PARTICLE VORTEX DUALITY: A REVIEW AND NUMERICAL INVESTIGATION

Abstract. We explore particle vortex duality, i.e. the idea that the 3D XY-Model and Abelian-Higgs Models are dual, with vortices in one model identified with particles in the other model. First we review the models and give some heuristic theoretical evidence for believing they describe the same physics. This culminates in a table to allow easy comparison between their phases. Next we discuss how we can use numerics to find the critical exponents of the models and check that they are the same. The Ising model is used to provide a natural introduction to numerical simulations, before moving on to simulating a model in the same universality class as the Abelian-Higgs Model. Finally we conclude with a demonstration that the Abelian-Higgs Model undergoes an inverted XY-Model type transition.

Contents

1. Introduction

Particle vortex (PV) duality was first proposed by Peskin [9] in the late '70s. The idea behind the duality is that the XY-Model and the Abelian-Higgs Model in fact represent the same underlying physics, up to some relabeling. In Statistical Field Theory we called this, perhaps less imposingly, universality. The idea here is the same; two seemingly different physical phenomena or theories turn out to be identical when we look at their field theoretic description. In particular, PV duality holds that the "relabeling" needed to turn the XY-Model to the Abelian-Higgs Model is particle $\leftrightarrow$ vortex. In other words, that which we choose to call a particle in the XY-Model is a vortex of the Abelian-Higgs Model, and vice versa (where here a vortex is defined by the winding of some local order

Date: May 4th, 2018.

parameter). In practice the PV duality means if we were interested in describing the vortices of the XY-Model we might look to the Abelian-Higgs Model, since the vortex degrees of freedom of a theory often lurk just out of sight. In this respect it seems clear that the PV duality has a lot of practical significance. But is it true, and if so how would we tell?

Just as with many other pieces of field theory, it is exceptionally difficult if not impossible to prove PV duality explicitly by showing that the partition function of the XY-Model is equal to the partition function of the Abelian-Higgs Model in some appropriate regime. We simply don't have many non-perturbative tools to work with. Instead the PV duality, and most other dualities, was "proven" with a lot of suggestive evidence. In the case of the PV duality this evidence took the form of: nearly rigorous arguments with some heuristic reasoning [9], the fact that duality can be shown explicitly in supersymmetric theories [3], numerical evidence [2], and a large number of systems where there is an underlying physical reason to expect the duality to hold, e.g. topological insulators [5].

In this essay we only have time to review one of these methods, and even then only to scratch the surface of what's been done. We'll choose to focus on the numerical evidence for the PV duality, but first we will give a review of the underlying models and some physical systems they describe. We'll use these physics to motivate the numerical results.

2. The Models

2.1. **The XY-Model: Review and Critical Exponents.** The XY-Model in $2+1$ dimensions has action given by

$$S_{XY} = \int \mathrm{d}^3x \; |\partial_\mu \phi|^2 - a|\phi|^2 - \frac{b}{2}|\phi|^4 + \cdots ,$$

where ϕ is a complex scalar field. In its nonrelativistic form it is used to describe a nonrelativistic superfluid. In this form the MFT description of the superfluid is known as the Gross-Pitaevskii Equation (GPE) and it is often used to numerically simulate nonrelativistic superfluids, producing images like Figure 1

Figure 1 provides a good introduction to what can happen in the XY-Model, but shouldn't be taken too seriously since it was produced with MFT instead of the full theory and describes the nonrelativistic theory. In particular, the equilibrium statistical mechanics of the nonrelativistic theory will be very different, since they will occur in two spatial dimensions with no time dependence. Thus they will be subject to the Kosterlitz-Thouless transition, something that the XY-Model in 2+1 dimensions is not [12].

However, provided we understand that the GPE is a MFT and don't take it too seriously, it will give us some insight into the XY-Model's basic behavior in at least one of its phases. In order to connect the complex scalar ϕ in our definition of the XY-Model to

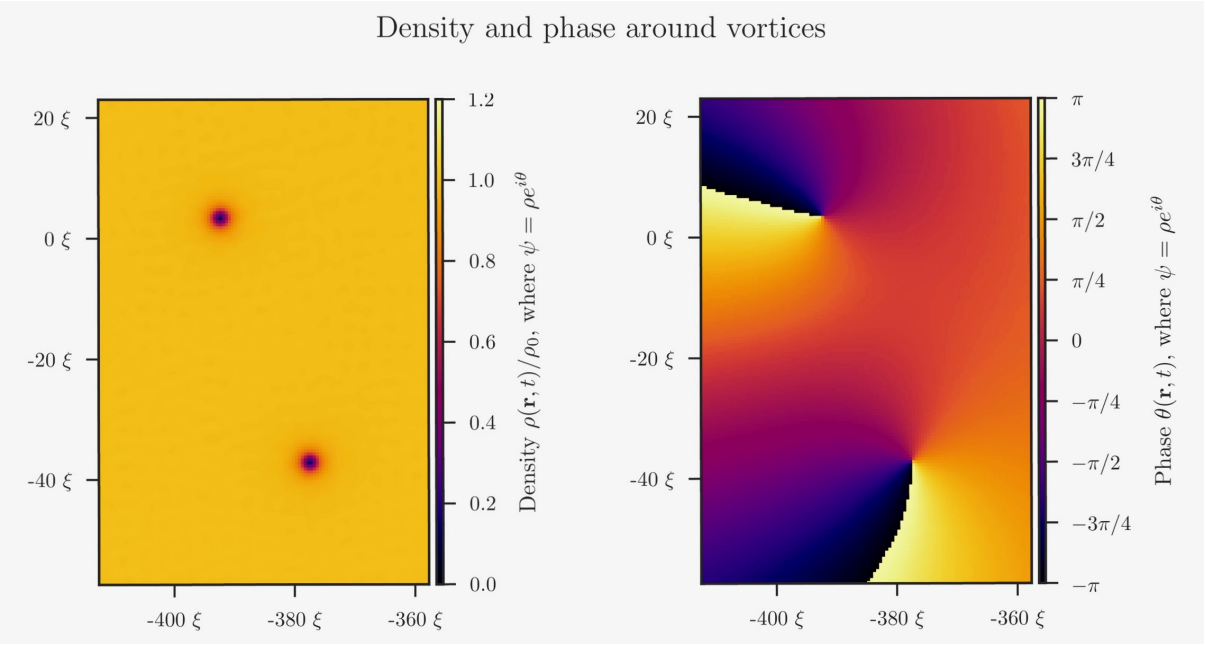

FIGURE 1. Two vortices of charge 2π orbiting one another. The vortices are part of a larger system that was previously simulated by the author to study vortex nucleation from a hydrofoil. They were simulated using a split-step Fourier method to solve the GPE. Note that the length scale is in healing lengths ξ, which for our purposes is just the typical size of a vortex.

the density and phase of the superfluid, we write $\phi = \rho e^{i\theta}$. Then ρ is the density of the superfluid and θ is its phase, as shown in the MFT Figure 1. The figure demonstrates: 1) that the fluctuations of ρ can occur around a nonzero value ρ_0 and 2) that we can have vortex-like excitations. As mentioned in the introduction a vortex is just the winding of some local order parameter, which in this case is θ. Indeed if we travel counterclockwise around the two singularities in θ we notice that the phase will increment by $+2\pi$; in general this increment will be some integer, called the "winding number", multiple of 2π. Why is the winding number quantized? We will investigate further.

$a < 0$:

The phase pictured in Figure 1 will arise when $a < 0$, i.e. when the system acquires a nonzero ground state. To see this consider that

$$\mathcal{H}_{XY} = |\partial_t \phi|^2 + |\vec{\nabla}\phi|^2 + a|\phi|^2 + \frac{b}{2}|\phi|^4 + \cdots .$$

Thus if $a < 0$ the potential $a|\phi|^2 + b|\phi|^4/2 = a\rho^2 + b\rho^4/2$ will be a Mexican hat potential. It will be minimized when

$$0 = \frac{\partial}{\partial \rho}\left(a\rho^2 + \frac{b}{2}\rho^4\right) = 2a\rho + 2b\rho^3 \quad \Longrightarrow \quad \rho_0 = \sqrt{\frac{-a}{b}}.$$

As an aside, we have required $b > 0$ in order that $\mathcal{H}_{XY}$ doesn't reach arbitrarily low energies by choosing fields with larger and larger values of $|\phi|^2$, and so $\rho_0 \in \mathbb{R}$. But then since $|\phi|$ acquires a vacuum expectation value of ρ_0, we will need to expand the vacuum as $\phi = (\tilde{\rho} + \rho_0)e^{i\theta}$, where $\tilde{\rho}$ will in fact be the quantity that remains small. Doing this expansion, we note that up to a constant term

$$\begin{aligned}
\mathcal{L}_{XY} =&(\partial_\mu\phi)(\partial^\mu\phi^*) - a|\phi|^2 - \frac{b}{2}|\phi|^4 + \cdots \\
=&(\partial_\mu\rho e^{i\theta})(\partial^\mu\rho e^{-i\theta}) - a\rho^2 - \frac{b}{2}\rho^4 + \cdots \text{ substituting } \psi = \rho e^{i\theta} \\
=&(\partial_\mu\rho)^2 + i[\rho(\partial_\mu\theta)(\partial^\mu\rho) - \rho(\partial^\mu\theta)(\partial_\mu\rho)] + \rho^2(\partial_\mu\theta)^2 - a\rho^2 - \frac{b}{2}\rho^4 + \cdots \\
=&(\partial_\mu\rho)^2 + \rho^2(\partial_\mu\theta)^2 - a\rho^2 - \frac{b}{2}\rho^4 + \cdots \\
\implies \mathcal{L}_{XY} =&(\partial_\mu\tilde{\rho})^2 + (\tilde{\rho} + \rho_0)^2(\partial_\mu\theta)^2 + 2a\tilde{\rho}^2 - 2b\rho_0\tilde{\rho}^3 - \frac{b}{2}\tilde{\rho}^4 + \cdots \text{ substituting } \rho = \rho_0 + \tilde{\rho}.
\end{aligned}$$

Now the mass term for $\tilde{\rho}^2$ has a negative coefficient, so the excitation governed by $\tilde{\rho}$ will have positive mass. This is as we would expect if we are expanding around the correct ground state.

The original theory had a $U(1)$ global symmetry $\phi \to e^{i\alpha}\phi$, where α is a constant. This $U(1)$ global symmetry had a current

$$\begin{aligned}
\alpha j^\mu =&\frac{\delta\mathcal{L}_{XY}}{\delta(\partial_\mu\phi)}\delta\phi + \frac{\delta\mathcal{L}_{XY}}{\delta(\partial_\mu\phi^*)}\delta\phi^* \\
=&(\partial^\mu\phi^*)(i\alpha\phi) + (\partial^\mu\phi)(-i\alpha\phi^*) \\
j^\mu =&i(\phi\partial^\mu\phi^* - \phi^*\partial^\mu\phi) \\
=&i(\rho e^{i\theta}\partial^\mu(\rho e^{-i\theta}) - \rho e^{-i\theta}\partial^\mu(\rho e^{i\theta})) \\
=&(\rho^2\partial^\mu\theta + \rho^2\partial^\mu\theta) = 2\rho^2\partial^\mu\theta.
\end{aligned}$$

Of course it makes sense that $j^\mu \propto \partial^\mu\theta$ since our global symmetry is just $\theta \to \theta + \alpha$. However, we break this rotational symmetry when we choose a vacuum expectation value for the field ϕ, i.e. $\langle\phi\rangle = \rho_0 e^{i\theta_0}$ for some $\theta_0 \in \mathbb{R}$. Thus, as discussed in Statistical Field Theory (SFT) we should expect to get a single Goldstone boson. This is exactly the field θ. We see this by noting that the Lagrangian expanded around the correct ground state has no terms of the form θ^2 and therefore the θ excitations will be massless excitations in our theory. We've already seen that when $a < 0$ the fluctuations of ρ will occur around a nonzero value ρ_0, our first observation about the phase shown in Figure 1. It turns out that the existence of the Goldstone boson θ will give rise to our second observation, that of vortex-like excitations.

We will now investigate these vortex excitations a little more. We note that θ need not be single valued. Indeed $\theta \to \theta + 2\pi$ will leave ϕ unchanged since $\phi = \rho e^{i\phi}$. This can be seen in Figure 1 by observing what appears to be dislocations associated with

each vortex. Along this line of dislocations one side has $\theta = -\pi$, while the other side has $\theta = \pi$. Since the sides differ by 2π they in fact give the same value of ϕ, and there is no physical dislocation in the wavefunction. However, these dislocations do allow the phase to increment by an integer multiple of 2π as we travel around a vortex excitation. In other words,

$$\oint \vec{\nabla}\theta \cdot \mathrm{d}\vec{x} = \Delta\theta = 2\pi n,$$

where n counts the number of times the phase winds around the vortex, or alternatively the number of dislocation lines emenating from a vortex.

There is one important caveat to θ not being single valued, however. We obviously don't want our physical field $\phi = \rho e^{i\theta}$ to be multivalued. So then suppose that $\theta = \theta_0$ and $\theta = \theta_0 + \sigma$ at some point. The fact that ϕ is single valued means that $\rho = \rho e^{i\sigma}$. If $\rho \neq 0$ this means that $\sigma = 2\pi n$ for $n \in \mathbb{Z}$. This is fine along a dislocation, as we have seen $\sigma = 2\pi$ at a dislocation line. But what about at the center of our vortex excitation? By looking at the right hand side of Figure 1 we note that at the center θ takes all values in $[-\pi, \pi]$. Then following through our logic above, this means that $\rho = 0$ at the center of a vortex if we want to keep ϕ single valued. Indeed, looking at the left hand side of Figure 1 we see that the density smoothly interpolates from $\rho = \rho_0$ outside of the vortex to a value of $\rho = 0$ at the center of the vortex. Suppose instead that ρ did not acquire this vacuum expectation value of ρ_0? Then there would be nothing to interpolate to zero from, and it would be much harder to spot the vortices. This provides the biggest distinction from the $a > 0$ case.

Before we move on to the $a > 0$ case we'd like to calculate the energy of some simple vortex configurations, since the vortices are the most interesting feature of this phase. We'll follow the discussion of Section 4.4.1 of [12]. Take the system to be static for the moment and to contain N vortices. Since we know from discussions in QFT that $T^{00} = (\partial_t \phi)^2 + |\vec{\nabla}\phi|^2 + a|\phi|^2 + b|\phi|^4/2 + \cdots$ for this model, then

$$\begin{aligned} E &= \int \mathrm{d}^2x \ |\vec{\nabla}\phi|^2 + a|\phi|^2 + \frac{b}{2}|\phi|^4 + \cdots \text{ neglecting time derivatives} \\ &= \int \mathrm{d}^2x \ \rho^2|\vec{\nabla}\theta|^2 + |\vec{\nabla}\rho|^2 + a\rho^2 + \frac{b}{2}\rho^4 + \cdots \\ &= \int \mathrm{d}^2x \ \left[\rho_0^2|\vec{\nabla}\theta|^2 + O(\tilde{\rho})\right]. \end{aligned}$$

If we look at Figure 1 we can see that the density very quickly approaches ρ_0 some distance l outside the center of the vortex. This means that we should be free to neglect the $\tilde{\rho}$ terms in the above integral, provided we lie outside the cores of the vortices. Then

$$E = E_{\text{cores}} + \rho_0^2 \int_{\text{outside cores}} \mathrm{d}^2x \ |\vec{\nabla}\theta|^2.$$

If we now take this to be the energy for this state, then the static field equation θ must satisfy in the MFT will be $\nabla^2\theta = 0$. Consider a single vortex, i.e. $N = 1$. From Figure

1 we might guess that $\theta(r,\vartheta) = \vartheta$ describes the winding of the phase, as it appears to increment continuously around the vortex. Indeed $\theta(r,\vartheta) = \vartheta$ satisfies $\nabla^2\theta = 0$. A vortex with winding number $n > 1$ will be described by $\theta(r,\vartheta) = n\vartheta$. But then $\vec{\nabla}\theta = (n/r)\vec{e}_\vartheta$ and the energy becomes

$$E = E_{\text{core}} + 2\pi\rho_0^2 \int_{r>l} r\mathrm{d}r \ \frac{n^2}{r^2} = E_{\text{core}} + 2\pi\rho_0^2 n^2 \ln\left(\frac{L}{l}\right),$$

where L is a cutoff to the system size. Thus the energy of a single vortex diverges logarithmically in the system size, i.e. the vortices are gapped. Also note that it is energetically more favorable to have $|n|$ vortices of sign $\mathrm{sgn}(n)$ than one vortex of sign n due to the fact that $E \propto n^2$.

It turns out that not only does the energy of a vortex diverge logarithmically, but that the interaction between vortices is also logarithmic. To make this clear, we will make some new definitions to rephrase the problem. Define $\vec{v} = \vec{\nabla}\theta$, it will have equation of motion $\vec{\nabla}\cdot\vec{v} = \nabla^2\theta = 0$ from the static MFT. Though the divergence of $\vec{v}$ is zero its curl will not be in general. For example let's consider the above case of $N = 1$, so that $\vec{v} = \vec{\nabla}\theta = (n/r)\vec{e}_\vartheta$. Then $\vec{\nabla}\times\vec{v} = r^{-1}\partial_r(nr/r)\hat{\vec{z}} := f(\vec{r})\hat{\vec{z}}$. This will be zero everywhere except the origin. To find its value at the origin we consider the following integral over some surface S containing the vortex at the origin

$$\begin{aligned}\int_S \mathrm{d}^2x \ f(\vec{r}) &= \int_S \mathrm{d}^2x \ \vec{z}\cdot(\vec{\nabla}\times\vec{v}) \\ &= \oint_C \mathrm{d}\vec{x}\cdot\vec{v} \text{ by Stokes} \\ &= \oint_C \mathrm{d}\vec{x}\cdot\vec{\nabla}\theta.\end{aligned}$$

But this integral is given by $2\pi n$ by definition of the vortex winding number. Since $f(\vec{x}) = 0$ except at the origin, we conclude that $f(\vec{r}) = 2\pi n\delta^{(2)}(\vec{r})$.

Suppose now that we generalize to the case of N vortices with positions $\vec{r}_i$ and winding numbers n_i, $1 \le i \le N$. We will still have $\vec{\nabla}\cdot\vec{v} = 0$, but now $\vec{\nabla}\times\vec{v} = 2\pi\hat{\vec{z}}\sum_i n_i\delta^{(2)}(\vec{r}-\vec{r}_i)$. This generalization should be obvious from the above example. Let's put this in a more familiar form by defining a vector $\vec{E}$ such that $E_i = \epsilon^{ij}v_j$. The effect of the ϵ^{ij} will just be to change the divergence to a curl and vice versa, since $\vec{\nabla}\times\vec{v} = \hat{\vec{z}}\epsilon^{ij}\nabla_i v_j$ for our $2d$ vectors. Then $\vec{E}$'s defining equations become

$$\vec{\nabla}\times\vec{E} = \vec{0} \ \text{ and } \ \vec{\nabla}\cdot\vec{E} = 2\pi\sum_i n_i\delta^{(2)}(\vec{r}-\vec{r}_i).$$

But these are just the electrostatic equations in two dimensions for a system of N point charges! Thus $\vec{E} = -\vec{\nabla}\chi$ for some potential χ.

From the above we see that $\nabla^2\chi$ must give rise to delta functions. If we were in three dimensions then $\chi \sim 1/r$. However, since we are in two dimensions we will need

$\chi \sim \log(r)$ to give rise to this kind of behavior. In particular

$$\chi(\vec{r}) = \sum_i n_i \log\left(\frac{|\vec{r} - \vec{r}_i|}{l}\right).$$

Now let's use this to find the interaction energy between vortices.

$$\begin{aligned}
E - E_{\text{cores}} =& \rho_0^2 \int_{\text{outside cores}} \mathrm{d}^2x \ |\vec{\nabla}\theta|^2 \\
=& \rho_0^2 \int \mathrm{d}^2x \ |\vec{v}|^2 \\
=& \rho_0^2 \int \mathrm{d}^2x \ |\vec{E}|^2 \text{ since } \epsilon^{ij}\epsilon_{ik} = \delta^j_k \\
=& \rho_0^2 \int \mathrm{d}^2x \ |\vec{\nabla}\chi|^2 \\
=& \rho_0^2 \int \mathrm{d}^2x \ \chi(\vec{\nabla} \cdot \vec{E}) \text{ integrating by parts} \\
=& 2\pi\rho_0^2 \sum_i n_i \chi(\vec{r}_i) \\
=& 2\pi\rho_0^2 \sum_{i \neq j} n_i n_j \log\left(\frac{|\vec{r}_j - \vec{r}_i|}{l}\right),
\end{aligned}$$

where we've dropped the $i = j$ terms, since these represent the energy of the core. Finally we do see the Coulomb force in two dimensions, as we expected from the electrostatic formulation. There is a logarithmically increasing attractive force between vortices of opposite sign. The vortices in this model are thus sometimes said to be "logarithmically confined".

$a > 0$:

If $a > 0$ the potential $a\rho^2 + b\rho^4/2$ will have a minimum at $\rho = 0$. Thus ρ will not acquire a vacuum expectation value. But this means the $U(1)$ symmetry will remain unbroken. In particular there will be no Goldstone bosons, as mentioned above. Without vortices this phase is rather uninteresting. However, we do note that since the $U(1)$ global symmetry is unbroken then ϕ will remain massive in this phase.

$a = 0$:

Provided we tune the other terms in our system properly, we expect to encounter a phase transition at this point. This is where ρ will begin to acquire a vacuum excitation value, and where we will start to notice vortices in our theory. Thus our system goes through a critical point at $a = 0$, all other terms being tuned appropriately.

Critical Exponents:

A lot of numerical work has been done by Hasenbush et al. [1] studying the critical behavior of the three-dimensional XY universality class. Over the course of several years they studied a number of systems that fall in the 3D XY universality class, both theoretically and numerically. Their studies have found that near the critical point, for our concerns $a = 0$ as discussed above, the critical exponents of the 3D XY-Model look like [1]:

$$\begin{aligned}
\nu &= 0.67155(27) \\
\eta &= 0.0380(4) \\
\alpha &= -\,0.0146(8) \\
\gamma &= 1.3177(5) \\
\beta &= 0.3485(2) \\
\delta &= 4.780(2),
\end{aligned}$$

where these critical exponents and their physical significance were all discussed in the Statistical Field Theory lectures (and of course we only needed ν and η to derive the rest). Hasenbusch et al. noted that their theoretical and numerical value of α differed significantly from the experimental value of $\alpha = -0.01056(38)$, obtained by studying the phase transition of systems that are expected to lie in the 3D XY universality class. For our purposes these are close enough. We expect the 3D XY-Model to have α that is negative and around -0.01, fairly close to zero.

There are several things we might observe about these exponents. Perhaps the most surprising is that α is less than zero, unlike almost all other values of α we encountered in SFT. But perhaps more fascinating than the values of the exponents is the way in which they were found. The critical exponents we met in SFT, the above included, were with reference to temperature. How did we get a temperature out of S_{XY}? As far as we could tell the parameter that was controlling our phase transition was a. How is this related to a critical temperature? Further, how did [1] even go about finding a system that gave them the correct exponents, and how do they verify that the system they are simulating is in the 3D XY-Model universality class? If they simulated a system on a lattice, how can this be derived from the continuum S_{XY} we considered? We will endeavor to at least touch on all of these questions when we discuss numerics. For now we continue to give heuristics supporting why we should expect the continuum XY-Model and the Abelian-Higgs Model to exhibit the same (up to a reversal in the sign of a) transition. It is when we decide to move beyond heuristics that we will need to return to numerics, and will be able to address some of these questions.

2.2. **The Abelian-Higgs Model.** The Abelian-Higgs model in $2+1$ dimensions has action given by

$$S_{AH} = \int \mathrm{d}^3x \ \left[-\frac{1}{4} f_{\mu\nu} f^{\mu\nu} + |D_\mu \psi|^2 - \tilde{a}|\psi|^2 - \frac{\tilde{b}}{2}|\psi|^4 + \cdots \right],$$

where ψ is another complex scalar field, and is coupled to a dynamical gauge field a_μ with strength e. Note that we'll choose to leave this coupling strength e in the model to make connection with the numerical simulations undertaken by [2]. Physically the Abelian-Higgs Model is the low-energy description of a superconductor near its critical point. The abelian gauge field a_μ is identified with the electromagnetic field in the superconductor, while the complex scalar field ψ is identified with the Cooper pair of electrons.

Before we launch into a discussion of the phase diagram similar to what we did with the XY-Model, we will first discuss some unique peculiarities of electrodynamics in 2+1 dimensions. This preliminary discussion will help us introduce a few key elements used to understand the phase diagram.

Electrodynamics in (2+1)d and a U(1) global symmetry

Recall that in $3+1$ dimensions the free theory of electromagnetism had two degrees of freedom, which we called polarizations of the photon. In $2+1$ dimensions it will have one less degree of freedom. Thus photons in this theory posses a single polarization state, and can be described by a real scalar field. We call this scalar field the dual photon, and will investigate it in the free theory. To do this we employ some handwaving about the parition function, based largely on [11]. The free theory partition function is given by

$$\mathcal{Z}_{\text{free}} = \int \mathcal{D}a \ \exp\left(-\frac{i}{4} \int \mathrm{d}^3x \ f_{\mu\nu} f^{\mu\nu} \right).$$

Looking at the action in the free theory we note that it is quadratic in f. Thus it would make our lives slightly easier if we could change our measure from $\mathcal{D}a$ to $\mathcal{D}f$. If we decide to ignore the a_μ dependence we will need to be careful about: 1) gauge fixing problems, and 2) any identities that follow from $f_{\mu\nu} = \partial_\mu a_\nu - \partial_\nu a_\mu$ which we might miss when we ignore the a_μ dependence. Recall that in QFT we added a term of the form $-\frac{1}{2\xi}(\partial^\mu a_\mu)^2$ to the Lagrangian in order to deal with 1). However, this term was mainly important for properly quantizing the theory, and equaled zero in the MFT anyway. If we want to ignore difficulties in quantizing, we might be safe ignoring 1) for the moment. As for 2) the only identity that we need to be careful to preserve is the Bianchi identity $\epsilon^{\mu\nu\rho}\partial_\mu f_{\nu\rho} = 0$. We can enforce this by using θ as a Lagrange multiplier. Then the free theory partition function with measure $\mathcal{D}f$ instead of $\mathcal{D}a$ will be given by,

$$\mathcal{Z}_{\text{free}} = \int \mathcal{D}\theta \mathcal{D}f \ \exp\left(i \int \mathrm{d}^3x \ \left[-\frac{1}{4} f_{\mu\nu} f^{\mu\nu} + \frac{e}{2\pi} \theta \epsilon^{\mu\nu\rho} \partial_\mu f_{\nu\rho} \right] \right).$$

We have chosen θ as the name for our Lagrange multiplier for a reason, namely that we expect it to be 2π-periodic. This will occur if, and only if, the above partition function is invariant under the transformation $\theta \to \theta + 2\pi$. In other words if

$$\begin{aligned}
\mathcal{Z}_{\text{free}} &= \int \mathcal{D}\theta\mathcal{D}f \ \exp\left(i\int \mathrm{d}^3x \ \left[-\frac{1}{4}f_{\mu\nu}f^{\mu\nu} + \frac{e}{2\pi}\theta\epsilon^{\mu\nu\rho}\partial_\mu f_{\nu\rho}\right]\right)\exp\left(i\int \mathrm{d}^3x \ e\epsilon^{\mu\nu\rho}\partial_\mu f_{\nu\rho}\right) \\
\implies 1 &= \exp\left(i\int \mathrm{d}^3x \ e\epsilon^{\mu\nu\rho}\partial_\mu f_{\nu\rho}\right) \\
\implies \frac{2\pi}{e}\mathbb{Z} &\ni \int \mathrm{d}^3x \ \epsilon^{\mu\nu\rho}\partial_\mu f_{\nu\rho}.
\end{aligned}$$

The requirement that $\int \mathrm{d}^3x \ \epsilon^{\mu\nu\rho}\partial_\mu f_{\nu\rho} \in (2\pi/e)\mathbb{Z}$ will in fact be enforced by the Dirac quantization condition. Physically the Dirac quantization condition ensures that magnetic flux must come in quanta with magnitude $2\pi\hbar/e$, where we will be using the units $\hbar = 1$. This can be interpreted topologically. Recall from quantum mechanics that a particle of charge e will acquire a phase given by $\gamma = e\oint_C \vec{A}\cdot \mathrm{d}\vec{x}$ if it traverses the closed path C with a background gauge field $\vec{A}$. Using Stoke's theorem we note $\gamma = e\int_S \vec{B}\cdot \mathrm{d}\vec{S}$. In two spatial dimensions $\vec{B}\cdot\mathrm{d}\vec{S} = B_z\mathrm{d}^2x = f_{12}\mathrm{d}^2x$. Thus the phase a particle of charge e will acquire is $\gamma = e\int \mathrm{d}^2x \ f_{12}$. Suppose now that we have an ifinitely thin solenoid with flux given by an integer multiple of $2\pi/e$, then $\gamma \in 2\pi\mathbb{Z}$ for any surface enclosing the solenoid. But then $e^{i\gamma} = 1$ and traveling around the solenoid won't be detectable by our particles of charge e. This means that the ends of the solenoid will look like two spatially separated magnetic monopoles! Thus even if our theory doesn't contain magnetic monopoles, they can be introduced by considering infinitely thin solenoids with flux quantized by $2\pi/e$. Further, it turns out that $2\pi/e$ is the smallest charge a magnetic monopole could posses in a theory with charge carriers of charge e and still have angular momentum be properly quantized. We conclude by noting that if we allow infinitely thin solenoid type solutions for $\vec{A}$, then we should ensure that magnetic flux comes in integer multiples of $2\pi\hbar/e$.

We now demonstrate that we should expect $\int \mathrm{d}^3x \ \epsilon^{\mu\nu\rho}\partial_\mu f_{\nu\rho} \in 2\pi\mathbb{Z}/e$ if the Dirac quantization condition is satisfied using one of the terms of this sum below.

$$\begin{aligned}
e\epsilon^{012}\int \mathrm{d}^3x \ \partial_0 f_{12} =& e\int \mathrm{d}t \ \partial_t\left(\int \mathrm{d}^2x \ f_{12}\right) \\
=& e\int \mathrm{d}t \ \partial_t\left(\frac{2\pi}{e}n(t)\right) \text{ by the Dirac quantization condition} \\
=& 2\pi(n_{\text{final}} - n_{\text{initial}}) \in 2\pi\mathbb{Z}.
\end{aligned}$$

The other terms follow similarily, and ensure that $e\int \mathrm{d}^3x \ \epsilon^{\mu\nu\rho}\partial_\mu f_{\nu\rho} \in 2\pi\mathbb{Z}$ is equivalent to the system obeying the Dirac quantization condition. Here we will choose to enforce the Dirac quantization condition, as is common in many field theories and as supported by our discussion above [11]. Thus the partition function will be invariant under $\theta \to \theta + 2\pi$ and we can conclude that θ will be 2π-periodic.

Having commented on the periodicity of θ we return to the partition function

$$\mathcal{Z}_{\text{free}} = \int \mathcal{D}\theta\mathcal{D}f \ \exp\left(i\int \mathrm{d}^3x \ \left[-\frac{1}{4}f_{\mu\nu}f^{\mu\nu} + \frac{e}{2\pi}\theta\epsilon^{\mu\nu\rho}\partial_\mu f_{\nu\rho}\right]\right)$$

in an attempt to integrate out the more complicated $f_{\mu\nu}$ terms. Since this action is quadratic in $f_{\mu\nu}$ we know from SFT that if we integrate out $f_{\mu\nu}$ it will be identical to simply setting $f_{\mu\nu}$ equal to the value it would obtain in MFT. This value is given by its equation of motion derived from the above, which is

$$0 = \partial_\mu \frac{\delta\mathcal{L}}{\delta(\partial_\mu f^{\nu\rho})} - \frac{\delta\mathcal{L}}{\delta(f^{\nu\rho})} = \frac{e}{2\pi}\epsilon^{\mu\nu\rho}(\partial_\mu\theta) + \frac{1}{2}f^{\nu\rho},$$

so $f^{\nu\rho}$ will have equation of motion $f^{\nu\rho} = e\epsilon^{\mu\nu\rho}(\partial_\mu\theta)/\pi$. Here's the single degree of freedom we mentioned! It turns out that the 2π-periodic θ will function as our dual photon.

As an aside note that this was only possible in $(2+1)$d since the Bianchi identity in $(3+1)$d cannot be written with $\epsilon^{\mu\nu\rho}$, indeed this tensor with three indices only makes sense in $(2+1)$d. Instead to enforce the Bianchi identity in $(3+1)$d we will need a Lagrange multiplier that itself has an index, i.e. a Lagrange multiplier that functions as a dual gauge field. Working this out will give us back the familiar two polarizations of the photon.

Having gained the insight that θ is the dual photon in $(2+1)$d dimensions, and is periodic, we can substitute $f^{\mu\nu}$ with its equation of motion back into the free action. This will give us

$$\begin{aligned}
\mathcal{Z}_{\text{free}} &= \int \mathcal{D}\theta \ \exp\left(i\int \mathrm{d}^3x \ \left[-\frac{e^2}{4\pi^2}\epsilon_{\mu\nu\rho}\epsilon^{\mu\nu\rho'}(\partial^\rho\theta)(\partial_{\rho'}\theta) + \frac{e^2}{2\pi^2}\theta\epsilon^{\nu\rho\mu}\epsilon_{\nu\rho\mu'}\partial_\mu\partial^{\mu'}\theta\right]\right) \\
&= \int \mathcal{D}\theta \ \exp\left(i\int \mathrm{d}^3x \ \left[-\frac{e^2}{2\pi^2}(\partial_\mu\theta)^2 + \frac{e^2}{\pi^2}\theta\partial^2\theta\right]\right) \\
&= \int \mathcal{D}\theta \ \exp\left(-i\int \mathrm{d}^3x \ \frac{3e^2}{2\pi^2}(\partial_\mu\theta)^2\right).
\end{aligned}$$

This free action looks a lot more straightforward than the original free action. In fact there is an obvious symmetry $\theta \to \theta + \text{const}$. Since θ is 2π-periodic this constant will be an element of S^1. Thus this will be a $U(1)$ global symmetry, since $U(1)$ is isomorphic to the circle group.

The $U(1)$ symmetry will have current given by

$$j^\mu = \frac{\delta\mathcal{L}}{\delta(\partial_\mu\theta)} = \frac{e^2}{\pi^2}\partial^\mu\theta,$$

where we have removed the factor of $3/2$ for clarity. Using the equation of motion for f again, we can rewrite this as

$$j^\mu = \frac{e}{2\pi}\epsilon^{\mu\nu\rho}f_{\nu\rho}.$$

But then the conservation of this current is just the Bianchi identity in $(2+1)$d! Clearly this will be a symmetry of the free theory. But by the definition of $f_{\mu\nu} = \partial_\mu a_\nu - \partial_\nu a_\mu$ the Bianchi identity is trivially satisfied for the full Abelian-Higgs Model as well. Thus the $U(1)$ global symmetry of the free theory will be a global symmetry for the interacting Abelian-Higgs Model also.

Before we go on to a discussion of the phase diagram let's summarize. Using the free theory we have shown that electrodynamics in $(2+1)$d is characterized by a single degree of freedom, the dual photon θ. This dual photon is 2π-periodic, provided we assume the Dirac quantization condition, which we will. Finally, the full Abelian-Higgs Model posesses a $U(1)$ global symmetry that can be thought of either as enforcing the Bianchi identity or as shifting the value of the dual photon by a constant. We now use all these facts to explore the phase diagram.

$\tilde{a} < 0$:

In the $\tilde{a} < 0$ phase ψ will acquire a vacuum expectation value, just as ϕ acquired a vacuum expectation value in the $a < 0$ of the XY-Model. This will break the $U(1)$ gauge symmetry $\psi \to e^{i\alpha(x)}\psi$ and $a_\mu \to a_\mu + \frac{1}{e}\partial_\mu \alpha(x)$ that the Abelian-Higgs Model posseses. Just as the breaking of the $U(1)$ global symmetry of the XY-Model led to vortices in its $a < 0$ phase, so will the breaking of the $U(1)$ gauge symmetry give rise to vortex solutions in this phase. However, these vortices will be rather different animals than the vortices we saw in the XY-Model. We'll see that their energy won't diverge logarithmically as it did for the vortices in the XY-Model. Further the vortices will carry a quantized magnetic flux.

Before we can show these differences we will need to find what parameter's winding actually defines the Abelian-Higgs vortices. To do this we make the redefinition $\psi = \rho e^{i\sigma}$. Then we can expand to covariant derivative

$$\begin{aligned}
|D_\mu \psi|^2 =& (\partial_\mu - iea_\mu)\psi(\partial^\mu + iea^\mu)\psi^* \\
=& |\partial_\mu \psi|^2 - iea_\mu(\psi\partial^\mu\psi^* - \psi^*\partial^\mu\psi) + e^2(a_\mu)^2|\psi|^2 \\
=& (\partial_\mu\rho)^2 + \rho^2(\partial_\mu\sigma)^2 - ea_\mu(2\rho^2\partial^\mu\sigma) + e^2(a_\mu)^2\rho^2 \\
=& (\partial_\mu\rho)^2 + \rho^2(\partial_\mu\sigma - ea_\mu)^2 \\
\implies \mathcal{L}_{AH} =& -\frac{1}{4}f_{\mu\nu}f^{\mu\nu} + (\partial_\mu\rho)^2 + \rho^2(\partial_\mu\sigma - ea_\mu)^2 - \tilde{a}\rho^2 - \frac{\tilde{b}}{2}\rho^4 + \cdots ,
\end{aligned}$$

where the substitution $\psi = \rho e^{i\sigma}$ was straightforward given that we did more or less the same calculation for the XY-Model. Note that the $U(1)$ gauge symmetry $\sigma \to \sigma + \alpha(x)$ and $a_\mu \to a_\mu + \frac{1}{e}\partial_\mu\alpha(x)$ is very clear in this form. It will turn out that the vortices in this phase are given by the winding of σ.

The fact that $\partial_\mu\sigma$ and ea_μ are packaged together in the same term should give us some hint that these vortices might pick up magnetic flux. To see this more explicitly let's look

at the equation of motion for a_μ

$$0 = -\partial_\mu f^{\mu\nu} - (-e)2\rho^2(\partial^\nu\sigma - ea^\nu).$$

Recall from electromagnetism that $\partial_\mu f^{\mu\nu} = J^\nu$, where we've used an uppercase J to distinguish the electromagnetic current from the $U(1)$ global current $j^\mu = (e\epsilon^{\mu\nu\rho}f_{\nu\rho})/2\pi$. Let's consider the case where there's no externally imposed current $J^\nu = 0$. Then the above equation reveals that $\partial^\mu\sigma = ea^\mu$ whenever $\rho \neq 0$. This relation is what will give the vortices a quantized magnetic flux. Just as with the vortices of the XY-Model we note that

$$\begin{aligned}
2\pi\mathbb{Z} \ni &\oint \partial_i\sigma \ \mathrm{d}x^i \text{ for topological reasons} \\
=&e\oint a_i \ \mathrm{d}x^i \text{ for } i = 1,2 \text{ by e.o.m.} \\
=&e\int_S (\partial_1 a_2 - \partial_2 a_1) \ \mathrm{d}^2x \text{ by Stokes} \\
=&e\int_S f_{12} \ \mathrm{d}^2x.
\end{aligned}$$

Thus in this phase the Dirac quantization condition will in fact fall neatly out of the fact that the winding number of a vortex must be an integer! This relationship also allows us to see that a vortex with winding number n will posses magnetic flux $2\pi n/e$.

Next we want to check the energy of a single vortex. We will consider a static vortex situated at the origin. If it has winding number n, then by analogy with the vortices in the XY-Model $\sigma(r,\vartheta) = n\vartheta$. We will further choose to work in the Lorenz gauge $\partial_\mu a^\mu = 0$. Then the equation of motion for a_μ becomes $\partial^2 a_\mu = 2e\rho^2(\partial_\mu\sigma - ea_\mu)$. Since $\partial_\mu\sigma = (n/r)\vec{e}_\vartheta$ then we can try the ansatz $a_\mu = (g(r)/r)\vec{e}_\vartheta$ which we note does in fact satisfy the Lorenz gauge condition. With these substitutions the equation of motion reduces to

$$\frac{\mathrm{d}^2 g}{\mathrm{d}r^2} - \frac{1}{r}\frac{\mathrm{d}g}{\mathrm{d}r} = 2e\rho^2(n - eg(r)).$$

We could try to choose $g(r) = n/e$, which would satisfy the above equation. But then $\partial^\mu\sigma = ea^\mu$ and we can see directly from the above integrals that this will give rise to $f_{12} = |\vec{\nabla}\times\vec{a}| = |\vec{\nabla}\times n/(er)\vec{e}_\vartheta| = (2\pi n/e)\delta^{(2)}(\vec{r})$. Having a δ function flux might lead to some problems with the energy. Thus we follow the lead of Nielsen and Olesen [8] and take the boundary condition $g(\infty) = n/e$ and $g(0) = 0$. With this choice we will eliminate the divergence of $f_{\mu\nu}$ at the origin.

Next, just as with the vortices in the XY-Model, we will need $\rho(0) = 0$ to eliminate problems with the multivaluedness of σ there. However, again just as with the XY-Model ρ will quickly interpolate from a value of zero at the center of the magnetic flux vortex to ρ_0 outside the vortex. We will call this interpolation length l as we did for the XY-Model. For $r \gg l$ the ρ^2 factor in the equation of motion will then become just ρ_0^2. It can then be checked by subsitution that $g(r) = n/e + c\sqrt{r}e^{-\sqrt{2}e\rho_0 r}$ for c some constant satisfies the

equation of motion to order $O(r^{-3/2}e^{-\sqrt{2}e\rho_0 r})$, provided $r \gg l$ so $\rho = \rho_0$. Looking at the exponential reveals that this will therefore be a valid description of $g(r)$ for $r \gg 1/(e\rho_0)$. Just as l represents the distance required for the density to reach its value at infinity, $1/(e\rho_0)$ represents the distance for the vector potential to reach its value at infinity. We will call $\xi := \max(l, 1/(e\rho_0))$ the core size of the vortex, since both the density and vector potential cores will be contained within ξ.

Having discussed some of the properties of the vortex solutions, we now want to find their energy. This means that we will need to find the energy momentum tensor for the Abelian-Higgs Model.

$$\begin{aligned} T^{\mu\nu}_{AH} =& 2\left.\frac{\delta \mathcal{L}_{AH}}{\delta g_{\mu\nu}}\right|_{g=\eta} - \eta^{\mu\nu}\mathcal{L}_{AH} \\ =& 2(\partial^\mu\rho)(\partial^\nu\rho) + 2\rho^2(\partial^\mu\sigma - ea^\mu)(\partial^\nu\sigma - ea^\nu) + f^{\mu\lambda}f_\lambda{}^\nu - \eta^{\mu\nu}\mathcal{L}_{AH}, \end{aligned}$$

where the variation of $f_{\alpha\beta}f^{\alpha\beta}$ with respect to the metric is a standard result from electromagnetism. Since we are considering a static vortex and $a^0 = 0$ the energy density becomes:

$$\begin{aligned} T^{00}_{AH} =& -\eta^{00}\mathcal{L}_{AH} \\ =& \frac{1}{4}f_{\mu\nu}f^{\mu\nu} + |\vec{\nabla}\rho|^2 + \rho^2|\vec{\nabla}\sigma - e\vec{a}|^2 + \tilde{a}\rho^2 + \frac{\tilde{b}}{2}\rho^4 + \cdots \\ =& \frac{1}{2}(f_{12})^2 + \frac{\rho^2}{r^2}(n - eg(r))^2 + |\vec{\nabla}\rho|^2 + \tilde{a}\rho^2 + \frac{\tilde{b}}{2}\rho^4 + \cdots \\ =& \frac{1}{2r^2}\left(\frac{\mathrm{d}g}{\mathrm{d}r}\right)^2 + \frac{\rho^2}{r^2}(n - eg(r))^2 + \rho \text{ terms}, \end{aligned}$$

where used the fact that $f_{12} = |\vec{\nabla} \times \vec{a}| = g'/r$. As with the XY-Model we note that for $r \gg \xi$ the ρ terms will be negligible. Then

$$E = E_{\text{core}} + 2\pi\int_\xi^\infty \left[\frac{1}{2r^2}\left(\frac{\mathrm{d}g}{\mathrm{d}r}\right)^2 + \frac{\rho_0^2}{r^2}(n - eg(r))^2\right] r\mathrm{d}r.$$

But we saw above that $g(r) = n/e + c\sqrt{r}e^{-\sqrt{2}e\rho_0 r}$ for $r \gg \xi$. Then the integrand will just look like $e^2\rho_0^2 r^{-1}e^{-r/\xi}$, which indeed has a finite integral. Since we were very careful in choosing $g(0) = 0$, there won't be any divergences near the origin and thus E_{core} will be finite as well.

The key to the magnetic flux vortex having finite energy is the same thing that gave it magnetic flux in the first place; the fact that $\partial_\mu\sigma$ and ea_μ are packaged together in the same term. In the XY-Model the energy looked like $\rho_0^2|\vec{\nabla}\theta|^2$ outside the vortex core. This meant that the n/r dependence of $\vec{\nabla}\theta$ was unable to be canceled by anything, and caused the energy to diverge at long distances. In the Abelian-Higgs Model the energy looks like $\rho_0^2|\vec{\nabla}\sigma - e\vec{a}|^2$ outside the vortex core. Thus even though $\vec{\nabla}\sigma$ still has a harmful n/r dependence, we can cancel that by taking $\vec{a} = n/(re)\vec{e}_\vartheta$ at large r. This takes care

of the divergent part of the energy, and leads to the Abelian-Higgs vortices having finite energy.

Having very thoroughly investigated the breaking of the $U(1)$ gauge symmetry, we might wonder what happens to the $U(1)$ global symmetry in this phase. We know that the photon will become massive because of the $e^2\rho_0^2 a_\mu a^\mu$ term that appears. Then the massiveness of a_μ means that all excitations will be gapped. But this means that there is no need to choose a vacuum expectation value for the dual photon θ. Thus the $U(1)$ global symmetry will remain unbroken in this phase.

$\tilde{a} > 0$:

In this phase the ψ fields will be massive and the $U(1)$ gauge symmetry will remain unbroken. This means that the photons will remain massless. Then our theory becomes gapless, and a vacuum expectation value will need to be chosen for the dual photon θ. Consequently the $U(1)$ global symmetry is spontaneously broken in this phase. The dual photon is precisely the Goldstone mode of this broken symmetry since the photon will remains massless. Note that the dual photon current $j^\mu = (e^2/\pi^2)\partial^\mu\theta$ bears a striking similarity to the global $U(1)$ symmetry current $j^\mu \propto \rho_0^2\partial^\mu\theta + O(\tilde{\rho})$ of the XY-Model's $a < 0$ phase.

We can make a further comment about the massive ψ fields. Any elementary ψ excitation will describe a particle with charge $\pm e$. In the MFT these particles will interact via Maxwell's equations in $(2+1)$d. This calculation is identical to the familiar one for QED, just with scalars instead of fermions. In particular, in the static case the particles will interact via the Coulomb force in two spatial dimensions. We've already noted in the $a < 0$ phase of the XY-Model that the Coulomb force in two spatial dimensions leads naturally to logarithmic confinement. The calculation here will be no different, except in this case the electric field will arise naturally and won't need to be introduced by a field transformation as with the XY-Model. Thus massive ψ particles of opposite charge will be logarithmically confined, just as the vortices in the $a < 0$ phase of the XY-Model.

$\tilde{a} = 0$:

Again we should expect a critical point to occur as we pass through $\tilde{a} = 0$. Numerical simulations show that this is the same quantum critical point as exists in the XY-Model [2], [11]. We will soon turn our attention to demonstrating this fact ourselves, but first we want to summarize the similarities between the models in a handy table.

2.3. **Comparison.** This table is meant to make comparison between the models easy. The key insight that allowed us to make this comparison in the first place is that the two models share a global $U(1)$ symmetry. In particular, in both models the current of

this global $U(1)$ symmetry has the form $j^\mu \propto \partial^\mu \theta$, where θ is some 2π-periodic real field. Matching the broken and unbroken phases of this nearly shared $U(1)$ global symmetry allows the similarities to become manifest.

XY-Model	**Abelian-Higgs Model**
Broken global $U(1)$ symmetry phase	
$a < 0$	$\tilde{a} > 0$
2π-periodic Goldstone modes θ	2π-periodic Goldstone modes θ
Vortices described by winding of θ	Particles given by ψ, charged
In MFT vortices interact via Coulomb force in 2d	In MFT particles interact via Coulomb force in 2d
Vortices of opposite charge/sign are logarithmically confined	Particles of opposite charge are logarithmically confined
Unbroken global $U(1)$ phase	
$a > 0$	$\tilde{a} < 0$
Particles described by field ϕ	Vortices by winding of σ
Particles have a finite mass	Vortices have a finite energy (mass in rest frame)

As a small note we've ignored the $U(1)$ gauge symmetry of the Abelian-Higgs Model. This shouldn't worry us too much because it merely described a redundancy in the model and not anything physical.

At any rate, the descriptions of the phases above look very similar if we assume that $\tilde{a} \sim -a$, and that vortices and particles of the opposing theories match to each other. Indeed the vortices of the XY-Model seem to have the same interaction as the particles of the Abelian-Higgs Model, and vice versa, in the appropriate phase. However, this was all heuristic. We could have been much more rigorous and taken more time with the interactions beyond the fourth order. This is attempted for example in [9], and can be taken further in supersymmetric theories [11]. One other thing that we didn't take time to verify was that the critical exponents of the Abelian-Higgs Model near the critical point $\tilde{a} = 0$ are actually the same as those of the XY-Model near the critical point $a = 0$, or that the Abelian-Higgs Model in fact goes through an inverted XY-type transition. We are interested in showing this for ourselves, since verifying that the models lie in the same universality class at the critical point would provide solid evidence to suggest that they are the equivalent away from the critical point. As such we dive right into numerics, using the Ising Model as our introduction.

3. Numerics

3.1. **Introduction: Metropolis and the Ising Model.** In the discussion of the critical exponents of the XY-Model we promised to discuss how we could go from a continuum field theory to one that was amenable to numerical analysis. Before we do that we must first establish how we can possibly use numerical analysis to analyze field theories. This introduction is motivated largely by Kardar [4].

The first principle of numerical analysis is that it is often easier and computationally cheaper to work with real numbers. As such we will firstly consider field theories in the SFT perspective to get rid of the prefactors of i. Recall we can just Wick rotate to take a QFT to a SFT, so this shouldn't lose us any generality. Secondly, we are primarily interested in computing the expectation values of certain operators, e.g.

$$\langle \mathcal{O} \rangle = \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \; \mathcal{O}(\phi, \partial_\mu \phi, \ldots) \; e^{-\beta F(\phi, \partial_\mu \phi, \ldots)},$$

for the case of a real scalar field. This seems hard. Let's start brainstorming how we might even attempt such an integral. Perhaps naïvely we could try to do this like we might a (simple) numerical integral. We could pick a discrete number of possible field configurations ϕ_i and sum $\mathcal{O}(\phi, \ldots) e^{-\beta F(\phi, \ldots)} / \mathcal{Z}$ over these discrete configurations. But the space of all possible configurations is enormous, and in the above case it is infinite dimensional! How can we even hope to choose a representative collection with only finitely many choices?

A better approach would be to forget the integral, and recall that we are looking to take an expectation value. Let $P(\phi) = e^{-\beta F(\phi, \ldots)} / \mathcal{Z}$ be the probability distribution for configurations. Now draw a representative sample of configurations $\mathcal{C}_n = \{\phi_i \mid 1 \leq i \leq n\}$ that are distributed according to P, i.e. if $P(\phi_j)$ is small then ϕ_j will be less likely to show up in our representative sample. Then our expectation value should be approximated by the mean according to this representative sample,

$$\langle \mathcal{O} \rangle \approx \frac{1}{|\mathcal{C}_n|} \sum_{\phi_i \in \mathcal{C}_n} \mathcal{O}(\phi_i).$$

Thus if we can find an algorithm that takes $P(\phi)$ and gives us our representative sample, $\mathcal{C}_n$, we will be able to actually compute these expectation values. In fact, since we likely won't know $\mathcal{Z}$ explicitly, we would prefer to find an algorithm that just takes $e^{-\beta F(\phi, \ldots)}$ and gives us $\mathcal{C}_n$. Thankfully the Metropolis algorithm does exactly this.

Metropolis algorithm

Having established the kind of question that we would like the Metropolis algorithm to resolve, we will endeavor to provide an introduction to the algorithm and its typical uses in statistical physics. For this introduction we will consider the Ising model, as it provides the simplest introduction to many of these concepts. The model has free energy

given by

$$F(\{s\}) = -J \sum_{\langle ij \rangle} s_i s_j,$$

where the sum is over nearest neighbor pairs on a lattice and where $s_i = +1$ or -1 for every site. We describe how the Metropolis algorithm works for the Ising model, and show the progression of a typical Ising state under the model. Then we move to a discussion of equilibration. We will use our equilibrated states to numerically find two of the critical exponents of the Ising model. All of these ideas will be useful when we move to the much more complicated task of finding the critical exponents of the Abelian-Higgs Model.

We will henceforth restrict ourselves to systems where the space of possible configurations is finite but large, as do [1] (XY Model), [2] (Abelian-Higgs Model), and [4] (spin models). For the Ising model this means making our lattice finite, but periodic. Then the space of possible configurations has size 2^N, where N is the number of lattice sites; large, but not infinite. The Metropolis algorithm then proceeds as follows for the simple case of the Ising model.

Metropolis algorithm: Ising Model

1. Choose an initial state $\{s\}$ drawn from a uniform distribution. That is to say, for each site assign it a value $+1$ or -1 with probability $1/2$, regardless of the values at the neighboring sites.
2. Next pick a site k at random. If the lattice is $N \times N$, then $k = (i, j)$ and this can be done by choosing two random numbers i and j between 1 and N from a uniform distribution.
3. Now flip the spin at the site k, $s_{i,j} \to -s_{i,j}$, and calculate the free energy change $\Delta F = F(\{s\}_{\text{final}}) - F(\{s\}_{\text{initial}})$. For the 2d Ising model this is a very fast calculation, since

   $$\Delta F = 2J s_{i,j}(s_{i+1,j} + s_{i-1,j} + s_{i,j+1} + s_{i,j-1})$$

   depends only on $k = (i, j)$'s nearest neighbors. Note that $i \pm 1$ and $j \pm 1$ are all taken mod N, as the lattice is taken to be periodic.
   (i) If ΔF is less than zero, keep the flipped spin at k.
   (ii) If $\Delta F > 0$, then choose a random number u from a uniform distribution between 0 and 1. If $u < e^{-\beta \Delta F}$ then keep the flipped spin. If $e^{-\beta \Delta F} \leq u$ then flip the spin again, taking it back to its original state.
4. Return to step 2 and pick a new site. If this has been done the appropriate number of times, we can exit here.

Note that step 3 is where the action happens, and it appears to be fairly intuitive. If changing the spin would lower the energy, then it makes sense that most systems will choose to do this. If most systems would choose to do this, then they should be more probable according to the Boltzmann distribution, and so we should be getting close to a representative state. Further, the noise in step 3 (ii) is governed by the temperature

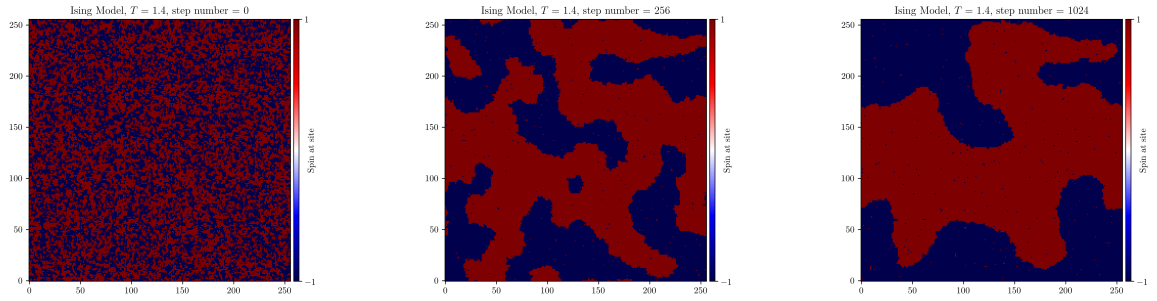

(A) Step number = 0 (B) Step number = 256 (C) Step number = 1024

FIGURE 2. Generic simulation of Ising model at $T = 1.4$, system equilibrates using Metropolis algorithm, starting from random intial state.

of the system and the size of the jump. If the temperature is very small these thermal changes will be unlikely, and the system will simply move towards a state of minimum energy. However, with larger temperature this step will introduce a lot of noise into the system. Again this agrees with our understanding of a system's temperature. For a more explicit proof that the Metropolis algorithm will produce representative states, see Kardar [4]. The proof relies mainly on linear algebra and the principle of detailed balance, as discussed in Lent term's Theoretical Physics of Soft Condensed Matter. It is not hard, but slightly beyond the scope of this essay.

Instead of proving that the Metropolis algorithm works, let's show it. We'll implement it for the Ising model in python. This is more or less straightforward using the above algorithm, so I won't provide much more detail. However, if the reader wants to read or run my code for themselves they can find it listed as `ising_time_series.py` at my GitHub repository [6]. It was inspired in part by [10], as well as [4].

Figure 2 shows an example simulation at $T = 1.4$. Since calculation of ΔF is very fast for the Ising model, running the simulation for this figure only took around 5 minutes on my laptop. It could likely be optimized further as per [10], but for the purposes of introducing the Metropolis algorithm this should be good enough.

From the discussions in SFT we expect the critical temperature of the Ising model to be roughly $T_c = 2.27$ in our dimensionless units, where $k_B = 1$. Thus at a temperature of $T = 1.4 < T_c$ we should expect to see the system developing domains with a small amount of noise. This is exactly what we see in Figure 2. The random initial state quickly turns into some rather complicated domains, which take a relatively longer time to coarsen and smooth out. At this step number we expect that the system will have reached a representative state, i.e. that it will be in a state $\{s\}$ with high probability $e^{-\beta F(\{s\})}/\mathcal{Z}$. We call the time it takes for a state to go from the totally random starting point to a representative state the equilibration time τ_{eq}.

As an aside, "time" here simply refers to the number of steps taken by our Metropolis algorithm. There is no physical mechanism arising from $F(\{s\})$ that sets a time scale.

If we wanted to assume that the system relaxes in an overdamped way towards its equilibrium state then we could write a stochastic PDE that would produce this. This is known as the Langevin equation and was discussed in Theoretical Physics of Soft Condensed Matter. It often produces coarsening dyanmics that are closer to what is seen in physical systems. To make our simulations behave more like this we can try to modify our Metropolis algorithm by using Langevin dynamics to propose new states (instead of choosing them randomly), and then using the Metropolis step to accept or reject them. This combination is known as the Metropolis-adjusted Langevin algorithm.

Critical Exponents

Despite the aside, for the moment we're not particularly interested in the way that our system gets to equilibrium. It could arrive at equilibrium in an unphysical seeming way, and this wouldn't bother us. We only care about what it does when it reaches equilibrium. This is because we're trying to get a representative sample of states and states at equilibrium are more likely. Thus we will gloss over some possible improvements, and simply compute any statistics of the model after states have equilibrated for an equilibration time τ_{eq}. We provide an algorithmic description of how this might work for computing the critical exponents α and β. In order to compute these we will need to determine the behavior of the specific heat, $C_V(T)$, and magnetization, $|\langle M\rangle|(T)$, near the critical temperature T_c.

Note that the algorithm below is less general than the more generic Metropolis algorithm, and essentially serves as a pseudocode description of `ising_crit_exp.py` found at my GitHub repository [6].

Critical Exponents: Ising Model

1. Choose a set of parameters; fix the size of the simulations to be tested and the number of temperature values to test. To avoid finite size effects larger simulation sizes are better, but small simulation sizes can often give instructive results. I chose an 8×8 system size. Since the critical temperature for the Ising Model is $T_c = 2.27$ we should cluster temperature values around T_c for better fits, as this is where we expect our quantities to diverge.
2. Pick a single temperature value T and choose a new random initial state, as in the Metropolis algorithm. Additionally create empty arrays G_{array} and M_{array} with length τ_{calc}. I have chosen $\tau_{\text{calc}} = 2^{10}$.
   (i) Equilibrate the random intial state by applying the Metropolis algorithm a number of times τ_{eq}. If the lattice is $N \times N$ each application will allow for N^2 spin flips. Unfortunately the time for the states to equilibrate is temperature dependent, $\tau_{\text{eq}} = \tau_{\text{eq}}(T)$, and diverges near T_c for an infinite system. Thankfully for a finite system finite size effects will stop this divergence. We will

discuss this later, but as an ad hoc solution I chose $\tau_{\text{eq}}(T) = \tau_0^{2-2.269\frac{|T-T_c|}{T_c}}$ with $\tau_0 = 2^9$. This way at T_c the equilibration time is much longer, 2^{18}.

(a) Once the system has been equilibrated we assume that it is a representative state $\{s\}_i$. Having reached such a state, we now measure the quantities we care about. For us this will be $G_i := \beta F(\{s\}_i)$ and $M_i := M(\{s\}_i)$, the magnetization. Add the measured values to the arrays we created $G_{\text{array}} = (G_1, G_2, \ldots)$ and $M_{\text{array}} = (M_1, M_2, \ldots)$.

(b) Having obtained a representative state after equilibration, we apply the Metropolis algorithm once more. Since we started from a representative state, we assume this application will take us to another representative state $\{s\}_{i+1}$. It is likely to be correlated slightly with $\{s\}_i$, but we expect that if we do this enough times that correlation will disappear [4].

(c) Having found the next equilibrated state, return to (a) and compute the quantities of interest. Continue through the loop of (a)-(c) for a number of times τ_{calc} until we have sampled our quantities of interest in sufficiently many representative states, and have totally filled G_{array} and M_{array}.

(ii) Now compute the specific heat $C_V = \beta^2\langle F^2\rangle - \beta^2\langle F\rangle^2 \approx \overline{G^2_{\text{array}}} - \overline{G_{\text{array}}}^2$ and the magnetization $\langle M\rangle \approx \overline{M_{\text{array}}}$ for this temperature value.

3. Return to step 2 and choose the next temperature value. Exit this loop when we have computed C_V and $\langle M\rangle$ for all temperature values.

A few comments on the algorithm. To have truly representative states in our arrays G_{array} and M_{array} we might want to start G_1 and M_1 in a different random state than G_2 and M_2. Then we would equilibrate both states, reaching two truly uncorrelated representative states, where we would then measure the quantities. This undoes any correlation between G_1 and G_2, since they are no longer separated by only a single Metropolis algorithm move. However, it would be prohibitively computationally costly to equilibrate $\tau_{\text{calc}} = 2^{10}$ systems for every temperature value, instead of only a single system. Thus we trust that if we apply steps 2(a)-2(c) enough times, we will get an appropriate average in 2(ii) [4].

Secondly, we could try to optimize our algorithm by lessening the temperature dependence of τ_{eq}. The reason it diverges near T_c is because $\tau_{\text{eq}} \sim \xi^z$, with z some positive exponent that depends on the algorithm used [7]. Thus as ξ diverges near the critical temperature, so does τ_{eq}. This makes some sense. Near the critical temperature spins begin to cluster together in chains that can extend across the whole simulation, since the characteristic length ξ is headed to infinity. These clusters are hard to change through the local updates that the Metropolis algorithm makes, so it stands to reason that it would take a lot more time to equilibrate the system. To try to lessen this dependence

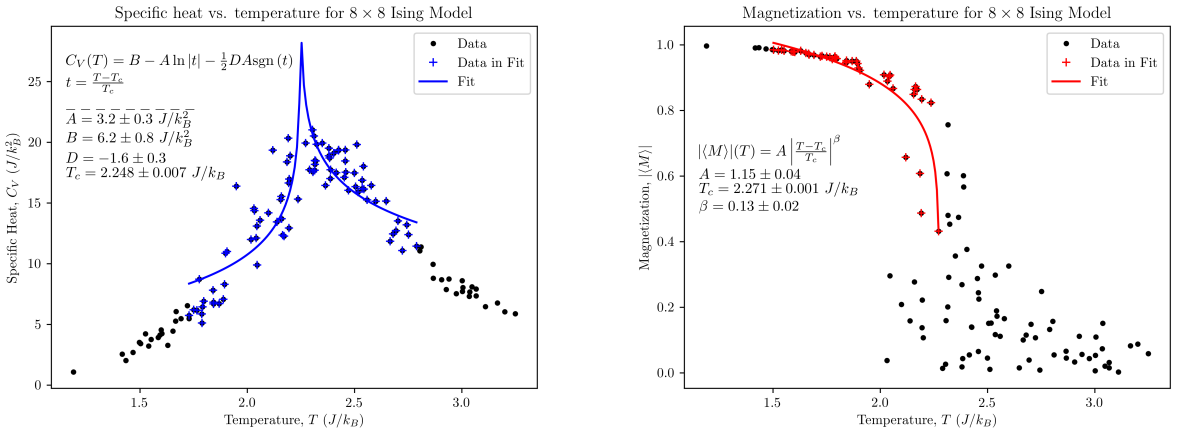

(A) Fit demonstrating $\alpha \simeq 0$.

(B) Fit demonstrating $\beta \simeq 0.125$

FIGURE 3. Scan of temperature axis for 8×8 Ising Model. Fits demonstrate that the critical exponents can be approximately determined, even for such a small system size. Using the value of T_c from each we can average to determine that $T_c = 2.260 \pm 0.008\ (J/k_B)$.

we might try to make an algorithm that flipped whole clusters at a time, instead of just flipping a single site like the Metropolis algorithm. For the Ising model these are known as cluster algorithms, the most famous being the Wolff algorithm. Indeed for the Ising model the Wolff algorithm almost entirely removes the temperature dependence of τ_{eq} [7]. However, the Wolff algorithm was designed to solve this specific problem for magnetic spin simulations, and so it lacks the generality of the Metropolis algorithm. As such we will stick to the Metropolis algorithm for this paper, since we will be applying it to the Abelian-Higgs Model which is not a magnetic spin type model.

Having commented on the algorithm for determining $C_V(T)$ and $|\langle M\rangle|(T)$ near $T = T_c$ and discussed possible improvements, it is now time to actually simulate it. We can see this in Figure 3, where the points are plotted. Note that the magnetization in (B) is a little noisy, this is likely because the simulation size was the relatively small 8×8 size. Though the specific heat is more clean, both plots are fairly amenable to fitting.

The fitting for the plots was done using `ising_fits.py` in my GitHub repository [6]. Since this fitting just used a standard python library for nonlinear fitting I will skip the details and just discuss the results of the fits. Previous studies have found that $\alpha = 0$ with a log divergence for the Ising model in two dimensions; we discussed this in SFT [12]. Thus following the lead of Halperin [2] we tried to fit the specific heat near T_c to

$$C_V(T) = -A\ln\left|\frac{T-T_c}{T_c}\right| - \frac{1}{2}DA\mathrm{sgn}\left(\frac{T-T_c}{T_c}\right) + B.$$

This fit was pretty good, as we can see in Figure 3(A). It leads us to believe that our Metropolis algorithm for the Ising model was correct and that $T_c \approx 2.25$. Fitting $|\langle M\rangle|$ near the critical point gives us that $\beta = 0.13 \pm 0.02$. Thus though this fit is rather crude

we see that the literature value of $\beta = 1/8$ is within our errorbars [12]. Finally, averaging the fitted value of T_c from both fits gives us that $T_c = 2.260 \pm 0.008$ (J/k_B). The literature value $T_c = 2.269$ is nearly within a single errorbar. Thus though our system size was small and we used the Metropolis algorithm instead of the more refined Wolff algorithm, we still got fairly accurate fits for α, β, and T_c. We therefore conclude our introduction to numerics using the Metropolis algorithm and the Ising Model.

3.2. **Numerics for a continuum QFT.** Having understood how to extract the critical exponents from a statistical field theory that has finitely many states, we now want to understand how to extract the critical exponents from a quantum field theory with infinitely many possible states. Since our aim is to find critical exponents of the Abelian-Higgs Model, we will use this as our example. Recall that its partition function is given by

$$\mathcal{Z}_{AH}(e, \tilde{a}) = \int \mathcal{D}a\mathcal{D}\rho\mathcal{D}\sigma \; e^{iS_{AH}}$$

$$S_{AH} = \int \mathrm{d}^3x \; \left[-\frac{1}{4} f_{\mu\nu} f^{\mu\nu} + (\partial_\mu \rho)^2 + \rho^2 (\partial_\mu \sigma - e a_\mu)^2 - \tilde{a}\rho^2 - \frac{\tilde{b}}{2}\rho^4 + \cdots \right]$$

Note that we have ignored any possible gauge fixing term, and that we have made the substitution $\psi = \rho e^{i\sigma}$. Especially in the resulting numerics having real variables will make our lives much easier. There still remains a factor of i in the exponent, which would cause difficulties if we tried to apply the Metropolis algorithm. Indeed the Metropolis step in this form would involve asking whether or not a random variable $u \in [0, 1]$ is less than or equal to the complex $e^{i\Delta S_{AH}}$, clearly not well defined. As such we will first need to understand what happens when we Wick rotate. This is covered very well using the illustrative example of quantum mechanics in pages of 148-150 of Tong's quantum Hall effect notes [11]. I can't really improve on the exposition there so I won't reproduce it here. The main takeaway is that if you take a QFT living on Minkowski space $\mathbb{R}^{1,d-1}$ and consider its Euclidean path integral formulated on $\mathbb{S}^1 \times \mathbb{R}^{d-1}$ with periodic time $\tau \in [0, \beta)$, then this will give you its thermal partition function. One needs to keep track of signs in the fermionic case, but this isn't a problem for our theory.

Continuum Abelian-Higgs to Lattice Superconductor Model (LSM)

To Wick rotate and obtain the Euclidean path integral we define $y^1 = it$, $y^2 = x^1$, and $y^3 = x^2$. The vectors $\vec{y}$ will live in $\mathbb{S}^1 \times \mathbb{R}^2$. Then $(\partial_\mu \rho)(\partial^\mu \rho) = (\partial_t \rho)^2 - \sum_{i=1,2} (\partial_{x^i} \rho^2) = -|\vec{\nabla}\rho|^2$, where the gradient is understood to be in the $\vec{y}$ variable. The gauge field will also need to be transformed. Define A_i to be the new gauge field and relabel its indices to be manifestly spatial. This means $(A_1, A_2, A_3) = (ia_0, a_1, a_2)$ given our redefined time

coordinate. Then $f_{0i} = i(\partial_1 A_{i+1} - \partial_{i+1} A_1)$ and we can deduce that

$$\begin{aligned} f_{\mu\nu} f^{\mu\nu} =& 2[(f_{01})^2 + (f_{02})^2 - (f_{12})^2] \\ =& -2[(\partial_1 A_2 - \partial_2 A_1)^2 + (\partial_1 A_3 - \partial_3 A_1)^2 + (\partial_2 A_3 - \partial_2 A_3)^2] \\ =& -2|\vec{\nabla} \times \vec{A}|^2, \end{aligned}$$

where again the gradient is understood to be in the $\vec{y}$ variable. With this change the path integral becomes

$$\mathcal{Z}_{AH}(e, \beta) = \int \mathcal{D}A \mathcal{D}\rho \mathcal{D}\sigma \; e^{-S_E}$$

$$S_E = \int \mathrm{d}^3 y \; \left[\frac{1}{2} |\vec{\nabla} \times \vec{A}|^2 + |\vec{\nabla}\rho|^2 + \rho^2 |\vec{\nabla}\sigma - e\vec{A}|^2 + \tilde{a}\rho^2 + \frac{\tilde{b}}{2}\rho^4 + \cdots \right],$$

where the coordinate y^1 will be β-periodic. Just as with the Ising model we'll now choose to make the spatial coordinates periodic, so that $\vec{y} \in (\mathbb{S}^1)^3$. We're now going to make some manipulations to bring the model more in line with the three-dimensional lattice superconductor model (LSM) studied by Halperin in [2].

Suppose that we are in the $\tilde{a} < 0$ phase with its accompanying magnetic flux vortices. Then $\rho = \rho_0 + \tilde{\rho}$ with $\rho_0^2 = -\tilde{a}/\tilde{b}$. We can check that up to a constant term the potential $\tilde{a}\rho^2 + \tilde{b}\rho^4/2 = \tilde{b}(\rho^2 - \rho_0^2)^2/2$ by expansion. But now suppose that we send $\tilde{b} \to \infty$. Then exciting modes where ρ^2 differs from ρ_0^2 will cost an infinite amount of energy, $\tilde{b}(\rho^2 - \rho_0^2)/2$, since $\rho_0^2 \neq \rho^2$. This will effectively "freeze out" fluctuations in the radial part of ψ and thus fix $\rho = \rho_0$. The Euclidean action for this frozen Abelian-Higgs Model then becomes:

$$S_E = \int \mathrm{d}^3 y \; \left[\frac{1}{2} |\vec{\nabla} \times \vec{A}|^2 + \rho_0^2 |\vec{\nabla}\sigma - e\vec{A}|^2 \right].$$

Note that since $\rho_0^2 = -\tilde{a}/\tilde{b}$ and we have sent $\tilde{b} \to \infty$, then to keep ρ_0^2 nonzero we will need to choose $\tilde{a}$ to be a very large negative value. This frozen Abelian-Higgs Model is the continuum version of the LSM studied by Halperin, where he assumes that $\rho_0^2 = \beta/2$ [2].

Here is where we really get into the question raised in our discussion of the critical exponents of the XY-Model; where did the temperature come from? It turns out it was lurking in $\rho_0^2 \propto \tilde{a}$ all along. We know that the critical point for the Abelian-Higgs Model should occur when $\tilde{a} = 0$. Thus if we assume that $\tilde{a} = \tilde{a}(T)$, as was μ^2 in the Ising model studied in SFT [12], we should have that $\tilde{a}(T_c) = 0$. In particular we might make the assumption that the scaling will be linear, i.e. $\tilde{a} \propto T_c - T$ as in the Ising model in SFT. Then $\rho_0^2 \propto (T - T_c)/\tilde{b}$. Since $\tilde{b}$ will be going to infinity this will be going to zero, and so we will need to choose T to be a large temperature value. But then as $\tilde{b}$ goes to infinity ρ_0^2 will becomes small like $\beta = 1/T$. This slightly dubious reasoning provides some justification for Halperin choosing $\rho_0^2 = \beta/2$.

At smaller values of T we might expect

$$\mathcal{Z}_{LSM}(e,\beta) = \int \mathcal{D}A\mathcal{D}\sigma \ \exp\left(-\int \mathrm{d}^3y \ \left[\frac{1}{2}|\vec{\nabla}\times\vec{A}|^2 + \frac{1}{2}\beta|\vec{\nabla}\sigma - e\vec{A}|^2\right]\right)$$

to lie in the same universality class as the full $\mathcal{Z}_{AH}(e,\beta)$. This is because we expect it to be identical to $\mathcal{Z}_{AH}(e,\beta)$ with radial oscillations frozen out and with large T.

The next big question to answer is how we can discretize the following partition function and stay in the same universality class. The answer is relatively simple, though there is one important change we must make! To discretize we simply define a periodic $N\times N\times N$ cubic lattice, discretizing space. Let's label the points on this lattice with tuples $(i,j,k)\in\{1,2,\ldots,N\}^3$. We then simply assign a value for σ on each lattice site, i.e. $\sigma_{i,j,k}$. The natural discretization of the vector valued real field $\vec{A}$ will likewise be just to assign it a value on each lattice site, $\vec{A}_{i,j,k}$. Since it will have three components we can also think of this as defining it on directed links between adjacent sites [2]. Finally the discretized $\vec{\nabla}$ will just be taken to be the lattice derivative, i.e. $\nabla_1\sigma_{i,j,k} = \sigma_{(i+1),j,k} - \sigma_{i,j,k}$. This is all fairly straightforward, and will nearly reproduce the same LSM written down by Halperin as equation (1) in [2]:

$$\mathcal{Z}_{LSM}(e,\beta) = \int_{-\pi}^{\pi}\frac{\mathrm{d}\sigma_j}{2\pi}\int_{-\infty}^{\infty}\mathrm{d}[\vec{A}_j]\sum_{[\vec{n}_j]=-\infty}^{\infty}\exp\left(-\sum_{j=1}^{N^3}\left[\frac{1}{2}|\vec{\nabla}\times\vec{A}_j|^2 + \frac{1}{2}\beta|\vec{\nabla}\sigma_j - 2\pi\vec{n}_j - e\vec{A}_j|^2\right]\right).$$

Note that apart from the natural discretization the only thing to have changed is the addition of the $2\pi\vec{n}$ term, where $\vec{n}\in\mathbb{Z}^3$. This isn't actually a fundamentally new field at all, it just addresses an issue that arose from discretization. Essentially this term serves to preserve the delicate topological concept of winding number that would have otherwise been destroyed by the rough discretization process. We will illustrate this as follows. Consider a vortex in the original Abelian-Higgs Model. It is a point in the two spatial dimensions of the model, but it also moves through time. Thus it will form a vortex line in the $(2+1)$d spacetime dimensions of the Abelian-Higgs Model. In the same way we expect that our Euclideanized and discretized theory will also contain vortex lines. Suppose now there is a vortex line of winding number $+1$ passing through $(1/2,1/2,0)$. Now define the path C to be $(0,0,0)\to(1,0,0)\to(1,1,0)\to(0,1,0)\to(0,0,0)$. In the continuum theory we know that

$$\oint \mathrm{d}\vec{x}\cdot\vec{\nabla}\sigma = 2\pi(+1)$$

from the winding number. But if we discretize then

$$\oint \mathrm{d}\vec{x}\cdot\vec{\nabla}\sigma = \vec{\nabla}_1\sigma_{000} + \vec{\nabla}_2\sigma_{100} - \vec{\nabla}_1\sigma_{110} - \vec{\nabla}_2\sigma_{010} = 0,$$

expanding out the lattice derivatives using our tuple notation. However, it is clear that $2\pi\oint \mathrm{d}\vec{x}\cdot\vec{n}\in 2\pi\mathbb{Z}$ for $\vec{n}\in\mathbb{Z}^3$. Thus when we discretize we need to pass from $\vec{\nabla}\sigma$ to $\vec{\nabla}\sigma_j - 2\pi\vec{n}_j$ in order to allow for vortices of arbitrary circulation between lattice points.

With some thought it should become clear that there are no other restrictions we can make on $\vec{n}$ besides it being integer.

Simulation of LSM

We've now gone most of the way to obtaining a SFT that we expect to be in the same universality class as the Abelian-Higgs Model, and that we can actually simulate. However, we follow Halperin in making a few modifications to the LSM in order to make something slightly easier to simulate. Firstly, instead of having σ take any real value in $[-\pi, \pi)$ we will choose a discrete variable that can take 100 evenly spaced values in $[-\pi, \pi)$. This will help keep our state space to a reasonable size. Now suppose we wanted to discretize the interval of values that $\vec{A}$ takes. Clearly we cannot do this, since it can take any value in the reals. In order to deal with this problem Halperin [2] decided to simulate a model that is dual to the LSM. We will give a heuristic derivation of the model, though Peskin [9] showed it more rigorously.

Since only $\vec{n}$ appears in $\mathcal{Z}_{LSM}$ and not its derivatives, then its equation of motion will be given by $\vec{\nabla}\sigma_j - 2\pi\vec{n}_j - e\vec{A}_j = 0$. Solving this for $\vec{A}_j$ reveals that $\vec{A}_j = (\vec{\nabla}\sigma_j - 2\pi\vec{n}_j)/e$. In order to subsitute this back into the equation of motion we will need to compute $\vec{\nabla} \times \vec{A}_j$. This would have been a problem in the continuum theory, since if there there were vortices then $\vec{\nabla} \times \vec{\nabla}\sigma$ would have diverged at the vortex core. However, on the lattice this term will simply vanish as we have demonstrated above. Then $\vec{\nabla} \times \vec{A}_j = -(2\pi/e)\vec{\nabla} \times \vec{n}_j$. The discretized action then becomes

$$S_E \sim \sum_{j=1}^{N^3} \frac{1}{2}\frac{4\pi^2}{e^2}|\vec{\nabla} \times \vec{n}_j|^2 + \frac{1}{2}\beta|\vec{\nabla}\sigma_j - 2\pi\vec{n}_j|^2.$$

In spite of replacing $\vec{A}$ in the first term, we will keep the second term to make sure the winding nature of vortices we worked so hard to preserve remains in the theory. We might expect this to lie in the same universality class, since it should still possess the vortex behavior of the LSM.

Next we will factor out an overall factor of $4\pi^2\beta/e^2$. At the most this should merely change our partition function by some constant. If we define $\beta' = 1/\beta$, then we will finally have a partition function given by

$$\mathcal{Z}_{\text{sim}}(e, \beta') = \sum_{[\vec{n}_j]=-\infty}^{\infty} \int_{-\pi}^{\pi} \frac{\mathrm{d}\sigma_j}{2\pi} \exp\left(-\sum_{j=1}^{N^3} \left[\frac{1}{2}\beta'|\vec{\nabla} \times \vec{n}_j|^2 + \frac{e^2}{8\pi^2}|\vec{\nabla}\sigma_j - 2\pi\vec{n}_j|^2\right]\right).$$

Following Halperin [2] and Peskin [9], we expect $\mathcal{Z}_{LSM}(e, \beta) \propto (2\pi\beta)^{-3N^3/2}\mathcal{Z}_{\text{sim}}(e, 1/\beta)$. Again, we won't derive this more rigorously as it was treated exhaustively by Peskin.

It is at this point where we make our own addition to the model. Though Halperin discretized σ, he allowed n to take any integer value. To keep this introductory we don't want to deal with this as it would leave our state space infinite. It is possible to modify the

Metropolis algorithm to incorporate an infinite state space, but that's more complicated than we want to address here. As such we will choose $\vec{n} \in \{-1, 0, +1\}^3$. With these choices to discretize $\vec{n}$ and σ, we have finally reduced ourselves to a finite state space.

We can now describe the slight changes that need to be made to the Metropolis algorithm we gave for the Ising model in order to explore $\mathcal{Z}_{\text{sim}}$. Most of the ajustments occur in the third step where we need to deal with our larger state space. This algorithm is a pseudocode description of the function `mcmove` defined in `LSM_crit_exp.py` found in my GitHub repository [6].

Metropolis algorithm: $\mathcal{Z}_{\text{sim}}$

1. Choose an initial state $\{\vec{n}, \sigma\}$ drawn from a uniform distribution.
2. Next pick a site $(i, j, k) \in \{1, 2, \ldots, N\}^3$ at random.
3. Now randomly choose a new value for σ and $\vec{n}$ at the site (i, j, k). That is to say choose uniformly from the 100 possible values for the discretized $\sigma \in [-\pi, \pi)$ and the 27 possible values for $\vec{n} \in \{-1, 0, +1\}^3$. Then compute $\Delta S_E = S_E(\{\vec{n}, \sigma\}_{\text{final}}) - S_E(\{\vec{n}, \sigma\}_{\text{initial}})$.

   The computation of ΔS_E can be made significantly faster by only computing the terms that depend on the site (i, j, k) instead of trying to compute all terms. Since anyone wanting to write code that will run in a reasonable amount of time will need to write out all the terms that depend on (i, j, k), we will write them out here in the interest of completeness. Below are the terms of S_E that depend only on (i, j, k). To find this we used $\vec{n}_{i,j,k} = (n^1_{i,j,k}, n^2_{i,j,k}, n^3_{i,j,k})$ and we reiterate that the lattice derivative is defined so $\nabla_1 \sigma_{i,j,k} = \sigma_{i+1,j,k} - \sigma_{i,j,k}$, likewise with the other components.

$$
\begin{aligned}
(S_E)_{i,j,k} =& \frac{e^2}{8\pi^2} \Big\{ [(\sigma_{i,j,k} - \sigma_{i-1,j,k}) - 2\pi n^1_{i-1,j,k}]^2 + [(\sigma_{i,j,k} - \sigma_{i,j-1,k}) - 2\pi n^2_{i,j-1,k}]^2 + \\
&+ [(\sigma_{i,j,k} - \sigma_{i,j,k-1}) - 2\pi n^3_{i,j,k-1}]^2 + [(\sigma_{i+1,j,k} - \sigma_{i,j,k}) - 2\pi n^1_{i,j,k}]^2 \\
&+ [(\sigma_{i,j+1,k} - \sigma_{i,j,k}) - 2\pi n^2_{i,j,k}]^2 + [(\sigma_{i,j,k} - \sigma_{i,j,k-1}) - 2\pi n^3_{i,j,k-1}]^2 \Big\} \\
&+ \frac{\beta'}{2} \Big\{ [(n^3_{i,j+1,k} - n^3_{i,j,k})^2 - (n^2_{i,j,k+1} - n^2_{i,j,k})]^2 + [(n^1_{i,j,k+1} - n^1_{i,j,k}) + (n^3_{i+1,j,k} - n^3_{i,j,k})]^2 \\
&+ [(n^2_{i+1,j,k} - n^2_{i,j,k}) - (n^1_{i,j+1,k} - n^1_{i,j,k})]^2 + [(n^1_{i-1,j,k+1} - n^1_{i-1,j,k}) - (n^3_{i,j,k} - n^3_{i-1,j,k})]^2 \\
&+ [(n^2_{i,j,k} - n^2_{i-1,j,k}) - (n^1_{i-1,j+1,k} - n^1_{i-1,j,k})]^2 + [(n^3_{i,j,k} - n^3_{i,j-1,k}) - (n^2_{i,j-1,k+1} - n^2_{i,j-1,k})]^2 \\
&+ [(n^2_{i+1,j-1,k} - n^2_{i,j-1,k}) - (n^1_{i,j,k} - n^1_{i,j-1,k})]^2 + [(n^3_{i,j+1,k-1} - n^3_{i,j,k-1}) - (n^2_{i,j,k} - n^2_{i,j,k-1})]^2 \\
&+ [(n^1_{i,j,k} - n^1_{i,j,k-1}) - (n^3_{i+1,j,k-1} - n^3_{i,j,k-1})]^2 \Big\} .
\end{aligned}
$$

Note that just as in the Ising model $i \pm 1, j \pm 1$, and $k \pm 1$ are defined mod N, as the lattice is periodic. Now using the above function we can efficiently find ΔS_E by noting that $\Delta S_E = (S_E)_{i,j,k}(\{\vec{n}, \sigma\}_{\text{final}}) - (S_E)_{i,j,k}(\{\vec{n}, \sigma\}_{\text{initial}})$ since we have only changed the variables at the site (i, j, k). We now use ΔS_E to take the Metropolis step.

(i) If ΔS_E is less than zero, keep the new values for σ and $\vec{n}$ at (i, j, k).

(ii) If $\Delta S_E \geq 0$, then choose a random number u from a uniform distribution between 0 and 1. If $u < e^{-\Delta S_E}$ then keep the new values for σ and $\vec{n}$ at the site (i, j, k). If $e^{-\Delta S_E} \leq u$ then stick with the original values.

4. Return to step 2 and continue changing sites. If this has been done the appropriate number of times, we can exit here.

We can see that the Metropolis algorithm for exploring $\mathcal{Z}_{\text{sim}}$ is not that different from the Metropolis algorithm for the Ising model. The only real difference was that the state space was (much) larger and the energy function was more complicated.

It would be hard to visualize these states. Unlike the Ising model which took place on a two-dimensional lattice and involved only a single parameter, this model takes place on a three-dimensional lattice and involves four parameters. This makes it less amenable to plotting than the Ising states in Figure 2. Thus we will move straight into evaluating various temperature dependent quantities of the model.

We won't be as explicit in describing the algorithm we used to find the various temperature dependent quantities as we were for the algorithm that found $C_V(T)$ and $|\langle M \rangle|(T)$ for the Ising model. This is because our algorithm is essentially identical. We: 1) pick a specific temperature, 2) generate a random state, 3) equilibrate the state for a time $\tau_{\text{eq}}(T)$ using the Metropolis algorithm described above, 4) sample the quantities of interest of the state for τ_{calc} successive applications of the Metropolis algorithm, 5) compute appropriate averages of the quantities of interest, and return to 1) with the next temperature value. Following Halperin, we chose to simulate all of our systems with $e^2 = 5$. Further we chose our equilibration time to be $\tau_{\text{eq}}(T) = 5000$ steps and our calculation time to be $\tau_{\text{calc}} = 5000$ steps. Since Halperin chose on the order of 10^4 steps for his simulation, this should be a reasonable choice [2].

Note that we chose to make $\tau_{\text{eq}}(T)$ a constant instead of the time dependent quantity it was for the Ising model. This is because simulating the LSM takes much longer computationally. Not only is ΔS_E more complicated, but because the system is three dimensional each Metropolis algorithm step across the whole lattice will take on the order of N^3 steps instead of the N^2 steps for the Ising model. Thus it would take a prohibitively long time to allow τ_{eq} to increase exponentially near T_c as we did for the Ising model. But this means we are likely to miss any critical slowing down behavior near T_c, and our states with temperature near T_c might not have enough time to equilibrate there. This lack of equilibration time means that we might obtain greater variance in our quantities of interest near T_c than we would otherwise, since we will be measuring the quantities against states that have not fully equilibrated. Thus we will need to take more samples of the temperature in order to get reasonable averages for fitting.

Having discussed the important parameters of the simulation we now move on to discussing some of its calculated quantities of interest, starting with its action. In SFT

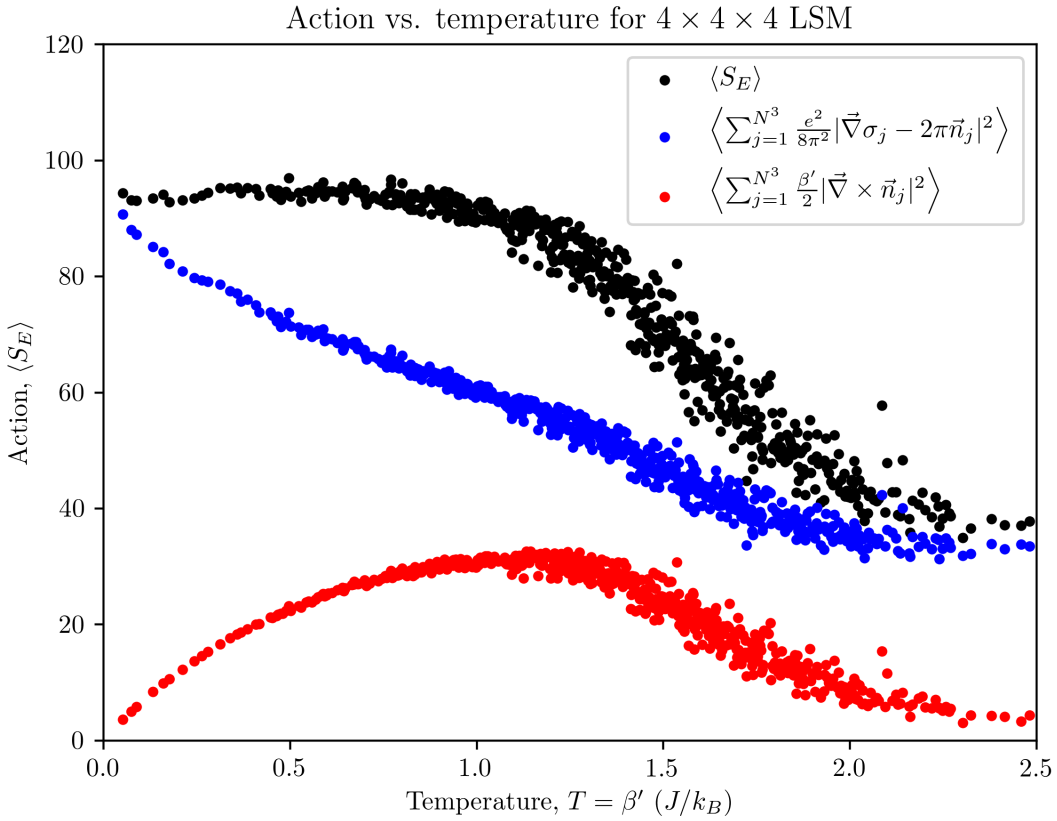

FIGURE 4. Scan of temperature axis for $4 \times 4 \times 4$ LSM. Full action shown in black, with covariant derivative piece and Maxwell piece in blue and red, respectively.

the partition functions we encountered looked like $\mathcal{Z} = \int \mathcal{D}\phi \, e^{-\beta E[\phi]}$, where the Euclidean action was given by $S_E = \beta E[\phi]$. Since we can't pull a β' factor out of S_E for $\mathcal{Z}_{\text{sim}}$, we will need to evaluate the expectation value of the action instead of the energy. By linearity we note that

$$\langle S_E \rangle = \left\langle \sum_j \frac{1}{2}\beta' |\vec{\nabla} \times \vec{n}_j|^2 \right\rangle + \left\langle \sum_j \frac{e^2}{8\pi^2} |\vec{\nabla}\sigma_j - 2\pi\vec{n}_j|^2 \right\rangle.$$

We call the first term on the right hand side the Maxwell term, since it came from the Maxwell term $|\vec{\nabla} \times \vec{A}_j|^2$ in the LSM. The second term will be called the covariant derivative term, since it arose from the covariant derivative in the Abelian-Higgs Model.

The plot of $\langle S_E \rangle$ vs. temperature T is shown as Figure 4. We obtained the plot by taking several hundred temperature values between $T = 0$ and $T = 2.5$ for an $N = 4$ cubic lattice. Recall that since Peskin showed $\mathcal{Z}_{LSM}(e, \beta) \propto \mathcal{Z}_{\text{sim}}(e, 1/\beta)$ then T for the LSM will be equivalent to β' for $\mathcal{Z}_{\text{sim}}$. Since we expect the LSM to be the model that lies in the same universality class as the Abelian-Higgs Model, we will plot all of our quantities vs. $T = \beta'$ instead of $T' = 1/\beta'$.

We can use the Maxwell term as an excellent sanity check on our simulation. Firstly, we note that as $\beta' \to \infty$ the Maxwell term approaches zero. This makes sense; $\beta' \to \infty$ should freeze out any fluctuations with $\vec{\nabla} \times \vec{n}_j \neq 0$ since any fluctuation of this form will begin to cost near infinite energy. We also note that for $\beta' \to 0$ the Maxwell term begins to approach zero as well. This is a little less intuitive. Near $\beta' = 0$ fluctuations with $\vec{\nabla} \times \vec{n}_j \neq 0$ will cost very little energy, and so we expect the states to take larger and larger values of $|\vec{\nabla} \times \vec{n}_j|^2$ as $\beta' \to 0$. In other words we might expect $|\vec{\nabla} \times \vec{n}_j|^2 \sim \beta'^{-c}$ for some positive c. But the Maxwell term looks like $\beta'|\vec{\nabla} \times \vec{n}_j|^2 \sim \beta'^{(1-c)}$, so provided $0 < c < 1$ we would expect it to approach zero as $\beta' \to 0$ but to do so slower than linearly. This does seem to be the case for $0.3 < \beta' < 1.3$, but for $\beta' < 0.3$ the descent of the Maxwell term seems instead to be linear. Why is this? It is simply because we've chosen $\vec{n} \in \{-1, 0, 1\}^3$. This means that $|\vec{\nabla} \times \vec{n}_j|^2$ has an upper bound (of $3(4)^2$ in fact). But this means that the divergence of $|\vec{\nabla} \times \vec{n}_j|^2$ is cut-off at a large positive value, so β'^{-c} stops describing its magnitude for small β'. Consequently, the Maxwell term $\beta'|\vec{\nabla} \times \vec{n}_j|^2 \sim \beta'$ for small β', and we obtain the noted linear descent. Having explained most of the features of the Maxwell term, we complete our sanity check and move on to the full action.

We note that the action displays many of the same properties that we saw displayed by the Ising model energy in SFT. It starts off at a level value at $T = \beta' = 0$, around 95, before plunging rapidly around $T = \beta' = 1.5$. It then hits a lower second asymptote at around 40 for $T = \beta'$ large. The Ising model energy had this same form, but with an increasing energy instead of a decreasing energy. Additionally, the slope of the Ising model's energy diverged near the critical temperature, where the energy crossed over from small to large. Indeed since the heat capacity C_V looks like the derivative of the energy, and since C_V exhibited logarithmic divergence near T_c, the slope of the energy for the Ising model had to diverge at T_c. Thus we might be slightly concerned that the lack of divergence of the slope of $\langle S_E \rangle$ indicates that the specific heat of the LSM will not diverge near T_c, which would contradict its being dual to the XY-Model. However, we shouldn't worry too much. Just as the Maxwell term's divergence near β' was limited by the finite state space of $\vec{n}$, we expect the divergence of C_V to be constrained by finite size effects. In fact Figure 3 of Halperin demonstrates that the height of C_V near T_c grows steeper with larger system sizes N [2]. Thus when it comes time for us to measure α we needn't worry about the slope of $\langle S_E \rangle$ not diverging near T_c, we just need to choose a larger system size than $4 \times 4 \times 4$ to capture this divergence. For now though we would like to measure β for our system.

We will continue to use the plot of energy in Figure 4 for guidance. Recall that for the Ising model the higher energy coincided with the disordered phase, and the lower energy the ordered phase. Thus we suspect that for $T < T_c \approx 1.5$ we will be in the disordered phase, and for $T > T_c \approx 1.5$ we will be in the ordered phase. This makes some sense. For $T = \beta'$ small we have no control over $|\vec{\nabla} \times \vec{n}_j|$, and it can take any value it likes. This

leads directly to a disordered phase. As $\beta' \to \infty$ we discussed above that $\beta'|\vec{\nabla} \times \vec{n}_j|^2 \to 0$, leading to a more ordered phase where there is control over the curl of $\vec{n}$.

This simple logic makes sense, but in order to actually measure the critical exponent β we need to know what the order parameter is. For the Abelian-Higgs Model the order parameter is given by the complex scalar field ψ. This follows from the standard fact that in the Landau-Ginzburg theory of superconductivity the order parameter is given by the complex scalar field ψ, and the Abelian-Higgs Model coincides with the low-energy description of a superconductor near its critical point. Now we know that $\psi = \rho e^{i\sigma}$ in our notation. Then since we froze out radial vibrations to arrive at the LSM, we suspect that the LSM will have an order parameter given by $\rho_0 e^{i\sigma}$. In what follows we will drop the ρ_0 and simply refer to the order parameter as $e^{i\sigma}$.

But if the order parameter is $\psi \propto e^{i\sigma}$, then how does this square with our earlier statement that the ordered and disordered phases should be characterized by the size of $|\vec{\nabla} \times \vec{n}_j|$? Recall that in the theoretical discussion of the Abelian-Higgs Model we showed that the covariant derivative term $\rho^2(\partial_\mu \sigma - e a_\mu)^2$ allowed for the gauge field to couple to the phase of the order parameter ψ. The same thing will happen for the LSM. The covariant derivative term $|\vec{\nabla}\sigma_j - 2\pi\vec{n}_j|^2$ will allow the integer winding numbers $\vec{n}$ to couple to the phase of $e^{i\sigma}$. Thus we might expect that changing $|\vec{\nabla} \times \vec{n}_j|$ will lead to a consequent change in σ through the covariant derivative term. We can see some hint of this in Figure 4. We've already seen that as we increase β' the Maxwell term decreases, but note that the blue covariant derivative term will also decrease. This is in spite of the fact that there is no β' factor in front of the covariant derivative term to freeze out vibrations with $\vec{\nabla}\sigma - 2\pi\vec{n} \neq 0$! Instead, since there is no accompanying β' in from of the covariant derivative term, all of the decrease in the covariant derivative term with β' must come from its coupling to $\vec{n}$. Thus the average size of $|\vec{\nabla} \times \vec{n}|$ appears to have a direct effect on the average size of $|\vec{\nabla}\sigma - 2\pi\vec{n}|$, and it stands to reason that in a disordered phase with $|\vec{\nabla} \times \vec{n}|$ large the order parameter $e^{i\sigma}$ will indeed look different than in an ordered phase with $|\vec{\nabla} \times \vec{n}| = 0$.

That $e^{i\sigma}$ does represent a good order parameter can be seen in Figure 5. The magnitude of its average at different temperatures $|\langle e^{i\sigma} \rangle|$ is displayed. It clearly has the same type of behavior, with the temperature axis inverted, that $|\langle M \rangle|$ did in Figure 3(B). In the disordered phase, $T < T_c$, it has a value of zero and jumps very quickly at $T = T_c$ to a nonzero value in the ordered phase. This is analogous to the Ising model having zero magnetization in the disordered phase, but jumping to a nonzero magnetization in the ordered phase. However, we do notice that there is a lot more spread around this jump than there was for the Ising model in Figure 3. We already mentioned that this was likely because of the possibility that states near T_c wouldn't fully equilibrate, due to our choosing to keep τ_{eq} temperature independent. As such we took many more data points and binned them into intervals on the temperature axis. Binning the data allowed us to

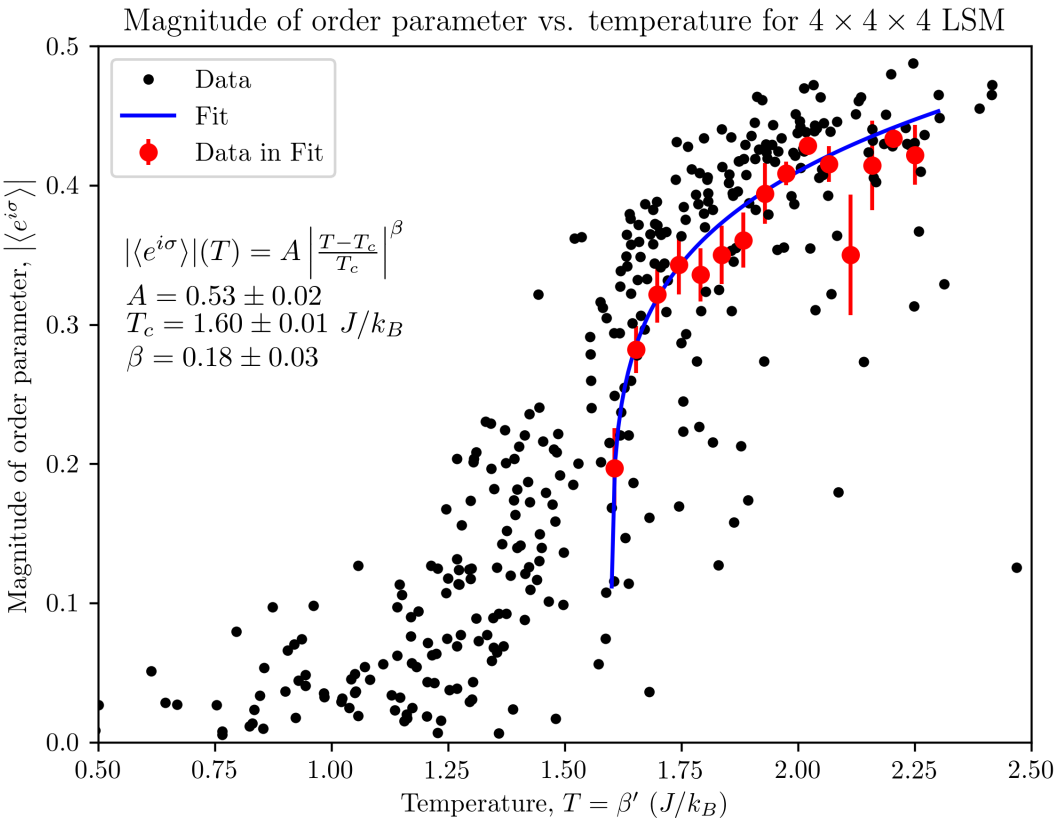

FIGURE 5. Determination of β for the LSM. The fit demonstrates that $T_c \approx 1.6$ and $\beta \approx 0.18$. However, β's value may be affected by finite size effects.

display the mean of the bins as the red data points, with their errorbars corresponding to the standard deviation of the mean value in the bins. These red data points give a clearer picture of the underlying behavior of the function, and allow a better fit to be made.

The fit is displayed in blue, showing that $|\langle e^{i\sigma}\rangle|$ did seem to behave like $|(T - T_c)/T_c|^{\beta}$ near T_c based on how well the fit appears to agree with the binned data in red. Reading off from the fit allows us to conclude that $T_c \approx 1.6$ and $\beta \approx 0.18$. This is certainly greater than the β of around 0.13 we found for the Ising model. However, it is nowhere near as large as the $\beta = 0.35$ value that the XY-Model possesses, as mentioned in our discussion of the XY-Model's critical exponents. At first glance this would see to be a bitter blow to the expectation that the LSM is dual to the XY-Model near the critical point. However, just as the derivative of $\langle S_E\rangle$ near T_c failed to diverge due to finite size effects, we suspect that finite size effects are at play here. As we mentioned, making N larger will cause C_V to becomes steeper near T_c, and we likewise expect that making N larger would lead to a steeper transition for $|\langle e^{i\sigma}\rangle|$ near T_c. This would in turn increase the critical exponent β. Unfortunately due to time constraints we were not able to test this. However, even

without showing $\beta_{LSM} = \beta_{XY}$ we can take heart in at least two facts. The first is that Halperin did not attempt to determine β for the LSM, so our result is at least not in conflict with this previous study of the LSM. The second is that our critical temperature $T_c \approx 1.6$ is extremely close to the critical temperature $T_c = 1.62$ that Halperin found [2]. This is in spite of the fact that we discretized the values $\vec{n}$ could take, and he did not. This makes us hopeful that we will be able to demonstrate the same phase transition he observed in $C_V(T)$, and show $\alpha \simeq 0$.

As mentioned in our discussion of the slope of $\langle S_E \rangle$, we will likely need to choose a larger system size than $4 \times 4 \times 4$ to capture the divergence of $C_V(T)$. We therefore chose to simulate a system with size $5 \times 5 \times 5$, this will have nearly twice the number of points as the $4 \times 4 \times 4$ system. Since the Metropolis algorithm's computation time scales with system size, this also means that simulating the system will take nearly twice as long. These time constraints forced us to take fewer temperature values. The temperature values we did take were clustered around the $T_c = \beta_c' \approx 1.6$ value that we found while fitting to find β. The plot of $C_V(T)$ vs. T for this simulation is shown in Figure 6, along with the appropriate fit. We have binned and averaged the data for our fit, just as we did our plot to find β.

Figure 6 mostly confirms the information found in Figure 3 of [2]. Firstly it appears to have a log divergence. This is suggested by the fit, which seems to fit the underlying data fairly well, despite the anticipated large spread. As with the Ising model it allows us to conclude that $\alpha \simeq 0$. We noted in the discussion of the critical exponents of the XY-Model that the XY-Model had a critical exponent $\alpha \approx -0.01$ which is very close to zero. Indeed Halperin takes the fact that his fits show a log divergence to support that the α value for the LSM is the same as the α value of the 3D XY-Model [2]. This is perhaps slightly dubious, but considering our value for β was off by 0.17 we will also consider obtaining a value $\alpha \simeq 0$ that is within -0.01 of the predicted value to be a success.

Secondly our value for $T_c = 1.51$ is again reasonably close to Halperin's value of $T_c = 1.62$. However, we note that it is less than the value $T_c \approx 1.6$ that we found in our fit for β where the system size was $N = 4$. This should not concern us a great deal. Figure 3 in Halperin reveals that increasing the system size N slightly lowers the critical temperature. The critical temperature does eventually approach an asymptotic limit as $N \to \infty$; one that is relatively close to its values for smaller N [2]. Thus the fact that our value for T_c is 0.09 less for the $N = 5$ system than it was for the $N = 4$ system is not surprising, and in fact provides more evidence that we are seeing the same transition that Halperin found.

Finally, it seems clear from looking at the binned data in red that the specific heat for $T = T_c + \Delta T$ is higher than that for $T = T_c - \Delta T$. This can be seen from our fit by noting that the fitting parametery $D < 0$. As does Halperin we note that this is the reverse of what is expected for the XY-Model [2]. This provides evidence that we are in fact seeing

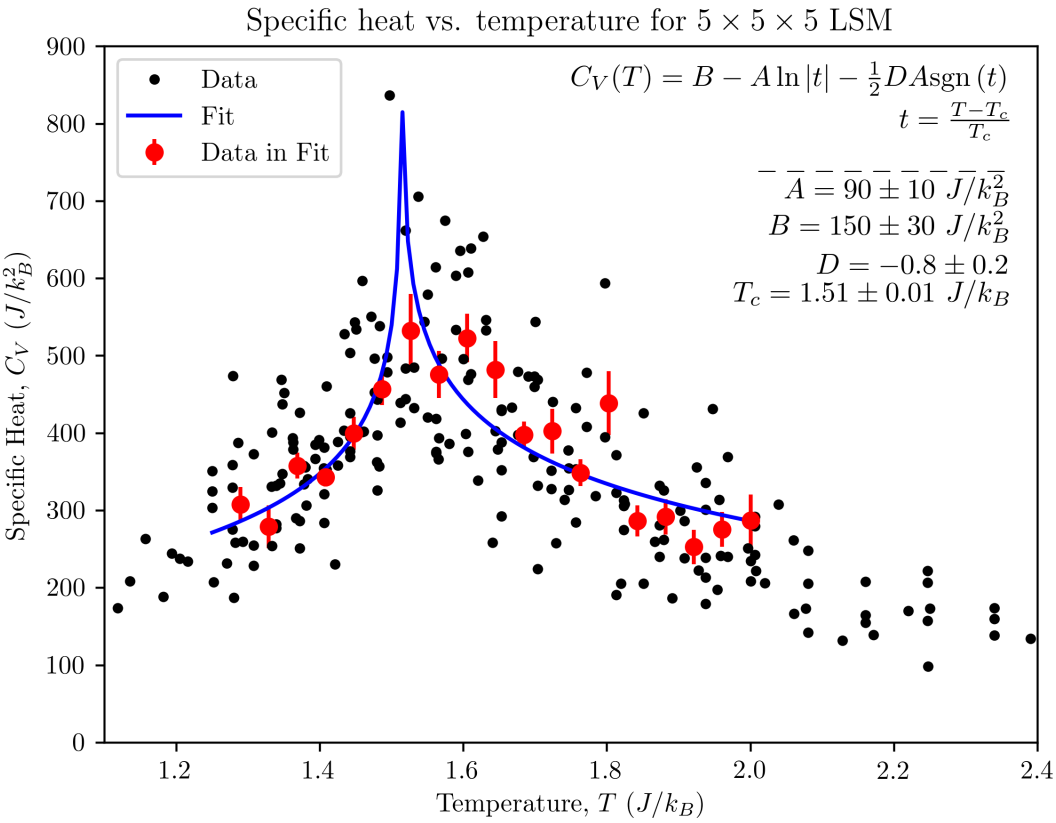

FIGURE 6. Scan of temperature axis for $5 \times 5 \times 5$ LSM. Binning was used to obtain the points with errorbars in red. Fit demonstrates that $\alpha \simeq 0$.

the same transition as the XY model, but with the temperature axis reversed around T_c. In particular, recall that for the Abelian-Higgs Model we suspected that $\tilde{a} \propto T_c - T$. If the same is true for the XY-Model, but with the temperature axis reversed around T_c, then we would expect $a \propto T - T_c \propto -\tilde{a}$. This is precisely what the table comparing the phases of the Abelian-Higgs Model and the XY-Model suggested!

We have thus at last come to a place where we have our own numerical evidence for the LSM undergoing an inverted XY-Model type transition. We therefore expect that the Abelian-Higgs Model will have an inverted XY-Model transition as well, since the arguments we gave in this section should mean that the LSM and the Abelian-Higgs Model will lie in the same universality class. Then our numerics do indeed suggest that the XY-Model and the Abelian-Higgs Model are dual, with $\tilde{a} \sim -a$ near the transition! This provides an excellent additional piece of evidence for believing the particle vortex duality that we motivated theoretically in the first section. We will conclude on this high note.

References

[1] M. Campostrini, M. Hasenbusch, A. Pelissetto, P. Rossi, and E. Vicari. Physical Review B **63** Iss. 21 (2001)

[2] C. Dasgupta and B. I. Halperin. Physical Review Letters **47** (1981) pg. 1556 - 1560

[3] A. Karch and D. Tong. Physical Review X **6** Iss. 3 (2016)

[4] M. Kardar. *Statistical Physics of Fields* (2007) pg. 114-117

[5] M. Metlitski and A. Vishwanath. Physical Review B **93** Iss. 24 (2016)

[6] S. Musser https://github.com/smusser/ParticleVortexEssay.git (2018)

[7] M. Newman and G. Barkema. *Monte Carlo Methods in Statistical Physics* (1999) ch. 3-4

[8] H. Nielsen and P. Olesen Nuclear Physics B **61** (1973) pg. 45-61

[9] M. Peskin. Annals of Physics **113** Iss. 1 (1978) pg. 122-152

[10] R. Rinet https://rajeshrinet.github.io/blog/2014/ising-model/ (2014)

[11] D. Tong. Quantum Hall Effect Notes.

[12] D. Tong. Statistical Field Theory Notes.