to separate them and put a number on each. The most useful single figure is the *fraction of the Tsirelson bound* reached; the most useful single identity is $\mathbf{S = 2\sqrt{2} \cdot v}$, which turns the vague question "how good is the state?" into a visibility v that I can measure directly.

---

3. Theory

3.1 The Bell state and the visibility identity

Every run begins by preparing the two qubits in the entangled Bell state

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle).$$

For a *perfect* Bell state, the correlation between a measurement of Alice's qubit along an angle θ_A and Bob's along θ_B depends only on the angle between them:

$$E(\theta_A, \theta_B) = \cos(\theta_A - \theta_B).$$

A real device never prepares a perfect state. To an excellent approximation the effect of all the small imperfections is to *scale down* every correlation by a common factor **v**, the visibility, without changing its shape:

$$E(\theta_A, \theta_B) = v \cdot \cos(\theta_A - \theta_B), \qquad 0 \le v \le 1.$$

Feeding this into the CHSH combination at the optimal angles gives the identity that governs the whole chapter, **S = 2√2 · v**. A perfect state (v = 1) reaches Tsirelson; a state at v = 0.9 reaches 90% of it. Measuring v therefore predicts the best S the machine can give *before* the CHSH test is even run — which is exactly what the calibration stage below does.

3.2 The CHSH inequality and the Tsirelson bound

For each pair of settings I record the **correlation**

$$E = P_{00} + P_{11} - P_{01} - P_{10},$$

which runs from +1 (results always agree) to −1 (always disagree). The CHSH quantity combines four of them, three added and one subtracted:

$$S = E(a, b) + E(a, b') + E(a', b) - E(a', b').$$

Any local-realistic theory obeys **|S| ≤ 2**. Quantum mechanics allows S to rise to the **Tsirelson bound, 2√2 ≈ 2.828**, and — this is a theorem, not an engineering limit — no higher. That gives a built-in sanity check I used throughout: any analysis that returns **S > 2.828 is wrong** (a normalization slip, a mislabelled outcome, a sign error), to be debugged and never celebrated.

3.3 Measuring along a tilted axis

The four CHSH settings are just four measurement axes in the X–Z plane. Measuring along the axis σ(θ) = cos θ · Z + sin θ · X is equivalent to first applying the rotation **Ry(−θ)** and then measuring in the ordinary Z basis. This one primitive builds both the calibration fringe and the CHSH runs. The angles that maximise S are

$$a = 0, \quad a' = \frac{\pi}{2}, \quad b = \frac{\pi}{4}, \quad b' = -\frac{\pi}{4},$$

which place three of the four Alice–Bob pairs 45° apart (large positive correlation) and the fourth, (a′, b′), 135° apart (large negative correlation) — the "three positive, one negative" pattern the CHSH formula rewards.

4. Experimental Design

4.1 The platform, and a free gift from the compiler

All programs ran on the **Tuna-17** backend of Quantum Inspire — a 17-qubit superconducting processor whose qubits are laid out as the distance-3 *rotated surface code* ("Surface-17"): nine *data* qubits and eight *ancilla* qubits, joined by 24 tunable couplers, each coupler connecting exactly one data qubit to one ancilla. Only those 24 connected pairs support a native two-qubit gate; any other pair is unusable without extra routing.

Two features of the platform shaped the whole design. First, **the compiler maps qubits one-to-one**: the label `q[i]` in my program *is* physical qubit Qi, and a program is never silently re-routed onto other qubits. Second, as a direct consequence, a two-qubit gate between two *unconnected* qubits is **rejected at compile time** rather than being patched with swaps. That rejection is a gift: it turns the compiler into a free **adjacency oracle**. Two trial submissions were enough to confirm the map — a gate on the connected pair (0, 1) compiled cleanly, while a gate on the unconnected pair (0, 8) was rejected — and the full 24-bond topology followed.

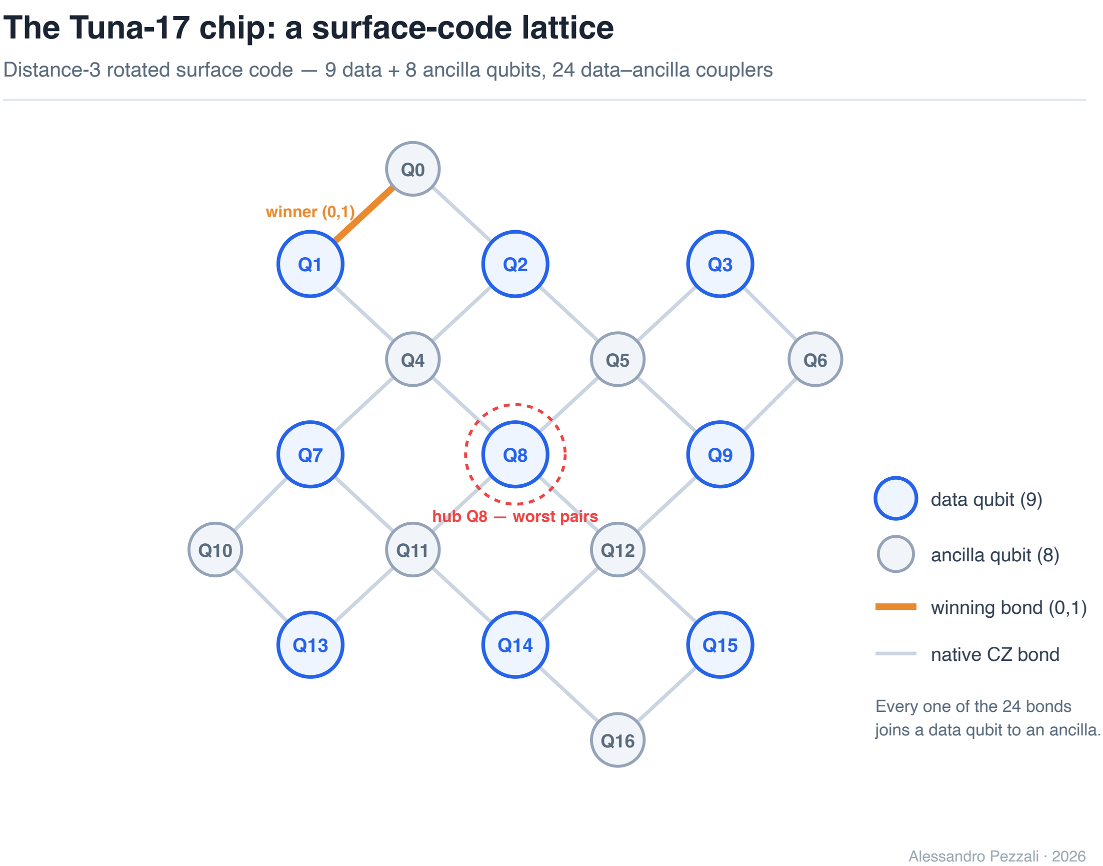

Figure 12. The Tuna-17 chip as a distance-3 rotated surface code: nine data qubits and eight ancilla qubits, 24 data–ancilla couplers. Every native two-qubit gate joins a data qubit to an ancilla. The pair finally used for the experiment, the corner bond **(0, 1)**, is highlighted.

The native two-qubit gate is the **controlled-Z (CZ)**, not the CNOT, so the Bell state is prepared with the CZ-native sequence `H, H, CZ, H` (equivalent to the familiar `H` + `CNOT`). Single-qubit rotations `Ry(θ)` are native but **quantized to the nearest multiple of π/36** (5°). Every angle this experiment needs is a multiple of $\pi/4 = 9 \times \pi/36$, so all of them land exactly on the hardware grid — the quantization costs nothing here.

4.2 Four phases, one continuous session

The experiment is built in four stages, run close together to keep the machine's calibration from drifting between them (Section 7 shows why that matters):

- **Phase A — pair scan.** No public calibration table exists for this chip, so the best pair of qubits has to be found empirically. I probed several connected bonds with two cheap proxies for state quality and kept the best.
- **Phase B — calibration fringe.** Rather than trust the nominal measurement angles, I swept Bob's angle through a full turn and fitted the resulting correlation curve to measure the true visibility v and any systematic angle offset φ_0.
- **Phase C — the definitive CHSH run.** The four CHSH settings, three blocks back-to-back.
- **Phase D — readout calibration.** The four computational basis states prepared and measured to build the readout error (confusion) matrix, run in the *same* session as Phase C.

The winning pair from Phases A–B was the corner bond **(0, 1)** — qubit Q0 (an ancilla) and qubit Q1 (a corner data qubit). With φ_0 measured to be negligible (Section 7.3), the CHSH settings used their nominal angles:

Party	Setting	Angle θ	Realised in the circuit by
Alice (q[0])	*a*	0	measure directly (no rotation)
Alice (q[0])	*a′*	π/2	apply **Ry(−π/2)** before measuring
Bob (q[1])	*b*	π/4	apply **Ry(−π/4)** before measuring
Bob (q[1])	*b′*	−π/4	apply **Ry(+π/4)** before measuring

Table 1. The four measurement settings and how each is realised on the winning bond (0, 1).

Crossing Alice's two settings with Bob's two gives the four CHSH configurations:

Run	Alice θ_A	Bob θ_B	Correlation
C1	0	π/4	E(a, b)
C2	0	−π/4	E(a, b′)
C3	π/2	π/4	E(a′, b)
C4	π/2	−π/4	E(a′, b′)

Table 2. The four CHSH configurations. All four share the same preparation block; only the two final rotations change. S = E1 + E2 + E3 − E4.

5. Methods

Hardware and software. All programs were written in cQASM 3.0 and executed on the Quantum Inspire **Tuna-17** backend, using two qubits and two classical bits. Every run used **512 shots** — enough to rank candidate pairs and to measure S to a few percent, and small enough to fit the whole campaign inside single calibration windows.

State preparation. Every run starts with the same CZ-native Bell preparation on the two qubits of the chosen bond:

```
H q[0]
H q[1]
CZ q[0], q[1]
H q[1]
```

This is the controlled-Z form of "`H` then `CNOT`": it leaves the pair in $|\Phi^+\rangle$.

Measurement settings. After preparation, each run applies its basis rotations (Table 1) and measures both qubits. The CHSH template is

```
H q[0]
H q[1]
CZ q[0], q[1]
H q[1]
Ry(-THETA_A) q[0]
Ry(-THETA_B) q[1]
b[0] = measure q[0]
b[1] = measure q[1]
```

with `Ry(0)` omitted for Alice's a setting.

Readout calibration (Phase D). The four basis states 00, 01, 10, 11 were prepared with `X` gates and measured, 512 shots each, to populate a 4×4 confusion matrix M whose column j is the measured distribution given that state j was prepared. The Phase-C probabilities are then corrected by p_true = $M^{-1} \cdot p$_meas (clipped to be non-negative and renormalized).

A platform convention worth recording. The result screens display each two-bit outcome in **reversed order — b[1]b[0], not b[0]b[1]** — as the dominant bars of the 01 and 10 calibration runs proved directly. Because every quantity in this notebook is built from the *parity* of the two bits (P00 +

P_{11} versus $P_{01} + P_{10}$), which is unchanged by swapping the two bits, none of the correlations or S values are affected. But the convention matters the moment one reads an individual outcome probability — for instance when assembling the confusion matrix — so it is logged here rather than left as a trap. (It joins Chapter 2's measurement sign flip as the second "the hardware taught me a hidden convention" note of the series.)

Session log. The definitive Phase-C and Phase-D runs were executed one after another in a single session on 23 July 2026:

Run	Job IDs	Result IDs	Timestamp (23/07/2026)
C block 1 (C1–C4)	1171572–1171575	1136678–1136681	20:27–20:31
C block 2 (C1–C4)	1171576–1171579	1136682–1136685	20:34–20:35
C block 3 (C1–C4)	1171581–1171584	1136687–1136690	20:43–20:46
D readout (D1–D4)	1171585–1171588	1136691–1136694	20:55–20:59

Table 3. Session log for the definitive runs (all Tuna-17, 512 shots). Program hashes are shared *per setting* across the three C blocks — C1 `1185542@215434c4`, C2 `1185544@a036ae76`, C3 `1185546@8a66fff6`, C4 `1185548@792bf1eb` — and per state for D (`1185550`, `1185552`, `1185554`, `1185556`). The Phase-A scan and tie-break ran earlier (22 July, jobs 1170535 and 1171296–1171338); the calibration fringe ran the same evening, 19:35–20:14 (jobs 1171562–1171571). Per-run details appear in the raw-data tables below and in the Appendix B screenshots.

---

6. Raw Data

6.1 Phase A — the pair scan

Six connected bonds, spread across the chip, were each probed twice at 512 shots: a **ZZ probe** (the bare Bell state, measured directly) reporting the agreement fraction $F = P_{00} + P_{11}$, and an **XX probe** (a rotation into the X basis before measuring) reporting the coherence correlation E_XX. The ranking score is $r = ½F + ½E\_XX$.

Rank	Bond	Roles	*F* (ZZ)	*E*_XX	*r*
1	(2, 5)	edge data – bulk ancilla	0.944	0.926	**0.935**
2	(0, 1)	ancilla – corner data	0.959	0.902	**0.931**
3	(15, 16)	corner data – ancilla	0.946	0.915	**0.930**
4	(7, 11)	edge data – bulk ancilla	0.912	0.891	0.902
5	(4, 8)	bulk ancilla – **centre data (Q8)**	0.904	0.868	0.886
6	(8, 12)	**centre data (Q8)** – bulk ancilla	0.894	0.856	0.875

Table 4. The six candidate bonds, ranked. Both bonds touching the most-connected qubit on the chip, the degree-4 centre data qubit Q8, come *last* — the opposite of the naive expectation.

6.2 Phase A — the tie-break

The top three of Table 4 are a statistical tie at 512 shots ($\sigma_r \approx 0.010$, spread only 0.005), so they were re-measured in a later session to compare their *stability*:

Bond	*r*, session 1	*r*, session 2	change over ~25 min
(0, 1)	0.931	0.931	**±0.004 (stable)**
(2, 5)	0.935	0.907	–0.028
(15, 16)	0.930	0.896	–0.034

Table 5. Stability of the top three bonds across two sessions (tie-break jobs 1171296–1171338, 22 July, 22:09–22:15). Bond (0, 1) barely moved while the other two each dropped by about 0.03 — the first sign of the calibration drift analysed in Section 7. Bond **(0, 1)** was chosen for the rest of the experiment on that stability, not on the highest single score.

6.3 Phase B — the calibration fringe

With Alice fixed at $\theta_A = 0$, Bob's angle θ_B was swept through nine angles from $-\pi$ to $+\pi$ in steps of $\pi/4$ (512 shots each), acquired as ten runs (B8 was queued twice), and the resulting correlation $E(\theta_B)$ was fitted to $v \cdot \cos(\theta_B - \varphi_0)$.

Point	θ_B	P(00)	P(01)	P(10)	P(11)	*E*
B1	−π	0.014	0.508	0.455	0.023	−0.926
B2	−3π/4	0.076	0.406	0.418	0.100	−0.648
B3	−π/2	0.260	0.195	0.264	0.281	+0.082
B4	−π/4	0.426	0.063	0.078	0.434	+0.719
B5	0	0.490	0.020	0.014	0.477	+0.933
B6	+π/4	0.432	0.080	0.074	0.414	+0.692
B7	+π/2	0.266	0.262	0.236	0.236	+0.004
B8	+3π/4	0.092	0.395	0.418	0.096	−0.625
B8′	+3π/4 *(re-run)*	0.076	0.418	0.408	0.098	−0.652
B9	+π	0.012	0.436	0.523	0.029	−0.918

Table 6. The measured fringe (probabilities in the platform's display order b[1]b[0]; *E* is parity-based and so unaffected by that order). The θ_B = 0 point returns *E* = +0.933, which is already the visibility the full fit recovers. Points B8, B8′ and B9 were acquired after a ~33-minute automatic-tuning stall (19:41 → 20:14); B8 was queued twice and both copies agree within statistics, so both were kept as a reproducibility check. That gap between the early and late points is what makes the fringe amplitude itself a small drift probe (Section 7.3).

The calibration fringe

$E(\theta) = v \cdot \cos(\theta - \phi_0)$ on bond (0,1) — fit: $v = 0.932$, $\phi_0 = -1.0° \pm 1.1°$

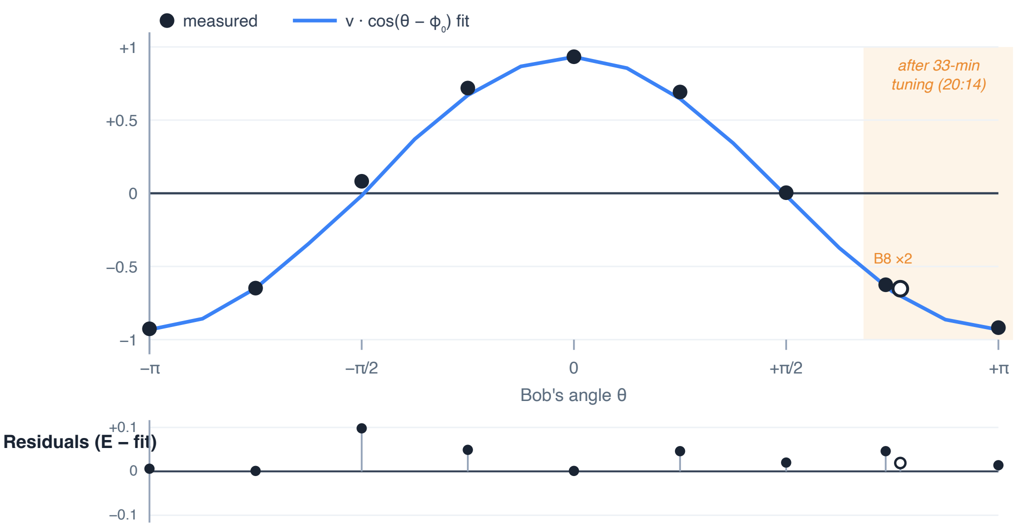

Alessandro Pezzali · 2026

Figure 14. The calibration fringe: ten measured correlations at nine angles with the fitted cosine $v \cdot \cos(\theta\_B - \varphi_0)$ and the fit residuals below. The amplitude of the curve is the visibility v; its horizontal shift is the systematic offset φ_0.

6.4 Phase C — the CHSH runs

Three blocks of the four configurations (Table 2) were run back-to-back. The four correlations in every block followed the designed pattern — three strongly positive legs and one strongly negative — with the negative leg **E4 = E(a′, b′)** consistently the *weakest in magnitude* (about −0.58, against roughly +0.64 for the three positive legs). The per-block CHSH values were

Block	Run	P(00)	P(01)	P(10)	P(11)	*E*	Block *S*
1	C1 E(a,b)	0.402	0.076	0.086	0.436	+0.676	
1	C2 E(a,b′)	0.428	0.076	0.102	0.395	+0.645	
1	C3 E(a′,b)	0.414	0.082	0.104	0.400	+0.628	
1	C4 E(a′,b′)	0.107	0.371	0.416	0.105	−0.575	**2.524**
2	C1 E(a,b)	0.379	0.094	0.100	0.428	+0.613	
2	C2 E(a,b′)	0.404	0.090	0.090	0.416	+0.640	
2	C3 E(a′,b)	0.385	0.098	0.094	0.424	+0.617	
2	C4 E(a′,b′)	0.115	0.383	0.400	0.102	−0.566	**2.436**
3	C1 E(a,b)	0.408	0.084	0.094	0.414	+0.644	
3	C2 E(a,b′)	0.434	0.076	0.096	0.395	+0.657	
3	C3 E(a′,b)	0.439	0.102	0.086	0.373	+0.624	
3	C4 E(a′,b′)	0.082	0.352	0.447	0.119	−0.598	**2.523**

Table 7. The twelve measured correlations (probabilities in display order b[1]b[0]) and the three per-block CHSH values, $S = E1 + E2 + E3 - E4$. In every block the negative leg E4 = E(a′, b′) is the weakest in magnitude (−0.575, −0.566, −0.598) — the one systematic feature that survives into the analysis (Section 7.4).

6.5 Phase D — readout calibration

The four basis states were prepared and measured (512 shots each). The diagonal of the confusion matrix — the probability that each prepared state is read back correctly — is

Prepared state	Readout fidelity
00	0.945
01	0.961
10	0.957
11	0.965

Table 8. Readout fidelities from the four basis-state preparations (jobs 1171585–1171588). Readout is the least accurate operation on the chip — around 4–5% error per state, far larger than the single- and two-qubit gate errors — which foreshadows where the gap to Tsirelson turns out to live.

The full confusion matrix M (columns = prepared state, rows = measured outcome, display order b[1]b[0]) is

Measured ↓ \ Prepared →	00	01	10	11
00	0.945	0.018	0.027	0.000
01	0.014	0.961	0.004	0.023
10	0.041	0.000	0.957	0.012
11	0.000	0.021	0.012	0.965

Table 9. The 4×4 readout confusion matrix. Each column sums to 1.000. The correction applied to the Phase-C data is p_true = $M^{-1} \cdot p$_meas, with the result clipped to be non-negative and renormalized.

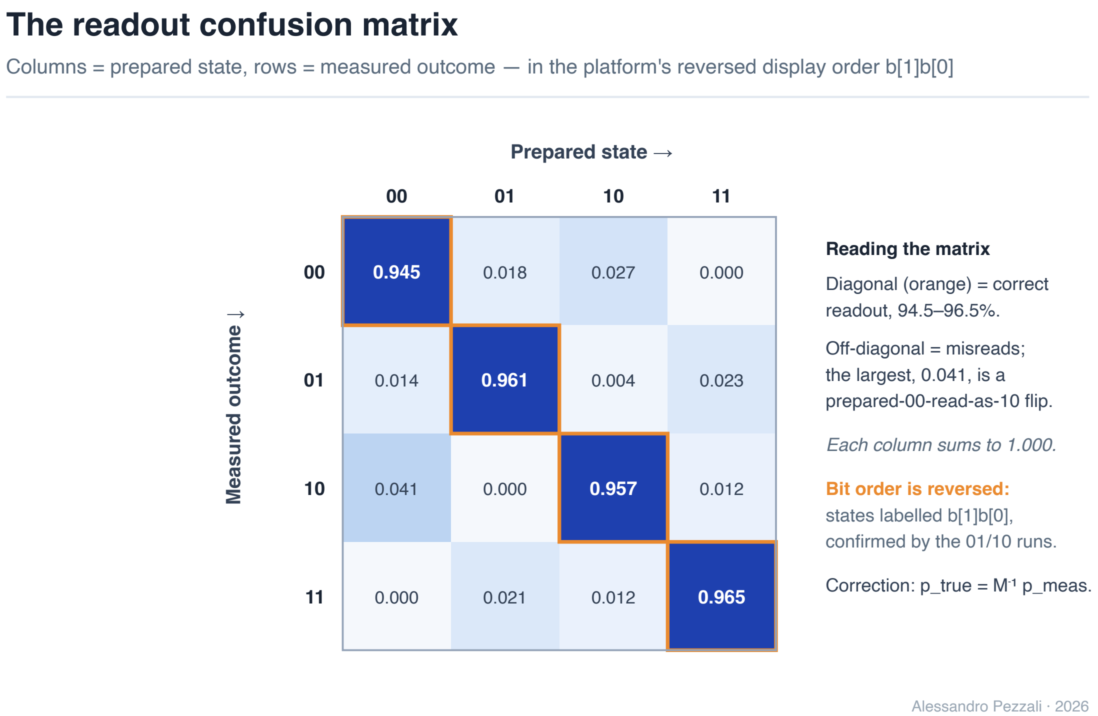

Figure 16. The 4×4 readout confusion matrix M: each column is the measured distribution for one prepared basis state. A perfect detector would be the identity; the off-diagonal weight is the readout error that Section 7.5 removes.

7. Analysis

7.1 Correlations

Each correlation is formed from the four outcome probabilities as $E = P_{00} + P_{11} - P_{01} - P_{10}$. Because this depends only on the *parity* of the two measured bits, the reversed b[1]b[0] display order (Section 5) leaves every correlation untouched — a convenient consequence of the CHSH quantity being built from parities rather than individual outcomes.

7.2 The pair scan: periphery beats centre

The clearest qualitative result of the chapter is in Table 4. The two bonds touching Q8 — the single most-connected qubit on the chip, the degree-4 data qubit at the centre of the surface code — are the two *worst* of the six, by a margin (0.886 and 0.875 against 0.935 at the top) far outside the ≈0.010 statistical error. The intuition that the most-connected qubit would be the best-calibrated is simply wrong here; on this device the quiet **periphery beats the busy centre**. The pair I carried forward, the corner bond (0, 1), was then chosen not for the top score but for the *stability* seen in Table 5 — a choice the drift analysis below justifies.

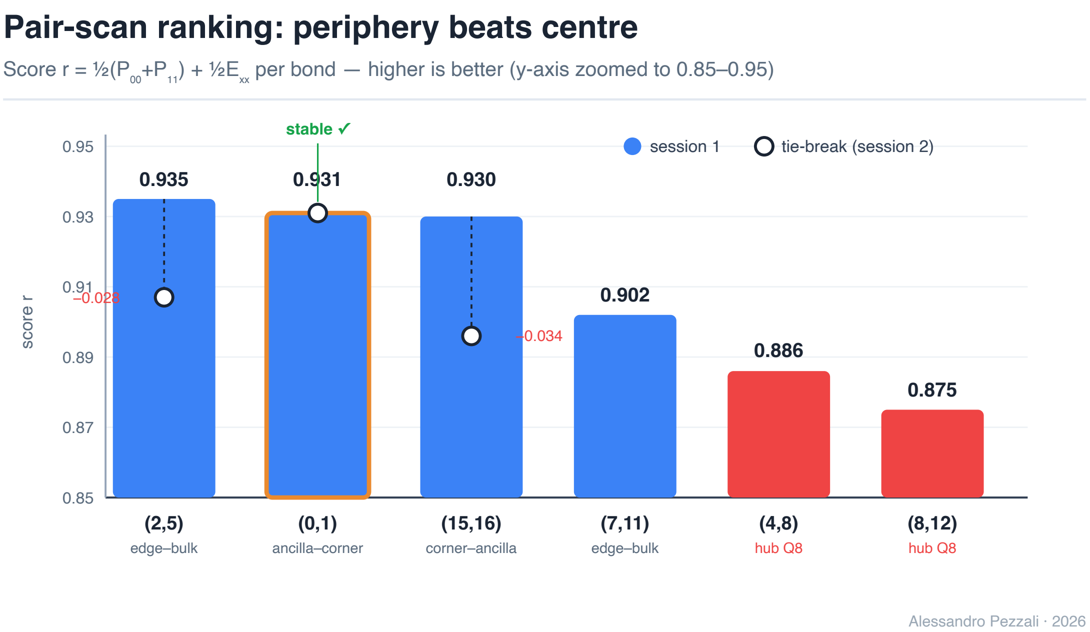

Figure 13. The six candidate bonds ranked by score r. The two bonds on the central hub qubit Q8 (shaded) sit at the bottom.

7.3 The fringe fit: visibility and offset

The fit of the ten-acquisition fringe (Figure 14) to $v \cdot \cos(\theta_B - \varphi_0)$ returned a visibility and a systematic offset of

$$v = 0.932, \qquad \phi_0 = -1.0° \pm 1.1°.$$

Two things follow. First, through **S = 2√2 · v** the visibility predicts a ceiling of **S_max = 2√2 × 0.932 ≈ 2.64 ± 0.05** for the CHSH run *before* it is performed. Second, the offset φ_0 is consistent with zero and, crucially, is **smaller than the 5° angle grid** the hardware quantizes to: there is no correction to apply, because the machine's nominal angles are already true to within the resolution it can even represent. The calibration's honest verdict was "nothing to correct" — so the CHSH run used the nominal angles unchanged.

The fringe carries one more clue. Its last three points (B8, B8′, B9) were taken after a ~33-minute auto-tuning stall; fitting only the earlier points gives a visibility of about **0.952**, while the full fit including the late points falls to **0.932**. That drop — seen *within a single fringe* — is the same calibration drift that Table 5 caught between sessions, here glimpsed on the ~half-hour timescale. It is why the whole campaign was kept to tight, back-to-back bursts.

7.4 The CHSH quantity

Averaging the three blocks of Table 7 gives the definitive raw value:

$$S_{\text{raw}} = \tfrac{1}{3}(2.524 + 2.436 + 2.523) = 2.494.$$

The scatter between blocks (a spread of about 0.05) is consistent with pure shot noise at 512 shots per setting, so there is **no significant drift in *S* across the ~20 minutes** of the three blocks. This is worth stating carefully, because drift *was* real elsewhere (Table 5, and the visibility slipping across the longer calibration window): it lives in the minutes-to-half- hour gaps *between* phases, not inside the tightly packed CHSH blocks themselves. Running Phase C as a single back-to-back burst is what kept it out of the final number.

One systematic feature survives into the analysis: the negative leg **E4 = E(a′, b′)** is the weakest correlation in all three blocks (−0.575, −0.566, −0.598; Table 7). The raw *S* of 2.494 also sits about 2.2σ below the fringe's promised 2.64 (combining σ_S = 0.039 with the ≈ 0.05 uncertainty on the fringe's own prediction), so *some* small cost is unaccounted for. E4 being the consistent weak leg points at the a′ = π/2 *setting* — but that is as far as the data reach: I cannot yet say whether the cause is the extra Ry(−π/2) rotation on Alice's qubit or simply inter-phase drift between the fringe and the CHSH run. The two remain undecided without the dedicated a′-leg diagnostic noted as a follow-up in the Discussion — and, importantly, the chapter's conclusion does not depend on which it is.

7.5 Readout mitigation: where the gap really is

Inverting the confusion matrix of Section 6.5 and applying it to the three Phase-C blocks lifts each of them substantially:

Block	*S*_raw	*S*_mitigated
1	2.524	2.769
2	2.436	2.669
3	2.523	2.770

Table 10. The three blocks before and after readout correction.

Averaging gives

$$S_{\text{mitigated}} = 2.736 \pm 0.047,$$

which is **96.7%** of the Tsirelson bound and, reassuringly, still *below* $2\sqrt{2}$ — the correction does not overshoot into the forbidden region, which would have signalled that it was over-correcting. This single step decomposes the gap:

- from raw to mitigated, $\Delta S \approx$ **0.24** (about 8.5 percentage points of Tsirelson) — this is the **readout error** alone;
- from mitigated to the Tsirelson bound, $\Delta S \approx$ **0.09** (about 3.3 points) — everything else: gate errors, decoherence, and the small a′-leg residual of Section 7.4.

One check that the inversion is not doing anything exotic. Treating readout as an independent error on each qubit, the mean two-qubit readout fidelity of 0.957 implies a per-qubit error of about 2.2%, which scales every correlation by $(1 - 2\varepsilon)^2 \approx 0.915$. Dividing the raw *S* by that single factor gives 2.73 — within 0.01 of the 2.736 the full matrix inversion returns. The confusion matrix is also very well conditioned (its condition number is 1.10), so the correction is a mild, well-behaved rescaling and not a delicate numerical operation that could invent structure out of noise.

In other words the detectors account for roughly *two-thirds* of the entire distance to the quantum ceiling. The entangled state Tuna-17 prepared was already good to about 97%; it was mainly the reading-out that made it look like 88%.

7.6 Uncertainty

Each correlation is estimated from 512 shots, with statistical uncertainty $\sigma_E \approx \sqrt{(1 - E^2)/N} \approx 0.034$ at $E \approx 0.64$. Combining four such correlations in quadrature gives about 0.068 per block, and averaging three independent blocks brings the raw figure to **$\sigma_S \approx$ 0.039**. Readout mitigation *amplifies* statistical noise (inverting *M* mixes the outcomes), so the mitigated value carries the slightly larger **$\sigma_S \approx$ 0.047** — wider despite being "corrected," which is expected and honest. Two systematic terms are small by comparison: angle quantization contributes ≤ 0.003 to *S* (second-order in a $\leq 2.5°$ error), and the between-phase drift is controlled by single-session execution.

That 0.047 propagates the shot noise of the Phase-C runs through the inversion. It does not include the statistical uncertainty of the confusion matrix itself, which was also measured at only 512 shots per column; propagating that as well by Monte Carlo widens the bar to about 0.055. Even at that width

S_mitigated sits 1.7σ below 2√2, so the check of Section 7.5 — that the correction does not overshoot into the forbidden region — survives the more conservative error bar.

8. Results

The experiment yields two numbers, and the distance between them is the point:

$$\boxed{S_{\mathrm{raw}} = 2.494 \pm 0.039} \qquad \text{and} \qquad S_{\mathrm{mitigated}} = 2.736 \pm 0.047.$$

The raw value exceeds the classical limit by

$$\frac{2.494 - 2}{0.039} \approx 12.7 \text{ standard deviations,}$$

so a local-realistic explanation of the four runs is excluded by a wide margin, while S stays safely below the quantum maximum — 88.2% of Tsirelson raw, 96.7% once the detectors are accounted for.

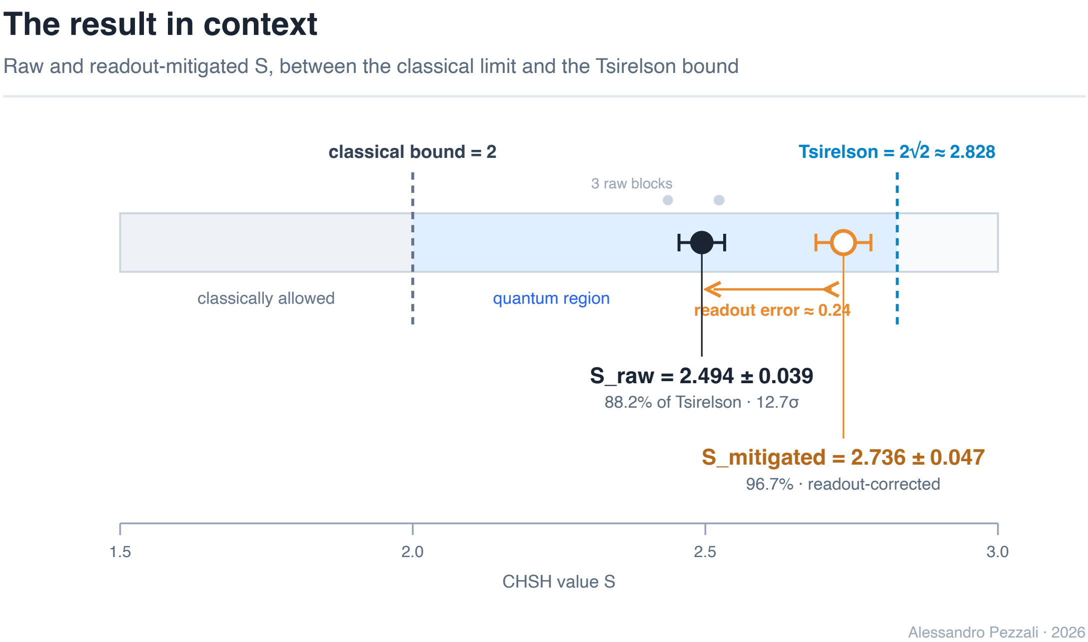

Figure 15. S_raw and S_mitigated on the CHSH axis, between the classical bound (2) and the Tsirelson bound (2√2 ≈ 2.828); the three raw blocks appear as faint points behind S_raw. Both values sit inside the quantum region, and the bracket between them is labelled with its physical meaning — readout error. The gap to the ceiling is mostly detector, not state.

9. Discussion

Six things came out of this campaign. It is worth stating them plainly before I return to the one that matters most:

1. **The compiler is a free adjacency oracle.** Because Tuna-17 maps qubits one-to-one and *rejects* a gate between unconnected qubits instead of rerouting it, the connectivity of the whole chip can be read off at compile time, for nothing. Two submissions were enough to confirm that the layout is the distance-3 rotated surface code.
2. **The periphery beats the centre.** The most-connected qubit on the chip made the *worst* pairs. Connectivity is not calibration quality; here the quiet corners beat the busy hub.
3. **Calibration drifts on a half-hour timescale.** Bond quality moved by about 0.03, and the fringe visibility slipped from 0.952 to 0.932, over tens of minutes — enough to matter, and the reason the experiment was run in tight back-to-back bursts.
4. **The nominal angles are already true.** The measured offset $\varphi_0 = -1.0° \pm 1.1°$ is smaller than the 5° grid the hardware can even represent. The calibration's verdict was the rare and satisfying "nothing to correct."
5. **The gap to Tsirelson is mostly in the detectors.** S_raw = 2.494 (88.2%) rose to S_mitigated = 2.736 (96.7%) once readout error was removed.
6. **The platform displays outcomes in reversed bit order (b[1]b[0]).** Harmless for parity-based quantities like S, but a genuine trap for anyone reading individual outcomes.

The redemption arc. Taken at face value, this chapter looks like a failure. I did everything more carefully than in Chapter 1 — searched for the best qubits, measured the true angles, controlled for drift — and the headline S came out *lower*, 88.2% against 91.4%. If the story stopped at the raw number, the lesson would be that care does not pay.

Phase D turns the story over. Correcting the *same* Phase-C data for the independently measured readout errors lifts S to 96.7% of the Tsirelson bound. The gap to the quantum ceiling was never mostly in the entangled state — what Tuna-17 prepared was already within about three percent of a perfect Bell pair. It was the detectors, reading that state out, that cost the missing points: roughly *two-thirds* of the entire distance to Tsirelson is readout error, and only the last third is gates, decoherence and whatever small residual sits in the a′ leg. Those last points cannot be bought back by simplifying the circuit — it is already minimal: a single entangling gate, with only q[0] and q[1] ever touched (the `qubit[17]` declaration is just the platform's requirement that the register match the chip, not fifteen idle operations). They are the chip's *intrinsic* cost — CZ and Ry gate error, decoherence, and spectator crosstalk from the fifteen idle neighbours, the last of which the corner-bond choice of Phase A already works to keep small. Closing that final third calls for a better two-qubit device, not a shorter program. The careful version did not fail; it *located the gap*, which is a more useful result than a slightly higher raw number would have been.

What is left undecided. The one loose thread is the ~2.2σ shortfall of the raw S below the fringe's own prediction, with $E_4 = E(a', b')$ the consistent weak leg. That points at the $a' = \pi/2$ *setting*, but I am careful not to over-read it: the data cannot separate a genuinely lossy $Ry(-\pi/2)$ rotation on Alice's qubit from ordinary drift between the fringe and the CHSH run. Deciding between them needs a dedicated sweep of Alice's rotation chain — a natural next measurement, and one that would change none of the conclusions above.

Honest caveats on the corrected number. *S*_mitigated is an *estimate of the state before readout*, not a better measurement of it, and it is reported strictly alongside *S*_raw, never in its place. Inverting the confusion matrix mixes the outcomes and so *widens* the error bar (0.039 → 0.047), and a fuller propagation that also accounts for the confusion matrix's own counting statistics widens it further still, to about 0.055 (Section 7.6): the corrected value is less certain, not more, despite looking cleaner. It also stays below 2√2 — the check that the correction is not manufacturing a super-quantum result out of noise. Had it crossed the Tsirelson bound, that would have been a sign of over-correction to report, not a discovery.

A recurring theme of the series. Three times now a real machine has taught me a convention I could not have guessed from the ideal theory: Chapter 2's global measurement sign flip, this chapter's happy $\varphi_0 \approx 0$, and now the reversed b[1]b[0] display order. None of them changed a headline number, because each notebook is built on parities and magnitudes that happen to be blind to exactly these conventions — but every one of them would have quietly corrupted a naive per-outcome analysis. Writing them down is much of the point of keeping a laboratory notebook. One expected quirk, tellingly, did *not* appear: these Ry-based measurements involve no S†·H chain anywhere, and no global sign flip showed up in the correlations — a clean null that pins Chapter 2's flip to that specific rotation chain rather than to the platform at large.

Scope and limits. This is a single-session campaign on one pair of qubits; it closes none of the loopholes of a rigorous Bell test. The readout correction rests on the assumption that the confusion matrix measured just after Phase C still described the detectors *during* it — the two ran minutes apart with no tuning event between, which is the best guarantee available here. The result should be read as a careful, fully-decomposed measurement of how close this processor comes to Tsirelson, not as a loophole-free Bell experiment.

---

10. Conclusion

I set out to see how close a freely accessible quantum processor could come to the Tsirelson bound, and to account for the gap that remained. On the corner bond (0, 1) of Tuna-17, the definitive CHSH run returned

$$S_{\text{raw}} = 2.494 \pm 0.039,$$

exceeding the classical limit by about 12.7 standard deviations — a clear violation — but reaching only 88.2% of the quantum maximum, and slightly less than Chapter 1. Measuring the readout errors separately and correcting for them raised the same data to

$$S_{\text{mitigated}} = 2.736 \pm 0.047,$$

or 96.7% of Tsirelson. The chapter's real result is not either number on its own but the distance between them: the entangled state this machine prepares is already close to ideal, and the largest single obstacle between a public quantum processor and the Tsirelson bound is not the state it makes but the detectors that read it out. Chasing Tsirelson, it turns out, is mostly chasing better measurement — and the three qubit-level surprises found along the way (the periphery beating the centre, the half-hour drift, the reversed bit order) are the kind of thing only a real device will tell you.

Appendix A — Source Code

All programs are written in cQASM 3.0 and were executed on the Tuna-17 backend at 512 shots, using the two qubits of the chosen bond. Every program declares `qubit[17]` because the hardware requires the register to match the chip, but only `q[0]` and `q[1]` (the winning bond) are ever operated on. The CZ-native Bell preparation block is identical throughout; only the final rotations change. The code is reproduced exactly as executed; inline comments have been translated from the original Italian shown in the Appendix B screenshots.

A.1 Adjacency oracle

A controlled-Z on a *connected* pair compiles and would run; on an *unconnected* pair it is rejected at compile time. This is the free topology probe of Section 4.1.

```
version 3.0

qubit[17] q

CZ q[0], q[1]      // (0,1) connected  -> compiles cleanly
// CZ q[0], q[8]   // (0,8) unconnected -> "qubit interactions prevent a 1-to-1 mapping
```

A.2 Phase A — pair-scan probes

Two probes per candidate bond (i, j). The **ZZ probe** measures the Bell state directly; the **XX probe** reads it out in the X basis — where the post-CZ `H·H` on `q[j]` cancels, so the only difference is which qubit keeps the trailing `H`.

```
version 3.0

qubit[17] q
bit[2] b

// CZ-native Bell preparation on bond (i, j)
H q[i]
H q[j]
CZ q[i], q[j]
H q[j]

// --- ZZ probe: last line above completes prep; measure directly ---
b[0] = measure q[i]
b[1] = measure q[j]
```

```
version 3.0

qubit[17] q
bit[2] b

H q[i]
H q[j]
CZ q[i], q[j]
H q[i]          // XX probe: X-basis readout (H·H on q[j] cancels)

b[0] = measure q[i]
b[1] = measure q[j]
```

Bonds scanned: (0,1), (2,5), (4,8), (7,11), (8,12), (15,16). The tie-break re-ran the top three: (0,1), (2,5), (15,16).

A.3 Phase B — calibration fringe

On the winning bond (0, 1), Alice is fixed at 0 and Bob's angle is swept via `Ry(-θ)` on `q[1]`.

```
version 3.0

qubit[17] q
bit[2] b

// Bell preparation on bond (0, 1)
H q[0]
H q[1]
CZ q[0], q[1]
H q[1]

// Bob's measurement angle
Ry(THETA_ARG) q[1]

b[0] = measure q[0]
b[1] = measure q[1]
```

with `THETA_ARG = -θ_B` for the nine angles (the θ_B = 0 angle omits the `Ry` line entirely):

Point	`θ_B`	`THETA_ARG`
B1	−π	`3.141592653589793`
B2	−3π/4	`2.356194490192345`
B3	−π/2	`1.5707963267948966`
B4	−π/4	`0.7853981633974483`
B5	0	*(Ry omitted)*
B6	+π/4	`-0.7853981633974483`
B7	+π/2	`-1.5707963267948966`
B8	+3π/4	`-2.356194490192345`
B9	+π	`-3.141592653589793`

A.4 Phase C — the four CHSH programs

The definitive runs, verbatim. Alice's $a = 0$ setting omits her `Ry`; her $a' = \pi/2$ setting applies `Ry(-π/2)`. Bob's $b = \pi/4$ uses `Ry(-π/4)`, his $b' = -\pi/4$ uses `Ry(+π/4)`.

```
version 3.0

qubit[17] q
bit[2] b

// C1 — E(a, b): a=0, b=+pi/4
H q[0]
H q[1]
CZ q[0], q[1]
H q[1]
Ry(-0.7853981633974483) q[1]
b[0] = measure q[0]
b[1] = measure q[1]
```

```
version 3.0

qubit[17] q
bit[2] b

// C2 — E(a, b'): a=0, b'=-pi/4
H q[0]
H q[1]
CZ q[0], q[1]
H q[1]
Ry(0.7853981633974483) q[1]
b[0] = measure q[0]
b[1] = measure q[1]
```

```
version 3.0

qubit[17] q
bit[2] b

// C3 — E(a', b): a'=pi/2, b=+pi/4
H q[0]
H q[1]
CZ q[0], q[1]
H q[1]
Ry(-1.5707963267948966) q[0]
Ry(-0.7853981633974483) q[1]
b[0] = measure q[0]
b[1] = measure q[1]
```

```
version 3.0

qubit[17] q
bit[2] b

// C4 — E(a', b'): a'=pi/2, b'=-pi/4
H q[0]
H q[1]
CZ q[0], q[1]
H q[1]
Ry(-1.5707963267948966) q[0]
Ry(0.7853981633974483) q[1]
b[0] = measure q[0]
b[1] = measure q[1]
```

A.5 Phase D — readout calibration

Each computational basis state is prepared with `X` gates and measured. Recall the platform displays the outcome as b[1]b[0], so `X q[1]` prepares the state shown as `10`.

```
version 3.0          version 3.0          version 3.0          version 3.0

qubit[17] q          qubit[17] q          qubit[17] q          qubit[17] q
bit[2] b             bit[2] b             bit[2] b             bit[2] b

// D1 — prep 00      // D2 — prep 10      // D3 — prep 01      // D4 — prep 11
                     X q[1]               X q[0]               X q[0]
                                                               X q[1]
b[0] = measure q[0]  b[0] = measure q[0]  b[0] = measure q[0]  b[0] = measure q[0]
b[1] = measure q[1]  b[1] = measure q[1]  b[1] = measure q[1]  b[1] = measure q[1]
```

Appendix B — Original Platform Screenshots

The original Quantum Inspire result screens are kept for provenance. Each shows the executed program, the circuit, the backend (Tuna-17), the shot count, the timestamp, and the measured outcome probabilities. The full run log follows; a representative set of screens is reproduced below it.

The first two screens are the **adjacency oracle of Finding 1** caught in the act, and they belong together. **M1** runs a controlled-Z on the *connected* pair (0, 1): the compiler accepts it, emits native gates with no inserted SWAPs, and it runs to a clean Bell histogram (00 and 11 dominant). **M2** is the identical gate on the *unconnected* pair (0, 8): the compiler rejects it outright — *"the following qubit interactions in the circuit prevent a 1-to-1 mapping: {(0, 8)}"* — so it never runs. This accept/reject duo is the entire basis for reading the chip's 24-bond connectivity off the compiler, for free.

Run	Job ID	Result ID	Timestamp (23/07/2026)
Fringe B1–B7	1171562–1171568	1136668–1136674	19:35:59 – 19:41:41
Fringe B8, B8′, B9	1171569, 1171571, 1171570	1136675, 1136677, 1136676	20:14:06 – 20:14:13
C block 1 (C1–C4)	1171572–1171575	1136678–1136681	20:27:46 – 20:31:37
C block 2 (C1–C4)	1171576–1171579	1136682–1136685	20:34:35 – 20:35:38
C block 3 (C1–C4)	1171581–1171584	1136687–1136690	20:43:54 – 20:46:00
D readout (D1–D4)	1171585–1171588	1136691–1136694	20:55:00 – 20:59:02

Table B1. Full provenance for the fringe, CHSH and readout runs. The Phase-A scan and tie-break (22 July, jobs 1170535 and 1171296–1171338) are logged in Section 6.

M1 — adjacency oracle, ACCEPT: CZ on the connected pair (0, 1) compiles and runs — Job 1170535, Result 1135679

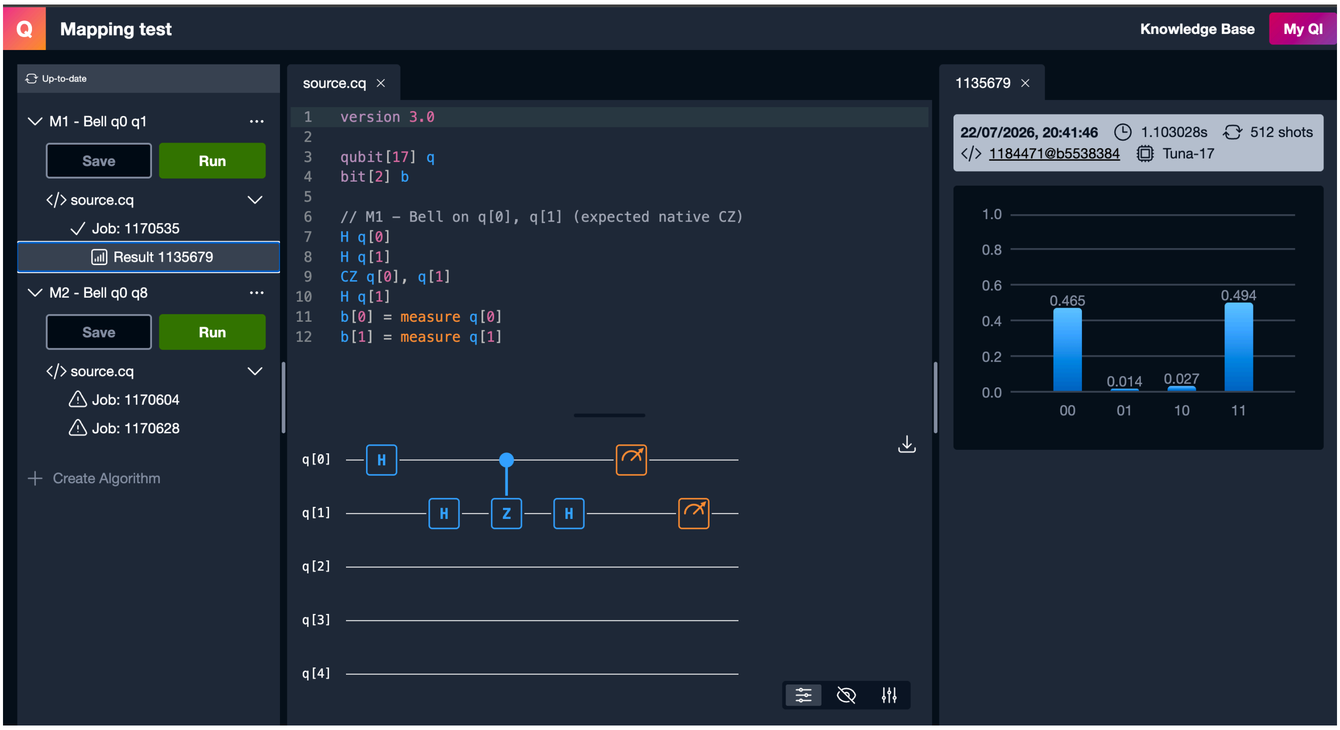

M2 — adjacency oracle, REJECT: CZ on the unconnected pair (0, 8) rejected at compile — Job 1170604

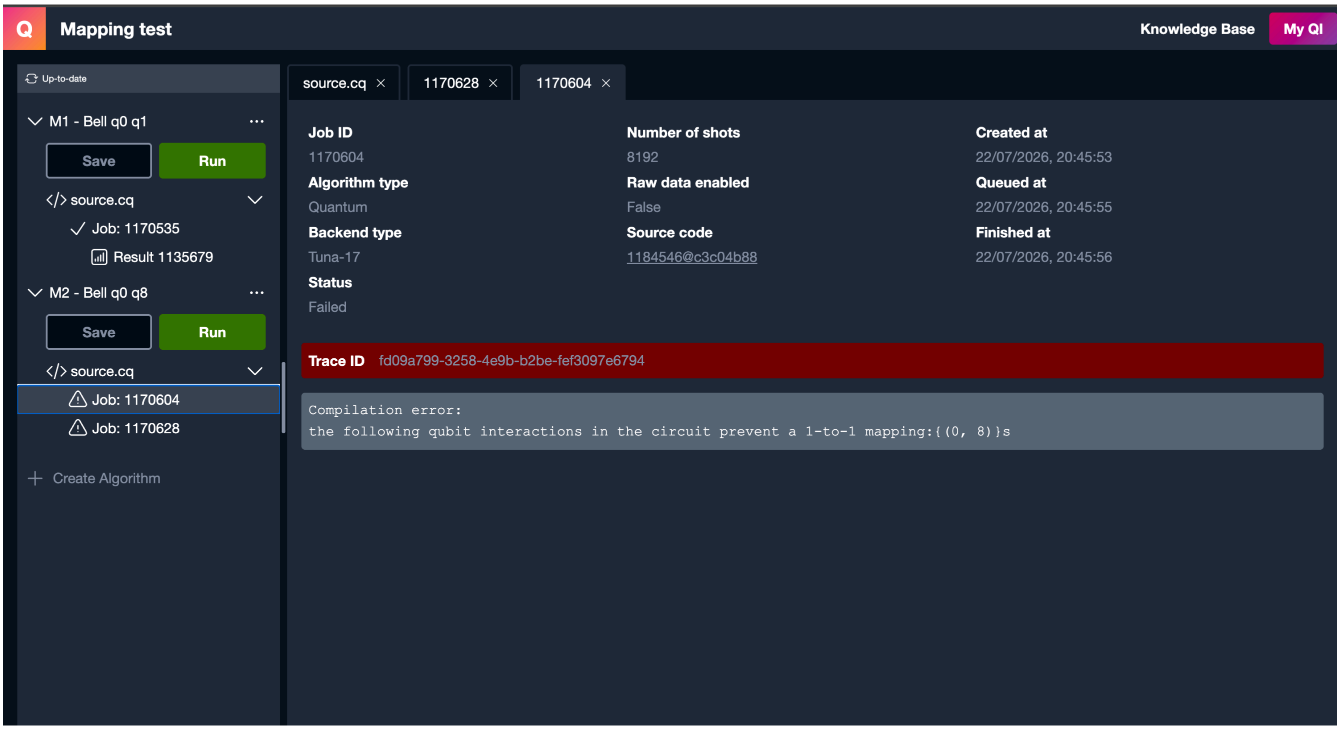

(The screen shows a requested shot count of 8192; the job failed at compile time, so no shots were ever executed and the number is immaterial — every run that produced data in this chapter used 512 shots.)

Fringe B5 (θ_B = 0) — Job 1171566, Result 1136672

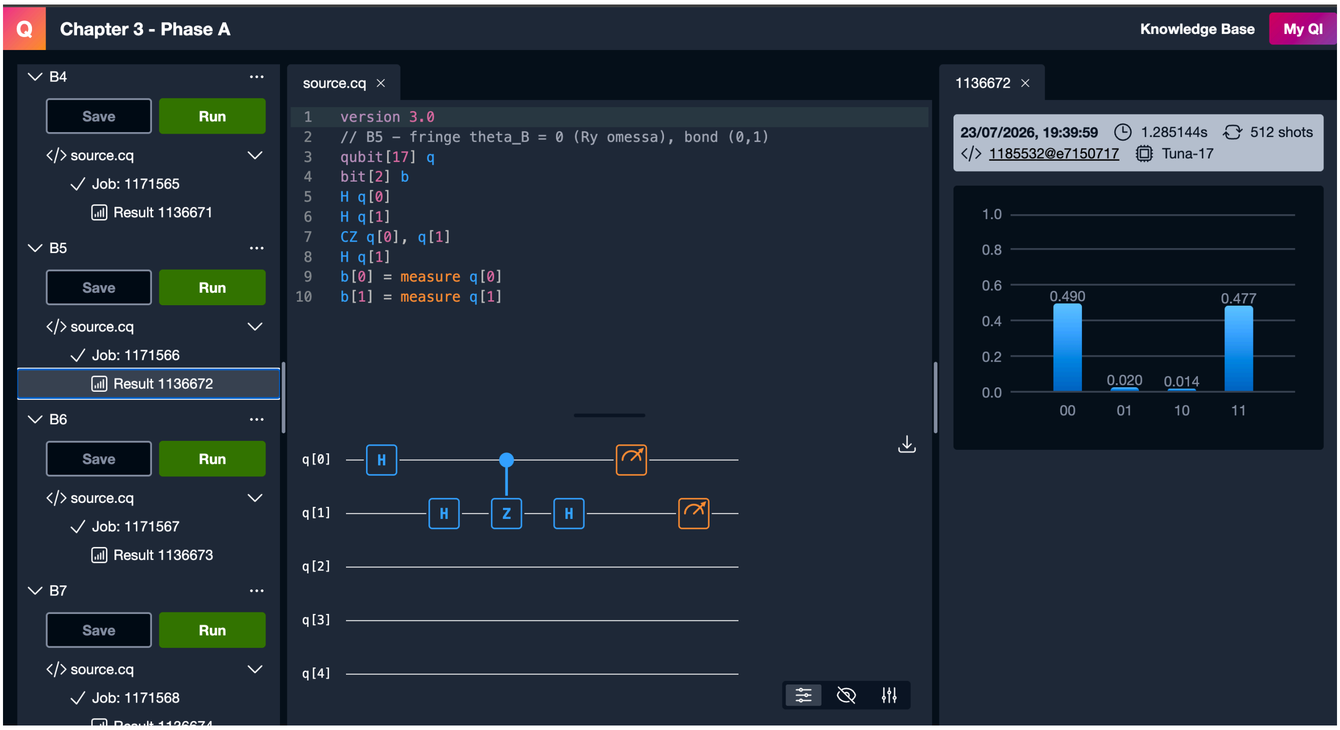

C1 — E(a, b), block 3 — Job 1171581, Result 1136687

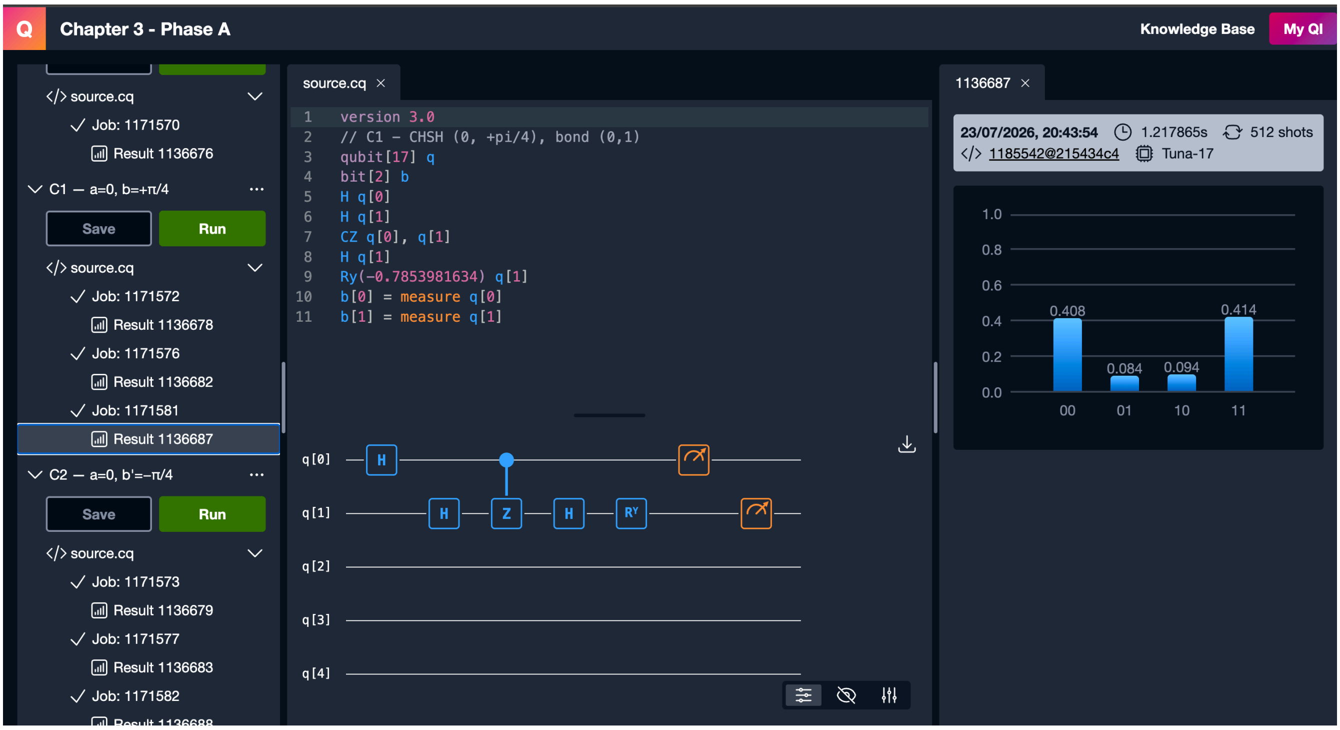

C4 — E(a′, b′), block 3, the negative leg — Job 1171584, Result 1136690

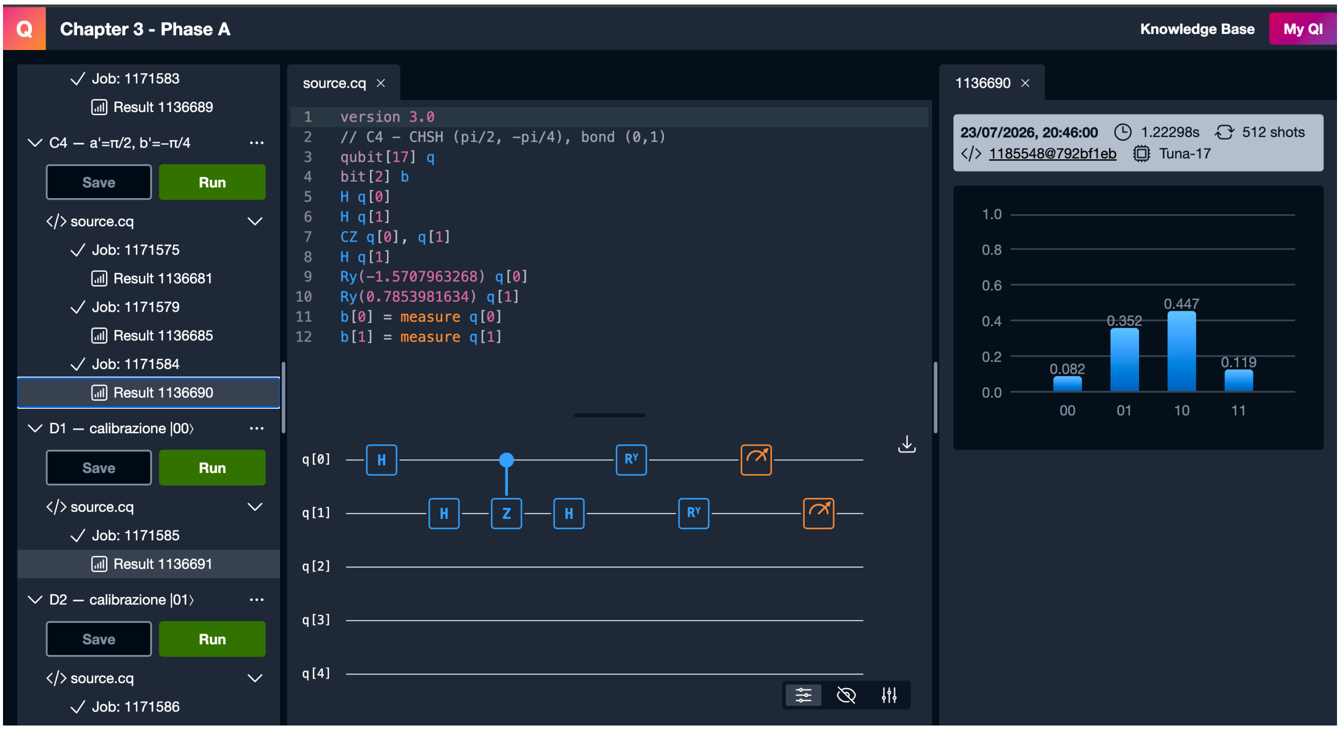

D1 — prep 00 — Job 1171585, Result 1136691

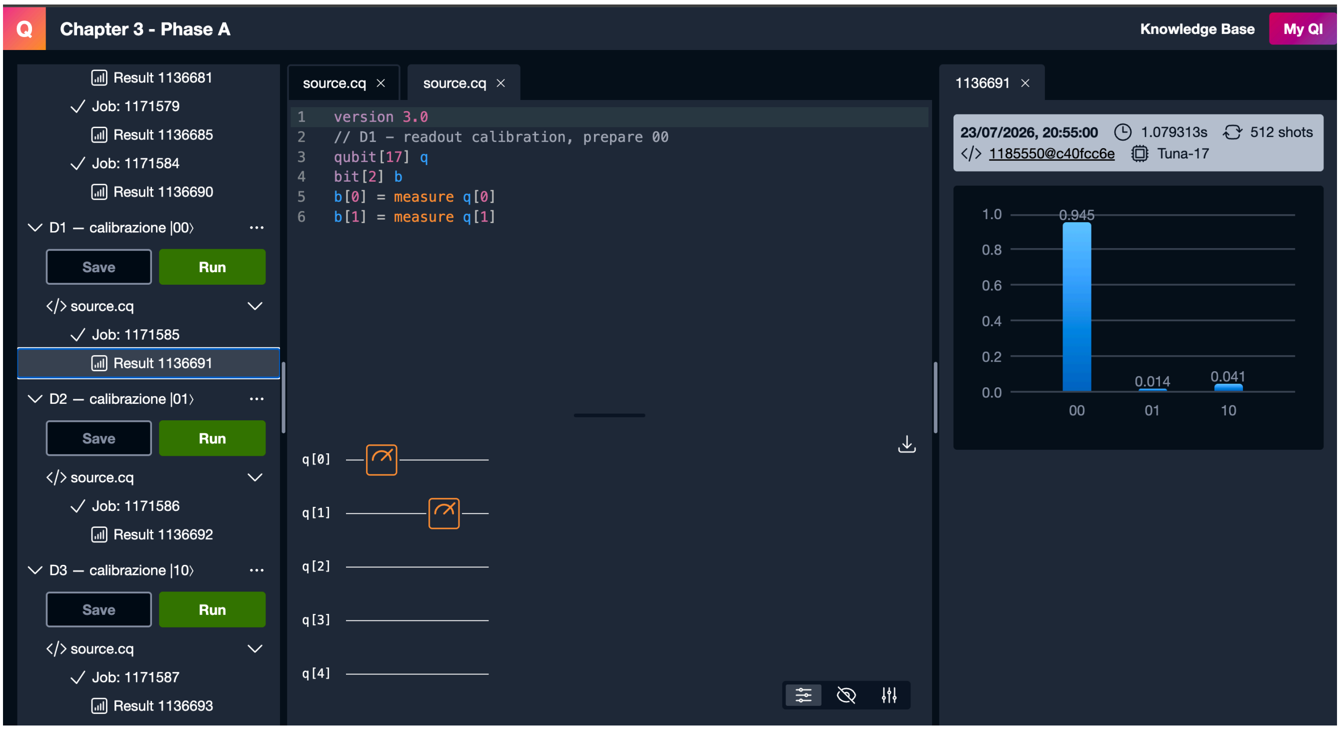

D4 — prep 11 — Job 1171588, Result 1136694

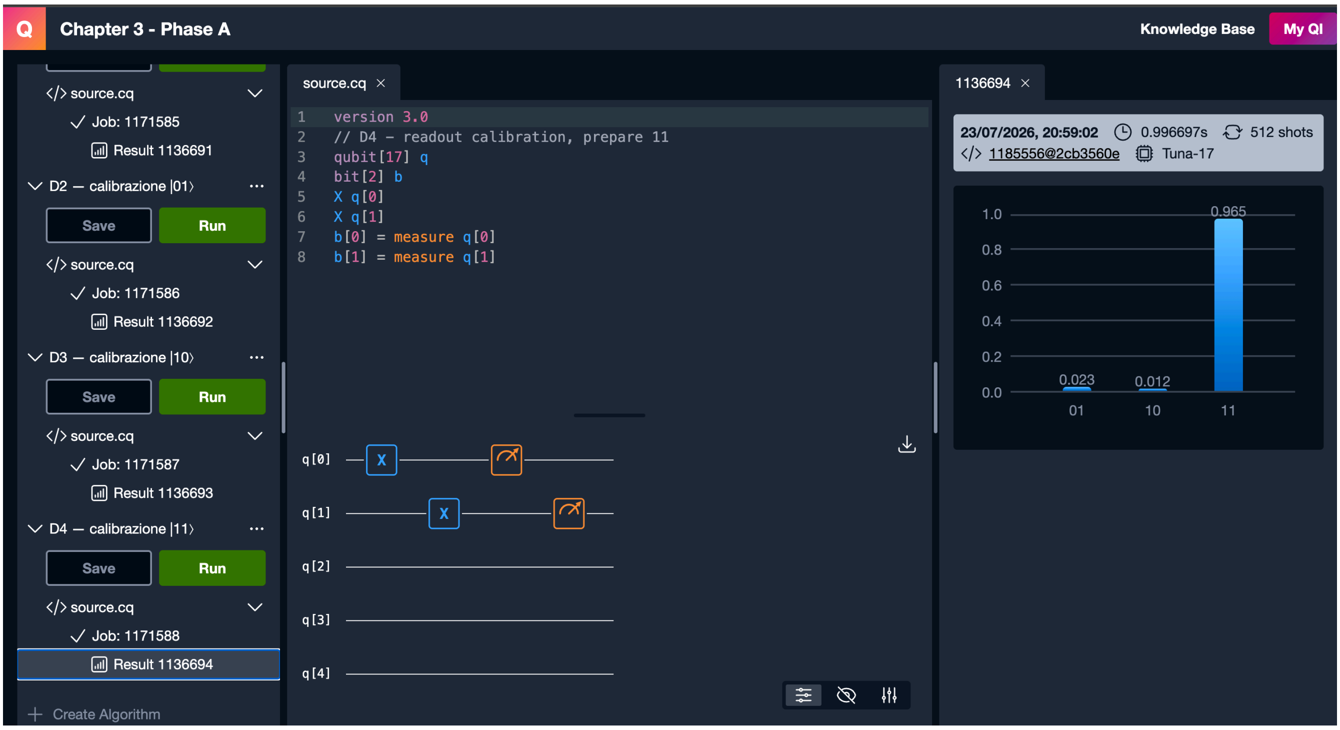

References

1. J. S. Bell, *On the Einstein Podolsky Rosen paradox*, Physics **1**, 195 (1964).
2. J. F. Clauser, M. A. Horne, A. Shimony, R. A. Holt, *Proposed experiment to test local hidden-variable theories*, Physical Review Letters **23**, 880 (1969).
3. B. S. Cirel'son (Tsirelson), *Quantum generalizations of Bell's inequality*, Letters in Mathematical Physics **4**, 93 (1980).
4. A. Pezzali, *Quantum Laboratory Notebook — Testing a Bell–CHSH Inequality on Real Quantum Hardware* (Chapter 1, companion notebook, 18 July 2026).
5. A. Pezzali, *Quantum Laboratory Notebook — A Three-Qubit GHZ / Mermin Test on Real Quantum Hardware* (Chapter 2, companion notebook, 19 July 2026).
6. Quantum Inspire — quantum computing platform documentation (Tuna-17 backend), QuTech.