

Measuring the broken phase sphaleron rate nonperturbatively

Guy D. Moore*

Department of Physics, McGill University, 3600 University Street, Montreal, PQ, Canada H3A 2T8

(Received 11 May 1998; published 19 November 1998)

We present details for a method to compute the broken phase sphaleron rate (rate of hot baryon number violation below the electroweak phase transition) nonperturbatively, using a combination of multicategorical and real time lattice techniques. The calculation includes the ‘‘dynamical prefactor,’’ which accounts for prompt recrossings of the sphaleron barrier. The prefactor depends on the hard thermal loops, getting smaller with increasing Debye mass; but for realistic Debye masses the effect is not large. The baryon number erasure rate in the broken phase is slower than a perturbative estimate by about $\exp(-3.6)$. Assuming the electroweak phase transition has enough latent heat to reheat the universe to the equilibrium temperature, baryon number is preserved after the phase transition if the ratio of (‘‘dimensionally reduced’’ thermal) scalar to gauge couplings λ/g^2 is less than 0.037. [S0556-2821(98)00523-2]

PACS number(s): 11.15.Ha, 11.15.Kc

I. INTRODUCTION

20 years ago, 't Hooft showed that baryon number is not a good quantum number in the standard model [1]. The reason involves the nontrivial vacuum structure of the SU(2) (weak) gauge group of the standard model. In any gauge theory, the vacuum is not unique; any gauge transformation of $\vec{A}=0$ has zero energy and is an acceptable vacuum. But SU(2) (and any simple gauge group) has the property that the space of three-dimensional (3D) gauge transformations is topologically nontrivial. A gauge transformation has an integer π_3 winding number associated with it. Since the winding number must be an integer, the space of smooth gauges, and also the space of vacua, is disconnected. The different connected components are characterized by their values of Chern-Simons number,

$$N_{CS} \equiv \frac{g^2}{32\pi^2} \int d^3x \epsilon_{ijk} \left(F_{ij}^a A_k^a - \frac{g}{3} f_{abc} A_i^a A_j^b A_k^c \right), \quad (1)$$

which is an integer for a vacuum configuration, though not necessarily for an excited state. Classically, for the gauge fields to change from one topological vacuum at time $t=t_i$ to another at time $t=t_f$, they must pass through excited states in the intervening time; to be specific,

$$dN_{CS}/dt = \frac{g^2}{32\pi^2} \int d^3x F_{\mu\nu}^a F_a^{\mu\nu}, \quad (2)$$

which is clearly a gauge invariant quantity (though its integral N_{CS} is not, because of the constant of integration). This means that it is possible to pass from a vacuum configuration to a gauge copy of that configuration via a path which cannot be smoothly deformed to remain always in vacuum. If we mod out the space of 3D configurations by the gauge transformations, the space we get will then have noncontractible loops.

This topological structure appears both in SU(2) (weak) and SU(3) (strong), where it is responsible for the physics of spontaneous chiral symmetry breaking. What 't Hooft noticed is that, because fermions couple to the weak SU(2) group of the standard model chirally, the anomaly relates N_{CS} to baryon number. If the gauge fields pass through some nonvacuum intermediate state from one topological vacuum to another (or around a noncontractible loop, if we think of configurations modulo gauge transformations), baryon number changes. Such changes are classically forbidden at zero temperature, so they only occur via quantum tunneling. Because the SU(2) gauge coupling is weak, and because the Higgs field breaks the symmetry, such processes are steeply exponentially suppressed, by $\sim \exp(-16\pi^2/g^2) \sim 10^{-170}$. Hence such processes are of no terrestrial phenomenological interest.

However, as a general rule, if a process only occurs in vacuum via quantum tunneling, then above some temperature it occurs much faster via thermal activation. (Chemistry and condensed matter physics are full of examples; annealing of crystal defects, for instance.) The same is true for baryon number changing processes in the standard model, although the ‘‘annealing temperature’’ is $O(100 \text{ GeV})$. In fact, there is a phase transition at $T_c \sim 100 \text{ GeV}$ in which the Higgs field loses its condensate, and above this transition baryon number violation is efficient [2,3]. It is quite possible that the baryon number of the universe originated in a cosmological electroweak phase transition, and in recent years this belief has driven the study of the electroweak phase transition.

It takes two things for a cosmological electroweak phase transition to generate the baryon asymmetry of the universe. First, there needs to be enough CP violation to generate at least the observed abundance of baryons during the transition. Second, there cannot be too much ‘‘annealing’’ of the baryon number after the phase transition, that is, baryon number violation must be inefficient enough after the phase transition that a good fraction of the baryons survive to the present day. This is a condition on the phase transition's strength.

*Email address: guymoore@physics.mcgill.ca

TABLE I. Perturbation theory versus nonperturbative Γ_d . Appearances of T_c^{-4} are really $(2.5g^2T_c)^{-4}$. $x \equiv \lambda/g^2$ is the ratio of the Higgs boson self-coupling to the gauge coupling, and $y \equiv m_H^2(T)/g^4T^2$ is the dimensionless Higgs boson mass squared. The error bars for the nonperturbative ϕ_0 are dominated by statistical errors in the determined value of T_c ; errors in the nonperturbative value of Γ are statistical errors from the Monte Carlo calculation. The last row includes the nonunity dynamical prefactor in the rate and should be taken as our most reliable estimate.

	$(x \equiv \lambda/g^2)$	$x=0.047$	$x=0.039$	$x=0.033$
two loop perturbative	$\phi(T_c)/gT_c$	1.360	1.568	1.789
	$B \equiv gE_{\text{sph}}/4\pi\phi$	1.643	1.633	1.626
	E_{sph}/T_c	28.08	32.20	36.55
	$-\ln(\Gamma_d T_c^{-4})$	22.27	25.39	28.82
	$-d \ln(\Gamma_d T_c^{-4})/dy$	860	920	1000
nonperturbative	$\phi(T_c)/gT_c$	1.38 ± 0.02	1.60 ± 0.01	1.82 ± 0.03
	$-\ln(\Gamma_d T_c^{-4})$ (excl. prefactor)	24.7 ± 0.4	28.3 ± 0.4	31.2 ± 0.6
	$-\ln(\Gamma_d T_c^{-4})$ (incl. prefactor)	25.9 ± 0.5	29.5 ± 0.5	32.4 ± 0.7

The minimal standard model fails both conditions badly [4,5]. However, well motivated extensions, such as the minimal supersymmetric extension of the standard model (MSSM) with a light right scalar top, appear to be viable. Recent studies of baryon number production during the phase transition appear to show that enough CP violation can hide in places with few low energy consequences to generate the observed abundance of baryon number, and maybe a little more (for recent work see, for instance, Refs. [6–8]). And the phase transition can be stronger. If the lightest scalar top quark is not very light, then perturbation theory can reliably relate the phase transition in the MSSM to the phase transition in the same effective theory used to study the standard model [9], which has been well analyzed numerically [5,10,11]. If the lightest scalar top quark is lighter still, the phase transition may be stronger and more exotic [12]. This system can also be studied by nonperturbative lattice techniques [13].

At present, the weakest link in our knowledge of baryon erasure after the phase transition is the relation between the strength of the phase transition, now known nonperturbatively, and the efficiency of baryon number violation after it is over, for which we have only a one loop calculation [3,14,15]. We know that the perturbation expansion at high temperature near the electroweak phase transition cannot be viewed as an expansion in \hbar , but at best as an expansion in a ratio of couplings. We also know that the two loop corrections in the perturbative expansion for the strength of the phase transition are not very small in the “interesting” range of couplings where the baryon number violation after the transition is close to the efficiency limit. So it would be nice to actually know how good the one loop calculation is, or to replace it with a fully nonperturbative investigation.

Very recently we have proposed a nonperturbative method to determine the rate of baryon number violation in the broken electroweak phase [16]. This paper will fill in all the details left out in that paper. Also, the calculation there was incomplete; it did not include a measurement of the “dynamical prefactor,” discussed below. This paper will complete this aspect of the calculation. It will also discuss

the importance to the broken phase sphaleron rate of hard thermal loops, which can modify the dynamical prefactor.

For the impatient reader, we will present the basic ideas and the results right now. To find the sphaleron rate nonperturbatively, we first define nonperturbatively a surface called the separatrix, which sits half way between distinct topological vacua. Sphaleron transitions which permanently change N_{CS} must cross this surface. To find the N_{CS} diffusion rate, we first compute the probability in the canonical ensemble to be in a narrow band about the separatrix; then we compute the mean inverse time to cross the band. The product is the probability flux across the separatrix. Then we compute a “dynamical prefactor,” which tells what fraction of crossings lead to permanent resettling about a different vacuum. All three quantities can be computed nonperturbatively on the lattice, using a combination of Monte Carlo and real time techniques. Including strong hard thermal loop effects modifies the dynamics in a way which lowers the dynamical prefactor, but for realistic parameter values the effect is minor. The N_{CS} diffusion constant is presented, and compared to an analytic estimate based on the two loop effective potential, in Table I.

The paper is structured as follows. In Sec. II, we outline the general idea of the calculation. Section III defines Chern-Simons number on the lattice, and the order parameter we will use, which is very closely related. It also discusses application of the definition to the symmetric phase case. Section IV tells how we go about things numerically. Section V presents numerical results and compares them to a “semi-two loop” analytic estimate, and to the erasure bound. The last section concludes. For readers who are allergic to details of numerical studies, we recommend reading Sec. II carefully, and perhaps the first subsection of Sec. IV, and then skipping to Sec. V.

II. BROKEN PHASE MEASUREMENT: GENERAL IDEA

We want a technique for determining the N_{CS} diffusion constant in the broken electroweak phase, where the rate is extremely small. The technique will be geared around the smallness of the rate and the fact that the system in finite volume will spend almost all of its time in a “neighbor-

hood” of a topological vacuum, in a sense to be made precise below. These assumptions can be checked *a posteriori*, and do not constitute a real limit to the technique in the broken phase. They will fail in the symmetric phase or when the phase transition is very weak, but in that case we can apply real time techniques [17–23], which can now produce quantitative results [24–26]. (We should mention here that the symmetric phase case is not completely settled; it has recently been argued that there are logarithmic corrections to the parametric scaling behavior [27], which are however too small to be seen over the noise and other systematics still present in [26].)

A. Thermodynamic approximations

Before we start to describe our approach to the calculation, we will specify the approximations to be made.

We treat the thermodynamics of the standard model, or whatever extension is of interest, in the dimensional reduction approximation [28–30], that is, as being well approximated by a three-dimensional, bosonic path integral with parameters carefully matched to those of the full theory. This is an excellent approximation and we have no regrets in making it. For a study of corrections to this approximation in the present context, see Ref. [31], which shows that the leading thermodynamic effects not included in the dimensional reduction procedure have a negligibly small effect on the sphaleron energy.

Conveniently, dimensional reduction is equivalent to treating the theory’s thermodynamics as equivalent to those of the classical bosonic theory [19], with certain mass counterterms. Similarly, we can treat the theory’s dynamics in a classical approximation. This should be valid in the infrared [23,32], with one serious complication. That is, the structure of divergent radiative corrections to unequal time correlators is much more complicated than that for equal time correlators. For the equal time correlators, which are all that matter to thermodynamics, the divergent radiative corrections are mass squared corrections for the Higgs and A_0 fields, which can be computed once and balanced by counterterms. For unequal times the linearly divergent radiative corrections, the hard thermal loops, have a more complicated structure. And they can significantly change the dynamical behavior on time scales longer than the inverse plasma frequency. The modifications can be important for the Chern-Simons number diffusion rate [33–35,27], as has recently been verified numerically for the symmetric phase dynamics [26]. In the current context they will modify the “dynamical prefactor,” but they have little bearing on those parts of the calculation which are thermodynamical.

We will also frequently make the approximation that the so called “heavy” degrees of freedom can be integrated out [30], including both the time component of the gauge fields and any squarks present, so the theory reduces to an effective theory for the minimal standard model [9]. This theory is specified by two parameters; $x \equiv \lambda/g^2$ the ratio of scalar to gauge self-couplings, and $m_H^2(T)$ the thermal Higgs boson mass squared Lagrangian parameter. $m_H^2(T)$ is a monotone increasing function of T , going from $O(g^2 T^2)$ at high temperatures to quite negative (symmetry breaking) in vacuum.

The phase transition occurs near where it is zero, at a critical temperature which in the context of dimensional reduction becomes a critical $m_H^2(T)$. At the tree level, $\lambda/g^2 = (m_H/m_W)^2/8$ but the radiative corrections are important and the relation between λ/g^2 and the ratio of physical zero temperature masses m_H/m_W is not simple, especially in extensions to the standard model with new light bosons. Also note that, for instance, if the MSSM right scalar top quark is too light, the reduction to an MSM-like effective theory is not very reliable; we should use an effective theory which contains the light squark and the gluons. Dropping heavy modes is not a necessary step for using our technique, it is merely convenient to reduce the numerical demands, which would be a few times larger if we include the A_0 fields and order 10 times larger if we include the squarks and QCD. We will discuss how we think light squarks would change our results in the conclusion. We discuss the matter of integrating out the A_0 field in more detail in Sec. IV B. It is not always appropriate to do so, and in particular we cannot when we are studying the influence of hard thermal loops.

We do keep the U(1) subgroup, which is often left out in electroweak studies. Its role in setting the sphaleron rate is probably almost entirely due to its effect on the strength of the phase transition and not a direct modification of the sphaleron, see Ref. [36], but the numerical cost of including it is small enough that dropping it is pointless. We use $\tan^2 \Theta_W = 0.32$, based on a one loop match between vacuum modified minimal subtraction scheme ($\overline{\text{MS}}$) and 3D thermal values using results in Ref. [30].

B. The separatrix

The idea of the separatrix between vacuum states is essential to our technique. Before introducing it, let us review what we expect the space of gauge-Higgs configurations to look like in the broken phase. The space of three-dimensional gauge-Higgs configurations is periodic, with a discrete set of vacua. To be more precise, we should consider the space of gauge-Higgs configurations modulo (all) gauge transformations. In this case all vacua coincide¹ but the space is not simply connected. Since the index of the Dirac operator remembers when we go around a noncontractible loop in configuration space, the relevant space of physical configurations is the universal cover of the space of configurations modulo (all) gauge transformations. The universal cover has a discrete set of vacua labeled by the index of the Dirac operator.² If the line connecting two vacua in the universal

¹If the global topology of space is multiply connected then Yang-Mills theory has a connected manifold of inequivalent vacua corresponding to different values for traces of certain noncontractible Wilson loops. The toroidal spaces we will consider are multiply connected; however, the (fundamental representation) Higgs condensate lifts the degeneracy of the would be gauge vacua. These complications will not be important for what we do.

²The universal cover is roughly the same as the space of configurations modulo small gauge transformations. But we prefer to think in terms of the universal cover of the space of configurations modulo all gauge transformations, because this is an explicitly gauge invariant approach.

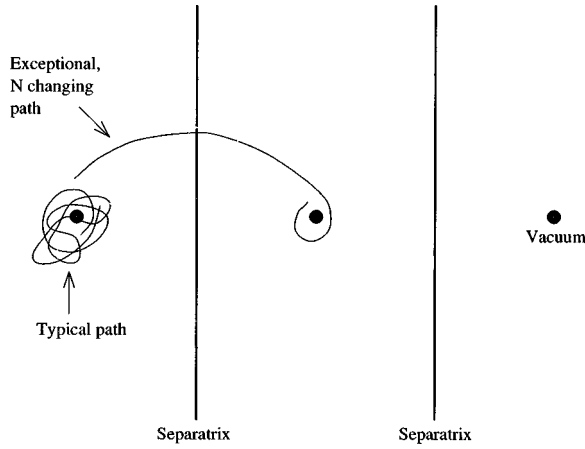


FIG. 1. Diagram showing the periodic vacuum structure, separatrix, typical path which stays near a vacuum, and exceptional path which crosses the separatrix and leads to permanent N_{CS} change.

cover projects to a winding one loop—or in more conventional language, if two vacua differ by 1 in Chern-Simons number—we will refer to them as neighboring vacua.

We expect that, in the broken phase, almost all of the weight of the canonical ensemble lies in states which are *in some sense* close to one of the vacua. The Hamiltonian evolution of a generic state in the ensemble will wander around in the neighborhood of one vacuum for an exponentially long time before it happens to make an excursion far enough away that it crosses to being nearest another vacuum, which it may find instead. We expect that this is how N_{CS} diffusion will occur.

To determine the rate of N_{CS} diffusion, we draw a (codimension 1) surface separating one vacuum from its neighbor, halfway between the minima in some sense, so that all of the well populated area near one minimum falls on one side and all the well populated area near the other minimum lies on the other side. This surface is called the separatrix between the vacua (see Fig. 1). To cross from being near one vacuum to being near another vacuum, a Hamiltonian trajectory must pass through the separatrix dividing them. We will assume that, after such a crossing, the trajectory almost never promptly continues to and crosses the next separatrix, but instead either settles around the new minimum for long enough that ergodicity “erases its memory,” or turns around and returns to the vacuum it started from. Then the flux of probability of the thermal ensemble through the separatrix is an *upper bound* on the diffusion rate for N_{CS} . It is an upper bound because of trajectories which cross the separatrix, turn around, and return to the original vacuum. These lead to flux of probability through the separatrix, but not to N_{CS} diffusion. To get the true diffusion rate, we need to find not only the flux through the separatrix, but the average number of crossings of the separatrix per long term change from being near one vacuum to being near another. We will call the reciprocal of this, the fraction of separatrix crossings which are associated with permanent N_{CS} change, the “dynamical prefactor.” If we know both the flux of probability across

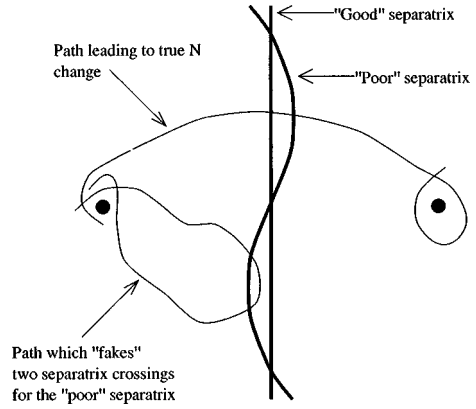


FIG. 2. Diagram of how a poor choice of separatrix can lead to overcounting the flux, and a small dynamical prefactor.

separatrices, and the dynamical prefactor, then we know the diffusion constant for N_{CS} ; the diffusion constant is

$$\gamma_d \equiv \lim_{t \rightarrow \infty} \frac{\langle [N_{CS}(t) - N_{CS}(0)]^2 \rangle}{t} = (\text{prefactor}) \times (\text{flux}). \tag{3}$$

What we really want is the diffusion constant per unit volume $\Gamma_d \equiv \gamma_d/V$. We should use a volume which is large enough to prevent finite size systematics but small enough that there are almost never two simultaneous N_{CS} changing events in different places.

A few points are in order here. First, the assumption that there are no prompt crossings of multiple separatrices is essential to the calculation. As we will see, it is also easy to test. Second, if we make a somewhat poor choice of the separatrix, so that there is some place where it bends nearer to one vacuum than the other, then much of the flux in that place will be of trajectories which double cross and return to their starting vacuum. The probability flux will be larger than with a better definition of the separatrix. However, the dynamical prefactor will be smaller, since these extra crossings do not lead to permanent N_{CS} change. The rate we determine is independent of exactly where we put the separatrix as long as the flux across it is exponentially small, and as long as we make a complete calculation, including the dynamical prefactor. We illustrate this point in Fig. 2. In practice we should look for a good choice of separatrix, since a poor choice of separatrix may make it harder to get good statistics for Γ , as most of the (numerical) effort will go towards studying trajectories which double cross rather than ones which really change N_{CS} permanently.

C. Calculation: perturbation theory

The existing perturbative calculations of the broken phase diffusion constant for N_{CS} are along the lines of the approach we just described. To allow a perturbative calculation, they make an additional assumption; that the separatrix is dominated by its saddle point. That is, they assume that gauge-

Higgs configurations on the separatrix are perturbatively small excitations about a background field which is the lowest energy point on the separatrix, which will be the Klinkhamer-Manton sphaleron [36]. The choice of a definition of the separatrix is then made perturbatively; a point is on the separatrix if the excitation in the unstable direction of the sphaleron is zero. The probability flux through the sphaleron can be computed in perturbation theory by comparing the free energy of all excitations of the sphaleron to the free energy of excitations about the naive vacuum, with the frequency of the unstable mode serving to convert a probability into a flux and the translational zero modes converting this into a flux per unit volume.

The probability flux through the separatrix has been computed in the above approximation at the one loop level [14,15]. However, extending the calculation beyond one loop raises severe technical problems. The sphaleron is not a spatially homogeneous background field, so the perturbative calculation must be done in real space with a numerically determined spectrum of fluctuations. The one loop calculation requires finding this spectrum, but the two loop calculation involves overlap integrals to compute the energy of their mutual interactions. There are also conceptual problems, because one of the fluctuation directions is unstable. At one loop it is excised from the sum over fluctuations; the other fluctuations set the probability to be near the separatrix and it turns that probability into a flux through the separatrix. But it is not clear how to separate it from the other modes at two loops. These problems obstruct a systematic improvement of the perturbative treatment.

There is also the problem of how to determine the dynamical prefactor perturbatively. Khlebnikov and Shaposhnikov argued that it should equal to 1 [37], but Arnold and McLerran made an estimate based on Landau damping which suggests that it is quite a bit less than 1 [3]. That argument has been more carefully developed by Arnold, Son, and Yaffe, who claim that, so long as the Higgs condensate gives a W mass which is parametrically $m_W \sim g^2 T$, the prefactor should be parametrically $O(\alpha_w)$ [33]. Their argument has recently been tested in the symmetric phase [26]. No one has used their picture to get a quantitative prediction of the dynamical prefactor within the context of the perturbative calculation of the sphaleron rate, although this should be possible in principle. These limitations of the perturbative approach, together with the generally spotty performance of perturbation theory for electroweak phenomena at temperatures near the electroweak phase transition, motivate a fully nonperturbative attack on the problem.

D. Topology and the lattice

The nonperturbative technique best suited to studying the sphaleron rate is the lattice. There is an apparent complication, though, which is that topology cannot in general be well defined on the lattice; the global structure of the space of three dimensional lattice gauge-Higgs configurations is different than in the continuum, and in particular it is simply connected, so there are not topologically distinct vacua. However, if we make the lattice spacing suitably small, then

the thermal ensemble is completely dominated by configurations in which *every* elementary plaquette is close to the identity. This means, roughly, that fields are perturbatively small at the lattice scale. This subspace of the space of lattice configurations does have the same topological structure as the continuum theory. Hence, if we excise a subspace of lattice configurations which carries an exponentially small weight, then we can talk about topology on the lattice.

The physical meaning of this is as follows. The space of continuum configurations permits sphaleronlike objects of arbitrarily small spatial extent. They also exist on the lattice, down to where their size is comparable to the lattice spacing; but at this point it becomes unclear how to define a smoothly interpolating continuum field, and the topological meaning is lost. Such a configuration in the analogous four-dimensional context is referred to as an “exceptional configuration.” However, the energy of spatially small, sphaleronlike field configurations rises linearly with inverse radius; so the Boltzmann suppression of lattice scale sphalerons is enormous and they essentially never occur. As long as the “genuine” sphalerons we study are comfortably larger than the lattice spacing, we have no problem. And when the genuine sphalerons we want to study are not comfortably larger than the lattice spacing, then obviously the lattice spacing is too coarse, and we should use a finer lattice.

The situation here is much better than it is in the 4D case considered in QCD. It is also true in four dimensions that if the gauge fields are smooth enough, then topology is well defined [38–40].³ However, in practice the fields may not generally be smooth enough. Instantons just larger than the lattice spacing typically do exist, because the instanton action is classically scale independent, and so only varies logarithmically with instanton size. Making the lattice finer does not eliminate lattice spacing sized instantons (exceptional configurations) very quickly; in fact, because their density goes as $\exp(-1/g^2)$ and g^2 varies logarithmically with lattice spacing, their density declines as an algebraic power of a . In the 3D context, though, the energy of a sphaleron of radius r goes as $1/(g^2 r)$ and the density of exceptional configurations varies with lattice spacing as $\exp[-(\text{coefficient})/(g^2 a T)]$, which falls off extremely rapidly as a is made small. This difference between the 3D and 4D cases is because the 3D theory is super-renormalizable; the coupling constant is $g^2 T$, which is dimensionful. Since $1/g^2$ appears in the exponent for the rate of any nonperturbative phenomenon, and since it must be accompanied by a T , on dimensional grounds a nonperturbative phenomenon which occurs at the lattice spacing scale must proceed at a rate which goes as $1/a$, leading to the exponentially fast rolloff in the density of exceptional configurations as $a \rightarrow 0$.

³In fact the 3D case is a subset of the 4D one, since the topology we are talking about is the second Chern class of a closed loop in 3D configuration space, which is equivalent to a periodic 4D lattice configuration with the spacing in the fourth dimension driven to zero.

E. A “nice” nonperturbative choice for the separatrix

Now that we have established that topological questions can have a meaning on the lattice, we ask how to choose a separatrix which is defined nonperturbatively and can be implemented on the lattice. A definition we would prefer for the separatrix is the “gradient flow” definition, suggested in Refs. [37,41,42]. It is defined in terms of gradient flow under the Hamiltonian. The Hamiltonian is a smooth function over the space of configurations, with degenerate global minima at the vacua. We believe that these are the only local minima of the Hamiltonian in 3D Yang-Mills Higgs theory, although we do not know a proof. Hence, following the direction of steepest descent of the energy (gradient flow) will lead, off a set of measure zero, to a vacuum configuration. A rigorous definition of “the vacuum closest to a configuration” is the vacuum arrived at by such gradient flow, which is also easily implemented on the lattice. A very sensible definition of the separatrix is then the boundary between the gradient flow basins of attraction of two neighboring vacua (vacua with N_{CS} differing by 1). Equivalently, it is the basin of attraction of the saddle point which sits between the two vacua, i.e., the sphaleron. If we mod out by all gauge transformations the two vacua are equivalent, but the separatrix can still be defined as the surface where two infinitesimally separated configurations on opposite sides will have macroscopically different gradient flow paths which, when spliced together, form a noncontractible loop.

Alternatively we could define the separatrix just in terms of the Yang-Mills field (connection) and the Yang-Mills term in the Hamiltonian. This choice has the added benefit that, as the configuration gradient flows under the Yang-Mills Hamiltonian in three dimensions, it becomes exceedingly smooth (meaning that all gauge invariant local measurables are slowly varying and the energy density is very small). Also, leaving out the Higgs fields evades the complication that the Higgs mass squared is renormalized by ultraviolet thermal excitations, which change during the gradient flow. The Yang-Mills gradient flow separatrix may not coincide with the Yang-Mills Higgs separatrix (which is not defined until we decide how to deal with the renormalization of the Higgs mass term in the Lagrangian). But if we perform a complete calculation, including the dynamical prefactor, then the exact choice of separatrix should not matter, as long as crossings are exponentially rare. Of course, a poor choice will make the calculation inefficient, since most crossings will not be associated with topology change. But we do not expect the Yang-Mills gradient flow separatrix to be a poor choice, and we will be able to check this belief when we study the dynamical prefactor.

F. Approach to computation

Now, we will outline how to compute the flux and the dynamical prefactor. To do so, we need more than just a definition of a separatrix. We need an order parameter N which takes a special value, say $N = \frac{1}{2}$, on the separatrix, and is smaller on one side and larger on the other (say, going from $N = 0$ at one vacuum to $N = 1$ at the other). Assume that we have such an N .

First, we should measure the probability over the canonical ensemble that N is within some small tolerance $\epsilon/2$ of $(\text{integer} + \frac{1}{2})$. This gives the probability to be very near a separatrix. Since the probability to be close to the separatrix is (expected to be) exponentially small, we will need to use a multicanonical reweighting technique to sample here accurately. The basic idea is presented early in Sec. IV.

Then we need to know the mean inverse time for crossing this narrow region, to turn the probability into a flux. This is $1/\epsilon$ times $\langle |dN/dt| \rangle$, the mean of the absolute value of the time derivative of N , where the averaging is over the ensemble restricted to the narrow band about the separatrix. We determine this by taking a canonical sample of states with $|N - 1/2| < \epsilon/2$, drawing momenta for each from the thermal ensemble,⁴ and performing a short segment of Hamiltonian evolution, comparing N with its value a short Hamiltonian time later. Thus, we can determine the flux.

To determine the dynamical prefactor, we take a canonically weighted sample of configurations at the separatrix and choose momenta for each, just as we do to find $\langle |dN/dt| \rangle$. Then we perform the Hamiltonian evolution both forward and backward in time, until the forward and backward histories both settle into the neighborhood of a vacuum. We count how many times the Hamiltonian trajectory crosses the separatrix before it settles semipermanently about a minimum. The prefactor is

$$\text{prefactor} = \sum_{\text{sample}} \frac{1}{\text{number of crossings}} \times \begin{cases} 1 & \text{final vacuum} \neq \text{initial vacuum,} \\ 0 & \text{final vacuum} = \text{initial vacuum.} \end{cases} \quad (4)$$

This is not the same as adding up the number of times the final vacuum differed from the initial one and dividing by the number of crossings of the separatrix. The reason is that the latter overcounts trajectories with many crossings, since the ensemble samples them more often than trajectories with fewer crossings. This is our recipe; given an order parameter N , we can determine both the flux of probability through the separatrix, and the dynamical prefactor, and hence Γ_d .

It remains to choose an order parameter. We want one such that the $N = \frac{1}{2}$ separatrix will be either the same as or quite close to the “good” choice of the gradient flow separatrix. Naively, we expect a good N to be the Chern-Simons number N_{CS} . Actually, this is not quite right, as we discuss at some length in the next section.

III. DEFINING LATTICE N_{CS}

We now have an idea for a nonperturbative definition of the separatrix, but we need an order parameter which varies

⁴It is essential here that the phase space is the tangent bundle of the configuration space, and that the thermal probability distribution is a product of a configuration space function and a (Gaussian) function over momentum space.




FIG. 3. Illustration of the pieces which make up a lattice gauge theory.

from 0 to 1 as it ranges between vacua and equals $\frac{1}{2}$ on a surface close to the Yang-Mills gradient separatrix. The literature generally considers the Chern-Simons number N_{CS} to be the best choice for this [37], so we will discuss how to define N_{CS} on the lattice; then we will define a different but closely related order parameter N , which gives a separatrix much closer to the gradient flow separatrix than N_{CS} does, and which we will actually use in the calculation of the broken phase diffusion rate.

A. The definition

The Chern-Simons number should be defined as some real valued function over the space of three-dimensional gauge connections. Let us review some properties which N_{CS} would have in the continuum, and see how closely we can preserve them on the lattice.

In the continuum, N_{CS} , defined on the space of gauge field configurations modulo *all* gauge transformations, has the following properties.

- (1) N_{CS} should be a continuous, multiply valued function, with the values separated by 1.
- (2) The multiple values correspond to the projections of different points on the universal cover, where N_{CS} is single valued. A noncontractible loop in the configuration space lifts to a line segment in the cover space, and the difference in N_{CS} between the end points is the winding number of the loop. (A more conventional way of saying this is that N_{CS} differs between two gauge copies by the winding number of the gauge transformation between them. But we prefer the above, gauge invariant statement.)
- (3) A vacuum configuration has N_{CS} modulo 1 equal zero.
- (4) Consider a path in configuration space parametrized by τ . The difference in N_{CS} between the beginning and end of the path should be

$$\Delta N_{CS} = \frac{g^2}{16\pi^2} \int d\tau \int d^3x \epsilon_{ijk} F_{ij}^a (D_\tau A_k)^a. \quad (5)$$

This is the same as saying that N_{CS} is the indefinite integral of

$$\frac{g^2}{32\pi^2} \int d^3x F_{\mu\nu}^a \tilde{F}_{\mu\nu}^a = \frac{g^2}{8\pi^2} \int d^3x E_i^a B_i^a. \quad (6)$$

How many of these properties can we preserve on the lattice? Not all of them; as we pointed out in Ref. [20], there is no local operator definition of $E_i^a B_i^a$ which is a total derivative. We also argue there that there are severe difficulties satisfying property (2); but this is because we demanded a singly valued, nongauge invariant definition of N_{CS} . In fact we can present a lattice definition of N_{CS} which satisfies everything but property (4) and the continuity requirement, provided we restrict to gauge fields which are smooth enough, in the sense of the last section, so that property (2) makes sense.

Before constructing a lattice definition of N_{CS} , we remind the reader how the lattice fields are defined (see Fig. 3). On a lattice, scalar fields are only defined at a discrete set of points, the lattice sites. The gauge field should be a connection, that is, it should be a rule which tells how to parallel transport fields along paths. We allow a path on the lattice to consist of a series of straight lines between nearest neighbor lattice sites. The connection is then defined by associating a group element $U \in SU(2)$ with each of these elementary straight lines (referred to as the links of the lattice). A small closed path, or the product of the U around the path, is called a plaquette; the 1×1 square is the elementary plaquette. When we do not specify the shape of a plaquette we mean an elementary plaquette. The product of the U around a plaquette (written U_\square , often just referred to as ‘‘a plaquette’’) will not in general be the identity; its failure, a curvature in the connection, is a field strength.

It is important that the field strength is not associated with a site of the lattice, but with a plaquette, which sits in between sites. Then, to make a lattice implementation of $\int \epsilon^{\mu\nu\alpha\beta} F_{\mu\nu} F_{\alpha\beta}$, we will have to do some averaging. The argument of the integral is a pseudoscalar and should perhaps be defined at the lattice sites; but the field strength is associated with a plaquette. To preserve cubic symmetry, $F_{\mu\nu}$ at a site will have to be the average over the four plaquettes in the μ, ν plane which touch the site.⁵ Because of this averaging process over things which do not live quite at a lattice site, the resulting lattice definition of $\epsilon^{\mu\nu\alpha\beta} F_{\mu\nu} F_{\alpha\beta}$ will not be a total derivative; and we cannot fix this problem by going to fancier definitions involving weighted averages of plaquettes of various shapes [20]. The problem is that the continuum proof that $F\tilde{F}$ is a total derivative relied on continuity of the fields, and this continuity is lost on the lattice. This makes it impossible to satisfy property (4). Note, however, that the lattice definition of $F_{\mu\nu} \tilde{F}_{\mu\nu}$ is gauge invariant.

⁵We are thinking of a path through configuration space as a four-dimensional lattice, where the fourth dimension is the path parameter, taken as a discrete rather than a continuous variable. Successive three-dimensional slices are configurations along the path. In a numerical setting, a path through the configuration space will always look like this, though making the spacing in the fourth direction arbitrarily small restores the continuity of the path.

However, when the gauge fields are weak at the lattice scale (meaning that the elementary plaquettes are close to the identity) and slowly varying (meaning that the departure from the identity is nearly the same between a plaquette and the parallel transport of a nearby plaquette in the same plane), then the lattice definition of $E_i^a B_i^a$ is approximately a total derivative. Further, we can pick a definition of N_{CS} so that property (4) is satisfied at least in one particular special direction. We choose it to be true in the cooling direction, that is, the direction of energy gradient flow. Here the energy of the lattice field is given by the standard Kogut-Susskind Hamiltonian [43], which is defined up to a multiplicative factor as

$$H_{KS}(U) \propto \sum_{\text{plaquettes } U_{\square}} \left(1 - \frac{1}{2} \text{Tr} U_{\square} \right). \quad (7)$$

The gradient of H_{KS} is to be understood in terms of the metric of the configuration space, which is the product of the Haar measure over each link matrix U .

We will call a path which follows the steepest descent of H_{KS} a cooling path, and we parametrize it with a cooling time τ , defined as $d\tau = d(\text{path length}) / [dH_{KS} / d(\text{path length})]$. The gauge fields evolve along the path according to [25]

$$\frac{\partial U}{\partial g t} = -D^{\alpha} U D^{\alpha} H_{KS}, \quad (8)$$

where D^{α} is the left acting covariant derivative, $D^{\alpha} U = i\tau^{\alpha} U$, and $D^{\alpha} H_{KS}$ is H_{KS} with U replaced with $D^{\alpha} U$. This is the gauge invariant lattice implementation of the continuum evolution

$$\frac{\partial A_i^a(x)}{\partial \tau} = - \frac{\partial H}{\partial A_i^a(x)}. \quad (9)$$

τ has dimensions of length squared.

We are defining N_{CS} so that condition (4) is true along the cooling path. Since cooling eventually leads to the vacuum off a set of measure zero, and since N_{CS} of the vacuum is by definition an integer, we get

$$N_{CS} \equiv \text{integer} - \int_0^{\infty} d\tau \frac{g^2}{8\pi^2} \int d^3x E_i^a B_i^a, \quad (10)$$

where we mean the lattice definition of $E_i^a B_i^a$ which we have described, and which is written down in Ref. [19]. This definition of N_{CS} is nice because the cooling path leads most quickly towards configurations with weak, slowly varying gauge fields, so the definition is minimally contaminated by the problems with the lattice definition of $E_i^a B_i^a$ which we have discussed.

Performing the integral in Eq. (10) takes an enormous numerical effort. Cooling the fields (following the cooling path) stably requires using a τ step size of less than $a^2/6$ (a the lattice spacing), or else the most ultraviolet excitations




FIG. 4. Original lattice and the coarsened version, which is just the solid lines and filled vertices. The link matrix between two vertices of the coarsened lattice is the product of the matrices on the two original links which make it up.

become unstable.⁶ But sufficient cooling may demand going to τ of order $1000a^2$. However, the cooling is by far the most efficient at removing ultraviolet excitations, and already by $\tau > a^2$ the fields are slowly varying. After this much cooling we lose almost no information if we drop some UV degrees of freedom by setting up a coarsened lattice. Define an even site as a site where all three coordinates are even numbers. In a scalar field theory, we would coarsen by dropping out all the lattice sites which are not even, leaving a lattice half as many points across in each direction. In a gauge theory, we also need to define the connections between the sites of the coarsened lattice; we define the connection between two neighboring even sites as the product of the two connections along the straight line between them. We illustrate the idea in Fig. 4. The remainder of the cooling then proceeds 2^5 times faster, 2^3 because the lattice is smaller and 2^2 because we can use a τ step size which is larger in physical units.

Of course we must check that this procedure produces the same answer as we get by not coarsening. But if we use an $O(a^2)$ improved definition of $E \cdot B$, for instance the one we developed in Ref. [21], then this is not a problem at all. The value of N_{CS} gets rescaled by an amount typically less than 1% and has an amount of noise added to it which is typically even smaller. For lattices more than about 28 sites across, we can even safely perform a second stage of coarsening after performing several coolings on the once coarsened configuration.

B. Application to the symmetric phase

As an application we discuss how to use this definition of N_{CS} to study N_{CS} diffusion in the symmetric phase, or pure Yang-Mills theory. In this case we want to track N_{CS} along the projection into configuration space of a Hamiltonian trajectory in phase space. That is, generating a Hamiltonian trajectory will give us a closely spaced series of points in configuration space, and we must associate a single valued function N_{CS} with this series. Our definition says that we should measure N_{CS} at each point by performing the integral

⁶We improve on this marginally by using alternately larger and smaller step sizes, analogous to what we did in Ref. [20] when quenching the E fields to enforce the Gauss law.

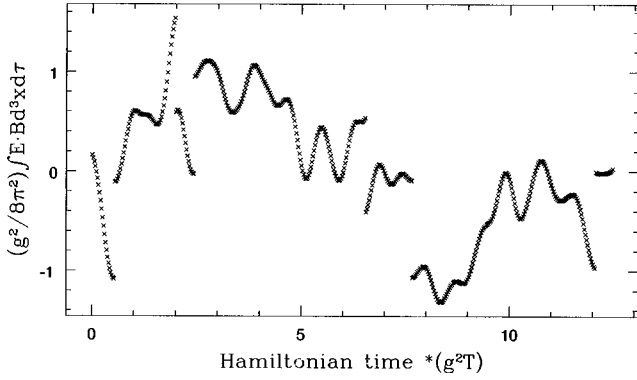


FIG. 5. $(g^2/8\pi^2)\int E_i^a B_i^a d^3x d\tau$ for a series of points on a Hamiltonian trajectory. It is clear where to adjust the integer part of Eq. (10) to keep N_{CS} (approximately) continuous.

in Eq. (10), that is, we must cool the configuration at each point along the Hamiltonian trajectory to the vacuum. We also have to choose the integer part of Eq. (10) somehow. At certain points, the value of the integral will abruptly jump by almost an integer. This happens whenever the Hamiltonian trajectory crosses a gradient flow separatrix. We should choose the integer part of N_{CS} to minimize the magnitude of the jump in N_{CS} at the separatrix. We illustrate with an example of real data from a simulation of Yang-Mills theory in Fig. 5. Note that the discontinuities in N_{CS} do not necessarily occur when N_{CS} is near $\pm\frac{1}{2}$; we will discuss this more in the next subsection.

In practice, even with lattice coarsening, the numerical costs of cooling every configuration are unbearable, but we can do better. We show how in Fig. 6. Every few steps, we cool a little, to a depth of $\tau \sim a^2$, and we thereby construct a cooled image of the Hamiltonian path, a technique explored by Ambjørn and Krasnitz [25]. We measure N_{CS} using Eq. (10) at occasional points along this path, interpolating N_{CS} in between by integrating $E_i^a B_i^a$ along the cooled image of the Hamiltonian path. The cooling has eliminated most of the UV excitations, so $E_i^a B_i^a$ along the cooled path is close to




FIG. 6. (Time, cooling time) plane, with curves used to track N_{CS} for a Hamiltonian trajectory. Every few Hamiltonian updates, we construct a cooled copy of the configuration, giving a parallel cooled path. Every few points on that path, we measure N_{CS} directly, and we fill in between by integrating $E \cdot B$ along the parallel cooled path. As long as the integrals along paths 1 and 2 always agree modulo an integer, within a small tolerance, the technique is topological.

behaving as a total derivative, especially if we use an $O(a^2)$ improved definition of $E_i^a B_i^a$. The interpolated value of N_{CS} is therefore almost what we would get by using the direct definition at each time. We only need it to be close enough to determine the integer part of Eq. (10) unambiguously, which we can if the value we get by integrating on path 1 in Fig. 6 and the value we get by integrating on path 2 in that figure differ by an integer plus a small remainder. We can think of this remainder as a calibration of the integration along the cooled path, so the approach is a ‘‘calibrated cooling method,’’ with the occasional cooling paths to the vacuum recalibrating the method of Ambjørn and Krasnitz [25] to make it topological.

We compared this approach of measuring N_{CS} to the ‘‘slave field’’ topological method [24], by evolving Yang-Mills theory on a 24^3 grid at $a = 1/(2g^2T)(\beta_L = 8)$ for a total time of $6000a$, tracking N_{CS} by each technique. For the technique we just described, we constructed a cooled image path with one point every $a/5$ time. The cooling depth to this path was $5a^2/8$, and we calibrated by cooling to the vacuum every $2a$ time. With this lattice spacing and this frequency of calibrating, the largest remainder we observed was 0.2 and the typical absolute value was less than 0.05. We present the results for N_{CS} in Fig. 7. We have offset N_{CS} measured by the slave field method by 5 to keep the curves from lying on top of each other. The agreement is outstanding, and the difference in the determined values of N_{CS} is white on long time scales.

To explain why the two methods have a white noise difference, we review briefly how the slave field method works. It tries to keep track of the ‘‘integer’’ part of Eq. (10) by assuming that the cooling path will end in a vacuum which has winding number zero in Coulomb gauge, and then adding up the number of large gauge transformations required to keep the system in Coulomb gauge during the Hamiltonian trajectory. However, this ignores the contribution to Eq. (10) from the integral, that is, the difference between N_{CS} of the configuration and of the vacuum it cools to. Also, the algorithm used to find Coulomb gauge sometimes gets trapped temporarily in a Gribov copy with a different winding number. But neither difference between the methods will grow without limit in time, so the difference between the two measurement methods for N_{CS} is white on long time scales. Hence, the derived diffusion constant will be the same within errors (caused by the white noise difference). Indeed, when we used the technique of Ref. [24] to extract Γ_d from each trajectory, the two methods of tracking N_{CS} gave the same answer within error ($\gamma_d = 0.0515 \pm 0.0077a^{-1}$ for the new method, versus $\gamma_d = 0.0516 \pm 0.0082a^{-1}$ for the slave method). As an aside, we mention that Hetrick and de Forcrand have used the method of cooling to resolve the Gribov problem, by defining Coulomb gauge (or, in their four-dimensional context, Landau gauge) as the gauge in which the vacuum configuration arrived at via cooling has the minimum value of $\int A^2$ [44].

C. Modifications for measuring Γ_d in the broken phase

The definition of N_{CS} which we just presented more accurately reproduces the continuum meaning of N_{CS} than any




FIG. 7. Left: N_{CS} measured by the slave field (upper curve) and “calibrated cooling method” (lower curve). Right: the difference between the curves (note scale). The difference is small and spectrally white, so the methods are in good agreement.

other we know. But we do not actually want all of the attributes of the continuum meaning of N_{CS} if we want to measure Γ_d in the broken phase by the separatrix method. The reason is that N_{CS} is only directly a measure of topology for vacuum configurations. There are contributions to N_{CS} in excited states which are uncorrelated to topology; for instance, N_{CS} does not vanish in Abelian gauge theory, even though that theory has no π_3 topological sectors. In the continuum Abelian theory, N_{CS} is given by

$$N_{CS} = \frac{g^2}{32\pi^2} \int d^3x \epsilon_{ijk} F_{ij} A_k, \quad (11)$$

and the mean square value of N_{CS} is

$$\langle N_{CS}^2 \rangle = \frac{g^4}{1024\pi^4} \int d^3x d^3y \epsilon_{ijk} \epsilon_{lmn} \times \langle F_{ij}(x) A_k(x) F_{lm}(y) A_n(y) \rangle. \quad (12)$$

Using Wick’s theorem and the momentum representation of the propagator, in a general covariant gauge, this becomes

$$\begin{aligned} \langle N_{CS}^2 \rangle &= \frac{g^4}{1024\pi^4} \int d^3x d^3y \int \frac{d^3p d^3q}{(2\pi)^6} e^{i(p+q)\cdot(x-y)} \\ &\times \epsilon_{ijk} \epsilon_{lmn} T^2 \left[\frac{4p_i p_l}{p^2 q^2} \left(\delta_{jm} + (\alpha - 1) \frac{p_j p_m}{p^2} \right) \right. \\ &\times \left. \left(\delta_{kn} + (\alpha - 1) \frac{q_k q_n}{q^2} \right) + (m \leftrightarrow n) \right] \\ &= \frac{g^4 T^2 V}{64\pi^4} \int \frac{d^3p}{(2\pi)^3} \frac{p^2}{(p^2)^2}, \end{aligned} \quad (13)$$

so N_{CS} will be Gaussian distributed with a linearly divergent variance. On the lattice, the UV divergence will be cut off by the lattice scale; the coefficient was found by Ambjorn and Krasnitz [19] and is

$$\langle N_{CS}^2 \rangle = (1.44 \times 10^{-5}) g^4 V T^2 / a. \quad (14)$$

The divergence occurs because, while the energy cost of storing N_{CS} in a UV mode grows linearly with p , the number of available states in which to store N_{CS} grows faster; entropy wins over energy. The same thing happens in SU(2) theory, because N_{CS} also contains the $\epsilon_{ijk} F_{ij} A_k$ term. The coefficient of the divergence in SU(2) is larger by 3, the dimension of the group. Since Yang-Mills theory in 3D is super-renormalizable, the UV decouples from the IR, where the genuine topology changing physics occurs, so this UV divergent, Gaussian contribution will appear as an additive correction to the IR contribution to N_{CS} . That is, to reasonable accuracy we can think of N_{CS} , defined in Eq. (10), as $N_{CS} = N_{CS}^{IR} + N_{CS}^{UV}$, where topological information is in N_{CS}^{IR} , and N_{CS}^{UV} is independent of N_{CS}^{IR} and Gaussian distributed.

Given a single IR field configuration with some particular value of N_{CS}^{IR} , different realizations of the UV excitations on top of the IR fields will then give a distribution of values of N_{CS} ; so the probability that N_{CS} will have a particular value x , $P_{NCS}(x)$, will be

$$P_{NCS}(x) = \int_{IR \text{ configs}} \int_{UV \text{ excit}} \delta(x - N_{CS}^{IR} - N_{CS}^{UV}), \quad (15)$$

or, defining $P_{NCS}^{IR}(x)$ and $P_{NCS}^{UV}(x)$ to be the probability distributions for the IR and UV components of N_{CS} ,

$$\begin{aligned}
P_{\text{NCS}}(x) &= \int dy \int dz P_{\text{NCS}}^{\text{IR}}(y) P_{\text{NCS}}^{\text{UV}}(z) \delta(x-y-z) \\
&= \int dy P_{\text{NCS}}^{\text{IR}}(y) P_{\text{NCS}}^{\text{UV}}(x-y). \quad (16)
\end{aligned}$$

The probability distribution for N_{CS} will be a convolution of the interesting IR distribution and a Gaussian noise distribution. Convolution with any periodic function with a Gaussian always degrades contrasts in that periodic function, enhancing the probability to be near the $N_{\text{CS}} = \frac{1}{2}$ separatrix relative to the case of no UV noise. The distortion will be small only if

$$\langle (N_{\text{CS}}^{\text{UV}})^2 \rangle \ll \left. \frac{d^2 \ln P_{\text{NCS}}^{\text{IR}}(x)}{dx^2} \right|_{x=1/2}. \quad (17)$$

Later, we will use a 40^3 lattice with $a=2/(5g^2T)$. Multiplying Eq. (14) by the group factor of 3 and plugging in numbers, $\langle (N_{\text{CS}}^{\text{UV}})^2 \rangle = 0.44$ for such a lattice, which is too big. So defining N_{CS} by Eq. (10) will not do.

Another way to state the above is that the separatrix one gets from the condition $N_{\text{CS}} = \frac{1}{2}$ is sensitive to ultraviolet excitations, which makes it “all wiggly”; it will have lots of “fingers” which stick out towards one or the other topological vacuum, and the problems we discussed in Sec. II, of there being many crossings of the separatrix which do not have to do with permanent N_{CS} change, will be severe.

The problem is that $N_{\text{CS}} = \frac{1}{2}$ is not particularly similar to the “good” gradient flow definition of the separatrix. We want an order parameter which is close to $N = \frac{1}{2}$ on the gradient flow separatrix. The $N = \frac{1}{2}$ separatrix does not need to correspond exactly with the gradient flow definition; is sufficient if the distribution of values of N on the gradient flow separatrix is narrow, preferably narrower than $d^2 P_N(x)/dx^2$. A slight change to Eq. (10) will do the trick; define

$$N \equiv \text{integer} + \frac{g^2}{8\pi^2} \int_{\tau_0}^{\infty} d\tau \int d^3x E_i^a B_i^a, \quad (18)$$

where we mean the lattice implementation of the integrals and E, B as before. In other words, we “precool” the configuration for cooling time τ_0 and then measure N_{CS} . The precooling is intended to remove UV excitations without affecting the underlying IR fields much.

Now consider N in the Abelian theory again. The theory is linear, so it is easy to analyze how cooling affects it. A particular transverse mode $A(k)$ evolves according to

$$\frac{dA(k)}{d\tau} = -\frac{\partial H}{\partial A(k)} = -k^2 A(k) \Rightarrow A(k, \tau) = e^{-k^2 \tau} A(k, 0). \quad (19)$$

The propagator in Landau gauge becomes

$$\langle A_i(k, \tau) A_j(l, \tau) \rangle = \left(\delta_{ij} - \frac{k_i k_j}{k^2} \right) \frac{e^{-2k^2 \tau}}{k^2} \delta(k+l), \quad (20)$$

and the variance of N is




FIG. 8. N_{CS} (wildly oscillating curve) and N (curve which stays near 0) during a broken phase Hamiltonian trajectory, in a 40^3 box with $a=2/(5g^2T)$ [and using $\tau_0=3.2/(g^2T)^2$]. N_{CS} has a lot of UV noise, which is absent in N .

$$\langle N^2 \rangle = \frac{g^4 T^2 V}{64\pi^4} \int \frac{d^3p}{(2\pi)^3} \frac{p^2 e^{-4p^2 \tau}}{(p^2)^2}. \quad (21)$$

So precooling removes the UV noise from the definition of N_{CS} .

For comparison, Fig. 8 shows N and N_{CS} for the same broken phase Hamiltonian trajectory; while N_{CS} varies wildly on a short time scale due to UV fluctuations, N is steady, and shows that the infrared fields never stray far from the vacuum. N will not meet all the conditions we set out for N_{CS} ; for instance it will violate condition (4) severely. However, it will be closer to a continuous function, since initial cooling removes UV excitation from the configuration, leaving weaker and more slowly varying fields for which the definition of $E_i^a B_i^a$ is less problematical. There is also less problem using coarsening with this definition. For instance, on a 28^3 grid, single coarsening after $\tau = (5/4)a^2$ and double coarsening after $\tau = 2.8(2a)^2$, and using an $O(a^2)$ improved definition of $E_i^a B_i^a$, the discontinuity in N across the gradient flow separatrix is 0.9870 ± 0.0029 (drawing configurations from broken phase Yang-Mills Higgs theory at $\beta_L=7$, see next section). To make N a continuous function modulo 1, the jump should have been 1. By rescaling N slightly, the discontinuity is removed almost altogether. The value of the discontinuity for larger lattices and deeper cooling is even closer to 1 with less noise.

Let us check that the new definition of N will have $N \equiv \frac{1}{2}$ on the gradient flow separatrix. A point right on the separatrix will gradient flow to the saddlepoint configuration and stick there.⁷ Perturbing the starting configuration slightly off the separatrix, it will cool to the saddle, miss slightly, and then slide off to a vacuum. The closer to the separatrix we start, the longer we will stick in the saddle before we slide out. There is some early τ (transient) contribution to

⁷This saddle point is not the same as the Klinkhamer-Manton sphaleron, because we are considering the Yang-Mills Hamiltonian only, in a finite volume. But we know such a saddle will exist by the same argument Manton originally made for the existence of the sphaleron [45].

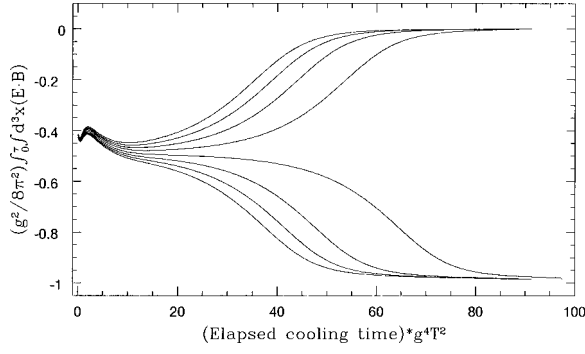


FIG. 9. Plot of $\int E \cdot B$ from 0 to τ , as a function of τ , for a series of points along a Hamiltonian trajectory as it passes through the gradient flow separatrix.

$\int \int E_i^a B_i^a d^3x d\tau$ while it is approaching the saddle, and then there is a contribution, almost exactly equal to $\frac{1}{2}$, as it rolls from the saddle to the vacuum. By choosing τ_0 large enough, we miss the transient and pick up only the $\frac{1}{2}$, and so the $N = \frac{1}{2}$ separatrix will correspond almost precisely with the gradient flow separatrix. We illustrate this with data from a Yang-Mills theory simulation in Fig. 9. The figure plots

$$\frac{g^2}{8\pi^2} \int_0^\tau d\tau' \int d^3x E_i^a B_i^a \quad (22)$$

(shifted so the vacuum will have integer N_{CS}) against τ for a series of points on a Hamiltonian trajectory as it goes through the separatrix. Each curve records the cooling process of a successive point on the Hamiltonian trajectory. We see that as the Hamiltonian trajectory approaches the (gradient) separatrix, the cooling path stays near the saddlepoint for longer and longer. When the Hamiltonian trajectory crosses the separatrix, the cooling paths roll out of the saddle towards the other side. The figure also shows the early τ transient. We want to choose a τ_0 large enough to eliminate this transient, so $N \cong \frac{1}{2}$ will hold for a configuration which starts almost exactly on the gradient flow separatrix.

We should make sure that the result for Γ_d will be independent of τ_0 , once τ_0 is large enough to eliminate the UV problems. We remind the reader that Γ_d is computed by choosing an $\epsilon \ll 1$, and computing the probability to be within $\epsilon/2$ of $\frac{1}{2}$,

$$P_\epsilon = \int_{(1-\epsilon)/2}^{(1+\epsilon)/2} P_N(x) dx. \quad (23)$$

Also, one computes the time rate of change of N , $\langle |dN/dt| \rangle$, evaluated for configurations with $N \cong \frac{1}{2}$. The rate Γ_d is then

$$\Gamma_d = \frac{1}{V} \frac{P_\epsilon}{\epsilon} \langle |dN/dt| \rangle, \quad (24)$$

times the dynamical prefactor.

Increasing the cooling time will decrease P_ϵ . The reason is that configurations which are near the saddle point at cooling time τ_0 are spreading out from each other as cooling time progresses, as Fig. 9 illustrates. The rate of the spreading is

given by the unstable frequency squared $(\omega_-)^2$ of the saddle.⁸ Increasing τ_0 makes them spread further before we measure their N_{CS} ; so the sample of states will be more diluted, by a factor of $\exp[-\Delta\tau_0/(\omega_-)^2]$. This reduces P_ϵ by the same factor. However, we measure dN/dt by choosing two neighboring configurations on a Hamiltonian trajectory and finding the difference in their values of N . The spread between these at cooling time τ_0 will also increase as we increase τ_0 , by the same amount; so $\langle |dN/dt| \rangle$ will go up by $\exp[\Delta\tau_0/(\omega_-)^2]$. Hence, Γ_d will be τ_0 independent. This means, however, that neither $\langle |dN/dt| \rangle$ nor $P_N(x)$ have simple physical interpretations.

There is a modification of the above reasoning if τ_0 is too short to eliminate the early transient; namely, a contribution to dN/dt due to the time evolution of the transient. In the complete calculation this will be compensated for because the dynamical prefactor will differ from the gradient flow value by a τ_0 dependent amount, which becomes non-negligible at the same time the transient contribution to dN/dt does. In a complete calculation, Γ_d will be independent of τ_0 , as we argued in the last section.

We end this section by discussing briefly why we choose to define N based only on the Yang-Mills fields and using the Yang-Mills Hamiltonian for the cooling, rather than including the Higgs field. Doing so is reasonable because N_{CS} should be defined as a function of the gauge fields alone. Also, cooling all the fields under the full Hamiltonian is problematic, because the UV fluctuations of the gauge and Higgs fields renormalize the Higgs mass squared [46]. The bare potential needs a large, negative mass squared counterterm. However, the UV fluctuations are the first casualty of the cooling process, and so they stop generating a thermal Higgs mass squared, early in the cooling. To keep the minimum of the Higgs potential from changing radically, one would need to vary the bare Higgs potential in a complicated way as the cooling progressed. Other important fluctuation induced effects are also lost; for instance, the cubic term which makes the phase transition first order disappears as we cool the excitations. Depending on how we handle the Higgs potential during cooling, we will either cause symmetry to break during cooling when we start in the symmetric phase, or cause it to be restored when we start in the broken phase; we cannot avoid both, because there is a range of temperatures where each phase is metastable. Cooling only the gauge fields avoids this complication.

IV. MONTE CARLO CALCULATIONS

Here we discuss details of performing the calculation of $P_N(x)$ and $\langle |dN/dt| \rangle$ using the definition of N from Eq. (18). We also discuss the computation of the dynamical prefactor, with and without ‘‘adding’’ hard thermal loops. Readers who do not do this kind of calculation will want to skip all but the first subsection and go to the results.

⁸Which does not equal $(\omega_-)^2$ of the Klinkhamer-Manton sphaleron.

A. Multicanonical Monte Carlo calculation: the idea

We need to compute the probability distribution P_N of Chern-Simons number [really, the Chern-Simons number of a precooled configuration, defined in Eq. (18)]. The probability density that $N=x$ is

$$P_N(x) = \lim_{\delta x \rightarrow 0} \frac{1}{Z \delta x} \int \mathcal{D}U \mathcal{D}u \mathcal{D}\Phi e^{-\beta H(U,u,\Phi)} \times \Theta[N(U) - x] \Theta[x + \delta x - N(U)], \quad (25)$$

where U, u, Φ are the SU(2) connection, the U(1) connection, and the Higgs field, and $N(U)$ is defined in Eq. (18). Here $\int \mathcal{D}U \mathcal{D}u \mathcal{D}\Phi$ means the integral over the value of each field at each lattice site, and Z is the value of the integral without the step functions. On an N^3 lattice, this is a $16N^3$ dimensional integral, which for $N=40$ is 1 024 000 dimensions. For this reason we turn to Monte Carlo integration. In a canonical Monte Carlo integration we generate a sample of configurations drawn with weight

$$e^{-\beta H(U,u,\Phi)} \mathcal{D}U \mathcal{D}u \mathcal{D}\Phi, \quad (26)$$

and replace the integral with a sum over that sample. This will not do in the present context, because we want to know $P_N(x)$ even where it is exponentially small. To get a good sampling there would require generating an exponentially large sample.

We evade this problem by doing a multicanonical Monte Carlo calculation [47]. We rewrite Eq. (25) as

$$P_N(x) = \lim_{\delta x \rightarrow 0} \frac{1}{Z \delta x} \int (\mathcal{D}U \mathcal{D}u \mathcal{D}\Phi e^{-\beta H(U,u,\Phi)} e^{f[N(U)]}) \times \{e^{-f[N(U)]} \Theta[N(U) - x] \Theta[x + \delta x - N(U)]\}, \quad (27)$$

with $f(x)$ some function we are free to choose. Now we generate a sample of configurations drawn with weight

$$e^{-\beta H(U,u,\Phi)} e^{f[N(U)]} \mathcal{D}U \mathcal{D}u \mathcal{D}\Phi, \quad (28)$$

and replace the integration in Eq. (27) with a sum over this sample, with the term in the second set of parenthesis as the argument of the sum. By choosing $f(x) \equiv -\ln P_N(x)$, we push the exponential suppression from the sampling into the integrand. The quality of the integration is now limited by how quickly we can generate a quality sample with this weight, and how well we can choose $f(x)$, which we must do by some form of bootstrapping.

A usual way to generate a canonical ensemble is by some Markov process. Given a configuration C_1 and a realization ξ of some random noise distribution, the process returns a new configuration $C_2 = M(C_1, \xi)$. Define the probability to return a particular C_2 as

$$\mathcal{P}(C_1, C_2) \equiv \int d\xi \delta(M(C_1, \xi), C_2). \quad (29)$$

If M satisfies detailed balance,

$$\frac{\mathcal{P}(C_1, C_2)}{\mathcal{P}(C_2, C_1)} = \exp[\beta(H(C_1) - H(C_2))], \quad (30)$$

then iterating the Markov process generates the canonical distribution. We get the best statistics if applying M is numerically cheap and if the return value differs as much as possible from the starting value.

If we have a Markov process which generates the canonical ensemble, we make it generate the multicanonical ensemble with weight function $f(C)$ by making the following modification. Given a configuration C_1 , generate $C_2 = M(C_1, \xi)$. Accept C_2 as the next configuration in the sequence with probability

$$\min[1, \exp(f(C_2) - f(C_1))] \quad (31)$$

and otherwise reject it and make C_1 the next configuration. This changes the detailed balance relation for the sequence to incorporate f into the weight. We may no longer want $M(C_1, \xi)$ to differ from C_1 by as much as possible, though, because that may make the reject rate very large. Instead we want $|f(M(C_1, \xi)) - f(C_1)| \sim 1$.

B. Note on algorithm

We need to do two kinds of things. The first is to evaluate the path integral for dimensionally reduced 3D Yang-Mills Higgs theory. The second is to study the dynamics of the (3+1)-dimensional classical theory. We comment briefly on the connection between the two; in particular we should compare the partition function of the classical theory to the path integral for the dimensionally reduced theory, a comparison first made by Ambjørn and Krasnitz [19].

The partition function of classical (3+1)-dimensional Yang-Mills Higgs theory (not worrying about the difference between lattice and continuum, which will not be important here) looks similar to

$$Z = \int \mathcal{D}A_i \mathcal{D}\Phi \mathcal{D}E_i \mathcal{D}\Pi \delta[(D_i E_i)^a + (g/2)(i\pi^\dagger \tau^a \Phi + \text{H.c.})] \times \exp[-H(A, \Phi, E, \Pi)/T], \quad (32)$$

$$H = \int d^3x \left(\frac{B^2}{2} + \frac{E^2}{2} + (D_i \Phi)^2 + \pi^2 \right), \quad (33)$$

where E , the electric field, is the conjugate momentum of A , and Π is the conjugate momentum of Φ . The delta function enforces the Gauss law. We can implement the Gauss law by means of an adjoint valued Lagrange multiplier A_0^a , giving [19]

$$Z = \int \mathcal{D}A_i \mathcal{D}\Phi \mathcal{D}E_i \mathcal{D}\Pi \mathcal{D}A_0 \times \exp\{iA_0^a [(D_i E_i)^a + (g/2)\Pi^\dagger \tau^a \Phi + \text{H.c.}]/T\} \times \exp[-H(A, \Phi, E, \Pi)]. \quad (34)$$

The integrals over E and Π are rendered Gaussian; doing them gives

$$Z = \int \mathcal{D}A_i \mathcal{D}\Phi \mathcal{D}A_0 \exp[-H(A, \Phi, A_0)/T], \quad (35)$$

$$H = \int d^3x \left(\frac{B^2}{2} + \frac{(D_i A_0)^2}{2} + (D_i \Phi)^2 + \frac{g^2}{4} A_0^2 \Phi^2 \right). \quad (36)$$

So the thermodynamics of the classical theory with the Gauss law is governed by the 3D dimensionally reduced path integral, but including the A_0 field at zero (bare) Debye mass. We could get the path integral without the A_0 field if we did not enforce the Gauss law. There are two choices; either we treat the dynamics dropping the Gauss law or we include the A_0 field in the 3D Monte Carlo parts of the calculation.

The conservative approach is to include the Gauss law in the dynamics, which means including an A_0 field in the thermodynamics. The easiest canonical Monte Carlo Markov process in this case is a short Hamiltonian trajectory starting with randomized but constraint respecting E and Π fields, with the multicanonical accept reject steps inserted between evolutions. This is a ‘‘constrained molecular dynamics’’ algorithm, and the problem of drawing E and Π from the thermal distribution respecting the constraints is addressed in Ref. [20]. Alternately we could use heat bath and overrelaxation updates on the system described in Eq. (35). Neither approach is terribly efficient.

The other option is to assume that Gauss’ Law is not very important to the dynamics, and not enforce it when we draw momenta E and Π from the thermal ensemble. This assumption is certainly justified with regards to measuring $\langle |dN/dt| \rangle$. What the combination of E and Π forced zero by the Gauss law would do if we did not set them zero is to gauge rotate the fields, but not the momenta. On long time scales that might be important, but at leading order in the length of a short Hamiltonian trajectory it only changes the gauge of the final configuration. Since N is a gauge invariant object, this does not matter. A less rigorous but more physically intuitive way to see the unimportance of the Gauss Law to $\langle |dN/dt| \rangle$ is to note that it is roughly the instantaneous value of $\int d^3x E_i^a B_i^a$. Now the magnetic field is transverse by the Bianchi identity, so only the transverse components of the electric field contribute to dN/dt . But the Gauss law only depends on the longitudinal components, so dN/dt should be the same whether or not we enforce it.

It is less clear whether the Gauss law will have a role in setting the dynamical prefactor, since it depends on longer time dynamics; but we are getting the dynamical prefactor wrong anyway if we do not enlarge the system somehow to account for hard thermal loop effects properly. We should deal with these two questions together.

The chief advantage of not enforcing the Gauss law is that there are very efficient update algorithms for the path integral without either the E , Π , or A_0 fields, for instance the one developed by Rummukainen *et al.* [5]. We adopt their lattice action, which is the standard Wilson 3D Yang-Mills Higgs action, but we add a noncompact U(1) field. We use their update, which must be extended to include the noncompact U(1) field. In the noncompact formulation, the U(1) gauge

field on a link is represented by a single real number $B_i(x)$, and the terms in the action which depend on the U(1) field are

$$H \supset \frac{a}{2g^2 z} \sum_x \left[\sum_{i>j} [B_i(x) + B_j(x+i) - B_j(x) - B_i(x+j)]^2 \right] - a \sum_x \sum_i \phi^\dagger(x) U_i(x) \exp[iaB_i(x)] \phi(x+i), \quad (37)$$

where a is the lattice spacing, $z = \tan^2 \Theta_w$, and $U_i(x)$ is the SU(2) connection on the i, x link. The compact formalism is the same except that $(\frac{1}{2})(\Sigma B)^2$ is replaced by $a^{-2}[1 - \cos(a\Sigma B)]$. In the compact case the energy, as a function of one B , is a trigonometric function. It is easy to perform an exact heat bath or over-relaxation step [10]. For the noncompact case the energy is the sum of a quadratic and a trigonometric function and it is not easy to perform an exact over-relaxation or heat bath update. We perform the update based just on the (much larger) quadratic term and include the trigonometric term by an accept reject step. The accept rate is quite high, so the cost to algorithm efficiency is low. We also occasionally gauge transform to bring the B fields towards Coulomb gauge so that the typical B is close to zero and the series expansion of $\exp[iaB_i(x)]$ converges quickly.

We always apply $O(a)$ improvement to the lattice action as described in Refs. [48, 49]. Whenever we refer to physical units in this paper they are always related to the lattice ones through $O(a)$ improved relations. The improvement is essential to achieving the numerical accuracy we want at reasonable lattice spacings, and it makes small spacing extrapolations of most quantities unnecessary. For instance, we have computed the jump in ϕ^2 at the equilibrium temperature for $a = 4/7g^2T$ ($\beta = 7$) and $a = 2/5g^2T$ ($\beta = 10$) to see how sensitive results are to varying a . The results are $(\Delta \phi^2/g^2T^2) = 2.51 \pm 0.02$ and 2.56 ± 0.03 respectively; the $[O(a^2)]$ errors are smaller than the statistical errors we will be able to achieve for Γ_d , so lattice spacing errors are under control.

We have also used the technique of Ref. [26] to study the influence of hard thermal loops on the dynamical prefactor. The idea is to add a large number of weakly interacting, ballistic charged particles to the lattice system, which reproduce the effects of the hard degrees of freedom left out when we set up our lattice. We refer the reader to Ref. [26] for details. This approach demands that we apply the Gauss law, and we must use the (Hamiltonian) thermalization algorithm as our Markov process. This is very inefficient, so we can only use the technique to study the dynamical prefactor, not the flux through the separatrix.

C. Finding T_c

At what temperature should we study the sphaleron rate in the broken phase? We choose the equilibrium temperature for the phase transition. This is appropriate if the latent heat liberated during the cosmological electroweak phase transition is sufficient to reheat the universe to T_c , which would be the case if the latent heat were large or the supercooling were small. It is in fact not clear whether this will be the

case, in the MSM or in more viable extensions. We will choose T_c anyway, because we then have a well specified question with only one free parameter λ/g^2 ; in addition, complete reheating may generically occur in supersymmetric extensions with a light top squark. (It also may not; see Ref. [50].)

In the effective 3D theory we actually find the critical Higgs mass squared, which is related through the dimensional reduction procedure to the critical temperature. Since this involves comparing the thermodynamical favorability of the broken and symmetric phases, it necessarily involves some multicanonical technique. We will use a particularly simple approach, somewhat similar to the one used in Ref. [51]. It is based on the fact that, in a very long rectangular box, values of $\int \Phi^\dagger \Phi / V \equiv \langle \phi^2 \rangle$ intermediate between the symmetric and broken phase values are obtained by having a mixed phase configuration, where part of the volume is in the symmetric phase and the rest is in the broken phase.

The free energy, as a function of $\langle \phi^2 \rangle$, is $-T \ln(P)$, with P the probability to have that value of $\langle \phi^2 \rangle$. In the range of intermediate $\langle \phi^2 \rangle$, the free energy will vary linearly with $\langle \phi^2 \rangle$, since a change of $\langle \phi^2 \rangle$ represents a change of how much bulk free energy comes from one phase and how much comes from the other phase. The slope of the linear relation tells the free energy difference between the two bulk phases. This linear regime breaks down where $\langle \phi^2 \rangle$ comes close to the value in one or the other phase, since the phase interfaces then get close enough together to interact.

Our approach is to add a $\langle \phi^2 \rangle$ dependent contribution to the Higgs boson mass, which we achieve by adding to the Hamiltonian a new nonlocal term

$$\eta N^3 \left(\frac{1}{2N^3} \sum \Phi^\dagger \Phi \right)^2. \quad (38)$$

Choosing η to be positive means that, if most of space is in the broken phase, the Higgs boson mass is heavier and the symmetric phase is favorable, whereas if most of space is in the symmetric phase, then the Higgs boson mass is smaller and the broken phase is favorable. The free energy is then a quadratic function, and the effective Higgs boson mass squared at its minimum, including the contribution from the η term, gives the equilibrium Higgs boson mass. The added term is simple enough that we can modify the canonical update to incorporate it, and Monte Carlo evolution will then naturally settle in a mixed phase configuration whose value of $\langle \phi^2 \rangle$ tells us the equilibrium Higgs boson mass,

$$m_H^2(\text{equilibrium}) = m_H^2(\langle \phi^2 \rangle = 0) + \eta \langle \phi^2 \rangle. \quad (39)$$

This approach can be viewed as a type of multicanonical Monte Carlo calculation with the η term as the multicanonical reweighting.

In practice we start with a long but very narrow box and a high value of η , to get a preliminary value. The narrowness is necessary to make it easy to nucleate a bubble of one phase in the other. To get the large volume limit, though, we need to go to a wider box; we necessarily need results in a regime where one phase cannot easily nucleate in the other.

We get an initial condition for a box an integer number of times wider in each short direction by extending the final configuration in the skinny box periodically. We also use a smaller value of η to improve the resolution of the determined m_H^2 . Our final values for m_H^2 are typically taken with a box of dimension, in physical units, of $(16/g^2 T)^2 \times (96/g^2 T)$, easily large enough to achieve the large volume regime.

D. Practicing with the symmetric phase in small volume

Before presenting the determination of the broken phase sphaleron rate, we will do a ‘‘practice run’’ on a system where we can get good statistics more quickly, which is the symmetric phase in a small enough volume to suppress topology change. This problem is almost certainly of no phenomenological significance, but it will let us test the τ_0 dependence of our technique and to study whether hard thermal loops do indeed modify the dynamical prefactor.

For the time being we drop the U(1) subgroup. We choose a very weak scalar self-coupling of $x=0.025$, and a large Higgs boson mass squared ($m_H^2 \approx 1.5g^4 T^2$), so we will be firmly in the symmetric phase. We use a 12^3 lattice with a lattice spacing of $a=1/(4g^2 T)$ ($\beta_g=16$), so the physical volume is $(3/g^2 T)^3$. This volume is small enough that topology change is suppressed and the broken phase techniques we are developing are applicable, but not so small that it will be hard to gather good statistics. It is also small enough that numerical costs are not overburdening, so we will enforce the Gauss law (and use the less efficient update algorithm).

We measure N by cooling for $\tau=(5/4)a^2$, coarsening once, and using an $O(a^2)$ improved definition of $E_i^a B_i^a$ during the subsequent cooling. The first thing we do is to determine the actual discontinuity in N across the gradient flow separatrix, which will differ from 1 because of lattice artifacts. To do this we generate an ensemble of configurations on the gradient flow separatrix. For each, we perform a short Hamiltonian evolution which crosses the separatrix, and we measure N at closely spaced intervals during the crossing to determine the discontinuity; it is $\Delta N = 0.982 \pm 0.005$, where the error is the standard deviation.⁹ This tells us how to re-scale all further measurements of N so that there will be an integer discontinuity at the separatrix (up to acceptably small noise).

Next, we measure the probability distribution for N . We only need to do this in the range $0 \leq N \leq \frac{1}{2}$, by periodicity; so we add 1 to all negative values of N and then put $N > \frac{1}{2}$ into range by $N \rightarrow 1 - N$. We use a continuous, piecewise linear reweighting function, with the widths of the linear pieces chosen by hand and the slopes of each segment determined by an automated bootstrapping procedure. We computed the probability distribution for two values of τ_0 ; $\tau_0 = 1.25a^2$ and $\tau_0 = 6.25a^2$. For the $\tau_0 = 1.25$ data we also recorded $N(\tau_0$

⁹It is important to compare the extrapolation based on a fit of the last few points on one side of the discontinuity with the value on the other, to remove errors from the time step size of the Hamiltonian evolution used to find the discontinuity.

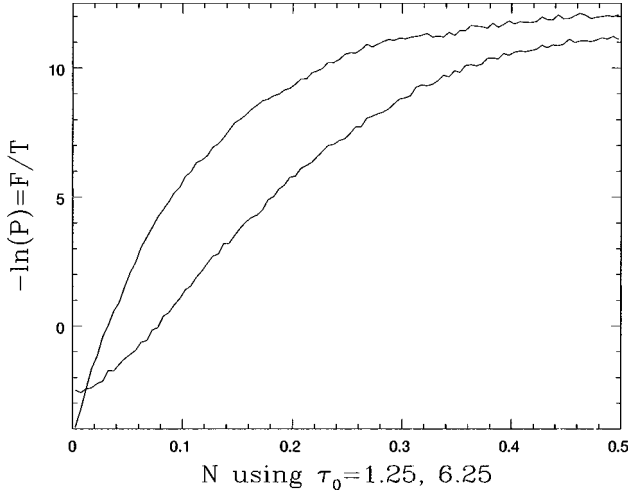


FIG. 10. Probability distribution of N for two values of τ_0 : $\tau_0 = 1.25a^2$ (the rounder curve which is lower at $N = \frac{1}{2}$) and $\tau_0 = 6.25a^2$ (the curve which is higher at $N = \frac{1}{2}$).

$= 6.25$); the data can therefore be used to get the probability distribution for either choice of τ_0 .

The probability distributions are compared in Fig. 10. The distributions are clearly different. Cooling longer before integrating $E_i^a B_i^a$ concentrates probability around $N=0$ and thins out the large N configurations. It is less likely to have N within some range of $\frac{1}{2}$ for the larger τ_0 . But this does not mean that the different τ_0 values give different values for the diffusion constant, as we still have not included $\langle |dN/dt| \rangle$ or the dynamical prefactor.

To get the dynamical prefactor and $\langle |dN/dt| \rangle$ we first need a sample of configurations very close to $N = \frac{1}{2}$. We get them by multicanonically sampling, not necessarily with the same reweighting function used to find the probability distribution of N . Then, for each, we choose momenta out of the appropriate distribution and perform a Hamiltonian evolution, with the algorithm of Ref. [18]. Once the Hamiltonian trajectory has settled into the neighborhood of a vacuum, we return to the configuration before we started the Hamiltonian evolution and reverse the sign of the momentum; then we evolve. Since momenta are odd and fields are even under time reversal, this computes the Hamiltonian trajectory in the backwards time direction. We determine $\langle |dN/dt| \rangle$ from the first time step in each time direction, and the dynamical prefactor from the number of $N = \frac{1}{2}$ crossings, as discussed in the previous section. It is also easy to find how many times the Hamiltonian trajectory crosses the gradient flow separatrix,




FIG. 11. A winding number changing section of a Hamiltonian trajectory in constrained volume Yang-Mills theory. After the crossing the system settles immediately into the neighborhood of a topological vacuum.

since $\int E_i^a B_i^a d^3x d\tau$ abruptly changes sign when it does; so we can also identify how well correlated crossings of the $N = \frac{1}{2}$ and gradient flow separatrices are. (We cannot directly determine the gradient flow separatrix prefactor if the correlation is not good, because our sampling procedure is for the $N = \frac{1}{2}$ separatrix and not the gradient flow one.)

Figure 11 gives an example of a Hamiltonian trajectory developed in this way, with $\tau_0 = 3.75a^2$. It is clear in this figure that once the trajectory settles in the neighborhood of a topological vacuum, in the sense that $N \sim 0$, then it stays there for some time. Since the maximal Lyapunov exponent of the Yang-Mills Higgs system is known to be about $0.3g^2T$ [52,53,19], the direction of the next permanent change in N_{CS} will surely be statistically independent from the previous one. This fact is essential to the whole approach; our most basic assumption is that the very long time N_{CS} diffusion is made up of a series of statistically independent steps, and we only need find out how frequently one of those steps is taken (which is what we are computing when we say we are computing γ_d).

Also note from Fig. 11 that there is no ‘‘overshoot’’ after reaching $N=0$, no sign that the trajectory is continuing in the direction of the next separatrix. This is true of all trajectories we have studied, both in finite volume and in the broken electroweak phase, which confirms the absence of prompt multiple crossings.

Our final results for this small volume system are presented in Table II. Note that both the probability to be near the separatrix and $\langle |dN/dt| \rangle$ vary quite a bit when we change τ_0 , but in opposite directions. Also, the shorter cooling leads to a smaller dynamical prefactor, though for the volume and cooling considered here the difference does not turn out to be large. The determined value of Γ for the two values of τ_0

TABLE II. Ingredients and results for the N_{CS} diffusion rate in a cubic, periodic volume $3/g^2T$ on a side. The measurements with two values of τ_0 disagree at 1.6σ .

quantity	value at $\tau_0 = 1.25a^2$	value at $\tau_0 = 6.25a^2$
$\ln P(N-0.5 < 0.05)$	-13.39 ± 0.22	-14.30 ± 0.18
$ dN/dt $	$(0.121 \pm 0.004)g^2T$	$(0.174 \pm 0.005)g^2T$
prefactor	0.52 ± 0.03	0.554 ± 0.024
$\ln[\Gamma/(\alpha_w T^4)]$	$-7.03 \pm .23$	-7.51 ± 0.19
Γ	$(0.0009 \pm 0.0002)\alpha_w^4 T^4$	$(0.0005 \pm 0.0001)\alpha_w^4 T^4$

differ at 1.6σ . This does not bother the author since it is the first statistical fluctuation above 1.5σ he has encountered since starting numerical work; he had one coming.

We can make a more stringent check of the τ_0 independence of the method by using the data set taken with $\tau_0 = 1.25a^2$ to determine the probability distribution of N for each value of τ_0 , since we recorded $N(\tau_0 = 6.25a^2)$ at each point while developing this data set. This data set gives a slightly different probability to be at large $N(\tau = 6.25a^2)$; the -14.30 in the table becomes -13.88 . Using this number, we get $\ln[\Gamma/(\alpha_w T)^4] = -7.09$. To compare to the $\tau_0 = 1.25a^2$ data we must remember that the probability distributions are now 100% correlated, so the expected difference is just from statistical errors in $\langle |dN/dt| \rangle$ and the prefactor. The difference in the logs of the rates is 0.06 ± 0.09 , so they do agree within error. The results for Γ_d are indeed τ_0 independent.

How do hard thermal loops change the rate? At the level of thermodynamics, hard thermal loops (HTL's) become just a Debye mass. They make the A_0 field heavier, which pushes it further to the regime where it decouples. So their thermodynamic influence is very small. Hence, they will barely change the probability distribution of N , which will have a well behaved large HTL strength limit. Similarly, as we have discussed, $\langle |dN/dt| \rangle$ does not depend on longitudinal E fields, and hence depends on the Debye mass only through its thermodynamic influence on the gauge fields. The flux through the separatrix should depend weakly on hard thermal loops and should have a good large HTL limit.

The dynamical prefactor is a completely different matter. It is a dynamical quantity which depends on unequal time correlators, potentially over quite long times. Hard thermal loops will change the time evolution of infrared degrees of freedom on all time scales longer than the inverse plasma frequency. Arnold, Son, and Yaffe (ASY) have argued [33] that hard thermal loops will suppress the baryon number violation rate by a factor parametrically of order $(g^4 T^2/m_D^2)$, because the number of crossings of the separatrix per permanent N_{CS} change will be of order $m_D^2/g^4 T^2$. That is, they predict that turning on hard thermal loops will reduce the dynamical prefactor by an amount linear in m_D^2 when it is large. Their arguments have recently been verified in the symmetric phase [26]. Now we need to check what the hard thermal loops do in the broken phase.

To address this question we include the hard thermal loops by using the technique developed in Ref. [26]. We add a large number of new ‘‘particle’’ degrees of freedom to the system, which propagate the hard thermal loop effects. For now we will be happy to know what the hard thermal loops do when they are extremely strong, so we can understand the parametric limit in which the ASY arguments should hold. With this in mind we put in particles of charge $Q = 0.1$ and number density $50/a^3$,¹⁰ which give a Debye mass of $43g^4 T^2$, an enormous number about 10 times the value in the MSM at a realistic value of g^2 . A Hamiltonian trajectory




FIG. 12. Hamiltonian trajectory through the separatrix in the symmetric phase in small volume, including hard thermal loops with $m_D^2 \sim 43g^4 T^2$. Plasma oscillations are clear, and there are numerous crossings of the separatrix.

crossing the separatrix is shown in Fig. 12. The qualitative features are indeed the same as Arnold, Son, and Yaffe predict, see Fig. 8 of Ref. [33]; plasma oscillations drive the system across the separatrix numerous times. The dynamical prefactor is correspondingly significantly smaller than without hard thermal loops; for these parameters it is about 0.16. Note, however, that we have had to add truly huge HTL effects to achieve this value, so at the realistic value the effect may not be too significant. We will study this question for the broken phase case in the next section.

How do our nonperturbative results compare to perturbation theory? We will not attempt to do a complete one loop calculation of the sphaleron rate in finite volume, but it is quite easy to compute the ‘‘sphaleron’’ energy, the energy of the saddle point between topological minima. We can use the technique of Ref. [54], or any other technique which can find a saddle point solution (we have one). We find $E = 27.77/N$, with N the linear dimension of the lattice. In our case that means $\beta E = 37$ and we would naively expect a rate suppression of $\exp(-37)$ before zero modes and the fluctuation determinant are included. This might be compared with the free energy difference between $N = 0$ and $N = \frac{1}{2}$, which is of order 15. We know that the complete inclusion of the zero modes and the fluctuation determinant is likely to make up some of this difference, but it certainly will not account for all of it. Nonperturbative physics is at work and it enhances the rate of sphaleron transitions in this case.

E. Broken phase rate

Now we will apply the same technology to the physically interesting case of large volume, broken phase $SU(2) \times U(1)$ Yang-Mills Higgs theory. Since the previous subsections already explained both the Monte Carlo update technique and the real time tools used to find $\langle |dN/dt| \rangle$ and the dynamical prefactor, we will just discuss here how this case is different from the small volume, symmetric phase calculation, and what we have to do differently to get it to work.

As we discussed, we will work at the critical temperature, which corresponds in the 3D language to the critical Higgs boson mass. This means that the broken phase, which we want to study, is actually only metastable; the symmetric

¹⁰See Ref. [26] for the implementation and the definitions of these quantities, and their relation to the Debye mass.

phase is equally thermodynamically preferable. In a small volume there can be fairly easy tunneling between the two, but the metastability becomes stronger as the volume becomes larger. We must choose a volume which is large enough that metastability is very strong and tunneling between the phases will not occur. This means that the physical size of the lattice we use must be significantly larger than the physical size of the sphaleron, which would have to be true anyway to keep the exponential tails of the sphaleron from seeing each other around the periodic boundary conditions.

The problem of tunneling to the symmetric phase is made worse because the sphaleron has a zero of the Higgs condensate at its core, so it looks somewhat similar to a symmetric phase bubble. To keep from nucleating to the symmetric phase we should use a volume big enough that the free energy of the state intermediate between phases is comparable to the free energy of the sphaleron. We have studied three values for $x \equiv \lambda/g^2$, $x=0.047$, 0.039 , and 0.033 ; for the former two we used a physical volume of $(16/g^2T)^3$ and for the latter we used $(13.33/g^2T)^3$. These were all sufficient to prevent nucleation of the symmetric phase, but for $x=0.047$, a volume of $(12.8/g^2T)^3$ was not. The volume requirement becomes less severe as the phase transition becomes stronger at smaller x , so for $x=0.039$ and 0.033 we used a good margin of excess volume. The volume requirement would also have been less severe if we worked below, rather than at, the equilibrium Higgs boson mass parameter.

The need for a large volume drives up numerical costs in two ways. One is obvious; we need to update a lot of lattice volume which is “dead weight” since the sphaleron is not sitting there. But the large volume also makes the multicanonical algorithm perform worse. Examining Fig. 10, we see the free energy rises roughly linearly with N at small values of N . This behavior is also expected analytically in the broken phase case, see, for instance, Ref. [55]. Naively, then, having the gauge fields being part way up the sphaleron at one place in the box is not thermodynamically favored over having them half as far up the sphaleron in two different places, at least for relatively small N . In fact the entropy from getting to choose the locations of the two places means that this may be slightly preferred to being further up the sphaleron in one place, for the range of N where the free energy is varying linearly with N . But at larger N the free energy for a single “near sphaleron” levels off while that for two continues to rise linearly, so near $N=\frac{1}{2}$ we prefer having a single sphaleron. Somewhere in between there is a mismatch in what kind of configuration is dominating the ensemble, and such a mismatch can reduce the efficiency of the Monte Carlo calculation.

To cut the numerical demand we integrate out A_0 fields (i.e., we do not enforce the Gauss law when we study dynamics), which allows us to use the very efficient update algorithm of Rummukainen *et al.* [5]. In fact a single step of the Rummukainen *et al.* update (one heat bath and four overrelaxation sweeps) is far too large an update of the fields; if the value of $f(N)$ typically changes by more than 1 under one update then the accept rate for updates becomes very small. Instead, we alternate between performing a “scaled back” version of the update in which only some “mod p

checkerboard” of the sites are updated, and an overrelaxation sweep updating only the Higgs fields. Since the overrelaxation sweep of the Higgs fields does not change the gauge fields, and since N depends on the gauge fields alone, this sweep can be automatically accepted. Doing this means that the evolution of the Higgs fields can be made “fast” in comparison to the update of the gauge fields. That ensures that the thermalization of the gauge fields is not gummed up by slow evolution of the Higgs fields. Even with the extra updates for the Higgs fields, and the blocking procedure for accelerating the measurement of N , most machine time is spent measuring N , and further improvements in the update algorithm will not help.

The results for the flux through the $N=\frac{1}{2}$ separatrix are given in Table I, and also in Ref. [16]. We have also measured the dynamical prefactor for the $x=0.039$ case, without added hard thermal loops or enforcement of the Gauss law. The value we get is 0.33 ± 0.05 , slightly lower than the value for the gradient flow separatrix, which is 0.40 ± 0.05 (we measure this by choosing a value of τ_0 large enough that the two separatrices almost coincide, so there is a one to one correspondence between crossing $N=\frac{1}{2}$ and crossing the gradient flow separatrix). It is not clear whether this represents some interesting dynamical behavior of the theory or whether it means that the Yang-Mills gradient flow separatrix is not the optimal divider between topological vacua.

Another interesting question is how the dynamical prefactor depends on hard thermal loops. Unfortunately, the numerical cost of using the “particles” technique of Ref. [26] is so high in this context that we have to cut a few more corners to make the calculation. We drop the U(1) factor, increase the lattice spacing from $a=2/(5g^2T)$ to $a=1/(2g^2T)$ ($\beta=8$), and reduce the volume to $(14/g^2T)^3$. To prevent nucleation to the symmetric phase, we work at a somewhat larger value of x , $x=0.042$, and below the equilibrium temperature, so the broken phase is more stable than the symmetric phase. We choose the temperature (the thermal Higgs boson mass, really) so the Higgs condensate is $\phi_0=1.7gT$, just more than enough to make the sphaleron rate too low to erase baryon number (we expect). We cannot directly measure the sphaleron rate with this set of parameters because the update including particles is expensive enough that the multicanonical calculations are prohibitive, but we can get a good sample of points on the separatrix by using a reweighting which favors $N \sim \frac{1}{2}$ strongly, and we can get sufficient statistics for the dynamical prefactor to make a good determination.

We can “shut off” the hard thermal loop contribution to the dynamics without changing their contribution to thermodynamics by not allowing the particles to move during the Hamiltonian evolutions used to determine the dynamical prefactor. Also, if the arguments of Arnold are correct [56], changing the velocity at which the particles move changes their contribution to the key part of the hard thermal loops, the $\omega \ll k \sim g^2T$ regime, linearly in the velocity. Hence we could explore very strong hard thermal loops either by putting in very many particles, or by making them move very fast. The numerical cost is the same but the memory costs favor the latter.




FIG. 13. Free energy (left) and $(\phi^\dagger \phi_{\text{broken}} - \phi^\dagger \phi_{\text{symm}})/(g^2 T^2)$ as functions of N at $x \equiv \lambda/g^2 = 0.039$, in a $(16/g^2 T)^3$ volume [at $\tau_0 = 3.6/(g^2 T)^2$]. In each case the upper curve is the broken phase and the lower curve is the symmetric phase. The plot of ϕ^2 shows that the volume used was large enough that the sphaleron did not bring us anywhere near a transition to the symmetric phase. The behavior of the two phases is completely different; in the broken phase there is a free energy barrier, and in the symmetric phase there is not.

We use a realistic total hard thermal loop strength, $m_D^2 = 5g^4 T^2$ from particles. [The standard model value is $m_D^2 = (11/6)g^2 T^2$, and since $g^2 \sim 0.40$ on performing the dimensional reduction calculation [30], this is just smaller than the value we used.] The dynamical prefactor changes only mildly, from 0.52 ± 0.05 to 0.40 ± 0.05 , when we turn the particles on. When we increase the particle velocities to 4 times the speed of light, the prefactor becomes 0.15 ± 0.03 . Hard thermal loops do indeed reduce the dynamical prefactor, which is already less than 1 without them. However, the parametric limit in which the reduction is large is not achieved for realistic parameter values. Also note that the value 0.52 is larger than we got at $x=0.039$ and $m_H^2 = m_{H,\text{crit}}^2$. We assume this is because the larger value was evaluated below the equilibrium temperature, where the Higgs condensate is larger and stiffer and the sphaleron is smaller and more energetic; its decay should be more vigorous and less susceptible to buffeting by large IR fields.

V. RESULTS

We present our results for the diffusion constant for N_{CS} in Table I. The first quoted value is without the dynamical prefactor. The dynamical prefactor is less than 1 for the ($N = \frac{1}{2}$) separatrix we have used, and also for the gradient flow separatrix; its value for the $x=0.039$ data is about 0.33, and we will take this to be representative of the other two values of x as well. It is not clear whether the prefactor is less than 1 because our separatrix is suboptimal, or because the dynamics are nontrivial in a way which often leads to multiple crossings.

We have found that the dynamical prefactor depends on the strength of hard thermal loops; but for the realistic Debye mass the effect is not strong. A reasonable estimate for the relevant dynamical prefactor for the cases of interest and for

the value of τ_0 we used to define the measurable N is around 0.3 ± 0.1 . We have included a row in the table where we include this (estimated) effect in the rate.

For amusement, we have also compared the broken and symmetric phase probability distributions for N . We present the results in Fig. 13, which also shows how ϕ^2 varies with N in each case. The two data sets were taken using identical values for all parameters [lattice spacing $a=2/5g^2 T$, $x=0.039$, critical temperature (i.e., critical Higgs mass), lattice volume $= (16/g^2 T)^3$, and $\tau_0 = 3.6/(g^2 T)^2$], but starting with a broken phase initial condition in one case and a symmetric phase initial condition in the other. A barrier to changing N_{CS} is clearly present in the broken phase case and clearly absent in the symmetric phase case.

We cannot use the data for the symmetric phase case to determine Γ_d in the symmetric phase. Although it would be straightforward to compute $\langle |dN/dt| \rangle$ and get the flux through the separatrix, the calculation of the dynamical prefactor is impossible. A Hamiltonian path through the separatrix does not settle into the neighborhood of a topological vacuum and stay there for a long time; it just continues to wander around, as we already saw in Fig. 5. But we can measure Γ_d in the symmetric phase with purely real time techniques; the value including hard thermal loops [and actually for pure SU(2) Yang-Mills theory] is $\Gamma_d = (29 \pm 6)\alpha_w^5 T^4$, or $-\ln \Gamma_d \approx 13.9$. (This result may get revised downwards somewhat when the issues involving logarithmic corrections [27] have been fully accounted for.) As expected, the broken phase rate at small $x = \lambda/g^2$ is enormously smaller; for $x=0.033$ the ratio is about 10^8 .

A. Comparison to perturbation theory

We want to compare the determined sphaleron rate to two things; perturbation theory and the value required to avoid

erasure of baryon number generated at the electroweak phase transition. One loop perturbation theory gives [3]¹¹

$$\Gamma_d = 4T^4 \frac{\omega_-}{g\phi_0} \left(\frac{\alpha_W}{4\pi}\right)^4 \left(\frac{4\pi\phi_0}{gT}\right)^7 \mathcal{N}_{\text{tr}} \mathcal{N}_{\text{rot}} \kappa e^{-\beta E_{\text{sph}}}. \quad (40)$$

Here ϕ_0 is the broken phase Higgs condensate expectation value, ω_- is the unstable frequency of the sphaleron, $\mathcal{N}_{\text{tr}} \mathcal{N}_{\text{rot}}$ are zero mode factors, κ is the one loop fluctuation determinant, and E_{sph} is the energy of the Klinkhamer-Manton sphaleron, using the tree level Hamiltonian. For small $\lambda/g^2 \equiv x$, $-T \ln \kappa$ equals the energy due to the one loop effective potential term, plus a modest correction [15]. We can guess that the dominant two loop corrections to Eq. (40) are absorbed by including the two loop effective potential terms in the Hamiltonian. So it seems reasonable to estimate the sphaleron rate by Eq. (40), but setting $\kappa=1$ and solving for the sphaleron energy using the two loop effective potential at the equilibrium temperature. One should also solve for the zero modes and ω_-/ϕ_0 at this value, but they are very weak functions of the effective potential [14]. We use the values from Ref. [14] at $x=(\lambda/g^2)=0.04$ for these, but solve for the sphaleron energy, $E_{\text{sph}}=4\pi B\phi_0/g$, numerically, using the two loop effective potential at T_c . We use the two loop potential presented in Ref. [57], without pieces from longitudinal gauge bosons (assumed integrated out). We also drop two loop terms proportional to λg^2 or λ^2 , because the perturbative determination of ϕ_0 is an expansion in λ/g^2 , and such terms contribute at the same or higher order as unknown three loop terms. (Including these two loop terms moves ϕ_0 closer to the nonperturbative value by an insignificant amount.) The ‘‘two loop’’ analytic sphaleron rate, also included in Table I, is about $\exp(2.5)$ times faster than the numerically determined nonperturbative rate, and falls further off when we include the dynamical prefactor.¹² The difference is more than can be explained by the difference in ϕ_0 , but it is not huge in the sense that it represents a change of less than 20% in the exponent. The difference gets smaller, relative to the exponent, as the sphaleron energy becomes larger.

B. Comparison to the erasure bound

We should compare the sphaleron rate to the limit set by requiring that baryon number not be erased. The rate at which sphalerons degrade baryon number is [3]

$$\frac{1}{N_B} \frac{dN_B}{dt} = -\frac{13N_F}{4} \Gamma_d T^{-3}, \quad (41)$$

where $N_F=3$ is the number of families, and the numerical factor $13N_F/4$ would be smaller in theories, such as super-

symmetry, in which additional degrees of freedom can store baryon number.¹³ Integrating from the end of the phase transition to the present day,

$$\ln[N_B/N_B(T_c)] = -\frac{13N_F}{4} \int_{t_0}^{\infty} \Gamma_d [T(t)] T^{-3}(t) dt, \quad (42)$$

where we have shown the dependence of Γ_d on T and of T on t .

Now $\ln \Gamma_d$ is very sensitive to $\sqrt{\langle \phi^\dagger \phi \rangle}$, and hence to T ; so we can approximate $\ln \Gamma_d(T) \approx \ln \Gamma_d(T_c) + (T - T_c)(d \ln \Gamma_d/dT)|_{T=T_c}$, and perform the integral:

$$\begin{aligned} \ln\{-\ln[N_B/N_B(T_c)]\} &= \ln(39/4) + \ln\left(\frac{\Gamma_d(T_c)}{T^4}\right) \\ &\quad - \ln\left(-\frac{d \ln \Gamma_d(T(t))}{T dt}\right)\Bigg|_{T=T_c}. \end{aligned} \quad (43)$$

By the chain rule,

$$\frac{d \ln \Gamma_d}{T dt} = \frac{d \ln \Gamma_d}{dy} \frac{dy}{dT} \frac{dT}{dt}, \quad (44)$$

where $y = m_H^2(T)/(g^4 T^2)$ is the dimensionless thermal Higgs boson mass squared.

We get dy/dT from the one loop correction to m_H^2 [57],

$$\frac{dy}{dT} \approx \frac{8\lambda + 4g_y^2 + g^2(3 + \tan^2 \Theta_W)}{8g^4 T}, \quad (45)$$

and we get $d \ln T/dt$ from the Friedmann equation in a radiation dominated universe,

$$\frac{1}{4t^2} = H^2 = \frac{8\pi G}{3} \frac{\pi^2 g_*}{30} T^4 \Rightarrow \frac{d \ln T}{dt} = -\sqrt{\frac{4\pi^3 g_*}{45}} \frac{T^2}{m_{\text{Pl}}}, \quad (46)$$

where g_* is the number of radiative degrees of freedom in the universe ($g_* = 106.75$ in the minimal standard model) and $m_{\text{Pl}} \approx 1.22 \times 10^{19}$ GeV is the Planck mass. Finally, we determine $d \ln \Gamma_d/dy$ perturbatively, by varying y slightly from the equilibrium value and recomputing the two loop sphaleron rate. The dependence is quite strong. We include it in Table I.

The most widely cited discussion of baryon number erasure after the phase transition makes the approximation that the baryon number violation rate after the phase transition is constant for about one Hubble time [58]. In fact, because Γ_d depends very strongly on y , which in turn depends strongly

¹¹The definition of Γ used in Refs. [3,14] is the response rate to a chemical potential, which is half the diffusion rate [37]; so Eq. (40) differs by a factor of 2 from the expressions in those references.

¹²There is no literature calculation of the dynamical prefactor, including hard thermal loops but in the perturbative context.

¹³Again, there is a factor of 2 difference from the reference because they write in terms of the response to a chemical potential, which is half the diffusion constant. This 2 cancels the other 2.

on T , most baryon number erasure occurs in the first 10^{-3} Hubble times after the phase transition. Hence the initial rate of baryon number violation, $\Gamma_d(T_c)$, which prevents wash-out is 10^3 times larger than assumed in Ref. [58], leading to a weaker bound on $\Gamma_d(T_c)$, roughly

$$-\ln[\Gamma_d(T_c)T_c^{-4}] > 30.4 - \ln(T_c/100 \text{ GeV}). \quad (47)$$

The values of g_* and dy/dT will both be larger in supersymmetric extensions of the standard model. Unless quite a number of supersymmetric partners have masses under 100 GeV, which now seems unlikely, g_* will not be too much larger; but if there are stops with supersymmetry breaking masses of less than 100 GeV, which is necessary to get a strong enough transition without violating the current experimental Higgs boson mass bound, then Tdy/dT gets extra contributions from top squarks which bring it up by about a factor of 2. The bound, Eq. (47), is weakened by about 1. Also note that, because Eq. (43) is for the double log of $N_B/N_B(T_c)$, failing to meet the bound by 1 means the baryon number is diminished by $\exp[\exp(1)] \approx 15$, and failure by 2 reduces baryon number by $\exp[\exp(2)] \approx 1600$; so the bound is quite sharp.

VI. CONCLUSION

We have shown how to define a lattice measurable which allows a (multicanonical Monte Carlo calculation) nonperturbative determination of the broken phase sphaleron barrier height. We have combined this with real time techniques to measure the diffusion constant of N_{CS} (and hence baryon number violation) in the broken electroweak phase nonperturbatively, including the dynamical prefactor. We find that the diffusion constant is smaller by about $\exp(-3.6)$ than in a perturbative estimate using the two loop effective potential and no wave function corrections (and assuming a dynamical prefactor of 1). The difference is too large to ascribe to the difference between the perturbative and nonperturbative value of ϕ_0 . However, it represents a shift in the exponent of less than 20% from the perturbative estimate.

We have also demonstrated that the physics of hard thermal loops does change the sphaleron rate in the broken phase, apparently consistently with the arguments of Arnold, Son, and Yaffe [33]. But to really achieve the parametric limit they discuss takes an unrealistically large Debye mass; for physical values of the parameters, the correction due to hard thermal loops is fairly minor.

Interpolating between the values of $x \equiv \lambda/g^2$ where we have measured, and including the estimate for the dynamical prefactor, we get a bound of about $(\lambda/g^2) \equiv x = 0.037$ in the standard model and $x = 0.039$ in the MSSM when it can be

perturbatively reduced to a standard-model-like effective theory. These are slightly looser than we quote in Ref. [16] because the measured value of the dynamical prefactor is smaller than our estimate there. Another convenient way to state our result is that the bound on the Higgs condensate after the phase transition is about $\phi_0 = 1.67gT$ in the MSM and $\phi_0 = 1.60gT$ in the MSSM (with the MSSM value weaker because the larger temperature dependence of the thermal Higgs boson mass makes the erasure rate fall off faster with time after the transition).

The bound is softened if the universe does not reheat to T_c during the phase transition, because the phase transition then ends at a lower temperature with a larger Higgs vacuum expectation value. Incomplete reheating may well be generic. Unfortunately, we do not have a nonperturbative measurement of the bubble nucleation action, which is required to determine definitively whether reheating happens. This is an interesting project which can perhaps be approached by techniques similar to what we have used here, i.e., defining a separatrix corresponding to the critical bubble (as the separatrix here corresponds to the sphaleron), and making a complete calculation including the dynamical prefactor.

We should comment on what we expect to be true below the equilibrium temperature and in the case of a light top squark. As the temperature falls below equilibrium and the Higgs condensate becomes larger, perturbation theory should become better, at least in the sense that the error in the exponent should get smaller compared to the magnitude of the exponent. It would be very surprising if the nonperturbative rate switches to being faster than the perturbative estimate, though. It is straightforward, though expensive, to repeat the analysis here for temperatures below the equilibrium point. Probably it will only be worth it after we have nonperturbative information on the bubble nucleation rate. In the case of the light top squark, we expect the most important differences to be in the effective potential, in which case perturbation theory should do as well as it does here, if we permit a ‘‘by hand’’ correction for the strength of the phase transition. We expect this because, while the top squark contributes at one loop to the effective potential, scalar interactions only influence wave function corrections (say, the Higgs gradient energy of the sphaleron) at two loops, and there is no direct interaction between a light right top squark and the SU(2) gauge fields.

ACKNOWLEDGMENTS

I would like to thank Jim Cline, Jim Hetrick, Kari Rummukainen, Misha Shaposhnikov, and Alex Krasnitz for discussions and correspondence. I also thank Berndt Müller and the North Carolina Supercomputing Center.

-
- [1] G. 't Hooft, Phys. Rev. Lett. **37**, 8 (1976).
 [2] V. Kuzmin, V. Rubakov, and M. Shaposhnikov, Phys. Lett. **155B**, 36 (1985).
 [3] P. Arnold and L. McLerran, Phys. Rev. D **36**, 581 (1987).

- [4] M. Gavela, P. Hernandez, J. Orloff, and O. Pene, Nucl. Phys. **B430**, 345 (1994); M. Gavela, P. Hernandez, J. Orloff, O. Pene, and C. Quimbay, *ibid.* **B430**, 382 (1994).
 [5] K. Kajantie, M. Laine, K. Rummukainen, and M. Shaposhni-

- kov, Nucl. Phys. **B466**, 189 (1996); Phys. Rev. Lett. **77**, 2887 (1996).
- [6] M. Carena, M. Quiros, A. Riotto, I. Vilja, and C. E. M. Wagner, Nucl. Phys. **B503**, 387 (1997).
- [7] H. Davoudiasl, K. Rajagopal, and E. Westphal, Nucl. Phys. **B515**, 384 (1998).
- [8] J. Cline, M. Joyce, and K. Kainulainen, Phys. Lett. B **417**, 79 (1998).
- [9] J. Cline and K. Kainulainen, Nucl. Phys. **B482**, 73 (1996); **B510**, 88 (1998); M. Losada, Phys. Rev. D **56**, 2893 (1997); M. Losada and G. Farrar, Phys. Lett. B **406**, 60 (1997); M. Losada, hep-ph/9612337; M. Laine, Nucl. Phys. **B481**, 43 (1996).
- [10] K. Kajantie, M. Laine, K. Rummukainen, and M. Shaposhnikov, Nucl. Phys. **B493**, 413 (1997).
- [11] H. Dosch, J. Kripfganz, A. Laser, and M. Schmidt, Phys. Lett. B **365**, 213 (1996); M. Gurtler, E. Ilgenfritz, J. Kripfganz, H. Perlt, and A. Schiller, Nucl. Phys. **B483**, 383 (1997); M. Gurtler, E. Ilgenfritz, and A. Schiller, Eur. Phys. J. C **1**, 363 (1998); Phys. Rev. D **56**, 3888 (1997).
- [12] D. Bödeker, P. John, M. Laine, and M. Schmidt, Nucl. Phys. **B497**, 387 (1997).
- [13] M. Laine and K. Rummukainen, Phys. Rev. Lett. **80**, 5259 (1998); Report No. CERN-TH-98-122, hep-lat/9804019.
- [14] L. Carson and L. McLerran, Phys. Rev. D **41**, 647 (1990); L. Carson, X. Li, L. McLerran, and R. Wang, *ibid.* **42**, 2127 (1990).
- [15] J. Baacke and S. Junker, Phys. Rev. D **49**, 2055 (1994); **50**, 4227 (1994).
- [16] G. D. Moore, Report No. MCGILL-97/36, hep-ph/9801204.
- [17] D. Grigorev and V. Rubakov, Nucl. Phys. **B299**, 248 (1988).
- [18] J. Ambjørn, T. Askgaard, H. Porter, and M. Shaposhnikov, Nucl. Phys. **B353**, 346 (1991).
- [19] J. Ambjørn and A. Krasnitz, Phys. Lett. B **362**, 97 (1995).
- [20] G. D. Moore, Nucl. Phys. **B480**, 657 (1996).
- [21] G. D. Moore, Nucl. Phys. **B480**, 689 (1996).
- [22] W. Tang and J. Smit, Nucl. Phys. **B482**, 265 (1996).
- [23] G. D. Moore and N. Turok, Phys. Rev. D **55**, 6538 (1997).
- [24] G. D. Moore and N. Turok, Phys. Rev. D **56**, 6533 (1997).
- [25] J. Ambjørn and A. Krasnitz, Nucl. Phys. **B506**, 387 (1997).
- [26] G. D. Moore, C. Hu, and B. Müller, Phys. Rev. D **58**, 045001 (1998).
- [27] D. Bödeker, Phys. Lett. B **426**, 351 (1998).
- [28] P. Ginsparg, Nucl. Phys. **B170**, 388 (1980); T. Applequist and R. Pisarski, Phys. Rev. D **23**, 2305 (1981); S. Nadkarni, *ibid.* **27**, 917 (1983); N. P. Landsmann, Nucl. Phys. **B322**, 498 (1989).
- [29] K. Farakos, K. Kajantie, K. Rummukainen, and M. Shaposhnikov, Nucl. Phys. **B425**, 67 (1994).
- [30] K. Kajantie, M. Laine, K. Rummukainen, and M. Shaposhnikov, Nucl. Phys. **B458**, 90 (1996).
- [31] G. D. Moore, Phys. Rev. D **53**, 5906 (1996).
- [32] D. Bödeker, Nucl. Phys. **B486**, 500 (1997).
- [33] P. Arnold, D. Son, and L. Yaffe, Phys. Rev. D **55**, 6264 (1997).
- [34] P. Huet and D. Son, Phys. Lett. B **393**, 94 (1997).
- [35] D. Son, Report No. UW/PT-97-19, hep-ph/9707351.
- [36] F. Klinkhamer and N. Manton, Phys. Rev. D **30**, 2212 (1984).
- [37] S. Khlebnikov and M. Shaposhnikov, Nucl. Phys. **B308**, 885 (1988).
- [38] M. Lüscher, Commun. Math. Phys. **85**, 39 (1982).
- [39] P. Woit, Phys. Rev. Lett. **51**, 638 (1983); Nucl. Phys. **B262**, 284 (1985).
- [40] A. Phillips and D. Stone, Commun. Math. Phys. **103**, 599 (1986).
- [41] E. Farhi, J. Goldstone, S. Gutmann, K. Rajagopal, and R. Singleton, Jr., Phys. Rev. D **51**, 4561 (1995).
- [42] G. D. Moore, Phys. Lett. B **412**, 359 (1997).
- [43] J. Kogut and L. Susskind, Phys. Rev. D **11**, 395 (1975).
- [44] J. Hetrick and Ph. de Forcrand, Nucl. Phys. B (Proc. Suppl.) **63**, 838 (1998).
- [45] N. Manton, Phys. Rev. D **28**, 2019 (1983).
- [46] K. Kajantie, K. Rummukainen, and M. Shaposhnikov, Nucl. Phys. **B407**, 356 (1993).
- [47] A. Berg and T. Neuhaus, Phys. Lett. B **331**, 249 (1991).
- [48] G. D. Moore, Nucl. Phys. **B493**, 439 (1997).
- [49] G. D. Moore, Nucl. Phys. **B523**, 569 (1998).
- [50] J. Cline and G. D. Moore, Report No. MCGILL-98-11, hep-ph/9806354.
- [51] F. Csikor, Z. Fodor, J. Hein, and J. Heitger, Phys. Lett. B **357**, 156 (1995).
- [52] B. Müller and A. Trayanov, Phys. Rev. Lett. **68**, 3387 (1992).
- [53] T. Biro and M. Thoma, Phys. Rev. D **54**, 3465 (1996).
- [54] M. Garcia Perez and P. van Baal, Nucl. Phys. **B429**, 451 (1994).
- [55] D. Diakonov, M. Polyakov, S. Sieber, J. Schaldach, and K. Goeke, Phys. Rev. D **49**, 6864 (1994).
- [56] P. Arnold, Phys. Rev. D **55**, 7781 (1997).
- [57] Z. Fodor and A. Hebecker, Nucl. Phys. **B432**, 127 (1994).
- [58] M. Shaposhnikov, Nucl. Phys. **B287**, 757 (1987).