A Extended discussion of DMFT

Using DMFT, the correlated behavior of the lattice model that describes a real material is related to an impurity model. In our study, we performed DMFT for the half-filled Hubbard model in the Bethe lattice (see the "DMFT convergence with SGS" section) using an approximate impurity ground state represented by a sum of Gaussian states. The Hamiltonian that describes the N-orbital Hubbard model is given by Eq. S1, where h is the hopping strength, μ is the chemical potential, and U is the Coulomb interaction strength.

$$\widehat{\mathbf{H}}_{\text{Hubbard}} = \widehat{\mathbf{H}}_0 + \widehat{\mathbf{H}}_{\text{int}} \tag{S1}$$

$$\widehat{\mathbf{H}}_0 = -\sum_{\langle i,j\rangle=1,\sigma}^{N} h(\hat{c}^\dagger_{i\sigma}\hat{c}_{j\sigma} + \text{h.c.}) - \sum_{i=1,\sigma}^{N} \mu\hat{c}^\dagger_{i\sigma}\hat{c}_{i\sigma} \tag{S2}$$

$$\widehat{\mathbf{H}}_{\text{int}} = \sum_{i}^{N} U\hat{c}^\dagger_{i\uparrow}\hat{c}_{i\uparrow}\hat{c}^\dagger_{i\downarrow}\hat{c}_{i\downarrow} \tag{S3}$$

Here, the brackets in $\langle i,j\rangle$ indicate that the hopping terms exist only between nearest neighbors. A notable distinction between the Hubbard model and the impurity model lies in the $\widehat{\mathbf{H}}_{\text{int}}$ term, which sums over all orbitals instead of being restricted to the impurity orbitals. While self-consistent embeddings exist for impurity clusters (e.g. the so-called cellular DMFT) and our methods do allow for implementation with more than one impurity orbital, we limit the scope of this work to the single impurity Anderson model. Consequently, all further equations presented will consider $N_I = 1$.

A key quantity required to enforce DMFT self-consistency is the impurity GF, G^R_{imp} in its time dependent (Eq. S4), real (Eq. S5), and Matsubara (imaginary) (Eq. S6) frequency-dependent forms ($\omega_n = \frac{(2n+1)\pi}{\beta}$ for fermions and β is inverse temperature).

$$G^R_{\text{imp}}(t) = -i\langle\Psi^{(0)}|\{\hat{d}_\sigma(t), \hat{d}^\dagger_\sigma\}|\Psi^{(0)}\rangle \tag{S4}$$

$$\begin{aligned} G^R_{\text{imp}}(\omega) = &\langle\Psi^{(0)}|\hat{d}_\sigma[(\omega + E^{(0)} + i\eta)\mathbf{I} - \widehat{\mathbf{H}}_{\text{imp}}]^{-1}\hat{d}^\dagger_\sigma|\Psi^{(0)}\rangle \\ &+ \langle\Psi^{(0)}|\hat{d}^\dagger_\sigma[(\omega - E^{(0)} + i\eta)\mathbf{I} + \widehat{\mathbf{H}}_{\text{imp}}]^{-1}\hat{d}_\sigma|\Psi^{(0)}\rangle \end{aligned} \tag{S5}$$

$$\mathcal{G}_{\text{imp}}(i\omega_n) = \int_{-\infty}^{\infty} \frac{d\omega}{\pi} \frac{-\text{Im}[G^R_{\text{imp}}(\omega)]}{i\omega_n - \omega} \tag{S6}$$

In Eq. S5, a small, positive constant η is added to ensure convergence of the Fourier transform from the time dependent form in Eq. S4.

Each of these forms is equivalent in terms of their dynamical information. However, in practice, the Matsubara GF tends to be smooth, while the time dependent

and frequency-dependent impurity GFs exhibit long-lived oscillations and sharp peaks, respectively. To ensure stable convergence to the self-consistent bath parameters during the numerical fitting procedures used in DMFT, it is most effective to work with the impurity GF on the Matsubara frequency axis. Although the forms mentioned refer to operators that exist within an exponential Hilbert space, Eq. S6 can also be expressed in a manner that separates the noninteracting component of the impurity GF — which includes the impurity on-site energy ($\nu_{11} = \epsilon_i - \mu$, with μ the chemical potential) and impurity-bath hybridization ($\Delta(i\omega_n)$) — and the local electron interactions of the impurity orbital, called the self energy ($\Sigma_{\text{imp}}(i\omega_n)$).

$$\mathcal{G}_{\text{imp}}(i\omega_n) = [i\omega_n - \epsilon_i + \mu - \Delta(i\omega_n) - \Sigma_{\text{imp}}(i\omega_n)]^{-1} \tag{S7}$$

The hybridization with the bath has an exact analytical form of

$$\Delta(i\omega_n) = \sum_{b=1}^{N_B} \frac{|V_b|^2}{i\omega_n - \epsilon_b} \tag{S8}$$

where V_b, ϵ_b are the bath hoppings and on-site energies. The self-energy of the impurity model comes from Dyson's equation

$$\mathbf{G} = \mathbf{G}^0 + \mathbf{G}^0\mathbf{\Sigma}\mathbf{G} \tag{S9}$$

where $\mathbf{G}^0$ is the non-interacting GF (which has an exact analytical form) and $\mathbf{\Sigma}$ is the self-energy. Here, the quantities in Eq. S9 are $2N \times 2N$ matrices, with each of the elements signifying all of the system's single-particle GFs with different spin and orbital creation/annihilation operators $[\hat{d}_\sigma^{(\dagger)}, \hat{c}_{1\sigma}^{(\dagger)}, ..., \hat{c}_{N_B\sigma}^{(\dagger)}]$.

As discussed in the main text, the impurity model is connected to the lattice model through its self-energy. In the case of the lattice model with infinite coordination, the lattice self-energy (Σ_{latt}) for each interacting orbital is independent of crystal momentum $\mathbf{k}$. This can be understood as each individual orbital being influenced by many-body interactions from the bulk in a uniform manner, regardless of its actual position within the lattice, due to the infinite coordination.

$$\Sigma_{\text{latt}}(i\omega_n, \mathbf{k}) \to \Sigma_{\text{latt,loc}}(i\omega_n) \tag{S10}$$

Using the self-consistent condition

$$\Sigma_{\text{imp}} \approx \Sigma_{\text{latt,loc}} \tag{S11}$$

the approximation of the lattice GF at the current iteration of DMFT can be evaluated with Eq. S12 using the assumption made with Eq. S11.

$$G_{\text{latt,loc}}(i\omega_n) \approx \int_{-\infty}^{\infty} dx \frac{\rho(x)}{i\omega_n - x + \mu - \Sigma_{\text{imp}}(i\omega_n)} \tag{S12}$$

$$\rho(x) = \frac{\sqrt{4h^2 - x^2}}{2\pi h^2} \tag{S13}$$

where $\rho(x)$ is the density of states for a Bethe lattice and h is the hopping strength of the Hubbard Hamiltonian. Returning to Eq. S7, we can obtain new bath parameters for the impurity model by minimizing

$$\mathcal{L}(V_b, \epsilon_b) = \left| \Delta(i\omega_n) - \left[\frac{1}{G_{\text{latt,loc}}(i\omega_n)} - i\omega_n + \epsilon_i - \mu + \Sigma_{\text{imp}}(i\omega_n) \right] \right|^2 . \tag{S14}$$

It is this fitting procedure that necessitates the dynamical quantities $(\mathcal{G}_{\text{imp}}, \Sigma_{\text{imp}}, G_{\text{latt,loc}}, \Delta)$ to vary smoothly with respect to their dependent variable, motivating the choice to perform DMFT using Matsubara frequencies. After some iterations of DMFT, if the change in V_b, ϵ_b is below some threshold, the impurity model's bath can be considered self-consistent in terms of describing the many-body effects on the atomic orbitals of the lattice model, thus concluding the DMFT loop.

B Gaussian states

A fermionic Gaussian Hamiltonian (FGH) for an N-orbital system has the form of Eq. S15, where $\mathbf{H}_{ij}^{\sigma\sigma'}$ will be referred to as elements of the *hopping matrix*.

$$\widehat{\mathbf{H}}_0 = \sum_{ij,\sigma\sigma'}^{N} \mathbf{H}_{ij}^{\sigma\sigma'} \hat{c}_{i\sigma}^{\dagger} \hat{c}_{j\sigma'} \tag{S15}$$

with $\sigma, \sigma' \in \{\uparrow, \downarrow\}$. While $\widehat{\mathbf{H}}_0$ is a linear operator that exists in the full Hilbert space $(2^{2N} \times 2^{2N})$, the hopping matrix $\mathbf{H}$ of size $2N \times 2N$ contains all of the information needed to characterize eigenstates of the FGH.

All fermionic Gaussian states (FGS) $|\phi\rangle$ can be described by their set of occupations within a skew-symmetric covariance matrix $\mathbf{M}$ that takes the form

$$[\mathbf{M}]_{pq} = -\frac{i}{2} \langle\phi| \left[\hat{\gamma}_p, \hat{\gamma}_q\right] |\phi\rangle \tag{S16}$$

where $[\hat{\gamma}_p, \hat{\gamma}_q]$ is the commutator of Majorana operators.

The magnitude of the overlap between two FGS can be computed with

$$\langle\phi_i|\phi_j\rangle = \sqrt{\mathrm{Pf}\left(\frac{\mathbf{M}_i + \mathbf{M}_j}{2}\right)} \tag{S17}$$

where Pf(.) is the Pfaffian. We assume $\langle\phi_i|\phi_j\rangle$ to always be positive, which can be guaranteed by gauge-fixing our states $|\phi_k\rangle$ on the level of a Fock basis element [1].

The matrix elements of any monomial of Majorana operators can be computed with Eq. S18, where $\hat{\boldsymbol{\gamma}}[\mathbf{x}] = \hat{\gamma}_1^{x_1}\hat{\gamma}_2^{x_2}...\hat{\gamma}_{2N}^{x_{2N}}$ and $x_k \in \{0, 1\}$.

$$\langle\phi_i|\hat{\boldsymbol{\gamma}}[\mathbf{x}]|\phi_j\rangle = \langle\phi_i|\phi_j\rangle \mathrm{Pf}(i\boldsymbol{\Delta}[\mathbf{x}]^*) \tag{S18}$$

Here, $\boldsymbol{\Delta}$ is a skew-symmetric matrix described by

$$\boldsymbol{\Delta} = (-2\mathbf{I} + i(\mathbf{M}_i - \mathbf{M}_j))(\mathbf{M}_i + \mathbf{M}_j)^{-1} \tag{S19}$$

and $\boldsymbol{\Delta}[\mathbf{x}]^*$ represents the submatrix with rows and columns supported by $\mathbf{x}$. $\boldsymbol{\Delta}^*$ denotes the complex conjugate. Eq. S17 and Eq. S18 are readily evaluated with all the information obtained from the hopping matrix, which scales linearly with the system size, allowing for efficient calculations using a basis of non-orthogonal FGS.

C Correlation Function and Corresponding Circuits

The correlation function we calculate (for a single impurity) is given as follows:

$$G^R_{\mathrm{imp}}(t) = -i \left\langle \Psi^{(0)} \right| \{\hat{d}(t), \hat{d}^\dagger\} \left| \Psi^{(0)} \right\rangle . \tag{S20}$$

Here $\hat{d}(t)$ represents the time evolved $\hat{d}$ operator in the Heisenberg picture, the state $|\Psi^{(0)}\rangle$ is the ground state of the impurity model, and $\hat{d}$ and $\hat{d}^\dagger$ are fermion annihilation and creation operators (for now, we neglect the spin indices σ), which are not unitary and cannot be directly implemented in a quantum circuit. However, the linear combinations $\hat{\gamma}_+ = \hat{d} + \hat{d}^\dagger$ and $\hat{\gamma}_- = i(\hat{d} - \hat{d}^\dagger)$ are unitary operators, and therefore can be applied on a quantum circuit. The Green's function in Eq. S20 can then be written in terms of $\hat{\gamma}_\pm$ as follows

$$\begin{aligned} G^R_{\mathrm{imp}}(t) &= \frac{1}{4} \left\langle \Psi^{(0)} \right| \{(\hat{\gamma}_+ + i\hat{\gamma}_-)(t), (\hat{\gamma}_+ - i\hat{\gamma}_-)\} \left| \Psi^{(0)} \right\rangle \\ &= \frac{1}{4} \left\langle \Psi^{(0)} \right| \{\hat{\gamma}_+(t), \hat{\gamma}_+\} \left| \Psi^{(0)} \right\rangle - \frac{i}{4} \left\langle \Psi^{(0)} \right| \{\hat{\gamma}_+(t), \hat{\gamma}_-\} \left| \Psi^{(0)} \right\rangle \\ &\quad + \frac{i}{4} \left\langle \Psi^{(0)} \right| \{\hat{\gamma}_-(t), \hat{\gamma}_+\} \left| \Psi^{(0)} \right\rangle + \frac{1}{4} \left\langle \Psi^{(0)} \right| \{\hat{\gamma}_-(t), \hat{\gamma}_-\} \left| \Psi^{(0)} \right\rangle \end{aligned} \tag{S21}$$

In addition to their unitarity, $\gamma_\pm$ are also Hermitian operators, which yields

$$\begin{aligned} G^R_{\text{imp}}(t) &= \frac{1}{4}\text{Re}\left\langle\Psi^{(0)}\right|\hat{\gamma}_+(t)\hat{\gamma}_+\left|\Psi^{(0)}\right\rangle - \frac{i}{4}\text{Re}\left\langle\Psi^{(0)}\right|\hat{\gamma}_+(t)\hat{\gamma}_-\left|\Psi^{(0)}\right\rangle \\ &\quad + \frac{i}{4}\text{Re}\left\langle\Psi^{(0)}\right|\hat{\gamma}_-(t)\hat{\gamma}_+\left|\Psi^{(0)}\right\rangle + \frac{1}{4}\text{Re}\left\langle\Psi^{(0)}\right|\hat{\gamma}_-(t)\hat{\gamma}_-\left|\Psi^{(0)}\right\rangle \\ &= \frac{1}{4}\sum_{a,b\in\{-,+\}} s_{a,b}\,\text{Re}\left\langle\Psi^{(0)}\right|\hat{\gamma}_a(t)\hat{\gamma}_b\left|\Psi^{(0)}\right\rangle, \end{aligned} \tag{S22}$$

where we defined $s_{++}=1$, $s_{+-}=-i$, $s_{-+}=i$, $s_{--}=1$. To calculate each term on a quantum computer, we approximate the ground state as a linear combination of the FGS:

$$\left|\psi^{(0)}\right\rangle = \sum_{k=1}^{\chi}\alpha_k^{(0)}\left|\phi_k\right\rangle = \sum_{k=1}^{\chi}\alpha_k^{(0)}U_k\left|0\right\rangle. \tag{S23}$$

Here $|\phi_k\rangle$ are FGS, and therefore U_k can be chosen as free fermion evolution operators [2]. With this, the Green's function can be calculated as follows:

$$G^R_{\text{imp}}(t) = \frac{1}{4}\sum_{a,b\in\{-,+\}} s_{ab}\text{Re}\left(\sum_{i,j=1}^{\chi}\alpha_i^*\alpha_j\left\langle 0\right|U_i^\dagger e^{itH}\hat{\gamma}_a e^{-itH}\hat{\gamma}_b U_j\left|0\right\rangle\right), \tag{S24}$$

where we applied the Heisenberg operator evolution $\hat{\gamma}_a(t) = e^{itH}\hat{\gamma}_a e^{-itH}$ with $\widehat{\mathbf{H}}_{\text{imp}}$ being simplified to a general Hamiltonian H — assumed to have a limited number of interaction terms — for notational convenience.

We use the Hadamard test to calculate each term using a quantum computer. This requires us to use an ancilla qubit in the $|+\rangle$ state, apply the unitary $U_i^\dagger\hat{\gamma}_a(t)\hat{\gamma}_b U_j$ in a controlled manner i.e. only for the component that has the ancilla qubit in $|1\rangle$ state, and finally measure the ancilla qubit on X direction for the real part, and Y direction for the imaginary part of the transition amplitude. This corresponds to the following circuit:

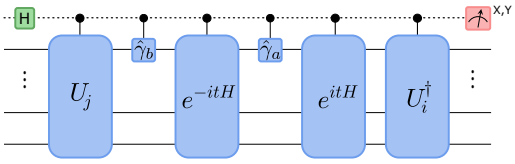

(S25)

We will apply several layers of simplification to this circuit, which will be explained in detail in the following subsections.

C.1 Controlled unitary simplification

First, we will show how to simplify the circuit given in Eq. S25 by half in depth for *any* Hamiltonian H, and *any* unitaries U_i, U_j, $\hat{\gamma}_a$ and $\hat{\gamma}_b$. We will frequently use the following equality:

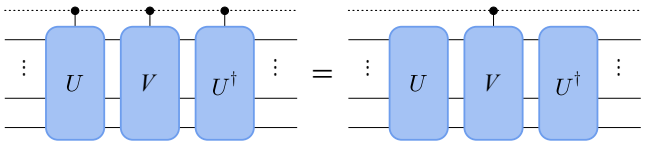

(S26)

To see that these two circuits are equal, let us follow what they implement for each state of the ancilla qubit. If the ancilla is in the $|0\rangle$ state, none of the unitaries are applied in the circuit on the left since they are controlled, and therefore the operation is the identity matrix I. In the circuit on the right, first U is applied. V is not applied since it is controlled, and then $U^\dagger$ is applied, yielding that the circuit is equivalent to $U^\dagger U = I$, which is the same as the circuit on the left-hand side. If the ancilla is in state $|1\rangle$, it can be seen that both circuits will apply $U^\dagger V U$, and therefore they are equal in this case as well.

A direct application of Eq. S26 on the circuit in Eq. S25 allows us to remove the controls on the time evolution operators by setting $V \equiv \hat{\gamma}_a$, and $U \equiv e^{-iHt}$, which yields

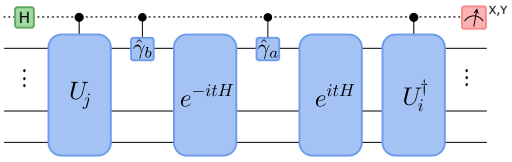

(S27)

To make use of Eq. S26 further, let us introduce an identity operator by using controlled U_i and controlled $U_i^\dagger$ as follows:

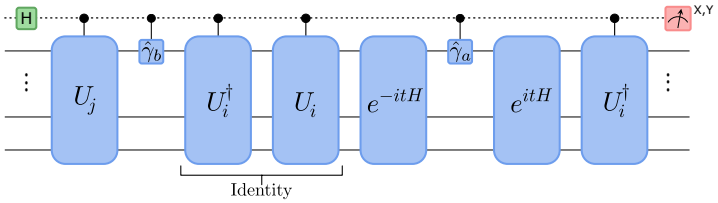

(S28)

We can lose the controls on the U_i we placed and the $U_i^\dagger$ at the end by using Eq. S26:

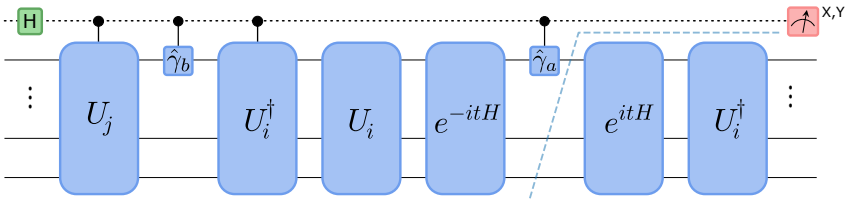

(S29)

Now it can be seen that $U_i^\dagger$ and e^{itH} at the end are outside the light cone of the ancilla measurement, and therefore can be discarded. In addition, the operators U_i and controlled $U_i^\dagger$ are equivalent to an anti-controlled U_i, i.e., it is U_i when the ancilla is in $|0\rangle$, and it is the identity when the ancilla is in $|1\rangle$. Applying these changes, we arrive at the following:

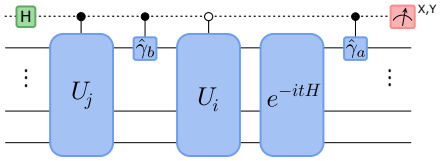

(S30)

Finally, anti-controlled U_i and controlled $\hat{\gamma}_b$ act on separate parts of the Hilbert space, and therefore commute. We can use this to change their order, to obtain the following simplified circuit for calculating the transition amplitude $\langle 0| U_i^\dagger \hat{\gamma}_a(t) \hat{\gamma}_b U_j |0\rangle$:

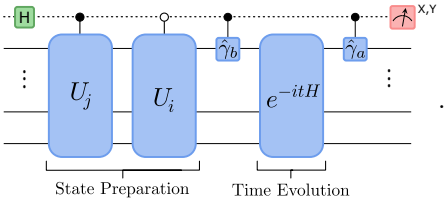

(S31)

Grouping the controlled U_j and anti-controlled U_i together allows us to separate the circuit into state preparation and time evolution circuits.

C.2 Partial compression of the second-order Trotter time evolution

In this subsection, we will focus on the time evolution operator $e^{-it\widehat{\mathbf{H}}_{\text{imp}}}$ for the impurity Hamiltonian. Let us separate the Hamiltonian into two parts as follows:

$$\widehat{\mathbf{H}}_{\text{imp}} = \widehat{\mathbf{H}}_2 + \widehat{\mathbf{H}}_4, \tag{S32}$$

where

$$\widehat{\mathbf{H}}_2 = \sum_{ij\sigma}^{N_I} \nu_{ij} \hat{d}^\dagger_{i\sigma} \hat{d}_{j\sigma} + \sum_{i}^{N_I} \sum_{b\sigma}^{N_B} \epsilon_{ib} \hat{c}^\dagger_{ib\sigma} \hat{c}_{ib\sigma} + \sum_{i}^{N_I} \sum_{b\sigma}^{N_B} V^i_b (\hat{d}^\dagger_{i\sigma} \hat{c}_{ib\sigma} + \text{h.c.}) \tag{S33}$$

which contains all the quadratic terms that are free fermionic, and

$$\widehat{\mathbf{H}}_4 = U \sum_{i}^{N_I} \hat{n}_{i\uparrow} \hat{n}_{i\downarrow} + U' \sum_{i\neq j}^{N_I} \sum_{\sigma\sigma'} \hat{n}_{i\sigma} \hat{n}_{j\sigma'}, \tag{S34}$$

with $\hat{n}_{i\sigma} = \hat{d}^\dagger_{i\sigma} \hat{d}_{i\sigma}$. This contains the quartic terms that correspond to the Coulomb interactions taking place on the impurity orbitals. For clarity, $\widehat{\mathbf{H}}_4$ contains only one term for a single impurity model ($\widehat{\mathbf{H}}_4(N_I = 1) = U \hat{n}_\uparrow \hat{n}_\downarrow$).

To implement $e^{-it\widehat{\mathbf{H}}_{\text{imp}}}$, we will use the second-order Trotter-Suzuki formula:

$$\begin{aligned} e^{-it\widehat{\mathbf{H}}_{\text{imp}}} = & \left(e^{-i\frac{t}{2r}\widehat{\mathbf{H}}_2} e^{-i\frac{t}{r}\widehat{\mathbf{H}}_4} e^{-i\frac{t}{2r}\widehat{\mathbf{H}}_2} \right)^r \\ & + \mathcal{O}\left\{ \left(\left|\left| [\widehat{\mathbf{H}}_2, [\widehat{\mathbf{H}}_2, \widehat{\mathbf{H}}_4]] \right|\right| + \left|\left| [\widehat{\mathbf{H}}_4, [\widehat{\mathbf{H}}_2, \widehat{\mathbf{H}}_4]] \right|\right| \right) \frac{t^3}{r^2} \right\}, \end{aligned} \tag{S35}$$

where $||.||$ refers to the spectral norm. We would like to note that in the generic case, this spectral norm would be $\mathcal{O}(N_I^4)$ simply because the number of Pauli strings in $\widehat{\mathbf{H}}_4$ after the Jordan-Wigner transformation is proportional to the square of the number of impurity orbitals (U' term in Eq. S34 is a sum over impurity pairs). This then yields that the Trotter error given in Eq. S35 is $\mathcal{O}(N_I^4 t^3 r^{-2})$, and is independent of N_B. For simplicity, we again set $N_I = 1$ and $N_B = 3$ for the following illustrations.

Let us now generate a circuit for the trotterized time evolution given in Eq. S35. $\widehat{\mathbf{H}}_2$ contains on-site terms on the bath orbitals and the impurity orbitals, as well as hoppings between the impurity orbitals and each bath orbital. However, each of these terms is limited within its spin sector: i.e., there is no hopping altering the spin of the electrons. Thus, $\widehat{\mathbf{H}}_2$ is actually two free fermionic Hamiltonians that act on effectively different orbitals, which are represented by different sets of qubits. Using this information, and the algebraic compression of free fermions on a generic lattice given in [2], we can then build the following circuit for a single Trotter step ($r = 1$) given in Eq. S35:

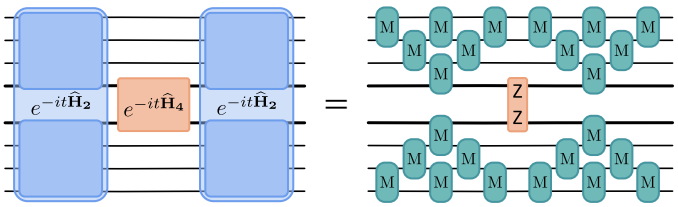

(S36)

where we used the fact that the interaction term in $\widehat{\mathbf{H}}_4$ is a density-density interaction which becomes a ZZ rotation after a Jordan-Wigner transformation. The bold qubit wires correspond to the impurity qubits, whereas the other thinner wires are the bath qubits. Note that we compressed the spin $\downarrow$ hoppings in the lower register into a triangle, and the spin $\uparrow$ hoppings in the upper register into an upside-down triangle. This is possible since the block rules are symmetric under the up-down parity [2–4].

Now let us compress one Trotter step partially. For certain free fermion matchgates, or TFXY blocks, we can use the block properties to reduce a Trotter step as follows:

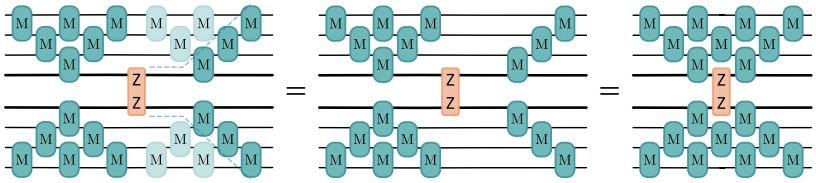

(S37)

The lighter-colored blocks on the left-most circuit act only on the bath qubits, and therefore can commute with the impurity term, which only acts on the impurity qubits. This allows the triangles on the left side of the ZZ gate to absorb them via the compression given in Thm. 1, Ref. [3], which leads to the circuit in the middle. Note that this cuts down the number of matchgates approximately by half for each Trotter step.

A similar partial compression can be applied to the subsequent Trotter steps as well ($r > 1$), which leads to an even greater simplification for the time evolution circuit. As an example, let us consider $r = 3$:

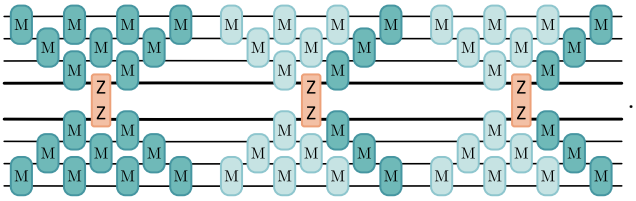

(S38)

As can be seen, the lighter-colored blocks on the third step can be absorbed by the second step via algebraic compression. The same can be done for the second Trotter step, which leads to the following circuit

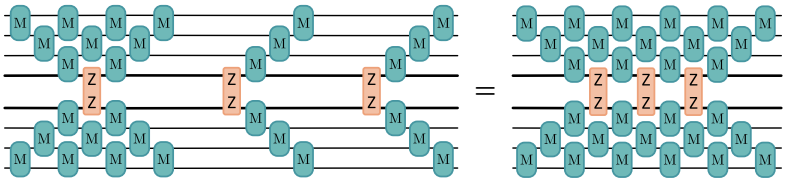

(S39)

Consequently, in cases where the impurity is small compared to the full system size, which is commonly observed in impurity models, the cost of additional Trotter steps is greatly minimized.

C.3 Controlled state preparation and further compression

In this subsection, we will show how to build the state preparation circuit at the beginning of the time evolution circuit, i.e., the controlled U_j and anti-controlled U_j. In this work, we are representing the ground state as a linear combination of FGS. This means that the unitaries U_i can be written as free fermionic evolution operators, and controlled U_i and anti-controlled U_j are controlled free fermionic evolutions with a single ancilla. Using the $Q-$compression algorithm [2], we can compress these unitaries into a diamond structure, and obtain the following circuit:

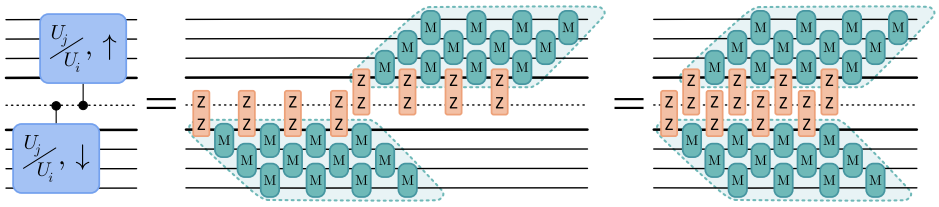

(S40)

Here, the middle qubit is the ancilla qubit (dashed line), the upper qubits are the spin $\uparrow$ qubits, and the lower qubits are the spin $\downarrow$ qubits. As the time evolution circuit contains two triangle (one reversed) structures, this circuit contains two diamond structures. Due to the conservation of number of spin $\uparrow$ and spin $\downarrow$ particles separately, we chose the FGS $|\phi_i\rangle$ as a tensor product of a FGS on each spin, i.e. $|\phi_i\rangle = |\phi_i^{\uparrow}\rangle \otimes |\phi_i^{\downarrow}\rangle$. This keeps the state preparation circuit separate and leads to two diamond structures, which is depicted in the middle of Eq. S40. Because the ZZ terms in each spin sector commute, we can combine the state prep into the compact form shown on the right of Eq. S40.

Combining this state preparation circuit and the simplified time evolution circuit in Eq. S39, we can implement the simplified Hadamard test circuit Eq. S31 which calculates the correlation $\langle 0| U_i^{\dagger} \hat{\gamma}_a(t) \hat{\gamma}_b U_j |0\rangle$ as the following:

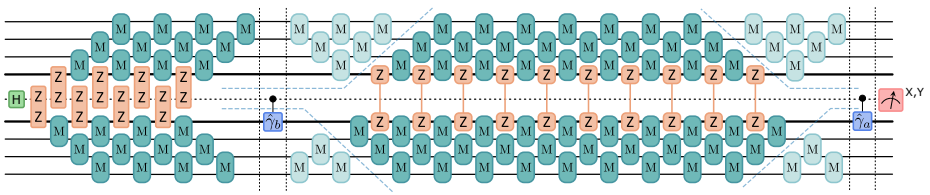

(S41)

Here we have illustrated $r = 10$. Note that the impurity ZZ term is split since the middle qubit is the ancilla qubit.

As illustrated by the lighter-colored matchgates in Eq. S41, this circuit can be further simplified. Firstly, those lighter-colored matchgates at the beginning of the time evolution circuit can pass the controlled $\hat{\gamma}_b$ operator, and get absorbed by the diamond structures in the state preparation circuit. Secondly, the lighter-colored matchgates at the end of the time evolution circuit can pass the controlled $\hat{\gamma}_b$ operator as well as

the measurement operation. Therefore, they do not affect the measurement result and can be discarded to reduce the gate count even further. In the end, the following is the final circuit structure we ran on hardware (for r=10):

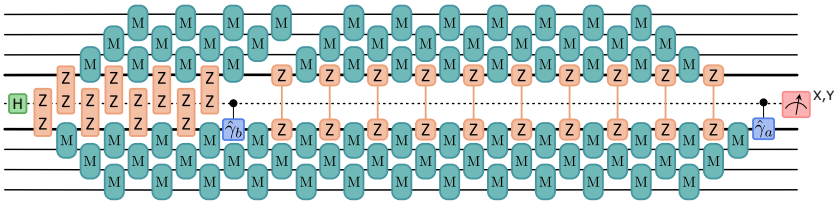

(S42)

C.4 Quantum resource estimation

We now decompose the quantum resource requirements for performing DMFT into three parts: the cost per circuit, the cost per evaluation of $G^R_{\text{imp}}(t)$, and the cost per DMFT self-consistency cycle. The first quantifies as the number of gates per circuit, while the last two result in the number of total circuit evaluations. Each circuit evaluation's runtime (in terms of wall time) may vary depending on the hardware and number of qubits. For our hardware runs shown in Fig. 7 of the main text, the total runtime was 127 seconds for the 200 circuit evaluations ($n_t = 20$ with 10 cycles of Pauli twirling).

C.4.1 Per-circuit cost in terms of CNOTs

Using the basic structure of the circuits in this work, we compute the two-qubit gate costs per circuit as a function of the total number of bath orbitals Λ, the total number of impurity orbitals N_I, and the number of Trotter steps r. We also make use of $N_q = 2(N_I + \Lambda)$, which is the number of qubits, neglecting the ancilla. The resource estimation we quote is for the most general case, which also scales the worst, where the impurity and bath qubits are fully connected. Estimation of the Trotter depths required for self-consistent solutions of the impurity solver, especially for systems beyond 30-40 qubits, is non-trivial and left for future work.

Preparing the initial states $|\phi_i\rangle, |\phi_j\rangle$ requires two (one for each spin sector) linear-depth product formula circuits as shown in Eq. S40. Each exponential block, a rotation around Z_iZ_j or $\alpha(X_iX_j + Y_iY_j) + b_iZ_i + b_jZ_j$, requires two CNOT gates. The state preparation then requires $N_q^2 = 4(N_I + \Lambda)^2$ two-qubit gates. The subsequent second-order Trotter steps each require an evolution of the impurity interaction and the compressed free-fermion evolution. The former includes all-to-all interactions between each orbital and spin configuration, requiring $2N_I(2N_I - 1)$ CNOTs. The free-fermion evolution reduces to a ladder of matchgates with length Λ and width N_I, and hopping within the impurity qubits requiring $N_I(N_I - 1)/2$ matchgates. One copy of the circuit is needed for each spin sector, and each matchgate again requires two CNOTs, giving a total cost per trotter step of $6N_I^2 + 4N_I(\Lambda - 1)$.

Next, there are an additional 2 CNOTs required to implement the controlled-Pauli operators $\hat{\gamma}_a$ and $\hat{\gamma}_b$, and an additional single layer of $\frac{N_q}{2} - 1$ matchgates in a single spin

sector to account for the elements which do not commute with $\hat{\gamma}_b$. Finally, there are additional circuit reductions due to the circuit light-cone of the measurement. For one spin sector, a full triangle of $\frac{N_q}{2}\left(\frac{N_q}{2}-1\right)$ CNOTs can be neglected, while in the other spin sector $\left(\frac{N_q}{2}-1\right)\left(\frac{N_q}{2}-2\right)$ CNOTs can be neglected. Thus, for circuits of Trotter depth $rN_I \geq \frac{N_q}{2}-1$, the final gate cost for circuits of the form shown in Eq. S42 is:

$$\begin{aligned}\mathrm{CNOT}(N_I,\Lambda,r) &= 4(N_I+\Lambda)^2 + r(6N_I^2+4N_I(\Lambda-1))+2 \\ &\quad -2(N_I+\Lambda)^2+4(N_I+\Lambda)-2 \\ &= 2N_I^2+2\Lambda^2+6N_I\Lambda+4N_I+4\Lambda \\ &\quad + r(6N_I^2+4N_I(\Lambda-1)).\end{aligned} \tag{S43}$$

The circuit depths required scale linearly in Trotter-depth and $N_I N_q$ rather than N_q^2. As $N_I \ll N_q$ in most cases of interest, this provides a quadratic reduction in gate costs. For this work, with $N_I = 1, N_B = 3$, the largest circuit used has $r = 18$ and 306 CNOT gates before additional generic compiling.

As a comparison, Ref. [5] estimates that a 16 qubit single-impurity model ($N_I = 1, \Lambda = 7$) using circuits of up to 6 Trotter steps and a CNOT depth of 254, which is approximately two thousand CNOT gates. Although that work computes expectations of the form $\langle\phi_i| U^\dagger(t_l)^\dagger \hat{A} U(t_m) |\phi_j\rangle$, and as such is not a direct comparison, our circuit structure for a Trotter depth of 6 for this model size requires only 354 CNOTs. In a similar study, [6] estimate the resources required to prepare the ground state using matrix product states compiled onto a quantum circuit. For 20 and 40 qubit models ($N_I = 3$ and $\Lambda = 7, 17$), they estimate that to prepare an initial quantum state with a fidelity of 0.99 to the ideal ground state, one would require circuits of approximately one thousand two-qubit gates in the 20 qubit case and four thousand two-qubit gates in the 40 qubit case. Using the method presented in this work, for the same total gate counts, both state preparation and time evolution with approximately 6 and 12 Trotter layers could be computed for the 20 and 40 qubit cases, respectively. Again, the comparison is not direct, as the SGS basis used in this work requires an additional factor of χ^2 circuit evaluations to produce expectations with respect to the ground state, compared to a single state preparation using matrix product state methods. Nevertheless, our CNOT count is amenable to today's noisy hardware, even for larger system sizes.

C.4.2 Per-Green's-function cost

Reconstructing one retarded impurity Green's function $G^R_{\mathrm{imp}}(t)$ requires $\chi^2 n_t$ circuit evaluations, where χ is the number of fermionic Gaussian states in the SGS and n_t is the number of measured time-domain samples. The χ^2 scaling reflects that each ordered pair of FGS in the SGS contributes an independent matrix element to the decomposition of G^R_{imp} in the SGS basis, with each matrix element requiring its own Hadamard-test circuit. Error-mitigation overhead — Pauli twirling and zero-noise extrapolation — multiplies this base count further: in the hardware demonstration

of Fig. 7 in the main text, 10 Pauli-twirling cycles brought the total per-data-point circuit count from $\chi^2 n_t = 20$ to 200.

Both χ and n_t can be bounded in advance for a given problem. The Bravyi-Gosset complexity result [7] bounds $\chi \sim \exp[O(b^3)]$ at fixed impurity size and additive precision $\epsilon = 2^{-b}$, while empirical experience with selective configuration interaction methods suggests polynomial growth in the active-space dimension for impurity sizes of practical interest. The number of measured time-domain samples n_t is, importantly, decoupled from the frequency resolution required of the final spectrum: the PSD method introduced in the Hardware Runs section uses the positive-definite extension of response functions to synthesize a long signal from a short measurement, rather than the much larger Fourier-transform sample count required to resolve sharp features. All $\chi^2 n_t$ circuit evaluations contributing to a single $G^R_{\mathrm{imp}}(t)$ are mutually independent and may be executed in parallel across multiple quantum processors or concurrently within a single device, making the χ^2 factor amenable to direct wall-clock reduction on multi-QPU infrastructure.

C.4.3 Per-DMFT-loop cost

The remaining factor in estimating total resource cost is the number of DMFT self-consistency iterations required to reach convergence, which is not known a priori for any given problem. Convergence is typically assessed by the change in the impurity model's bath parameters between successive iterations falling below a chosen tolerance. In the DMFT literature [8–10], one can typically expect a few iterations to converge away from phase boundaries, with the rate set by the mixing of parameters between successive iterations [11].

Near phase boundaries — for example, the Mott transition in the half-filled single-band Hubbard model — the convergence rate slows substantially, and tens to hundreds of iterations are not uncommon. This critical slowing down reflects the marginal stability of the DMFT self-consistency map as the system approaches the transition. Acceleration via mixing schemes can substantially reduce the iteration count in this regime, but does not eliminate the underlying slowdown.

For DMFT scans across interaction strength, a "warm start" — initializing each parameter point with the converged solution from a nearby point — typically reduces the iteration count to a handful per parameter step, and to a single iteration in favorable cases when the parameter step is small relative to the local stability of the fixed-point. The total iteration count for a full phase diagram is therefore typically only modestly larger than that for a single converged solution. As a working estimate, a conservative upper bound on the per-problem iteration count for multi-orbital impurity systems near critical points is $\mathcal{O}(100)$, while warm-started scans away from criticality may require only $\mathcal{O}(10)$ or fewer.

D Quantum error mitigation and signal processing

Obtaining the hardware results presented as the magenta "De-noised & Extended" curve in Fig. 7(c) of the main text involved several levels of refinement to achieve a signal from which fruitful analysis could occur. A combination of error mitigation

performed on the hardware and classical post-processing was required to extract meaningful correlation functions in the frequency domain. In this section, we will explore in more detail the techniques used and the motivations behind them.

D.1 Gate-based error mitigation

Once the measurement of our quantum circuits occurs, the information lost to qubit dephasing during runtime cannot be fully recovered with classical post-processing techniques. To mitigate some of that dephasing, we employed dynamical decoupling (DD) [12]. At the cost of a few single-qubit Pauli operations, DD acts as a deterrent to dephasing due to qubit idling. While some active qubits undergo long two-qubit operations, those that are idle will become susceptible to coupling to the environment, leading to errors. DD will act as a deterrent to these interactions by inserting a series of time dependent pulse sequences (which together amount to the identity operation) on the idle qubits. In practice, the single-qubit gate operations of a DD pulse sequence are orders of magnitude shorter than two-qubit gate operations, resulting in the benefits of including DD greatly outweighing the single-qubit gate costs.

For our hardware runs, we employed the $XY4$ pulse sequence described by Fig. S1. Here, τ is the total idle time of the qubit, and X, Y are single-qubit Pauli operations. In the time period for which we performed our hardware runs, IBM's runtime service gives access to time dependent information in transpiled circuits, which scans for idle periods. We took advantage of the Qiskit Research [13] Python package to implement the DD sequence.

Fig. S1 XY4 pulse sequence. The purple f "gates" represent the delay periods between the X and Y Pauli gates, with τ being the total idle time.

Although qubit dephasing can be alleviated by the inclusion of DD, coherent errors due to the bias of the noisy quantum hardware are still prevalent. During the execution of a quantum circuit, if a specific gate introduces an error, say from miscalibration, repeatedly executing this gate results in compounding (or coherent) errors. To counteract coherent errors, we used an error mitigation technique called Pauli twirling, sometimes referred to as randomized compiling [14, 15].

Pauli twirling is similar to DD, such that it introduces a few (relatively short) single-qubit Pauli operations into our circuits. A randomly selected set of Pauli gates sandwiches all two-qubit gates as a means to locally transform the state of the qubits before the two-qubit gate operation, yet keeps the logical operation of the two-qubit gate unchanged. Several instances of the same circuit are generated, each with different randomly implemented Pauli-twirled two-qubit gates, and the resulting circuits are averaged over. This is essentially changing the coherent errors in our circuit due to the repeated use of low-fidelity two-qubit gate operations into statistical errors. For our hardware results, we averaged over 10 Pauli-twirled circuits per data point in Fig. 7 of the main text.

D.2 Post-selection and rescaling

Conveniently, the impurity Hamiltonian $\widehat{\mathbf{H}}_{\text{imp}}$ is particle-conserving, which allows us to filter out the measurements received in error in a procedure called post-selection. For our hardware runs in Fig. 7 of the main text, we used a half-filled impurity model with $N_I = 1$ and $N_B = 3$, meaning any shots that did not obey this particle content were certainly produced in error, and could be disregarded. Each data point corresponds to 40,000 shots (4000 shots $\times$ 10 twirls per data point), and the highest shot acceptance ratio per data point was 99.1% (first data point), with the lowest being 68.3% (18^{th} data point).

These post-selected results are further improved by rescaling the data based on 1- and 2-qubit gate error rates, ϵ_{1q} and ϵ_{2q} respectively, of `ibm_sherbrooke` at the time of our hardware runs as reported by IBM's quantum compute resources [16]. The rescaling formula is given by Eq. S44. Here, N_{1q} and N_{2q} are the numbers of one- and two-qubit gates in a circuit, respectively. $\hat{O}_{\text{noisy}}$ in our case is the raw measurement of the correlator at a particular time, and $\hat{O}_{\text{rescaled}}$ is the mitigated result using rescaling.

$$\hat{O}_{\text{rescaled}} = \frac{\hat{O}_{\text{noisy}}}{(1-\epsilon_{2q})^{N_{2q}}(1-\epsilon_{1q})^{N_{1q}}} \tag{S44}$$

It is important to note that all error rates reported by IBM are estimates based on randomized benchmarking [17], and are thus not a perfect representation of the true errors of our circuits. Nevertheless, we can consider these estimates to be the best possible without much effort in completely characterizing the errors within the circuits we ran. Further, as shown by Fig. S2, rescaling has a noticeable impact on recovering much of the signal that is lost due to depolarizing noise within the deeper circuits we ran on hardware.

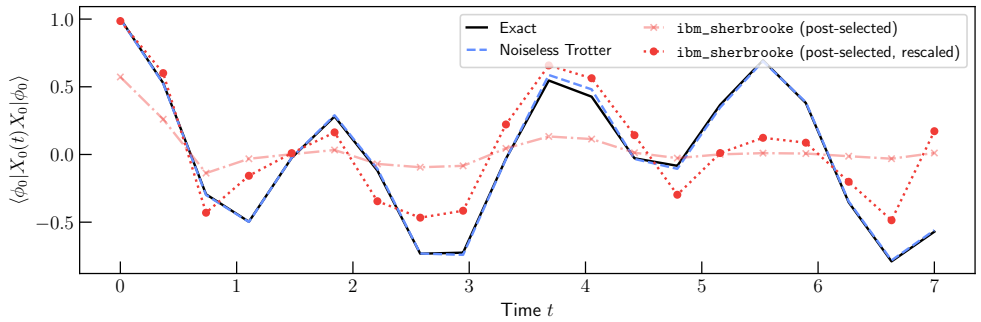

Fig. S2 Effect of rescaling based on gate error rates. While post-selection discards known errors in our data, the amplitude of the signal is quite damped, even at early times (dashed red line, crosses). Using rescaling, much of this dampening can be mitigated (dotted red line, circles). All data is compared to the evaluation of the correlator using ED (black line).

E Data tables and additional results

As discussed in Sec. C.4, the number of circuit evaluations scales quadratically with the rank χ of the SGS. Thus, the practicality of a quantum simulation using the SGS crucially depends on χ remaining manageable for multi-orbital and large-bath impurity models, where quantum implementations offer the greatest advantage. Here, we concretely demonstrate that χ remains small across the system sizes considered in this work.

Supplementing the analysis in Fig. 3(a) of the main text, table S1 reports the range of χ over 10 randomly sampled impurity model parameterizations to achieve the relative energy error of $\leq 10^{-2}$. As described in the main text, we use particle-hole symmetric parameterizations, with baths uniquely sampled for each impurity in the multi-orbital setting. These data are also randomly sampled over interaction strengths U ranging from 0 to 9.

Overall, the SGS offers a greatly compacted representation of the impurity model's ground state where it matters: as the dimension of the particle-selected Hilbert space reaches 400, the fraction of the space required to achieve energy errors $\leq 10^{-2}$ drops below one percent. This advantage extends naturally to the multi-orbital setting.

N_I	N_B per N_I	$\dim(\mathcal{H}_N^{ps})$	χ (min–max)
1	1	4	2–3
1	2	9	1–2
1	3	36	1–6
1	4	100	1–3
1	5	400	1–6
1	6	1225	1–5
1	7	4900	1–6
1	8	15876	1–4
1	9	63504	1–4
2	1	36	2–4
2	2	400	1–4
2	3	4900	1–7
2	4	63504	1–4
3	1	400	3–12
3	2	15876	1–8

Table S1 Dimension of the particle-selected Hilbert space $\dim(\mathcal{H}_N^{ps})$ and rank window χ of the SGS for the system sizes in Fig. 3(a) of the main text. For the randomly sampled bath parameterization and interaction strengths, the required χ to achieve a relative energy error $\leq 10^{-2}$ becomes an increasingly small fraction of $\dim(\mathcal{H}_N^{ps})$ as the system size grows.

Expanding on this demonstration, Fig. S3 runs the subspace search until the termination criteria: adding another Gaussian state to the SGS produces a negligible change in ground state energy. Fig. S3 shows the resulting average relative energy error and average infidelity over the 10 sampled models. The infidelity is computed by expanding

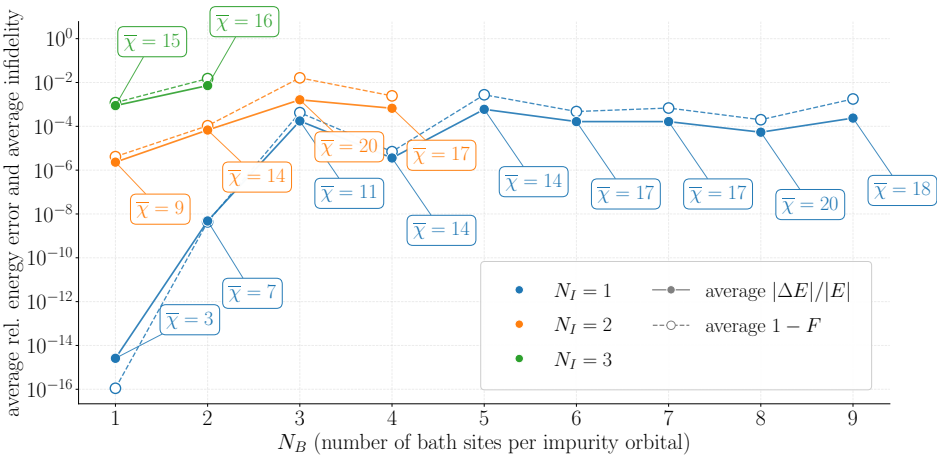

Fig. S3 Average rank and errors of the converged SGS for the system sizes considered in Fig. 3(a) of the main text. Rather than fixing an energy error bound to determine the rank, the subspace search runs to convergence: adding another state to the SGS yields no further reduction in ground state energy. Averaged over 10 randomly sampled impurity models parameterized by (U, V_b^i, ϵ_b^i), the average relative ground state energy error (solid line) and infidelity (dashed line) remain low despite the relatively small average SGS rank ($\overline{\chi}$ annotations) across all system sizes considered.

the SGS as a state vector within the ground state's particle sector. This demonstration shows that the infidelity tracks closely with the energy error upon convergence.

Accompanying the results in Fig. 4 of the main text, table S2 reports the SGS rank for the DMFT-converged impurity ground state across our 1-band Hubbard model calculations on a Bethe lattice. The rank reflects the structure of the phase diagram: at small, perturbative U, a few FGS suffice by construction, and at large U the atomic limit again admits a compact representation. Consequently, χ is peaked near the Mott transition.

U	χ
0.010	4
0.898	17
1.786	17
2.673	20
3.561	23
4.449	24
5.337	21
6.224	16
7.112	14
8.000	13

Table S2 Rank of the SGS for the DMFT results in Fig. 4 of the main text. At weak interaction, few Gaussian states suffice to represent the ground state. Similarly, the rank decreases at large U as the system approaches the atomic limit, which admits a compact representation. The rank peaks near $U \approx 4.449$, reflecting the critical behavior of the self-consistently determined impurity model near the Mott transition.

References

[1] Boutin, S. & Bauer, B. Quantum impurity models using superpositions of fermionic gaussian states: Practical methods and applications. *Phys. Rev. Res.* **3**, 033188 (2021). URL https://link.aps.org/doi/10.1103/PhysRevResearch.3.033188.

[2] Kökcü, E. *et al.* Algebraic compression of free fermionic quantum circuits: Particle creation, arbitrary lattices and controlled evolution. *arXiv preprint* (2025). URL https://doi.org/10.48550/arXiv.2303.09538.

[3] Kökcü, E. *et al.* Algebraic compression of quantum circuits for hamiltonian evolution. *Phys. Rev. A* **105**, 032420 (2022). URL https://doi.org/10.1103/PhysRevA.105.032420.

[4] Camps, D. *et al.* An algebraic quantum circuit compression algorithm for hamiltonian simulation. *SIAM Journal on Matrix Analysis and Applications* **43**, 1084–1108 (2022). URL https://doi.org/10.1137/21M1439298.

[5] Jamet, F., Agarwal, A. & Rungger, I. Quantum subspace expansion algorithm for green's functions. *arXiv preprint* (2022). URL https://arxiv.org/abs/2205.00094.

[6] Jamet, F. *et al.* Anderson impurity solver integrating tensor network methods with quantum computing. *APL Quantum* **2** (2025). URL https://doi.org/10.1063/5.0245488.

[7] Bravyi, S. & Gosset, D. Complexity of Quantum Impurity Problems. *Communications in Mathematical Physics* **356**, 451–500 (2017). URL https://doi.org/10.1007/s00220-017-2976-9.

[8] Caffarel, M. & Krauth, W. Exact diagonalization approach to correlated fermions in infinite dimensions: Mott transition and superconductivity. *Physical review letters* **72**, 1545 (1994).

[9] Gull, E. *et al.* Continuous-time monte carlo methods for quantum impurity models. *Rev. Mod. Phys.* **83**, 349–404 (2011). URL https://link.aps.org/doi/10.1103/RevModPhys.83.349.

[10] Georges, A., Kotliar, G., Krauth, W. & Rozenberg, M. J. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Reviews of Modern Physics* **68**, 13 (1996).

[11] Johnson, D. D. Modified broyden's method for accelerating convergence in self-consistent calculations. *Phys. Rev. B* **38**, 12807–12813 (1988). URL https://link.aps.org/doi/10.1103/PhysRevB.38.12807.

[12] Viola, L., Knill, E. & Lloyd, S. Dynamical decoupling of open quantum systems. *Phys. Rev. Lett.* **82**, 2417–2421 (1999). URL https://doi.org/10.1103/PhysRevLett.82.2417.

[13] developers, T. Q. R. & contributors. Qiskit research (v0.0.2) (2023). URL https://github.com/qiskit-community/qiskit-research?tab=readme-ov-file.

[14] Hashim, A. *et al.* Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor. *Phys. Rev. X* **11**, 041039 (2021). URL https://link.aps.org/doi/10.1103/PhysRevX.11.041039.

[15] Wallman, J. J. & Emerson, J. Noise tailoring for scalable quantum computation via randomized compiling. *Phys. Rev. A* **94**, 052325 (2016). URL https://link.aps.org/doi/10.1103/PhysRevA.94.052325.

[16] Quantum processing units, IBM Quantum Platform, https://quantum.ibm.com/services/resources.

[17] McKay, D. C. *et al.* Benchmarking quantum processor performance at scale. *arXiv preprint* (2023). URL https://doi.org/10.48550/arXiv.2311.05933.

Supplementary Information Figure Legends

Figure 1. XY4 pulse sequence. The purple f "gates" represent the delay periods between the X and Y Pauli gates, with τ being the total idle time.

Figure 2. Effect of rescaling based on gate error rates. While post-selection discards known errors in our data, the amplitude of the signal is quite damped, even at early times (dashed red line, crosses). Using rescaling, much of this dampening can be mitigated (dotted red line, circles). All data is compared to the evaluation of the correlator using ED (black line).

Figure 3. Average rank and errors of the converged SGS for the system sizes considered in Fig. 3(a) of the main text. Rather than fixing an energy error bound to determine the rank, the subspace search runs to convergence: adding another state to the SGS yields no further reduction in ground state energy. Averaged over 10 randomly sampled impurity models parameterized by (U, V_b^i, ϵ_b^i), the average relative ground state energy error (solid line) and infidelity (dashed line) remain low despite the relatively small average SGS rank ($\overline{\chi}$ annotations) across all system sizes considered.