

Kicked-Ising Online Quantum Reservoirs: Operating Regimes and Hardware Tests

Ege Yilmaz^{1,2}, Ibrahim Shehzad², Francesco Tacchino²

¹*Department of Mathematics and Computer Science University of Basel (Switzerland)*

²*IBM Research – Zurich, 8803 Rüschlikon, Switzerland*

Quantum reservoir computing processes temporal data by driving a fixed quantum system and training only a classical readout. Recent large-scale quantum extreme learning machine experiments with kicked-Ising circuits show that these digital quantum feature maps can be useful when operated in regimes balancing task performance and observable variability with diagnostics for concentration and dynamical complexity [1]. However, these QELMs are feedforward feature maps, whereas online QRC requires a recurrent quantum state with fading memory of the input stream.

We convert the hardware-tested kicked-Ising family used in dual-unitary chaos simulations and large-scale QELMs [1], [2] into an online QRC architecture with repeated readout and reset. We focus on *observable memory*, namely the requirement that temporal information persists in hidden qubits, remains accessible through a sufficiently rich observable space, and survives finite-shot sampling and hardware distortions. We use Quantum Volterra Theory (QVT) [3] to diagnose the hidden memory channel, and observable-subset sweeps to determine how much of that memory is accessible through measured hardware features.

The register contains $N = M + R$ qubits, with M hidden memory qubits and R visible readout qubits measured and reset at each time step, following the memory-channel construction of [3]. The scalar input u_t , scaled by κ , modulates the kicked-Ising evolution, where (J, h, b) denote the Ising coupling, longitudinal field, and transverse kick. Measuring and resetting the visible subsystem while carrying the hidden subsystem forward induces the input-dependent memory channel $\rho_{t+1}^M = \mathcal{P}_{u_t}(\rho_t^M)$. The measured Pauli expectation values define the feature vector ϕ_t , used by a trained linear readout.

Following the operating-regime philosophy of [1], we select parameters by comparing exact, noiseless diagnostics on the same reservoirs. Memory capacity (MC) quantifies linear reconstruction of past inputs; observable variability ν , the average temporal standard deviation of measured Pauli expectation values, measures resolvability. To avoid high-memory but dynamically trivial reservoirs, we also track the second Rényi entropy S_2 as a scrambling proxy and tensor-network bond-dimension diagnostics as a compressibility test.

Figure 2 shows a representative (J, b) regime map for the $(N, R) = (8, 4)$ reservoir. The map is averaged over scanned h values, and the MC, observable-variability ν , and S_2 channels are normalized independently. The corresponding $(6, 4)$ and $(10, 4)$ maps, not shown, exhibit the same qualitative tradeoff geometry as the displayed $(8, 4)$ map, while ν remains in a comparable useful range and the raw MC and S_2 scales increase along the checked $R = 4$ sequence. Variability is largest toward one side of the grid, MC forms a broad intermediate band, and S_2 increases toward the dual-unitary corner of the (J, b) plane, where the dynamics become more strongly entangling. For this h -averaged map, the grid maxima are $\text{MC} = 10.545$, $\nu = 0.0874$, and $S_2 = 3.692$, while the selected point has $\text{MC} = 8.01$, $\nu = 0.0404$, and $S_2 = 1.65$. We choose $J = \pi/8$ and $b = 3\pi/16$ as a representative compromise point in this useful operating band, where the dynamics are rich enough to encode history while remaining locally observable.

At this (J, b) point, $h = 11\pi/32$ is selected from the full scan tables because it gives high MC and variability across inspected sizes and lies inside a stable high-memory lobe rather than an isolated optimum. With (J, h, b) fixed, κ is the remaining tuning axis. Figure 2 shows that small κ preserves memory but weakens variability, while larger κ improves visibility and eventually degrades memory. We use $\kappa = 0.33$ because it keeps MC, variability, and system-size trends simultaneously acceptable for hardware.

After fixing the operating point, we use QVT to diagnose the induced hidden-memory channel $\mathcal{P}_u$ [3]. For the null-input channel $\mathcal{P}_0$, we compute the fixed point, subleading eigenvalue magnitude




Figure 1: Online QRC architecture. The kicked-Ising family is made recurrent by measuring and resetting a visible register while carrying a hidden register forward. Input κu_t modulates the dynamics, inducing $\rho_{t+1}^M = \mathcal{P}_{u_t}(\rho_t^M)$.




Figure 2: Regime selection. Left, representative noiseless $(N, R) = (8, 4)$ RGB map. MC, observable variability ν , and S_2 are independently normalized over the grid and mapped to the red, green, and blue channels shown in the triangular legend. The black cross marks the selected point $J = \pi/8$, $b = 3\pi/16$. Right, noiseless input-scale sweep at the selected Hamiltonian point, showing the memory–variability tradeoff.

$|\lambda_2|$, and memory time $n_M = -1/\log |\lambda_2|$. Two trace-distance witnesses quantify non-unitality and input response. The first tests whether $\mathcal{P}_0$ moves the maximally mixed state, excluding the unital case associated with trivial Volterra kernels; the second tests whether an encoded input perturbs the memory fixed point.

At the selected parameter point, $(N, R) = (6, 4), (8, 4), (10, 4)$ all have $|\lambda_2| < 1$, nonzero non-unitality, and nonzero input response. Figure 3 separates intrinsic and accessible memory. The QVT memory time grows mainly with hidden-memory size M over the checked range. Recovered MC instead depends on how much of the visible operator space is measured: at fixed $R = 4$, larger hidden registers reach higher MC ceilings, but this advantage appears only after enough Pauli observables are included. Since R sets the available visible Pauli space, scaling M alone is insufficient.

We implement the selected $(N, R) = (6, 4)$ reservoir on IBM superconducting hardware with repeated mid-circuit measurements and resets. The visible register is measured in X , Y , and Z Pauli bases with 10^5 shots per basis, then reinitialized before the next input step; the resulting expectation values are hardware features while the hidden register carries the recurrent state.

Figure 4 summarizes the hardware-accessible temporal signal. The left panel compares delay-wise MC contributions MC_k with nonlinear autoregressive moving-average (NARMA)- $(k + 1)$ prediction scores for the same feature families. The right panel compares total MC from noiseless exact features, finite-shot circuit simulations, and hardware expectation values. The gap between the 10^5 -shot finite-shot simulation and the hardware estimate indicates device-level distortions beyond sampling noise. Thus the selected regime has short-delay accessible memory, while longer-horizon recovered memory remains precision-limited.




Figure 3: Intrinsic and accessible memory. Left, QVT memory time n_M versus hidden memory size M , showing roughly linear growth over the inspected range. Right, noiseless MC computed from exact expectation values versus the number of Pauli expectation values included in the readout at fixed $R = 4$; larger hidden registers reach higher ceilings but require larger observable sets before saturating.




Figure 4: Hardware memory and temporal-task performance. Left, bars show delay-wise memory-capacity contributions MC_k , while dashed curves show NARMA- $(k+1)$ test R^2 , for exact, finite-shot circuit simulation, and hardware features. Right, total MC versus shot budget for finite-shot circuit simulations, the noiseless exact reference, and the 10^5 -shot hardware run.

We introduce an online QRC architecture by converting the kicked-Ising circuit family used in large-scale QELM and dual-unitary chaos-simulation experiments into a reservoir with repeated readout and reset. Exact simulations identify a compromise regime balancing MC, observable variability, and many-body dynamics. QVT diagnostics indicate a fading-memory channel with memory time increasing with hidden-register size over the checked range, while observable-subset results show that larger hidden memories require sufficiently large measured observable sets to become accessible. Hardware experiments on the selected $(N, R) = (6, 4)$ reservoir show short-delay memory surviving mid-circuit measurement and reset, with longer horizons limited by finite-shot precision and device-level distortions. Overall, these results suggest that scaling online QRC is not simply a matter of increasing system size or scrambling strength. It also requires engineering observable memory, where a fading hidden-state memory remains recoverable from finite-shot measurements.

References

- [1] T. Dao et al., “Breaking concentration barriers for quantum extreme learning on digital quantum processors,” *arXiv preprint arXiv:2603.13005*, 2026. arXiv: 2603.13005.
- [2] L. E. Fischer et al., “Dynamical simulations of many-body quantum chaos on a quantum computer,” *Nature Physics*, vol. 22, pp. 302–307, 2026.
- [3] F. Hu, S. A. Khan, N. T. Bronn, et al., “Overcoming the coherence time barrier in quantum machine learning on temporal data,” *Nature Communications*, vol. 15, p. 7491, 2024.