



Gravitational Waves from a Supercooled Tachyonic Phase Transition

Masterarbeit

zur Erlangung des akademischen Grades
Master of Science (M.Sc.) in Physik an der
GOETHE-UNIVERSITÄT FRANKFURT AM MAIN

vorgelegt von
Morten Alexander Will

am
August 27, 2026

Erstgutachterin: Prof. Dr. Laura Sagunski
Zweitgutachter: Dr. Daniel Schmitt

ABSTRACT

Classically scale-invariant extensions of the Standard Model can substantially modify the thermal history of the early Universe and lead to prolonged periods of supercooling. In the $U(1)_{B-L}$ model considered in this work, sufficiently strong supercooling can persist to the QCD scale, where QCD symmetry breaking destabilises the scalar field responsible for $B - L$ symmetry breaking. If bubble percolation is inefficient, the transition instead proceeds through a tachyonic instability and classical rolling of the scalar field.

This thesis investigates the nonlinear stage of this QCD-induced tachyonic phase transition and the associated production of gravitational waves using classical-statistical lattice simulations with *CosmoLattice*. The simulations follow the exponential growth of unstable scalar modes into the regime of backreaction, fragmentation, and nonlinear mode coupling. Gravitational waves are produced predominantly during the rapid onset of this nonlinear evolution, when large spatial gradients and anisotropic stresses develop.

Across the simulated range, the gravitational-wave peak amplitude follows the inverse-square scaling $\Omega_{\text{GW}}^{\text{peak}} \propto (H/m_{\varphi, \text{QCD}})^2$ predicted analytically, while the characteristic physical momentum remains of order $m_{\varphi, \text{QCD}}$. The spectra moreover exhibit an approximately universal shape around the peak, allowing a numerical template to be mapped onto the physical $U(1)_{B-L}$ parameter space. After reheating and cosmological redshifting, the resulting stochastic gravitational-wave backgrounds may be detectable by future observatories. In particular, detectable regions for BBO, DECIGO, and the Einstein Telescope (ET) already occur within the parameter range directly calibrated by the lattice simulations, while LISA sensitivity requires a modest extrapolation towards transition scales closer to the Hubble radius.

ZUSAMMENFASSUNG

Klassisch skaleninvariante Erweiterungen des Standardmodells können die thermische Geschichte des frühen Universums wesentlich verändern und zu langen Phasen starker Unterkühlung führen. In dem hier betrachteten $U(1)_{B-L}$ -Modell kann diese Unterkühlung bis zur QCD-Skala andauern, wo die QCD-Symmetriebrechung das Skalarfeld destabilisiert, das für die Brechung der $B - L$ Symmetrie verantwortlich ist. Falls die Perkolation von Vakuumblasen ineffizient ist, erfolgt der Übergang stattdessen durch eine tachyonische Instabilität und klassisches Rollen des Skalarfeldes.

Diese Arbeit untersucht die nichtlineare Phase dieses QCD-induzierten tachyonischen Phasenübergangs sowie die dabei erzeugten Gravitationswellen mithilfe klassisch-statistischer Gittersimulationen mit *CosmoLattice*. Die Simulationen verfolgen das exponentielle Wachstum instabiler Skalarmoden bis in den Bereich von Backreaction, Fragmentierung und nichtlinearer Modenkopplung. Der Großteil der Gravitationswellen wird unmittelbar nach Einsetzen dieser nichtlinearen Dynamik erzeugt, da hier starke räumliche Gradienten und anisotrope Spannungen entstehen.

Im simulierten Parameterbereich folgt die maximale Gravitationswellenamplitude der analytisch erwarteten Skalierung $\Omega_{\text{GW}}^{\text{peak}} \propto (H/m_{\varphi, \text{QCD}})^2$, während der charakteristische physikalische Impuls von der Größenordnung $m_{\varphi, \text{QCD}}$ ist. Darüber hinaus weisen die Spektren im Bereich ihres Maximums eine näherungsweise universelle Form auf, sodass ein numerisch bestimmtes Spektral-Template auf den physikalischen $U(1)_{B-L}$ -Parameterraum übertragen werden kann. Nach Reheating und kosmologischer Rotverschiebung können die resultierenden stochastischen Gravitationswellenhintergründe für zukünftige Observatorien messbar sein. Insbesondere liegen für BBO, DECIGO und das Einstein-Teleskop (ET) nachweisbare Bereiche bereits innerhalb des direkt durch die Gittersimulationen abgedeckten Parameter-raums, während für LISA eine moderate Extrapolation zu Übergangsskalen näher an der Hubble-Skala erforderlich ist.

CONTENTS

1	INTRODUCTION	1
2	SCALAR-FIELD DYNAMICS AND GRAVITATIONAL WAVES IN AN EXPANDING UNIVERSE	5
2.1	Cosmological Evolution and Relevant Scales	5
2.1.1	Homogeneous and isotropic expansion	5
2.1.2	Radiation and vacuum domination	7
2.1.3	Physical and comoving scales	8
2.2	Scalar Fields in an Expanding Universe	10
2.2.1	Homogeneous scalar-field evolution	10
2.2.2	Linear scalar-field fluctuations	12
2.2.3	From linear growth to nonlinear dynamics	15
2.3	Gravitational Waves from Scalar-Field Sources	16
2.3.1	Tensor perturbations and anisotropic stress	16
2.3.2	Gravitational-wave spectra	17
2.3.3	Cosmological redshifting	19
3	SUPERCOOLING AND TACHYONIC SYMMETRY BREAKING IN THE $U(1)_{B-L}$ MODEL	23
3.1	The Classically Scale-Invariant $U(1)_{B-L}$ Extension	23
3.1.1	Field content and scalar potential	23
3.1.2	Radiative symmetry breaking and parameter matching	25
3.2	Finite-Temperature Potential and Supercooling	27
3.2.1	Thermal stabilisation of the symmetric phase	27
3.2.2	The thermal barrier and vacuum domination	30
3.3	The QCD-Induced Exit from Supercooling	32
3.3.1	Transmission of QCD symmetry breaking through the Higgs portal	32
3.3.2	Bubble percolation and classical rolling	35
3.4	Linear Onset of the Tachyonic Transition	37
3.4.1	Homogeneous rolling and scalar fluctuations	38
3.4.2	Instability band and characteristic scales	38

3.4.3	Breakdown of the linear regime	39
3.5	Analytic Expectations for Gravitational-Wave Production	40
3.5.1	The characteristic source scale	40
3.5.2	Scaling of the peak amplitude	41
3.5.3	Reheating and redshifting	42
4	NUMERICAL IMPLEMENTATION WITH <i>CosmoLattice</i>	45
4.1	Classical Lattice Description	45
4.1.1	Discretised equations of motion	45
4.1.2	Validity of the classical-statistical approximation	46
4.2	Simulation Potential and Scale Matching	47
4.2.1	Hierarchy of scales in the physical model	47
4.2.2	Surrogate scalar potential	48
4.2.3	Matching through μ/H	49
4.3	Initial Conditions and Simulation Parameters	50
4.3.1	Homogeneous field and vacuum fluctuations	51
4.3.2	Lattice size, resolution, and time step	52
4.4	Extraction of the Gravitational-Wave Spectrum	53
4.4.1	Anisotropic stress and TT projection	54
4.4.2	Spectral energy density	55
4.5	Numerical Consistency Checks	56
4.5.1	Spatial resolution and finite volume	56
4.5.2	Initial conditions and surrogate parameters	58
4.5.3	Time step and simulation duration	61
5	NUMERICAL RESULTS AND GRAVITATIONAL-WAVE PHENOMENOLOGY	63
5.1	Nonlinear Scalar-Field Dynamics	63
5.1.1	Tachyonic growth and backreaction	63
5.1.2	Energy transfer and field fragmentation	67
5.2	Gravitational-Wave Production	70
5.2.1	Time evolution of the signal	70
5.2.2	Peak amplitude and characteristic momentum	72
5.2.3	Spectral fit	76
5.3	Mapping to the Physical Model	80
5.3.1	Dependence on the model parameters	80
5.3.2	Reheating and redshifting	82

5.4	Observational Prospects	83
5.4.1	Benchmark spectra	84
5.4.2	Detector sensitivities and signal-to-noise ratios	86
5.5	Uncertainties and Limitations	90
6	CONCLUSIONS AND OUTLOOK	93
	ACKNOWLEDGMENTS	97
	ACRONYMS	99
	BIBLIOGRAPHY	101
	USE OF AI-TOOLS	109

1 INTRODUCTION

The Standard Model of particle physics and Einstein’s theory of General Relativity provide remarkably successful descriptions of fundamental particles and gravity. Their experimental successes include the discovery of the Higgs boson [1, 2] and the first direct observation of gravitational waves [3]. Nevertheless, important questions remain unanswered. The Standard Model does not account for phenomena such as neutrino masses and dark matter, nor does it explain the observed matter–antimatter asymmetry [4, 5]. These shortcomings provide strong motivation to search for physics beyond the Standard Model (BSM).

The early Universe provides a natural laboratory for such new physics. Observations of the cosmic microwave background and primordial light-element abundances constrain the standard cosmological history back to Big Bang nucleosynthesis [6, 7], while substantially less is known about earlier and hotter epochs. Primordial gravitational waves offer a possible observational window onto this regime. Once produced, they interact only very weakly with matter and can propagate essentially freely through the subsequent cosmic history. A stochastic gravitational-wave background may therefore retain information about energetic processes and fundamental physics in the early Universe [8].

Cosmological phase transitions are among the most widely studied possible sources of such a background. As the Universe cools, changes in the equilibrium state of quantum fields can alter the symmetries realised by the thermal plasma. If a transition happens rapidly enough that the system cannot adjust smoothly to the changing conditions, the resulting inhomogeneities and anisotropic stresses can source gravitational radiation. First-order phase transitions are a familiar example, where bubbles of a new phase form and expand through the Universe. However, strongly inhomogeneous field dynamics can also arise through other mechanisms [8].

Neither the electroweak nor the quantum chromodynamics (QCD) transition is strongly first order within the Standard Model for the observed particle masses [9, 10]. Extensions of the Standard Model can, however, modify the thermal history substantially and give rise to much more violent transitions. This makes gravitational waves a particularly interesting probe of BSM physics. Future observatories will extend this search over a broad frequency range, with space-based detectors such as the Laser Interferometer Space Antenna (LISA) probing the

millihertz regime and next-generation ground-based detectors such as the Einstein Telescope (ET) probing higher frequencies [11–14].

One class of theories in which the cosmological thermal history can differ considerably from that of the Standard Model is provided by classically scale-invariant models. In these theories, explicit mass parameters are absent from the classical scalar potential and physical scales can instead arise through quantum effects such as the Coleman–Weinberg mechanism [15]. Besides providing an alternative perspective on the origin of symmetry-breaking scales, classical scale invariance can have important consequences at finite temperature. Symmetry breaking may be delayed far below the scale associated with the broken phase at zero temperature, leading to a strongly supercooled Universe.

The classically scale-invariant $U(1)_{B-L}$ extension considered in this work provides an example of such a thermal history [16]. Previous work has shown that sufficiently strong supercooling can delay the transition until the QCD epoch, where QCD dynamics can trigger an exit from the metastable phase [17]. In part of the parameter space this exit does not proceed through conventional bubble percolation but through a tachyonic instability of the scalar field. The first stage of this process can be treated using linear perturbation theory [18], but this description breaks down once the fluctuations become large. The subsequent nonlinear, sub-Hubble dynamics require a real-time lattice treatment.

This thesis studies this nonlinear stage and the gravitational waves it produces. The aim is to understand how the scalar field evolves as its fluctuations grow, interact and become highly inhomogeneous, and how this evolution generates gravitational waves. This is investigated using classical-statistical lattice simulations with *CosmoLattice* [19–21]. The resulting gravitational-wave signal is then related back to the underlying physical model and its possible observational consequences. The analysis shows that the nonlinear evolution contains a robust gravitational-wave signal whose main properties can be characterised across the relevant parameter range, providing a basis for connecting the microscopic transition dynamics with future gravitational-wave observations.

Throughout this thesis, natural units are used,

$$\hbar = c = k_B = 1. \quad (1.1)$$

The metric convention is

$$\eta_{\mu\nu} = \text{diag}(1, -1, -1, -1). \quad (1.2)$$

Greek indices $\mu, \nu, \dots \in \{0, 1, 2, 3\}$ denote spacetime components, while Latin indices $i, j, \dots \in \{1, 2, 3\}$ denote spatial components. Three-vectors are written in bold, for example $\mathbf{x}$ and $\mathbf{k}$, and repeated indices are summed unless stated otherwise. The reduced Planck

mass is defined as $M_{\text{Pl}} = (8\pi G)^{-1/2}$. A dot denotes differentiation with respect to cosmic time t , whereas a prime denotes differentiation with respect to conformal time η . Comoving wavenumbers are denoted by k , with the corresponding physical wavenumber $k_{\text{phys}} = k/a$. Where required, a subscript $\star$ denotes quantities evaluated around the epoch of gravitational-wave production, while a subscript 0 denotes present-day quantities.

The remainder of this thesis is organised as follows. Chapter 2 introduces the cosmological, scalar-field and gravitational-wave framework required for the subsequent analysis. Chapter 3 describes the supercooled thermal history of the classically scale-invariant $U(1)_{B-L}$ model and the onset of the tachyonic transition. Chapter 4 presents the lattice implementation and the numerical consistency tests. The resulting nonlinear dynamics and their gravitational-wave phenomenology are then studied in Chapter 5. Finally, Chapter 6 summarises the conclusions and discusses the main limitations and possible directions for future work.

2 SCALAR-FIELD DYNAMICS AND GRAVITATIONAL WAVES IN AN EXPANDING UNIVERSE

2.1 COSMOLOGICAL EVOLUTION AND RELEVANT SCALES

The scalar fields considered in this work evolve in an expanding cosmological background. The purpose of this section is not to review the complete thermal history of the Universe, but to introduce the background quantities and scales that enter the scalar-field and gravitational-wave equations. In particular, we will distinguish between radiation- and vacuum-dominated expansion and describe how physical wavelengths evolve relative to the Hubble scale.

For standard treatments of homogeneous cosmology and the thermal history of the early Universe, we refer to Refs. [22, 23].

2.1.1 HOMOGENEOUS AND ISOTROPIC EXPANSION

On sufficiently large scales, the Universe is well described as spatially homogeneous and isotropic. Assuming vanishing spatial curvature, the corresponding Friedmann-Lemaître-Robertson-Walker (FLRW) metric is

$$ds^2 = dt^2 - a^2(t) d\mathbf{x}^2 = a^2(\eta) (d\eta^2 - d\mathbf{x}^2), \quad (2.1)$$

where t denotes cosmic time, $\mathbf{x}$ are comoving spatial coordinates and $a(t)$ is the scale factor. In the second expression, we have introduced conformal time through

$$d\eta = \frac{dt}{a(t)}. \quad (2.2)$$

The expansion rate is characterised by the Hubble parameter

$$H(t) \equiv \frac{\dot{a}(t)}{a(t)}, \quad (2.3)$$

where a dot denotes a derivative with respect to cosmic time. Equivalently, the conformal or comoving Hubble parameter is

$$\mathcal{H}(\eta) \equiv \frac{a'(\eta)}{a(\eta)} = aH, \quad (2.4)$$

with a prime denoting a derivative with respect to conformal time. The quantity H measures the fractional change of the scale factor per unit time and therefore defines the characteristic expansion timescale H^{-1} .

Homogeneity and isotropy imply that the background energy–momentum tensor takes the perfect–fluid form

$$T^\mu{}_\nu = \text{diag}(\rho, -p, -p, -p), \quad (2.5)$$

for a comoving observer, where ρ and p are the total energy density and pressure.

For the metric in Eq. (2.1), the Einstein equations reduce to the Friedmann equations

$$H^2 = \frac{\rho}{3M_{\text{Pl}}^2}, \quad (2.6)$$

$$\frac{\ddot{a}}{a} = -\frac{\rho + 3p}{6M_{\text{Pl}}^2}, \quad (2.7)$$

where $M_{\text{Pl}} = (8\pi G)^{-1/2} \simeq 2.435 \times 10^{18} \text{ GeV}$ is the reduced Planck mass [24]. The first Friedmann equation relates the expansion rate directly to the total energy density. The second determines whether the expansion accelerates or decelerates.

Covariant conservation of the energy–momentum tensor gives the continuity equation

$$\dot{\rho} + 3H(\rho + p) = 0. \quad (2.8)$$

For a separately conserved component i , it is useful to introduce the equation-of-state parameter

$$w_i = p_i/\rho_i. \quad (2.9)$$

If w_i is approximately constant, Eq. (2.8) implies

$$\rho_i \propto a^{-3(1+w_i)}. \quad (2.10)$$

The two cases relevant for this work are radiation, for which $w_{\text{rad}} = 1/3$, and vacuum energy, for which $w_{\text{vac}} = -1$. Consequently,

$$\rho_{\text{rad}} \propto a^{-4}, \quad \rho_{\text{vac}} \simeq \text{const.} \quad (2.11)$$

These different scaling behaviours allow a constant vacuum contribution to eventually dominate over an initially larger radiation density.

2.1.2 RADIATION AND VACUUM DOMINATION

During radiation domination, the energy density is carried predominantly by a thermal bath of relativistic particles. Its energy density can be written as

$$\rho_{\text{rad}}(T) = \frac{\pi^2}{30} g_*(T) T^4, \quad (2.12)$$

where $g_*(T)$ denotes the effective number of relativistic degrees of freedom contributing to the energy density. The corresponding Hubble parameter is

$$H_{\text{rad}}(T) = \sqrt{\frac{\pi^2 g_*(T)}{90}} \frac{T^2}{M_{\text{Pl}}}. \quad (2.13)$$

Thus, in a radiation-dominated Universe, the temperature of the thermal bath provides a direct measure of the expansion rate.

The entropy density of the relativistic plasma is

$$s(T) = \frac{2\pi^2}{45} g_{*s}(T) T^3, \quad (2.14)$$

where $g_{*s}(T)$ denotes the effective number of entropic degrees of freedom. In the absence of entropy production, the entropy in a comoving volume is conserved,

$$sa^3 = \text{const.}, \quad (2.15)$$

and therefore

$$g_{*s}(T) a^3 T^3 = \text{const.} \quad (2.16)$$

If g_{*s} changes only slowly, this reduces to the familiar relation $T \propto a^{-1}$. Together with $\rho_{\text{rad}} \propto T^4$, this reproduces the scaling $\rho_{\text{rad}} \propto a^{-4}$ [22, 23, 25, 26].

A qualitatively different regime arises when an approximately constant vacuum energy density becomes dominant,

$$\rho_{\text{vac}} \gtrsim \rho_{\text{rad}}. \quad (2.17)$$

Since vacuum energy has pressure $p_{\text{vac}} \simeq -\rho_{\text{vac}}$, the expansion accelerates. The Hubble parameter is then approximately constant,

$$H \simeq H_{\text{vac}} \equiv \sqrt{\frac{\rho_{\text{vac}}}{3M_{\text{Pl}}^2}}, \quad (2.18)$$

and the scale factor evolves as

$$a(t) \simeq a_i e^{H_{\text{vac}}(t-t_i)}. \quad (2.19)$$

For exactly constant vacuum energy this is de Sitter expansion. If the vacuum energy changes slowly, the expansion is instead described as quasi-de Sitter.

The presence of a radiation bath is compatible with vacuum domination, but the radiation becomes increasingly subdominant. Provided that no entropy is injected into the bath, its temperature and energy density continue to scale as

$$T \propto a^{-1}, \quad \rho_{\text{rad}} \propto a^{-4}. \quad (2.20)$$

After $N_e = \ln(a/a_i)$ e-folds of vacuum-dominated expansion, one therefore has

$$T = T_i e^{-N_e}, \quad \rho_{\text{rad}} = \rho_{\text{rad},i} e^{-4N_e}. \quad (2.21)$$

Vacuum domination thus rapidly dilutes and cools the thermal bath while leaving the Hubble parameter nearly unchanged. This separation between the temperature scale and the expansion scale is important for strongly supercooled transitions: the temperature can become very small even though H remains fixed by the vacuum energy [17, 27].

2.1.3 PHYSICAL AND COMOVING SCALES

The spatial coordinates $\mathbf{x}$ in Eq. (2.1) are comoving coordinates. A perturbation may therefore be decomposed into Fourier modes labelled by a constant comoving wavevector $\mathbf{k}$, with $k = |\mathbf{k}|$. Cosmological expansion does not change k , but it redshifts the corresponding physical wavenumber,

$$k_{\text{phys}}(t) = \frac{k}{a(t)}. \quad (2.22)$$

The associated physical wavelength is

$$\lambda_{\text{phys}}(t) = \frac{2\pi}{k_{\text{phys}}(t)} = \frac{2\pi a(t)}{k}. \quad (2.23)$$

Physical wavelengths therefore grow in proportion to the scale factor, whereas physical wavenumbers decrease as a^{-1} .

The Hubble parameter defines both an expansion timescale H^{-1} and a characteristic physical length scale

$$R_H = H^{-1}, \quad (2.24)$$

usually referred to as the Hubble radius or Hubble horizon. Strictly speaking, the Hubble radius is not in general identical to the particle horizon, which is defined by the complete causal history of the spacetime. Nevertheless, H^{-1} is the relevant local scale for comparing the evolution of a perturbation with the cosmological expansion.

A mode is called sub-Hubble when

$$k_{\text{phys}} \gg H \quad \Longleftrightarrow \quad k \gg aH, \quad (2.25)$$

and super-Hubble when

$$k_{\text{phys}} \ll H \quad \Longleftrightarrow \quad k \ll aH. \quad (2.26)$$

The transition between the two regimes is conventionally referred to as Hubble crossing and occurs around

$$k = aH = \mathcal{H}. \quad (2.27)$$

For sub-Hubble modes, the characteristic oscillation or propagation timescale k_{phys}^{-1} is much shorter than the expansion timescale. For super-Hubble modes, the expansion of the background cannot be neglected over one mode timescale.

The evolution of the comoving Hubble radius $\mathcal{H}^{-1}$ differs between radiation and vacuum domination. During radiation domination, $H_{\text{rad}} \propto a^{-2}$, and hence

$$(\mathcal{H}_{\text{rad}})^{-1} = \frac{1}{aH_{\text{rad}}} \propto a. \quad (2.28)$$

The comoving Hubble radius grows, allowing previously super-Hubble modes to enter the Hubble scale. During vacuum domination, $H_{\text{vac}} \simeq \text{const.}$, so that

$$(\mathcal{H}_{\text{vac}})^{-1} = \frac{1}{aH_{\text{vac}}} \propto a^{-1}. \quad (2.29)$$

The comoving Hubble radius instead shrinks. Since $k_{\text{phys}} \propto a^{-1}$ while H remains approximately constant, a fixed comoving mode may evolve from the sub-Hubble to the super-Hubble regime.

These relations provide the background for the scalar-field dynamics considered in the following section. In particular, the comparison of k/a with H determines whether spatial gradients or cosmological expansion dominate the evolution of a given scalar-field mode.

2.2 SCALAR FIELDS IN AN EXPANDING UNIVERSE

Scalar fields in the early Universe evolve under the combined influence of their potential, spatial gradients and the cosmological expansion. Their dynamics may initially be described in terms of a homogeneous background and small fluctuations. Once the fluctuations become sufficiently large, however, this linear description breaks down and the full nonlinear field evolution must be followed. This section introduces the homogeneous and linearised evolution before outlining the transition to nonlinear dynamics.

The treatment of scalar fields in curved spacetime and their linearised fluctuations follows the standard discussions in Refs. [22, 23, 28]. For the nonlinear and classical-statistical regimes, we refer to Refs. [29, 30].

2.2.1 HOMOGENEOUS SCALAR-FIELD EVOLUTION

Consider a real scalar field φ that is minimally coupled to gravity. Its action is

$$S_\varphi = \int d^4x \sqrt{-g} \left[\frac{1}{2} g^{\mu\nu} \partial_\mu \varphi \partial_\nu \varphi - V(\varphi) \right], \quad (2.30)$$

where g is the determinant of the metric and $V(\varphi)$ is the scalar potential. Minimal coupling means that the field couples to the spacetime geometry only through the metric appearing in the kinetic term and the volume element.

Varying the action with respect to φ gives the covariant Klein–Gordon equation,

$$\frac{1}{\sqrt{-g}} \partial_\mu (\sqrt{-g} g^{\mu\nu} \partial_\nu \varphi) + \frac{\partial V}{\partial \varphi} = 0. \quad (2.31)$$

For the spatially flat FLRW metric introduced in Eq. (2.1), this becomes

$$\ddot{\varphi} + 3H\dot{\varphi} - \frac{1}{a^2} \nabla^2 \varphi + \frac{\partial V}{\partial \varphi} = 0. \quad (2.32)$$

The expansion of the Universe enters through the term $3H\dot{\varphi}$. This term reduces the kinetic energy of the field and is therefore commonly referred to as Hubble friction. The gradient

term is suppressed by a^{-2} , reflecting the redshifting of physical momenta during the expansion. For a homogeneous field,

$$\varphi(t, \mathbf{x}) = \bar{\varphi}(t), \quad (2.33)$$

the gradient term vanishes and the field equation reduces to

$$\ddot{\bar{\varphi}} + 3H\dot{\bar{\varphi}} + \frac{dV(\bar{\varphi})}{d\bar{\varphi}} = 0. \quad (2.34)$$

This equation resembles the motion of a particle in the potential $V(\bar{\varphi})$, subject to a time-dependent friction force. The field is accelerated by the slope of the potential, while the cosmological expansion opposes its motion.

The energy–momentum tensor obtained from Eq. (2.30) is

$$T_{\mu\nu}^{(\varphi)} = \partial_\mu \varphi \partial_\nu \varphi - g_{\mu\nu} \left[\frac{1}{2} g^{\alpha\beta} \partial_\alpha \varphi \partial_\beta \varphi - V(\varphi) \right]. \quad (2.35)$$

For the homogeneous field, its energy density and pressure are

$$\rho_{\bar{\varphi}} = \frac{1}{2} \dot{\bar{\varphi}}^2 + V(\bar{\varphi}), \quad (2.36)$$

$$p_{\bar{\varphi}} = \frac{1}{2} \dot{\bar{\varphi}}^2 - V(\bar{\varphi}). \quad (2.37)$$

With these, the scalar equation of motion is equivalent to the continuity equation

$$\dot{\rho}_{\bar{\varphi}} + 3H(\rho_{\bar{\varphi}} + p_{\bar{\varphi}}) = 0. \quad (2.38)$$

Thus, an expanding background continuously removes kinetic energy from the homogeneous field.

The scalar dynamics and the cosmological expansion are coupled through the Friedmann equations introduced in Sec. 2.1. In the presence of additional background components with energy density ρ_{other} , one has

$$H^2 = \frac{1}{3M_{\text{Pl}}^2} [\rho_{\text{other}} + \rho_{\bar{\varphi}}]. \quad (2.39)$$

The scalar evolution therefore depends on H , while H itself depends on the scalar energy density. The field dynamics and the cosmological expansion are generally coupled and must be determined self-consistently.

If the scalar energy density is dominated by its potential energy,

$$\dot{\bar{\varphi}}^2 \ll V(\bar{\varphi}), \quad (2.40)$$

then

$$p\bar{\varphi} \simeq -\rho\bar{\varphi}. \quad (2.41)$$

An approximately constant scalar potential therefore behaves as vacuum energy and produces accelerated expansion if it dominates the total energy density.

Note that in the applications considered later, radiative and thermal corrections are incorporated by replacing the classical potential with an effective potential,

$$V(\bar{\varphi}) \longrightarrow V_{\text{eff}}(\bar{\varphi}, T). \quad (2.42)$$

The temperature dependence then changes the force acting on the field as the thermal bath evolves.

2.2.2 LINEAR SCALAR-FIELD FLUCTUATIONS

A perfectly homogeneous field is an idealisation. Quantum and thermal fluctuations provide small spatially varying deviations from the homogeneous configuration. We therefore decompose the field as

$$\varphi(t, \mathbf{x}) = \bar{\varphi}(t) + \delta\varphi(t, \mathbf{x}), \quad \langle \delta\varphi(t, \mathbf{x}) \rangle = 0, \quad (2.43)$$

where $\bar{\varphi}(t)$ denotes the spatially homogeneous expectation value and $\delta\varphi$ contains the inhomogeneous fluctuations.

Expanding the potential force about the homogeneous field gives

$$\left. \frac{\partial V}{\partial \varphi} \right|_{\varphi=\bar{\varphi}+\delta\varphi} = \frac{dV(\bar{\varphi})}{d\bar{\varphi}} + \frac{d^2V(\bar{\varphi})}{d\bar{\varphi}^2} \delta\varphi + \mathcal{O}(\delta\varphi^2). \quad (2.44)$$

Keeping only terms up to first order in $\delta\varphi$ reproduces Eq. (2.34) for the background and gives

$$\delta\ddot{\varphi} + 3H\delta\dot{\varphi} - \frac{1}{a^2}\nabla^2\delta\varphi + m_{\text{eff}}^2(t)\delta\varphi = 0, \quad (2.45)$$

where

$$m_{\text{eff}}^2(t) \equiv \left. \frac{d^2V(\bar{\varphi})}{d\bar{\varphi}^2} \right|_{\bar{\varphi}=\bar{\varphi}(t)} \quad (2.46)$$

is the time-dependent curvature of the potential along the background trajectory. For a thermal effective potential, the replacement

$$m_{\text{eff}}^2(t) = \left. \frac{\partial^2 V_{\text{eff}}(\bar{\varphi}, T)}{\partial \bar{\varphi}^2} \right|_{\bar{\varphi}=\bar{\varphi}(t), T=T(t)} \quad (2.47)$$

is understood.

At linear order, different Fourier modes evolve independently. For a quantum field, the fluctuation operator may be expanded as

$$\delta \hat{\varphi}(t, \mathbf{x}) = \int \frac{d^3 \mathbf{k}}{(2\pi)^3} \left[\hat{a}_{\mathbf{k}} u_{\mathbf{k}}(t) e^{i\mathbf{k} \cdot \mathbf{x}} + \hat{a}_{\mathbf{k}}^\dagger u_{\mathbf{k}}^*(t) e^{-i\mathbf{k} \cdot \mathbf{x}} \right]. \quad (2.48)$$

The functions $u_{\mathbf{k}}(t)$ are the mode functions and contain the time evolution of fluctuations with comoving wavenumber $k = |\mathbf{k}|$. The creation and annihilation operators satisfy

$$[\hat{a}_{\mathbf{k}}, \hat{a}_{\mathbf{q}}^\dagger] = (2\pi)^3 \delta^{(3)}(\mathbf{k} - \mathbf{q}). \quad (2.49)$$

Substitution into Eq. (2.45) gives one equation for each comoving wavenumber,

$$\ddot{u}_{\mathbf{k}} + 3H\dot{u}_{\mathbf{k}} + \omega_{\mathbf{k}}^2(t)u_{\mathbf{k}} = 0, \quad (2.50)$$

with effective frequency

$$\omega_{\mathbf{k}}^2(t) = \frac{k^2}{a^2(t)} + m_{\text{eff}}^2(t). \quad (2.51)$$

When $\omega_{\mathbf{k}}^2 > 0$ and the frequency is sufficiently large compared with the expansion rate, the mode behaves as a damped oscillator. By contrast, if $\omega_{\mathbf{k}}^2 < 0$, the mode equation instead admits exponentially growing and decaying solutions for modes in the range

$$0 \leq \frac{k}{a(t)} \leq \sqrt{-m_{\text{eff}}^2(t)}, \quad (2.52)$$

signalling an instability. This instability band is illustrated in Fig. 2.1.

The initial mode amplitudes depend on the state of the scalar field. Even in its lowest-energy state, a quantum field possesses vacuum fluctuations. If the field is instead in thermal equilibrium, its modes are additionally populated by particles from the thermal bath. The average number of particles in a bosonic mode is described by the occupation number

$$n_k^{\text{th}} = \frac{1}{\exp(\omega_k/T) - 1}, \quad (2.53)$$




Figure 2.1: Schematic illustration of the relation between the curvature of the effective potential and the stability of scalar-field fluctuations. (a) As the effective mass squared m_{eff}^2 , changes from positive to negative, the origin changes from a stable local minimum to an unstable local maximum. (b) For $m_{\text{eff}}^2 < 0$, the fluctuation modes obey the dispersion relation in Eq. (2.51). Modes satisfying $k_{\text{phys}}^2 < -m_{\text{eff}}^2$ therefore have $\omega_k^2 < 0$ and are exponentially amplified, whereas modes with $k_{\text{phys}}^2 > -m_{\text{eff}}^2$ remain oscillatory. The boundary of the instability band is given by $k_{\text{phys}} = \sqrt{-m_{\text{eff}}^2}$.

evaluated in an initial regime where the mode frequency is positive and well-defined.

The occupation number provides a useful measure of how strongly a mode is populated. For $n_k \ll 1$, quantum fluctuations are important and the particle interpretation is intrinsically quantum. For $n_k \gg 1$, many quanta occupy the same mode. The relative importance of quantum commutators then becomes small, and the mode behaves approximately as a classical wave. A highly occupied quantum field may consequently be represented by an ensemble of classical field configurations with appropriately chosen random initial conditions.

The occupation number is most clearly defined when the background changes slowly and the modes behave approximately as particles with a well-defined energy. During a rapidly varying or unstable stage, this interpretation becomes less precise. The field correlation functions and mode amplitudes nevertheless remain well-defined and continue to describe the growth of the fluctuations.

Throughout this discussion, scalar perturbations of the metric have been neglected. This is appropriate while gravitational feedback on the scalar modes remains small and their dynamics are dominated by the potential, gradients and the background expansion. Tensor perturbations sourced by the scalar field will be introduced separately in Sec. 2.3.

2.2.3 FROM LINEAR GROWTH TO NONLINEAR DYNAMICS

The linear decomposition is valid only while the fluctuations remain small enough that terms of quadratic and higher order in $\delta\varphi$ can be neglected. Beyond linear order, the expansion of the potential force reads

$$\left. \frac{\partial V}{\partial \varphi} \right|_{\varphi=\bar{\varphi}+\delta\varphi} = \frac{dV}{d\bar{\varphi}} + \frac{d^2V}{d\bar{\varphi}^2} \delta\varphi + \frac{1}{2} \frac{d^3V}{d\bar{\varphi}^3} \delta\varphi^2 + \frac{1}{6} \frac{d^4V}{d\bar{\varphi}^4} \delta\varphi^3 + \dots \quad (2.54)$$

The linear approximation fails once the omitted terms become comparable to the linear contribution. This may occur when the root-mean-square fluctuation

$$\delta\varphi_{\text{rms}} = \sqrt{\langle \delta\varphi^2 \rangle}, \quad (2.55)$$

approaches the characteristic field-space scale over which the curvature of the potential changes. Equivalently, the approximation must break down once the energy stored in the fluctuations becomes comparable to the energy of the homogeneous mode.

The first important nonlinear effect is backreaction. The growth of the fluctuations modifies the force acting on the homogeneous field. Their energy density also contributes to the Friedmann equation and can alter the expansion rate. This means that the background and fluctuation equations can no longer be solved independently.

Nonlinear interactions also couple Fourier modes with different wavenumbers. Modes amplified during the linear stage scatter from one another and redistribute energy across momentum space. This process is referred to as rescattering. It can populate modes that were not directly amplified during the initial instability. This means that the independent mode evolution of Eq. (2.50) is no longer valid.

As rescattering proceeds, an increasing fraction of the scalar energy may be stored in spatial gradients. The initially homogeneous field then develops spatial structure and may break up into inhomogeneous configurations. This loss of coherence of the homogeneous mode is commonly called fragmentation.

Backreaction, rescattering and fragmentation are intrinsically nonlinear and cannot be captured by evolving a collection of independent linear modes. Their study requires the full real-time field equation in Eq. (2.32) to be solved without expanding in $\delta\varphi$. In practice, this is achieved using classical-statistical lattice simulations, which discretise the field in space and evolve its nonlinear dynamics in real time. Such simulations are particularly well motivated when the relevant modes have developed large occupation numbers and therefore behave approximately as classical fields [8].

2.3 GRAVITATIONAL WAVES FROM SCALAR-FIELD SOURCES

Spatially inhomogeneous scalar fields can generate tensor perturbations of the spacetime metric. The relevant source is not the scalar energy density itself, but its anisotropic spatial stress. This establishes a direct connection with the nonlinear dynamics discussed in the preceding section: a homogeneous scalar field does not source gravitational waves, whereas fragmentation and the formation of spatial gradients generally do.

For general reviews of cosmological gravitational-wave production and propagation, we refer to Refs. [8, 31]. The treatment of gravitational waves sourced directly by nonlinear scalar fields follows Refs. [32, 33].

2.3.1 TENSOR PERTURBATIONS AND ANISOTROPIC STRESS

Gravitational waves are described by transverse and traceless tensor perturbations of the spatial metric. For a spatially flat FLRW background, the line element may be written as

$$ds^2 = dt^2 - a^2(t) [\delta_{ij} + h_{ij}(t, \mathbf{x})] dx^i dx^j, \quad (2.56)$$

where the tensor perturbation satisfies

$$\partial_i h_{ij} = 0, \quad h_{ii} = 0. \quad (2.57)$$

These conditions leave two independent degrees of freedom, corresponding to the two gravitational-wave polarisations.

Expanding the Einstein equations to linear order in h_{ij} gives

$$\ddot{h}_{ij} + 3H\dot{h}_{ij} - \frac{1}{a^2} \nabla^2 h_{ij} = \frac{2}{M_{\text{Pl}}^2} \Pi_{ij}^{\text{TT}}. \quad (2.58)$$

The second term describes Hubble friction, while the gradient term governs the propagation of the gravitational waves. The source Π_{ij}^{TT} is the transverse-traceless (TT) part of the spatial energy–momentum tensor. To obtain Π_{ij}^{TT} , one removes the isotropic trace, corresponding to the pressure contribution, from the spatial stress tensor and projects the remaining tensor onto the part that is transverse to the wavevector.

A homogeneous field $\Phi(t, \mathbf{x}) = \bar{\varphi}(t)$ has no spatial gradients and consequently $\Pi_{ij}^{\text{TT}} = 0$. Its homogeneous motion may dominate the energy density and affect the cosmological expansion, but it cannot directly generate gravitational waves. A source arises only once the scalar develops anisotropic spatial structure. Processes such as fragmentation, rescattering and the growth of

gradient energy therefore provide the physical mechanisms that allow for gravitational-wave production from nonlinear scalar-field dynamics.

2.3.2 GRAVITATIONAL-WAVE SPECTRA

Once produced, gravitational waves carry energy. For wavelengths well inside the Hubble radius, their energy density may be obtained by averaging over several oscillations,

$$\rho_{\text{GW}} = \frac{M_{\text{Pl}}^2}{4} \langle \dot{h}_{ij} \dot{h}_{ij} \rangle. \quad (2.59)$$

The brackets denote a spatial average over scales larger than the gravitational-wave wavelength but smaller than the characteristic scale of the background evolution.

It is convenient to express the energy density per logarithmic interval in wavenumber as

$$\Omega_{\text{GW}}(k, t) \equiv \frac{1}{\rho_{\text{tot}}(t)} \frac{d\rho_{\text{GW}}(t)}{d \ln k}, \quad (2.60)$$

where

$$\rho_{\text{tot}} = 3M_{\text{Pl}}^2 H^2, \quad (2.61)$$

is the total energy density of the Universe. The total gravitational-wave abundance is obtained by integrating the spectrum,

$$\Omega_{\text{GW}}(t) = \int d \ln k \Omega_{\text{GW}}(k, t). \quad (2.62)$$

Following the convention used throughout this work, the physical frequency associated with a comoving wavenumber k is

$$f(t) = \frac{k}{2\pi a(t)}. \quad (2.63)$$

Since $d \ln f = d \ln k$, the spectrum can equivalently be expressed as a function of frequency. A simple estimate illustrates which sources produce sizeable gravitational-wave signals. Consider a source concentrated around a characteristic comoving wavenumber k , corresponding to the physical scale

$$k_{\text{phys}} = \frac{k}{a}. \quad (2.64)$$

Suppose that the source is sub-Hubble, $k_{\text{phys}} \gtrsim H$, and varies on a characteristic timescale $\Delta t \sim k_{\text{phys}}^{-1}$. The scale factor and Hubble parameter then change only weakly during one period of the source.

A tensor mode varying on this timescale has

$$\dot{h}_{ij} \sim k_{\text{phys}} h_{ij}, \quad \ddot{h}_{ij} \sim k_{\text{phys}}^2 h_{ij}, \quad \frac{1}{a^2} \nabla^2 h_{ij} \sim k_{\text{phys}}^2 h_{ij}, \quad (2.65)$$

because once emitted, a sub-Hubble gravitational wave oscillates approximately as $h_{ij} \propto e^{ik_{\text{phys}}t}$, and the differential operator on the left-hand side of Eq. (2.58) is characterised by the scale k_{phys}^2 . The Hubble-friction term is smaller by a factor of order H/k_{phys} and may be neglected in this estimate.

Consequently, a rough estimate is

$$\dot{h}_{ij} \sim \frac{\Pi_{ij}}{M_{\text{Pl}}^2 k_{\text{phys}}}, \quad (2.66)$$

up to numerical factors of order unity. The resulting gravitational-wave energy density scales as

$$\rho_{\text{GW}} \sim M_{\text{Pl}}^2 \dot{h}_{ij} \dot{h}_{ij} \sim \frac{\Pi_{ij}^2}{M_{\text{Pl}}^2 k_{\text{phys}}^2}. \quad (2.67)$$

Dividing by the total energy density and using $\rho_{\text{tot}} = 3M_{\text{Pl}}^2 H^2$ yields

$$\Omega_{\text{GW}} \sim \left(\frac{\Pi_{ij}}{\rho_{\text{tot}}} \right)^2 \left(\frac{H}{k_{\text{phys}}} \right)^2. \quad (2.68)$$

To separate the total source energy from its anisotropic component, we schematically write

$$\Pi_{ij} \sim \epsilon_{\text{an}} \rho_{\text{source}}, \quad (2.69)$$

where ρ_{source} is the energy density carried by the source and ϵ_{an} parametrises the fraction that contributes to anisotropic stress. The estimate then becomes

$$\Omega_{\text{GW}} \sim \epsilon_{\text{an}}^2 \left(\frac{\rho_{\text{source}}}{\rho_{\text{tot}}} \right)^2 \left(\frac{H}{k_{\text{phys}}} \right)^2. \quad (2.70)$$

A strong signal therefore requires a source that carries a substantial fraction of the total energy density, develops sizeable anisotropic stress and varies on scales that are not far below the Hubble radius. A large homogeneous scalar energy density alone is insufficient because it corresponds to $\epsilon_{\text{an}} = 0$.

2.3.3 COSMOLOGICAL REDSHIFTING

After the source has become inactive, gravitational waves satisfy the homogeneous form of Eq. (2.58). Sub-Hubble modes then propagate approximately freely. Their physical wavenumber and amplitude scale as

$$\frac{k}{a} \propto a^{-1}, \quad h_{ij} \propto a^{-1}, \quad (2.71)$$

and their energy density consequently redshifts as radiation,

$$\rho_{\text{GW}} \propto a^{-4}. \quad (2.72)$$

Let the subscript $\star$ denote the end of gravitational-wave production and the subscript 0 the present day. The present-day physical frequency is

$$f_0 = \frac{k_\star}{2\pi a_0} = \frac{k_\star}{2\pi a_\star} \frac{a_\star}{a_0}. \quad (2.73)$$

The gravitational-wave energy density evolves according to

$$\rho_{\text{GW},0} = \rho_{\text{GW},\star} \left(\frac{a_\star}{a_0} \right)^4. \quad (2.74)$$

Using $\rho_{\text{GW},\star} = \Omega_{\text{GW},\star} 3M_{\text{Pl}}^2 H_\star^2$ and $\rho_{\text{crit},0} = 3M_{\text{Pl}}^2 H_0^2$, this gives

$$\Omega_{\text{GW},0} = \Omega_{\text{GW},\star} \left(\frac{H_\star}{H_0} \right)^2 \left(\frac{a_\star}{a_0} \right)^4. \quad (2.75)$$

The scale-factor ratio can be related to the temperature of the radiation bath through entropy conservation. Using Eq. (2.16), one obtains

$$\frac{a_\star}{a_0} = \frac{T_0}{T_\star} \left[\frac{g_{*,s}(T_0)}{g_{*,s}(T_\star)} \right]^{1/3}, \quad (2.76)$$

where $g_{*,s}$ counts the effective relativistic degrees of freedom contributing to the entropy density. Substitution into Eq. (2.73) gives the general frequency-redshift relation

$$f_0 = \frac{k_\star}{2\pi a_\star} \frac{T_0}{T_\star} \left[\frac{g_{*,s}(T_0)}{g_{*,s}(T_\star)} \right]^{1/3}. \quad (2.77)$$

Similarly, the present-day gravitational-wave abundance is

$$\Omega_{\text{GW},0} = \Omega_{\text{GW},\star} \left[\frac{g_{*,s}(T_0)}{g_{*,s}(T_\star)} \right]^{4/3} \left(\frac{T_0}{T_\star} \right)^4 \left(\frac{H_\star}{H_0} \right)^2. \quad (2.78)$$

The standard numerical relations follow if the Universe is radiation dominated at the end of gravitational-wave production. In this case,

$$H_\star = \sqrt{\frac{\pi^2 g_*(T_\star)}{90}} \frac{T_\star^2}{M_{\text{Pl}}}, \quad (2.79)$$

where $g_*(T)$ counts the relativistic degrees of freedom contributing to the energy density.

For temperatures well above the masses of the relevant particle species, one may approximate

$$g_*(T_\star) \simeq g_{*,s}(T_\star) \equiv g_* \quad (2.80)$$

Using Eq. (2.79) in Eq. (2.77), and writing the physical production scale relative to the Hubble scale, gives

$$\begin{aligned} f_0 &= \frac{k_\star}{2\pi a_\star H_\star} H_\star \frac{T_0}{T_\star} \left[\frac{g_{*,s}(T_0)}{g_*} \right]^{1/3} \\ &= \frac{k_\star}{2\pi a_\star H_\star} \sqrt{\frac{\pi^2 g_*}{90}} \frac{T_\star T_0}{M_{\text{Pl}}} \left[\frac{g_{*,s}(T_0)}{g_*} \right]^{1/3} \\ &= \frac{100^{1/6}}{2\sqrt{90}} \frac{T_0}{M_{\text{Pl}}} g_{*,s}(T_0)^{1/3} \frac{k_\star}{a_\star H_\star} \frac{T_\star}{\text{GeV}} \left(\frac{g_*}{100} \right)^{1/6} \text{GeV}. \end{aligned} \quad (2.81)$$

Inserting [26, 34]

$$T_0 = 2.7255 \text{ K} \simeq 2.35 \times 10^{-13} \text{ GeV}, \quad g_{*,s}(T_0) \simeq 3.91, \quad (2.82)$$

together with the reduced Planck mass and the conversion factor

$$1 \text{ GeV} = 1.519 \times 10^{24} \text{ Hz}, \quad (2.83)$$

yields [8, 35]

$$f_0 = 2.62 \times 10^{-8} \text{ Hz} \frac{k_\star}{a_\star H_\star} \frac{T_\star}{\text{GeV}} \left(\frac{g_*}{100} \right)^{1/6}. \quad (2.84)$$

For the amplitude, it is useful to relate the gravitational-wave density to the radiation density today. Radiation domination at production gives

$$\rho_{\text{tot},\star} = \rho_{\text{rad},\star} = \frac{\pi^2}{30} g_{\star} T_{\star}^4. \quad (2.85)$$

Combining this relation with entropy conservation gives

$$\Omega_{\text{GW},0} = \Omega_{\text{rad},0} \frac{g_{\star}}{g_{\star}(T_0)} \left[\frac{g_{\star,s}(T_0)}{g_{\star}} \right]^{4/3} \Omega_{\text{GW},\star}. \quad (2.86)$$

Using $g_{\star}(T_0) \simeq 3.36$ [26] and the present radiation abundance $h^2 \Omega_{\text{rad},0} \approx 4.15 \times 10^{-5}$ [26, 34] then yields [8, 35]

$$h^2 \Omega_{\text{GW},0} = 1.67 \times 10^{-5} \left(\frac{100}{g_{\star}} \right)^{1/3} \Omega_{\text{GW},\star}, \quad (2.87)$$

where $H_0 = 100h \text{ km s}^{-1} \text{ Mpc}^{-1}$, such that $h \approx 0.67$ [6].

Equations (2.84) and (2.87) assume that production is followed by radiation domination and adiabatic expansion. If the source operates during vacuum domination or reheating is delayed, the spectrum must first be evolved through this intermediate period [8, 17].

3 SUPERCOOLING AND TACHYONIC SYMMETRY BREAKING IN THE $U(1)_{B-L}$ MODEL

3.1 THE CLASSICALLY SCALE-INVARIANT $U(1)_{B-L}$ EXTENSION

The model considered in this chapter extends the Standard Model gauge group according to

$$SU(3)_c \times SU(2)_L \times U(1)_Y \rightarrow SU(3)_c \times SU(2)_L \times U(1)_Y \times U(1)_{B-L}, \quad (3.1)$$

where $B - L$ denotes baryon number minus lepton number. Quarks carry $B - L$ charge $1/3$, while leptons carry charge -1 . Gauging this symmetry introduces an additional gauge boson, denoted by Z' .

The defining assumption of the model is classical scale invariance. That is, that the classical Lagrangian contains no dimensionful parameters. This distinguishes the model from an ordinary polynomial symmetry-breaking theory with a potential of the form

$$V(\varphi)_{\text{poly}} = -\frac{1}{2}\mu_\varphi^2\varphi^2 + \frac{1}{4}\lambda_\varphi\varphi^4, \quad (3.2)$$

where the mass parameter μ_φ already fixes the symmetry-breaking scale at tree level. In the classically scale-invariant theory, no such scale is present initially. It must instead emerge through quantum corrections.

The field content and symmetry-breaking structure described in this section follow the original classically scale-invariant $U(1)_{B-L}$ construction of Refs. [16, 36]. The parameter-matching prescription and renormalisation-group (RG) conventions are adopted from Ref. [17].

3.1.1 FIELD CONTENT AND SCALAR POTENTIAL

In addition to the Standard Model fields, the theory contains a complex scalar field Φ , three right-handed neutrinos N_R^i , with $i \in \{1, 2, 3\}$, and the gauge boson Z' associated with $U(1)_{B-L}$. The new fields' transformation behaviour is shown in Table 3.1.

3 Supercooling and Tachyonic Symmetry Breaking in the $U(1)_{B-L}$ Model

Field	$SU(3)_c$	$SU(2)_L$	$U(1)_Y$	$U(1)_{B-L}$
N_R^i	1	1	0	-1
Φ	1	1	0	+2

Table 3.1: Quantum numbers of the new fields under the gauge group $SU(3)_c \times SU(2)_L \times U(1)_Y \times U(1)_{B-L}$. Table adapted from Ref. [16].

The three right-handed neutrinos are required to cancel the gauge and gravitational anomalies associated with $U(1)_{B-L}$. They also permit the generation of light neutrino masses through the type-I seesaw mechanism [37]. In the following, their Yukawa couplings are assumed to be sufficiently small that they do not significantly affect the scalar potential or the symmetry-breaking dynamics.

The scalar fields may be parametrised as

$$H = \begin{pmatrix} G^+ \\ \frac{1}{\sqrt{2}}(h + iG^0) \end{pmatrix}, \quad \Phi = \frac{1}{\sqrt{2}}(\varphi + iG), \quad (3.3)$$

where H is the Standard Model Higgs doublet and φ denotes the radial component, responsible for breaking the new gauge symmetry. The field G becomes the longitudinal degree of freedom of the massive Z' boson.

Classical scale invariance excludes quadratic terms for both H and Φ . The most general renormalisable scalar potential is consequently

$$V_{\text{tree}}(\Phi, H) = \lambda_h (H^\dagger H)^2 + \lambda_\varphi (\Phi^\dagger \Phi)^2 - \lambda_p (H^\dagger H)(\Phi^\dagger \Phi), \quad (3.4)$$

where λ_h and λ_φ are the scalar self-couplings and $\lambda_p > 0$ is the Higgs-portal coupling. Along the two radial directions, the potential becomes

$$V_{\text{tree}}(\varphi, h) = \frac{\lambda_\varphi}{4}\varphi^4 + \frac{\lambda_h}{4}h^4 - \frac{\lambda_p}{4}\varphi^2 h^2. \quad (3.5)$$

The portal term is essential. It replaces the negative Higgs mass term that is inserted by hand in the Standard Model. As long as $\varphi = 0$, the portal interaction does not generate a quadratic term for h . Once φ develops a vacuum expectation value, however, the same interaction becomes a negative Higgs mass term and triggers electroweak symmetry breaking.

3.1.2 RADIATIVE SYMMETRY BREAKING AND PARAMETER MATCHING

At tree level, the potential in Eq. (3.5) contains only quartic terms and does not single out a preferred energy scale. The $B-L$ scale is generated radiatively through the Coleman–Weinberg mechanism [15]. Quantum fluctuations introduce a renormalisation scale into the effective potential, allowing a non-zero vacuum expectation value to emerge through dimensional transmutation.

Along the $B-L$ direction, where $h = 0$, the tree-level potential is

$$V_{\text{tree}}(\varphi) = \frac{\lambda_\varphi}{4} \varphi^4. \quad (3.6)$$

The dominant one-loop correction is generated by the Z' gauge boson, whose field-dependent mass is

$$m_{Z'}(\varphi) = 2g_{B-L}\varphi, \quad (3.7)$$

where g_{B-L} is the new gauge coupling. Neglecting the smaller scalar and right-handed-neutrino contributions, the Coleman–Weinberg correction is

$$V_{\text{CW}}(\varphi) = \frac{3}{64\pi^2} m_{Z'}^4(\varphi) \left[\ln \left(\frac{m_{Z'}^2(\varphi)}{\bar{\mu}^2} \right) - \frac{5}{6} \right]. \quad (3.8)$$

The full zero-temperature potential along this direction is therefore

$$V_{\text{eff}}(\varphi) = V_{\text{tree}}(\varphi) + V_{\text{CW}}(\varphi). \quad (3.9)$$

Although the tree-level potential alone has no preferred non-zero field value, the logarithm in Eq. (3.8) changes its shape and generates a minimum at $\varphi = v_\varphi$, thereby breaking both classical scale invariance and the $U(1)_{B-L}$ gauge symmetry. The Z' gauge boson consequently acquires the physical mass

$$m_{Z'} = 2g_{B-L}v_\varphi. \quad (3.10)$$

Choosing the renormalisation scale to coincide with the physical symmetry-breaking scale, $\bar{\mu} = m_{Z'}$, and requiring the effective potential to be stationary at $\varphi = v_\varphi$ gives

$$\lambda_\varphi(\bar{\mu} = m_{Z'}) = \frac{g_{B-L}^4}{\pi^2}. \quad (3.11)$$

This relation illustrates how Coleman–Weinberg symmetry breaking works. The tree-level quartic coupling is of the same order as the one-loop gauge correction. The radiative contribution can therefore compete with the classical potential and generate a non-trivial minimum. It

3 Supercooling and Tachyonic Symmetry Breaking in the $U(1)_{B-L}$ Model

also justifies neglecting further scalar-loop corrections: since $\lambda_\varphi \sim g_{B-L}^4$, such contributions scale as $\lambda_\varphi^2 \sim g_{B-L}^8$ and are subleading compared with the one-loop gauge contribution.

The mass of the radial scalar around the true minimum is

$$m_\varphi^2 = \left. \frac{\partial^2 V_{\text{eff}}}{\partial \varphi^2} \right|_{\varphi=v_\varphi} = \frac{6}{\pi^2} g_{B-L}^4 v_\varphi^2 = \frac{3}{2\pi^2} g_{B-L}^2 m_{Z'}^2. \quad (3.12)$$

For a small gauge coupling g_{B-L} , the radial scalar is therefore lighter than the Z' boson.

Once φ settles at v_φ , the portal part of the potential becomes

$$V_{\text{tree}}(\varphi = v_\varphi, h) \supset -\frac{\lambda_p}{4} v_\varphi^2 h^2. \quad (3.13)$$

This has precisely the form of a negative Higgs mass term. Comparing it with the Standard Model Higgs potential,

$$V_{\text{SM}}(h) = -\frac{1}{2} \mu_{h,\text{SM}}^2 h^2 + \frac{\lambda_h}{4} h^4, \quad (3.14)$$

gives the matching condition at leading order

$$\lambda_p(\bar{\mu} = m_Z) \simeq 2\lambda_h \left(\frac{v_h}{v_\varphi} \right)^2 \simeq 2 \left(\frac{\mu_{h,\text{SM}}}{v_\varphi} \right)^2. \quad (3.15)$$

Note that we take λ_h to be its Standard Model value. This is justified by the hierarchy $v_\varphi \gg v_h$, which suppresses mixing between the two scalar states, so that mixing-induced corrections to the Higgs-sector parameters are small.

These relations allow the model to be parametrised by the two independent quantities

$$\{g_{B-L}, m_{Z'}\}. \quad (3.16)$$

The corresponding vacuum expectation value follows from Eq. (3.10), while λ_φ is fixed by the Coleman–Weinberg condition (3.11). The couplings are then evolved to the electroweak scale, where the measured Higgs vacuum expectation value fixes the portal coupling through Eq. (3.15). Note that the kinetic-mixing coupling $\tilde{g}$ between the $U(1)_Y$ and $U(1)_{B-L}$ gauge fields is set to zero at the electroweak scale $\bar{\mu} = m_Z$, because although it appears in the RG evolution equations responsible for running the couplings, it has little influence on the scalar dynamics considered here.

Since collider constraints generally require $v_\varphi \gg v_h$, the portal coupling is small [38, 39]. The two scalar directions are therefore only weakly mixed, and their symmetry breaking can be understood as an approximately two-stage process,

$$(\varphi, h) = (0, 0) \longrightarrow (v_\varphi, 0) \longrightarrow (v_\varphi, v_h). \quad (3.17)$$

The second step is controlled by the portal-induced Higgs mass and ordinarily occurs rapidly once $B - L$ symmetry has been broken. Delaying the $U(1)_{B-L}$ transition therefore also delays electroweak symmetry breaking, apart from the additional QCD-induced effects discussed in Sec. 3.3. This connection is central to the modified thermal history developed in the following sections.

3.2 FINITE-TEMPERATURE POTENTIAL AND SUPERCOOLING

The vacuum potential introduced in Sec. 3.1 determines the final symmetry-breaking pattern, but the transition towards this vacuum takes place in the hot plasma of the early Universe. Interactions with the thermal bath modify the effective potential, particularly near the origin of field space. In the classically scale-invariant model, these corrections can keep the $B - L$ symmetry restored long after the non-zero vacuum has become energetically preferred.

In this section, we introduce the finite-temperature potential only to the extent needed to understand this thermal history. In particular, the full potential is required to determine the existence of the barrier, the duration of supercooling and the onset of vacuum domination. The later analysis of the tachyonic transition focuses on a regime of extreme supercooling in which the QCD-induced contribution dominates the curvature around the origin. Thermal terms can then be neglected, and the scalar dynamics are described by a simpler RG-improved potential.

The finite-temperature field-theory conventions used in this section follow Ref. [25]. The model-specific effective potential and supercooled thermal history of the classically scale-invariant $U(1)_{B-L}$ model follow Ref. [17]. See also Refs. [40–42] for earlier analyses of strongly supercooled transitions in this model. The discussion of vacuum domination and thermal inflation follows Refs. [43, 44].

3.2.1 THERMAL STABILISATION OF THE SYMMETRIC PHASE

At finite temperature, the effective potential along the $B - L$ direction is approximated by

$$V_{\text{eff}}(\varphi, T) = V_{\text{tree}}(\varphi) + V_{\text{CW}}(\varphi, T) + V_T(\varphi, T) + V_{\text{daisy}}(\varphi, T). \quad (3.18)$$

3 Supercooling and Tachyonic Symmetry Breaking in the $U(1)_{B-L}$ Model

As in the zero-temperature analysis, we retain the dominant contribution from the Z' gauge boson. Contributions from the scalar self-interaction are suppressed by the Coleman–Weinberg relation $\lambda_\varphi \sim g_{B-L}^4$, while the right-handed-neutrino Yukawa couplings are assumed to be small. The tree-level and Coleman–Weinberg contributions were introduced in Sec. 3.1. We therefore concentrate here on the two temperature-dependent terms.

The thermal contribution can be written as

$$V_T(\varphi, T) = \frac{3T^4}{2\pi^2} J_B \left(\frac{m_{Z'}^2(\varphi)}{T^2} \right), \quad (3.19)$$

where J_B is the bosonic thermal function. It accounts for the free-energy contribution of the Z' excitations in the thermal bath. Its explicit integral representation is not needed here.

The infrared dynamics relevant for the phase transition are dominated by the bosonic zero-frequency mode, while the non-zero Matsubara modes are heavy at these scales. Near the transition, the effective mass of the zero mode becomes small, enhancing infrared fluctuations and eventually spoiling the perturbative expansion. Resumming the leading contributions from the hard thermal modes generates a thermal mass for the zero mode, screening these infrared fluctuations and extending the regime of perturbative validity. This gives the daisy contribution

$$V_{\text{daisy}}(\varphi, T) = -\frac{T}{12\pi} \left[\left(m_{Z'}^2(\varphi) + \Pi_{Z'}(T) \right)^{3/2} - m_{Z'}^3(\varphi) \right], \quad (3.20)$$

where

$$\Pi_{Z'}(T) = 4g_{B-L}^2 T^2, \quad (3.21)$$

is the thermal mass of the longitudinal Z' mode [45, 46]. Only this longitudinal degree of freedom is resummed, since the transverse gauge modes do not acquire a perturbative thermal mass.

Equations (3.18)–(3.21) define the one-loop, daisy-resummed potential used in the following. The couplings entering this potential are evaluated at the field- and temperature-dependent scale

$$\bar{\mu}(\varphi, T) = \max\{m_{Z'}(\varphi), \pi T, \Lambda_{\text{QCD}}\}. \quad (3.22)$$

This choice accounts for the two distinct regimes encountered during strong supercooling. Near the origin, the relevant scale is set by the thermal bath, whereas near the true vacuum it is set by the large field-dependent Z' mass. The running is frozen below the characteristic

confinement scale of QCD Λ_{QCD} , where the perturbative RG evolution of the Standard Model couplings is no longer reliable. For the RG evolution, it is taken to be

$$\Lambda_{\text{QCD}} = 100 \text{ MeV}. \quad (3.23)$$

The effect of the thermal corrections is most transparent from the curvature of the potential at the origin. Expanding the full potential for small φ gives

$$\left. \frac{\partial^2 V_{\text{eff}}(\varphi, T)}{\partial \varphi^2} \right|_{\varphi=0} \simeq g_{B-L}^2 T^2. \quad (3.24)$$

Equivalently, the potential near the origin contains the positive contribution

$$V_{\text{eff}}(\varphi, T) \supset \frac{1}{2} g_{B-L}^2 T^2 \varphi^2. \quad (3.25)$$

Interactions with the hot plasma therefore give φ a positive temperature-dependent mass. At sufficiently high temperature, this term dominates the potential and pins the field to $\varphi = 0$. The $U(1)_{B-L}$ symmetry is then thermally restored.

As the Universe cools, the thermal mass decreases. The radiatively generated minimum at non-zero φ consequently reappears, initially as a local minimum. The critical temperature T_c is defined by the degeneracy condition

$$V_{\text{eff}}(0, T_c) = V_{\text{eff}}(v_\varphi(T_c), T_c), \quad (3.26)$$

where $v_\varphi(T)$ denotes the non-zero minimum at temperature T . For $T < T_c$, the broken phase has lower free energy and is therefore the true equilibrium state.

Nevertheless, the field does not immediately move to the new minimum. The origin remains locally stable because its curvature is still positive. Consequently, the symmetric minimum and the broken minimum are separated by a maximum of the potential: a thermal barrier. The qualitative evolution of $V_{\text{eff}}(\varphi, T)$ is illustrated in Fig. 3.1.

This behaviour differs from that of a conventional polynomial potential, such as Eq. (3.2). In that case, the negative tree-level mass competes directly with the positive thermal mass. Once the temperature falls far enough, the curvature at the origin changes sign and the symmetric phase becomes classically unstable. Here, there is no negative tree-level mass parameter. The curvature in Eq. (3.24) remains positive at every finite temperature and only vanishes asymptotically as $T \rightarrow 0$. The barrier can therefore survive at temperatures many orders of magnitude below v_φ .

3 Supercooling and Tachyonic Symmetry Breaking in the $U(1)_{B-L}$ Model

As long as the thermal barrier persists and no additional contribution destabilises the origin, the field can leave the symmetric phase only through bubble nucleation rather than by classical rolling. If the tunnelling rate is small, the Universe remains in the symmetric phase long below T_c . This is the origin of the unusually strong supercooling found in classically scale-invariant models.




Figure 3.1: Shifted effective potential $\Delta V_{\text{eff}}(\varphi, T) \equiv V_{\text{eff}}(\varphi, T) - V_{\text{eff}}(0, T)$ for various temperatures. At high temperature, the origin is the only minimum. A second minimum subsequently develops and becomes degenerate with the origin at $T_c \simeq 313$ GeV. At lower temperatures, the broken minimum is energetically preferred, although the field may remain trapped behind the shrinking thermal barrier. A fiducial $(g_{B-L}, m_{Z'}) = (0.1, 100 \text{ GeV})$ has been chosen here.

3.2.2 THE THERMAL BARRIER AND VACUUM DOMINATION

Below the critical temperature, the symmetric phase is no longer the global minimum. It can nevertheless remain long-lived if the tunnelling rate is small compared with the cosmic expansion rate. The Universe then stays in the false vacuum at $\varphi = 0$, even though the state at $\varphi = v_\varphi(T)$ has lower free energy.

The energy stored in the false vacuum relative to the true vacuum is

$$\Delta V_{\text{eff}}(T) = V_{\text{eff}}(0, T) - V_{\text{eff}}(v_\varphi(T), T). \quad (3.27)$$

Below T_c , this difference is positive. As long as the field remains trapped at the origin, this energy difference contributes approximately as a cosmological constant.

By contrast, the energy density of the thermal bath is

$$\rho_{\text{rad}}(T) = \frac{\pi^2}{30} g_*(T) T^4, \quad (3.28)$$

where g_ρ denotes the effective number of relativistic degrees of freedom contributing to the energy density. During the cosmological expansion, the radiation density decreases approximately as a^{-4} , whereas the false-vacuum energy remains nearly constant. Vacuum domination therefore begins at the temperature T_i at which the two contributions become equal,

$$\Delta V_{\text{eff}}(T_i) = \rho_{\text{rad}}(T_i) = \frac{\pi^2}{30} g_*(T_i) T_i^4. \quad (3.29)$$

For $T < T_i$, the Friedmann equation becomes

$$H^2 = \frac{\rho_{\text{rad}}(T) + \Delta V_{\text{eff}}(T)}{3M_{\text{Pl}}^2} \simeq \frac{\Delta V_{\text{eff}}(T)}{3M_{\text{Pl}}^2}. \quad (3.30)$$

The temperature dependence of the vacuum-energy difference is weak in the strongly supercooled regime. The Hubble parameter therefore becomes approximately constant. The scale factor then grows approximately exponentially, while the remaining radiation continues to redshift. This intermediate period of accelerated expansion is known as thermal inflation. In contrast to primordial slow-roll inflation, it is not driven by a slowly evolving scalar field. The field is instead held almost fixed in a metastable thermal state, and the accelerated expansion is driven by its false-vacuum energy.

For a strongly supercooled transition, a typical hierarchy is

$$T_c \gg T_i \gg T_{\text{exit}}, \quad T_{\text{exit}} = \begin{cases} T_p, & \text{bubble-driven exit,} \\ T_{\text{roll}}, & \text{rolling-driven exit,} \end{cases} \quad (3.31)$$

where T_p denotes the percolation temperature. The vacuum-energy difference at T_i is then well approximated by its zero-temperature value,

$$\Delta V_{\text{eff}}(T_i) \simeq \Delta V_{\text{eff}}(0) = \frac{3}{128\pi^2} m_{Z'}^4. \quad (3.32)$$

The onset of thermal inflation can then be estimated as

$$T_i \simeq \left(\frac{90}{128\pi^4 g_*(T_i)} \right)^{1/4} m_{Z'}, \quad (3.33)$$

while the Hubble parameter during vacuum domination is approximately

$$H \simeq \frac{1}{\sqrt{128}\pi} \frac{m_{Z'}^2}{M_{\text{Pl}}}. \quad (3.34)$$

These expressions make the relevant scales transparent. In the numerical analysis, however, T_i and H are obtained from the full temperature-dependent effective potential.

The plasma temperature continues to fall approximately as $T \propto a^{-1}$ during thermal inflation. Because the barrier prevents the field from following the changing equilibrium minimum, the Universe may cool far below both T_c and the characteristic $B-L$ scale without completing the transition. This is the supercooled state relevant for the remainder of this chapter. Its eventual exit depends on whether QCD-enhanced bubble nucleation reaches percolation before the thermal barrier disappears completely. This is discussed in the next section.

3.3 THE QCD-INDUCED EXIT FROM SUPERCOOLING

If the $B-L$ transition remains incomplete down to the QCD epoch, the thermal history changes qualitatively. Electroweak symmetry is still unbroken, so the Standard Model quarks have not acquired their usual Higgs-generated masses. The formation of the QCD quark condensate can then induce a small Higgs expectation value, which in turn modifies the $B-L$ potential through the portal interaction. In this way, QCD chiral symmetry breaking provides an exit from the supercooled phase.

The QCD-induced transmission of symmetry breaking through the Higgs portal follows Refs. [17, 47, 48]. The model-specific distinction between a bubble-driven and rolling-driven exit from supercooling follows Ref. [17].

3.3.1 TRANSMISSION OF QCD SYMMETRY BREAKING THROUGH THE HIGGS PORTAL

As long as the $B-L$ scalar remains trapped at $\varphi = 0$, the portal interaction does not generate a negative Higgs mass parameter. The Higgs field therefore also remains close to the origin $(\varphi, h) = (0, 0)$ and electroweak symmetry is still restored when the plasma reaches the QCD epoch. In particular, the Higgs-generated quark masses vanish. The QCD sector is conse-

quently close to the chiral limit with six light quark flavours, rather than the theory with the physical quark masses encountered in the standard thermal history.

We denote the temperature at which the quark condensate forms by T_{QCD} . This should be distinguished from Λ_{QCD} , introduced in Sec. 3.2.1 as the lower cut-off for the perturbative RG evolution. Both scales are of order 100 MeV, but they play different roles. In the present scenario, where all six quark flavours are effectively massless, the QCD transition is expected to occur at a somewhat lower temperature than for physical quark masses, with $T_{\text{QCD}} \sim \mathcal{O}(100 \text{ MeV})$. Its precise value is not important for the dynamics considered here [17, 49].

Below this temperature, QCD chiral symmetry is broken by the formation of quark condensates,

$$\langle \bar{q}_i q_i \rangle \neq 0. \quad (3.35)$$

Through the Yukawa interactions, these condensates generate a term linear in the Higgs field [50, 51]. At $\varphi = 0$, the relevant Higgs potential is

$$V_h(h) = \frac{\lambda_h}{4} h^4 + \sum_i \frac{y_i}{\sqrt{2}} \langle \bar{q}_i q_i \rangle h. \quad (3.36)$$

Since the condensates are negative, the linear term displaces the Higgs from the origin. Minimising the potential gives a finite vacuum expectation value

$$v_{h,\text{QCD}} = \left[-\frac{1}{\sqrt{2}\lambda_h} \sum_i y_i \langle \bar{q}_i q_i \rangle \right]^{1/3} \simeq \left[-\frac{y_t}{\sqrt{2}\lambda_h} \langle \bar{t} t \rangle \right]^{1/3}. \quad (3.37)$$

The top-quark contribution dominates because y_t is the largest Yukawa coupling. The precise value of $v_{h,\text{QCD}}$ depends on strongly coupled QCD dynamics and cannot be calculated reliably using perturbation theory [17, 52]. Following Refs. [17, 47, 48], we therefore adopt the representative value

$$v_{h,\text{QCD}} = 100 \text{ MeV}, \quad (3.38)$$

unless stated otherwise.

This QCD-induced Higgs expectation value is much smaller than the present electroweak value $v_h \simeq 246 \text{ GeV}$. Nevertheless, it is sufficient to affect the extremely shallow $B - L$ potential. Substituting $h = v_{h,\text{QCD}}$ into the portal interaction term of the tree-level potential in Eq. (3.5) gives

$$V_{\text{QCD}}(\varphi) \equiv -\frac{\lambda_p}{4} v_{h,\text{QCD}}^2 \varphi^2 = -\frac{1}{2} m_{\varphi,\text{QCD}}^2 \varphi^2, \quad (3.39)$$

3 Supercooling and Tachyonic Symmetry Breaking in the $U(1)_{B-L}$ Model

where

$$m_{\varphi, \text{QCD}}^2 \equiv \frac{\lambda_p}{2} v_{h, \text{QCD}}^2 \quad (3.40)$$

denotes the positive magnitude of the QCD-induced mass parameter. The corresponding contribution to the effective mass squared is negative.

The physical effect of QCD symmetry breaking is most clearly seen from the curvature of the effective potential at the origin. Approximating the formation of the condensate as occurring instantaneously at T_{QCD} , one obtains

$$m_{\varphi, \text{eff}}^2(T) \equiv \left. \frac{\partial^2 V_{\text{eff}}(\varphi, T)}{\partial \varphi^2} \right|_{\varphi=0} = \begin{cases} g_{B-L}^2 T^2, & T > T_{\text{QCD}}, \\ g_{B-L}^2 T^2 - m_{\varphi, \text{QCD}}^2, & T < T_{\text{QCD}}. \end{cases} \quad (3.41)$$

This equation summarises the entire mechanism. At early times, the positive thermal mass dominates and stabilises the symmetric phase. After QCD chiral symmetry breaking, the portal transmits a negative, temperature-independent contribution to the $B - L$ scalar. As the Universe continues to cool, the thermal term decreases as T^2 , while the QCD-induced term remains approximately constant.

Immediately below T_{QCD} , the curvature may still be positive. In that case, QCD symmetry breaking lowers the thermal barrier but does not remove it at once. Continued supercooling eventually brings the two terms in Eq. (3.41) into balance. This defines the rolling temperature,

$$T_{\text{roll}} = \frac{m_{\varphi, \text{QCD}}}{g_{B-L}} = \sqrt{\frac{\lambda_p}{2}} \frac{v_{h, \text{QCD}}}{g_{B-L}}. \quad (3.42)$$

Using the portal matching relation in Eq. (3.15) and $m_{Z'} = 2g_{B-L}v_\varphi$, this may be written as

$$T_{\text{roll}} = 2\sqrt{\lambda_h} v_h \frac{v_{h, \text{QCD}}}{m_{Z'}} \simeq 180 \text{ GeV} \frac{v_{h, \text{QCD}}}{m_{Z'}}. \quad (3.43)$$

For the large Z' masses relevant to the lattice study, T_{roll} can lie far below T_{QCD} . The Universe then continues to inflate after the QCD condensate has formed. At $T = T_{\text{roll}}$, the curvature at the origin vanishes and the thermal barrier disappears. Below this temperature,

$$m_{\varphi, \text{eff}}^2(T) < 0, \quad (3.44)$$

so the symmetric configuration is no longer a local minimum.

The thermal terms are indispensable for describing the trapping and for determining T_{roll} . Once the system enters the regime where $m_{\varphi, \text{QCD}}^2 \gg g_{B-L}^2 T^2$, however, they become sub-

leading. The later tachyonic dynamics can therefore be studied using the RG-improved zero-temperature potential supplemented by Eq. (3.39), with the scale choice $\bar{\mu} = m_{Z'}(\varphi)$.

3.3.2 BUBBLE PERCOLATION AND CLASSICAL ROLLING

The QCD-induced contribution lowers the thermal barrier as the temperature approaches T_{roll} . The supercooled phase can therefore end in two different ways. If tunnelling becomes efficient while the barrier is still present, bubbles of the broken phase nucleate, expand and eventually percolate. If this does not happen before the barrier disappears, the scalar is released from the origin and the transition proceeds through classical rolling.

To distinguish these possibilities, one must account not only for the instantaneous bubble nucleation rate, but also for the subsequent growth of the bubbles in the expanding background. The thermal nucleation rate is obtained from the three-dimensional bounce action,

$$\Gamma(T) \simeq T^4 \left(\frac{S_3(T)}{2\pi T} \right)^{3/2} \exp \left[-\frac{S_3(T)}{T} \right]. \quad (3.45)$$

Here, $S_3(T)$ is evaluated numerically from the finite-temperature effective potential [53, 54]. Its exponential appearance makes the nucleation rate highly sensitive to the thermal barrier. After QCD chiral symmetry breaking, the barrier rapidly shrinks and $S_3(T)/T$ decreases until it vanishes at T_{roll} .

The probability for a point to remain in the false vacuum is written as

$$P(T) = \exp[-I(T)], \quad (3.46)$$

where $I(T)$ is the integrated nucleation history, including both the formation and expansion of the bubbles. Percolation is defined by

$$I(T_p) = 0.34, \quad (3.47)$$

with T_p denoting the percolation temperature [55, 56].

Since the thermal barrier ceases to exist at T_{roll} , the two possible exits are classified by evaluating the percolation integral at this temperature:

$$\begin{aligned} I(T_{\text{roll}}) \geq 0.34 &\implies \text{bubble-driven exit at } T_p \geq T_{\text{roll}}, \\ I(T_{\text{roll}}) < 0.34 &\implies \text{exit through classical rolling.} \end{aligned} \quad (3.48)$$

In the first case, the QCD-induced contribution accelerates an otherwise strongly suppressed first-order transition. A percolation temperature $T_p \geq T_{\text{roll}}$ exists, and expanding bubbles convert the Universe to the broken phase while the origin is still locally stable.

As the tunnelling suppression becomes stronger, however, efficient bubble nucleation begins increasingly close to T_{roll} . The interval between the onset of appreciable nucleation and the disappearance of the barrier then becomes too short for the bubbles to occupy a sufficiently large fraction of space. The barrier disappears before the percolation threshold is reached, and the field is released from the origin throughout the remaining false-vacuum regions.

Ref. [17] evaluated this criterion across the $(g_{B-L}, m_{Z'})$ parameter space and found a sizeable region in which

$$I(T_{\text{roll}}) < 0.34, \quad (3.49)$$

so that bubble percolation cannot complete the transition before the thermal barrier vanishes. An adapted version of their parameter-space result is shown schematically in Fig. 3.2. The bubble-driven region lies generally towards larger gauge couplings, while stronger supercooling at smaller couplings favours the rolling-driven exit. Note that the precise location of this boundary is subject to theoretical uncertainties in the calculation of the nucleation rate. In particular, higher-order thermal corrections and fluctuation effects can become relevant for strongly supercooled transitions and modify the predicted percolation temperature [57, 58]. Including these effects would therefore shift the boundary between the bubble-driven and rolling-driven regions, without changing the qualitative distinction between the two possible exit mechanisms considered here.

A further condition is required to ensure that the Universe can return to a radiation-dominated state after the transition. Once the field reaches the broken minimum, the amplified fluctuations may be interpreted as excitations of the radial scalar with mass m_φ , given in Eq. (3.12). Through the Higgs portal, these excitations can transfer their energy to the Standard Model via

$$\varphi \longrightarrow hh. \quad (3.50)$$

This decay is kinematically allowed only when [42, 59]

$$m_\varphi > 2m_h. \quad (3.51)$$

The present analysis assumes that reheating is dominated by this channel. Parameter points satisfying $m_\varphi < 2m_h$ are therefore not considered, since the transfer of the scalar energy density to the Standard Model cannot be established within this approximation. This region is also highlighted in Fig. 3.2.




Figure 3.2: Schematic parameter-space classification in the $(m_{Z'}, g_{B-L})$ plane distinguishing the region where bubble percolation completes the transition from the region satisfying $I(T_{\text{roll}}) < 0.34$, where the exit from supercooling is driven by classical rolling. The region $m_\varphi < 2m_h$ is not included because the decay $\varphi \rightarrow hh$, assumed to reheat the Standard Model sector, is kinematically forbidden. Adapted from Ref. [17].

The rolling region in Fig. 3.2 is the part of the model parameter space relevant for the lattice study. There, the system does not fragment through bubble nucleation. Instead, the scalar remains approximately homogeneous until the origin becomes unstable and then begins to evolve through a region of negative curvature. The subsequent nonlinear dynamics rapidly release the false-vacuum energy and reheat the Universe. In the prompt-reheating approximation adopted here, the reheating temperature is therefore taken to be $T_{\text{rh}} \simeq T_i$.

We neglect a possible intermediate matter-dominated period following the transition. The post-transition equation of state can depend non-trivially on the scalar-field dynamics and its subsequent thermalisation [60]. Determining whether such a period arises in the present tachyonic transition requires dedicated simulations and is beyond the scope of this work.

3.4 LINEAR ONSET OF THE TACHYONIC TRANSITION

In the rolling region identified in Sec. 3.3.2, bubble percolation does not complete the transition before the thermal barrier disappears. The $B - L$ scalar therefore remains approximately homogeneous until T_{roll} , when the curvature at the origin changes sign. The generic scalar fluctuation equation introduced in Chapter 2 can now be specialised to the QCD-modified effective potential of the model.

The model-specific onset of the tachyonic transition follows Ref. [17]. The general properties of tachyonic amplification and the subsequent fragmentation of the scalar field are reviewed in Refs. [61, 62]. See also Ref. [63] for a modern analysis of the linear regime.

3.4.1 HOMOGENEOUS ROLLING AND SCALAR FLUCTUATIONS

Once the QCD-induced mass term dominates over the remaining thermal corrections, the effective potential is approximated by

$$V_{\text{eff}}(\varphi) = \frac{\lambda_\varphi(\bar{\mu})}{4}\varphi^4 - \frac{\lambda_p(\bar{\mu})}{4}v_{h,\text{QCD}}^2\varphi^2 + V_{\text{CW}}(\varphi, \bar{\mu}), \quad \bar{\mu} = m_{Z'}(\varphi). \quad (3.52)$$

Here and below, the running couplings are understood to be frozen for $\bar{\mu} < \Lambda_{\text{QCD}}$, as prescribed in Eq. (3.22). Replacing the classical potential in Eq. (2.34) according to Eq. (2.42), the homogeneous field obeys

$$\ddot{\varphi} + 3H\dot{\varphi} + \frac{\partial}{\partial\varphi} \left[\frac{\lambda_\varphi(\bar{\mu})}{4}\varphi^4 - \frac{\lambda_p(\bar{\mu})}{4}v_{h,\text{QCD}}^2\varphi^2 + V_{\text{CW}}(\varphi, \bar{\mu}) \right] = 0. \quad (3.53)$$

The negative quadratic term first releases the field from the origin, while the radiatively generated part of the potential determines its subsequent motion towards the broken minimum. Using the decomposition in Eq. (2.43), the general mode equation (2.50) is controlled by the curvature derived in Eq. (3.41). Near the origin,

$$\ddot{u}_k + 3H\dot{u}_k + \left[\frac{k^2}{a^2} + g_{B-L}^2 T^2 - m_{\varphi,\text{QCD}}^2 \right] u_k \simeq 0. \quad (3.54)$$

The instability begins when the positive thermal contribution is overtaken by the QCD-induced mass. At sufficiently low temperatures below T_{roll} , once $g_{B-L}^2 T^2 \ll m_{\varphi,\text{QCD}}^2$, this reduces to

$$\ddot{u}_k + 3H\dot{u}_k + \left[\frac{k^2}{a^2} - m_{\varphi,\text{QCD}}^2 \right] u_k \simeq 0. \quad (3.55)$$

3.4.2 INSTABILITY BAND AND CHARACTERISTIC SCALES

From Eqs. (2.51) and (3.55), the effective frequency becomes negative for

$$0 \leq \frac{k}{a} < m_{\varphi,\text{QCD}}. \quad (3.56)$$

This is the model-specific form of the instability condition introduced in Sec. 2.2 and illustrated in Fig. 2.1.

Neglecting the expansion during the initial growth, the unstable modes grow with the approximate rate

$$|\omega_k| \simeq \sqrt{m_{\varphi, \text{QCD}}^2 - \frac{k^2}{a^2}}. \quad (3.57)$$

Long-wavelength modes are therefore amplified most efficiently. Towards the edge of the band, the gradient contribution nearly compensates the negative curvature and the growth rate approaches zero. The resulting spectrum is thus expected to peak somewhat below the upper boundary $k/a = m_{\varphi, \text{QCD}}$.

The QCD-induced mass fixes the characteristic scales at the onset of the transition,

$$t_\star \sim R_\star \sim m_{\varphi, \text{QCD}}^{-1}. \quad (3.58)$$

Consequently, the ratio $m_{\varphi, \text{QCD}}/H$ compares both the growth time and the fluctuation length scale with the Hubble scale. For $m_{\varphi, \text{QCD}} \gg H$, the instability develops rapidly and the relevant structures are well inside the horizon. For $m_{\varphi, \text{QCD}} \sim H$, the characteristic fluctuation scale approaches the Hubble radius.

The perturbative evolution of super-Hubble fluctuations during thermal inflation was studied in Ref. [18]. The present work instead focuses on the subsequent sub-Hubble evolution, for which spatial gradients and nonlinear interactions become essential.

3.4.3 BREAKDOWN OF THE LINEAR REGIME

The independent mode evolution in Eq. (3.55) is valid only while the fluctuations remain small. Their exponential amplification eventually makes the higher-order terms in Eq. (2.54) important. The fluctuations then backreact on the rolling background, different momentum modes interact through rescattering, and the initially homogeneous field fragments into an inhomogeneous configuration.

As discussed in Chapter 2, this stage cannot be captured by a collection of independent linear modes. The full field equation (2.32) must instead be evolved without expanding around a homogeneous background. The resulting nonlinear dynamics, including the growth of gradient energy and the anisotropic stresses sourcing gravitational waves, are studied using *CosmoLattice* in the following chapter [19, 20].

3.5 ANALYTIC EXPECTATIONS FOR GRAVITATIONAL-WAVE PRODUCTION

Before turning to the lattice simulations, it is useful to estimate the characteristic scale and amplitude of the gravitational-wave signal. The aim is not to predict the numerical prefactor or the detailed spectral shape, but to identify the model parameter that should control the result.

The model-specific estimates in this section follow Ref. [17]. The general scaling of gravitational-wave production from nonlinear scalar-field dynamics is discussed in Refs. [32, 64]. For a broader review of cosmological gravitational-wave sources, see Ref. [8].

3.5.1 THE CHARACTERISTIC SOURCE SCALE

The initial tachyonic band derived in Eq. (3.56) is bounded by $k/a = m_{\varphi, \text{QCD}}$. Modes close to this boundary grow more slowly, so the scalar and gravitational-wave spectra need not peak exactly at the edge of the band. Nevertheless, the QCD-induced mass sets the only relevant microphysical scale at the onset of the transition. The characteristic physical wavenumber is therefore, to order of magnitude,

$$\frac{k_{\star}}{a_{\star}} \sim m_{\varphi, \text{QCD}}, \quad (3.59)$$

up to a numerical factor to be determined from the simulations.

The associated source size $R_{\star}$ and duration $t_{\star}$ are both of order $m_{\varphi, \text{QCD}}^{-1}$. Their relation to the Hubble radius is therefore controlled by

$$\frac{m_{\varphi, \text{QCD}}}{H}. \quad (3.60)$$

Large values correspond to short-lived, deeply sub-Hubble dynamics, whereas $m_{\varphi, \text{QCD}}/H$ approaching unity places the characteristic structures close to the cosmological horizon. This ratio is consequently the central dimensionless parameter used to organise the lattice simulations.

At leading order, the parameter matching of Sec. 3.1 gives

$$\frac{m_{\varphi, \text{QCD}}}{H} \simeq \sqrt{512} \pi g_{B-L} \frac{M_{\text{Pl}} \mu_{h, \text{SM}} v_{h, \text{QCD}}}{m_{Z'}^3}, \quad (3.61)$$

where the running of the couplings has been neglected. The numerical mapping to the physical model will instead use the RG-improved parameters.

3.5.2 SCALING OF THE PEAK AMPLITUDE

For a causal source with characteristic physical wavenumber $k_\star/a_\star$, a generic estimate was derived in Sec. 2.3, repeated here for convenience:

$$\Omega_{\text{GW}} \sim \left(\frac{\rho_{\text{source}}}{\rho_{\text{tot}}} \right)^2 \left(\frac{a_\star H_\star}{k_\star} \right)^2. \quad (3.62)$$

During supercooling, the energy density is dominated by the false-vacuum energy of the scalar sector. Once the transition begins, this energy is transferred into the rolling field and its fluctuations. For an order-of-magnitude estimate, one may therefore take

$$\frac{\rho_{\text{source}}}{\rho_{\text{tot}}} \sim 1. \quad (3.63)$$

Combining this with Eq. (3.59) gives

$$\Omega_{\text{GW}}^{\text{peak}} \sim \left(\frac{H_\star}{m_{\varphi, \text{QCD}}} \right)^2. \quad (3.64)$$

Note that when $m_{\varphi, \text{QCD}}/H \sim 1$, the assumptions underlying this sub-Hubble estimate are only marginally satisfied. This equation should therefore be understood as a parametric scaling. The overall numerical proportionality factor, which incorporates the anisotropic-stress efficiency, source coherence, and other order-unity factors omitted in this estimate, together with the precise peak position and spectral shape, must be extracted from the lattice.

Equation (3.64) also explains why the classically scale-invariant potential is especially favourable for gravitational-wave production. In an ordinary polynomial symmetry-breaking potential with no additional hierarchy, the negative mass parameter μ_φ is typically comparable to the temperature scale of the transition. Taking an electroweak-scale transition as an example, $\mu_\varphi \sim T_{\text{EW}}$, such that the characteristic source scale is $k_\star/a_\star \sim T_{\text{EW}}$. During radiation domination, the Hubble rate is instead set by $H \sim T_{\text{EW}}^2/M_{\text{Pl}}$. Even if the scalar supplied the entire source energy, one would therefore obtain

$$\Omega_{\text{GW}}^{\text{peak}} \sim \left(\frac{T_{\text{EW}}}{M_{\text{Pl}}} \right)^2 \ll 1. \quad (3.65)$$

For $T_{\text{EW}} \sim \mathcal{O}(100 \text{ GeV})$, this corresponds to a suppression of order 10^{-33} – 10^{-32} before accounting for the additional efficiency factor.

3 Supercooling and Tachyonic Symmetry Breaking in the $U(1)_{B-L}$ Model

The hierarchy is very different in the present model. Classical scale invariance is broken only weakly by the QCD-induced Higgs expectation value, so

$$m_{\varphi, \text{QCD}} \ll v_{\varphi}. \quad (3.66)$$

The curvature governing the instability can therefore be much smaller than the underlying symmetry-breaking scale and, across part of the parameter space, only moderately larger than H . The amplified structures can then approach the Hubble radius rather than lying many orders of magnitude inside it. This is the origin of the comparatively large signal expected from the tachyonic transition.

The lattice results will test the scaling in Eq. (3.64), determine its numerical coefficient and establish where around $m_{\varphi, \text{QCD}}$ the spectrum actually peaks.

3.5.3 REHEATING AND REDSHIFTING

To relate the gravitational-wave spectrum at production to its present-day frequency, the redshifting relation derived in Sec. 2.3 requires the temperature and Hubble parameter at the end of gravitational-wave production. We assume that the scalar energy released during the tachyonic transition is transferred rapidly to the Standard Model and thermalises within less than a Hubble time. Under this instantaneous-reheating approximation, $T_{\star}$ and $H_{\star}$ from Eq. (2.84) are identified with the reheating temperature T_{rh} and the corresponding Hubble parameter H_{rh} .

Energy conservation gives

$$\frac{\pi^2}{30} g_* T_{\text{rh}}^4 = \Delta V_{\text{eff}}(0), \quad (3.67)$$

where $\Delta V_{\text{eff}}(0)$ is the zero-temperature energy difference between the false and true vacua, and the relativistic degrees of freedom are well approximated by their Standard Model value $g_* = 106.75$. The reheating temperature is therefore

$$T_{\text{rh}} = \left[\frac{30 \Delta V_{\text{eff}}(0)}{\pi^2 g_*} \right]^{1/4}. \quad (3.68)$$

Using Eq. (3.32), this becomes approximately

$$T_{\text{rh}} \simeq \left[\frac{90}{128 \pi^4 g_*} \right]^{1/4} m_{Z'}. \quad (3.69)$$

In the numerical analysis, Eq. (3.67) is instead solved using the full RG-improved vacuum-energy difference.

3.5 Analytic Expectations for Gravitational-Wave Production

The Hubble parameter immediately after reheating follows from the Friedmann equation,

$$H_{\text{rh}} = \sqrt{\frac{\Delta V_{\text{eff}}(0)}{3M_{\text{Pl}}^2}} \simeq \frac{1}{\sqrt{128}\pi} \frac{m_{Z'}^2}{M_{\text{Pl}}}. \quad (3.70)$$

Equations (3.68) and (3.70) provide the model-dependent quantities entering the frequency-redshifting equation of Sec. 2.3.

4 NUMERICAL IMPLEMENTATION WITH *CosmoLattice*

4.1 CLASSICAL LATTICE DESCRIPTION

The linear treatment developed in Sec. 3.4 ceases to be valid once the tachyonically amplified fluctuations backreact on the rolling field and different Fourier modes begin to interact. The subsequent evolution must therefore be obtained from the full nonlinear field equation (2.32). We simulate this stage with *CosmoLattice*, a framework for the real-time evolution of classical fields in an expanding Universe [19–21].

Only the components of *CosmoLattice* used in this work are described here. These are the evolution of a real scalar field, the self-consistent expansion of the homogeneous background, the simultaneous evolution of tensor perturbations, and the extraction of scalar and gravitational-wave observables in Fourier space.

4.1.1 DISCRETISED EQUATIONS OF MOTION

The scalar field is represented on a periodic cubic lattice of comoving side length L , containing N^3 lattice sites. The comoving lattice spacing is therefore

$$\delta x = \frac{L}{N}. \quad (4.1)$$

A lattice site is labelled by $\mathbf{n} = (n_1, n_2, n_3)$, and the continuous field $\varphi(\mathbf{x}, t)$ is replaced by the set of values $\varphi_{\mathbf{n}}(t)$.

Spatial derivatives are approximated by finite differences. For the nearest-neighbour lattice Laplacian,

$$\left(\nabla_{\text{L}}^2 \varphi\right)_{\mathbf{n}} = \frac{1}{\delta x^2} \sum_{i=1}^3 (\varphi_{\mathbf{n}+\hat{\mathbf{e}}_i} - 2\varphi_{\mathbf{n}} + \varphi_{\mathbf{n}-\hat{\mathbf{e}}_i}), \quad (4.2)$$

where $\hat{e}_i$ denotes one lattice step in the i -th spatial direction. The full nonlinear scalar equation is then represented by

$$\ddot{\varphi}_{\mathbf{n}} + 3H\dot{\varphi}_{\mathbf{n}} - \frac{1}{a^2} \left(\nabla_{\mathbf{L}}^2 \varphi \right)_{\mathbf{n}} + \left. \frac{\partial V_{\text{sim}}}{\partial \varphi} \right|_{\varphi=\varphi_{\mathbf{n}}} = 0. \quad (4.3)$$

Unlike the linear mode equation discussed in Sec. 3.4, this equation retains the complete local potential force. Backreaction, rescattering and fragmentation therefore arise directly from the evolution rather than being added as separate effects. The simulation potential V_{sim} and its relation to the physical $U(1)_{B-L}$ potential are introduced in Sec. 4.2.

The expansion of the Universe is evolved simultaneously with the scalar field. The homogeneous background is sourced by spatial averages over the lattice,

$$\langle X \rangle_V \equiv \frac{1}{N^3} \sum_{\mathbf{n}} X_{\mathbf{n}}, \quad (4.4)$$

so that the Hubble constraint becomes

$$H^2 = \frac{\langle \rho_{\text{tot}} \rangle_V}{3M_{\text{Pl}}^2}. \quad (4.5)$$

Here, ρ_{tot} contains the scalar kinetic, gradient and potential contributions. The field configuration therefore determines the background expansion, which in turn enters the local field equation through the Hubble-friction and gradient terms. This gives a self-consistent evolution of the nonlinear scalar dynamics in an expanding background.

Scalar observables such as volume averages and power spectra are evaluated directly from the evolving lattice configuration. The simultaneous evolution of the tensor perturbations and the extraction of the gravitational-wave spectrum are discussed separately in Sec. 4.4.

4.1.2 VALIDITY OF THE CLASSICAL-STATISTICAL APPROXIMATION

The lattice equations describe classical fields, whereas the initial fluctuations originate from a quantum state. The connection is provided by the classical-statistical approximation. The initial quantum two-point functions are represented by an ensemble of classical field configurations with random phases and amplitudes. Each configuration is then evolved with the classical equations of motion, and observables are obtained from spatial or ensemble averages. As discussed in Sec. 2.2, this approximation is reliable for bosonic modes with large occupation numbers,

$$n_k \gg 1. \quad (4.6)$$

The vacuum fluctuations used to initialise the simulation do not initially satisfy this condition. They serve only as seeds for the instability. Tachyonic amplification subsequently increases the amplitudes and occupation numbers of the unstable modes exponentially. Once these modes are highly occupied, their statistical fluctuations dominate over their quantum commutators and their evolution is accurately approximated by classical fields.

The approximation is required only for the modes that dominate the nonlinear dynamics and the gravitational-wave source. Ultraviolet modes outside the instability band may remain close to their vacuum state and should not be interpreted as a classical particle population. At sufficiently late times, classical rescattering can also transfer energy towards the lattice cut-off and introduce an unphysical ultraviolet dependence. To avoid this, we stop the simulations before this effect can influence the observables of interest and check the robustness of our results by varying the lattice parameters in Sec. 4.5.

4.2 SIMULATION POTENTIAL AND SCALE MATCHING

The physical effective potential introduced in Chapter 3 contains a hierarchy of scales that cannot be represented simultaneously on a finite lattice. We therefore replace its ultraviolet structure by a simpler symmetry-breaking potential while preserving the infrared quantities that control the onset of the tachyonic transition. The central assumption is that the nonlinear dynamics and the resulting gravitational-wave spectrum are primarily determined by the curvature near the origin relative to the Hubble scale.

4.2.1 HIERARCHY OF SCALES IN THE PHYSICAL MODEL

In the rolling region of the physical $U(1)_{B-L}$ model, the relevant scales obey [17]

$$H \lesssim m_{\varphi, \text{QCD}} \ll m_{\varphi} \lesssim v_{\varphi}. \quad (4.7)$$

Here, H is the approximately constant Hubble rate during vacuum domination, $m_{\varphi, \text{QCD}}$ is the magnitude of the negative curvature generated by the QCD-induced Higgs expectation value, m_{φ} is the scalar mass about the broken minimum, and v_{φ} is the $U(1)_{B-L}$ symmetry-breaking scale. These quantities were introduced in Chapter 3.

The smallest dynamically relevant momentum at the onset of the transition is of order H , whereas resolving the subsequent oscillations about the true minimum requires momenta at least of order m_{φ} . A lattice with comoving side length L and spacing δx covers approximately the range

$$k_{\text{IR}} \simeq \frac{2\pi}{L}, \quad k_{\text{UV}} \simeq \frac{\pi}{\delta x} = \frac{N}{2} k_{\text{IR}}. \quad (4.8)$$

To include Hubble-scale modes while resolving the true-vacuum mass would therefore require

$$\frac{k_{\text{IR}}}{a} \lesssim H, \quad \frac{k_{\text{UV}}}{a} \gtrsim m_\varphi, \quad (4.9)$$

and hence, up to factors of order unity,

$$N \gtrsim 2 \frac{m_\varphi}{H}. \quad (4.10)$$

The physical hierarchy makes this requirement computationally inaccessible.

The same hierarchy also produces a severe separation of timescales. The time step must be short enough to follow oscillations with period m_φ^{-1} , while the transition and expansion are followed over timescales comparable to H^{-1} . In addition, the field must traverse the large field-space interval between the origin and v_φ . A direct implementation of the complete RG-improved potential would therefore be prohibitively expensive.

4.2.2 SURROGATE SCALAR POTENTIAL

To isolate the infrared dynamics, the physical potential is replaced by

$$V_{\text{sim}}(\varphi) = -\frac{\mu^2}{2}\varphi^2 + \frac{\lambda}{4}\varphi^4 + V_0, \quad (4.11)$$

where $\mu^2 > 0$ and $\lambda > 0$. The origin has negative curvature,

$$V''_{\text{sim}}(0) = -\mu^2, \quad (4.12)$$

and the two symmetry-breaking minima lie at

$$\pm v_{\text{sim}} = \pm \frac{\mu}{\sqrt{\lambda}}. \quad (4.13)$$

The scalar mass about either minimum is

$$m_{\text{sim}}^2 = V''_{\text{sim}}(v_{\text{sim}}) = 2\mu^2. \quad (4.14)$$

The minimal constant shift that makes the potential non-negative is

$$V_0 = \frac{\mu^4}{4\lambda}, \quad (4.15)$$

for which $V_{\text{sim}}(\pm v_{\text{sim}}) = 0$ and $V_{\text{sim}}(0) = V_0$. Since V_0 is independent of φ , it does not enter the scalar equation of motion. It therefore does not alter the instability band, the scalar force or the anisotropic stress generated by spatial gradients.

The constant does, however, contribute to the homogeneous energy density. Near the initial configuration $\varphi \simeq 0$, it represents the false-vacuum energy driving the expansion. Neglecting the initial kinetic and gradient energies,

$$H_{\text{sim}}^2 \simeq \frac{V_0}{3M_{\text{Pl}}^2} = \frac{\mu^4}{12\lambda M_{\text{Pl}}^2}. \quad (4.16)$$

Consequently,

$$\left(\frac{\mu}{H}\right)_{\text{sim}} \simeq \sqrt{12\lambda} \frac{M_{\text{Pl}}}{\mu}. \quad (4.17)$$

Once the target value of μ/H is specified, the constant shift introduces no additional dynamical freedom.

4.2.3 MATCHING THROUGH μ/H

Sufficiently far below T_{roll} , the physical mode equation near the origin is given by Eq. (3.55). The corresponding simulation equation has the same form,

$$\ddot{u}_k + 3H\dot{u}_k + \left(\frac{k^2}{a^2} - \mu^2\right)u_k \simeq 0. \quad (4.18)$$

The two systems are therefore related by matching the negative curvature in units of the Hubble rate,

$$\left(\frac{\mu}{H}\right)_{\text{sim}} \longleftrightarrow \left(\frac{m_{\varphi, \text{QCD}}}{H}\right)_{\text{physical}}. \quad (4.19)$$

This matching preserves the properties associated with the initial tachyonic instability. In particular, the simulation reproduces the width of the instability band in Hubble units,

$$0 \leq \frac{k}{aH} < \frac{\mu}{H} \leftrightarrow \frac{m_{\varphi, \text{QCD}}}{H}, \quad (4.20)$$

as well as the growth rate relative to the expansion. It consequently preserves the characteristic source size and duration relative to the Hubble scale,

$$HR_{\star} \sim Ht_{\star} \sim \frac{H}{\mu} \leftrightarrow \frac{H}{m_{\varphi, \text{QCD}}}. \quad (4.21)$$

These are the quantities that determine whether the transition is deeply sub-Hubble or occurs close to the Hubble scale. They also control the parametric peak-amplitude scaling derived in Sec. 3.5.

The matching does not reproduce the complete physical potential. In the surrogate model, the true-vacuum mass is only $m_{\text{sim}} = \sqrt{2}\mu$, whereas physically $m_\varphi \gg m_{\varphi, \text{QCD}}$. Likewise, v_{sim} is not identified with the physical symmetry-breaking scale v_φ . The logarithmic Coleman–Weinberg shape, the running of the couplings, and the steep rise of the potential near the physical minimum are therefore omitted. The simulation also does not describe ultraviolet modes associated with m_φ , the detailed high-frequency oscillations about the physical vacuum, or the subsequent perturbative decay into Standard Model particles.

The lattice results should thus be interpreted as a description of the infrared, nonlinear dynamics generated by the initial curvature flip. Absolute physical energy scales, the reheating temperature and the redshifting of the gravitational-wave spectrum are restored only when the simulation results are mapped back to the $U(1)_{B-L}$ model.

4.3 INITIAL CONDITIONS AND SIMULATION PARAMETERS

The simulations are formulated in terms of the dimensionless program variables used by *CosmoLattice*,

$$\tilde{\varphi} = \frac{\varphi}{\varphi_{\text{ref}}}, \quad d\tilde{\eta} = a^{-\alpha} M_{\text{ref}} dt, \quad d\tilde{\mathbf{x}} = M_{\text{ref}} d\mathbf{x}. \quad (4.22)$$

Here, φ_{ref} and M_{ref} correspond, respectively, to the quantities conventionally denoted by $f_\star$ and $\omega_\star$ in the *CosmoLattice* literature. For the surrogate potential in Eq. (4.11), we choose

$$\varphi_{\text{ref}} = \frac{\mu}{\sqrt{\lambda}} = v_{\text{sim}}, \quad M_{\text{ref}} = \mu, \quad \alpha = 1. \quad (4.23)$$

Field amplitudes are therefore measured in units of the symmetry-breaking scale v_{sim} , while momenta and inverse times are measured in units of the tachyonic mass scale μ . Since $\alpha = 1$, the numerical time coordinate is conformal time, with $d\tilde{\eta} = \mu d\eta$.

The main simulations cover

$$\left(\frac{\mu}{H}\right)_{\text{sim}} = 10^3, 10^4, \dots, 10^9. \quad (4.24)$$

This range is chosen such that the characteristic scalar-field scale remains well inside the Hubble radius. In this regime, the estimate derived in Eq. (3.64) suggests an approximately power-law dependence of the peak amplitude on the matched ratio $m_{\varphi, \text{QCD}}/H$. Keeping $(\mu/H)_{\text{sim}} \gg 1$

therefore allows the numerical peak amplitudes to be fitted directly as a function of μ/H before mapping them onto the physical model. Close to $\mu/H \sim 1$, the characteristic source scale approaches the Hubble radius and the sub-Hubble estimate underlying this scaling is no longer expected to hold.

For a fixed value of μ/H , the parameters μ and λ are not determined separately. Equation (4.17) constrains only their combination, so different pairs (μ, λ) realise the same matched ratio. For each benchmark, these parameters are chosen to avoid unnecessarily large or small numerical values. Their individual values should therefore be regarded as a convenient numerical parametrisation rather than additional physical parameters. To what extent they influence the gravitational-wave signal is tested in Sec. 4.5.

4.3.1 HOMOGENEOUS FIELD AND VACUUM FLUCTUATIONS

The homogeneous component is initialised exactly at the origin and at rest,

$$\bar{\varphi}_{\text{init}} = 0, \quad \dot{\bar{\varphi}}_{\text{init}} = 0. \quad (4.25)$$

This is the maximally symmetric initial condition for the surrogate potential and avoids introducing an arbitrary homogeneous displacement. The subsequent departure from the unstable maximum is instead seeded entirely by the initial field fluctuations. Since the two branches of the symmetric surrogate potential are dynamically equivalent, the sign towards which the field eventually evolves is irrelevant for the observables considered here.

The inhomogeneous component is initialised as a Gaussian random field representing vacuum fluctuations,

$$\varphi(\mathbf{x}, \eta_{\text{init}}) = \delta\varphi(\mathbf{x}, \eta_{\text{init}}). \quad (4.26)$$

In Fourier space, *CosmoLattice* samples random amplitudes and phases such that the ensemble reproduces the vacuum two-point function. The field momenta are initialised consistently with the same state, so that $\dot{\bar{\varphi}}_{\text{init}} = 0$ applies only to the homogeneous mode.

For a scalar with positive initial mass squared, the corresponding vacuum frequency would be

$$\omega_{k,\text{init}}^2 = k^2 + a_{\text{init}}^2 m_{\text{init}}^2. \quad (4.27)$$

In the present case, however, $V''_{\text{sim}}(0) = -\mu^2 < 0$. For the unstable modes this expression does not define a real positive-frequency vacuum state. *CosmoLattice* therefore initialises the field using

$$\omega_{k,\text{init}} = |\mathbf{k}|, \quad (4.28)$$

corresponding to massless vacuum fluctuations.

It is useful to make explicit how this prescription depends on the numerical choice of μ and λ . For the massless initial state the dimensionless power per logarithmic momentum interval is

$$\Delta_{\delta\varphi}(k) = \frac{k^2}{4\pi^2 a_{\text{init}}^2}. \quad (4.29)$$

Using $k = \mu \tilde{k}$, $\tilde{\varphi} = \varphi / v_{\text{sim}}$, and $v_{\text{sim}} = \mu / \sqrt{\lambda}$, the corresponding spectrum in program units becomes

$$\tilde{\Delta}_{\delta\varphi}(\tilde{k}) \equiv \frac{\Delta_{\delta\varphi}(k)}{v_{\text{sim}}^2} = \frac{\lambda \tilde{k}^2}{4\pi^2 a_{\text{init}}^2}. \quad (4.30)$$

At fixed μ/H , Eq. (4.17) implies

$$\lambda = \frac{\mu^2}{12M_{\text{Pl}}^2} \left(\frac{\mu}{H} \right)^2, \quad (4.31)$$

so that $\lambda \propto \mu^2$ when the matched ratio is held fixed. The dimensionless initial vacuum power therefore grows with the absolute value of μ , even though the infrared tachyonic dynamics is controlled by the same value of μ/H .

No additional ultraviolet cut-off is imposed on the initial fluctuations in the main simulations. All modes represented on the lattice are therefore populated according to the same vacuum prescription. This avoids introducing an arbitrary additional momentum scale into the initial state. As will be shown in Sec. 4.5, imposing a finite cut-off changes only the initialisation-sensitive ultraviolet tail, while the gravitational-wave peak relevant for the physical analysis is essentially unaffected.

A fixed random seed is used throughout the main simulation set to make the comparison between different benchmark points reproducible. The scale factor is initially normalised to

$$a_{\text{init}} = 1, \quad (4.32)$$

so that physical and comoving momenta coincide at the initial time. It is subsequently evolved self-consistently from the lattice-averaged energy density and pressure.

4.3.2 LATTICE SIZE, RESOLUTION, AND TIME STEP

All main simulations use a periodic cubic lattice with $N^3 = 512^3$ sites and infrared momentum

$$\tilde{k}_{\text{IR}} \equiv \frac{k_{\text{IR}}}{\mu} = 0.1. \quad (4.33)$$

The infrared scale also sets the momentum resolution,

$$\Delta\tilde{k} = \tilde{k}_{\text{IR}} = 0.1. \quad (4.34)$$

The characteristic initial instability band $k/a < \mu$ is consequently resolved by approximately ten radial momentum intervals.

Since $k_{\text{IR}} = 2\pi/L$, the dimensionless box length and lattice spacing are

$$\tilde{L} = \mu L = \frac{2\pi}{\tilde{k}_{\text{IR}}} = 20\pi, \quad \delta\tilde{x} = \frac{\tilde{L}}{N} \simeq 0.123. \quad (4.35)$$

Following the notation of the *CosmoLattice* literature, the Nyquist momentum in each spatial direction is

$$\tilde{k}_{\text{UV}} = \frac{N}{2} \tilde{k}_{\text{IR}} = 25.6. \quad (4.36)$$

This is the maximum momentum that can be represented independently in every Cartesian direction.

The equations are evolved with the second-order staggered-leapfrog algorithm using the program time step

$$\delta\tilde{\eta} = 0.01. \quad (4.37)$$

Together with Eq. (4.35), this gives

$$\frac{\delta\tilde{\eta}}{\delta\tilde{x}} \simeq 0.081, \quad (4.38)$$

well below the usual three-dimensional stability estimate $\delta\tilde{\eta}/\delta\tilde{x} < 1/\sqrt{3}$. The simulations are evolved until $\tilde{\eta}_{\text{end}} = 150$.

The parameters common to the main simulations are summarised in Table 4.1.

4.4 EXTRACTION OF THE GRAVITATIONAL-WAVE SPECTRUM

The continuum description of gravitational waves sourced by scalar inhomogeneities was introduced in Sec. 2.3. Here, we only summarise its implementation on the lattice. A detailed discussion of the corresponding algorithms and lattice projectors can be found in Ref. [21].

Parameter	Value
Program scales	$\varphi_{\text{ref}} = v_{\text{sim}}, M_{\text{ref}} = \mu, \alpha = 1$
Init. hom. field	$\tilde{\varphi}_{\text{init}} = 0, \dot{\tilde{\varphi}}_{\text{init}} = 0$
Lattice size	$N^3 = 512^3$
IR-cut-off	$\tilde{k}_{\text{IR}} = 0.1$
Box length	$\tilde{L} = 20\pi$
Lattice spacing	$\delta\tilde{x} \approx 0.123$
UV-cut-off	$\tilde{k}_{\text{UV}} = 25.6$
Time step	$\delta\tilde{\eta} = 0.01,$
Final time	$\tilde{\eta}_{\text{end}} = 150$
Evolution algorithm	Staggered leapfrog

Table 4.1: Parameters used for the main lattice simulations. Quantities carrying a tilde are expressed in the program units defined in Eq. (4.22).

4.4.1 ANISOTROPIC STRESS AND TT PROJECTION

For the single canonical scalar field considered here, the part of the spatial stress tensor relevant for gravitational-wave production can be written as

$$\Pi_{ij}^{\text{eff}} = \partial_i \varphi \partial_j \varphi. \quad (4.39)$$

In principle, one could Fourier transform Π_{ij}^{eff} at every time step, project it onto its TT part, and evolve the physical tensor perturbation h_{ij} . This would, however, require repeated forward and inverse Fourier transforms throughout the simulation. *CosmoLattice* instead makes use of the linearity of the tensor equation and evolves an auxiliary symmetric tensor u_{ij} sourced by the unprojected stress. The lattice form of this equation in conformal time is

$$u_{ij}'' + 2\mathcal{H}u_{ij}' - \nabla_L^2 u_{ij} = \frac{2}{M_{\text{Pl}}^2} \Pi_{ij}^{\text{eff}}, \quad (4.40)$$

where the source is evaluated locally from the scalar field at each lattice site. The auxiliary tensor can therefore be evolved entirely in position space alongside φ .

The physical tensor perturbation is reconstructed only when a gravitational-wave output is required. After Fourier transforming u_{ij} ,

$$h_{ij}(\mathbf{k}, \eta) = \Lambda_{ij,lm}^L(\mathbf{k}) u_{lm}(\mathbf{k}, \eta), \quad (4.41)$$

where $\Lambda_{ij,lm}^L$ is the lattice TT projector. Its structure is analogous to the continuum projector,

$$\Lambda_{ij,lm}^L = P_{il}^L P_{jm}^L - \frac{1}{2} P_{ij}^L P_{lm}^L, \quad P_{ij}^L = \delta_{ij} - \hat{k}_{L,i} \hat{k}_{L,j}, \quad (4.42)$$

but the unit vector $\hat{\mathbf{k}}_{\text{L}}$ is constructed from the lattice momentum associated with the discrete derivative operator rather than from the continuum momentum $\mathbf{k}$. This ensures that the projected tensor is transverse with respect to the same discretisation used to evolve the fields. The main simulations use the choice `GWprojectorType = 1`. Further details on the available lattice projectors and their numerical properties are given in Ref. [21].

The tensor perturbations are evolved only to linear order. Their energy density remains small compared with that stored in the scalar field, so their backreaction on the scalar dynamics and on the expansion is neglected.

4.4.2 SPECTRAL ENERGY DENSITY

At every spectral output time, the auxiliary tensor u_{ij} and its conjugate momentum are transformed to Fourier space and projected according to Eq. (4.41). The quantity entering the gravitational-wave energy density is then the squared time derivative of the projected field,

$$h'_{ij}(\mathbf{k}, \eta) h'^{*}_{ij}(\mathbf{k}, \eta). \quad (4.43)$$

This is the lattice counterpart of the quantity appearing in the gravitational-wave energy density in Eq. (2.59). Because the system is statistically isotropic, the gravitational-wave spectrum depends only on the magnitude of the comoving momentum, rather than on its direction. We therefore group Fourier modes with similar momentum magnitude k . In *CosmoLattice*, this is done by sorting the points of the reciprocal lattice into spherical momentum shells and averaging Eq. (4.43) over all modes within each shell. Accounting for the number of lattice modes in each shell and the appropriate normalisation then gives a discrete estimate of $\Omega_{\text{GW}}(k, \eta)$ as defined in Eq. (2.60). The precise lattice expression depends on the convention used to define the power spectrum and can be found in Ref. [21].

The output momentum is expressed in the program units introduced in Sec. 4.3,

$$\tilde{k} = \frac{k}{\mu}. \quad (4.44)$$

At a given simulation time, the physical momentum is $k_{\text{phys}} = k/a$. It is therefore convenient to express the numerical spectrum in terms of

$$\frac{k_{\text{phys}}}{\mu} = \frac{\tilde{k}}{a}, \quad (4.45)$$

or, when comparing directly with the Hubble scale,

$$\frac{k_{\text{phys}}}{H} = \frac{\tilde{k}}{a} \left(\frac{\mu}{H} \right)_{\text{sim}}. \quad (4.46)$$

Through the matching relation of Sec. 4.2.3, the latter quantity maps directly onto the corresponding physical ratio involving $m_{\varphi, \text{QCD}}/H$.

For the main analysis, the gravitational-wave spectrum is evaluated at the end of the lattice evolution, after the rapid production stage has largely saturated. The adequacy of the chosen final time is checked explicitly in Sec. 4.5. The subsequent conversion from the production spectrum to the present-day frequency and abundance does not depend on the lattice implementation and uses the redshift relations derived in Sec. 2.3.3. The resulting spectra and their mapping to the physical $U(1)_{B-L}$ parameter space are presented in Chapter 5.

4.5 NUMERICAL CONSISTENCY CHECKS

Before using the lattice results to determine the dependence of the gravitational-wave signal on μ/H , we verify that the quantities of interest are sufficiently insensitive to the main numerical choices. All tests are performed for the benchmark

$$\left(\frac{\mu}{H} \right)_{\text{sim}} = 10^3. \quad (4.47)$$

Unless stated otherwise, all remaining parameters are kept at their standard values from Sec. 4.3. In particular, the same random seed is used for every comparison, so that differences between the runs can be attributed directly to the parameter being varied. The intrinsic dependence on the random realisation itself is tested separately below.

To quantify changes in the peak amplitude, we define the relative deviation from a reference simulation as

$$\Delta_{\Omega}(X; X_{\text{ref}}) \equiv \frac{\left| \Omega_{\text{GW}}^{\text{peak}}(X) - \Omega_{\text{GW}}^{\text{peak}}(X_{\text{ref}}) \right|}{\Omega_{\text{GW}}^{\text{peak}}(X_{\text{ref}})}. \quad (4.48)$$

Here, X denotes the parameter being varied and X_{ref} the better-resolved or reference result.

4.5.1 SPATIAL RESOLUTION AND FINITE VOLUME

We first test the spatial resolution by keeping $\tilde{k}_{\text{IR}} = 0.1$ fixed and varying $N = 64, 128, 256, 512$. The physical volume is therefore unchanged, while increasing N decreases the lattice spacing and pushes the ultraviolet cut-off towards larger momenta.




Figure 4.1: Spatial-resolution comparison for $(\mu/H)_{\text{sim}} = 10^3$. Left: gravitational-wave spectra at $\tilde{\eta} = 150$ for different lattice spacings. Right: corresponding peak amplitudes as a function of $\delta\tilde{x}$, with the colours matching those of the spectra in the left panel. The point at $\delta\tilde{x} \approx 0.12$, corresponding to $N = 512$, is the resolution used for the main simulations.

The right panel of Fig. 4.1 shows a clear convergence of the peak amplitude as the lattice spacing is reduced. Between the two highest resolutions,

$$\Delta_{\Omega}(0.25; 0.12) \approx 1.9\%. \quad (4.49)$$

The choice $\delta\tilde{x} \approx 0.12$ therefore provides percent-level convergence while remaining computationally manageable.

The full spectra illustrate how this convergence is reached. For coarse lattices, the ultraviolet cut-off lies relatively close to the dynamically relevant scales and the spectrum develops a more pronounced peak structure. As the resolution is increased, the cut-off moves towards larger k_{phys}/μ and the shape at lower momenta changes progressively less. At sufficiently high resolution, a second feature also becomes visible in the ultraviolet. This feature is related to the initial vacuum population of high-momentum modes rather than to the tachyonic amplification itself. As will be discussed later in this section, it does not affect the convergence of the main gravitational-wave peak used in the analysis.

Next, we test the finite-volume dependence by keeping the spatial resolution δx fixed while changing the size of the simulation box L . We use

$$(N, \tilde{k}_{\text{IR}}) = (64, 0.8), (256, 0.2), (512, 0.1), \quad (4.50)$$

for which $N\tilde{k}_{\text{IR}}$ is constant and hence $\delta\tilde{x}$ and $\tilde{k}_{\text{UV}}$ remain unchanged. The corresponding box sizes in program units are $\tilde{L} = 2.5\pi, 10\pi, 20\pi$.




Figure 4.2: Finite-volume comparison at fixed spatial resolution. Left: gravitational-wave spectra at $\tilde{\eta} = 150$ for different box sizes. Right: corresponding peak amplitudes as a function of $\tilde{L}$, with matching colours between the two panels. The point at $\tilde{L} = 20\pi$, corresponding to $(N, \tilde{k}_{\text{IR}}) = (512, 0.1)$, is the volume used for the main simulations.

Increasing the box size extends the accessible spectrum towards smaller momenta, as expected. The smallest volume produces a visibly different peak, whereas the two larger boxes agree closely. In particular,

$$\Delta_{\Omega}(10\pi; 20\pi) \approx 1.3\%. \quad (4.51)$$

The peak amplitude therefore shows clear finite-volume convergence.

Some dependence on the volume remains in the ultraviolet tail. Since the spatial resolution is identical in all three runs, this cannot be attributed to a changing ultraviolet cut-off. Instead, a larger box contains more long-wavelength modes and provides denser sampling of momentum space, which can feed into the later nonlinear evolution. The main peak is considerably more stable than this high-momentum region. Since it is the peak that enters the subsequent physical analysis, $\tilde{L} = 20\pi$ is sufficiently large for the purposes of this work.

4.5.2 INITIAL CONDITIONS AND SURROGATE PARAMETERS

The matching procedure of Sec. 4.2.3 fixes the ratio μ/H , but not μ and λ separately. To verify that the physical part of the spectrum does not depend on this remaining freedom, we repeat the benchmark simulation for seven values of μ , spanning

$$2.435 \times 10^7 \leq \mu \leq 2.435 \times 10^{13}, \quad (4.52)$$

in steps of one order of magnitude. For each value, λ is adjusted so that $(\mu/H)_{\text{sim}} = 10^3$ remains unchanged.




Figure 4.3: Gravitational-wave spectra for seven different choices of μ , with λ adjusted such that $(\mu/H)_{\text{sim}} = 10^3$ in every case. The value $\mu = 2.435 \times 10^{11}$ GeV is used for the corresponding benchmark in the main simulation set.

Figure 4.3 shows two clearly different behaviours. The part of the spectrum containing the tachyonically generated peak is almost unchanged over a wide range of μ . Around the chosen value $\mu = 2.435 \times 10^{11}$, varying μ by one order of magnitude in either direction changes the peak amplitude by less than 1%.

By contrast, the far-ultraviolet part develops a second peak that becomes increasingly prominent for larger μ . This behaviour originates from the μ -dependence of the vacuum initialisation derived in Sec. 4.3.1.

For modes inside the tachyonic band this dependence rapidly becomes unimportant, since their amplitudes grow exponentially and subsequently enter the nonlinear regime, largely erasing their memory of the precise vacuum seed. High-momentum modes outside the instability band are not amplified in the same way and can retain a visible dependence on their initial amplitudes. They do not represent the highly occupied classical field modes for which the lattice calculation is intended. We therefore identify the physical gravitational-wave peak with the stable low-momentum part of the spectrum and do not use the ultraviolet feature in the subsequent analysis.

The treatment of the initial ultraviolet modes was tested separately by comparing the reference initialisation without a momentum cut-off to a simulation with

$$\frac{k_{\text{cut}}}{\mu} = 4. \quad (4.53)$$

As discussed in Sec. 4.3.1, changing this cut-off directly modifies the vacuum population of modes outside the tachyonic band, but should have little effect on the modes that undergo the subsequent instability.

This expectation is borne out by the simulations. The gravitational-wave spectra are essentially indistinguishable throughout the region containing the tachyonically generated peak. The relative difference of the peak amplitudes is only

$$\Delta_{\Omega}(k_{\text{cut}}/\mu = 4; \text{no cut-off}) \simeq 10^{-5} = 0.001\%. \quad (4.54)$$

Slight differences remain only in the ultraviolet tail, where the initialisation-dependent second spectral feature discussed above appears. The negligible change of the physical peak shows that its amplitude is not set by the arbitrary treatment of the initial high-momentum vacuum modes. We therefore use the simpler prescription without an imposed k_{cut} for the main simulations.

A final dependence of the initial conditions arises from the random realisation of the vacuum fluctuations. To estimate its importance, three benchmark simulations with $\mu/H \in \{10^3, 10^5, 10^9\}$ were each repeated at least five times with different random seeds while keeping all physical and numerical parameters unchanged. Denoting the sample mean and standard deviation of the resulting peak amplitudes by $\overline{\Omega}_{\text{GW}}^{\text{peak}}$ and s_{Ω} , respectively, we find

$$\frac{s_{\Omega}}{\overline{\Omega}_{\text{GW}}^{\text{peak}}} \simeq 5 - 10\%. \quad (4.55)$$

This scatter is noticeably larger than the residual dependence on spatial resolution, volume, time step, or the treatment of the ultraviolet initial modes. It is not a numerical convergence problem, but reflects the fact that a finite simulation volume contains only one stochastic realisation of the initial vacuum fluctuations. Different seeds produce slightly different initial spatial configurations, which are subsequently amplified by the tachyonic instability and can lead to appreciable differences in the final anisotropic stress and hence in the gravitational-wave amplitude. In the main simulation set we therefore use the same random seed for every value of μ/H . This does not remove the finite-realisation uncertainty of the absolute peak amplitude, but it avoids introducing an independent stochastic variation at each benchmark and makes the relative dependence on μ/H considerably cleaner. The approximately 9% scatter found here should accordingly be kept in mind as a representative stochastic uncertainty of a single-realisation prediction.

We also varied the homogeneous initial field value. Moving the field away from the origin systematically lowers the gravitational-wave signal. For example, the relative difference between




Figure 4.4: Peak gravitational-wave amplitude for different numerical time steps. The point at $\delta \tilde{\eta} = 0.01$ corresponds to the value used in the main simulations.

$\bar{\varphi}_{\text{init}} = 10^{-6} v_{\text{sim}}$ and $\bar{\varphi}_{\text{init}} = 0$ is approximately 25%, while the relative difference between $\bar{\varphi}_{\text{init}} = 10^{-12} v_{\text{sim}}$ and $\bar{\varphi}_{\text{init}} = 0$ is on the order of 10^{-5} . A displaced homogeneous mode begins rolling down the potential before the fluctuations have experienced the same amount of tachyonic growth, reducing the relative importance of fragmentation and the associated anisotropic stresses. This provides additional motivation for the choice $\bar{\varphi}_{\text{init}} = 0$.

4.5.3 TIME STEP AND SIMULATION DURATION

The temporal discretisation is tested by varying

$$\delta \tilde{\eta} = 0.005, 0.01, 0.015, 0.02. \quad (4.56)$$

The peak amplitude rapidly converges as the time step is reduced. The difference between the fiducial result and the finer $\delta \tilde{\eta} = 0.005$ simulation is only

$$\Delta_{\Omega}(0.01; 0.005) \simeq 0.01\%. \quad (4.57)$$

The uncertainty associated with the temporal discretisation is therefore much smaller than that from the finite spatial resolution and volume.

Finally, we check that the standard endpoint $\tilde{\eta}_{\text{end}} = 150$ is sufficiently late for the gravitational-wave signal to have saturated. The benchmark simulation is extended to $\tilde{\eta} = 275$, and the evolution of the peak amplitude is shown in Fig. 4.5.




Figure 4.5: Evolution of the peak gravitational-wave amplitude for an extended simulation up to $\tilde{\eta} = 275$. The standard simulations are stopped at $\tilde{\eta} = 150$.

After its rapid initial growth, the peak gradually approaches a constant value. Comparing the standard endpoint directly with the end of the extended run gives

$$\Delta_{\Omega}(275; 150) \simeq 0.5\%. \quad (4.58)$$

Thus, extending every simulation substantially beyond $\tilde{\eta} = 150$ would change the final peak amplitude only at the percent level while considerably increasing the computational cost.

Taken together, these tests show that the gravitational-wave peak used in the subsequent analysis is numerically robust. Residual effects from spatial resolution and finite volume are at the percent level, while those associated with the time step, the choice of (μ, λ) at fixed μ/H , and the initial ultraviolet cut-off are smaller still. The largest variation instead comes from the stochastic realisation of the initial vacuum fluctuations, for which the seven-seed comparison gives a scatter of approximately 9% in the peak amplitude. The ultraviolet tail of the spectrum is less universal and retains a stronger dependence on the vacuum initialisation and numerical setup and is therefore not used to infer the physical scaling relations in the following chapter.

5

NUMERICAL RESULTS AND GRAVITATIONAL-WAVE PHENOMENOLOGY

5.1 NONLINEAR SCALAR-FIELD DYNAMICS

We first turn to the scalar-field dynamics of the transition. The benchmark simulation with $\mu/H = 1000$ will be used throughout this section to illustrate the main stages of the evolution. The other simulations follow the same qualitative sequence.

5.1.1 TACHYONIC GROWTH AND BACKREACTION

The beginning of the lattice evolution closely follows the linear picture developed in Sec. 3.4.2. The field starts close to the unstable origin, with small fluctuations superimposed on an almost homogeneous configuration. For the simplified lattice potential introduced in the previous chapter, the tachyonic band corresponds to $k_{\text{phys}}/\mu < 1$. Modes inside this band therefore begin to grow exponentially, with the longest wavelengths growing most efficiently, as expected from Eq. (3.57).

Figure 5.1 shows how this instability develops in position space. Alongside the absolute value of the spatial average $\langle \tilde{\varphi} \rangle$, we plot the root variance

$$\sigma_{\varphi} = \sqrt{\langle \tilde{\varphi}^2 \rangle - \langle \tilde{\varphi} \rangle^2}, \quad (5.1)$$

which measures the typical departure of the field from its spatial mean.

Initially, the average field remains very close to the origin, while the variance rises rapidly. This is precisely the regime in which the unstable modes grow almost independently and the linear treatment remains a useful guide. Around $\tilde{\eta} \simeq 10\text{--}20$, however, the fluctuations become comparable to the symmetry-breaking scale. The variance overshoots and begins to oscillate, while the average field itself starts to move appreciably towards the broken minimum.

This marks the onset of the backreaction discussed in Sec. 3.4.3. Once the fluctuations are no longer small, the higher-order terms neglected by the linear mode equations become important.




Figure 5.1: Evolution of the spatial average $|\langle \varphi \rangle|/v_{\text{sim}}$ and root variance $\sigma_{\varphi}/v_{\text{sim}}$ for the benchmark simulation $\mu/H = 1000$, shown as functions of the numerical conformal time $\tilde{\eta} = \mu\eta$. Note that the absolute value of the spatial average is shown here. The field average actually rolls towards the negative minimum $-v_{\text{sim}}$.

Different momentum modes interact and the amplified fluctuations modify the evolution of the spatially averaged field. From this point onwards, the full lattice evolution rather than a collection of independent mode equations is required.

At later times, $\langle \varphi \rangle$ approaches $-v_{\text{sim}}$, while the variance gradually decreases. It is worth noting that the average field does not immediately settle at exactly $-v_{\text{sim}}$. This is a consequence of the remaining inhomogeneities, rather than a shift of the minimum of the simulation potential.

This can be seen directly from the field decomposition introduced in Eq. (2.43). At any time we may write $\varphi = \langle \varphi \rangle + \delta\varphi$, with $\langle \delta\varphi \rangle = 0$. Here this is simply an exact separation into the lattice mean and the deviations from it, so $\delta\varphi$ is no longer assumed to be perturbatively small. It follows immediately that $\langle \delta\varphi^2 \rangle = \sigma_{\varphi}^2$, with σ_{φ} defined in Eq. (5.1).

For the quartic simulation potential introduced in the previous chapter, the spatially averaged first derivative contains the cubic moment $\langle \varphi^3 \rangle$. Expanding this about the spatial mean gives

$$\langle \varphi^3 \rangle = \langle \varphi \rangle^3 + 3\langle \varphi \rangle \sigma_{\varphi}^2 + \langle \delta\varphi^3 \rangle. \quad (5.2)$$

Once the mean field starts evolving only slowly, the spatially averaged equation of motion approximately requires the mean potential force to vanish

$$\left\langle \frac{\partial V_{\text{sim}}(\varphi)}{\partial \varphi} \right\rangle \simeq 0. \quad (5.3)$$

Neglecting a small skewness $\langle \delta\varphi^3 \rangle$, this gives

$$\frac{\langle \varphi \rangle^2}{v_{\text{sim}}^2} + 3 \frac{\sigma_\varphi^2}{v_{\text{sim}}^2} \simeq 1. \quad (5.4)$$

The late-time values in Fig. 5.1, $|\langle \varphi \rangle|/v_{\text{sim}} \simeq 0.94$ and $\sigma_\varphi/v_{\text{sim}} \simeq 0.20$, satisfy this relation remarkably well. The residual fluctuations therefore provide a simple explanation for why the spatial average remains slightly above $-v_{\text{sim}}$. As the inhomogeneous component is further damped, the average can continue to move slowly towards the vacuum value.

The development of these fluctuations in position space is illustrated in Fig. 5.2. Starting from an almost featureless initial configuration, the tachyonic instability first produces coherent structures on comparatively large length scales, reflecting the preferential growth of the infrared modes. Once the fluctuations become large and the evolution turns nonlinear, the field fragments further and structure develops on progressively shorter scales. The snapshots also show that the transition is not spatially uniform: different regions initially move towards opposite branches of the symmetry-breaking potential and remain there for some time. At late times, however, the negative branch dominates almost the entire volume, with only small regions of the opposite sign remaining. These local dynamics are fully compatible with the smooth evolution of the spatial average in Fig. 5.1, which captures the net preference for the negative minimum but not the substantial inhomogeneity of the field during the transition.

The same evolution is particularly transparent in momentum space. Figure 5.3 shows the scalar power spectrum $\Delta_\varphi/v_{\text{sim}}^2$ at successive times.

At early times, the expected tachyonic growth is clearly visible. Modes with $k_{\text{phys}}/\mu < 1$ are rapidly amplified, whereas modes outside the tachyonic band initially remain largely unaffected. The momentum dependence closely follows the linear picture discussed in Sec. 3.4.2: the growth is strongest in the infrared and becomes progressively weaker towards $k_{\text{phys}}/\mu = 1$, where the tachyonic growth rate vanishes. In addition to this infrared structure, the initial spectrum contains a second peak at much larger momenta. As discussed in Sec. 4.3, this ultraviolet feature originates from the initialisation of the lattice modes rather than from the tachyonic instability. It is therefore not used in the quantitative analysis below.




Figure 5.2: Two-dimensional slices through the centre of the simulation volume at selected numerical conformal times $\tilde{\eta}$ for the benchmark simulation $\mu/H = 1000$. The initially small fluctuations rapidly develop into long-wavelength structures through the tachyonic instability. As the evolution becomes nonlinear, increasingly short-wavelength structure appears and the field fragments. Different regions initially evolve towards opposite branches of the symmetry-breaking potential, before the negative branch comes to dominate the volume at late times. Note that the colour scale is adjusted separately in each panel to make the spatial structure visible throughout the evolution.

Finally, the spectra also show a mild decrease of power in the original tachyonic band after the main growth has saturated. By this stage the field has moved away from the region of negative curvature, so these modes are no longer tachyonically amplified. At the same time, nonlinear rescattering redistributes power towards shorter wavelengths and the expansion damps the remaining fluctuations. This is the momentum-space counterpart of the decreasing variance in Fig. 5.1. The system has therefore left the linear tachyonic-growth stage and entered the nonlinear, fragmented regime whose energy distribution is discussed next.




Figure 5.3: Evolution of the scalar power spectrum $\Delta_\phi/v_{\text{sim}}^2$ for the benchmark simulation $\mu/H = 1000$. The colour scale denotes the numerical conformal time $\tilde{\eta}$. The initial tachyonic band is given by $k_{\text{phys}}/\mu < 1$.

5.1.2 ENERGY TRANSFER AND FIELD FRAGMENTATION

The onset of the nonlinear regime is also clearly reflected in the energy budget of the scalar field. Figure 5.4 shows the lattice-averaged kinetic, gradient and potential energy densities, together with their sum, for the same $\mu/H = 1000$ benchmark considered above. All energy densities are normalised by $\mu^2 v_{\text{sim}}^2$.

At the beginning of the simulation the field is still close to the origin and is almost homogeneous and at rest. The gradient and kinetic contributions are therefore negligible, and essentially the entire energy density is stored in the potential. The initial value of 0.25 follows




Figure 5.4: Evolution of the lattice-averaged kinetic, gradient and potential energy densities, together with the total scalar-field energy density, for the benchmark simulation $\mu/H = 1000$. All quantities are shown in units of $\mu^2 v_{\text{sim}}^2$ as functions of the numerical conformal time $\tilde{\eta} = \mu\eta$.

directly from the normalisation of the simulation potential introduced in the previous chapter. With the true vacuum shifted to zero, the potential energy at the origin is

$$V_{\text{sim}}(0) = \frac{1}{4} \mu^2 v_{\text{sim}}^2, \quad (5.5)$$

and hence $V_{\text{sim}}(0)/(\mu^2 v_{\text{sim}}^2) = 1/4$. This is the simulation analogue of the vacuum-energy-dominated state from which the tachyonic transition begins.

Once the instability develops, the initially stored potential energy is rapidly converted into field motion and spatial gradients. This occurs at the same time as the sharp rise of the variance and the amplification of the $k_{\text{phys}}/\mu < 1$ modes seen in the previous subsection. As the field falls away from the unstable origin, the potential contribution drops steeply, while the kinetic energy rises. At the same time, the exponentially growing inhomogeneous modes generate a sizeable gradient contribution. The growth of the gradient energy is particularly important here: unlike the kinetic energy, it cannot be produced by purely homogeneous rolling and therefore directly signals the development of spatial structure in the field.

Around the end of the initial tachyonic growth stage, the three energy components undergo sizeable oscillations and repeatedly exchange energy. This coincides with the breakdown of the

linear description identified in Sec. 3.4.3. The fluctuations have by then become large enough to backreact on the mean field, while nonlinear mode coupling redistributes energy between different momentum scales. The large gradient contribution shows that the initially almost homogeneous configuration has fragmented into an inhomogeneous scalar field. Together with the broadening of the power spectrum in Fig. 5.3, this provides a useful diagnostic of the transition into the nonlinear regime.

The oscillations of the potential energy during this stage have a simple interpretation. Different parts of the field configuration roll through and oscillate around the broken minimum, repeatedly exchanging potential and kinetic energy. Because these oscillations are no longer coherent across the whole lattice, a substantial fraction of the energy remains stored in spatial gradients rather than returning to a homogeneous condensate.

After $\tilde{\eta} \sim 30$, the rapid oscillations subside and the energy components evolve much more smoothly. Kinetic and gradient energy remain sizeable throughout the rest of the simulation, confirming that the field does not return to a homogeneous configuration. Instead, nonlinear rescattering and the cosmic expansion gradually redistribute and dilute the energy stored in the fragmented field. This is consistent with the slow decrease of the variance and the redistribution of spectral power shown in Fig. 5.1.

The total energy density decreases monotonically once the transition begins. This does not represent a loss of energy from the numerical evolution, but is the expected consequence of evolving the scalar field in an expanding background. Using the kinetic and gradient contributions,

$$\rho = E_K + E_G + E_V, \quad p = E_K - \frac{1}{3}E_G - E_V, \quad (5.6)$$

the spatially averaged continuity equation can be written as

$$\frac{d\langle\rho\rangle}{dt} = -3H\langle\rho + p\rangle = -6H\langle E_K\rangle - 2H\langle E_G\rangle \leq 0. \quad (5.7)$$

While the field remains at the origin, both terms on the right-hand side are negligible and the vacuum-like initial energy density remains approximately constant. Once tachyonic growth generates kinetic and gradient energy, however, the expanding background continuously redshifts these contributions. The gradual decline of $\langle\rho\rangle$ in Fig. 5.4 is therefore the expected Hubble dilution of the scalar-field energy after the vacuum energy has been released.

5.2 GRAVITATIONAL-WAVE PRODUCTION

Having established the scalar-field dynamics of the transition, we now turn to the resulting gravitational-wave signal. As described in Sec. 4.4, the inhomogeneous scalar configuration sources tensor perturbations through its anisotropic stress. We first follow the build-up of the signal during the nonlinear evolution, before extracting its characteristic peak amplitude and momentum scale and constructing a spectral fit that can later be mapped onto the physical $U(1)_{B-L}$ model.

5.2.1 TIME EVOLUTION OF THE SIGNAL

The evolution of the gravitational-wave spectral energy density is shown in Fig. 5.5 for the same benchmark $\mu/H = 1000$ considered in the previous section.




Figure 5.5: Evolution of the gravitational-wave spectral energy density for the benchmark simulation $\mu/H = 1000$, shown as a function of the physical momentum k_{phys}/μ . The colour scale denotes the numerical conformal time $\tilde{\eta}$. Most of the signal is generated during the rapid growth and fragmentation of the scalar field, after which the spectral shape evolves only slowly.

At the beginning of the simulation the gravitational-wave signal is negligible, as expected for the nearly homogeneous initial field configuration. Once the tachyonic fluctuations have grown sufficiently to generate sizeable spatial gradients, the spectrum rises rapidly over several orders of magnitude. This occurs around the same time as the sharp increase of the field

variance and the conversion of potential energy into kinetic and gradient energy shown in Sec. 5.1. This is expected since the instability itself does not produce gravitational waves simply because the scalar amplitude grows, but because the growing inhomogeneities generate the anisotropic stresses that enter the tensor equation of motion.

The initial enhancement appears predominantly at comparatively small momenta, reflecting the long-wavelength structures produced by the tachyonic instability. As the scalar evolution becomes nonlinear, the gravitational-wave spectrum broadens towards larger momenta. This follows the development of shorter-scale structure seen both in the scalar power spectrum and in the real-space snapshots of the previous section. At late times the overall spectral shape changes only mildly, indicating that the main period of gravitational-wave production has ended.

A separate feature is visible in the ultraviolet part of the spectrum. In analogy with the corresponding high-momentum structure of the scalar sector discussed in Sec. 4.3, this region is not used to characterise the physical peak of the signal in the analysis below.

The duration of gravitational-wave production is even clearer from the integrated energy density. Figure 5.6 shows the ratio $\rho_{\text{GW}}/\rho_{\text{tot}}$, obtained from the spectrum as described in Sec. 4.4.




Figure 5.6: Evolution of the total gravitational-wave energy density relative to the total energy density for the benchmark simulation $\mu/H = 1000$. The rapid rise coincides with the tachyonic growth and onset of nonlinear scalar dynamics, while the subsequent flattening indicates that gravitational-wave production has largely ceased.

The gravitational-wave energy fraction remains extremely small during the initial, approximately homogeneous stage and then rises sharply around $\tilde{\eta} \sim 10\text{--}20$. By $\tilde{\eta} \sim 20$, the signal has already reached a substantial fraction of its final value. This identifies the tachyonic growth and the onset of fragmentation as the dominant period of gravitational-wave production. The close timing with the rapid build-up of gradient energy in Fig. 5.4 further supports this interpretation.

Production does not stop completely once the initial instability has saturated. The ratio $\rho_{\text{GW}}/\rho_{\text{tot}}$ continues to increase more slowly during the nonlinear stage, as the fragmented scalar configuration continues to carry time-dependent anisotropic stresses. Eventually the curve approaches an approximately constant value, with only small residual variations at late times. At this point the scalar source has become inefficient and the gravitational-wave background can be regarded as effectively frozen-in apart from the subsequent cosmological redshifting, which will be discussed in Sec. 5.3.

The final spectrum is therefore generated in two qualitatively distinct stages: a short and efficient burst associated with tachyonic amplification and the onset of fragmentation, followed by a longer period of weaker production during the nonlinear relaxation of the scalar field. In the following, we use the late-time spectrum after this saturation to extract its characteristic amplitude and momentum scale.

5.2.2 PEAK AMPLITUDE AND CHARACTERISTIC MOMENTUM

We next extract the peak amplitude and characteristic momentum of the gravitational-wave spectra across the full set of simulations. These are the two quantities that will be needed most directly when mapping the numerical results back onto the physical $U(1)_{B-L}$ model.

Figure 5.7 shows the peak spectral energy density as a function of μ/H . For each simulation, $h^2\Omega_{\text{GW}}^{\text{peak}}$ is taken to be the largest value reached by the physical part of the gravitational-wave spectrum over all stored time steps. The first momentum bin at $k = k_{\text{IR}}$ and the initialisation-dependent ultraviolet peak discussed in Sec. 4.3 are excluded from this procedure.

The numerical consistency tests in Sec. 4.5 showed that the largest variation of the peak amplitude comes from the stochastic realisation of the initial vacuum fluctuations. In particular, the random-seed comparison gave a scatter of approximately 5 – 10% for three separate μ/H benchmarks. We therefore assign a conservative relative uncertainty of 10% to every peak value in the main simulation set. These uncertainties are used as weights in the fit below.

The data follow a remarkably clean power law over the entire simulated range,

$$h^2\Omega_{\text{GW}}^{\text{peak}} = c \left(\frac{\mu}{H} \right)^s, \quad (5.8)$$




Figure 5.7: Peak gravitational-wave spectral energy density as a function of μ/H . A relative uncertainty of 10% is assigned to each simulation point, motivated by the random-realisation dependence found in Sec. 4.5. The solid line shows the best-fitting power law $h^2 \Omega_{\text{GW}}^{\text{peak}} = c(\mu/H)^s$.

with

$$s = -2.004 \pm 0.005, \quad c = 1.479 \pm 0.105. \quad (5.9)$$

The uncertainties quoted here are the formal parameter uncertainties returned by the fitting procedure. They should therefore be distinguished from the approximately 10% realisation uncertainty of an individual simulation identified above.

The fitted exponent has a particularly simple interpretation. In Sec. 3.5.2, the analytic estimate gave $\Omega_{\text{GW}}^{\text{peak}} \propto (H/m_{\varphi, \text{QCD}})^2$, up to an efficiency factor and numerical coefficients that could not be fixed analytically. The lattice construction was designed precisely such that

$$\frac{\mu}{H} \longleftrightarrow \frac{m_{\varphi, \text{QCD}}}{H},$$

as discussed in Sec. 4.2.3. The analytic expectation therefore translates directly into $h^2 \Omega_{\text{GW}}^{\text{peak}} \propto (\mu/H)^{-2}$. The result $s = -2.004 \pm 0.005$ confirms this scaling with very high accuracy. The

simulations thus provide the numerical normalisation that was left undetermined by the order-of-magnitude estimate, giving approximately

$$h^2 \Omega_{\text{GW}}^{\text{peak}} \simeq 1.5 \left(\frac{H}{\mu} \right)^2. \quad (5.10)$$

Including the assigned 10% uncertainties gives $\chi^2/\text{dof} = 0.37$. The small value is consistent with the very smooth scaling visible in Fig. 5.7 together with the deliberately conservative uncertainty estimate. Since the 10% errors represent a characteristic realisation-to-realisation scatter rather than independent statistical measurement uncertainties for every point, the goodness-of-fit value should not be over-interpreted.

We now turn to the position of the gravitational-wave peak. Unlike its amplitude, the exact momentum bin containing the maximum can fluctuate slightly at late times. Rather than choosing the peak position at a single output time, we therefore average it over the saturated part of the evolution. We define $\tilde{\eta}_{\text{sat}} = 120$, from which point onwards the total gravitational-wave energy density changes by less than 2%, and use $\tilde{\eta}_{\text{end}} = 150$ as the final time of the standard simulations. For each benchmark we then calculate

$$\left\langle \left(\frac{k_{\text{phys}}}{\mu} \right)_{\text{peak}} \right\rangle_{\tilde{\eta} \in [\tilde{\eta}_{\text{sat}}, \tilde{\eta}_{\text{end}}]}. \quad (5.11)$$

The uncertainty assigned to each averaged value combines the temporal variation of the peak bin with the finite momentum resolution,

$$\sigma_{\text{peak}} = \sqrt{\sigma_{\text{time}}^2 + \sigma_{\text{bin}}^2}, \quad \sigma_{\text{bin}} = 0.05. \quad (5.12)$$

Here, σ_{time} is the standard deviation of $(k_{\text{phys}}/\mu)_{\text{peak}}$ over the averaging interval. The constant binning uncertainty corresponds to half of the momentum spacing $\Delta(k_{\text{phys}}/\mu) \simeq 0.1$ used in the main simulations.

In contrast to the amplitude, the characteristic momentum shows no clear dependence on μ/H . A constant fit to the points in Fig. 5.8, including the uncertainties described above, gives

$$\left\langle \left(\frac{k_{\text{phys}}}{\mu} \right)_{\text{peak}} \right\rangle_{\tilde{\eta} \in [\tilde{\eta}_{\text{sat}}, \tilde{\eta}_{\text{end}}]} = 1.30 \pm 0.05. \quad (5.13)$$

This confirms the analytic expectation of Sec. 3.5.1 that the characteristic gravitational-wave momentum is set by the tachyonic mass scale, up to a numerical factor of order unity. Through the matching in Sec. 4.2.3, the corresponding physical result is therefore $k_{\text{phys}}^{\text{GW}} \sim m_{\varphi, \text{QCD}}$.




Figure 5.8: Late-time characteristic momentum of the gravitational-wave spectrum as a function of μ/H . For each simulation, $(k_{\text{phys}}/\mu)_{\text{peak}}$ is averaged over $\tilde{\eta} \in [120, 150]$. The error bars combine the temporal variation of the peak with the finite momentum resolution. The solid line shows the best-fitting constant.

It is also noteworthy that the gravitational-wave peak lies slightly above $k_{\text{phys}}/\mu = 1$, even though μ sets the boundary of the initial scalar tachyonic band. There is no requirement for these scales to coincide exactly. The tensor source is quadratic in the scalar gradients, so a gravitational-wave mode receives contributions from convolutions of scalar modes at different momenta. Together with the subsequent nonlinear redistribution of scalar power, this naturally allows the characteristic gravitational-wave momentum to be somewhat larger than the scale delimiting the original instability. The gravitational-wave peak is therefore not expected to coincide with either the most strongly amplified scalar mode or the boundary of the original instability band. The measured value $(k_{\text{phys}}/\mu)_{\text{peak}} \simeq 1.3$ is instead a natural consequence of converting a scalar source concentrated at momenta of order μ into a tensor source that is quadratic in the scalar field.

The two fits therefore give a compact numerical characterisation of the signal. The peak amplitude follows the inverse-square scaling predicted analytically, $h^2 \Omega_{\text{GW}}^{\text{peak}} \propto (H/\mu)^2$, while its characteristic momentum remains fixed at approximately $(k_{\text{phys}}/\mu)_{\text{peak}} \simeq 1.3$. These relations provide the main input for mapping the lattice spectra onto the physical $U(1)_{B-L}$ model in the following section.

5.2.3 SPECTRAL FIT

The previous subsection showed that the peak amplitude depends strongly on μ/H , while the characteristic momentum remains approximately fixed in units of μ . We now ask whether the shape of the spectrum is similarly universal.

Figure 5.9 compares the final gravitational-wave spectra of all seven simulations at $\tilde{\eta} = \tilde{\eta}_{\text{end}} = 150$, after normalising each spectrum by its own peak amplitude.




Figure 5.9: Final gravitational-wave spectra for all simulated values of μ/H at $\tilde{\eta} = \tilde{\eta}_{\text{end}}$, normalised by their respective peak amplitudes. The spectra follow an approximately common shape throughout the region containing the physical peak, before the initialisation-dependent ultraviolet feature becomes important.

Once the large difference in amplitude is removed, the spectra agree rather closely over the momentum range relevant for the tachyonically generated signal. In particular, the position and width of the broad peak and the subsequent decrease towards larger momenta vary only mildly with μ/H . The agreement breaks down further into the ultraviolet, where the initialisation-dependent feature discussed in Sec. 4.3 becomes important. Since the purpose of the fit is to describe the physical peak rather than this ultraviolet structure, we restrict the analysis to the conservative interval

$$\frac{k_{\text{IR}}}{\mu} < \frac{k_{\text{phys}}}{\mu} < 4. \quad (5.14)$$

To parametrise this common shape, we adopt a template inspired by previous studies of gravitational-wave spectra generated through tachyonic particle production [65–69]. We write

$$\frac{\Omega_{\text{GW}}(\kappa)}{\Omega_{\text{GW}}^{\text{peak}}} = \mathcal{S}(\kappa; \kappa_{\text{peak}}, \gamma, p), \quad \kappa \equiv \frac{k_{\text{phys}}}{\mu}, \quad (5.15)$$

with

$$\mathcal{S} = \frac{1}{\mathcal{N}} \frac{x^p}{1 + x^p \exp[\gamma(x - 1)]}, \quad x = \frac{\kappa}{\kappa_s}. \quad (5.16)$$

The form of the template has a simple interpretation. At sufficiently small momenta it behaves as $\mathcal{S} \propto \kappa^p$, so that p controls the rise towards the peak. At large x , the denominator is dominated by the exponential term and the spectrum falls approximately as $\exp(-\gamma x)$. The parameter γ therefore controls the high-momentum fall-off.

Rather than treating the scale κ_s as an independent fit parameter, we parametrise the template directly in terms of its physical peak position κ_{peak} . For given γ and p , the maximum of the unnormalised shape occurs at $x = x_{\text{m}}$, where

$$x_{\text{m}} = \frac{p+1}{\gamma} W \left[\frac{\gamma}{p+1} \left(\frac{p}{\gamma} e^\gamma \right)^{1/(p+1)} \right], \quad (5.17)$$

and W denotes the principal branch of the Lambert W function. The scale appearing in Eq. (5.16) is then fixed by requiring the maximum to occur at the fitted peak position,

$$\kappa_s = \frac{\kappa_{\text{peak}}}{x_{\text{m}}}. \quad (5.18)$$

Thus, κ_s is only an auxiliary scale: varying κ_{peak} shifts the physical maximum directly, while γ and p determine the shape around it.

The remaining normalisation $\mathcal{N}$ is likewise not fitted independently. It is simply the value of the unnormalised template at its maximum,

$$\mathcal{N} = \frac{x_{\text{m}}^p}{1 + x_{\text{m}}^p \exp[\gamma(x_{\text{m}} - 1)]}, \quad (5.19)$$

so that

$$\mathcal{S}(\kappa_{\text{peak}}) = 1, \quad (5.20)$$

by construction. The template therefore contains only three independent shape parameters,

$$\{\kappa_{\text{peak}}, \gamma, p\}. \quad (5.21)$$

For each simulation, the measured spectrum is first divided by its peak amplitude and only the momentum interval in Eq. (5.14) is retained. The three parameters are then obtained from a nonlinear least-squares fit to the logarithm of the normalised spectrum. The fitting uncertainties quoted below are obtained from the covariance matrix of this procedure.

A representative example is shown in Fig. 5.10 for $\mu/H = 10^8$.




Figure 5.10: Representative fit to the final gravitational-wave spectrum for $\mu/H = 10^8$. The points show the normalised numerical data within the fitting region, while the solid curve denotes the best-fitting spectral template.

For this benchmark, the best-fitting parameters are

$$\kappa_{\text{peak}} = 1.27 \pm 0.05, \quad \gamma = 1.4 \pm 0.3, \quad p = 0.57 \pm 0.07. \quad (5.22)$$

The template captures both the broad maximum and the subsequent decrease of the numerical spectrum well, despite the remaining bin-to-bin fluctuations. It is also reassuring that the fitted $\kappa_{\text{peak}} \simeq 1.27$ agrees closely with the characteristic value $(k_{\text{phys}}/\mu)_{\text{peak}} \simeq 1.3$ obtained independently in the previous subsection. The two methods therefore identify the same characteristic gravitational-wave scale.

Repeating the fit for every simulated value of μ/H gives no clear systematic evolution of the shape parameters. Their means and standard deviations across the simulation set are

$$\bar{\kappa}_{\text{peak}} = 1.21 \pm 0.09, \quad \bar{\gamma} = 1.43 \pm 0.18, \quad \bar{p} = 0.52 \pm 0.05. \quad (5.23)$$

The corresponding relative standard deviations are approximately 7%, 13% and 9%, respectively. These quoted spreads arise solely from the variation of the best-fit values between simulations. They do not include the fitting uncertainty of each individual spectrum. They should therefore be viewed as an estimate of how much the spectral shape changes across the simulated parameter range rather than as a complete uncertainty on the template parameters. The relatively small scatter, together with the absence of a systematic trend with μ/H , motivates using these averaged values as a single description of the spectral shape. Figure 5.11 compares the resulting template with all normalised numerical spectra in the fitting region.




Figure 5.11: Final normalised gravitational-wave spectra for all simulated values of μ/H within the fitting region, compared with the spectral template evaluated using the averaged parameters in Eq. (5.23). Each numerical spectrum and the template are normalised to unity at their respective peaks.

The averaged template provides a good qualitative description of all simulations around the peak. Agreement of the peak heights is imposed by the normalisation, but the width of the maximum and the decrease towards larger momenta are reproduced without any further adjustment. The residual differences are modest compared with the many orders of magnitude by which the absolute gravitational-wave amplitude changes across the simulation set.

Combining this result with the amplitude scaling obtained in the previous subsection gives the approximate factorised form

$$h^2 \Omega_{\text{GW}} \simeq 1.5 \times (H/\mu)^2 \mathcal{S}(\kappa; \bar{\kappa}_{\text{peak}}, \bar{\gamma}, \bar{p}). \quad (5.24)$$

Within the range studied on the lattice, the dependence on μ/H is thus carried predominantly by the overall amplitude, while the spectrum expressed in units of the tachyonic scale retains an approximately universal shape. This provides the spectral template used below when mapping the numerical results onto the physical $U(1)_{B-L}$ model.

5.3 MAPPING TO THE PHYSICAL MODEL

The lattice study has so far been formulated in terms of the surrogate potential and the dimensionless ratio μ/H . The purpose of the present section is to translate these numerical results back to the physical $U(1)_{B-L}$ model. The matching introduced in Sec. 4.2.3 was constructed precisely for this purpose: the curvature scale μ of the surrogate potential represents the QCD-induced curvature $m_{\varphi, \text{QCD}}$ of the physical potential, so that

$$\left(\frac{\mu}{H}\right)_{\text{sim}} \longleftrightarrow \left(\frac{m_{\varphi, \text{QCD}}}{H}\right)_{\text{phys}}. \quad (5.25)$$

The numerical scaling relations obtained in the previous section can therefore be interpreted directly as functions of the physical model parameters.

5.3.1 DEPENDENCE ON THE MODEL PARAMETERS

As discussed in Sec. 3.1, the model is parametrised by the gauge coupling g_{B-L} and gauge-boson mass $m_{Z'}$. For every pair $(g_{B-L}, m_{Z'})$, the symmetry-breaking scale and the remaining couplings are fixed by the matching conditions and subsequently evolved using the RG equations. The QCD-induced curvature then follows from Eq. (3.40), while the Hubble rate during the strongly supercooled stage is determined by the false-vacuum energy through Eq. (3.30). Their ratio $m_{\varphi, \text{QCD}}/H$ is therefore fixed once the model parameters have been specified. The full RG-improved model is used when evaluating this quantity numerically. The value of this ratio across the relevant parameter space is shown in Fig. 5.12. Moving towards larger $m_{Z'}$ or smaller g_{B-L} decreases $m_{\varphi, \text{QCD}}/H$, whereas smaller $m_{Z'}$ or larger g_{B-L} move the characteristic instability scale further inside the Hubble radius. The lattice simulations cover this hierarchy directly through the corresponding range of μ/H demarcated by the red line in Fig. 5.12.

Once $m_{\varphi, \text{QCD}}/H$ is known, the production amplitude follows from the scaling measured in Sec. 5.2.2. Since the fitted exponent agrees extremely well with the inverse-square behaviour anticipated analytically in Sec. 3.5.2, we use the simplified relation

$$h^2 \Omega_{\text{GW}, \star}^{\text{peak}} \simeq 1.5 \left(\frac{H}{m_{\varphi, \text{QCD}}} \right)^2. \quad (5.26)$$




Figure 5.12: Ratio of the QCD-induced curvature scale $m_{\varphi, \text{QCD}}$ to the Hubble parameter in the $(m_{Z'}, g_{B-L})$ plane. The red line corresponds to the smallest value of μ/H directly probed in the simulations. Figure adapted from Ref. [17].

Here and below, the subscript $\star$ denotes quantities at the end of gravitational-wave production. The small difference between the fitted exponent -2.004 and -2 is irrelevant for the parameter-space mapping compared with the numerical uncertainties discussed previously.

The characteristic momentum follows independently from the late-time peak position. The simulations give

$$\left(\frac{k_{\text{phys}, \star}}{m_{\varphi, \text{QCD}}} \right)_{\text{peak}} \simeq 1.3, \quad (5.27)$$

where the matching in Eq. (5.25) has been used. This fixes the order-unity coefficient that was left undetermined in the analytic estimate of Sec. 3.5.1.

The full production spectrum is obtained from the common spectral shape determined in Sec. 5.2.3. We again introduce

$$\kappa \equiv \frac{k_{\text{phys}, \star}}{m_{\varphi, \text{QCD}}}, \quad (5.28)$$

which is simply the physical-model equivalent of $\kappa = k_{\text{phys}}/\mu$ used in the lattice analysis. The spectrum can then be written as

$$\Omega_{\text{GW}, \star}(\kappa) = \Omega_{\text{GW}, \star}^{\text{peak}} \mathcal{S}(\kappa), \quad (5.29)$$

where $\mathcal{S}$ is the template introduced in the previous subsection, using the representative parameters

$$\bar{\kappa}_{\text{peak}} = 1.21, \quad \bar{\gamma} = 1.43, \quad \bar{p} = 0.52. \quad (5.30)$$

The fitted peak position is consistent with the independently extracted value in Eq. (5.27). The latter provides the direct characteristic-momentum estimate, while the averaged fit is used to construct the complete spectrum.

Thus, for each pair $(g_{B-L}, m_{Z'})$, the full RG-improved model fixes $m_{\varphi, \text{QCD}}/H$, and the lattice results then determine the amplitude, characteristic momentum and shape of the gravitational-wave signal at production. No separate simulation is required for each point in the physical parameter space.

5.3.2 REHEATING AND REDSHIFTING

The final step is to evolve the production spectrum to the present day. As already assumed in Sec. 3.5.3, we take reheating to be effectively instantaneous: the false-vacuum energy released during the transition is transferred to the Standard Model and thermalises in less than a Hubble time. The end of gravitational-wave production can therefore be identified with the onset of radiation domination,

$$T_{\star} = T_{\text{rh}}, \quad H_{\star} = H_{\text{rh}}, \quad (5.31)$$

and no intermediate matter-dominated epoch is included. The reheating temperature is calculated from the RG-improved vacuum-energy difference using the relation derived in Sec. 3.5.3.

The general redshift formulae have already been derived in Sec. 2.3.3 and need only be specialised to the present source. Using $k_{\text{phys}, \star} = \kappa m_{\varphi, \text{QCD}}$ in Eq. (2.84) gives

$$f_0(\kappa) = 2.62 \times 10^{-8} \text{ Hz } \kappa \frac{m_{\varphi, \text{QCD}}}{H} \frac{T_{\text{rh}}}{\text{GeV}} \left(\frac{g_{\star}}{100} \right)^{1/6}. \quad (5.32)$$

This relation converts the dimensionless momentum coordinate of the lattice spectrum directly into an observed frequency.

In particular, the peak frequency is

$$f_0^{\text{peak}} = 2.62 \times 10^{-8} \text{ Hz } \kappa_{\text{peak}} \frac{m_{\varphi, \text{QCD}}}{H} \frac{T_{\text{rh}}}{\text{GeV}} \left(\frac{g_{\star}}{100} \right)^{1/6}. \quad (5.33)$$

An important improvement over analytic estimates used in previous work is that the order-unity factor κ_{peak} no longer has to be assumed. It is now determined directly from the nonlinear lattice evolution.

The gravitational-wave amplitude redshifts according to Eq. (2.87). Combining this with the lattice amplitude and spectral template gives

$$h^2 \Omega_{\text{GW},0}(f_0) \simeq 5.58 \times 10^{-5} \left(\frac{100}{g_\star} \right)^{1/3} \left(\frac{H}{m_{\varphi,\text{QCD}}} \right)^2 \mathcal{S}[\kappa(f_0)], \quad (5.34)$$

where

$$\kappa(f_0) = \frac{f_0}{2.62 \times 10^{-8} \text{ Hz}} \frac{H}{m_{\varphi,\text{QCD}}} \frac{\text{GeV}}{T_{\text{rh}}} \left(\frac{100}{g_\star} \right)^{1/6}. \quad (5.35)$$

Equations (5.34) and (5.35) therefore provide the complete map from a point in the $U(1)_{B-L}$ parameter space to its present-day gravitational-wave spectrum.

The leading parameter dependence is particularly transparent if the running of the couplings is neglected. As shown already in Sec. 3.5,

$$\frac{m_{\varphi,\text{QCD}}}{H} \propto g_{B-L} v_{h,\text{QCD}} m_{Z'}^{-3}, \quad T_{\text{rh}} \propto m_{Z'}. \quad (5.36)$$

Together with the lattice scaling, this gives approximately

$$h^2 \Omega_{\text{GW},0}^{\text{peak}} \propto g_{B-L}^{-2} m_{Z'}^6 v_{h,\text{QCD}}^{-2}, \quad (5.37)$$

and

$$f_0^{\text{peak}} \propto g_{B-L} v_{h,\text{QCD}} m_{Z'}^{-2}. \quad (5.38)$$

The inverse dependence of the peak frequency on $m_{Z'}$ may appear counter-intuitive at first. A larger $m_{Z'}$ raises the false-vacuum energy and hence the reheating temperature, which by itself would shift the signal to higher frequencies. At the same time, however, $m_{\varphi,\text{QCD}}/H$ decreases much more rapidly, as $m_{Z'}^{-3}$. The characteristic fluctuation scale therefore moves closer to the Hubble radius, and this effect dominates over the increase in the reheating scale, giving the net behaviour $f_0^{\text{peak}} \propto m_{Z'}^{-2}$.

These simple scalings are useful for understanding the trends across the parameter space, while the numerical spectra presented in the following section use the full RG-improved values of $m_{\varphi,\text{QCD}}/H$ and T_{rh} together with the nonlinear lattice results derived above.

5.4 OBSERVATIONAL PROSPECTS

The present-day spectra obtained in the previous section can now be compared with the sensitivities of future gravitational-wave observatories. We first illustrate the range of pos-

sible signals using a small set of representative benchmark points and then quantify their detectability through signal-to-noise ratios across the physical parameter space.

5.4.1 BENCHMARK SPECTRA

Figures 5.13 and 5.14 show two complementary sets of representative spectra together with the power-law-integrated (PLI) sensitivity curves of several proposed gravitational-wave experiments. The first keeps $g_{B-L} = 10^{-2}$ fixed while varying $m_{Z'}$, whereas the second fixes $m_{Z'} = 200$ TeV and varies the gauge coupling g_{B-L} . The corresponding model parameters and peak properties are collected in Table 5.1.

The spectral curves shown here are deliberately restricted to the region around the peak for which the numerical template was fitted in Sec. 5.2.3. More precisely, only the part of the spectrum lying approximately within one order of magnitude of the peak amplitude is retained. Outside this range the template is set to zero rather than extrapolated into spectral regions that are not constrained by the lattice fit. The same prescription is used in the signal-to-noise ratio calculation below.




Figure 5.13: Present-day gravitational-wave spectra for four benchmark points with $g_{B-L} = 10^{-2}$, compared with the power-law-integrated sensitivities of several proposed gravitational-wave observatories. Only the peak region described by the numerical spectral fit is shown. The template is set to zero outside this range. Increasing $m_{Z'}$ enhances the peak amplitude while shifting the characteristic frequency towards lower values.

The two benchmark sequences illustrate the parameter dependence derived in Sec. 5.3.2. At fixed g_{B-L} and $v_{h,\text{QCD}}$, increasing $m_{Z'}$ decreases $m_{\varphi,\text{QCD}}/H$ and moves the characteristic




Figure 5.14: Present-day gravitational-wave spectra for four benchmark points with $m_{Z'} = 200$ TeV and varying g_{B-L} , compared with the power-law-integrated sensitivities of several proposed gravitational-wave observatories. Only the peak region described by the numerical spectral fit is shown. The template is set to zero outside this range. Increasing g_{B-L} suppresses the peak amplitude while shifting the characteristic frequency towards higher values.

g_{B-L}	$m_{Z'}$ [TeV]	$m_{\varphi, \text{QCD}}/H$	$h^2 \Omega_{\text{GW},0}^{\text{peak}}$	f_0^{peak} [Hz]
0.01	100	1.1×10^4	4.3×10^{-13}	3.25
0.01	200	1.4×10^3	2.7×10^{-11}	0.82
0.01	300	4.3×10^2	3.0×10^{-10}	0.37
0.01	400	1.8×10^2	1.7×10^{-9}	0.21
0.004	200	5.7×10^2	1.7×10^{-10}	0.33
0.006	200	8.5×10^2	7.6×10^{-11}	0.49
0.008	200	1.1×10^3	4.2×10^{-11}	0.65
0.01	200	1.4×10^3	2.7×10^{-11}	0.82

Table 5.1: Model parameters and present-day peak properties of the benchmark spectra shown in Figs. 5.13 and 5.14.

source scale closer to the Hubble radius, making gravitational-wave production substantially more efficient. Between $m_{Z'} = 100$ TeV and 400 TeV, the peak amplitude increases from approximately 4×10^{-13} to 2×10^{-9} , while the peak frequency decreases from about 3.3 Hz to 0.2 Hz.

The complementary behaviour under variation of g_{B-L} is illustrated in Fig. 5.14. At fixed $m_{Z'} = 200$ TeV, increasing the gauge coupling increases $m_{\varphi, \text{QCD}}/H$ and therefore suppresses

the gravitational-wave amplitude. Over the range $g_{B-L} = 4 \times 10^{-3} - 10^{-2}$, the peak amplitude decreases from 1.71×10^{-10} to 2.70×10^{-11} , while the peak frequency shifts from 0.325 Hz to 0.817 Hz. Together, the two benchmark sequences illustrate the approximate parameter scalings

$$h^2 \Omega_{\text{GW},0}^{\text{peak}} \propto g_{B-L}^{-2} m_{Z'}^6, \quad f_0^{\text{peak}} \propto g_{B-L} m_{Z'}^{-2}, \quad (5.39)$$

for fixed $v_{h,\text{QCD}}$.

Note that some of the benchmark sequences lie below $m_{\varphi,\text{QCD}}/H = 1000$ and therefore use an extrapolation of the peak-amplitude scaling beyond the range directly calibrated by the lattice simulations. This limitation is discussed in more detail below in the context of the SNR analysis.

The predicted signals consequently occupy the decihertz-to-few-hertz range. This makes the planned Big Bang Observer (BBO) [70] and DECIGO (DECI-hertz Interferometer Gravitational wave Observatory) [71–73] particularly well matched to the benchmark spectra, while LISA [11, 12, 74] and the Einstein Telescope (ET) [14] probe the lower- and higher-frequency sides of the relevant range, respectively. Since only the reliably fitted peak region is shown, the absence of a plotted signal outside these intervals should not be interpreted as a prediction of a vanishing physical gravitational-wave spectrum. Rather, no assumption is made here about the poorly constrained spectral tails.

The comparison with PLI curves provides a useful first indication of detectability, but it does not make full use of the known spectral shape or the observation time of each experiment. We therefore turn to the signal-to-noise ratio for a more quantitative comparison.

5.4.2 DETECTOR SENSITIVITIES AND SIGNAL-TO-NOISE RATIOS

A PLI sensitivity curve provides a convenient visual benchmark by summarising the sensitivity of an experiment to a family of power-law spectra. The predicted signal in this work is instead a peaked spectrum with a specific frequency dependence. The signal-to-noise ratio (SNR) provides a more direct measure of its detectability by integrating the predicted spectrum against the detector sensitivity over the relevant observing band.

For the present analysis, the SNR is calculated schematically as

$$\rho_{\text{SNR}} = \left[T_{\text{obs}} \int_{f_{\text{min}}}^{f_{\text{max}}} df \left(\frac{h^2 \Omega_{\text{GW},0}(f)}{h^2 \Omega_{\text{sens}}(f)} \right)^2 \right]^{1/2}, \quad (5.40)$$

up to the detector-dependent correlation factor associated with the adopted sensitivity convention. The implementation contains an additional factor of 2 inside the square root for BBO, DECIGO and ET.

For LISA, BBO and DECIGO we assume an observation time of four years and a detection threshold $\rho_{\text{SNR}} = 10$, while for ET we use five years and $\rho_{\text{SNR}} = 5$. The integration ranges are 10^{-7} – 10^{-1} Hz for LISA, 10^{-3} –1 Hz for BBO and DECIGO, and 1–10 Hz for ET [75].

At each point in the $(m_{Z'}, g_{B-L})$ plane, the RG-improved model fixes $m_{\varphi, \text{QCD}}/H$ and the reheating temperature, from which the present-day signal is constructed as in Sec. 5.3.2. Importantly, the SNR is calculated using the same restricted spectral template as in the benchmark comparison above. Only the peak region constrained by the lattice fit contributes to the integral; the template is set to zero outside this range. This avoids introducing an additional extrapolation of the spectral shape into the infrared or ultraviolet tails.

The resulting SNRs are shown in Fig. 5.15. The colour scale shows ρ_{SNR} on a logarithmic scale. Blue curves denote the adopted detectability thresholds, while the red line marks $m_{\varphi, \text{QCD}}/H = 1000$, the smallest value directly simulated in this work.

The four panels show that the predicted signal can be detected by all four experiments over parts of the physically relevant parameter space. BBO and DECIGO again provide particularly broad coverage, reflecting their strong sensitivity in the decihertz range where many of the predicted spectra peak. For BBO, every point for which the restricted template overlaps the detector band already exceeds the adopted SNR threshold, and consequently no separate blue threshold contour is present. DECIGO and ET also reach detectable regions within the part of parameter space directly covered by the lattice simulations. LISA probes a complementary lower-frequency region, but its detectability threshold lies slightly beyond the $m_{\varphi, \text{QCD}}/H = 1000$ contour. A detectable LISA signal therefore requires a modest extrapolation of the numerical amplitude scaling beyond the range directly simulated in this work.

A new feature compared with an SNR calculation based on an unrestricted spectral template is the appearance of white areas inside the otherwise allowed parameter space. These occur when the non-zero part of the fitted template lies entirely outside the frequency interval used for a particular experiment. The integrand in Eq. (5.40) then vanishes and the SNR is exactly zero in the adopted prescription. These white regions should therefore not be interpreted as a prediction that no gravitational-wave power exists at those frequencies. They simply indicate that the part of the spectrum that is sufficiently well constrained by the present lattice fit does not overlap the detector band. In this sense, the SNR analysis deliberately avoids claiming sensitivity based on extrapolated spectral tails.

It is useful to distinguish three regimes in $m_{\varphi, \text{QCD}}/H$. First, for

$$\frac{m_{\varphi, \text{QCD}}}{H} \geq 1000, \quad (5.41)$$




Figure 5.15: Signal-to-noise ratio across the $(m_{Z'}, g_{B-L})$ parameter space for LISA, BBO, DECIGO and ET. The colour scale shows ρ_{SNR} on a logarithmic scale. Blue curves indicate the adopted detection thresholds, $\rho_{\text{SNR}} = 10$ for LISA, BBO and DECIGO and $\rho_{\text{SNR}} = 5$ for ET. No blue contour appears for BBO because all parameter points with non-zero SNR in the displayed region already lie above its detection threshold. The red curve marks $m_{\varphi, \text{QCD}}/H = 1000$, corresponding to the lower edge of the range directly tested by the lattice simulations. In addition to the bubble-percolation and $m_\varphi < 2m_h$ regions, the region $m_{\varphi, \text{QCD}}/H < 1$ is excluded from the analysis.

the amplitude scaling and spectral behaviour are directly supported by the simulation set. The red contour in Fig. 5.15 marks the boundary of this simulation-calibrated region. Within this range, BBO, DECIGO and ET already possess parameter points above their adopted detection thresholds. The LISA threshold instead lies slightly on the extrapolated side of the red contour. Its projected sensitivity is therefore encouraging, but establishing a detectable LISA region requires extending the fitted peak-amplitude scaling somewhat below the smallest value of $m_{\varphi, \text{QCD}}/H$ directly simulated here.

For

$$1 \leq \frac{m_{\varphi, \text{QCD}}}{H} < 1000, \quad (5.42)$$

the figure uses an extrapolation of the inverse-square peak-amplitude scaling and the common spectral template. This extrapolation is already relevant close to the red boundary: in particular, the LISA detection contour is reached only after moving a short distance into this region. The required extension is therefore considerably milder than an extrapolation all the way towards $m_{\varphi, \text{QCD}}/H \sim 1$, but it should nevertheless be distinguished from the parameter space directly calibrated by the lattice simulations. This region is nevertheless physically interesting because decreasing $m_{\varphi, \text{QCD}}/H$ moves the characteristic source scale closer to the Hubble radius and rapidly enhances the predicted gravitational-wave amplitude. At the same time, the uncertainty of the prediction increases. The simulations only establish the scaling down to $m_{\varphi, \text{QCD}}/H = 1000$, while the analytic argument of Sec. 3.5.2 assumes a source well inside the Hubble radius. Predictions between the red line and the $m_{\varphi, \text{QCD}}/H = 1$ boundary should therefore be regarded as indicative rather than quantitatively calibrated.

Finally, we do not consider the region

$$\frac{m_{\varphi, \text{QCD}}}{H} < 1. \quad (5.43)$$

Here the characteristic instability length $m_{\varphi, \text{QCD}}^{-1}$ becomes larger than the Hubble radius, so the sub-Hubble picture underlying both the analytic estimate and the lattice matching no longer applies. Previous analyses have likewise excluded this regime, since the classical-rolling description becomes questionable as the relevant dynamics approach the horizon scale [17, 18]. In addition, the inverse-square relation would otherwise predict an ever-growing gravitational-wave amplitude as $m_{\varphi, \text{QCD}}/H$ is reduced, even though the assumptions used to derive this behaviour have already broken down. Results in this region would therefore have no controlled interpretation within the present framework and are masked in Fig. 5.15.

Taken together, the benchmark spectra and the SNR scan show that the QCD-induced tachyonic transition can produce signals accessible to several future gravitational-wave observatories. In particular, for BBO, DECIGO, and ET, detectable regions are already present within the range directly calibrated by the nonlinear lattice simulations. Extending the prediction towards the strongest signals closer to the Hubble scale, however, will require dedicated simulations at smaller $m_{\varphi, \text{QCD}}/H$ rather than a continued extrapolation of the sub-Hubble scaling.

5.5 UNCERTAINTIES AND LIMITATIONS

The results of this chapter should be interpreted in light of several numerical and physical approximations. The numerical uncertainties around the physical gravitational-wave peak are comparatively small. As shown in Sec. 4.5, the residual dependence on lattice resolution, volume, time step and simulation duration is at the percent level, while the dominant numerical variation comes from the finite stochastic realisation of the initial vacuum fluctuations, which produces a scatter of roughly 10% in the peak amplitude. The initialisation-sensitive ultraviolet part of the spectrum is less robust and has therefore been excluded from the quantitative analysis. Likewise, the spectral template is used only over the momentum range in which it was fitted.

A more important limitation is the use of the surrogate lattice potential. The matching preserves the tachyonic curvature in Hubble units through $\mu/H \leftrightarrow m_{\varphi, \text{QCD}}/H$, but not the full hierarchy of scales or the detailed Coleman–Weinberg shape of the physical potential. The analysis therefore assumes that the nonlinear dynamics relevant for the gravitational-wave peak are controlled predominantly by this matched infrared scale. The clean scaling found across the simulation set supports this assumption, but it cannot demonstrate that the additional scales of the full model have no effect.

The range of the simulations provides a related limitation. The amplitude scaling and common spectral shape are directly tested only for $m_{\varphi, \text{QCD}}/H \geq 10^3$. Predictions below this value rely on extrapolation and become increasingly uncertain as the characteristic source approaches the Hubble scale. The region $m_{\varphi, \text{QCD}}/H < 1$ is not considered, since both the sub-Hubble analytic estimate and the lattice matching cease to provide a controlled description there. Consequently, the detectable LISA region requires a modest extrapolation, while the BBO, DECIGO and ET forecasts already include detectable parameter points within the directly simulated range. Note that the projected detector sensitivities and adopted SNR thresholds introduce an additional observational uncertainty. They should therefore be interpreted as forecasts rather than sharp experimental boundaries.

The initial conditions and field content also simplify the physical system. The scalar is initialised exactly at the unstable origin with vacuum fluctuations, rather than by explicitly evolving the preceding supercooled phase. Moreover, only the radial scalar degree of freedom is simulated. Possible effects from the Z' gauge field, the Higgs sector and their backreaction are therefore absent. The classical-statistical treatment is well motivated for the strongly amplified infrared modes, but not for the vacuum-like ultraviolet modes, providing another reason not to interpret the high-momentum feature physically.

A further uncertainty concerns residual finite-temperature effects during the onset of the transition. The lattice treatment assumes that, once the barrier disappears, the scalar undergoes essentially undamped tachyonic growth. In a thermal environment, however, plasma interactions can generate friction and modify the effective potential, potentially suppressing the instability or allowing bubble nucleation to compete with homogeneous rolling [76]. Whether such effects are important in the present model requires a dedicated finite-temperature real-time treatment and is not captured by the single-field simulations used here.

Finally, the mapping to present-day observables introduces model and cosmological uncertainties. The calculation depends on the RG-improved effective potential and on the representative choice $v_{h,\text{QCD}} = 100 \text{ MeV}$. It also assumes instantaneous reheating into radiation domination. A delayed decay or an intermediate non-radiation-dominated epoch would modify both the frequency and amplitude of the redshifted spectrum.

The main conclusions should therefore be viewed as robust statements about the tachyonic gravitational-wave signal within the simulated regime: a characteristic momentum of order $m_{\varphi,\text{QCD}}$, an approximately universal spectral shape around the peak, and an inverse-square amplitude scaling. Improving the quantitative prediction would require simulations closer to the Hubble scale and, ultimately, a treatment including more of the physical potential, field content and reheating dynamics.

6 CONCLUSIONS AND OUTLOOK

This thesis has investigated the nonlinear dynamics and gravitational-wave production of a QCD-triggered tachyonic phase transition in a strongly supercooled, classically scale-invariant $U(1)_{B-L}$ extension of the Standard Model. In the regime considered here, bubble percolation fails before the thermal barrier disappears. The symmetry-breaking scalar is instead released from the origin and develops a tachyonic instability. While the initial growth of fluctuations can be described perturbatively, their rapid amplification eventually leads to backreaction, mode coupling and fragmentation, requiring a nonlinear real-time treatment.

Using classical-statistical lattice simulations with *CosmoLattice*, this work has followed the nonlinear stage of the transition using the surrogate scalar potential

$$V_{\text{sim}}(\varphi) = -\frac{\mu^2}{2}\varphi^2 + \frac{\lambda}{4}\varphi^4 + V_0, \quad (6.1)$$

which replaces the full radiatively generated potential while preserving the infrared tachyonic dynamics through the curvature scale μ . This avoids the prohibitively large hierarchy of scales in the physical model and permits the nonlinear evolution to be studied systematically as a function of μ/H . The simulations show a clear transition from the initial tachyonic growth to a strongly inhomogeneous scalar configuration. Potential energy is rapidly converted into kinetic and gradient energy, and the resulting anisotropic stresses source gravitational waves. Most of the signal is generated during the onset of fragmentation, followed by a weaker period of production during the subsequent nonlinear evolution.

Across the simulated range, the gravitational-wave signal exhibits a particularly simple dependence on the hierarchy μ/H . The peak amplitude follows

$$h^2\Omega_{\text{GW}}^{\text{peak}} \simeq 1.5 \left(\frac{H}{\mu}\right)^2, \quad (6.2)$$

in excellent agreement with the analytic inverse-square expectation. The characteristic physical momentum is likewise set directly by the tachyonic scale,

$$k_{\text{phys}}^{\text{peak}} \simeq 1.3 \mu. \quad (6.3)$$

After normalising by the peak amplitude and expressing momenta in units of μ , the spectra obtained for different values of μ/H collapse onto an approximately universal shape around the peak. The lattice results can therefore be summarised by a simple spectral template in which the dominant parameter dependence enters through the overall amplitude and characteristic momentum scale.

These dimensionless results can be mapped directly onto the physical $U(1)_{B-L}$ model by matching the curvature of the surrogate potential to the QCD-induced curvature of the physical potential,

$$\left(\frac{\mu}{H}\right)_{\text{sim}} \longleftrightarrow \left(\frac{m_{\varphi, \text{QCD}}}{H}\right)_{\text{phys}}. \quad (6.4)$$

The numerical scaling relations therefore translate into predictions for the peak amplitude, characteristic frequency and spectral shape throughout the physical parameter space. After reheating and cosmological redshifting, the resulting signals extend across the frequency ranges targeted by several planned observatories. In particular, parts of the parameter space are detectable by BBO, DECIGO and ET while remaining entirely within the range directly calibrated by the lattice simulations. LISA probes a complementary lower-frequency region, although its projected detectable region lies slightly beyond the smallest value of $m_{\varphi, \text{QCD}}/H$ simulated here and therefore requires a modest extrapolation of the measured scaling. The results nevertheless demonstrate that the nonlinear tachyonic transition can constitute an observationally relevant gravitational-wave source, while the nonlinear treatment improves previous estimates based only on the linear dynamics.

Several qualifications remain important. The lattice simulations replace the full radiatively generated potential by a simpler surrogate potential matched through the infrared tachyonic scale, and only the radial scalar degree of freedom is evolved. The quantitative predictions are therefore best controlled in the directly simulated regime $m_{\varphi, \text{QCD}}/H \geq 10^3$. Approaching the Hubble scale requires extrapolation precisely where the sub-Hubble approximation becomes less reliable. In addition, the treatment idealises the state of the field at the beginning of the nonlinear evolution and neglects possible finite-temperature damping, additional field interactions and a non-instantaneous reheating history. These approximations do not alter the main numerical conclusions within the simulated regime, but they limit the precision with which the results can presently be extended throughout the physical parameter space.

The most immediate direction for future work is therefore to extend the simulations towards smaller $m_{\varphi, \text{QCD}}/H$. This would directly test whether the inverse-square amplitude scaling and approximately universal spectral shape persist as the characteristic length scale approaches the Hubble radius. Such simulations are particularly important for establishing the LISA

reach without relying on extrapolation and for determining how the signal behaves before entering the regime in which the present sub-Hubble description ceases to apply.

A more complete treatment should also progressively restore the physics omitted from the simplified lattice model. This includes evolving the full effective potential, incorporating the additional scalar and gauge degrees of freedom, and studying whether interactions with the residual thermal plasma modify or suppress the tachyonic instability. Likewise, following the energy transfer into a thermal bath would allow the reheating history and subsequent redshifting of the signal to be determined dynamically rather than imposed through instantaneous reheating. Finally, connecting the nonlinear lattice evolution consistently to the preceding perturbative stage would provide a unified description of gravitational-wave production across both super- and sub-Hubble scales.

More broadly, the results illustrate how strongly supercooled theories can give rise to symmetry-breaking dynamics qualitatively different from conventional bubble-driven first-order phase transitions. In classically scale-invariant models, the unusually shallow scalar potential allows the characteristic instability scale to approach cosmological scales, making tachyonic fragmentation an efficient source of gravitational radiation. The lattice results obtained here provide a quantitative characterisation of this nonlinear signal and a basis for refining its predictions as future gravitational-wave observatories begin to probe the thermal history of the early Universe.

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to Laura Sagunski for giving me the opportunity to complete this thesis in the DMGW group. I am also grateful to all members of the group for their support, helpful discussions, and the welcoming working environment. I would especially like to thank Daniel Schmitt for his guidance, expertise, patience, and valuable input throughout this work.

I would also like to acknowledge the professors and staff at Goethe-University who helped make my time there both enjoyable and rewarding.

I am deeply grateful to my family for the many ways in which they have supported me throughout my studies. In particular, I would like to thank my parents, David and Susanne, and my sister, Rebecca, for their constant encouragement and support.

Finally, I would like to thank Kathi for her unwavering support, understanding, and companionship throughout this thesis and beyond.

ACRONYMS

BBO	Big Bang Observer
BSM	Beyond the Standard Model
DECIGO	DECI-hertz Interferometer Gravitational wave Observatory
ET	Einstein Telescope
FLRW	Friedmann–Lemaître–Robertson–Walker
IR	Infrared
LISA	Laser Interferometer Space Antenna
PLI	Power-Law Integrated
QCD	Quantum Chromodynamics
RG	Renormalisation-Group
SNR	Signal-to-Noise Ratio
TT	Transverse Traceless
UV	Ultraviolet

BIBLIOGRAPHY

- [1] G. Aad et al. “Observation of a New Particle in the Search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC”. In: *Physics Letters B* 716 (2012), pp. 1–29. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020). arXiv: [1207.7214](https://arxiv.org/abs/1207.7214) [hep-ex].
- [2] S. Chatrchyan et al. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”. In: *Physics Letters B* 716 (2012), pp. 30–61. DOI: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021). arXiv: [1207.7235](https://arxiv.org/abs/1207.7235) [hep-ex].
- [3] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Physical Review Letters* 116.6 (2016), p. 061102. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). arXiv: [1602.03837](https://arxiv.org/abs/1602.03837) [gr-qc].
- [4] Y. Fukuda et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Physical Review Letters* 81 (1998), pp. 1562–1567. DOI: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562). arXiv: [hep-ex/9807003](https://arxiv.org/abs/hep-ex/9807003).
- [5] A. D. Sakharov. “Violation of CP Invariance, C Asymmetry, and Baryon Asymmetry of the Universe”. In: *Pisma Zhurnal Eksperimentalnoi i Teoreticheskoi Fiziki* 5 (1967), pp. 32–35. DOI: [10.1070/PU1991v034n05ABEH002497](https://doi.org/10.1070/PU1991v034n05ABEH002497).
- [6] Planck Collaboration. “Planck 2018 Results. VI. Cosmological Parameters”. In: *Astronomy & Astrophysics* 641 (2020), A6. DOI: [10.1051/0004-6361/201833910](https://doi.org/10.1051/0004-6361/201833910). arXiv: [1807.06209](https://arxiv.org/abs/1807.06209) [astro-ph.CO].
- [7] F. Iocco et al. “Primordial Nucleosynthesis: From Precision Cosmology to Fundamental Physics”. In: *Physics Reports* 472 (2009), pp. 1–76. DOI: [10.1016/j.physrep.2009.02.002](https://doi.org/10.1016/j.physrep.2009.02.002). arXiv: [0809.0631](https://arxiv.org/abs/0809.0631) [astro-ph].
- [8] C. Caprini and D. G. Figueroa. “Cosmological Backgrounds of Gravitational Waves”. In: *Classical and Quantum Gravity* 35.16 (July 2018), p. 163001. ISSN: 1361-6382. DOI: [10.1088/1361-6382/aac608](https://doi.org/10.1088/1361-6382/aac608). arXiv: [1801.04268](https://arxiv.org/abs/1801.04268) [astro-ph.CO].
- [9] K. Kajantie et al. “Is There a Hot Electroweak Phase Transition at $m_H \gtrsim m_W$?” In: *Physical Review Letters* 77 (1996), pp. 2887–2890. DOI: [10.1103/PhysRevLett.77.2887](https://doi.org/10.1103/PhysRevLett.77.2887). arXiv: [hep-ph/9605288](https://arxiv.org/abs/hep-ph/9605288).

- [10] Y. Aoki et al. “The Order of the Quantum Chromodynamics Transition Predicted by the Standard Model of Particle Physics”. In: *Nature* 443 (2006), pp. 675–678. DOI: [10.1038/nature05120](https://doi.org/10.1038/nature05120). arXiv: [hep-lat/0611014](https://arxiv.org/abs/hep-lat/0611014).
- [11] P. Amaro-Seoane et al. *Laser Interferometer Space Antenna*. 2017. arXiv: [1702.00786](https://arxiv.org/abs/1702.00786) [[astro-ph.IM](#)].
- [12] P. Auclair et al. “Cosmology with the Laser Interferometer Space Antenna”. In: *Living Reviews in Relativity* 26 (2023), p. 5. DOI: [10.1007/s41114-023-00045-2](https://doi.org/10.1007/s41114-023-00045-2). arXiv: [2204.05434](https://arxiv.org/abs/2204.05434) [[astro-ph.CO](#)].
- [13] M. Punturo et al. “The Einstein Telescope: A Third-Generation Gravitational Wave Observatory”. In: *Classical and Quantum Gravity* 27 (2010), p. 194002. DOI: [10.1088/0264-9381/27/19/194002](https://doi.org/10.1088/0264-9381/27/19/194002).
- [14] B. Sathyaprakash et al. “Scientific Objectives of Einstein Telescope”. In: *Classical and Quantum Gravity* 29.12 (2012), p. 124013. DOI: [10.1088/0264-9381/29/12/124013](https://doi.org/10.1088/0264-9381/29/12/124013). arXiv: [1206.0331](https://arxiv.org/abs/1206.0331) [[gr-qc](#)].
- [15] S. R. Coleman and E. J. Weinberg. “Radiative Corrections as the Origin of Spontaneous Symmetry Breaking”. In: *Physical Review D* 7.6 (1973), pp. 1888–1910. DOI: [10.1103/PhysRevD.7.1888](https://doi.org/10.1103/PhysRevD.7.1888).
- [16] S. Iso, N. Okada, and Y. Orikasa. “Classically Conformal $B - L$ Extended Standard Model”. In: *Physics Letters B* 676.1-3 (2009), pp. 81–87. DOI: [10.1016/j.physletb.2009.04.046](https://doi.org/10.1016/j.physletb.2009.04.046). arXiv: [0902.4050](https://arxiv.org/abs/0902.4050) [[hep-ph](#)].
- [17] D. Schmitt and L. Sagunski. “QCD-Sourced Tachyonic Phase Transition in a Supercooled Universe”. In: *Journal of Cosmology and Astroparticle Physics* 02.02 (2025), p. 075. DOI: [10.1088/1475-7516/2025/02/075](https://doi.org/10.1088/1475-7516/2025/02/075). arXiv: [2409.05851](https://arxiv.org/abs/2409.05851) [[hep-ph](#)].
- [18] T. Caldas Cifuentes. “Scalar-Induced Gravitational Waves from Supercool(ed) Cosmological Phase Transitions”. MA thesis. Frankfurt am Main, Germany: Goethe University Frankfurt, Mar. 5, 2026.
- [19] D. G. Figueroa et al. “CosmoLattice: A Modern Code for Lattice Simulations of Scalar and Gauge Field Dynamics in an Expanding Universe”. In: *Computer Physics Communications* 283 (2023), p. 108586. DOI: [10.1016/j.cpc.2022.108586](https://doi.org/10.1016/j.cpc.2022.108586). arXiv: [2102.01031](https://arxiv.org/abs/2102.01031) [[astro-ph.CO](#)].
- [20] D. G. Figueroa et al. “The Art of Simulating the Early Universe. Part I: Integration Techniques and Canonical Cases”. In: *Journal of Cosmology and Astroparticle Physics* 04 (2021), p. 035. DOI: [10.1088/1475-7516/2021/04/035](https://doi.org/10.1088/1475-7516/2021/04/035). arXiv: [2006.15122](https://arxiv.org/abs/2006.15122) [[astro-ph.CO](#)].

- [21] J. Baeza-Ballesteros et al. “The Art of Simulating the Early Universe. Part II. Non-canonical Cases & Gravitational Waves”. In: *Journal of Cosmology and Astroparticle Physics* 2026.06 (June 2026), p. 087. DOI: [10.1088/1475-7516/2026/06/087](https://doi.org/10.1088/1475-7516/2026/06/087).
- [22] D. Baumann. *Cosmology*. Cambridge University Press, 2022. ISBN: 978-1-108-83807-8. DOI: [10.1017/9781108937092](https://doi.org/10.1017/9781108937092).
- [23] E. W. Kolb and M. S. Turner. *The Early Universe*. Vol. 69. Frontiers in physics. Boulder, Col.: Westview Press, 1994, XLII, 547 S. ISBN: 0201626748.
- [24] P. J. Mohr et al. “CODATA Recommended Values of the Fundamental Physical Constants: 2022”. In: *Reviews of Modern Physics* 97.2 (2025), p. 025002. DOI: [10.1103/RevModPhys.97.025002](https://doi.org/10.1103/RevModPhys.97.025002). arXiv: [2409.03787](https://arxiv.org/abs/2409.03787) [hep-ph].
- [25] M. Laine and A. Vuorinen. *Basics of Thermal Field Theory: A Tutorial on Perturbative Computations*. Vol. 925. Lecture Notes in Physics. Cham: Springer, 2016. DOI: [10.1007/978-3-319-31933-9](https://doi.org/10.1007/978-3-319-31933-9). arXiv: [1701.01554](https://arxiv.org/abs/1701.01554) [hep-ph].
- [26] L. Husdal. “On Effective Degrees of Freedom in the Early Universe”. In: *Galaxies* 4.4 (2016), p. 78. DOI: [10.3390/galaxies4040078](https://doi.org/10.3390/galaxies4040078). arXiv: [1609.04979](https://arxiv.org/abs/1609.04979) [astro-ph.CO].
- [27] D. M. Schmitt. “Gravitational Wave Probes of New Physics: Echoes from a Supercooled Universe”. PhD thesis. Frankfurt am Main, Germany: Johann Wolfgang Goethe-Universität Frankfurt am Main, 2025.
- [28] N. D. Birrell and P. C. W. Davies. *Quantum Fields in Curved Space*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 1982. DOI: [10.1017/CBO9780511622632](https://doi.org/10.1017/CBO9780511622632).
- [29] J. Berges. “Introduction to Nonequilibrium Quantum Field Theory”. In: *AIP Conf. Proc.* 739.1 (2004), pp. 3–62. DOI: [10.1063/1.1843591](https://doi.org/10.1063/1.1843591). arXiv: [hep-ph/0409233](https://arxiv.org/abs/hep-ph/0409233).
- [30] M. A. Amin et al. “Nonperturbative Dynamics of Reheating After Inflation: A Review”. In: *Int. J. Mod. Phys. D* 24.01 (2015), p. 1530003. DOI: [10.1142/S0218271815300037](https://doi.org/10.1142/S0218271815300037). arXiv: [1410.3808](https://arxiv.org/abs/1410.3808) [hep-ph].
- [31] P. Binetruy et al. “Cosmological Backgrounds of Gravitational Waves and eLISA/NGO: Phase Transitions, Cosmic Strings and Other Sources”. In: *JCAP* 06 (2012), p. 027. DOI: [10.1088/1475-7516/2012/06/027](https://doi.org/10.1088/1475-7516/2012/06/027). arXiv: [1201.0983](https://arxiv.org/abs/1201.0983) [gr-qc].
- [32] J.-F. Dufaux et al. “Theory and Numerics of Gravitational Waves from Preheating after Inflation”. In: *Physical Review D* 76 (2007), p. 123517. DOI: [10.1103/PhysRevD.76.123517](https://doi.org/10.1103/PhysRevD.76.123517). arXiv: [0707.0875](https://arxiv.org/abs/0707.0875) [astro-ph].

- [33] J. Garcia-Bellido, D. G. Figueroa, and A. Sastre. “A Gravitational Wave Background from Reheating after Hybrid Inflation”. In: *Phys. Rev. D* 77 (2008), p. 043517. DOI: [10.1103/PhysRevD.77.043517](https://doi.org/10.1103/PhysRevD.77.043517). arXiv: [0707.0839](https://arxiv.org/abs/0707.0839) [hep-ph].
- [34] D. J. Fixsen. “The Temperature of the Cosmic Microwave Background”. In: *The Astrophysical Journal* 707.2 (2009), pp. 916–920. DOI: [10.1088/0004-637X/707/2/916](https://doi.org/10.1088/0004-637X/707/2/916). arXiv: [0911.1955](https://arxiv.org/abs/0911.1955) [astro-ph.CO].
- [35] M. Kamionkowski, A. Kosowsky, and M. S. Turner. “Gravitational Radiation from First-Order Phase Transitions”. In: *Phys. Rev. D* 49 (1994), pp. 2837–2851. DOI: [10.1103/PhysRevD.49.2837](https://doi.org/10.1103/PhysRevD.49.2837). arXiv: [astro-ph/9310044](https://arxiv.org/abs/astro-ph/9310044).
- [36] S. Iso, N. Okada, and Y. Orikasa. “The Minimal $B - L$ Model Naturally Realized at the TeV Scale”. In: *Physical Review D* 80.11 (2009), p. 115007. DOI: [10.1103/PhysRevD.80.115007](https://doi.org/10.1103/PhysRevD.80.115007). arXiv: [0909.0128](https://arxiv.org/abs/0909.0128) [hep-ph].
- [37] R. N. Mohapatra and A. Y. Smirnov. “Neutrino Mass and New Physics”. In: *Annual Review of Nuclear and Particle Science* 56 (2006), pp. 569–628. DOI: [10.1146/annurev.nucl.56.080805.140534](https://doi.org/10.1146/annurev.nucl.56.080805.140534). arXiv: [hep-ph/0603118](https://arxiv.org/abs/hep-ph/0603118).
- [38] ATLAS Collaboration. “Search for New High-Mass Phenomena in the Dilepton Final State Using 36 fb^{-1} of Proton–Proton Collision Data at $\sqrt{s} = 13 \text{ TeV}$ with the ATLAS Detector”. In: *Journal of High Energy Physics* 10 (2017), p. 182. DOI: [10.1007/JHEP10\(2017\)182](https://doi.org/10.1007/JHEP10(2017)182). arXiv: [1707.02424](https://arxiv.org/abs/1707.02424) [hep-ex].
- [39] M. Escudero, S. J. Witte, and N. Rius. “The Dispirited Case of Gauged $U(1)_{B-L}$ Dark Matter”. In: *Journal of High Energy Physics* 08 (2018), p. 190. DOI: [10.1007/JHEP08\(2018\)190](https://doi.org/10.1007/JHEP08(2018)190). arXiv: [1806.02823](https://arxiv.org/abs/1806.02823) [hep-ph].
- [40] R. Jinno and M. Takimoto. “Probing a Classically Conformal $B - L$ Model with Gravitational Waves”. In: *Physical Review D* 95.1 (2017), p. 015020. DOI: [10.1103/PhysRevD.95.015020](https://doi.org/10.1103/PhysRevD.95.015020). arXiv: [1604.05035](https://arxiv.org/abs/1604.05035) [hep-ph].
- [41] C. Marzo, L. Marzola, and V. Vaskonen. “Phase Transition and Vacuum Stability in the Classically Conformal $B - L$ Model”. In: *European Physical Journal C* 79 (2019), p. 601. DOI: [10.1140/epjc/s10052-019-7076-x](https://doi.org/10.1140/epjc/s10052-019-7076-x). arXiv: [1811.11169](https://arxiv.org/abs/1811.11169) [hep-ph].
- [42] J. Ellis, M. Lewicki, and V. Vaskonen. “Updated Predictions for Gravitational Waves Produced in a Strongly Supercooled Phase Transition”. In: *Journal of Cosmology and Astroparticle Physics* 11 (2020), p. 020. DOI: [10.1088/1475-7516/2020/11/020](https://doi.org/10.1088/1475-7516/2020/11/020). arXiv: [2007.15586](https://arxiv.org/abs/2007.15586) [astro-ph.CO].

- [43] D. H. Lyth and E. D. Stewart. “Thermal Inflation and the Moduli Problem”. In: *Physical Review D* 53 (1996), pp. 1784–1798. DOI: [10.1103/PhysRevD.53.1784](https://doi.org/10.1103/PhysRevD.53.1784). arXiv: [hep-ph/9510204](https://arxiv.org/abs/hep-ph/9510204).
- [44] T. Barreiro et al. “Some Aspects of Thermal Inflation: The Finite-Temperature Potential and Topological Defects”. In: *Physical Review D* 54.2 (1996), pp. 1379–1392. DOI: [10.1103/PhysRevD.54.1379](https://doi.org/10.1103/PhysRevD.54.1379). arXiv: [hep-ph/9602263](https://arxiv.org/abs/hep-ph/9602263).
- [45] M. E. Carrington. “The Effective Potential at Finite Temperature in the Standard Model”. In: *Physical Review D* 45 (1992), pp. 2933–2944. DOI: [10.1103/PhysRevD.45.2933](https://doi.org/10.1103/PhysRevD.45.2933).
- [46] P. B. Arnold and O. Espinosa. “The Effective Potential and First-Order Phase Transitions: Beyond Leading Order”. In: *Physical Review D* 47 (1993), pp. 3546–3573. DOI: [10.1103/PhysRevD.47.3546](https://doi.org/10.1103/PhysRevD.47.3546). arXiv: [hep-ph/9212235](https://arxiv.org/abs/hep-ph/9212235). Erratum: Phys. Rev. D 50, 6662 (1994).
- [47] S. Iso, P. D. Serpico, and K. Shimada. “QCD–Electroweak First-Order Phase Transition in a Supercooled Universe”. In: *Physical Review Letters* 119.14 (2017), p. 141301. DOI: [10.1103/PhysRevLett.119.141301](https://doi.org/10.1103/PhysRevLett.119.141301). arXiv: [1704.04955 \[hep-ph\]](https://arxiv.org/abs/1704.04955).
- [48] L. Sagunski, P. Schicho, and D. Schmitt. “Supercool Exit: Gravitational Waves from QCD-Triggered Conformal Symmetry Breaking”. In: *Physical Review D* 107.12 (2023), p. 123512. DOI: [10.1103/PhysRevD.107.123512](https://doi.org/10.1103/PhysRevD.107.123512). arXiv: [2303.02450 \[hep-ph\]](https://arxiv.org/abs/2303.02450).
- [49] J. Braun and H. Gies. “Chiral Phase Boundary of QCD at Finite Temperature”. In: *Journal of High Energy Physics* 06 (2006), p. 024. DOI: [10.1088/1126-6708/2006/06/024](https://doi.org/10.1088/1126-6708/2006/06/024). arXiv: [hep-ph/0602226](https://arxiv.org/abs/hep-ph/0602226).
- [50] E. Witten. “Cosmological Consequences of a Light Higgs Boson”. In: *Nuclear Physics B* 177 (1981), pp. 477–488. DOI: [10.1016/0550-3213\(81\)90182-6](https://doi.org/10.1016/0550-3213(81)90182-6).
- [51] B. von Harling and G. Servant. “QCD-Induced Electroweak Phase Transition”. In: *Journal of High Energy Physics* 01 (2018), p. 159. DOI: [10.1007/JHEP01\(2018\)159](https://doi.org/10.1007/JHEP01(2018)159). arXiv: [1711.11554 \[hep-ph\]](https://arxiv.org/abs/1711.11554).
- [52] D. Bödeker. “Remarks on the QCD–Electroweak Phase Transition in a Supercooled Universe”. In: *Physical Review D* 104.11 (2021), p. L111501. DOI: [10.1103/PhysRevD.104.L111501](https://doi.org/10.1103/PhysRevD.104.L111501). arXiv: [2108.11966 \[hep-ph\]](https://arxiv.org/abs/2108.11966).
- [53] A. D. Linde. “Fate of the False Vacuum at Finite Temperature: Theory and Applications”. In: *Physics Letters B* 100 (1981), pp. 37–40. DOI: [10.1016/0370-2693\(81\)90281-1](https://doi.org/10.1016/0370-2693(81)90281-1).

- [54] A. D. Linde. “Decay of the False Vacuum at Finite Temperature”. In: *Nuclear Physics B* 216 (1983), pp. 421–445. DOI: [10.1016/0550-3213\(83\)90072-X](https://doi.org/10.1016/0550-3213(83)90072-X). Erratum: *Nuclear Physics B* 223, 544 (1983).
- [55] A. H. Guth and E. J. Weinberg. “Cosmological Consequences of a First-Order Phase Transition in the $SU(5)$ Grand Unified Model”. In: *Physical Review D* 23 (1981), pp. 876–885. DOI: [10.1103/PhysRevD.23.876](https://doi.org/10.1103/PhysRevD.23.876).
- [56] J. Ellis, M. Lewicki, and J. M. No. “On the Maximal Strength of a First-Order Electroweak Phase Transition and Its Gravitational Wave Signal”. In: *Journal of Cosmology and Astroparticle Physics* 04 (2019), p. 003. DOI: [10.1088/1475-7516/2019/04/003](https://doi.org/10.1088/1475-7516/2019/04/003). arXiv: [1809.08242](https://arxiv.org/abs/1809.08242) [hep-ph].
- [57] M. Kierkla et al. “Finite-temperature bubble nucleation with shifting scale hierarchies”. In: *Journal of High Energy Physics* 2025.7 (2025), p. 153. DOI: [10.1007/JHEP07\(2025\)153](https://doi.org/10.1007/JHEP07(2025)153). arXiv: [2503.13597](https://arxiv.org/abs/2503.13597) [hep-ph].
- [58] M. Kierkla et al. “Theoretical uncertainties in reconstructing model parameters with gravitational waves from supercooled phase transitions”. In: (2026). arXiv: [2607.18233](https://arxiv.org/abs/2607.18233) [hep-ph].
- [59] S. Kawai and N. Okada. “Reheating Consistency Condition on the Classically Conformal $U(1)_{B-L}$ Higgs Inflation Model”. In: *Physical Review D* 108.1 (2023), p. 015013. DOI: [10.1103/PhysRevD.108.015013](https://doi.org/10.1103/PhysRevD.108.015013). arXiv: [2303.00342](https://arxiv.org/abs/2303.00342) [hep-ph].
- [60] H. Mansour, Y. Gouttenoire, and F. Kahlhoefer. “Can the universe be matter-dominated after a supercooled first-order phase transition?” In: (2026). arXiv: [2607.19469](https://arxiv.org/abs/2607.19469) [hep-ph].
- [61] G. N. Felder et al. “Dynamics of Symmetry Breaking and Tachyonic Preheating”. In: *Physical Review Letters* 87 (2001), p. 011601. DOI: [10.1103/PhysRevLett.87.011601](https://doi.org/10.1103/PhysRevLett.87.011601). arXiv: [hep-ph/0012142](https://arxiv.org/abs/hep-ph/0012142).
- [62] G. N. Felder, L. Kofman, and A. D. Linde. “Tachyonic Instability and Dynamics of Spontaneous Symmetry Breaking”. In: *Physical Review D* 64 (2001), p. 123517. DOI: [10.1103/PhysRevD.64.123517](https://doi.org/10.1103/PhysRevD.64.123517). arXiv: [hep-th/0106179](https://arxiv.org/abs/hep-th/0106179).
- [63] N. Koivunen, E. Tomberg, and H. Veermäe. “The Linear Regime of Tachyonic Preheating”. In: *Journal of Cosmology and Astroparticle Physics* 07 (2022), p. 028. DOI: [10.1088/1475-7516/2022/07/028](https://doi.org/10.1088/1475-7516/2022/07/028). arXiv: [2201.04145](https://arxiv.org/abs/2201.04145) [astro-ph.CO].
- [64] J.-F. Dufaux et al. “Gravity Waves from Tachyonic Preheating after Hybrid Inflation”. In: *Journal of Cosmology and Astroparticle Physics* 03 (2009), p. 001. DOI: [10.1088/1475-7516/2009/03/001](https://doi.org/10.1088/1475-7516/2009/03/001). arXiv: [0812.2917](https://arxiv.org/abs/0812.2917) [astro-ph].

- [65] C. S. Machado et al. “Audible Axions”. In: *Journal of High Energy Physics* 2019.1 (2019), p. 53. DOI: [10.1007/JHEP01\(2019\)053](https://doi.org/10.1007/JHEP01(2019)053). arXiv: [1811.01950](https://arxiv.org/abs/1811.01950) [hep-ph].
- [66] C. S. Machado et al. “Gravitational Wave Probes of Axionlike Particles”. In: *Physical Review D* 102.7 (2020), p. 075033. DOI: [10.1103/PhysRevD.102.075033](https://doi.org/10.1103/PhysRevD.102.075033). arXiv: [1912.01007](https://arxiv.org/abs/1912.01007) [hep-ph].
- [67] C. Gerlach, D. Schmitt, and P. Schwaller. *Audible Axion Magnetogenesis: Linking Intergalactic Magnetic Fields and Gravitational Waves*. 2026. arXiv: [2605.21092](https://arxiv.org/abs/2605.21092) [hep-ph].
- [68] E. Madge et al. “Audible Axions with a Booster: Stochastic Gravitational Waves from Rotating ALPs”. In: *SciPost Physics* 12.5 (2022), p. 171. DOI: [10.21468/SciPostPhys.12.5.171](https://doi.org/10.21468/SciPostPhys.12.5.171). arXiv: [2111.12730](https://arxiv.org/abs/2111.12730) [hep-ph].
- [69] C. Gerlach, D. Schmitt, and P. Schwaller. “Supercooled Audible Axions”. In: *Journal of High Energy Physics* 9 (2025), p. 185. DOI: [10.1007/JHEP09\(2025\)185](https://doi.org/10.1007/JHEP09(2025)185). arXiv: [2504.05386](https://arxiv.org/abs/2504.05386) [hep-ph].
- [70] J. Crowder and N. J. Cornish. “Beyond LISA: Exploring Future Gravitational Wave Missions”. In: *Physical Review D* 72.8 (2005), p. 083005. DOI: [10.1103/PhysRevD.72.083005](https://doi.org/10.1103/PhysRevD.72.083005). arXiv: [gr-qc/0506015](https://arxiv.org/abs/gr-qc/0506015).
- [71] N. Seto, S. Kawamura, and T. Nakamura. “Possibility of Direct Measurement of the Acceleration of the Universe Using 0.1-Hz Band Laser Interferometer Gravitational Wave Antenna in Space”. In: *Physical Review Letters* 87.22 (2001), p. 221103. DOI: [10.1103/PhysRevLett.87.221103](https://doi.org/10.1103/PhysRevLett.87.221103). arXiv: [astro-ph/0108011](https://arxiv.org/abs/astro-ph/0108011).
- [72] S. Sato et al. “The Status of DECIGO”. In: *Journal of Physics: Conference Series* 840.1 (2017), p. 012010. DOI: [10.1088/1742-6596/840/1/012010](https://doi.org/10.1088/1742-6596/840/1/012010).
- [73] S. Kawamura et al. “Current Status of Space Gravitational Wave Antenna DECIGO and B-DECIGO”. In: *Progress of Theoretical and Experimental Physics* 2021.5 (2021), 05A105. DOI: [10.1093/ptep/ptab019](https://doi.org/10.1093/ptep/ptab019). arXiv: [2006.13545](https://arxiv.org/abs/2006.13545) [gr-qc].
- [74] T. Robson, N. J. Cornish, and C. Liu. “The Construction and Use of LISA Sensitivity Curves”. In: *Classical and Quantum Gravity* 36.10 (2019), p. 105011. DOI: [10.1088/1361-6382/ab1101](https://doi.org/10.1088/1361-6382/ab1101). arXiv: [1803.01944](https://arxiv.org/abs/1803.01944) [astro-ph.HE].
- [75] M. Breitbach et al. “Dark, Cold, and Noisy: Constraining Secluded Hidden Sectors with Gravitational Waves”. In: *JCAP* 07 (2019), p. 007. DOI: [10.1088/1475-7516/2019/07/007](https://doi.org/10.1088/1475-7516/2019/07/007). arXiv: [1811.11175](https://arxiv.org/abs/1811.11175) [hep-ph].

Bibliography

- [76] T. P. Dutka, T. H. Jung, and C. S. Shin. “What Happens When Supercooling Is Terminated by Curvature Flipping of the Effective Potential?” In: *Journal of High Energy Physics* 2025.05 (2025), p. 182. doi: [10.1007/JHEP05\(2025\)182](https://doi.org/10.1007/JHEP05(2025)182). arXiv: [2412.15864](https://arxiv.org/abs/2412.15864) [hep-ph].

USE OF AI-TOOLS

ChatGPT (GPT-5.6 Sol, OpenAI) was used as a supportive tool during the writing of this thesis. Its use included assistance with structuring and organising parts of the manuscript, generating ideas and inspiration, and improving wording and readability. All AI-generated suggestions and text were critically reviewed, adapted where necessary, and integrated into the thesis under the author's responsibility. The scientific content, analysis, interpretation of results, and final conclusions remain the responsibility of the author.

EHRENWÖRTLICHE ERKLÄRUNG / STATUTORY DECLARATION

Erklärung nach § 39 (15) Prüfungsordnung 2020 für den BAMA-Studiengang Physik:

Hiermit erkläre ich, dass ich die Arbeit selbstständig und ohne Benutzung anderer als der angegebenen Quellen und Hilfsmittel verfasst habe. Alle Stellen der Arbeit, die wörtlich oder sinngemäß aus Veröffentlichungen oder aus anderen fremden Texten entnommen wurden, sind von mir als solche kenntlich gemacht worden. Ferner erkläre ich, dass die Arbeit nicht - auch nicht auszugsweise - für eine andere Prüfung verwendet wurde.

Morten Will,

Frankfurt am Main, August 27, 2026