Head-to-head comparison of analog quantum simulations

WHITEPAPER

Summary

Analog quantum simulation is progressing rapidly, with many quantum-computing manufacturers demonstrating results that highlight the strengths of their quantum processing units (QPUs). Some of these demonstrations have significant overlap, providing a unique opportunity for apples-to-apples comparisons between different QPUs in different modalities. Here we compare several recent simulations of quantum dynamics in the XY Hamiltonian and transverse-field Ising model.

Quantum computers (QCs) of increasing capability are being built in diverse modalities; the most prominent are superconducting qubits (flux, transmon, fluxonium, etc.), neutral atom arrays, and trapped ions. These differ significantly in both qubit coherence times and energy scales, which together determine how long a simulation can run. These QPUs all have different strengths and weaknesses, and every QC manufacturer presents their most impressive work—so how can we compare different QPUs on equal footing?

Three essential characteristics of a quantum simulation are the available connectivity, the family of realizable Hamiltonians, and the normalized coherence time. Here we consider recent publications of leading QC manufacturers and compare these characteristics where possible.

Qubit connectivity

Superconducting chips have a prefabricated geometry. Google's is square [1], IBM's is heavy-hex (a sparse subgraph of a square lattice) [2], and D-Wave Quantum Inc.'s is the much denser Zephyr™ graph [3], which enables flexible embedding of higher-dimensional and complex interaction geometries [4]. Neutral atom arrays are laid out in 2D [5, 6]; while 3D is possible in principle, no practical demonstration has been achieved. Furthermore, the two-body coupling terms are based on a van der Waals interaction depending on the distance between two atoms—their sign and magnitude cannot be tuned beyond that. Trapped ion companies (Quantinuum [7], IonQ) boast all-to-all interaction, but running multiple two-body interactions simultaneously introduces interference, control complexity, and reduced energy scale, which is why all-to-all connectivity is rarely used in practice.

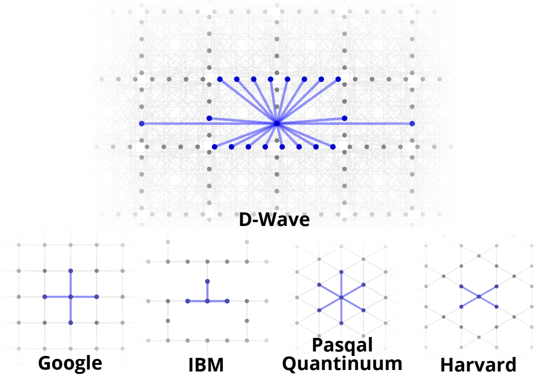

Figure 1: Local qubit connectivity for QPU experiments: Zephyr (D-Wave superconducting annealing [3]); square (Google superconducting gate [1]); heavy-hex (IBM superconducting gate [2]); triangular (Pasqal neutral atom [6] and Quantinuum trapped ion [7]); kagome (Harvard neutral atom [5]).

Realizable Hamiltonians

D-Wave's QPUs natively realize the transverse-field Ising model (TFIM). Additionally, recent analog-digital demonstrations realize an effective isotropic XY Hamiltonian, also known as the XX Hamiltonian, via a reduction called the "rotating wave approximation" in the weak coupling limit. Taking the simplest possible form (no detunings) to ensure cross-platform comparability, we have:

$$H = \frac{\mathcal{J}}{4}\sum_{\langle i,j\rangle}(\sigma_i^x\sigma_j^x + \alpha\sigma_i^y\sigma_j^y) \tag{1}$$

with $\alpha = 1$, which was also recently simulated on two other types of QPU [2, 5]. Related models include the general XY Hamiltonian, which allows $\alpha \neq 1$, and the full Heisenberg model, which also has $\sigma_i^z\sigma_j^z$ terms. The sum is over $\langle i,j\rangle$, meaning all coupled pairs in the geometry. Gate-model systems are in principle capable of simulating any of these models in discrete time via basis rotation [2], but this comes at a cost of effective depth (see Fig. 2)—superconducting gate-based systems can also use an analog XY Hamiltonian natively, allowing deeper simulations [1]. The native Hamiltonian of neutral-atom systems can be reduced to a TFIM [6] or XX Hamiltonian [5], similar to D-Wave's capabilities.

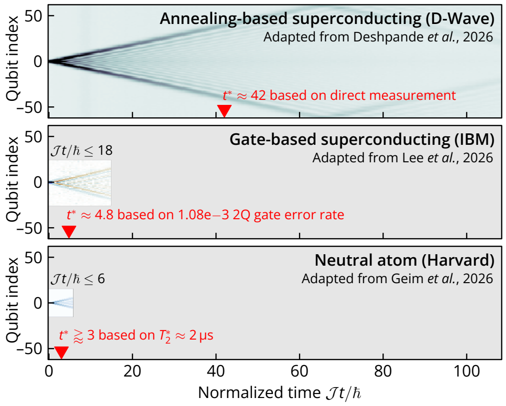

Figure 2: Comparing effective depth: ballistic transport of an excitation in an isotropic XY chain. Data have equal scale for space and normalized time. Top: D-Wave Advantage2™ QPU ([3] Fig. 4c) in a length $L = 124$ chain with periodic boundaries. Middle: IBM Boston ([2] Fig. 2B), $L = 50$ with open boundaries. Bottom: QuEra-style neutral atom array ([5] Fig. 2e), $L = 31$ with open boundaries. Estimates of t^* are shown; for IBM we consider gate error with 115 gates per $\Delta t = 0.6$ Trotter step, and for Harvard we use $2\,\mu s \approx T_2^* \leq T_2$ since no value of T_2 is available.

Normalized coherence time and effective system size

For Hamiltonian (1) in the clean 1D chain (nearest-neighbor couplings and no disorder), a single excitation propagates ballistically—linearly in time—with a constant velocity of $\mathcal{J}/\hbar$. The normalized time of a simulation is $\mathcal{J}t/\hbar$, where t is the physical time. It tells us very simply how far an excitation can propagate in the experiment.

IBM [2] and Harvard [5] experiments demonstrating the same phenomenon are compared in Fig. 2: D-Wave's simulations are roughly $5\times$ and $20\times$ deeper, respectively, highlighting a substantial advantage in achievable dynamical depth. Google [1] demonstrated 2D simulations on 70 qubits out to $\mathcal{J}/\hbar = 110$, limited in their approach by exponential postselection overhead. In the transverse-field Ising model, Pasqal presented plausibly coherent "short-time" data out to $\mathcal{J}t/\hbar \approx 6$, and Quantinuum demonstrated discretized dynamics to $\mathcal{J}t/\hbar = 5$ [7].

We define the normalized coherence time of a simulation as

$$t^* = \mathcal{J}T_2/\hbar, \qquad (2)$$

where T_2 is the effective combined dephasing time. For IBM's discrete-time (Trotterized) dynamics we estimate T_2 via two-qubit gate error, taking into account the number of gates required per time step. For the Harvard experiments T_2 is not given, but we can use the lower bound $t^* \geq JT_2^*/\hbar$ for the related dephasing time T_2^*.

As for size, a high qubit count is only meaningful if the qubits can exchange information within the experimental time scale.

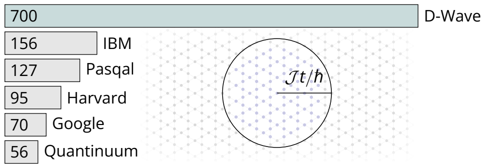

Figure 3: Comparing effective size: the number of qubits reachable from a central qubit in $\mathcal{J}t/\hbar$ lattice steps (Harvard example shown). D-Wave Advantage2 QPU, Google Willow, IBM Boston, and Quantinuum H2 reach all qubits within the normalized time. D-Wave's effective system size assumes 400 qubits reserved for state preparation and measurement, with the remaining $>$700 in a $\sim$1200-qubit research system used for the quantum simulation.

For example, Ref. [5] shows an array of 271 neutral atoms, but only 95 can be reached from a central qubit in six lattice steps ($\mathcal{J}t/\hbar \leq 6$). To track the effective system size, we compare the number of computational qubits reachable from a central site within $\mathcal{J}t/\hbar$ lattice steps, as shown in Fig. 3. Again, D-Wave shows a clear advantage over the other simulations in effective system size.

Taken together, these metrics suggest that D-Wave's annealing-based platform currently offers a leading approach for large-scale, deep analog quantum simulations of many-body dynamics.

References

[1] A. Lunkin, N. S. Ticea, S. Kumar, C. Miao, J. Choi, et al., *Evidence for a two-dimensional quantum glass state at high temperatures*, Jan. 2026, `arXiv:2601.01309 [quant-ph]`.

[2] Y.-T. Lee, K. Kumaran, B. Pokharel, A. Scheie, C. L. Sarkis, D. A. Tennant, T. Humble, A. Schleife, A. Kandala, and A. Banerjee, *Benchmarking quantum simulation with neutron-scattering experiments*, Mar. 2026, `arXiv:2603.15608 [quant-ph]`.

[3] R. Deshpande, M. Kheirkhah, C. Rich, R. Harris, J. Raymond, et al., *Analog-Digital Quantum Computing with Quantum Annealing Processors*, Mar. 2026, `arXiv:2603.15534 [quant-ph]`.

[4] A. D. King, A. Nocera, M. M. Rams, J. Dziarmaga, R. Wiersema, et al., "Beyond-classical computation in quantum simulation," Science **388**, 199–204 (2025).

[5] A. A. Geim, N. U. Koyluoglu, S. J. Evered, R. Sahay, S. H. Li, et al., *Engineering quantum criticality and dynamics on an analog-digital simulator*, Feb. 2026, `arXiv:2602.18555 [quant-ph]`.

[6] L. Leclerc, S. Julià-Farré, G. S. Freitas, G. Villaret, B. Albrecht, et al., *One-to-one quantum simulation of the low-dimensional frustrated quantum magnet TmMgGaO$_4$ with 256 qubits*, Mar. 2026, `arXiv:2603.20372 [quant-ph]`.

[7] R. Haghshenas, E. Chertkov, M. Mills, W. Kadow, S.-H. Lin, et al., "Digital quantum magnetism on a trapped-ion quantum computer," Nature, `10.1038/s41586-026-10445-3` (2026).