CHAPTER 04 · THE MACHINES · SUPERCONDUCTING

Hardware deep dive — *superconducting qubits*.

Three machines, one physics stack — Willow, Heron r2, and Zuchongzhi 3.0. By the end of this chapter you'll know what makes a transmon tick, which three numbers actually decide whether your circuit runs, and how to talk to a real quantum backend through Qiskit Runtime.

WHAT YOU'LL LEARN

INSIDE A TRANSMON

A non-linear LC oscillator, cooled to *15 millikelvin*.

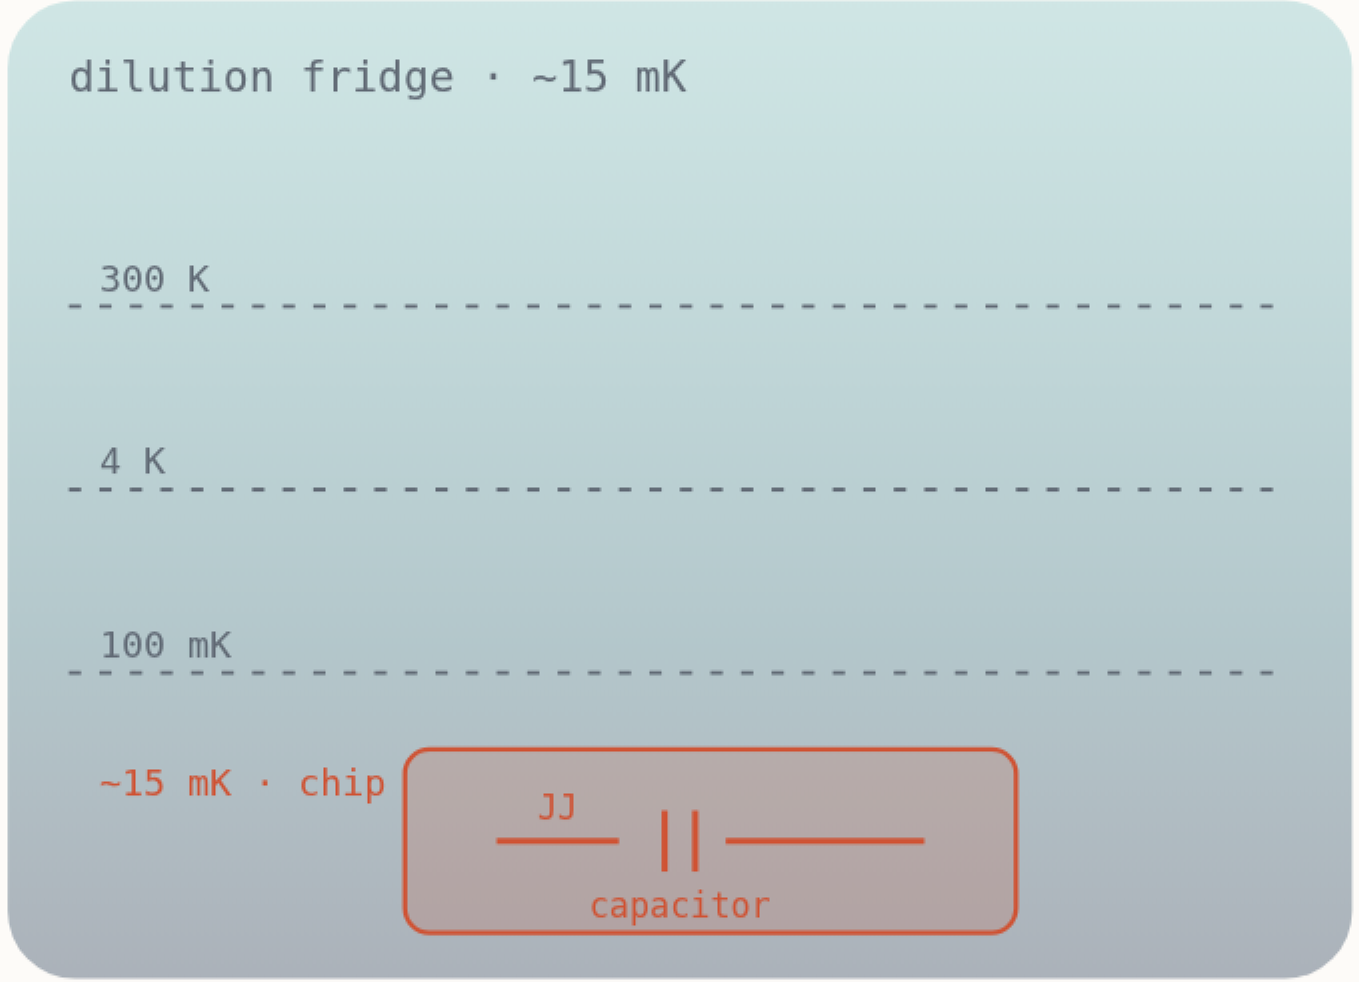

THREE INGREDIENTS

Josephson junction

Two superconductors split by a thin insulator — Cooper pairs tunnel coherently, the non-linearity that makes the qubit a qubit.

Shunt capacitor

Suppresses charge noise — the "transmon" insight (2007). Trades anharmonicity for stability.

Dilution fridge

15 millikelvin — colder than deep space — so thermal photons don't randomly flip the qubit.

Anharmonicity makes $|0\rangle \rightarrow |1\rangle$ a different energy than $|1\rangle \rightarrow |2\rangle$ — microwave pulses can address one transition without leaking to the next.

THE NUMBERS THAT MATTER

Qubit count is marketing. *Three numbers* decide your circuit.

● **T 1 · ENERGY RELAXATION**

~100 µs

How long the qubit holds |1⟩ before decaying to |0⟩. 2026 leaders cluster at 80–150 microseconds on the best transmons.

● **T 2 · DEPHASING**

~100 µs

How long the qubit keeps its phase. Often shorter than T_1. Bounds the depth of any phase-sensitive algorithm.

● **2Q GATE FIDELITY**

99.5–99.9%

Probability an entangling gate does what it claims. 99% → 99.9% is two orders of magnitude in achievable circuit depth.

Right question: *how deep can the circuit go before noise wins?* Not *how many qubits*.

— VENDOR SPOTLIGHT · DECEMBER 2024

Google Willow — the first *below-threshold* demonstration.

● **WHAT SHIPPED**

105 transmon qubits

Surface-code distance scaled 3 → 5 → 7. Logical error rate dropped *exponentially* as code distance grew — the threshold theorem confirmed on real silicon. Published in *Nature*, December 2024.

● **CAVEAT**

RCS is not a computation

The companion random-circuit-sampling benchmark — "5 minutes vs 10^{25} years classically" — is impressive physics but produces no useful output. The QEC result is the headline that matters.

- ● 105 qubits
- ● Surface code d=3,5,7
- ● Below threshold ✓
- ● Nature · Dec 2024

— VENDOR SPOTLIGHT · IBM ROADMAP

Heron r2 → Loon → Kookaburra — betting on *connectivity*, not count.

● 2024 · HERON R2

156 qubits

Mid-circuit measurement, dynamic circuits, tunable couplers. Production silicon today.

● 2025 · LOON

c-couplers

Cross-chip long-range connectivity — the prerequisite for leaving surface codes.

● 2026 · KOOKABURRA

qLDPC memory

First module targeting Gross code `[[144,12,12]]` — 12 logical from 144 physical.

● 2028–2029 · STARLING

~200 logical

10^8 gate operations on logical qubits — IBM's stated first useful FTQC machine.

Surface code costs ~1,452 physical qubits per logical at distance 25. Gross code: 144 physical → 12 logical. *An order of magnitude in overhead* — if the engineering lands.

VENDOR SPOTLIGHT · USTC · MARCH 2025

Zuchongzhi 3.0 — highest reported *2Q fidelity* in the group.

● **WHAT SHIPPED**

105 qubits · 99.62% 2Q fidelity

Pan Jianwei's USTC lab. Published in *Physical Review Letters*, March 2025. Zuchongzhi 3.2 (Dec 2025) added below-threshold on distance-7 surface code.

● **HONEST FRAMING**

10^{15}× RCS — a benchmark

Like Willow's: random-circuit sampling is not a useful computation. The 99.62% gate fidelity is the result that moves the field — that is the number that buys circuit depth.

● 105 qubits

● 99.62% 2Q fidelity

● PRL · Mar 2025

● Below-threshold (3.2)

THE 2026 SUPERCONDUCTING SCORECARD

Same physics. *Three bets*.

● GOOGLE WILLOW

105 qubits

2Q fidelity ≈ 99.7%
T_1 ≈ 100 µs
Connectivity grid (nearest-neighbor)
Headline below-threshold surface code, Nature Dec 2024

● IBM HERON R2

156 qubits

2Q fidelity ≈ 99.7%
T_1 ≈ 250 µs
Connectivity heavy-hex + tunable couplers
Headline mid-circuit measurement; qLDPC path via Loon → Kookaburra

● ZUCHONGZHI 3.0

105 qubits

2Q fidelity 99.62%
T_1 ≈ 70 µs
Connectivity grid (nearest-neighbor)
Headline highest reported 2Q fidelity; below-threshold on 3.2 (Dec 2025)

All three need millikelvin fridges. All three are transmons. The *bets* are: which code geometry, which connectivity, and which manufacturing process scales first.

FRAMEWORK ANCHOR · QISKIT 1.X RUNTIME

Two primitives — *Sampler* and *Estimator*.

● REAL BACKEND · SAMPLER RETURNS COUNTS · ESTIMATOR RETURNS EXPECTATION VALUES

```
from qiskit import QuantumCircuit
from qiskit.quantum_info import SparsePauliOp
from qiskit_ibm_runtime import QiskitRuntimeService, SamplerV2, EstimatorV2

service = QiskitRuntimeService()
backend = service.backend("ibm_brisbane")   # real 127-qubit Eagle

qc = QuantumCircuit(2)
qc.h(0); qc.cx(0, 1); qc.measure_all()      # Bell state

# Sampler: counts -> {'00': ~512, '11': ~512}
sampler = SamplerV2(backend)
counts = sampler.run([qc], shots=1024).result()[0].data.meas.get_counts()

# Estimator:  ~ +1 on ideal Bell state
obs = SparsePauliOp.from_list([("ZZ", 1.0)])
ev  = EstimatorV2(backend).run([(qc.remove_final_measurements(inplace=False), obs
```

● WHY PRIMITIVES

No manual job loop

You describe *what* you want (samples or expectation values), not *how* to fetch a job, poll, retry, parse. Runtime handles queueing, error mitigation knobs, and session resumption.

Same API, sim or QPU

Swap `backend` for an Aer simulator and the rest of the code is unchanged — develop local, run remote.

— TRY IT YOURSELF

Run a 4-qubit GHZ on *ibm_brisbane* — and meet your noise.

● THE EXERCISE

~10 minutes · free IBM Quantum tier

1. `pip install qiskit qiskit-ibm-runtime`
2. Sign up at `quantum.ibm.com`, save your token.
3. Build a 4-qubit GHZ: H on q0, then CX(0→1), CX(1→2), CX(2→3), measure all.
4. Run with `SamplerV2(backend="ibm_brisbane")`, 4096 shots.
5. Run the same circuit on `AerSimulator()`, 4096 shots — ideal.
6. Compute Hellinger fidelity between the two histograms.

● WHAT YOU SHOULD SEE

Ideal: $|0000\rangle + |1111\rangle \approx 50/50$

Hardware histogram will show those two peaks *plus* a long tail of bit-flip and dephasing errors. Hellinger fidelity typically lands 0.75–0.90 on a 4-qubit GHZ today.

If $F < 0.8$

You are watching the T_1, T_2, and 99.7% gate fidelity from slide 3 actually *do* something. The numbers are not abstract — they're the histogram tail.

Next chapter — trapped ions and neutral atoms. Different physics, all-to-all connectivity, fidelities north of 99.9%.