FFT-Accelerated Worldline Monte Carlo for Fermionic Lattice Models

A Determinant-Free Approach with $O(N \log N)$ Complexity

Deqian Kong[1], Shi Feng[2,3], Jianwen Xie[4], Ying Nian Wu[1]

Mar. 16, APS Global Physics Summit 2026

[1]Statistics & Data Science, UCLA [2]Physics, TU Munich [3]MCQST [4]Lambda, Inc.

Outline

1. The computational bottleneck in quantum simulation
2. Key idea: joint sampling of worldlines & auxiliary fields
3. Momentum-space transfer kernel & convolution structure
4. **Two core MCMC moves:**
   - FFBS: Forward Filtering Backward Sampling for worldlines
   - Permutations: resample worldline topology to recover fermion statistics
5. Fast auxiliary field updates
6. Benchmark on 2D Hubbard model

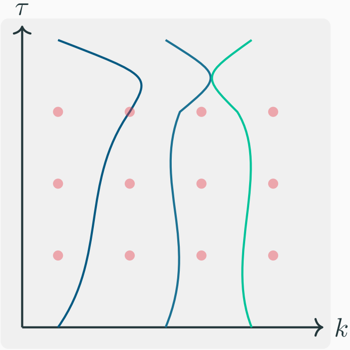

worldlines + aux. fields

DQMC in 2026: Algorithmic Status

Current status

- Half filling on bipartite lattices: unbiased finite-T benchmarks are now high-precision (PRL **134**, 016503 (2025); PRB **111**, 035123 (2025)).
- Besides the sign problem, core bottleneck is still determinant linear algebra; fast-update/submatrix schemes reduce prefactors.

DQMC in 2026: Algorithmic Status

Current status

- Half filling on bipartite lattices: unbiased finite-T benchmarks are now high-precision (PRL **134**, 016503 (2025); PRB **111**, 035123 (2025)).
- Besides the sign problem, core bottleneck is still determinant linear algebra; fast-update/submatrix schemes reduce prefactors.

Method	Cost per sweep	Representative refs
BSS-DQMC	$O(\beta N^4)$	Blankenbecler *et al.*, PRD **24**, 2278 (1981); D. Scalapino, R. Sugar, PRB **24**, 4295 (1981)

DQMC in 2026: Algorithmic Status

Current status

- Half filling on bipartite lattices: unbiased finite-T benchmarks are now high-precision (PRL **134**, 016503 (2025); PRB **111**, 035123 (2025)).
- Besides the sign problem, core bottleneck is still determinant linear algebra; fast-update/submatrix schemes reduce prefactors.

Method	Cost per sweep	Representative refs
BSS-DQMC	$O(\beta N^4)$	Blankenbecler *et al.*, PRD **24**, 2278 (1981); D. Scalapino, R. Sugar, PRB **24**, 4295 (1981)
Low-rank update	$O(\beta N^3)$	Hirsch, PRB **31**, 4403 (1985)

DQMC in 2026: Algorithmic Status

Current status

- Half filling on bipartite lattices: unbiased finite-T benchmarks are now high-precision (PRL **134**, 016503 (2025); PRB **111**, 035123 (2025)).
- Besides the sign problem, core bottleneck is still determinant linear algebra; fast-update/submatrix schemes reduce prefactors.

Method	**Cost per sweep**	**Representative refs**
BSS-DQMC	$O(\beta N^4)$	Blankenbecler *et al.*, PRD **24**, 2278 (1981); D. Scalapino, R. Sugar, PRB **24**, 4295 (1981)
Low-rank update	$O(\beta N^3)$	Hirsch, PRB **31**, 4403 (1985)
Delayed/submatrix batching	Similar asymptotic scaling, smaller prefactor	Nukala *et al.*, PRB **80**, 195111 (2009); Sun, Xu, PRB 109, 235140 (2024)

DQMC in 2026: Algorithmic Status

Current status

- Half filling on bipartite lattices: unbiased finite-T benchmarks are now high-precision (PRL **134**, 016503 (2025); PRB **111**, 035123 (2025)).
- Besides the sign problem, core bottleneck is still determinant linear algebra; fast-update/submatrix schemes reduce prefactors.

Method	Cost per sweep	Representative refs
BSS-DQMC	$O(\beta N^4)$	Blankenbecler *et al.*, PRD **24**, 2278 (1981); D. Scalapino, R. Sugar, PRB **24**, 4295 (1981)
Low-rank update	$O(\beta N^3)$	Hirsch, PRB **31**, 4403 (1985)
Delayed/submatrix batching	Similar asymptotic scaling, smaller prefactor	Nukala *et al.*, PRB **80**, 195111 (2009); Sun, Xu, PRB 109, 235140 (2024)
FFT-FFBS	$\boldsymbol{O(\beta N \log N)}$	**D. Kong, et al., arXiv:2510.13866**

The Determinant Bottleneck in DQMC

Standard approach (Blankenbecler, Scalapino, Sugar 1981):

1. Trotter decompose: $e^{-\beta H} = \prod_\ell e^{-\Delta\tau H}$
2. HS-decouple: introduce auxiliary fields $\Sigma = \{s_{\mathbf{r}\ell}\}$ to linearize
3. Integrate out fermions analytically $\Rightarrow$ determinant

$$Z = \sum_{\Sigma} \underbrace{\prod_{\sigma} \det M_\sigma(\Sigma)}_{O(N^2)\sim O(N^3)\text{ per update}} e^{-S_{\text{aux}}(\Sigma)}$$

The Determinant Bottleneck in DQMC

Standard approach (Blankenbecler, Scalapino, Sugar 1981):

1. Trotter decompose: $e^{-\beta H} = \prod_\ell e^{-\Delta\tau H}$
2. HS-decouple: introduce auxiliary fields $\Sigma = \{s_{\mathbf{r}\ell}\}$ to linearize
3. Integrate out fermions analytically $\Rightarrow$ determinant

$$Z = \sum_{\Sigma} \underbrace{\prod_{\sigma} \det M_\sigma(\Sigma)}_{O(N^2)\sim O(N^3)\text{ per update}} e^{-S_{\text{aux}}(\Sigma)}$$

Key Idea: Don't integrate out the fermions. Sample k-space worldlines *jointly* with HS field (Σ).

Key Idea: Joint Worldline–Auxiliary Field Sampling

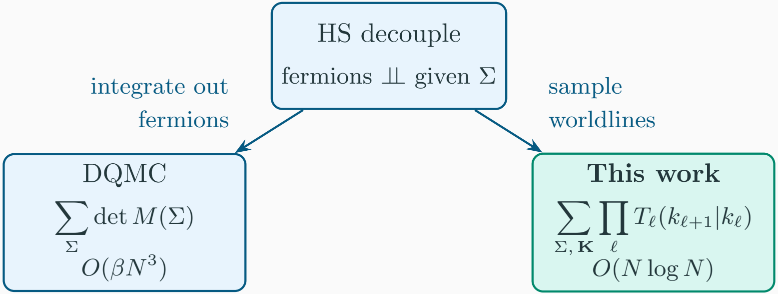

Once decoupled by Σ, each fermion is **independent**:

$$\pi(\mathbf{K}_\uparrow, \mathbf{K}_\downarrow, \Sigma, P_\uparrow, P_\downarrow) \;\propto\; P_{\mathrm{HS}}(\Sigma) \cdot \mathrm{sgn}(P_\uparrow)\,\mathrm{sgn}(P_\downarrow) \cdot \prod_{\sigma,\,p,\,\ell} \left\langle k^{(p)}_{\ell+1,\sigma} \middle| T_{\ell,\sigma}[\Sigma] \middle| k^{(p)}_{\ell,\sigma} \right\rangle$$

Three-component Gibbs:

$\mathbf{K} \mid \Sigma$ (FFBS) $\longleftrightarrow$ $\Sigma \mid \mathbf{K}$ (exact conditionals) $\longleftrightarrow$ P_σ (permutation moves)

The Transfer Kernel and Convolution by FFT in k-Space

Fermi Hubbard model:

$$H = K + U \sum_{\mathbf{r}} \left(n_{\mathbf{r}\uparrow} - \tfrac{1}{2} \right) \left(n_{\mathbf{r}\downarrow} - \tfrac{1}{2} \right), \qquad K = -t \sum_{\langle \mathbf{r}, \mathbf{r}' \rangle} (c^\dagger_{\mathbf{r}} c_{\mathbf{r}'} + \text{h.c.})$$

After HS decoupling at slice ℓ, $V_\ell(\mathbf{r})$ is the site-diagonal one-body potential from Σ_ℓ.

The Transfer Kernel and Convolution by FFT in k-Space

Fermi Hubbard model:

$$H = K + U \sum_{\mathbf{r}} \left(n_{\mathbf{r}\uparrow} - \tfrac{1}{2}\right)\left(n_{\mathbf{r}\downarrow} - \tfrac{1}{2}\right), \qquad K = -t \sum_{\langle \mathbf{r},\mathbf{r}' \rangle} (c^\dagger_{\mathbf{r}} c_{\mathbf{r}'} + \text{h.c.})$$

After HS decoupling at slice ℓ, $V_\ell(\mathbf{r})$ is the site-diagonal one-body potential from Σ_ℓ.

Single-particle propagator via symmetric Trotter splitting:

$$T_\ell = e^{-\frac{\Delta\tau}{2}K}\; e^{-\Delta\tau V_\ell}\; e^{-\frac{\Delta\tau}{2}K}$$

The Transfer Kernel and Convolution by FFT in k-Space

Fermi Hubbard model:

$$H = K + U \sum_{\mathbf{r}} \left(n_{\mathbf{r}\uparrow} - \tfrac{1}{2}\right)\left(n_{\mathbf{r}\downarrow} - \tfrac{1}{2}\right), \qquad K = -t \sum_{\langle \mathbf{r},\mathbf{r}' \rangle} (c^\dagger_{\mathbf{r}} c_{\mathbf{r}'} + \text{h.c.})$$

After HS decoupling at slice ℓ, $V_\ell(\mathbf{r})$ is the site-diagonal one-body potential from Σ_ℓ.

Single-particle propagator via symmetric Trotter splitting:

$$T_\ell = e^{-\frac{\Delta\tau}{2}K}\, e^{-\Delta\tau V_\ell}\, e^{-\frac{\Delta\tau}{2}K}$$

Insert $\sum_{\mathbf{r}} |\mathbf{r}\rangle \langle \mathbf{r}|$ resolutions of identity twice:

$$\langle k' | T_\ell | k \rangle = \frac{1}{V} \underbrace{e^{-\frac{\Delta\tau}{2}\varepsilon_{k'}}}_{D(k')} \underbrace{\sum_{\mathbf{r}} W_\ell(\mathbf{r})\, e^{i(k-k')\cdot\mathbf{r}}}_{\widehat{W}_\ell(k-k') \leftarrow \text{DFT of } W_\ell} \underbrace{e^{-\frac{\Delta\tau}{2}\varepsilon_k}}_{D(k)}$$

$$\varepsilon_k = -2t \textstyle\sum_\alpha \cos k_\alpha, \quad W_\ell(\mathbf{r}) = e^{-\Delta\tau V_\ell(\mathbf{r})}.$$

The Transfer Kernel and Convolution by FFT in k-Space

Fermi Hubbard model:

$$H = K + U \sum_{\mathbf{r}} \left(n_{\mathbf{r}\uparrow} - \tfrac{1}{2}\right)\left(n_{\mathbf{r}\downarrow} - \tfrac{1}{2}\right), \qquad K = -t \sum_{\langle \mathbf{r},\mathbf{r}' \rangle} (c_{\mathbf{r}}^{\dagger} c_{\mathbf{r}'} + \text{h.c.})$$

After HS decoupling at slice ℓ, $V_\ell(\mathbf{r})$ is the site-diagonal one-body potential from Σ_ℓ.

Single-particle propagator via symmetric Trotter splitting:

$$T_\ell = e^{-\frac{\Delta\tau}{2} K}\, e^{-\Delta\tau V_\ell}\, e^{-\frac{\Delta\tau}{2} K}$$

Insert $\sum_{\mathbf{r}} |\mathbf{r}\rangle \langle \mathbf{r}|$ resolutions of identity twice:

$$\left\langle k' \middle| T_\ell \middle| k \right\rangle = \frac{1}{V} \underbrace{e^{-\frac{\Delta\tau}{2}\varepsilon_{k'}}}_{D(k')} \underbrace{\sum_{\mathbf{r}} W_\ell(\mathbf{r})\, e^{i(k-k')\cdot\mathbf{r}}}_{\widehat{W}_\ell(k-k') \leftarrow \text{DFT of } W_\ell} \underbrace{e^{-\frac{\Delta\tau}{2}\varepsilon_k}}_{D(k)}$$

$\varepsilon_k = -2t \sum_\alpha \cos k_\alpha, \quad W_\ell(\mathbf{r}) = e^{-\Delta\tau V_\ell(\mathbf{r})}.$

Dependence on $k - k'$ $\Longrightarrow$ **convolution** $\Longrightarrow$ **FFT in** $O(N \log N)$

Two Core MCMC Moves for Worldlines: Path and Permutation

FFBS: Resample shape **Permutation: Resample topology**

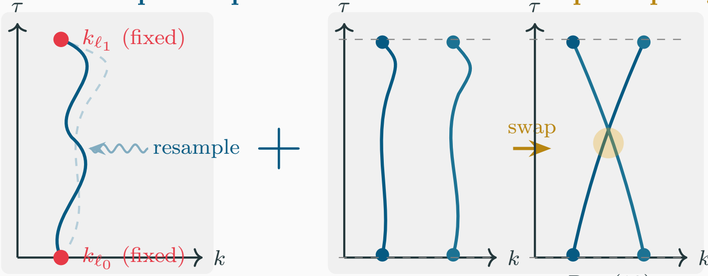

Continuous Path

$P = \mathrm{id}$ $P = (12)$

Discrete Sampling of Boundary Connection

Both are essential for ergodicity. FFBS explores continuous momenta along each worldline. Permutation moves explore how worldlines connect at $\tau = \beta$, encoding $\mathrm{sgn}(P)$.

Block Update for Worldlines by Forward Filtering Backward Sampling

Forward Filtering Backward Sampling (FFBS)

- Fix k_{ℓ_0}, k_{ℓ_1} and update the whole interior block.

Block Update for Worldlines by Forward Filtering Backward Sampling

Forward Filtering Backward Sampling (FFBS)

- Fix k_{ℓ_0}, k_{ℓ_1} and update the whole interior block.
- **Forward Filtering for target conditional**:

$$p(k_{\ell_0+1:\ell_1-1} \mid k_{\ell_0}, k_{\ell_1}, \Sigma) \propto \prod_{\ell=\ell_0}^{\ell_1-1} T_\ell(k_{\ell+1} \mid k_\ell; \Sigma)$$

  stores weights of all partial paths (exact). Sample interior paths with probability proportional to the weight.

Block Update for Worldlines by Forward Filtering Backward Sampling

Forward Filtering Backward Sampling (FFBS)

- Fix k_{ℓ_0}, k_{ℓ_1} and update the whole interior block.
- **Forward Filtering for target conditional**:

$$p(k_{\ell_0+1:\ell_1-1} \mid k_{\ell_0}, k_{\ell_1}, \Sigma) \propto \prod_{\ell=\ell_0}^{\ell_1-1} T_\ell(k_{\ell+1} \mid k_\ell; \Sigma)$$

  stores weights of all partial paths (exact). Sample interior paths with probability proportional to the weight.
- **Backward Sampling the tractable Block distribution. No extra Metropolis-Hasting accept/reject.**

Block Update for Worldlines by Forward Filtering Backward Sampling

Forward Filtering Backward Sampling (FFBS)

- Fix k_{ℓ_0}, k_{ℓ_1} and update the whole interior block.
- **Forward Filtering for target conditional**:

$$p(k_{\ell_0+1:\ell_1-1} \mid k_{\ell_0}, k_{\ell_1}, \Sigma) \propto \prod_{\ell=\ell_0}^{\ell_1-1} T_\ell(k_{\ell+1} \mid k_\ell; \Sigma)$$

  stores weights of all partial paths (exact). Sample interior paths with probability proportional to the weight.
- **Backward Sampling the tractable Block distribution. No extra Metropolis-Hasting accept/reject.**

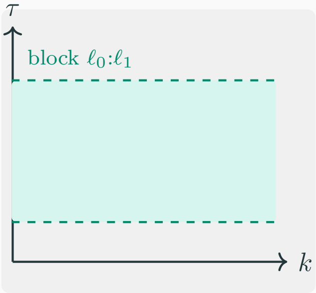

FFBS refs: Carter & Kohn (1994); Fruhwirth-Schnatter (1994); Doucet–Godsill–Andrieu (2000).

Block Update for Worldlines by Forward Filtering Backward Sampling

Forward Filtering Backward Sampling (FFBS)

- Fix k_{ℓ_0}, k_{ℓ_1} and update the whole interior block.
- **Forward Filtering for target conditional**:

$$p(k_{\ell_0+1:\ell_1-1} \mid k_{\ell_0}, k_{\ell_1}, \Sigma) \propto \prod_{\ell=\ell_0}^{\ell_1-1} T_\ell(k_{\ell+1} \mid k_\ell; \Sigma)$$

  stores weights of all partial paths (exact). Sample interior paths with probability proportional to the weight.

- **Backward Sampling the tractable Block distribution. No extra Metropolis-Hasting accept/reject.**

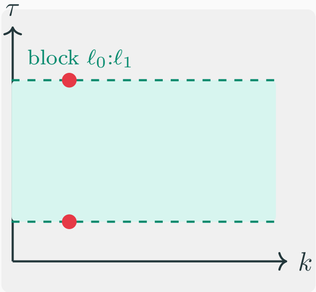

FFBS refs: Carter & Kohn (1994); Fruhwirth-Schnatter (1994); Doucet–Godsill–Andrieu (2000).

Block Update for Worldlines by Forward Filtering Backward Sampling

Forward Filtering Backward Sampling (FFBS)

- Fix k_{ℓ_0}, k_{ℓ_1} and update the whole interior block.
- **Forward Filtering for target conditional**:

$$p(k_{\ell_0+1:\ell_1-1} \mid k_{\ell_0}, k_{\ell_1}, \Sigma) \propto \prod_{\ell=\ell_0}^{\ell_1-1} T_\ell(k_{\ell+1} \mid k_\ell; \Sigma)$$

  stores weights of all partial paths (exact). Sample interior paths with probability proportional to the weight.
- **Backward Sampling the tractable Block distribution. No extra Metropolis-Hasting accept/reject.**

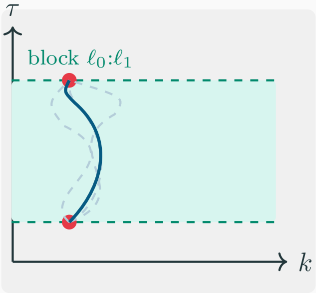

FFBS refs: Carter & Kohn (1994); Fruhwirth-Schnatter (1994); Doucet–Godsill–Andrieu (2000).

Block Update for Worldlines by Forward Filtering Backward Sampling

Forward Filtering Backward Sampling (FFBS)

- Fix k_{ℓ_0}, k_{ℓ_1} and update the whole interior block.
- **Forward Filtering for target conditional**:

$$p(k_{\ell_0+1:\ell_1-1} \mid k_{\ell_0}, k_{\ell_1}, \Sigma) \propto \prod_{\ell=\ell_0}^{\ell_1-1} T_\ell(k_{\ell+1} \mid k_\ell; \Sigma)$$

  stores weights of all partial paths (exact). Sample interior paths with probability proportional to the weight.
- **Backward Sampling the tractable Block distribution. No extra Metropolis-Hasting accept/reject.**

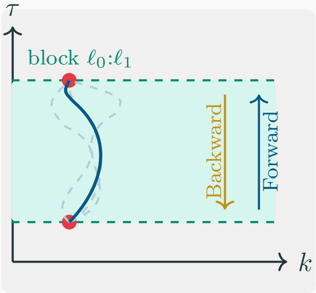

FFBS refs: Carter & Kohn (1994); Fruhwirth-Schnatter (1994); Doucet–Godsill–Andrieu (2000).

Block Update for Worldlines by Forward Filtering Backward Sampling

Forward Filtering Backward Sampling (FFBS)

- Fix k_{ℓ_0}, k_{ℓ_1} and update the whole interior block.
- **Forward Filtering for target conditional**:

$$p(k_{\ell_0+1:\ell_1-1} \mid k_{\ell_0}, k_{\ell_1}, \Sigma) \propto \prod_{\ell=\ell_0}^{\ell_1-1} T_\ell(k_{\ell+1} \mid k_\ell; \Sigma)$$

  stores weights of all partial paths (exact). Sample interior paths with probability proportional to the weight.
- **Backward Sampling the tractable Block distribution. No extra Metropolis-Hasting accept/reject.**

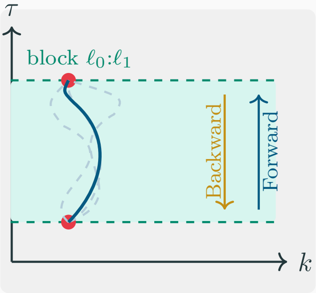

Forward: total weight of all partial paths to slice ℓ ending at k. **Backward**: sample full interior paths from those weights.

FFBS refs: Carter & Kohn (1994); Fruhwirth-Schnatter (1994); Doucet–Godsill–Andrieu (2000).

Forward Filtering Backward Sampling (FFBS): Block Sampling of Worldlines

Goal: Given fixed endpoints k_{ℓ_0}, k_{ℓ_1}, resample $\{k_{\ell_0+1}, \ldots, k_{\ell_1-1}\}$ from the *exact* conditional.

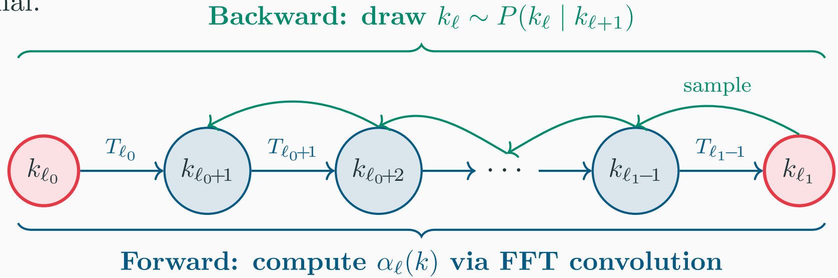

1. **Forward:** Propagate $\alpha_{\ell+1}(k) = D(k) \sum_{k'} \frac{\widehat{W}_\ell(k-k')}{V} D(k')\, \alpha_\ell(k')$
   $[\text{convolution} \Rightarrow \text{FFT} \sim O(N \log N)]$

Forward Filtering Backward Sampling (FFBS): Block Sampling of Worldlines

Goal: Given fixed endpoints k_{ℓ_0}, k_{ℓ_1}, resample $\{k_{\ell_0+1}, \ldots, k_{\ell_1-1}\}$ from the *exact* conditional.

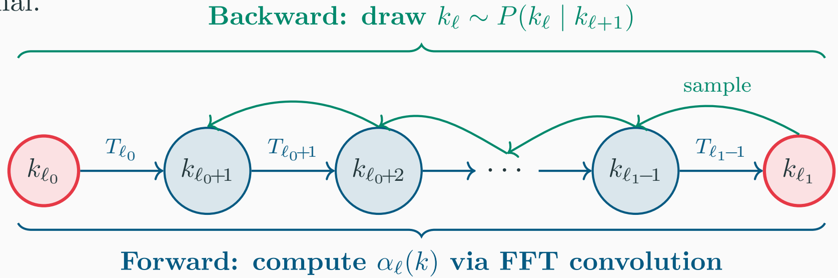

1. **Forward:** Propagate $\alpha_{\ell+1}(k) = D(k) \sum_{k'} \frac{\widehat{W}_\ell(k-k')}{V} D(k')\, \alpha_\ell(k')$
   [convolution $\Rightarrow$ FFT $\sim O(N \log N)$]
2. **Backward:** Draw k_ℓ from $P(k_\ell \mid k_{\ell+1}) \propto \alpha_\ell(k_\ell) \cdot \langle k_{\ell+1} | T_\ell | k_\ell \rangle$

This is an **exact Gibbs step**: draws from the true conditional $P(\text{block} \mid \text{endpoints}, \Sigma)$.

FFBS Block Placement: Scramble End Points

FFBS conditions on block endpoints $\Rightarrow$ endpoints not resampled within that block.

FFBS Block Placement: Scramble End Points

FFBS conditions on block endpoints $\Rightarrow$ endpoints not resampled within that block.

Requirement: Every time slice must appear as an interior point in at least one block.

FFBS Block Placement: Scramble End Points

FFBS conditions on block endpoints ⇒ endpoints not resampled within that block.

Requirement: Every time slice must appear as an interior point in at least one block. **Simplest solution:** Two passes offset by $B/2$.

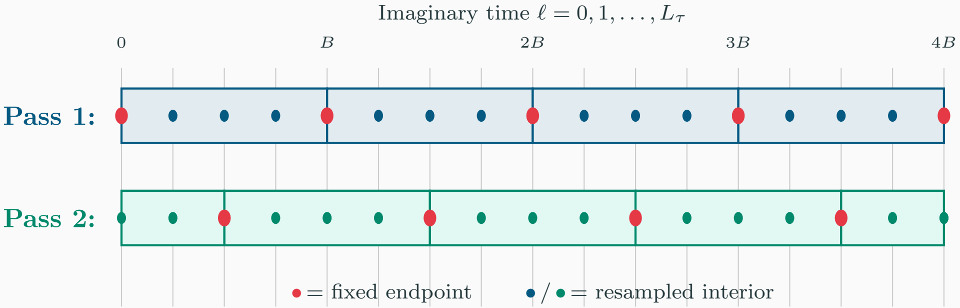

Any scheme satisfying the coverage requirement works (random blocks, sliding windows, etc.).

Permutation Updates for Fermionic Statistics

FFBS resamples momenta along *each* worldline, but never changes *which worldline connects to which* at $\tau = \beta$. The fermionic partition function requires summing over permutations weighted by $\mathrm{sgn}(P)$.

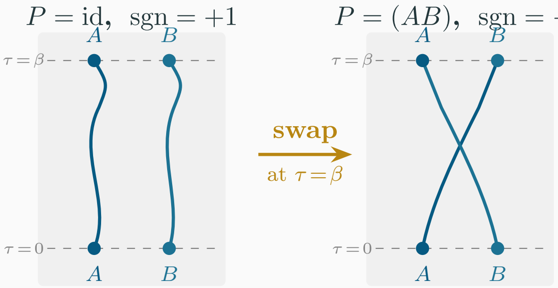

Each closes on itself

Worldlines exchange endpoints

Note:

FFBS *never* reconnects worldline endpoints.

Without permutation moves:
$\Rightarrow$ stuck in identity sector
$\Rightarrow$ miss all fermionic physics

Implementation:

Propose pair swap at $\tau = \beta$.
Accept/reject by Metropolis.

Auxiliary Field Updates: Cut-Invariant Form

Note: sampled k-space worldlines provide only $n_k(\ell)$ (momentum-basis diagonal), not directly $n(\mathbf{r}, \ell)$, which depends on off-diagonal elements in k-space. So we use the the following form to get site-wise HS conditionals.

Auxiliary Field Updates: Cut-Invariant Form

Note: sampled k-space worldlines provide only $n_k(\ell)$ (momentum-basis diagonal), not directly $n(\mathbf{r},\ell)$, which depends on off-diagonal elements in k-space. So we use the the following form to get site-wise HS conditionals.

$$Z_\sigma(\Sigma) = \mathrm{Tr}\left[\prod_{j=0}^{L_\tau-1} T_{\sigma,j}\right] = \mathrm{Tr}\left[\underbrace{T_{\sigma,L_\tau-1}\cdots T_{\sigma,\ell+1}}_{\hat{\beta}_{\sigma,\ell+1}} \, {\color{red}T_{\sigma,\ell}} \, \underbrace{T_{\sigma,\ell-1}\cdots T_{\sigma,0}}_{\hat{\alpha}_{\sigma,\ell}}\right] = \sum_{\mathbf{r}} C_{\sigma,\ell}(\mathbf{r}){\color{red}e^{\sigma\lambda s_{\mathbf{r}\ell}}}$$

Auxiliary Field Updates: Cut-Invariant Form

Note: sampled k-space worldlines provide only $n_k(\ell)$ (momentum-basis diagonal), not directly $n(\mathbf{r},\ell)$, which depends on off-diagonal elements in k-space. So we use the the following form to get site-wise HS conditionals.

$$Z_\sigma(\Sigma) = \mathrm{Tr}\left[\prod_{j=0}^{L_\tau-1} T_{\sigma,j}\right] = \mathrm{Tr}\left[\underbrace{T_{\sigma,L_\tau-1}\cdots T_{\sigma,\ell+1}}_{\hat{\beta}_{\sigma,\ell+1}}\, {\color{red}T_{\sigma,\ell}}\, \underbrace{T_{\sigma,\ell-1}\cdots T_{\sigma,0}}_{\hat{\alpha}_{\sigma,\ell}}\right] = \sum_{\mathbf{r}} C_{\sigma,\ell}(\mathbf{r}) {\color{red}e^{\sigma\lambda s_{\mathbf{r}\ell}}}$$

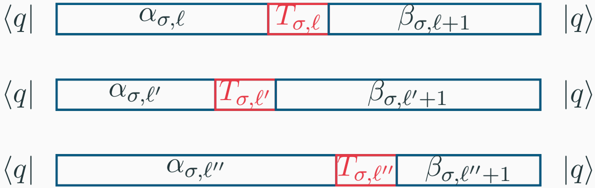

Cut invariance: move the red slice factor to any time slice.

Auxiliary Field Updates: Cut-Invariant Form

Note: sampled k-space worldlines provide only $n_k(\ell)$ (momentum-basis diagonal), not directly $n(\mathbf{r},\ell)$, which depends on off-diagonal elements in k-space. So we use the the following form to get site-wise HS conditionals.

$$Z_\sigma(\Sigma) = \mathrm{Tr}\left[\prod_{j=0}^{L_\tau-1} T_{\sigma,j}\right] = \mathrm{Tr}\left[\underbrace{T_{\sigma,L_\tau-1}\cdots T_{\sigma,\ell+1}}_{\hat{\beta}_{\sigma,\ell+1}} \; {\color{red}T_{\sigma,\ell}} \; \underbrace{T_{\sigma,\ell-1}\cdots T_{\sigma,0}}_{\hat{\alpha}_{\sigma,\ell}}\right] = \sum_{\mathbf{r}} C_{\sigma,\ell}(\mathbf{r}) {\color{red}e^{\sigma\lambda s_{\mathbf{r}\ell}}}$$

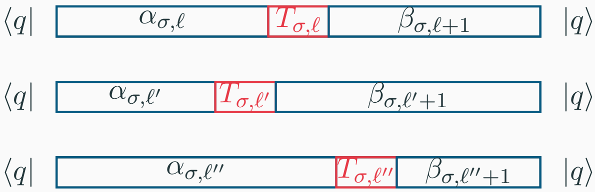

Cut invariance: move the red slice factor to any time slice.

Real space conditional

$$P(s_{\mathbf{r}\ell} = s \mid \mathrm{rest}) = \frac{w_s}{w_{+1} + w_{-1}}$$

$$w_s = \prod_{\sigma}\left[\Gamma_{\sigma,\ell} + C_{\sigma,\ell}(\mathbf{r})e^{\sigma\lambda s}\right]$$

$$\Gamma_{\sigma,\ell} = \sum_{\mathbf{x}\neq\mathbf{r}} C_{\sigma,\ell}(\mathbf{x})e^{\sigma\lambda s_{\mathbf{x}\ell}}$$

Algorithm Summary

1. Precompute potential weights $\widehat{W}_{\sigma,\ell}(q) = \sum_{\mathbf{r}} W_{\sigma,\ell}(\mathbf{r})\, e^{iq\cdot\mathbf{r}}$ $O(\beta \cdot N \log N)$

Total per-sweep complexity: $O(\beta N \log N)$

Each step satisfies detailed balance $\Rightarrow$ composition preserves $\pi(\mathbf{K}, \Sigma, P)$.

Algorithm Summary

1. Precompute potential weights $\widehat{W}_{\sigma,\ell}(q) = \sum_{\mathbf{r}} W_{\sigma,\ell}(\mathbf{r})\, e^{iq\cdot\mathbf{r}}$ $O(\beta \cdot N \log N)$

2. FFBS worldline update: resample worldlines via FFT $O(\beta \cdot N \log N)$

Particles independent given $\Sigma \Rightarrow$ fully parallelizable across particles and spins.

Total per-sweep complexity: $O(\beta N \log N)$

Each step satisfies detailed balance $\Rightarrow$ composition preserves $\pi(\mathbf{K}, \Sigma, P)$.

Algorithm Summary

1. Precompute potential weights $\widehat{W}_{\sigma,\ell}(q) = \sum_{\mathbf{r}} W_{\sigma,\ell}(\mathbf{r})\, e^{iq\cdot\mathbf{r}}$ $O(\beta \cdot N \log N)$

2. FFBS worldline update: resample worldlines via FFT $O(\beta \cdot N \log N)$

Particles independent given $\Sigma \Rightarrow$ fully parallelizable across particles and spins.

3. Permutation updates: propose pair swaps at $\tau = \beta$ $O(N)$

Total per-sweep complexity: $O(\beta N \log N)$

Each step satisfies detailed balance $\Rightarrow$ composition preserves $\pi(\mathbf{K}, \Sigma, P)$.

Algorithm Summary

1. Precompute potential weights $\widehat{W}_{\sigma,\ell}(q) = \sum_{\mathbf{r}} W_{\sigma,\ell}(\mathbf{r})\, e^{iq\cdot\mathbf{r}}$ $O(\beta \cdot N \log N)$

2. FFBS worldline update: resample worldlines via FFT $O(\beta \cdot N \log N)$

Particles independent given $\Sigma \Rightarrow$ fully parallelizable across particles and spins.

3. Permutation updates: propose pair swaps at $\tau = \beta$ $O(N)$

4. Auxiliary field Gibbs: sweep $s_{\mathbf{r}\ell}$ via cut-invariant conditionals $O(\beta N \log N)$

Exact Bernoulli sampling at each spacetime point; recompute $C_{\sigma,\ell}$ between slices.

Total per-sweep complexity: $O(\beta N \log N)$

Each step satisfies detailed balance $\Rightarrow$ composition preserves $\pi(\mathbf{K}, \Sigma, P)$.

Results: Weak Coupling ($U/t = 1$, $\beta t = 32$, 32×32)

Charge: Sharp Fermi surface in $\langle n_{\mathbf{k}} \rangle$, discontinuity at $\varepsilon(\mathbf{k}) = 0$. **Spin:** Staggered AF $\langle S_0^z S_{\mathbf{r}}^z \rangle$; $S(\mathbf{k})$ peaks at $\mathbf{Q} = (\pi, \pi)$, grows with L.

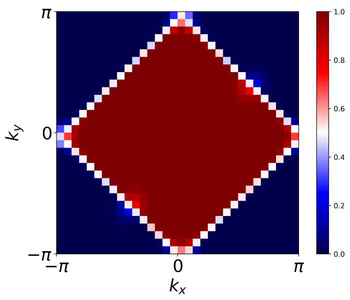

Momentum distribution $n_{\mathbf{k}}$

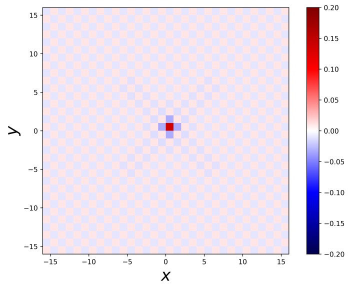

Real-space $C(\mathbf{r}) = \langle S_0^z S_{\mathbf{r}}^z \rangle$

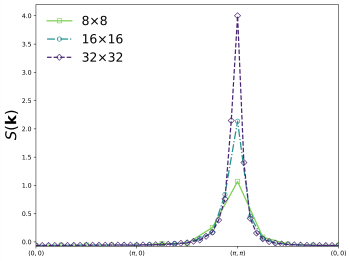

Spin structure factor $S(\mathbf{k})$

All results match DQMC benchmarks (Hirsch 1985, Tomas et al. 2012).

Results: Strong Coupling ($U/t = 20$, $\beta t = 32$, 32×32)

Mott insulator: No sharp FS, $n_{\mathbf{k}} \to 1$ as $U/t \to \infty$, double occupancy suppressed. **Heisenberg magnet:** $J_{\text{eff}} = 4t^2/U$; staggered $\langle S_0^z S_{\mathbf{r}}^z \rangle$ with exponential decay; $S(\mathbf{k})$ AF peak at (π, π) grows with L.

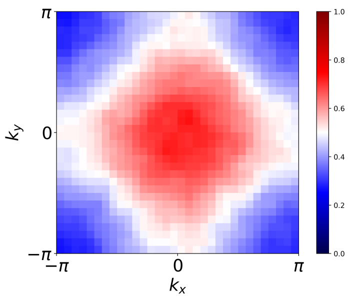

Momentum distribution $n_{\mathbf{k}}$

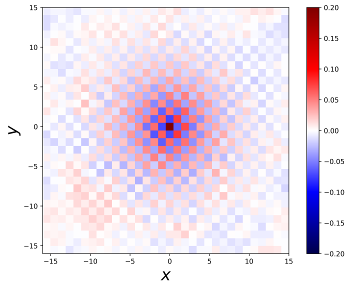

Real-space $C(\mathbf{r}) = \langle S_0^z S_{\mathbf{r}}^z \rangle$

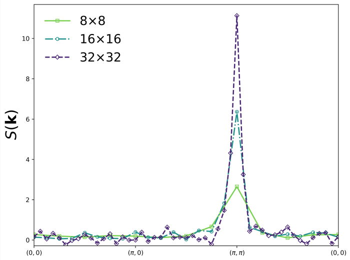

Spin structure factor $S(\mathbf{k})$

Both limits recovered: Fermi liquid $\xrightarrow{U/t}$ Mott insulator. Matches White et al. 1989, Moreo et al. 1990.

Method Comparison

	Fast-Update DQMC	**FFT-FFBS**
Fermion treatment	Integrate out (det)	Sample worldlines
Per-sweep cost	$O(\beta N^3)$	$O(\beta N \log N)$
Block updates	Expensive	FFBS via FFT
Auxiliary HS field	Metropolis	Exact Gibbs
Parallelization	Limited	Easy GPU Implementation
Long-range interaction	Increases to $O(\beta N^4)$	Remains $O(\beta N \log N)$

The key advantage: replacing N^3 per sweep with $N \log N$ convolutions.

D. Kong, *et al.*, arXiv:2510.13866

Summary

Determinant-free QMC with $O(N \log N)$ scaling

Key ingredients:

1. **Joint sampling:** Keep k-space fermion worldlines as explicit degrees of freedom alongside auxiliary fields — avoid the $O(N^2) \sim O(N^3)$ determinant update entirely
2. **FFT-accelerated FFBS:** The k-space transfer kernel is a convolution, enabling exact block updates of worldlines in $O(N \log N)$. Two offset sweeps ensure full ergodicity along each worldline.

Thank you!

Acknowledgment

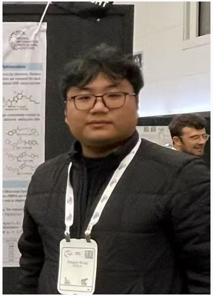

Deqian Kong

UCLA

Jianwen Xie

Lambda, Inc.

Ying Nian Wu

UCLA