Parallel software for lattice $\mathcal{N} = 4$ supersymmetric Yang–Mills theory

David Schaich^a, Thomas DeGrand^b

^a*Department of Physics, Syracuse University, Syracuse, New York 13244, United States*

^b*Department of Physics, University of Colorado, Boulder, Colorado 80309, United States*

Abstract

We present new parallel software, `SUSY LATTICE`, for lattice studies of four-dimensional $\mathcal{N} = 4$ supersymmetric Yang–Mills theory with gauge group $SU(N)$. The lattice action is constructed to exactly preserve a single supersymmetry charge at non-zero lattice spacing, up to additional potential terms included to stabilize numerical simulations. The software evolved from the MILC code for lattice QCD, and retains a similar large-scale framework despite the different target theory. Many routines are adapted from an existing serial code [1], which `SUSY LATTICE` supersedes. This paper provides an overview of the new parallel software, summarizing the lattice system, describing the applications that are currently provided and explaining their basic workflow for non-experts in lattice gauge theory. We discuss the parallel performance of the code, and highlight some notable aspects of the documentation for those interested in contributing to its future development.

Keywords: Lattice gauge theory, Supersymmetric Yang–Mills, Monte Carlo methods, Parallel computing

PACS: 11.15.Ha, 12.60.Jv, 02.70.Uu

1. Introduction: Supersymmetry on the lattice

For decades, lattice field theory has played important roles in many areas of physics, from statistical and condensed-matter physics to high-energy particle physics. Within the domain of particle physics, the formulation of gauge theories on a space-time lattice provides a non-perturbative definition of such systems, with large-scale numerical computations making crucial contributions to the phenomenology of quantum chromodynamics (QCD) and the search for new physics beyond the standard model (BSM).¹

An area where progress has been more difficult is the lattice formulation of supersymmetric theories [2]. This is unfortunate, because supersymmetry is an extremely important theoretical tool, and a common ingredient in BSM physics. The simplifications resulting from supersymmetry have aided analytic investigations of many fascinating non-perturbative phenomena such as confinement, dynamical symmetry breaking, and dualities between pairs of supersymmetric gauge theories as well as between gauge and gravity theories [3]. Complementary studies of these issues through first-principles lattice calculations, in addition to potential phenomenological applications to supersymmetric extensions of the standard model, motivate continuing research into supersymmetric lattice gauge theory.

Briefly stated, the essential obstacle to lattice supersymmetry is the fact that supersymmetry transformations extend the algebra of space-time translations and rotations. The anticommutators of supersymmetry generators produce infinitesimal translations, $\{Q_\alpha, Q_\alpha^\dagger\} \propto \sigma_{\alpha\dot{\alpha}} \cdot P$, which do not exist in discrete space-time. This problem cannot be evaded by defining a “discrete supersymmetry” based on finite lattice translations, because lattice finite-difference operators do not satisfy the Leibniz rule [4, 5]. Naive lattice

Email addresses: `dschaich@syr.edu` (David Schaich), `thomas.degrand@colorado.edu` (Thomas DeGrand)

¹The proceedings of the annual International Symposium on Lattice Field Theory provide comprehensive reviews and may be found at <http://pos.sissa.it/cgi-bin/reader/family.cgi?code=lattice>.

discretization thus breaks supersymmetry at the classical level, leading to the generation of supersymmetry-violating operators whose couplings must be fine-tuned to recover the supersymmetric theory of interest. Since there are typically many such operators, such fine-tuning is generally impractical.²

For certain gauge theories with extended supersymmetry, however, significant progress was made in the past decade through the construction of lattice formulations in which a subset of the supersymmetry algebra is exactly preserved at non-zero lattice spacing.³ The historical development and essential features of these constructions are thoroughly reviewed in Ref. [2]. The systems for which this approach is possible include $\mathcal{N} = (2, 2)$, $(4, 4)$ and $(8, 8)$ supersymmetric Yang–Mills (SYM) theories in two dimensions, $\mathcal{N} = 4$ SYM in three dimensions, and the especially interesting case of $\mathcal{N} = 4$ SYM in four dimensions. For all of these theories, the method of topological twisting provides a change of variables under which the supercharges form Kähler–Dirac multiplets of antisymmetric p -index tensors $(Q, Q_\mu, Q_{\mu\nu}, \dots)$ with $p = 0, 1, \dots, d$ in d dimensions. The nilpotent 0-form supercharges Q anticommute to form a sub-algebra that can be exactly preserved on the lattice [17, 18]. This procedure does not depend on the gauge group, and in this article we consider $SU(N)$ gauge theories with relatively small numbers of colors $N = 2, 3$ and 4 .

More recently, these constructions have been used as the basis for practical numerical lattice calculations, which have passed several non-trivial consistency checks and begun to make novel contributions to our physical understanding [19, 20, 21, 22]. Especially in the most interesting case of four-dimensional $\mathcal{N} = 4$ SYM, these lattice studies rapidly become computationally expensive, a challenge familiar from lattice QCD. To address this challenge, we have developed new parallel software for lattice $\mathcal{N} = 4$ SYM, **SUSY LATTICE**, which we introduce in this article. Our starting point was the highly-optimized MIMD Lattice Computation (MILC) software for lattice QCD [23], although the significantly different target theory required substantial changes throughout the code. Many central routines are adapted from the serial code presented in Ref. [1], which this software supersedes, and we have also added some new capabilities that extend the scope of the supported computations.

In this work we present an overview of **SUSY LATTICE**, aiming to provide useful information both for particle theorists with limited Monte Carlo experience as well as for lattice gauge theorists more familiar with QCD-like systems. We begin in the next section by briefly reviewing the lattice formulation of $\mathcal{N} = 4$ SYM, and the basic workflow of lattice investigations. Summarizing discussions recently published in Ref. [21], we describe the degrees of freedom of the discretized system, the A_4^* lattice structure, and the lattice action implemented by **SUSY LATTICE**, including potential terms that softly break the otherwise exactly preserved 0-form supersymmetry Q . In Section 2.2 we explain the rational hybrid Monte Carlo (RHMC) algorithm for importance sampling of gauge field configurations and introduce a few important measurements, in particular the phase of the complex Pfaffian of the fermion operator. For users interested in obtaining and running **SUSY LATTICE** without modification, in Section 3 we provide a quick-start guide including a high-level overview of the available measurements, along with some discussion of the scaling and parallel performance of the code. Finally, in Section 4 we highlight some notable aspects of the documentation for those who may be interested in extending **SUSY LATTICE** with additional measurements or features. We conclude with a discussion of some potential directions for future development, and invite all who wish to contribute to do so through the publicly accessible version control repository [49].

2. The discretized theory, and methodology of lattice calculations

In this section we first write down the lattice formulation of $\mathcal{N} = 4$ SYM in four dimensions, and then review the essential ingredients of lattice gauge theory calculations. We omit the details of the relevant topological twisting of continuum $\mathcal{N} = 4$ SYM [24, 25], which is discussed in more detail in Refs. [2, 21]. Instead, in Section 2.1 we simply introduce the degrees of freedom of the lattice theory, describe the A_4^*

²One notable exception is $\mathcal{N} = 1$ supersymmetric Yang–Mills theory in four dimensions, where only the gaugino mass needs to be fine-tuned [6], or controlled by working with Ginsparg–Wilson lattice fermions [7, 8, 9].

³There are, of course, other approaches that lie beyond the scope of this article [10, 11, 12, 13, 14, 15, 16].

lattice structure, and write down the lattice action. We explain the various terms in the lattice action, emphasizing their invariance under lattice gauge transformations and (except for certain potential terms) preservation of the 0-form twisted supersymmetry $\mathcal{Q}$.

Given the lattice action S , we can carry out numerical computations of operator expectation values

$$\langle \mathcal{O} \rangle = \frac{1}{\mathcal{Z}} \int [dX] \mathcal{O} e^{-S[X]} \quad \mathcal{Z} = \int [dX] e^{-S[X]}, \quad (1)$$

where X is a placeholder for the lattice fields to be defined below. Section 2.2 summarizes how the discretized path integral is stochastically evaluated through importance-sampling Monte Carlo. In broadest terms, we use the rational hybrid Monte Carlo (RHMC) algorithm to obtain a sequence of gauge field configurations with the appropriate probability distribution. We then measure observables of interest on a set of n such configurations, so that $\langle \mathcal{O} \rangle = \frac{1}{n} \sum_i \mathcal{O}_i$, with statistical uncertainties from the finite sample size and systematic uncertainties from working in a finite discretized space-time. In Section 2.2 we will discuss some important observables, including the low-lying eigenvalues and Pfaffian of the fermion operator.

2.1. Lattice variables, lattice structure and lattice action

Continuum $\mathcal{N} = 4$ SYM is a theory of a gauge field, four Majorana fermion fields and six scalar fields, all of which transform in the adjoint representation of the $SU(N)$ gauge group, and are related to each other by 16 fermionic supersymmetry charges. Working in euclidean space-time, the twisting procedure regroups the ten bosons into a five-component complexified gauge field $\mathcal{A}_a$ with $a = 0, \dots, 4$, while the 16 fermions are assigned to the multiplet $(\eta, \psi_a, \chi_{ab})$, with χ_{ab} antisymmetric. The supersymmetry charges are similarly converted to $(\mathcal{Q}, \mathcal{Q}_a, \mathcal{Q}_{ab})$.

At this point we can move onto the lattice [17, 18, 26, 27, 21], by defining complexified gauge links $\mathcal{U}_a(n)$ that are elements of the algebra $\mathfrak{gl}(N, \mathbb{C})$,

$$\mathcal{U}_a(n) = \sum_{C=0}^{N^2-1} T^C \mathcal{U}_a^C(n), \quad (2)$$

with complex coefficients $\mathcal{U}_a^C(n)$. The T^C are antihermitian generators of $\mathfrak{u}(N)$, with normalization $\text{Tr}[T^A T^B] = -\delta^{AB}$. The fermion fields $\eta(n)$, $\psi_a(n)$ and $\chi_{ab}(n)$ are defined in the same way, as required by supersymmetry. As a consequence, the fields transform in the adjoint representation of $U(N)$,

$$\begin{aligned} \mathcal{U}_a(n) &\rightarrow G(n) \mathcal{U}_a(n) G^\dagger(n + \hat{\mu}_a) & \eta(n) &\rightarrow G(n) \eta(n) G^\dagger(n) \\ \bar{\mathcal{U}}_a(n) &\rightarrow G(n + \hat{\mu}_a) \bar{\mathcal{U}}_a(n) G^\dagger(n) & \psi_a(n) &\rightarrow G(n) \psi_a(n) G^\dagger(n + \hat{\mu}_a) \\ \chi_{ab}(n) &\rightarrow G(n + \hat{\mu}_a + \hat{\mu}_b) \chi_{ab}(n) G^\dagger(n), \end{aligned} \quad (3)$$

for lattice gauge transformation $G \in U(N)$. Speaking informally, we will say that the link $\mathcal{U}_a(n)$ points “from” site $n + \hat{\mu}_a$ (on the right) “to” site n (on the left), and similarly for the other fields.

Because we have five gauge links going out from each lattice site, we cannot discretize the theory on the familiar hypercubic lattice. Instead we must use the A_4^* lattice, the four-dimensional lattice where each site has ten symmetric nearest neighbors connected by five linearly dependent basis vectors. (The A_d^* lattice in d dimensions has coordination number $2(d+1)$; the most familiar example is the triangular lattice A_2^* , while A_3^* can be deformed into the body-centered cubic lattice.) In `SUSY LATTICE` we employ an extremely convenient representation of the A_4^* lattice, which simply adds a body-diagonal link $\hat{\mu}_4 = (-1, -1, -1, -1)$ to the usual four hypercubic basis vectors $\hat{\mu}_\nu$. Converting from this abstract hypercubic basis to physical space-time requires some care, as discussed in Ref. [21].

Now we can write down the full lattice action (repeated indices summed except where noted),

$$S = S_{exact} + S_{closed} + S_{stab} \quad (4)$$

$$S_{exact} = \frac{N}{2\lambda} \sum_n a^4 \text{Tr} \left[-\bar{\mathcal{F}}_{ab}(n) \mathcal{F}_{ab}(n) + \frac{c_2}{2} \left(\bar{\mathcal{D}}_a^{(-)} \mathcal{U}_a(n) \right)^2 - \chi_{ab}(n) \mathcal{D}_{[a}^{(+)} \psi_{b]}(n) - \eta(n) \bar{\mathcal{D}}_a^{(-)} \psi_a(n) \right] \quad (5)$$

$$S_{closed} = -\frac{N}{8\lambda} \sum_n a^4 \text{Tr} \left[\epsilon_{abcde} \chi_{de}(n + \hat{\mu}_a + \hat{\mu}_b + \hat{\mu}_c) \bar{\mathcal{D}}_c^{(-)} \chi_{ab}(n) \right] \quad (6)$$

$$S_{stab} = \frac{N}{2\lambda} \mu^2 \sum_{n, c} a^4 \left(\frac{1}{N} \text{Tr} [\mathcal{U}_c(n) \bar{\mathcal{U}}_c(n)] - 1 \right)^2 + \kappa \sum_{\mathcal{P}} a^4 |\det \mathcal{P} - 1|^2. \quad (7)$$

These require some further definitions. The field strengths $\mathcal{F}_{ab}$ and $\bar{\mathcal{F}}_{ab}$ in the first term are

$$\begin{aligned} \mathcal{F}_{ab}(n) &= \mathcal{U}_a(n) \mathcal{U}_b(n + \hat{\mu}_a) - \mathcal{U}_b(n) \mathcal{U}_a(n + \hat{\mu}_b) = \mathcal{D}_a^{(+)} \mathcal{U}_b(n) \\ \bar{\mathcal{F}}_{ab}(n) &= \bar{\mathcal{U}}_a(n + \hat{\mu}_b) \bar{\mathcal{U}}_b(n) - \bar{\mathcal{U}}_b(n + \hat{\mu}_a) \bar{\mathcal{U}}_a(n) = \bar{\mathcal{D}}_a^{(+)} \bar{\mathcal{U}}_b(n), \end{aligned} \quad (8)$$

while the forward and backward finite-difference operators are [26, 27]

$$\begin{aligned} \mathcal{D}_a^{(+)} f_b(n) &= \mathcal{U}_a(n) f_b(n + \hat{\mu}_a) - f_b(n) \mathcal{U}_a(n + \hat{\mu}_b) \\ \bar{\mathcal{D}}_a^{(-)} f_a(n) &= f_a(n) \bar{\mathcal{U}}_a(n) - \bar{\mathcal{U}}_a(n - \hat{\mu}_a) f_a(n - \hat{\mu}_a) \\ \bar{\mathcal{D}}_a^{(+)} f_b(n) &= \bar{\mathcal{U}}_a(n + \hat{\mu}_b) f_b(n) - f_b(n + \hat{\mu}_a) \bar{\mathcal{U}}_a(n) \\ \bar{\mathcal{D}}_c^{(-)} f_{ab}(n) &= f_{ab}(n + \hat{\mu}_c) \bar{\mathcal{U}}_c(n) - \bar{\mathcal{U}}_c(n + \hat{\mu}_a + \hat{\mu}_b) f_{ab}(n). \end{aligned} \quad (9)$$

In the final term of Eq. 7, $\mathcal{P}$ is the oriented product of links around a fundamental plaquette of the lattice. For the remainder of this article we will work in lattice units where we set the lattice spacing $a = 1$.

With these definitions it is straightforward to confirm that the lattice action S is gauge invariant: all terms in Eq. 4 form closed loops, and $G^\dagger G = G G^\dagger = 1$ for gauge transformation $G \in \text{U}(N)$. With the exception of S_{stab} (discussed below), S also preserves the 0-form supersymmetry $\mathcal{Q}$, so that $\mathcal{Q}S = 0$ when $\mu = 0$ and $\kappa = 0$. Under the action of $\mathcal{Q}$ the fields transform as

$$\begin{aligned} \mathcal{Q} \mathcal{U}_a(n) &= \psi_a(n) & \mathcal{Q} \psi_a(n) &= 0 \\ \mathcal{Q} \chi_{ab}(n) &= -\bar{\mathcal{F}}_{ab}(n) & \mathcal{Q} \bar{\mathcal{U}}_a(n) &= 0 \\ \mathcal{Q} \eta(n) &= d(n) & \mathcal{Q} d(n) &= 0 \end{aligned} \quad (10)$$

with d a bosonic auxiliary field that maintains off-shell supersymmetry. In Eq. 4 we replaced d by its equation of motion $d(n) \rightarrow \bar{\mathcal{D}}_a^{(-)} \mathcal{U}_a(n)$. From Eq. 10 it is clear that $\mathcal{Q}^2 = 0$ on every field, and it is not hard to see that $\mathcal{Q}$ acting on S_{exact} vanishes. The action of $\mathcal{Q}$ on S_{closed} also vanishes due to a lattice Bianchi identity, $\epsilon_{abcde} \bar{\mathcal{D}}_c \bar{\mathcal{F}}_{de} = 0$.

The other fifteen supersymmetries $\mathcal{Q}_a$ and $\mathcal{Q}_{ab}$ are broken on the lattice, and must be restored in the continuum limit to recover full $\mathcal{N} = 4$ SYM. This requirement is analogous to the restoration of $\text{SO}(4)$ euclidean Lorentz symmetry in the continuum limit of lattice QCD calculations, where only a discrete subgroup of $\text{SO}(4)$ is preserved at non-zero lattice spacing.

Finally, we define the parameters in the lattice action, Eq. 4. The overall factor outside most of the terms in S involves the number of colors N and the 't Hooft coupling $\lambda = g^2 N$, where g is the gauge coupling. The coefficient c_2 in Eq. 5 takes the value $c_2 = 1$ classically, but may be shifted by quantum effects, potentially requiring fine-tuning to recover $\mathcal{N} = 4$ SYM in the continuum limit. Preliminary investigations suggest that the necessary fine-tuning may be negligible in practice [22], and we typically fix $c_2 = 1$.

S_{stab} with its two coefficients μ^2 and κ is used to stabilize numerical computations. The ‘‘bosonic mass’’ μ regulates flat directions, lifting a bosonic zero mode present due to the periodic boundary conditions

(BCs) in all four directions. We lift the corresponding fermionic zero mode by imposing antiperiodic (thermal) temporal BCs for the fermions. Alternately, a fermion mass term

$$m \sum_n a^4 \left\{ \text{Tr} \left[\overline{\mathcal{D}}_b^{(-)} \mathcal{U}_b(n) \right] \text{Tr} \left[\mathcal{U}_a(n) \overline{\mathcal{U}}_a(n) \right] - \text{Tr} [\eta(n)] \text{Tr} [\psi_a(n) \overline{\mathcal{U}}_a(n)] \right\} \quad (11)$$

could be added to the action, but this is not currently implemented in `SUSY LATTICE`. Finally, when $\kappa > 0$ the plaquette determinant term projects the product of links around each plaquette $\mathcal{P}$ from $\mathfrak{gl}(N, \mathbb{C})$ to $\mathfrak{sl}(N, \mathbb{C})$, effectively reducing the gauge group from $U(N)$ to the target $SU(N)$. Specifically, $\kappa \geq 0.5$ forbids monopole condensation in the $U(1)$ sector, protecting the lattice system from these unphysical strong-coupling lattice artifacts [21].

Non-zero μ and κ softly break the 0-form supersymmetry $\mathcal{Q}$, which is otherwise exactly preserved even at non-zero lattice spacing: $\lim_{(\mu, \kappa) \rightarrow (0, 0)} \mathcal{QS} = 0$. (The fermion mass term in Eq. 11 is constructed in a $\mathcal{Q}$ -invariant way, but would break a global shift symmetry under $\eta \rightarrow \eta + c\mathbb{I}$ with c a constant Grassmann parameter.) The bosonic mass μ must be tuned to zero in the continuum limit to recover $\mathcal{N} = 4$ SYM. It may not be necessary to tune κ to zero, since its $\mathcal{Q}$ -breaking effects are confined to the $U(1)$ sector that decouples in the continuum limit where $U(N) = SU(N) \otimes U(1)$. In the next subsection we will discuss ways to monitor the effects of non-zero μ and κ in numerical calculations, along with ways to check the restoration of the broken supersymmetries $\mathcal{Q}_a$ and $\mathcal{Q}_{ab}$.

2.2. Rational hybrid Monte Carlo algorithm and important observables

Given the lattice formulation of $\mathcal{N} = 4$ SYM written above, we wish to compute operator expectation values through importance sampling Monte Carlo evaluation of the path integral in Eq. 1. Let us divide the lattice action into its bosonic and fermionic parts,

$$S[\mathcal{U}, \overline{\mathcal{U}}, \Psi] = S_B[\mathcal{U}, \overline{\mathcal{U}}] + \Psi^T \mathcal{D}[\mathcal{U}, \overline{\mathcal{U}}] \Psi \quad (12)$$

where $\Psi \equiv (\eta, \psi_a, \chi_{ab})^T$ is the Kähler–Dirac fermion field and $\mathcal{D}[\mathcal{U}, \overline{\mathcal{U}}]$ is the (antisymmetric) fermion operator. Eq. 1 then becomes

$$\langle \mathcal{O} \rangle = \frac{1}{\mathcal{Z}} \int [d\mathcal{U}][d\overline{\mathcal{U}}][d\Psi] \mathcal{O} e^{-S[\mathcal{U}, \overline{\mathcal{U}}, \Psi]} = \frac{1}{\mathcal{Z}} \int [d\mathcal{U}][d\overline{\mathcal{U}}] \mathcal{O} e^{-S_B[\mathcal{U}, \overline{\mathcal{U}}]} \text{pf } \mathcal{D}[\mathcal{U}, \overline{\mathcal{U}}], \quad (13)$$

where the Pfaffian $\text{pf } \mathcal{D}[\mathcal{U}, \overline{\mathcal{U}}]$ results from Gaussian integration over Ψ .

Because we wish to employ $[e^{-S_B} \text{pf } \mathcal{D}]$ as a Boltzmann weight, both the bosonic action and Pfaffian should be real and positive. This condition is satisfied for S_B . Unfortunately, for a given gauge field configuration the Pfaffian is generically complex, $\text{pf } \mathcal{D} = |\text{pf } \mathcal{D}| e^{i\alpha}$. We proceed by considering the phase-quenched path integral

$$\langle \mathcal{O} \rangle_{pq} = \frac{1}{\mathcal{Z}_{pq}} \int [d\mathcal{U}][d\overline{\mathcal{U}}] \mathcal{O} e^{-S_B[\mathcal{U}, \overline{\mathcal{U}}]} |\text{pf } \mathcal{D}[\mathcal{U}, \overline{\mathcal{U}}]| \quad \mathcal{Z}_{pq} = \int [d\mathcal{U}][d\overline{\mathcal{U}}] e^{-S_B[\mathcal{U}, \overline{\mathcal{U}}]} |\text{pf } \mathcal{D}[\mathcal{U}, \overline{\mathcal{U}}]|. \quad (14)$$

So long as $\langle e^{i\alpha} \rangle_{pq}$ is statistically non-zero, the true expectation values can be reconstructed via phase reweighting,

$$\langle \mathcal{O} \rangle = \frac{\langle \mathcal{O} e^{i\alpha} \rangle_{pq}}{\langle e^{i\alpha} \rangle_{pq}}. \quad (15)$$

The phase of the Pfaffian is therefore a crucial observable to monitor in phase-quenched calculations, and we will discuss its measurement in detail in Section 4.4. In practice, we find that the Pfaffian of lattice $\mathcal{N} = 4$ SYM is very nearly real and positive, so that $\langle \mathcal{O} \rangle_{pq} \approx \langle \mathcal{O} \rangle$ and we do not suffer from an insurmountable sign problem [21]. Accordingly, we will omit the $_{pq}$ subscripts in the remainder of this article.

Explicitly evaluating the Pfaffian is far too computationally expensive to be a part of our configuration generation algorithm. The standard approach we use instead employs a set of bosonic pseudofermions Φ , of the same dimensionality as Ψ , which are introduced through the mathematical identity

$$|\text{pf } \mathcal{D}| = |\det \mathcal{D}|^{1/2} = \left(\det[\mathcal{D}^\dagger \mathcal{D}] \right)^{1/4} \propto \int [d\Phi^\dagger][d\Phi] \exp \left[-\Phi^\dagger \left(\mathcal{D}^\dagger \mathcal{D} \right)^{-1/4} \Phi \right]. \quad (16)$$

The negative fractional power is achieved by the use of the rational hybrid Monte Carlo (RHMC) algorithm to generate an appropriate distribution of gauge configurations [28].

There are three main steps to the RHMC algorithm. First the pseudofermion field is set to $\Phi(n) = (\mathcal{D}^\dagger \mathcal{D})^{1/8} g(n)$ at each lattice site, with $g(n)$ being $16N^2$ -component random vectors drawn from a Gaussian distribution. We also define random Gaussian $\Pi_a(n) \in \mathfrak{gl}(N, \mathbb{C})$ to serve as fictitious momenta conjugate to the gauge fields $\mathcal{U}_a(n)$, leading to an effective Hamiltonian

$$H = \frac{1}{2} \sum_n \text{Tr} [\Pi_a^2(n)] + S_B[\mathcal{U}, \bar{\mathcal{U}}] + \Phi^\dagger \left(\mathcal{D}^\dagger \mathcal{D} \right)^{-1/4} \Phi. \quad (17)$$

The core of the algorithm is molecular dynamics (MD) evolution of Π_a and $\mathcal{U}_a$ along a trajectory of length τ in a fictitious MD time. Since the pseudofermion field already has the proper distribution, Φ is kept fixed along this trajectory. The equation of motion for the links is simplified by the fact that they live in $\mathfrak{gl}(N, \mathbb{C})$ rather than $\text{SU}(N)$. The MD evolution involves inexact integration of Hamilton's equations, which does not conserve the effective Hamiltonian, $\Delta H \neq 0$. The final step of the RHMC algorithm is a Metropolis–Rosenbluth–Teller test: the new gauge field values produced by the MD evolution are accepted with probability $P = \min(1, e^{-\Delta H})$, otherwise the original field values are restored to produce a new configuration identical to the starting configuration. So long as the MD integration scheme is area preserving (symplectic) and reversible (symmetric), the acceptance test makes the algorithm exact by stochastically correcting for the integration errors.

The above summary of RHMC is almost identical to that of the more familiar HMC algorithm [29]. The main difference between the two appears in the MD evolution, each step of which requires solving

$$\frac{\delta \mathcal{U}_a}{\delta t} = \frac{\delta H}{\delta \Pi_a} = \Pi_a \quad \frac{\delta \Pi_a}{\delta t} = -\frac{\delta H}{\delta \mathcal{U}_a} = -\frac{\delta S_B}{\delta \mathcal{U}_a} - \frac{\delta}{\delta \mathcal{U}_a} \Phi^\dagger \left(\mathcal{D}^\dagger \mathcal{D} \right)^{-1/4} \Phi. \quad (18)$$

In order to take the Lie derivative of the pseudofermion term in the effective Hamiltonian, we approximate the rational power of the matrix inverse by a series of P partial fractions,

$$\left(\mathcal{D}^\dagger \mathcal{D} \right)^{-1/4} = \alpha_0 + \sum_{i=1}^P \alpha_i \left(\mathcal{D}^\dagger \mathcal{D} + \beta_i \right)^{-1} \quad (19)$$

(and similarly for the rational power $(\mathcal{D}^\dagger \mathcal{D})^{1/8}$ needed to initialize Φ). The fermionic contribution to the force $\frac{\delta H}{\delta \mathcal{U}_a}$ then becomes

$$\frac{\delta}{\delta \mathcal{U}_a} \Phi^\dagger \left(\mathcal{D}^\dagger \mathcal{D} \right)^{-1/4} \Phi = \sum_{i=1}^P \alpha_i \left[(\mathcal{D}^\dagger \mathcal{D} + \beta_i)^{-1} \Phi \right]^\dagger \frac{\delta (\mathcal{D}^\dagger \mathcal{D})}{\delta \mathcal{U}_a} \left[(\mathcal{D}^\dagger \mathcal{D} + \beta_i)^{-1} \Phi \right], \quad (20)$$

and all the $(\mathcal{D}^\dagger \mathcal{D} + \beta_i)^{-1} \Phi$ can be efficiently determined by a multi-shift conjugate gradient (CG) inverter [30]. These inversions dominate the cost of lattice computations.

In Sections 3 and 4 we will provide some further details about the rational approximations and MD integration schemes currently provided by SUSY LATTICE. For now we remark that optimal shifts β_i and amplitudes α_i are computed offline by minimizing the relative approximation error within a given spectral range $[\lambda_{\text{low}}, \lambda_{\text{high}}]$, with $\lambda_{\text{low}} > 0$ [31]. Keeping the errors negligible across larger spectral ranges requires

adding more terms to the partial fraction approximation, increasing computational costs. At a minimum, we must demand that the smallest and largest eigenvalues of $\mathcal{D}^\dagger \mathcal{D}$ fall within the spectral range we use, $\lambda_{low} < \lambda_{min} \ll \lambda_{max} < \lambda_{high}$. The extremal eigenvalues of the squared fermion operator are therefore, like the Pfaffian, important quantities to monitor during RHMC evolution. Fortunately such eigenvalue computations are much cheaper than Pfaffian measurements, and in **SUSY LATTICE** we carry them out using the PReconditioned Iterative Multi-Method Eigensolver (PRIMME) library [32].

By repeating the steps of the RHMC algorithm to accumulate some number (typically thousands) of MD time units, we generate an ensemble of gauge configurations. We have already discussed some important observables to measure using these configurations, namely the Pfaffian and low-lying eigenvalues of the fermion operator. A simpler set of common lattice observables consists of products of gauge links around closed loops on the lattice. Since the links are elements of $\mathfrak{gl}(N, \mathbb{C})$, they are not unitarized and the simplest possible observable $\frac{1}{N} \text{Tr} [\mathcal{U}_a(x) \bar{\mathcal{U}}_a(x)]$ is non-trivial. The plaquette $\text{Tr} [\mathcal{U}_b(x) \mathcal{U}_a(x + \hat{\mu}_b) \bar{\mathcal{U}}_b(x + \hat{\mu}_a) \bar{\mathcal{U}}_a(x)]$ is also straightforward to compute even on the A_4^* lattice with five basis vectors at each lattice site, and its determinant provides information about how far the links are from $\mathfrak{sl}(N, \mathbb{C})$. The average link-trace, plaquette, Polyakov loop $\text{Tr} [\prod_{N_t} \mathcal{U}_t]$, and all contributions to the lattice action are measured and printed after every trajectory.

Further observables are typically computed only on stored configurations. These include larger Wilson loops oriented along the principal axes of the lattice, which can be used to monitor the restoration of the $\mathcal{Q}_a$ and $\mathcal{Q}_{ab}$ supersymmetries, as discussed in Ref. [21]. To monitor the $\mathcal{Q}$ supersymmetry that is broken only by non-zero μ and κ in S_{stab} (Eq. 7), we measure violations of the Ward identity

$$\mathcal{QO} = \text{Tr} \left[\sum_b (\mathcal{U}_b \bar{\mathcal{U}}_b - \bar{\mathcal{U}}_b \mathcal{U}_b) \sum_a \mathcal{U}_a \bar{\mathcal{U}}_a \right] - \text{Tr} \left[\eta \sum_a \psi_a \bar{\mathcal{U}}_a \right] = 0, \quad (21)$$

which requires computing the fermion bilinear $\eta \psi_a$. In **SUSY LATTICE** we compute $\eta \psi_a$ by inverting the fermion operator $\mathcal{D}$ on Gaussian-distributed stochastic (“noisy”) sources, a standard technique used to compute the chiral condensate in lattice QCD.⁴ The static potential is also extracted from Wilson loops, but in this case it is advantageous to measure loops for all possible spatial separations $\vec{r}$, not just the “on-axis” loops that lie along the principal axes of the lattice. To do so, we gauge fix to Coulomb gauge and compute

$$W(\vec{r}, t) = \text{Tr} \left[P(\vec{x}, t, t_0) P^\dagger(\vec{x} + \vec{r}, t, t_0) \right], \quad (22)$$

where $P(\vec{x}, t, t_0)$ is a product of temporal links $\mathcal{U}_t$ at spatial location $\vec{x}$, which extends from timeslice t_0 to timeslice $t_0 + t$.

3. Running the code: High-level overview

In this section we summarize how to obtain **SUSY LATTICE**, compile the provided programs, and run the resulting executables on a generic GNU/Linux system. In the next section we will describe the code itself in greater detail, providing information necessary to extend it with additional measurements or other features.

SUSY LATTICE is developed in a publicly accessible version control repository [49], so that one way to download the most recent version of the code is to run

```
$ git clone https://github.com/daschaich/susy.git
```

Alternately, a tarball containing the version of the code discussed in this article may be found in the CPC program library. After obtaining the code, move to the **4dSYM/susy** application directory. In addition to a **README** and other files to be discussed in the next section, this directory contains two sample Makefiles:

⁴The structure of the fermion operator allows us to obtain two estimates of $\eta \psi_a$ for each stochastic source, by simultaneously computing both the η components of $\mathcal{D}^{-1} \psi_a$ and the ψ_a components of $\mathcal{D}^{-1} \eta$.

`Make_scalar` uses `gcc` to compile serial executables, while `Make_mpi` uses `mpicc` to compile parallel (MPI) executables. These Makefiles may need minor modifications on a given system. (For example, `Make_mpi` is set up for the USQCD clusters at Fermilab.) They are both wrappers that load architecture-independent information from the main `Make_template`, where all the available executables are defined.

To compile RHMC evolution for serial (“scalar”) execution, simply run

```
$ make -f Make_scalar susy_hmc
```

This will produce an executable named `susy_hmc`, one of several executables that are actively maintained and well tested:

- `susy_hmc` runs the RHMC algorithm for gauge configuration generation, along with inexpensive basic measurements (e.g., of the plaquette, link-trace and Polyakov loop) that can be used to monitor the RHMC evolution.
- `susy_meas` carries out more expensive measurements on saved gauge configurations, as discussed at the end of the previous section. These include both on-axis and gauge-fixed Wilson loop computations, with optional stout smearing. Measurements of the $\eta\psi_a$ fermion bilinear from inversions using a given number of stochastic sources are also optional, and dominate the computational cost if performed.
- `susy_eig` is an interface to PRIMME [32] that computes the smallest and largest eigenvalues of $\mathcal{D}^\dagger\mathcal{D}$ for a given gauge configuration, along with the corresponding eigenvectors.
- `susy_phase` computes the complex Pfaffian of the fermion operator $\mathcal{D}[\mathcal{U},\overline{\mathcal{U}}]$ for a given gauge configuration, using the algorithm described in Section 4.4.

An additional `susy_hmc_meas` executable carries out the same measurements as `susy_meas` in the course of RHMC evolution.

In the `testsuite/scalar` directory we provide a script named `run_tests` that will compile each of these executables, run it in serial using a fixed set of input options, and check the resulting output against reference files distributed with the code. The `run_tests` script takes two optional command-line arguments. Running

```
$ ./run_tests <N> <target>
```

will compile, run and check the `susy_${target}` executable for gauge group $U(\$N)$, where `$N` must be either 2, 3 or 4. Alternately, if `run_tests` is executed with no arguments it will check every executable listed above, and then test that the Pfaffian computation can be split into multiple checkpointed pieces. Running every test in this manner will take several hours, with the bulk of the time spent computing the Pfaffian. Similar tests of the parallel executables can be carried out by customizing the `run_tests` script in the `testsuite/mpi` directory to use the appropriate Makefile and `mpirun` for the MPI system to be tested. So long as the parallel tests are run using two cores, the reference output should be reproduced.

All regular SUSY LATTICE output is written to `stdout`, and is typically redirected to a single plain-text file. Each measurement appears on a line that starts with a corresponding label, which can easily be selected by offline analyses using `awk`, `python`, `perl`, etc. The in-line documentation specifies the labels currently used, along with other details of the output.

In the next two subsections we will summarize various options that users may specify when using SUSY LATTICE. Some of these options have to be fixed when compiling, while others are read in by the executables when they are run. We will then conclude this section by discussing the scaling and parallel performance of SUSY LATTICE.

3.1. Compile-time options

The executables listed above are produced by compiling with appropriate combinations of macro definitions: `-DHMC_ALGORITHM` for RHMC evolution; `-DWLOOP` for gauge-fixed Wilson loop computations; `-DSTOUT` for stout smearing; `-DBILIN` for fermion bilinear measurements using stochastic sources; `-DEIG` for eigenvalue measurements using PRIMME; and `-DPHASE` to compute the complex Pfaffian. In addition, several other features of the code must be fixed at the time of compilation. For example, the gauge group $U(N)$

is chosen by uncommenting one (and only one) definition of the number of colors $\text{NCOL} = N$ and corresponding dimension of the adjoint fermion representation $\text{DIMF} = \text{NCOL}^2$ at the top of the header file `include/su3.h`. The default gauge group is $U(2)$. Gauge groups $U(3)$ and $U(4)$ are also well tested and available through this header file. Larger values of NCOL require further extensions of the code, as we will discuss in Section 4.1. In practice, computational costs increase rapidly as the number of colors grows. In Fig. 1, we observe the costs of RHMC configuration generation increasing $\propto N^5$, while the Pfaffian calculation scales $\propto \text{DIMF}^3 = \text{NCOL}^6$ as expected. This behavior makes it unlikely that computations with $N \geq 5$ will be practical in the immediate future. Even so, because deviations from the large- N limit of $N = 4$ SYM go as $1/N^2$, investigations with $N \leq 4$ should access the large- N regime up to few-percent effects that may be comparable to statistical uncertainties.

Figure 1: Total computational costs for $8^3 \times 24$ RHMC generation of a single molecular dynamics time unit, and $2^3 \times 4$ Pfaffian measurements, for gauge groups $U(2)$, $U(3)$ and $U(4)$, on log-log axes. The red dashed line illustrates the cost scaling $\propto N^5$ observed for RHMC configuration generation, while the blue dashed line shows the expected N^6 scaling for Pfaffian measurements.

Other important compile-time options may be found in the header file `susy/defines.h`. The simplest of these is `PBC`, which controls the temporal boundary conditions (BCs) of the fermions. This flag can be set to 1 for periodic BCs, or -1 for antiperiodic (thermal) BCs. The default antiperiodic BCs are recommended, to lift a fermionic zero mode that may otherwise destabilize the computations. The coefficient c_2 in S_{exact} (Eq. 5) may be changed by redefining the flag `C2` in this file. We recommend leaving this coefficient set to its classical value $c_2 = 1$.

Finally, the order P of the rational approximation in Eq. 19 is set by the flag `DEGREE` in `defines.h`. Five options are supported, each of which provides rational approximations that reproduce $(\mathcal{D}^\dagger \mathcal{D})^{-1/4}$ and $(\mathcal{D}^\dagger \mathcal{D})^{1/8}$ up to fractional errors of at most 2×10^{-5} across a given spectral range:

- `DEGREE 5` is set up for the spectral range $[0.1, 50]$
- `DEGREE 6` is set up for the spectral range $[0.02, 50]$
- `DEGREE 8` is set up for the spectral range $[10^{-3}, 50]$
- `DEGREE 9` is set up for the spectral range $[10^{-4}, 45]$
- `DEGREE 15` is set up for the spectral range $[10^{-7}, 1000]$

A smaller `DEGREE` allows the CG inverter to converge more quickly, but if the eigenvalues of $\mathcal{D}^\dagger \mathcal{D}$ exceed the spectral range of the rational approximation, significant numerical errors may result. A `DEGREE` of 8 or less suffices for our typical computations.

3.2. Run-time options

When run, each executable must be given a set of input parameters that specify the lattice volume, the gauge configuration to load, and other pertinent information. All possible input parameters for every executable discussed above are summarized in Fig. 2. Because this figure combines input options for multiple executables, it cannot be used as an input file itself. Sample input files for each executable are provided in the `testsuite` directory. Although comments and whitespace are ignored, the options must come in the order expected by the executable, and each variable must be immediately preceded by its name as shown in Fig. 2.

Figure 2: All possible input parameters that may be read in by a `SUSY LATTICE` executable, with brief in-line explanations.

```

0  prompt 0  # If non-zero, prompt user to manually input each parameter
1  nx 6      # Lattice volume is nx*ny*nz*nt
2  ny 6
3  nz 6
4  nt 6
5  iseed 41  # Random number generator seed
6
7  warms 0    # Number of trajectories without expensive measurements
8  trajecs 10 # Number of trajectories with expensive measurements
9  traj_length 1 # Trajectory length (tau)
10 nstep 10   # Fermion steps per trajectory; step_size = traj_length / nstep
11 nstep_gauge 10 # Gauge steps per fermion step
12 traj_between_meas 10 # How many trajectories to skip between expensive measurements
13
14 lambda 1.0 # 't Hooft coupling
15 kappa_u1 0.5 # Plaquette determinant coupling
16 bmass 0.5 # Bosonic mass parameter (mu)
17 fmass 0.0 # Mass shift in fermion operator
18
19 # The next two lines must only be included when compiling with -DSTOUT
20 Nstout 1 # Number of stout smearing steps
21 rho 0.1 # Stout smearing parameter
22
23 max_cg_iterations 5000 # Maximum number of CG iterations in each inversion
24 error_per_site 1e-5 # "Linear" CG stopping condition (squared to compare with |r|^2)
25
26 # The next line must only be included when compiling with -DBILIN
27 nsrc 3 # Number of stochastic sources for fermion bilinear calculations
28
29 # The next three lines must only be included when compiling with -DEIG
30 Nvec 100 # Number of eigenvalues and eigenvectors to calculate
31 eig_tol 1e-8 # Eigensolver tolerance (convergence criterion)
32 maxIter 5000 # Maximum number of eigensolver iterations
33
34 # The next two lines must only be included when compiling with -DPHASE
35 ckpt_load -1 # If positive, load checkpointed pfaffian computation
36 # from config.Q$ckpt_load and config.diag$ckpt_load
37 ckpt_save -1 # If positive, checkpoint pfaffian computation
38 # to config.Q$ckpt_save and config.diag$ckpt_save
39
40 reload_serial config # How to set up lattice at start:
41 # only fresh, random, continue and reload_serial supported
42 no_gauge_fix # Only coulomb_gauge_fix and no_gauge_fix supported
43 save_serial new-config # What to do with lattice at end: only forget and save_serial supported

```

While the in-line explanations in Fig. 2 should suffice to explain most options, some further comments are called for. First, the lines starting with `reload_serial`, `warms`, `trajecs` and `save_serial` determine the overall structure of an RHMC computation. In this case, the configuration `config` would be loaded

from disk, the RHMC algorithm would be run for ten trajectories (with trajectory length of $\tau = 1$ MD time unit between accept/reject tests set by `traj_length`), and the resulting configuration would be saved to disk as `new-config`. A trajectory length of $\tau = 1$ is a reasonable default. Although longer trajectories decrease autocorrelations between subsequent configurations in the Markov chain, they also require smaller steps in the MD evolution, to keep $|\Delta H| \lesssim 1$ and maintain reasonable acceptance. The details of optional checkpointing in the Pfaffian computation will be discussed in Section 4.4. Finally, we emphasize that the `fmass` parameter in the input is *not* the Q -invariant coefficient m that appears in Eq. 11. Instead this parameter simply shifts the squared fermion operator

$$\mathcal{D}^\dagger \mathcal{D} \longrightarrow \mathcal{D}^\dagger \mathcal{D} + \text{fmass}^2. \quad (23)$$

This shift is not needed when the fermionic fields are subject to antiperiodic temporal BCs, and we recommend using compile-time option `PBC = -1` and `fmass = 0`.

3.3. Performance and scaling

The performance advantages from parallelization are dramatic, and greatly extend the applicability of `SUSY LATTICE` compared to the serial code presented in Ref. [1]. For typical small-volume calculations (e.g., $4^3 \times 12$ gauge configuration generation with $N = 2$), `SUSY LATTICE` is roughly two orders of magnitude faster than the serial code when each is run in its standard production environment (including single-GPU acceleration of the serial code’s multi-shift CG inverter [33], which significantly improves its performance). More importantly, `SUSY LATTICE` can be scaled up to larger lattice volumes and larger gauge groups than the serial code can handle – the restriction to a single GPU limits the serial code to lattice volumes of at most $8^3 \times 16$ even for the smallest $N = 2$.

Efficiently studying larger lattice volumes and larger N requires that `SUSY LATTICE` performs well when running across many cores. We observe good parallel scaling from `SUSY LATTICE` (thanks in large part to the work done over many years to optimize the MILC code), although further performance improvements may be feasible. First consider the strong scaling shown in Fig. 3, which considers the time to solution for a fixed computation spread across various numbers of cores. The computation under consideration is to generate a single MD time unit with the RHMC algorithm, for a $16^3 \times 32$ -site lattice volume and either gauge group $U(2)$ or $U(3)$. The time to solution steadily decreases as we use up to 512 cores (32 nodes) of the USQCD pi0 cluster at Fermilab.

In particular, up to around 128 cores (eight 16-core pi0 nodes) the scaling is roughly optimal, following a straight line with slope -1 on the logarithmic axes of Fig. 3. The local volume on 128 cores is $4^2 \times 8^2$ lattice sites per core, which decreases to $4^3 \times 8$ and 4^4 sites per core on 256 and 512 cores, respectively. The performance deteriorates on more than 128 cores, as communication between all these cores begins to take longer than the computations that each core has to carry out. Although lattice $\mathcal{N} = 4$ SYM involves more computation per site than lattice QCD with staggered or Wilson quarks, it also requires much more intercore communication. In particular, the structure of the S_{closed} term in Eq. 6 introduces dozens of intercore data transfers (“gathers”) into each application of the fermion operator (the matrix–vector operation in the CG algorithm).

We can also consider weak scaling, in which the problem size increases in tandem with the number of cores being used. In Fig. 4 we plot the total computational cost in core-hours for Pfaffian measurements with gauge group $U(2)$ and lattice volumes $3^3 \times 4$, $3^3 \times 6$ and $3^3 \times 8$ on 2, 3 and 4 cores, respectively. That is, the local volume of $3^3 \times 2$ sites per core is fixed, and all computations use part of a single 8-core node in the HEP-TH cluster at the University of Colorado. The straight line on the logarithmic axes of Fig. 4 is a fit to power-law scaling, with a power of 2.86(7). The Pfaffian computation requires $\mathcal{O}(V^3)$ operations, so this power of approximately 3 indicates optimal weak scaling of the code.

4. Modifying the code: Underlying details

In this section we provide more information about elements of the code itself, highlighting notable aspects of the documentation for the benefit of those who may be interested in contributing to future de-

Figure 3: Strong scaling of the time to solution for U(2) and U(3) $16^3 \times 32$ RHMC generation of a single molecular dynamics time unit on the USQCD pi0 cluster at Fermilab, on log-log axes. The dashed lines illustrate optimal scaling (a power of -1), which the computations follow fairly well through 128 cores (a local volume of $4^2 \times 8^2$ lattice sites per core), with deteriorating performance on 256 and 512 cores (local volumes $4^3 \times 8$ and 4^4).

Figure 4: Weak scaling of the total computational cost for U(2) Pfaffian measurements with fixed local volume $3^3 \times 2$, on log-log axes with a power-law fit. Since the Pfaffian computation requires $\mathcal{O}(V^3)$ operations, the power of approximately 3 indicates optimal scaling of the code.

velopment. `SUSY LATTICE` evolved from a version of the MILC QCD code that was extended by DeGrand, Shamir and Svetitsky (DSS) to handle $SU(N)$ gauge groups with arbitrary N and fermions transforming in either the fundamental, adjoint, two-index symmetric or two-index antisymmetric representations. Compared to the DSS code, further changes were required to add the fifth link required by the A_4^* lattice structure, to implement Kähler–Dirac fermions and the $\mathcal{N} = 4$ SYM lattice action, and to convert the gauge links from $SU(N)$ to $\mathfrak{gl}(N, \mathbb{C})$. Even so, the basic structure of the code continues to resemble MILC, and much of the information in the MILC manual remains relevant [23]. After describing the directory tree, in the subsequent subsections we will address issues that make `SUSY LATTICE` different from a typical parallel QCD code.

The five subdirectories within the `4dSYM` directory are set up analogously to the MILC code:

- `libraries` contains simple single-processor routines, for example $N \times N$ matrix multiplication.
- `include` holds header files for procedures common to typical applications, as well as important compile-time definitions such as the number of colors N and basic N -dependent data structures.
- `generic` implements procedures common to typical applications, such as data layout, internode communications, and input/output (I/O).
- `susy` is currently the only application directory, containing code for four-dimensional $\mathcal{N} = 4$ SYM.
- `testsuite` provides reference input and output files for all actively maintained executables, along with scripts to test any or all of these executables, in serial or in parallel, for $U(N)$ gauge theory with $N = 2, 3$ or 4 colors.

The `2dSYM` directory with the same structure is set up for two-dimensional $\mathcal{N} = (2, 2)$ SYM theory, which we briefly discuss in the appendix.

4.1. Gauge group and fermion representation

To generalize the MILC code to handle arbitrary $SU(N)$, the DSS code replaced the hardwired “3” for the number of colors in $SU(3)$ with a macro definition `NCOL`. This involved modifications to most single-processor library routines (e.g., $\text{NCOL} \times \text{NCOL}$ matrix multiplication). The bosonic fields in `SUSY LATTICE` are $\text{NCOL} \times \text{NCOL}$ complex matrices, just as the gauge fields in non-supersymmetric $SU(\text{NCOL})$ lattice gauge theory would be. However, because the lattice $\mathcal{N} = 4$ SYM link variables are elements of $\mathfrak{gl}(N, \mathbb{C})$ rather than $SU(N)$, they are not unitarized in the course of the simulation.

The DSS code allows for fermions transforming in representations other than the fundamental. The dimensionality of the fermion representation is set by the macro definition `DIMF`. While fermions in the adjoint representation can be naturally encoded in a set of $\text{NCOL} \times \text{NCOL}$ complex matrices, the DSS code remains more general by making the fermionic fields `DIMF`-component vectors in color space. We retain this setup in `SUSY LATTICE`, with `DIMF = NCOL2` fixed for the $U(\text{NCOL})$ adjoint representation. This approach forces us to keep two sets of bosonic variables in the code. One set, called `su3_matrix_f`, denotes $\text{NCOL} \times \text{NCOL}$ complex matrices and is used in the bosonic part of the code. The other set, called `su3_matrix`, denotes `DIMF` $\times$ `DIMF` matrices and appears in the routines involving (pseudo)fermions, such as the calculation of the pseudofermion action through Eq. 16. An `su3_matrix` $\mathcal{V}_a$ is composed from an `su3_matrix_f` $\mathcal{U}_a$ by

$$\mathcal{V}_a^{AB} = \text{Tr} \left[\lambda_{ij}^A (\mathcal{U}_a)_{jk} \lambda_{ki}^B \right], \quad (24)$$

where the λ^A matrices are generators of $U(N)$ defined in the file `susy/setup_lambda.c`. Whenever the gauge configuration $\mathcal{U}_a$ is updated, the “adjoint links” $\mathcal{V}_a$ must be regenerated by the `fermion_rep()` routine. Of course, one must keep a complete set of library routines for matrix operations involving either `su3_matrix` or `su3_matrix_f` variables.

As mentioned in Section 3.1, certain structures and routines are currently implemented only for $\text{NCOL} \leq 4$. These include the `anti_hermitmat` structure in `include/su3.h`, which explicitly lists the $\text{NCOL}(\text{NCOL} + 1)/2$ non-trivial components of $\text{NCOL} \times \text{NCOL}$ anti-hermitian complex matrices. This structure is used for stout smearing, by routines that set up random gauge configurations, and by four library routines. In addition,

monitoring the restoration of the $\mathcal{Q}_a$ and $\mathcal{Q}_{ab}$ supersymmetries requires inverting $\text{NCOL} \times \text{NCOL}$ gauge link matrices, which is implemented using the cofactor matrix for $\text{NCOL} \leq 4$. While it is straightforward to extend these structures and routines to $\text{NCOL} \geq 5$, the rapidly growing computational costs make such work unlikely to prove useful at present.

4.2. Data layout and communications

The MILC code lays out its dynamical variables in two ways. “Site-wise” variables are packaged into a `site` structure defined globally in the header file `susy/lattice.h`. The lattice itself is an array of such `sites`, each of which contains, inter alia, five gauge links $\mathcal{U}_a$ and five adjoint links $\mathcal{V}_a$. This construction dates from the earliest version of the MILC code, when this layout was optimal for the compilers and computing architectures then in use. More recently the trend has been to define each lattice variable as a separate “field-wise” array assigned a contiguous block of memory on each node; see Section 5.3 of Ref. [23] for further discussion. These field-wise variables are also globally defined in `susy/lattice.h`, and allocated in `susy/setup.c` after the lattice volume is set at run time (Fig. 2).

Because both site-wise and field-wise variables coexist in `SUSY LATTICE`, we must maintain two sets of internode communication (“gather”) routines, with each set handling either site-wise or field-wise data layout. In addition, the macro `FORALLSITES` is provided to simplify looping over the lattice volume. This macro keeps track of both the linear index of field-wise variables along with a pointer to the corresponding `site`, making it easy to use the two types of variables in combination.

A new structure in `SUSY LATTICE` that does not appear in the MILC code is called `Twist_Fermion`, and bundles together all the fermion variables $\Psi = (\eta, \psi_a, \chi_{ab})$, for convenience.

The MILC code was written to deal with four-dimensional gauge fields discretized on hypercubic lattices. Even using the convenient representation of the A_4^* lattice discussed in Section 2.1, in which four of the five gauge links at each site are the usual hypercubic objects, accommodating the fifth link $\hat{\mu}_4 = (-1, -1, -1, -1)$ required rewriting various low-level routines. For example, utilities for reading and writing binary gauge configuration files had to be extended to include the fifth link.

In addition, in the MILC code each `site` in the lattice carries a flag labeling its parity in the sense of an even/odd checkerboard: the parity is

$$p(n) = \text{mod} \left(\sum_{i=0}^3 n_i, 2 \right). \quad (25)$$

The first four links connect “even” ($p = 0$) sites to “odd” ($p = 1$) sites, and vice versa. This checkerboarding is widely used in ordinary QCD codes. For example, staggered fermions are often defined to live on only one parity, both at the level of HMC evolution and in measurements. In `SUSY LATTICE` it is hard to use checkerboarding, since the fifth link connects sites of the same parity.

The only context in which we make use of the parity flag is parallel gauge fixing for static potential analyses. A typical gauge choice is one which maximizes some set of links. The only option currently implemented in `SUSY LATTICE` is Coulomb gauge, which maximizes $\sum_{i=0}^3 \text{Tr} [\mathcal{U}_i(n)]$ at every site. The formula for the gauge transformation, $\mathcal{U}_a(n) \rightarrow G(n) \mathcal{U}_a(n) G^\dagger(n + \hat{\mu}_a)$, motivates considering an iterative sequence of gauge transformations G that are alternately taken to be the identity on even or on odd sites. Then for the first four links the maximization becomes linear in the G s, reducing to multiplication on the left for links at sites of the parity for which G is nontrivial, and on the right for links at sites of the other parity. The fifth link always has to be multiplied on both sides, and cannot be included in the maximization without substantially rewriting the gauge fixing code. This is not an issue for Coulomb gauge, which includes only the first three links in the gauge condition.

Next, we must implement antiperiodic temporal BCs for the fermions differently than in standard QCD codes where matter fields are defined only on the sites. In the QCD context, such BCs can be included simply by negating the temporal gauge links on the last timeslice of the lattice. Since the ψ_a and χ_{ab} fermionic variables in `SUSY LATTICE` connect different sites, this trick doesn’t work. Instead, we include

in each `site` the necessary factors of $\text{PBC} = \pm 1$ (discussed in Section 3.1) to be applied to the fermion fields whenever they cross the temporal boundary.⁵ This is the most straightforward implementation of antiperiodic temporal BCs, but is unlikely to be the most efficient. These BC factors are set up in `susy/setup_offset.c` after the lattice has been laid out.

The primary purpose of this `setup_offset.c` file is to speed up the internode communications mentioned above. Both site-wise and field-wise gathers come in two varieties. The usual gathers can move data only along specific paths that are tracked by a global static `gather_array` set up by routines in `generic/com_mpi.c`. For example, these include the paths connecting nearest-neighbor sites in the four spacetime directions. In contrast, “general” gathers handle an arbitrary lattice displacement four-vector, but take much longer to run. Every general gather must determine which nodes will exchange data, while this information is precomputed for the usual gathers. In addition, only one general gather can run at a time. Many of the usual gathers may be carried out simultaneously, which we use to overlap communications with computations.

In order to use the usual gathers even for the several dozen two- and three-link paths required by the Q -closed part of the action (Eq. 6), in `susy/setup_offset.c` we add all of these paths to the `gather_array`. We also add all five A_4^* link paths to the `gather_array`, in addition to the four spacetime directions set up by default in `generic/com_mpi.c`. We keep track of all these new gathers in the arrays `gq_offset` and `goffset`, respectively. Although the first four A_4^* link paths duplicate the four spacetime directions, the latter are set up to allow even/odd checkerboarding as in MILC, while the former are not. In the rare event that checkerboarding is needed, the standard MILC-like gather paths must be used. Otherwise we recommend gathering via the `goffset` array.

Avoiding general gathers in the Q -closed term significantly improves the performance of `SUSY LATTICE`, by roughly a factor of two. Overlapping communications with computations in this term is also beneficial when running on many nodes. For $U(2)$ 12^4 gauge generation on 64 cores (2 nodes) of the USQCD bc cluster at Fermilab, with local volume $3^2 \times 6^2$, such overlapping reduced costs by an additional $\sim 30\%$.

4.3. Rational hybrid Monte Carlo

As summarized in Section 2.2, the molecular dynamics (MD) evolution at the heart of the RHMC algorithm is responsible for the bulk of computing in typical lattice calculations. Lattice QCD researchers have developed many clever techniques to make the MD evolution more efficient, several of which we import to `SUSY LATTICE` from the DSS code. Our default MD algorithm is a second-order Omelyan integrator [34, 35], with separate time scales [36] for the gauge and fermion force terms in Eq. 18. To help tune run parameters the average and maximum forces along each trajectory are printed in the output, separately for the gauge and fermion terms.

While developing `SUSY LATTICE`, it was useful first to implement a simple leapfrog integrator, against which our current algorithm could be tested. In case similar testing is necessary in the future, we have retained this leapfrog integrator in the source code distribution. It can be used by replacing `update_o.o` with `update_leapfrog.o` in `susy/Make_template`. The leapfrog integrator has only a single time scale, and for simplicity this is set by `nstep` in Fig. 2, ignoring `nstep_gauge`.

The five different rational approximations discussed in Section 3.1 are set up in `susy/setup_rhmc.c`, simply by copying in output from the Remez algorithm implemented in Ref. [31]. For a given spectral range, we find the smallest degree P in Eq. 19 that produces a maximum fractional approximation error less than 2×10^{-5} , a number set in previous work [1]. While it is straightforward to add more approximations, those already available should provide good efficiency for foreseeable computations. The shifts β_i and amplitudes α_i chosen at compile time are used by a standard multi-shift CG inverter [30]. We make it possible to redefine P , the shifts and the amplitudes during the computation, so that the same inversion routine also provides the basic CG used to compute the fermion bilinear $\eta\psi_a$.

⁵We are free to take any of the five links to define the temporal direction, so we use the same $\hat{\mu}_3 = (0, 0, 0, 1)$ as MILC.

4.4. Parallel Pfaffian computation

Computing the Pfaffian of the fermion operator $\mathcal{D}$ is a notoriously hard problem, but is necessary to assess the potential sign problem in our system. In this section we present the algorithm we use to compute the Pfaffian in parallel, and discuss the checkpointing that allows these measurements to be split up into a sequence of computations.

The Pfaffian is formally

$$\text{pf } \mathcal{D} \equiv \int [d\Psi] \exp [-\Psi_i \mathcal{D}_{ij} \Psi_j] = \frac{1}{N! 2^N} \epsilon_{\alpha_0 \beta_0 \dots \alpha_{N-1} \beta_{N-1}} \mathcal{D}_{\alpha_0 \beta_0} \dots \mathcal{D}_{\alpha_{N-1} \beta_{N-1}}, \quad (26)$$

where the indices run over the lattice volume, the 16 fermion fields $\Psi = (\eta, \psi_a, \chi_{ab})$, and the N^2 generators of $U(N)$, so that each Ψ_i is a Grassmann variable. The Pfaffian is only well-defined if $\mathcal{D}$ is an antisymmetric matrix of even dimension, which is the case for our fermion operator. Up to a sign, about which we care a great deal, the Pfaffian is the square root of the determinant, $\det \mathcal{D} = [\text{pf } \mathcal{D}]^2$.

We can compute the Pfaffian by considering the analogue of LU factorization for antisymmetric matrices, $\mathcal{D} = L T L^T$, where L is lower-triangular and T is a trivial ($\text{pf } T = 1$) tridiagonal matrix composed of 2×2 blocks $T_{i, i \pm 1} = \pm 1$. As a consequence,

$$\text{pf } \mathcal{D} = \text{pf } [L T L^T] = (\det L) \text{pf } T = \det L = \prod_i L_{ii}. \quad (27)$$

There exist several algorithms to compute $\text{pf } \mathcal{D}$ based on this factorization [37, 38, 39, 40]. All scale as $\mathcal{O}(N_\Psi^3)$, where $N_\Psi = 16 N^2 L^3 N_t$ is the (even) number of elements in the fermionic fields for $U(N)$ gauge theory on an $L^3 \times N_t$ lattice.

The algorithms described in Refs. [39, 40] aim to minimize the number of floating-point operations required to determine the Pfaffian. Our main concern is to parallelize the computation, and in this context it is easier to consider the algorithms in Refs. [37, 38]. These employ procedures analogous to Gram-Schmidt orthogonalization in 2×2 blocks, based on applying the fermion operator itself, for which we already possess an efficient parallel implementation.

Specifically, the algorithm in the appendix of Ref. [38] constructs the upper-triangular matrix $Q = L^{-1}$ for which

$$Q \mathcal{D} Q^T = T \implies (\det Q) \text{pf } \mathcal{D} = 1 \implies \text{pf } \mathcal{D} = (\det Q)^{-1} = \left(\prod_i Q_{ii} \right)^{-1}. \quad (28)$$

The construction begins by initializing Q to the unit matrix, $q^{(i)} = e_i$, where $q^{(i)}$ denotes the i th column of Q and e_i is the i th basis vector. We cycle over pairs of columns, using the even column $q^{(i)}$ to normalize the odd column $q^{(i+1)} \rightarrow q^{(i+1)} / \left(q_a^{(i+1)} \mathcal{D}_{ab} q_b^{(i)} \right)$. We then propagate each pair of columns through all subsequent columns $j \geq i + 2$,

$$q^{(j)} \rightarrow q^{(j)} - q^{(i)} \left(q_a^{(i+1)} \mathcal{D}_{ab} q_b^{(j)} \right) + q^{(i+1)} \left(q_a^{(i)} \mathcal{D}_{ab} q_b^{(j)} \right) \quad (29)$$

in order to obtain $Q \mathcal{D} Q^T = T$. In components, this condition is

$$\left(q_a^{(i)} \mathcal{D}_{ab} q_b^{(j)} \right) = \begin{cases} 1 & \text{if } i = j + 1 \text{ and } j \text{ is even} \\ -1 & \text{if } i = j - 1 \text{ and } j \text{ is odd} \\ 0 & \text{otherwise} \end{cases}, \quad (30)$$

which we have explicitly checked for the matrix Q produced by this algorithm. (This test has since been removed from `SUSY LATTICE`, since it is incompatible with the optional checkpointing discussed below.) Note that (as in Eq. 26) there is no complex conjugation in any of the inner products $\left(q_a^{(i)} \mathcal{D}_{ab} q_b^{(j)} \right)$. We retain the indices a and b being summed to emphasize this.

Algorithm 1 restates the description in the previous paragraph more compactly. Obviously this al-

Algorithm 1 Construction of upper-triangular Q with $\det Q = \prod_i Q_{ii} = (\text{pf } \mathcal{D})^{-1}$

```

1   $Q \leftarrow \text{diag}\{1, 1, \dots, 1\}$ 
2
3  for  $i = 0$  to  $N - 2$  by 2 do
4       $q^{(i+1)} \leftarrow q^{(i+1)} / \left( q_a^{(i+1)} \mathcal{D}_{ab} q_b^{(i)} \right)$ 
5      for  $j = i + 2$  to  $N - 1$  by 1 do
6           $q^{(j)} \leftarrow q^{(j)} - q^{(i)} \left( q_a^{(i+1)} \mathcal{D}_{ab} q_b^{(j)} \right) + q^{(i+1)} \left( q_a^{(i)} \mathcal{D}_{ab} q_b^{(j)} \right)$ 
7      end for
8  end for

```

gorithm will break down if $\left(q_a^{(i+1)} \mathcal{D}_{ab} q_b^{(i)} \right) = 0$ for any even column $q^{(i)}$. Although this issue was not encountered in studies of the two-dimensional Wess–Zumino model considered by Ref. [38], it can occur for lattice $\mathcal{N} = 4$ SYM in four dimensions. In order to guarantee that none of these inner products vanish, we must choose an appropriate basis for the matrices $\mathcal{D}$ and Q , which is determined by the order we cycle over the fermionic field components. In particular, different η^i , ψ_a^i and χ_{ab}^i must be considered on each site, before looping over the $i = 0, \dots, N^2 - 1$ generators of $U(N)$. For the two-dimensional $\mathcal{N} = (2, 2)$ SYM theory discussed in the appendix, this step suffices to make all matrix elements non-zero. For $\mathcal{N} = 4$ SYM we also have to gather three χ_{ab}^i components from different lattice sites, instead of cycling over all the fields on a given site. These extra gathers are a fairly small addition to the cost of each application of the fermion operator $\mathcal{D}$.

Up to this point it has gone without saying that the sparse $N_\Psi \times N_\Psi$ matrix $\mathcal{D}$ is never explicitly constructed or stored in memory. Instead, its action on a fermion vector (the matrix–vector or “matvec” operation) is implemented by the `fermion_op` routine in `susy/utilities.c`. This is standard practice in lattice gauge theory. In contrast, we must allocate the upper-triangular matrix Q in order to compute it through Algorithm 1, even though we only care about the resulting diagonal components Q_{ii} . The two nested loops in Algorithm 1 require $(N_\Psi/2)^2$ matvec operations to compute $\mathcal{D}q^{(i)}$ and $\mathcal{D}q^{(j)}$ for steadily changing $q^{(j)}$. The cost of each matvec operation is proportional to N_Ψ , producing the expected $\mathcal{O}(N_\Psi^3)$ scaling of the overall computation.

For example, on a small $4^3 \times 6$ lattice with gauge group $U(2)$, $N_\Psi = 24,576$ and Algorithm 1 requires roughly 151 million matvec operations. We minimize the time to solution in this case by running on 16 of the 32 cores on a single node in the USQCD `ds` cluster at Fermilab, resulting in a local volume of $2^3 \times 3$ sites per core. This setup computes 217 matvecs per second on average, leading to a total runtime of about 8 days.

To avoid running such long jobs, we allow the Pfaffian computation to be checkpointed, by saving the partially-computed Q to disk and reloading it in a subsequent job. Checkpointing is controlled by the `ckpt_save` and `ckpt_load` options in Fig. 2. Since Q requires a large amount of memory (about 10 GB in the $4^3 \times 6$ case discussed above), we reduce disk space usage by discarding each column of Q after its diagonal element Q_{ii} has been determined through Algorithm 1. Both the known diagonal elements and remaining columns that still need to be considered (specified by `ckpt_save` and `ckpt_load`) are saved to disk in separate scratch files. In addition, we save only the non-zero elements of each column, which significantly reduces disk usage during the early stages of the computation. This special I/O is implemented in `generic/io_phase.c`.

5. Summary and discussion of potential future developments

In this paper we have presented new parallel software for lattice studies of four-dimensional $\mathcal{N} = 4$ SYM theory with gauge group $SU(N)$. This system has several unusual features compared to more familiar QCD-like lattice gauge theories, in particular the non-hypercubic A_4^* lattice with five basis vectors symmetrically spanning four dimensions, the Kähler–Dirac fermions spread out across both sites and links, and the complexified gauge links living in $\mathfrak{gl}(N, \mathbb{C})$, requiring approximate projection to the target $SU(N)$ gauge

theory. We reviewed the main elements of the system in Section 2, also discussing how the lattice action exactly preserves a subset of the supersymmetry algebra, up to two soft breaking terms that stabilize numerical computations.

`SUSY LATTICE` itself uses rational hybrid Monte Carlo importance sampling to compute various observables as discussed in Section 2.2. The present code allows a number of different $\mathcal{N} = 4$ SYM studies using larger lattices than could be considered by the serial code presented in Ref. [1], which `SUSY LATTICE` supersedes. In Section 3 we provided a quick-start guide explaining how to obtain, compile and run this publicly available software [49]. Since `SUSY LATTICE` is based on the MILC code for lattice QCD, it is relatively easy to modify, extend and improve. Simulations of $\mathcal{N} = 4$ SYM are in their infancy, and we encourage researchers to adapt the code as they wish, with the hope of making algorithmic improvements. To assist such efforts, in Section 4 we discussed some notable aspects of the documentation, including details of the adjoint representation for arbitrary N , data layout, internode communications, MD integrators, and the algorithm we use to compute the Pfaffian in parallel.

While all the applications discussed in this paper are stable and well tested, we are currently extending `SUSY LATTICE` to carry out further physics projects, including investigations of scalar correlation functions and of the RG blocking scheme recently proposed by Ref. [22]. Algorithmic improvements are also likely to be explored in the future. As one example, our MD integrator might be made more efficient by introducing additional time scales to be used by different terms in the partial fraction expansion (Eq. 19), along the lines of the algorithm described in Ref. [41]. At a purely practical level, we have observed uncomfortably slow thermalization in RHMC runs started from a `fresh` or `random` gauge configuration, especially for larger lattice volumes and larger values of N . Such systems may need to run for more than 2000 MD time units before data taking can begin, introducing a significant overhead cost. It should be possible to address this issue by designing better initial configurations.

Finally, `SUSY LATTICE` should serve as a convenient starting point for parallelized numerical studies of similar lattice theories with extended supersymmetry. Promising examples include two- and three-dimensional systems (potentially involving matter fields) with lattice actions constructed through topological twisting [1, 42, 43, 44]. We have already carried out such an extension for two-dimensional $\mathcal{N} = (2, 2)$ SYM, which we briefly describe in the appendix. Much of the parallel framework provided by `SUSY LATTICE` should even remain useful for lattice formulations based on different approaches, such as that considered in Ref. [45].

Acknowledgments

We thank our collaborators Simon Catterall, Poul Damgaard and Joel Giedt for many instructive discussions of $\mathcal{N} = 4$ SYM on and off the lattice, and Aarti Veernala for helping us compare results from `SUSY LATTICE` with those from the serial code [1]. The development of `SUSY LATTICE` is supported in part through the Scientific Discovery through Advanced Computing (SciDAC) program funded by the U.S. Department of Energy (DOE) Office of Science, Office of High Energy Physics under Award Number DE-SC0008669, in addition to support through Award Numbers DE-SC0010005 (TD) and DE-SC0009998 (DS). Numerical calculations were carried out on the HEP-TH cluster at the University of Colorado and on the DOE-funded USQCD facilities at Fermilab.

Appendix: Two-dimensional $\mathcal{N} = (2, 2)$ supersymmetric Yang–Mills theory

Although the main focus of `SUSY LATTICE` is four-dimensional $\mathcal{N} = 4$ SYM, we have also adapted the code to consider the similar $\mathcal{N} = (2, 2)$ SYM theory in two dimensions. We have implemented this extension as a separate directory tree mirroring that for $\mathcal{N} = 4$ SYM, rather than guaranteeing that the four-dimensional code will reduce to $\mathcal{N} = (2, 2)$ SYM upon setting the input parameters `ny = 1` and `nz = 1` (Fig. 2). This approach is much more computationally efficient, allowing for two dimensions to be dropped completely, at the cost of some code duplication.

The lattice action of two-dimensional $\mathcal{N} = (2, 2)$ SYM is quite similar to that of $\mathcal{N} = 4$ SYM:

$$\begin{aligned}
S &= S_{exact} + S_{stab} \\
S_{exact} &= \frac{NLN_t}{2\lambda} \sum_n a^4 \text{Tr} \left[-\overline{\mathcal{F}}_{ab}(n) \mathcal{F}_{ab}(n) + \frac{c_2}{2} \left(\overline{\mathcal{D}}_a^{(-)} \mathcal{U}_a(n) \right)^2 - \chi_{ab}(n) \mathcal{D}_{[a}^{(+)} \psi_{b]}(n) - \eta(n) \overline{\mathcal{D}}_a^{(-)} \psi_a(n) \right] \\
S_{stab} &= \frac{NLN_t}{2\lambda} \mu^2 \sum_{n, c} a^4 \left(\frac{1}{N} \text{Tr} [\mathcal{U}_c(n) \overline{\mathcal{U}}_c(n)] - 1 \right)^2 + \kappa \sum_{\mathcal{P}} a^4 |\det \mathcal{P} - 1|^2,
\end{aligned} \tag{31}$$

which nearly matches Eqs. 4–7. The main difference is that there are now only two gauge links at each site (forming the familiar square lattice), which forbids the S_{closed} term in Eq. 6. The Kähler–Dirac multiplet $\Psi = (\eta, \psi_a, \chi_{ab})$ involves only $2^d = 4$ rather than 16 fermion fields. Finally, the 't Hooft coupling λ is now dimensionful, and has to be divided by the lattice volume $L \times N_t$ [1].

The lower dimensionality and reduced field content of $\mathcal{N} = (2, 2)$ SYM compared to $\mathcal{N} = 4$ SYM makes the system more tractable numerically. Parallelization only becomes crucial for lattices larger than 64^2 , for which the overall factor of $\frac{NLN_t}{2\lambda}$ tends to be rather large. Since this system has already been fairly well explored using serial codes [46, 1, 19, 47, 48], we are not currently carrying out any $\mathcal{N} = (2, 2)$ SYM physics projects using SUSY LATTICE. Instead, our main interest is using this simpler and cheaper system as a testing ground for observables and algorithms ultimately intended for our $\mathcal{N} = 4$ investigations. For example, it was helpful to debug our parallel Pfaffian algorithm (Section 4.4) in $\mathcal{N} = (2, 2)$ SYM before setting up this difficult computation in four dimensions.

Since the 2dSYM portion of SUSY LATTICE was adapted from the $\mathcal{N} = 4$ SYM code discussed in the body of this paper, much of the information presented is directly applicable to this case. The executables, input parameters and compile-time options are the same as discussed in Section 3; the directory structure and most low-level routines are the same as discussed in Section 4. The main differences are the straightforward square lattice and the absence of the $\mathcal{Q}$ -closed term, which greatly simplify the lattice action, forces and internode communications.

References

- [1] S. Catterall and A. Joseph, “An object oriented code for simulating supersymmetric Yang–Mills theories”, *Comput. Phys. Commun.* **183** (2012) 1336–1353 [[arXiv:1108.1503](#)].
- [2] S. Catterall, D. B. Kaplan and M. Unsal, “Exact lattice supersymmetry”, *Phys. Rept.* **484** (2009) 71–130 [[arXiv:0903.4881](#)].
- [3] J. Terning, *Modern Supersymmetry: Dynamics and Duality*. Oxford University Press, Oxford, 2006.
- [4] P. Dondi and H. Nicolai, “Lattice Supersymmetry”, *Nuovo Cim.* **A41** (1977) 1.
- [5] M. Kato, M. Sakamoto and H. So, “Taming the Leibniz Rule on the Lattice”, *JHEP* **0805** (2008) 057 [[arXiv:0803.3121](#)].
- [6] G. Bergner, I. Montvay, G. Münster, U. D. Özgürel and D. Sandbrink, “Towards the spectrum of low-lying particles in supersymmetric Yang–Mills theory”, *JHEP* **1311** (2013) 061 [[arXiv:1304.2168](#)].
- [7] J. Giedt, R. Brower, S. Catterall, G. T. Fleming and P. Vranas, “Lattice super-Yang–Mills using domain wall fermions in the chiral limit”, *Phys. Rev.* **D79** (2009) 025015 [[arXiv:0810.5746](#)].
- [8] M. G. Endres, “Dynamical simulation of $\mathcal{N} = 1$ supersymmetric Yang–Mills theory with domain wall fermions”, *Phys. Rev.* **D79** (2009) 094503 [[arXiv:0902.4267](#)].

- [9] JLQCD Collaboration: S. Kim, H. Fukaya, S. Hashimoto, H. Matsufuru, J. Nishimura and T. Onogi, “Lattice study of 4d $\mathcal{N} = 1$ super Yang–Mills theory with dynamical overlap gluino”, *PoS Lattice* **2011** (2011) 069 [[arXiv:1111.2180](#)].
- [10] T. Ishii, G. Ishiki, S. Shimasaki and A. Tsuchiya, “ $\mathcal{N} = 4$ Super Yang–Mills from the Plane Wave Matrix Model”, *Phys. Rev.* **D78** (2008) 106001 [[arXiv:0807.2352](#)].
- [11] G. Ishiki, S.-W. Kim, J. Nishimura and A. Tsuchiya, “Deconfinement phase transition in $\mathcal{N} = 4$ super Yang–Mills theory on $R \times S^3$ from supersymmetric matrix quantum mechanics”, *Phys. Rev. Lett.* **102** (2009) 111601 [[arXiv:0810.2884](#)].
- [12] G. Ishiki, S.-W. Kim, J. Nishimura and A. Tsuchiya, “Testing a novel large- N reduction for $\mathcal{N} = 4$ super Yang–Mills theory on $R \times S^3$ ”, *JHEP* **0909** (2009) 029 [[arXiv:0907.1488](#)].
- [13] M. Hanada, S. Matsuura and F. Sugino, “Two-dimensional lattice for four-dimensional $\mathcal{N} = 4$ supersymmetric Yang–Mills”, *Prog. Theor. Phys.* **126** (2011) 597–611 [[arXiv:1004.5513](#)].
- [14] M. Honda, G. Ishiki, J. Nishimura and A. Tsuchiya, “Testing the AdS/CFT correspondence by Monte Carlo calculation of BPS and non-BPS Wilson loops in 4d $\mathcal{N} = 4$ super-Yang–Mills theory”, *PoS Lattice* **2011** (2011) 244 [[arXiv:1112.4274](#)].
- [15] M. Honda, G. Ishiki, S.-W. Kim, J. Nishimura and A. Tsuchiya, “Direct test of the AdS/CFT correspondence by Monte Carlo studies of $\mathcal{N} = 4$ super Yang–Mills theory”, *JHEP* **1311** (2013) 200 [[arXiv:1308.3525](#)].
- [16] M. Hanada, Y. Hyakutake, G. Ishiki and J. Nishimura, “Holographic description of quantum black hole on a computer”, *Science* **344** (2014) 882–885 [[arXiv:1311.5607](#)].
- [17] D. B. Kaplan and M. Unsal, “A Euclidean lattice construction of supersymmetric Yang–Mills theories with sixteen supercharges”, *JHEP* **0509** (2005) 042 [[hep-lat/0503039](#)].
- [18] M. Ünsal, “Twisted supersymmetric gauge theories and orbifold lattices”, *JHEP* **0610** (2006) 089 [[hep-th/0603046](#)].
- [19] S. Catterall, R. Galvez, A. Joseph and D. Mehta, “On the sign problem in 2D lattice super Yang–Mills”, *JHEP* **1201** (2012) 108 [[arXiv:1112.3588](#)].
- [20] S. Catterall, P. H. Damgaard, T. Degrand, R. Galvez and D. Mehta, “Phase Structure of Lattice $\mathcal{N} = 4$ Super Yang–Mills”, *JHEP* **1211** (2012) 072 [[arXiv:1209.5285](#)].
- [21] S. Catterall, D. Schaich, P. H. Damgaard, T. DeGrand and J. Giedt, “ $\mathcal{N} = 4$ supersymmetry on a space-time lattice”, *Phys. Rev.* **D90** (2014) 065013 [[arXiv:1405.0644](#)].
- [22] S. Catterall and J. Giedt, “Real space renormalization group for twisted lattice $\mathcal{N} = 4$ super Yang–Mills”, *JHEP* (in press, 2014) [[arXiv:1408.7067](#)].
- [23] MILC Collaboration: A. Bazavov, C. Bernard, T. Burch, T. DeGrand, C. DeTar, J. Foley, S. Gottlieb, U. Heller, J. Hetrick, L. Levkova, C. McNeile, K. Orginos, J. Osborn, K. Rummukainen, B. Sugar and D. Toussaint, “The MILC Code” (2011), <http://www.physics.utah.edu/~detar/milc/>.
- [24] N. Marcus, “The other topological twisting of $\mathcal{N} = 4$ Yang–Mills”, *Nucl. Phys.* **B452** (1995) 331–345 [[hep-th/9506002](#)].
- [25] A. Kapustin and E. Witten, “Electric-Magnetic Duality And The Geometric Langlands Program”, *Commun. Num. Theor. Phys.* **1** (2007) 1–236 [[hep-th/0604151](#)].

- [26] S. Catterall, “From Twisted Supersymmetry to Orbifold Lattices”, *JHEP* **0801** (2008) 048 [[arXiv:0712.2532](#)].
- [27] P. H. Damgaard and S. Matsuura, “Geometry of Orbifolded Supersymmetric Lattice Gauge Theories”, *Phys. Lett.* **B661** (2008) 52–56 [[arXiv:0801.2936](#)].
- [28] M. A. Clark and A. D. Kennedy, “Accelerating dynamical fermion computations using the rational hybrid Monte Carlo (RHMC) algorithm with multiple pseudofermion fields”, *Phys. Rev. Lett.* **98** (2007) 051601 [[hep-lat/0608015](#)].
- [29] S. Duane, A. Kennedy, B. Pendleton and D. Roweth, “Hybrid Monte Carlo”, *Phys. Lett.* **B195** (1987) 216–222.
- [30] B. Jegerlehner, “Krylov space solvers for shifted linear systems”, [hep-lat/9612014](#).
- [31] M. A. Clark and A. D. Kennedy, “AlgRemez”, <http://github.com/mikeaclark/AlgRemez> (2005).
- [32] A. Stathopoulos and J. R. McCombs, “PRIMME: preconditioned iterative multimethod eigensolver—methods and software description”, *ACM Trans. Math. Softw.* **37** (2010) 21.
- [33] R. Galvez and G. van Anders, “Accelerating the solution of families of shifted linear systems with CUDA”, [arXiv:1102.2143](#).
- [34] I. P. Omelyan, I. M. Mryglod and R. Folk, “Optimized Verlet-like algorithms for molecular dynamics simulations”, *Phys. Rev. E* **65** (2002) 056706.
- [35] I. P. Omelyan, I. M. Mryglod and R. Folk, “Symplectic analytically integrable decomposition algorithms: classification, derivation, and application to molecular dynamics, quantum and celestial mechanics simulations”, *Comput. Phys. Commun.* **151** (2003) 272–314.
- [36] C. Urbach, K. Jansen, A. Shindler and U. Wenger, “HMC algorithm with multiple time scale integration and mass preconditioning”, *Comput. Phys. Commun.* **174** (2006) 87–98 [[hep-lat/0506011](#)].
- [37] DESY–Munster Collaboration: I. Campos, A. Feo, R. Kirchner, S. Luckmann, I. Montvay, G. Münster, K. Spanderen and J. Westphalen, “Monte Carlo simulation of SU(2) Yang–Mills theory with light gluinos”, *Eur. Phys. J.* **C11** (1999) 507–527 [[hep-lat/9903014](#)].
- [38] S. Catterall and S. Karamov, “A Lattice study of the two-dimensional Wess-Zumino model”, *Phys. Rev.* **D68** (2003) 014503 [[hep-lat/0305002](#)].
- [39] M. Wimmer, “Algorithm 923: Efficient Numerical Computation of the Pfaffian for Dense and Banded Skew-Symmetric Matrices”, *ACM Trans. Math. Softw.* **38** (2012) 30 [[arXiv:1102.3440](#)].
- [40] J. Rubow and U. Wolff, “A Factorization algorithm to compute Pfaffians”, *Comput. Phys. Commun.* **182** (2011) 2530–2532 [[arXiv:1102.3576](#)].
- [41] M. Lüscher and S. Schaefer, “openQCD” (2012), <http://luscher.web.cern.ch/luscher/openQCD/>.
- [42] A. Joseph, “Lattice formulation of three-dimensional $\mathcal{N} = 4$ gauge theory with fundamental matter fields”, *JHEP* **1309** (2013) 046 [[arXiv:1307.3281](#)].
- [43] A. Joseph, “Supersymmetric quiver gauge theories on the lattice”, *JHEP* **1401** (2014) 093 [[arXiv:1311.5111](#)].

- [44] A. Joseph, “Two-dimensional $\mathcal{N} = (2, 2)$ lattice gauge theories with matter in higher representations”, *JHEP* **1407** (2014) 067 [[arXiv:1403.4390](#)].
- [45] S. Matsuura and F. Sugino, “Lattice formulation for 2d $\mathcal{N} = (2, 2)$, $(4, 4)$ super Yang–Mills theories without admissibility conditions”, *JHEP* **1404** (2014) 088 [[arXiv:1402.0952](#)].
- [46] M. Hanada and I. Kanamori, “Absence of sign problem in two-dimensional $\mathcal{N} = (2, 2)$ super Yang–Mills on lattice”, *JHEP* **1101** (2011) 058 [[arXiv:1010.2948](#)].
- [47] D. Mehta, S. Catterall, R. Galvez and A. Joseph, “Supersymmetric gauge theories on the lattice: Pfaffian phases and the Neuberger 0/0 problem”, *PoS Lattice 2011* (2011) 078 [[arXiv:1112.5413](#)].
- [48] R. Galvez, S. Catterall, A. Joseph and D. Mehta, “Investigating the sign problem for two-dimensional $\mathcal{N} = (2, 2)$ and $\mathcal{N} = (8, 8)$ lattice super Yang–Mills theories”, *PoS Lattice 2011* (2011) 064 [[arXiv:1201.1924](#)].
- [49] <http://github.com/daschaich/susy>