

Expanded ensemble method for bubble nucleation

Jaakko Hällfors^{*} and Kari Rummukainen[†]

Department of Physics and Helsinki Institute of Physics, PL64, FI-00014 University of Helsinki, Finland



(Received 3 April 2025; accepted 13 May 2025; published 17 June 2025)

In the absence of impurities and boundary effects, first order phase transitions are initiated by the nucleation of critical bubbles. In thermally driven transitions many systems can remain metastable for an extended time, possibly tens of orders of magnitude longer than typical microscopic timescales. In standard Markov chain Monte Carlo simulations of these systems the probability of critical bubbles can be too suppressed for the transition to happen in any practical simulation time. The computation can be accelerated by using modified sampling methods, for example multicanonical or Wang-Landau sampling. However, even using these methods, there remains a condensation barrier which dramatically reduces the efficiency at large volumes. We present a novel sampling method, the method of expanded ensembles, which very effectively circumvents the condensation barrier and enables efficient simulations at large volumes.

DOI: [10.1103/b61n-q1qq](https://doi.org/10.1103/b61n-q1qq)

I. INTRODUCTION

In the Standard Model of particle physics the high-temperature electroweak phase transition is a smooth crossover [1,2]. However, a first order electroweak phase transition is possible in many extensions of the Standard Model. In this case the phase transition does not happen immediately as the Universe cools down to the critical temperature; the system first supercools, and the potential develops a new minimum. Initially the fields are trapped in a metastable phase, separated from the stable minimum by an energy barrier. The phase transition transpires through the nucleation of bubbles of the stable phase. These bubbles then expand and coalesce to complete the transition into the stable phase. Since bubble nucleation is a violent nonequilibrium event, it could provide conditions for a successful baryogenesis [3] and generate primordial gravitational waves [4,5].

The planned Laser Interferometer Space Antenna should be particularly suitable for detecting gravitational waves originating from first order phase transitions [6,7]. Since the goal is to use the gravitational wave background for differentiating between the various models, we must be able to predict the spectra accurately. In this work we focus on one of the key parameters, the nucleation rate, which

essentially describes how frequently bubbles of the stable phase appear [5].

Given a prospective model, the nucleation rate can be computed using either perturbative or nonperturbative methods. A classical prescription for the nucleation rate computation was first presented by Langer in [8] with the additional effect of quantum tunneling described by Affleck in [9]. A similar approach for quantum fields at zero temperature was derived by Callan and Coleman in [10,11] and extended to finite temperature by Linde in [12].

For a phase transition at the electroweak scale, the computational problem can be simplified through dimensional reduction; when the temperature is high, heavy scales can be integrated out, leaving us with a three-dimensional effective field theory (EFT). This EFT only contains the bosonic infrared modes, and its thermodynamics are much more amicable to analysis on the lattice [13]. Dimensional reduction also transforms the nucleation rate computation into a problem in (almost) classical statistical field theory. In fact, the standard nonperturbative approach is conceptually most similar to Langer's formulation.

Simply put, the lattice computation of the nucleation rate can be factored into a statistical and a dynamical part [14]. The statistical factor gives the probability of critical bubbles relative to the metastable phase, while the dynamical factor describes how often the critical bubbles actually evolve into the stable phase. The dynamical factor is computed using Langevin dynamics for the critical bubbles and is in principle easy to perform. The statistical part involves simple Markov chain Monte Carlo (MCMC) sampling of the probability distribution. Due to the extreme probability suppression of the critical configurations interpolating the two phases, multicanonical methods are required.

^{*}Contact author: jaakko.hallfors@helsinki.fi

[†]Contact author: kari.rummukainen@helsinki.fi

While the multicanonical method solves the most glaring part of the sampling problem it does not remove all obstacles. Firstly, the periodicity of the lattice results in so-called shape transitions, where the phase interface must morph between different geometric shapes in order to pass from one phase to the other [14,15]. A less obvious and more computationally relevant barrier is given by the condensation transition, where the preferred configuration changes from delocalized fluctuations of the metastable phase into a single localized bubble of the stable phase. In the MCMC computation these transitions appear as barriers to the Markov chain that slow down its exploration of the phase space. Since we need to obtain the probability difference between the regions separated by the barrier, simulations on larger lattices become increasingly arduous. The desire for improved tunneling has been expressed in [16], with improvements presented in [17,18]. In this work we develop a new method, one that eliminates the condensation barrier and is easily applicable to current systems of interest.

To understand the shortcomings of the standard approach, we begin by describing the sampling problem in Sec. II. We pay particular attention to the effect of the multicanonical weighting to clarify the type of problem it *does not solve*. With these ideas in mind, in Sec. III we investigate the tunneling problem in more detail, briefly with respect to shape transitions, but mostly focusing on the computationally relevant condensation transition and the associated transition barrier. We also consider the effect of using so-called weighted order parameters to manipulate the location of the condensation transition. In Sec. IV we describe the method of expanded ensembles and use weighted order parameters to construct a novel set of subensembles to bypass the condensation barrier. Finally, in Sec. V the effectiveness of the new method is verified for the case of a simple two-dimensional scalar field. We show that our new method eliminates the characteristic exponential slowing down near the condensation transition. Future applications and possible variants of the method are discussed in Sec. VI.

II. THE NUCLEATION RATE AND THE MULTICANONICAL METHOD

The nonperturbative computation of the bubble nucleation rate presents a special type of the sampling problem. Let us consider a system supercooled below the critical temperature of a first order phase transition: the system is initially in the metastable phase, and the tunneling to the stable phase is suppressed by the free energy barrier between the phases. The transition to the stable phase happens along the trajectory with the smallest increase in free energy, and the critical bubble is the configuration with the maximum of the free energy on this trajectory, i.e., a saddle point of the free energy landscape.

The energy budget of the bubble consists of the favorable formation of the stable phase and the expense of forming a phase interface. Since the former is proportional to volume and the latter to surface area, there exists a critical size (or radius) where bubble growth eventually becomes energetically favorable. This defines the size of the critical bubble.

Langer's formalism consists of investigating the probability flux through this saddle point [8]. In particular, the nucleation rate factorizes into a statistical part (statistical probability of the saddle-point configurations) and a dynamical part (the probability flux through the saddle point). Measuring the dynamical part requires a real-time approach, whereas the statistical part can be measured using equilibrium thermodynamics.

For weakly coupled gauge field theories (Standard Model-like) both the statistical and the dynamical part can be measured in lattice simulations [14,16]. The procedure is essentially equivalent to the earlier nonperturbative sphaleron rate computation of [19].

Directing our focus to the statistical factor, let us suppose a system described by a state s , and the canonical probability distribution $P_{\text{ca}}(s) \propto \exp[-H(s)]$. We assume that the system has two phases A (metastable) and B (stable), separated by a region of suppressed phase space. Here we take that an order parameter is a real scalar function $\mathcal{O}(s)$ such that it obtains distinct values, $\mathcal{O}_A$ and $\mathcal{O}_B$, in the two phases. The choice of the order parameter $\mathcal{O}$ is not unique, but a typical candidate usually exists for a given system (e.g., magnetization for the Ising model).

As the system traverses from one phase to the other, the order parameter obtains values $\mathcal{O}_A < \mathcal{O}(s) < \mathcal{O}_B$ in the intermediate configurations. This gives $P_{\text{ca}}(\mathcal{O})$ the characteristic two-peaked structure sketched in Fig. 1. The minimum between the peaks at $\mathcal{O}_{\text{critical}}$ contains configurations at and near the saddle point. It is precisely the value of $P_{\text{ca}}(\mathcal{O}_{\text{critical}})$ relative to the metastable peak that we seek.

The obvious hindrance to the MCMC sampling of P_{ca} results from the fact that the suppression of the bubble configurations can be extreme, even of the order $\sim e^{-100}$. Typical MCMC methods will never visit these configurations, since they are not designed to sample suppressed




FIG. 1. Sketch of $P(\mathcal{O})$ below the critical temperature.

parts of the phase space. The *multicanonical method*, presented first in [20], is however particularly suitable for first order phase transitions. Other modified sampling methods include Wang-Landau [21], density of states [22] and the application of Jarzynski's theorem [23]. The multicanonical method has been successfully applied to the critical bubble calculation in models related to the Standard Model [14,16,17], and to $SU(N)$ gauge theory [18]. The method is based on sampling an alternative probability distribution where (optimally) the suppression does not exist.

To elaborate, let us define the *multicanonical probability distribution* for the state s by

$$P_{\text{muca}}(s) := P_{\text{ca}}(s)P_w(\mathcal{O}(s)) \quad (1)$$

$$\propto \exp[-(H(s) + W(\mathcal{O}(s)))], \quad (2)$$

where W is called the *weight function* and depends on s only through $\mathcal{O}$. W is to be chosen such that the probability distribution of the order parameter in the multicanonical ensemble is uniform. From the form above it is easy to see that the optimal choice is simply to take $P_w(\mathcal{O}) = 1/P_{\text{ca}}(\mathcal{O})$. The problem, of course, is that $P_{\text{ca}}(\mathcal{O})$ is the unknown function we want to sample. Luckily, a simple iterative process can yield a sufficiently good approximation for W . We detail our method in Appendix B.

Given a good guess for W , our multicanonical probability distribution is now almost flat with respect to $\mathcal{O}$. The weighting has increased the probability of the intermediate states (relative to the peaks) as to eliminate the canonical probability barrier. To recover expectation values in the canonical ensemble one notes that for any measurable A the canonical expectation value $\langle A \rangle$ can be rewritten as

$$\langle A \rangle = \frac{1}{Z} \int \mathcal{D}s A(s) \exp[-H(s)] \quad (3)$$

$$= \frac{1}{Z} \int \mathcal{D}s A(s) \exp[W(s)] \exp[-H(s) - W(s)] \quad (4)$$

$$= \frac{\langle A(s) \exp[W(s)] \rangle_{\text{muca}}}{\langle \exp[W(s)] \rangle_{\text{muca}}}, \quad (5)$$

where the multicanonical ensemble should be by construction easy to sample. That is, the canonical expectation value is recovered by weighting each measurement of the multicanonical ensemble with its respective weight. Given a well-behaved W the canonical expectation values obtained through the reweighting are always formally correct. If the choice of W is poor however, obtaining accurate results may still require arbitrarily long runs. A canonical MCMC computation is easily transformed into a multicanonical one by appending to each regular update the multicanonical accept-reject step.

It is important to think about the choice of $\mathcal{O}$ carefully, since the effect of the multicanonical method boils down the following statement: *Using $W(\mathcal{O})$ to flatten $P_{\text{ca}}(\mathcal{O})$ corresponds to equalizing the probability mass between canonical ensembles at each fixed value of the order parameter.*

If there exist probability barriers “orthogonal” to change in the order parameter (e.g., multiple metastable states for a fixed value of $\mathcal{O}$) MCMC sampling of the multicanonical distribution can still be very difficult. In the next section we will see where such barriers can render the standard application of the multicanonical method ineffective.

III. THE TUNNELING PROBLEM IN FIRST ORDER PHASE TRANSITIONS

A. Shape transitions

Shape transitions give an intuitive example of such orthogonal transitions. In the typical case the simulation volume is taken to be periodic, which leads to distinct minimal energy configurations for the phase interface. For values near the metastable peak the minimal energy configuration is given by a bubble, where a spherical clump of the stable phase is surrounded by the metastable phase. As the value of the order parameter increases, at some point it becomes preferable (for the fixed value of the order parameter) to rather form a cylindrical configuration. We have sketched these regions in Fig. 2 for a two-dimensional system.




FIG. 2. Sketch of the different phase interface shapes for a two-dimensional system. In the three-dimensional case there would in addition exist a cylindrical configuration [24]. If we were not at the critical value, one of the two peaks would be larger than the other such that the two-peaked distribution would be skewed toward the stable peak.

The transition barrier arises from the fact that for the bubble to morph into the cylinder (and vice versa) it must pass through intermediate deformed configurations (here elongated bubbles) of greater free energy cost. The probability of deformed configurations is suppressed for any value of the order parameter, and therefore the multicanonical weighting does not remove the transition barrier. These shape transitions happen (in this sense) in an orthogonal direction to change in $\mathcal{O}$. For the nucleation problem, however, they are not relevant, since they are merely undesirable artifacts of the periodic boundary conditions.

B. The condensation barrier

For the nucleation problem the relevant barrier is given by the condensation transition (see [25,26] for treatment in the Ising model). The condensation transition is located near the value of the order parameter where the preferred configuration changes from delocalized fluctuations of the metastable phase into a single localized bubble of the stable phase. For sufficiently large volumes this shows as a kink in $P(\mathcal{O})$, as sketched in Fig. 3. On both sides of the transition point there exist metastable branches for both small bubbles and large bulk fluctuations [25]. In Fig. 3 the critical bubble would be at the minimum of the bubble branch off to the right. For a poor choice of order parameter this is not necessarily the case; the critical bubble can be smaller than the bubble at the condensation transition, in which case the critical bubbles cannot be identified. In this work we presume that this is not a problem.

It is now convenient to imagine a measurable λ that could distinguish whether a configuration contains a bubble or not. This has been constructed for the two-dimensional Ising model [27], but we will keep the discussion on a heuristic level. The probability distribution of λ in the




FIG. 3. Sketch of the structure of $P(\mathcal{O})$ near the condensation transition. The dotted lines describe metastable branches of the bubble and fluctuation configurations.




FIG. 4. Sketch of $P(\mathcal{O})$ near the condensation transition with the hidden metastability in the direction of λ . Subscripts b and f denote the bubble and fluctuation branches respectively.

transition region would be bimodal, with peaks corresponding to the bubble and fluctuation branches. At the condensation transition these peaks are of equal volume, while near it one of the peaks, that corresponding to the stable branch, dominates.

To see why the barrier exists in the multicanonical ensemble, we have sketched the transition region in Fig. 4. As we recall from earlier, the effect of the multicanonical weighting is to equalize the “area” under the probability distribution $P(\lambda)$ for each value of the order parameter. As the condensation transition happens in the direction of changing λ , this multiplicative weighting does not affect the bimodal nature of $P(\lambda)$, and the barrier between the two branches remains. This is further illustrated in Fig. 5, where we sketch the probability distribution




FIG. 5. The probability distribution of Fig. 4 from above. The dashed black line indicates the location of the condensation barrier in the multicanonical ensemble. If there existed an order parameter that varied perpendicular to the dashed line, the corresponding multicanonical simulation should have no trouble smoothing out the barrier.

of Fig. 4 from above. We emphasize that in the context of field theories the present discussion on the condensation transition is a heuristic one, based on analogous results in the context of the two-dimensional Ising model [25,27,28].

The height of the barrier is essentially given by the suppression between the peaks of $P(\lambda)$ at $\mathcal{O}_{\text{cond}}$. If the physical size of the corresponding bubble at $\mathcal{O}_{\text{cond}}$ is large, the intermediate states between the branches become increasingly suppressed. Therefore for the MCMC computation it is desirable to have the condensation transition occur at the smallest possible physical bubble size. To this end, the choice of order parameter is essential. It would be preferable for the order parameter to have a small variance in the metastable phase,¹ since this results in the condensation transition occurring at smaller physical bubble size.

In this sense it is sometimes said that an order parameter with small variance can “see” smaller bubbles. Of course, for larger system volumes there exists more bulk phase and therefore more bulk fluctuations, resulting in the condensation transition shifting to ever larger physical bubble size.

One way to alleviate this problem is to find alternative order parameters of low variance with examples constructed in [17,18]. More complicated measurables may also be considered, but they can become too computationally expensive for use in the multicanonical weighting or otherwise difficult to interpret, possibly foliating the phase space in a manner that inefficiently locates the critical bubbles. In addition, it would be a tedious affair to invent such order parameters for every system of interest. We will now show how the typical order parameter can be used to construct a set of related low-variance order parameters.

C. Measuring in subvolumes

Consider what happens when we only measure the order parameter in some subvolume of the lattice. We assume the following usually attainable properties of the order parameter: it is able to resolve a large bubble on the lattice and is a volume average of some local measurable of the field values s_i . That is, it can be written as

$$\mathcal{O}(s) = \sum_i f(s_i), \quad (6)$$

where f is a function of the value s_i at site i , and possibly some of its nearby neighbors. When the order parameter takes this form, it remains meaningful when measured in a subvolume of the lattice. Such modified measurables will be known as *weighted order parameters*.

The values of a weighted order parameter only depend on the state of the system in the subvolume considered.

¹This is to be taken relative to the values the order parameter takes in the two phases. E.g., simple scaling will not have an effect as it does not change the foliation of the phase space in a nontrivial manner.

This results in some useful features. Heuristically, the probability distribution $P(\mathcal{O}^w)$ of a weighted order parameter $\mathcal{O}^w$ (in a subvolume $V^w \subset V$) should quite closely match that of $\mathcal{O}$, provided that $\mathcal{O}$ is considered in a system of volume $V' \approx V^w$ (boundary effects aside). Most importantly, the condensation transition occurs at a similar physical bubble size as in a system of volume V' . Through a choice of small V^w it is then possible to “see” smaller bubbles (within the subvolume) than when using the regular order parameter.

When a weighted order parameter is used for the multicanonical weighting, deviations from the metastable minimum are only encouraged within V^w . Furthermore, for sufficiently large values of $\mathcal{O}^w$ the system is almost certainly in the bubble branch of $P(\mathcal{O}^w)$. If $\mathcal{O}^w$ is restricted to a suitable interval, the bubbles on the bubble branch should fit inside the subvolume and be similar to those found in the regular multicanonical ensemble. As the condensation transition of the weighted ensemble occurs at small physical bubble size, the corresponding condensation barrier is also small, and the transition between the two branches of $P(\mathcal{O}^w)$ should be essentially unhindered.

Of course, this alone does not solve the initial problem of obtaining the probability of large bubbles. In the typical problematic scenario, no particular choice of subvolume V^w will be small enough to reduce the condensation barrier for the corresponding $\mathcal{O}^w$ while also being large enough to fit critical bubbles of desired size. This is where the method of expanded ensembles comes in: we can use different subvolumes in different parts of $P(\mathcal{O})$ in a manner that results in the condensation barrier being restricted only to the small and unproblematic subvolumes.

IV. METHOD OF EXPANDED ENSEMBLES

Introduced in [29], the method of expanded ensembles is most commonly found in the form of the familiar simulated tempering, where one considers additional subensembles that vary in their temperature. In general, however, the subensembles can use near arbitrary probability densities. To elucidate the usefulness of this property in the context of our multicanonical simulations, consider as a starting point the now familiar system from Sec. II described by the multicanonical probability density

$$P(s) \propto \exp[-\{H(s) + W(\mathcal{O}(s))\}]. \quad (7)$$

The problem of $P(s)$ was the fact that the multicanonical weighting, based on the full-volume order parameter, is unable to efficiently encourage the formation of bubbles.

Using the expanded ensemble method, we can instead consider N multicanonical probability distributions $\{P(s, i)\}$ with

$$P(s, i) \propto \exp[-\{H(s) + W_i(\mathcal{O}_i^w(s)) + \eta_i\}], \quad (8)$$

where $\mathcal{O}_i^w$ are different order parameters, W_i their respective weight functions, and η_i are constants. Now, in addition to the regular updates modifying the configuration s , we can also propose to change the subensemble from i to j . A symmetric proposal is easy to fashion with, say,

$$q(i, j) = \begin{cases} 1/2 & \text{for } j = i \pm 1 \\ 0 & \text{otherwise} \end{cases}. \quad (9)$$

The resulting acceptance probability for the ensemble update is then simply

$$\mathcal{P}(i \rightarrow j) = \min(1, P(s, j)/P(s, i)) \quad (10)$$

$$= \min(1, \exp[-(W_j(\mathcal{O}_j^w(s)) - W_i(\mathcal{O}_i^w(s)) - (\eta_j - \eta_i))]). \quad (11)$$

The constants η_i should be chosen such that this acceptance probability is on average symmetric.

To see the merit in this ordeal, recall the weighted order parameters from earlier. We reasoned that these are much more effective at creating small bubbles than the full-volume order parameter. The idea is then to use the weighted order parameters to construct additional multicanonical ensembles where bubbles are easily created and destroyed. The ensemble update can then feed these configurations into the original multicanonical ensemble, essentially removing the need to cross the condensation barrier. The actual process of constructing these subensembles is only slightly more involved and will be described next.

A. Choosing subensembles for the condensation barrier

The subensembles $P(s, i)$ will be multicanonical ensembles using different weighted order parameters. We present the subensembles as ordered (in the index i), since they are only designed to overlap with their “neighboring” subensembles. An illustration of the subensembles is shown in Fig. 6. The first and last subensemble use the regular full-volume order parameter, and are restricted to the metastable peak and bubble branch respectively. These are the subensembles of interest, since they sample parts of the original

probability distribution $P(\mathcal{O})$. The goal of the Monte Carlo computation is to recover the relative probability difference between these two subensembles, as it is the quantity that the standard multicanonical method is unable to determine efficiently due to the condensation barrier. The intermediate subensembles are too constrained to correspond to physical ensembles (translation invariance is broken) and only serve as a bridge connecting the physically interesting endpoints.

We choose the first weighted order parameter such that the condensation barrier for the given subvolume is small. The values of this weighted order parameter are restricted onto an interval so that the largest value corresponds to a bubble that still fits inside the subvolume. This is to prevent deformed bubbles; to obtain the overlap between the subensembles, we want the largest bubbles of the preceding subensemble to be the typical small (relative to the subvolume) bubbles of the next one. For the first weighted subensemble the lower limit can be taken as the location of the metastable peak, $\mathcal{O}_m$, though this is not particularly important.

For the next subensemble we choose the subvolume to be some constant factor larger than the preceding one. We choose the upper limit again so that the bubbles fit inside the subvolume, while the lower limit is chosen such that the small bubbles (relative to the subvolume) of this subensemble are of similar size to the largest bubbles of the preceding one. This is not a very precise statement, and it appears sufficient to estimate the lower limit using the value of the order parameter in the metastable phase. That is, we estimate the next interval from the previous one by assuming that outside of the weighted subvolume the system is in the metastable phase. This yields the estimate

$$\mathcal{O}_{n+1}^w V_{n+1}^w \approx \mathcal{O}_n^w V_n^w + \mathcal{O}_m(V_{n+1}^w - V_n^w) \quad (12)$$

and therefore to obtain an overlap of $q \in (0, 1)$ the lower limit is

$$\mathcal{O}_{n+1}^{w, \text{low}} = (q\mathcal{O}_n^{w, \text{low}} + (q-1)\mathcal{O}_n^{w, \text{high}}) \frac{V_n^w}{V_{n+1}^w}. \quad (13)$$

For example, $q = \frac{1}{2}$ should yield a lower limit that (approximately) corresponds to the midpoint of the previous subensemble with respect to the full-volume order parameter.




FIG. 6. Sketch of the subensembles of the expanded ensemble. The light areas mark the sites used in the computation of the weighted order parameter. The first and last subensembles compute the order parameter in the full volume. Values near the metastable peak are colored blue, while the red blob represents a bubble of the stable phase.

We use $q = \frac{2}{3}$ for our examples, but stress that this choice has no particular significance.

The process is repeated until the subvolumes approach the system volume. The last subensemble, now again using the full-volume order parameter, is finally attached to the end of the chain. Note that the upper limit for the weighted order parameter can usually be set to a constant for all subensembles, since it corresponds to setting the upper limit for the ratio V_b^w/V_i^w . This is the quantity that determines whether the bubble fits inside a given subvolume and is effectively independent of V_i^w . Furthermore, it is easily estimated from the histogram of a standard multicanonical simulation performed at small, manageable volume. This iterative determination of the subensembles is easily automated and requires relatively little information of the system at hand.

As a result the weighted order parameters following the first one are always restricted to their respective bubble branches. That is, the condensation barrier only exists in the first weighted subensemble where the small subvolume guarantees that the condensation barrier is weak and *independent of the full volume*. Due to the imposed limits, in the other weighted subensembles the bubbles can only shrink and grow, not evaporate. In practice, depending on, e.g., the overlap parameter q and the location of the condensation transition on the order parameter interval of interest, more than one weighted subensemble might contain the transition barrier. Choosing sufficiently small subvolumes for the first few subensembles still results in the barrier being negligible.

Only a few quantities need to be known in order to apply the method: the value of $\mathcal{O}$ in the metastable phase, $\mathcal{O}_m$, and some upper limit for $\mathcal{O}^{\max}$ which ensures that bubbles fit inside the chosen subvolumes. From these it is easy enough to construct the limits for the unweighted subensembles as well. The initial subvolume, the factor of volume increase between subvolumes, and the overlap parameter appear to require little to no tuning.

This simple construction comes with a small caveat: the subensembles should only contain bubbles well below the volume of the critical bubble. If this volume is exceeded in the weighted subensembles the expanding bubble can slip outside of the subvolume (since bubble growth is now the energetically preferable direction), and the system can end up entirely into the stable phase, leaving the Markov chain stuck. This could be avoided by adding further boundaries to the weighted ensembles that prevent the full-volume order parameter from going past the minimum.

Fortunately, this is typically not a problem since we only need to find the probability difference between the metastable peak and any part of the well-resolved bubble branch. The rest of the branch (including the critical value) is easily investigated using the standard multicanonical method. This result can then be simply glued onto the previously obtained section of the bubble branch. Since the

critical bubble is not hidden by the bulk fluctuations (in a sensible setup) the subcritical part of the bubble branch is always available.

V. A PRACTICAL EXAMPLE: THE TWO-DIMENSIONAL SCALAR FIELD

To demonstrate the method we consider a simple scalar field ϕ on a periodic two-dimensional lattice with the discretized Hamiltonian

$$H(\phi) = \sum_x \sum_{\hat{\mu}} (\phi(x + \hat{\mu}) - \phi(x))^2 + \sum_x V(\phi(x)), \quad (14)$$

where

$$V(\phi) = \lambda\phi^4 - m^2\phi^2 \quad (15)$$

and the corresponding probability density is given by

$$P(\phi) \propto \exp[-\beta H(\phi)]. \quad (16)$$

The parameters are chosen so that the expectation value of ϕ is nonzero:

$$\frac{1}{V} \left\langle \sum_x \phi(x) \right\rangle = \pm v \neq 0. \quad (17)$$

Let us consider the first order phase transition from $\langle \phi \rangle = +v$ to $-v$. As the potential is symmetric with respect to $\phi \rightarrow -\phi$, the system is effectively at the critical point, and the critical bubble is infinitely large. This is not essential for the argument for it is sufficient to obtain the probability density on the subcritical part of the bubble branch. $\langle \phi \rangle$ is a natural choice for the order parameter, and when used for multicanonical weighting the condensation barrier becomes apparent even for small volumes. We emphasize that this is most likely a nonoptimal choice, since its variance is easily reduced, e.g., by the method of [17].² Of course, these deficiencies make it optimal for demonstration purposes. Schematically, the order parameter probability distribution is similar to what we presented in Fig. 2.

We take the parameters to be $\beta = 5$, $m = 1$ and $\lambda = 1$, because this results in a condensation barrier that increases conveniently fast as a function of the system volume. This is a simple balancing act: if the barrier is too strong only a small range of system volumes are feasible using the standard multicanonical method. On the other hand, if it is too weak, very large lattices are needed in order to observe the resulting slowing down. We note that the parameters are not implied to carry any particular physical significance or relation to the continuum limits usually desired in a typical lattice field theory computation. The

²In our case this would correspond to $\mathcal{O}(s) := \frac{1}{V} \sum_x (\phi(x) - v)^2$.




FIG. 7. On the left (right) we show the weighted (full-volume) order parameter intervals and the resulting probability distributions from a run with $L = 24$. The solid lines in the upper panels and filled histograms in the lower panels correspond to the full-volume subensembles. Note that the intervals for the full-volume order parameter are estimates based on Eq. (12).

simulation program was written using the HILA lattice simulation framework [30].

We perform Monte Carlo simulations of the model using the standard multicanonical method and the method of expanded ensembles. For both methods the fields are updated by proposing Gaussian changes $\delta\phi$ on even and odd lattices sites with the standard Metropolis-Hastings algorithm and a multicanonical accept-reject step. The expanded ensemble simulation also performs an ensemble swap update and a shift update. The shift update is used to counter the breaking of translation invariance in the weighted subensembles and consists of a simple spatial shift of the field configurations. All in all, one update sweep then consists of the following steps:

- (a) Multicanonical update of the field configuration
- (b) Ensemble swap update
- (c) Shift update.

The standard method only uses step (a). Twenty-five such sweeps are performed between each measurement.

The shift proposal is generated from uniform distribution into directions $\{\hat{x}, -\hat{x}, \hat{y}, -\hat{y}\}$. The magnitude of the shift is obtained as values drawn from $U([0, L/2])$ rounded to the nearest integer. Again as $H(\phi)$ is unchanged (translation invariance), the update is accepted with the probability

$$P_{\text{shift}} = \exp[-(W_i(\mathcal{O}_i^w(\phi')) - W_i(\mathcal{O}_i^w(\phi)))], \quad (18)$$

where ϕ' is the linearly shifted ϕ . This update is trivial in the first and last subensembles using the translation invariant full-volume order parameter.

For the weighting we use spherical subvolumes with the initial weight radius as $r_{\text{init}} = 5$ and obtain successive weight radii through multiplication by $\sqrt{4/3}$. This corresponds to about a 1/3 increase in V^w on the lattice (deviations result from the fact that the lattice is discrete). It must be emphasized that these numbers carry no particular significance to the specific system at hand. Perhaps the only important requirement is that the regular multicanonical simulation with $V \approx V_{\text{init}}^w$ should face a negligible transition barrier for the expanded ensemble method to be effective. The subensemble intervals resulting from this procedure for $L = 24$ are drawn in Fig. 7, along with the histograms from the resulting run.

In Fig. 8 we show the comparison between the Markov chains and obtained probability distributions for $L = 32$. It is easy to see that the tunneling rate between the two parts of the distribution is much improved using the expanded ensemble method. Additionally, as should be, the resulting probability distributions are in good agreement where they overlap.

We also compare the probability distributions in more detail in Fig. 9, where we have computed the ratio of the probability densities on both sides of the barrier. In particular we compare the values at the metastable peak $\mathcal{O}_m$, and at $\mathcal{O}_r = -0.4$ such that the desired ratio is given by $\delta P = P(\mathcal{O}_m)/P(\mathcal{O}_r)$. All data have been normalized with respect to the values from the expanded ensemble method. The error bars are obtained from simple jackknife analysis with each chain divided into ten sections. Note that since $\mathcal{O}$ is a volume average, the same numerical value for different lattice sizes corresponds to a different sized bubble and




FIG. 8. In the topmost panel we show how the histograms from the two methods match for $L = 32$ (filled histogram is for the expanded ensemble). The two bottom panels have sections of the corresponding Markov chains. For the expanded ensemble only measurements corresponding to the first and last ensembles have been included.




FIG. 9. Probability differences and their error estimates normalized with respect to the expanded ensemble method. The results agree, but the statistical uncertainties in the multicanonical simulations grow rapidly as the volume increases.




FIG. 10. Section of the Markov chain of $\mathcal{O}$ for $L = 32$ in the expanded ensemble. The gray line includes all subensembles, and we have used red dots to highlight those that are in either the first or last subensemble.

therefore different relative probability. We see that both methods agree well with the expanded ensemble method also yielding much reduced error estimates.

To quantify the improvement in the tunneling rate we must first define the tunneling event. As the description of the probability distribution consisting of a metastable peak and a bubble branch is not a strict one for finite volume, we construct a simple heuristic definition. First we arbitrarily choose two regions based on the full-volume order parameter: one for the metastable peak and another for the bubble branch. To the metastable peak belong values with $\mathcal{O} < \mathcal{O}_m$ and to the bubble branch $\mathcal{O}_b < \mathcal{O}$.

The Markov chain is to hold a label that indicates the region it has visited last, with the change in label counted as a tunneling event. This definition works for both methods when we additionally require that the Markov chain of the expanded ensemble is in the first or last subensemble (those that sample the same probability distribution as the standard method). We have illustrated this in Fig. 10, where we display the full Markov chain highlighting the points that are in the two physically interesting subensembles. We define a normalized tunneling rate $\bar{R}_t$ as the number of tunneling events (N_t) per 10^8 MCMC update sweeps. In Fig. 11 we show the normalized tunneling rate $\bar{R}_t$ for the expanded ensemble method and the regular multicanonical method. We have multiplied both $\bar{R}_t$'s by $\propto V^{3/2}$ to factor out some of the volume scaling. We expect that in the absence of the condensation barrier the rate should be $\propto V^\alpha$. With the $V^{3/2}$ -scaling the normalized tunneling rate of the expanded ensemble method is consistent with constant while the standard multicanonical method suffers from severe exponential slowing down.

We note that $\bar{R}_t$ does not vary very smoothly as a function of L . Aside from statistical errors, there are always some variations in the iterated weight function. We have checked that the multicanonical distributions for all our




FIG. 11. Appropriately scaled normalized tunneling rate for different two-dimensional lattice sizes. The L -axis is scaled as L^2 , the system volume. In addition we have split the horizontal axis in order to fit the largest volumes. This results in the two panels having different horizontal scaling. We have picked the threshold values $\mathcal{O}_m = -0.63$ and $\mathcal{O}_b = -0.41$ to define a crossing event. The error estimates are based on ten jackknife samples.

runs are approximately flat (with the probability density between the two branches differing by at most a factor of 2). For the expanded ensemble these variations are expected since the number and limits of the intermediate ensembles depend on the full system volume. Additionally, the number of tunneling events is relatively small for the larger systems ($N_t \approx 150$ for $L = 128$).

Naïve extrapolation of the tunneling times of the standard multicanonical method suggests that our improved procedure is able to investigate system volumes that would otherwise be completely out of reach.

VI. DISCUSSION

We have presented a novel combination of the multicanonical and expanded ensemble methods for the non-perturbative determination of the bubble nucleation rate. Based on the results obtained for the toy model, the condensation barrier faced by the standard multicanonical method is eliminated. That is, with the new method the tunneling times do not increase exponentially with the system volume, but are reduced to a relatively mild scaling, consistent with $\propto V^{3/2}$. The new method is trivially generalized for the three-dimensional case and the typical order parameters used in many systems of physical interest. Preliminary testing suggests that sampling of the critical bubbles is also sped up in the cubic anisotropy model (see [24] for the standard method), although there the exponential slowing down is less apparent for the range of volumes investigated so far.

For future applications, provided that the condensation barrier can now be efficiently circumvented, the problem of

obtaining the statistical probability of very large bubbles becomes much more approachable. Such nucleation rate computations at low supercooling would allow for an interesting comparison to the typical thin-wall approximation. Unfortunately the two-dimensional toy model of this work does not yield useful results, since the lattice spacing is most likely too large for modeling the finite thickness of the phase interface.

Although of limited use in the nucleation rate computation, it would be interesting to investigate whether weighted ensembles could also be used for bypassing shape barriers discussed earlier. Through the use of appropriately deformed weighted subvolumes it may be possible to create subensembles where the otherwise unlikely deformed phase interfaces are preferred. These could then efficiently interpolate between the different geometric configurations.

It is possible to make the weighted order parameters translation invariant and do without the shift update. One way to achieve this is to maximize the weighted order parameter over all possible center points of the weighted area, that is, to maximize the convolution of the order parameter field and the weighting envelope. This can be done efficiently by Fourier transforming the envelope and the order parameter field and then multiplying them. The inverse Fourier transform of this product field then yields the desired convolution.³

ACKNOWLEDGMENTS

We acknowledge support from the ERC Grant CoCoS No. 101142449, and the Research Council of Finland Grant No. 354572.

DATA AVAILABILITY

The data that support the findings of this article are openly available [31].

APPENDIX A: IMPLEMENTATION OF THE ALGORITHM

The probability density of expanded ensemble takes the form

$$P(s, i) = \exp[-H(s, i)], \quad (\text{A1})$$

where

$$H(s, i) = H(s) + W_i(\mathcal{O}_i^w(s)) + \eta_i, \quad (\text{A2})$$

$H(s)$ being the original Hamiltonian of interest. First the weight functions W_i are iterated separately for each subensemble. After this the relative weights η_i are iterated

³We would like to thank the anonymous referee for pointing out this solution.

with the subensemble weight functions fixed. Separating the weight function from the relative weights is inconsequential but convenient as it is then easier to estimate the time taken by the weight iteration process. The first round of weight iteration tunes the subensemble histograms flat, and the relative weights aim to do the same for the distribution of the index i . The iteration method is detailed in Appendix B.

APPENDIX B: ITERATION OF THE MULTICANONICAL WEIGHT FUNCTION

As the probability distribution $P(\mathcal{O})$ is *a priori* unknown, the multicanonical weight function $W(\mathcal{O})$ must be obtained through an iterative process. Plenty of methods exist, but we have resorted to a very simple heuristic approach. Note that during the actual MCMC computation the weight function is fixed and the simulation is formally correct for any finite W . That is, W only needs to be a good enough approximation. This can be verified afterward by the observation that the resulting multicanonical histogram of the order parameter is approximately flat.

The chosen order parameter interval is divided into bins of equal width with locations of the bin centers $\{\mathcal{O}_i^b\}$ corresponding to the values of the order parameter at these points. The weight function is obtained for other values through a simple linear interpolation. That is,

$$W(\mathcal{O}) = W(\mathcal{O}_i^b) + (\mathcal{O} - \mathcal{O}_i^b) \frac{W(\mathcal{O}_{i+1}^b) - W(\mathcal{O}_i^b)}{\mathcal{O}_{i+1}^b - \mathcal{O}_i^b}, \quad (\text{B1})$$

where i is such that $\mathcal{O} \in (\mathcal{O}_i^b, \mathcal{O}_{i+1}^b)$. What happens outside of the chosen interval depends on the end goal. We want to restrict the system to the interval, and one solution is to set $W(\mathcal{O})$ to some excessive value in the forbidden region. This basically guarantees that the Markov chain never leaves the interval. An alternative option is to have the weight function increase rapidly outside of the chosen interval, e.g., by replacing the quotient in Eq. (B1) by some large positive (negative) constant, essentially constructing steep slopes on the right (left) side of the interval. This is convenient since if the chain is somehow initialized outside of the desired region, it will be driven there instead of just getting stuck.

The weight function is iterated by a pseudo Markov chain, where the weights are updated “on the fly” according to Algorithm 1.

Algorithm 1.

```

 $W(\mathcal{O}) = \text{constant};$ 
while  $C > C_{\min}$  do
  Perform MCMC updates;
  Measure order parameter value and find corresponding bin;
  Add  $C$  to the bin value (increasing  $W$ );
  Add the bin to the list of visited bins;
  if All bins have been visited then
     $C = C/1.5;$ 
    Empty the list of visited bins;
  end
end

```

Recalling that a larger value of W decreases the multicanonical probability density, it is easy to imagine why this works. The often visited regions are decreased in probability until they become unlikely enough that other regions are visited as well. C is decreased each time the interval is fully covered which makes future adjustments to the weights finer. The process is considered complete when C is decreased below $C_{\min}$, the value of which should be chosen such that W does not change appreciably over the time it takes to visit all bins. The iteration procedure is convenient in that it requires little to no adjustments and is very easy to implement for arbitrary systems. We visualize the iteration of the weight function for the two-dimensional system in Fig. 12.




FIG. 12. The evolution of the weight function W during the iteration process. We shift W such that its maximum is always at zero. This is inconsequential as only relative differences matter for the multicanonical update.

- [1] K. Kajantie, M. Laine, K. Rummukainen, and M. E. Shaposhnikov, Is there a hot electroweak phase transition at $m_H \gtrsim m_W$?, *Phys. Rev. Lett.* **77**, 2887 (1996).
- [2] M. D’Onofrio and K. Rummukainen, Standard model cross-over on the lattice, *Phys. Rev. D* **93**, 025003 (2016).
- [3] V. A. Rubakov and M. E. Shaposhnikov, Electroweak baryon number non-conservation in the early Universe and in high-energy collisions, *Phys. Usp.* **39**, 461 (1996).
- [4] C. Caprini, M. Chala, G. C. Dorsch, M. Hindmarsh, S. J. Huber, T. Konstandin, J. Kozaczuk, G. Nardini, J. M. No, K. Rummukainen, P. Schwaller, G. Servant, A. Tranberg, and D. J. Weir, Detecting gravitational waves from cosmological phase transitions with LISA: An update, *J. Cosmol. Astropart. Phys.* **03** (2020) 024.
- [5] M. Hindmarsh, M. Lüben, J. Lumma, and M. Pauly, Phase transitions in the early universe, *SciPost Phys. Lecture Notes* **024** (2021).
- [6] Amaro-Seoane *et al.*, Laser interferometer space antenna, [arXiv:1702.00786](https://arxiv.org/abs/1702.00786).
- [7] C. Caprini, M. Hindmarsh, S. Huber, T. Konstandin, J. Kozaczuk, G. Nardini, J. M. No, A. Petiteau, P. Schwaller, G. Servant, and D. J. Weir, Science with the space-based interferometer eLISA. II: Gravitational waves from cosmological phase transitions, *J. Cosmol. Astropart. Phys.* **04** (2016) 001.
- [8] J. S. Langer, Statistical theory of the decay of metastable states, *Ann. Phys. (N.Y.)* **54**, 258 (1969).
- [9] I. Affleck, Quantum-statistical metastability, *Phys. Rev. Lett.* **46**, 388 (1981).
- [10] S. Coleman, Fate of the false vacuum: Semiclassical theory, *Phys. Rev. D* **15**, 2929 (1977).
- [11] C. G. Callan and S. Coleman, Fate of the false vacuum. II. First quantum corrections, *Phys. Rev. D* **16**, 1762 (1977).
- [12] A. D. Linde, Decay of the false vacuum at finite temperature, *Nucl. Phys.* **B216**, 421 (1983).
- [13] K. Farakos, K. Kajantie, K. Rummukainen, and M. Shaposhnikov, 3D physics and the electroweak phase transition: A framework for lattice Monte Carlo analysis, *Nucl. Phys. B* **442**, 317 (1995).
- [14] G. D. Moore and K. Rummukainen, Electroweak bubble nucleation, nonperturbatively, *Phys. Rev. D* **63**, 045002 (2001).
- [15] T. Neuhaus and J. S. Hager, 2D crystal shapes, droplet condensation, and exponential slowing down in simulations of first-order phase transitions, *J. Stat. Phys.* **113**, 47 (2003).
- [16] O. Gould, S. Güyer, and K. Rummukainen, First-order electroweak phase transitions: A nonperturbative update, *Phys. Rev. D* **106**, 114507 (2022).
- [17] O. Gould, A. Kormu, and D. J. Weir, A nonperturbative test of nucleation calculations for strong phase transitions, *Phys. Rev. D* **111**, L051901 (2025).
- [18] K. Rummukainen, R. Seppä, and D. J. Weir, Resolving the critical bubble in SU(8) deconfinement transition, *Proc. Sci. LATTICE2024* (2025) 434.
- [19] G. D. Moore, Measuring the broken phase sphaleron rate nonperturbatively, *Phys. Rev. D* **59**, 014503 (1998).
- [20] B. A. Berg and T. Neuhaus, Multicanonical algorithms for first order phase transitions, *Phys. Lett. B* **267**, 249 (1991).
- [21] F. Wang and D. P. Landau, Efficient, multiple-range random walk algorithm to calculate the density of states, *Phys. Rev. Lett.* **86**, 2050 (2001).
- [22] E. Bennett, B. Lucini, D. Mason, M. Piai, E. Rinaldi, and D. Vadacchino, The density of states method for symplectic gauge theories at finite temperature, [arXiv:2409.19426](https://arxiv.org/abs/2409.19426).
- [23] M. Caselle, G. Costagliola, A. Nada, M. Panero, and A. Toniato, Jarzynski’s theorem for lattice gauge theory, *Phys. Rev. D* **94**, 034503 (2016).
- [24] G. D. Moore, K. Rummukainen, and A. Tranberg, Non-perturbative computation of the bubble nucleation rate in the cubic anisotropy model, *J. High Energy Phys.* **04** (2001) 017.
- [25] K. Binder, Theory of the evaporation/condensation transition of equilibrium droplets in finite volumes, *Physica A (Amsterdam)* **319**, 99 (2003).
- [26] K. Binder and M. H. Kalos, “Critical clusters” in a super-saturated vapor: Theory and Monte Carlo simulation, *J. Stat. Phys.* **22**, 363 (1980).
- [27] A. Nußbaumer, E. Bittner, and W. Janke, Free-energy barrier at droplet condensation, *Prog. Theor. Phys. Suppl.* **184**, 400 (2010).
- [28] A. Nußbaumer, E. Bittner, and W. Janke, Monte Carlo study of the droplet formation-dissolution transition on different two-dimensional lattices, *Phys. Rev. E* **77**, 041109 (2008).
- [29] A. P. Lyubartsev, A. A. Martsinovski, S. V. Shevkunov, and P. N. Vorontsov-Velyaminov, New approach to Monte Carlo calculation of the free energy: Method of expanded ensembles, *J. Chem. Phys.* **96**, 1776 (1992).
- [30] HILA lattice simulation framework, <https://github.com/CFT-HY/HILA>.
- [31] J. Hällfors and K. Rummukainen, Data for paper “Expanded ensemble method for bubble nucleation”, Version v1 (2025), [10.5281/zenodo.15423156](https://zenodo.org/record/15423156).