Table of Contents

Bounded Set Theory and Physics: A Companion Volume

Working Draft — 2026

Companion to: *The Axiom of Finite Bounds* (AFB paper)

Imports from: AFB Parts III–XIII only. Does not depend on Parts XIV–XV.

This document is a companion to *The Axiom of Finite Bounds* (the main paper, hereafter AFB). It develops the physical applications of Bounded Set Theory in order of logical dependence, building exclusively on the formal core established in Parts III–XIII of the AFB paper. Nothing in this volume modifies or extends the foundations; everything imports from them.

This volume does not import from, depend on, or presuppose Parts XIV–XV of the AFB paper. Those Parts make motivational and structural arguments about the relationship between BST and physics. This volume replaces that motivation with construction.

Governing principles

1. Dependency order is strict. Each Part imports only from earlier Parts of this volume and from the AFB paper (Parts III–XIII only). No forward references.

2. Simple before complex. Classical mechanics before quantum mechanics. Electromagnetism before Yang-Mills. Flat spacetime before curved spacetime. Single particles before many-body systems. Each advanced topic is a generalisation of a simpler one already built.

3. The maturity gradient is stated honestly. Each Part is classified:

- **Tier 1 (Theorem-ready):** The mathematics is essentially complete. What remains is writing the proofs, not discovering them.
- **Tier 2 (Strategy clear, key lemmas open):** The proof architecture is identified and the tools are available from the AFB paper, but specific lemmas still need proving.
- **Tier 3 (Research programme):** The framework is defined and the questions are well-posed, but the core theorems are open.

4. "Defines a finite analogue" comes before "recovers." The Type I–IV recovery taxonomy from the AFB paper (§9.1) governs all claims. Each claim is tagged.

5. The holographic bound is motivation, not theorem.

Part 0: Imports from the AFB Paper — Mathematical Completeness Note

This Part lists every mathematical tool the companion volume uses and traces it to a specific section of the AFB paper. It serves as both a reference interface and a completeness argument: the claim is that the AFB paper, as of its current version, provides all the fundamental mathematical infrastructure required for basic physics. No new axioms, definitions, or foundational extensions are needed. The companion volume assembles existing tools into physical constructions.

0.1 Foundational logic and set theory

Tool	AFB Source	Used by (this volume)
Bounded quantification (BFOL)	Part III, §3.1	Everything
Axiom of Finite Bounds (AFB)	Part III, §3.3	Everything
BST axioms (7 axioms)	Part IV	Everything
Choice as constructive theorem	Theorem 4.1	Statistical mechanics, QM
Bounded Separation	Part IV, §4.2	All subset constructions
Bounded induction (BI-BST)	Part VI, §6.2	Recursive algorithms, inductive proofs

0.2 Number systems

Tool	AFB Source	Used by (this volume)
ℕ_B(k), ℤ_B(k), ℚ_B(k^2)	Part VIII, §§8.1–8.6	Discrete models, counting
ℝ_B(k) (approximate ordered field)	Part VIII, §§8.7, Theorem 8.13	All continuous physics
ℂ_B(k^4) (approxi	Part VIII, §8.8, Theorem 8.16	QM, E&M, gauge theory

Tool	AFB Source	Used by (this volume)
mate field)		
$\mathbb{H}_B(k^8)$ (quaternions)	Part VIII, §8.8.4	Spinor theory, rotations
$\mathbb{O}_B(k^{16})$ (octonions)	Part VIII, §8.8.4	Exceptional groups (conditional)
Rounding function ρ_k	Definition 8.14	All $\mathbb{R}_B(k)$ arithmetic
Approximate field error $O(1/k^2)$	Theorem 8.13, Theorem 8.16	Error tracking in all computations

0.3 Analysis

Tool	AFB Source	Used by (this volume)
Bounded continuity	Part IX, §9.3	Continuous fields and potentials
Bounded differentiation	Part IX, §9.3	All equations of motion
Bounded integration (Riemann sums)	Part IX, §9.3	Actions, partition functions, expectation values
Bounded transcendental functions (exp, sin, cos, ln)	Part IX, §9.3	Quantum phases, Boltzmann factors, rotations
IVT, EVT, Heine-Borel (Type II)	Part IX, §9.3	Existence of extrema, compactness arguments
Bounded DFT, FFT, Parseval, convolution	Part IX, §9.3.2	Wave mechanics, spectral methods, momentum space

Tool	AFB Source	Used by (this volume)
Bounded ODEs (Euler, RK4)	Part IX, §9.3.3, Theorems 9.7–9.8	Newton's laws, Hamilton's equations, Schrödinger
Systems of ODEs	Part IX, §9.3.3.3, Definition 9.4	Multi-particle dynamics, coupled oscillators
Bounded Gronwall inequality	Part IX, §9.3.3.4, Theorem 9.9	Stability analysis, error propagation
Symplectic integrators	Part IX, §9.3.3.4	Hamiltonian dynamics, energy conservation
Bounded PDEs (heat, wave, Poisson)	Part IX, §9.3.3.5, Theorems 9.10–9.11	Diffusion, wave propagation, electrostatics
Discrete Laplacian	Part IX, §9.3.3.5, Definition 9.6	PDEs, lattice field theory
Spectral PDE methods via DFT	Part IX, §9.3.3.5, Theorem 9.12	Efficient PDE solvers
Bounded probability spaces	Part IX, §9.3.4, Definition 9.7	Statistical mechanics, quantum measurement
Expectation, variance, covariance	Part IX, §9.3.4, Definition 9.8	All statistical quantities
Bayes' theorem	Part IX, §9.3.4, Theorem 9.13	Bayesian inference, state update
Chebyshev inequality	Part IX, §9.3.4, Theorem 9.15	Concentration bounds
Bounded Markov chains	Part IX, §9.3.4, Definition 9.10	Stochastic processes, Monte Carlo
Weak law of large numbers	Part IX, §9.3.4, Theorem 9.16	Statistical estimation

Tool	AFB Source	Used by (this volume)
Shannon entropy, KL divergence	Part IX, §9.3.4, Definitions 9.11–9.12	Thermodynamic entropy, information theory
Jensen's inequality	Part IX, §9.3.4, Theorem 9.17	Gibbs' inequality, free energy bounds

0.4 Complex analysis

Tool	AFB Source	Used by (this volume)
k-holomorphic functions	Part X, §10.2	Complex potentials, conformal maps
Bounded contour integration	Part X, §10.3	Residue calculations, scattering amplitudes
Bounded Cauchy theorem	Part X, §10.3	Complex analysis applications
Bounded Dolbeault cohomology	Part X, §10.4	Hodge theory
Preliminary bounded Kähler geometry	Part X, §10.5	Hodge programme (Tier 2)

0.5 Linear algebra and functional analysis

Tool	AFB Source	Used by (this volume)
Finite-dimensional vector spaces over $\mathbb{C}\_B(k^4)$	Part XI, §11.1	All state spaces, configuration spaces
Bounded norms, norm equivalence	Part XI, §11.1, Theorem 11.1	Distance, convergence

Tool	AFB Source	Used by (this volume)
Bounded linear operators (finite matrices)	Part XI, §11.2, Definition 11.3	All operators in QM and field theory
Operator norm (computable)	Part XI, §11.2, Definition 11.5	Operator bounds, perturbation theory
Adjoint, Hermitian, unitary, normal operators	Part XI, §11.2, Definition 11.4	QM observables, time evolution, symmetries
Dual spaces, Hahn-Banach (constructive)	Part XI, §11.3	Functionals, covectors
Spectral theorem (exact diagonalisation)	Part XI, §11.4, Theorem 11.4	Energy levels, measurement outcomes
Weyl eigenvalue stability	Part XI, §11.4, Theorem 11.5	Perturbation theory
Davis-Kahan eigenspace stability	Part XI, §11.4, Theorem 11.6	Spectral gap stability, adiabatic evolution
Finite Hilbert spaces	Part XI, §11.5, Definition 11.8	All quantum state spaces
Inner product, Cauchy-Schwarz	Part XI, §11.5, Theorem 11.7	Quantum amplitudes, overlaps
Gram-Schmidt	Part XI, §11.5, Theorem 11.8	Orthonormal bases, orthogonal polynomials

Tool	AFB Source	Used by (this volume)
Riesz representation	Part XI, §11.5, Theorem 11.9	Bra-ket formalism
Tensor products	Part XI, §11.6, Definition 11.10	Composite quantum systems, tensor fields
Higher tensor products, (r,s)-tensors	Part XI, §11.6, Definitions 11.11–11.15	Metric, curvature, stress-energy tensors
Contraction, index raising/lowering	Part XI, §11.6, Definitions 11.12–11.14	All tensor computations
Symmetric and antisymmetric tensors	Part XI, §11.6, Definition 11.15	Bosons/fermions, differential forms
Matrix exponential (spectral)	Part XI, §11.7, Definition 11.16	Time evolution, Lie groups, flow maps
exp(iH) unitary for Hermitian H	Part XI, §11.7, Theorem 11.11(iv)	Quantum unitarity
Linear ODE solution via matrix exp	Part XI, §11.7, Theorem 11.12	All linear dynamical systems
Wedge product	Part XI, §11.8, Definition 11.17	Differential forms, electromagnetism
Exterior algebra Λ(V*)	Part XI, §11.8, Definition 11.18	Form calculus
Hodge star	Part XI, §11.8, Definition 11.19	Maxwell’s equations, Hodge duality
Connection to cochain	Part XI, §11.8 + Part VII	Gauge theory, cohomology

Tool	AFB Source	Used by (this volume)
complexes ($d^2 = 0$)		

0.6 Representation theory

Tool	AFB Source	Used by (this volume)
Group actions, orbits, stabilisers	Part XIII, §13.1	Symmetry analysis
Orbit-stabiliser, Burnside	Part XIII, §13.1, Theorems 13.1–13.2	Counting distinct states
Finite representations	Part XIII, §13.2, Definition 13.3	Particle classification
Maschke's theorem (constructive)	Part XIII, §13.2, Theorem 13.3	Decomposition of state spaces
Characters, orthogonality	Part XIII, §13.3, Theorem 13.4	Selection rules, multiplicities
Schur's lemma	Part XIII, §13.3, Theorem 13.5	Irreducibility criteria
Irreducible decomposition	Part XIII, §13.4, Theorem 13.6	Particle multiplets, spectral decomposition
Dimension formula	Part XIII, §13.4, Theorem 13.7	Counting irreducibles

0.7 Complexity theory

Tool	AFB Source	Used by (this volume)
BST-P, BST-NP	Part XII, §§12.2–12.4	Computational complexity of physical problems
Polynomial hierarchy	Part XII, §12.3	Complexity classification of verification

0.8 Completeness assessment

The tools listed above collectively provide:

For classical mechanics: Phase spaces as ℝ_B(k)^{2n} (§0.2). Hamilton's equations as ODE systems (§0.3, Definition 9.4). Energy conservation by Gronwall and symplectic integration (§0.3). Variational principles as finite optimisation over finite path spaces (§0.2 + §0.3). Noether's theorem as algebraic manipulation of finite Lagrangians.

For electromagnetism: Exterior algebra for F = dA, dF = 0, d*F = J (§0.5). Discrete Laplacian for Poisson's equation (§0.3). Fourier analysis for plane waves and spectral methods (§0.3). Lorentz group as finite matrix group over ℝ_B(k) (§0.5 + §0.6).

For quantum mechanics: Finite Hilbert spaces (§0.5). Spectral theorem for energy levels (§0.5). Matrix exponential for time evolution (§0.5). Born rule via probability theory (§0.3). Tensor products for composite systems (§0.5). Representation theory for symmetry classification (§0.6).

For statistical mechanics: Finite probability spaces and partition functions (§0.3). Entropy and Boltzmann distribution (§0.3). Markov chains for stochastic dynamics (§0.3). Jensen's inequality for free energy bounds (§0.3).

For gauge theory: Matrix groups SU(N)_B over ℂ_B(k^4) (§0.5 + §0.6). Representations for matter content (§0.6). Exterior algebra for curvature forms (§0.5). Finite sums for partition functions (§0.3).

For gravity: Tensor algebra for metrics and curvature (§0.5). Discrete Laplacian and finite PDEs (§0.3). Regge calculus uses ℝ_B(k) trigonometry, finite Cayley-Menger determinants, and finite sums for the action (§0.2 + §0.3 + §0.5).

For condensed matter: Finite lattice models are immediate applications of finite Hilbert spaces (§0.5), finite sums (§0.3), and spectral theory (§0.5). These are among the strongest applications because the physics is already finite.

No missing foundations identified. Every mathematical tool that basic physics requires traces to a specific definition, theorem, or construction in Parts III–XIII of the AFB paper. The companion volume assembles these tools; it does not need to extend them.

Two categories of material are NOT in the AFB paper and are NOT needed there:

- **Specific physics constructions** (Lagrangians, Hamiltonians, gauge connections, Regge geometries) — these are applications that belong in this companion volume.
- **Specific function families** (Bessel, Legendre, Hermite, spherical harmonics) — these are instances of orthogonal polynomial families constructible by Gram-Schmidt (Theorem 11.8) from specific inner products, and belong in companion sections where specific physics problems require them.

Two items are trivially available but worth noting:

- **The Dirac delta** on a finite grid with spacing h is the Kronecker function δ_h(x) = (1/h) if x = 0, 0 otherwise. No distribution theory needed.
- **Green's functions** for any discrete linear operator L are the columns of the matrix L^{-1}, computed by Part XI finite linear algebra.

0.9 Experimental grounding

The companion volume's claim is not merely that BST *can* construct physics, but that the physics it constructs is the physics that experiments have verified. Each Part opens with a plain-language orientation, a summary of the key experimental evidence, and a bridge explaining why the BST construction suffices to reproduce that evidence. The following table provides an overview.

Part	Physics	Key experimental confirmations	Best precision achieved
I	Classical mechanics	Planetary orbits; spacecraft navigation (Voyager Neptune flyby); conservation laws in collision experiments; normal mode frequencies of mechanical systems	~1 part in 10^9 (Voyager trajectory)
II	Waves, special relativity, electromagnetism	Seismic wave imaging; muon time dilation; GPS relativistic corrections; Hertz's confirmation of EM waves; diffraction patterns; Fourier optics	~38 μs/day (GPS); 20 orders of magnitude in EM frequency
III	Statistical mechanics and thermodynamics	Gas laws (Boyle 1662–); laws of thermodynamics (never violated); Ising model critical exponents; Boltzmann distribution (blackbody radiation, spectroscopy); Monte Carlo methods (Los Alamos 1953–)	Critical exponents to ~0.1%
IV	Quantum mechanics	Hydrogen spectrum (Rydberg series); double-slit interference; Bell inequality violation (Hensen et al. 2015); quantum teleportation (1997); spectral gaps in quantum materials	Hydrogen Lamb shift: 12 significant figures
V	Gauge theory and Yang-Mills	Proton mass by lattice QCD; asymptotic freedom across 3 decades of energy; W and Z boson discovery and masses; Wilson loop area law and confinement; mass gap (numerical evidence)	Lattice QCD hadron masses: sub-percent
VI	Simplicial geometry and	*Indirect:* supports lattice QCD (Part V), quantum Hall effect (Part VIII), Regge calculus (Part VII)	Via Parts V, VII, VIII

Part	Physics topology	Key experimental confirmations	Best precision achieved
VII	Gravity	Mercury perihelion precession; LIGO gravitational waves (2015); EHT black hole shadow (2019); Hulse-Taylor binary pulsar; CMB power spectrum (Planck satellite)	Pulsar orbital decay: 0.2% over 30 years
VIII	Condensed matter	Ising critical temperatures and exponents; magnon dispersion (neutron scattering); Mott metal-insulator transition; integer quantum Hall effect; BCS superconducting gap; phonon spectra of crystals	Quantum Hall quantisation: 1 part in 10^9
IX	Towards the Standard Model	Higgs boson discovery (2012, 125.1 GeV); electron g−2 anomalous magnetic moment; W/Z masses to 0.01%; CKM matrix unitarity; asymptotic freedom of the strong coupling	Electron g−2: 1 part in 10^{12}

The pattern across all Parts is the same: the experimental predictions that have been verified to the highest precision were computed by finite methods — numerical ODE integration, finite matrix diagonalisation, finite sums of Feynman diagrams, lattice Monte Carlo. The infinite-dimensional continuum formulations ($\mathbb{R}$, $L^2(\mathbb{R})$, smooth gauge fields on $\mathbb{R}^4$) are theoretical frameworks from which finite predictions are extracted. BST works directly with the finite predictions. Nothing in the experimental record requires the infinite scaffolding. What it requires is sufficient precision — and BST parameterises precision explicitly through k.

Part I: Classical Mechanics in Bounded Set Theory

Maturity: Tier 1

AFB imports: ℝ_B(k) and approximate field arithmetic (Part VIII, Theorem 8.13); bounded functions (Part VII); bounded difference quotients (Part IX, Definition 9.2); bounded ODE systems (Part IX, Definition 9.4, Theorems 9.7–9.9); symplectic integrators (Part IX, §9.3.3.4); bounded tensor products (Part XI, Definition 11.10)

The physics

Classical mechanics describes how objects move under forces. It is the physics of thrown balls, orbiting planets, vibrating strings, and colliding billiard balls — every system where quantum effects are negligible and speeds are far below light. Its mathematical framework was developed by Newton (1687), reformulated by Lagrange (1788) and Hamilton (1833), and remains the foundation on which all subsequent physics is built.

What experiments confirm

Classical mechanics is not an approximation or a theoretical proposal. It is among the most precisely verified physical theories:

Planetary orbits. The predictions of Newtonian gravity, refined by perturbation theory, tracked the planets to arcsecond precision for over two centuries. The residual discrepancy in Mercury's perihelion precession (43 arcseconds per century) was explained by general relativity — confirming both the precision of the Newtonian computation and the need for its successor.

Spacecraft navigation. The trajectories of interplanetary spacecraft (Voyager, Cassini, New Horizons) are computed by numerical integration of Hamilton's equations — the same ODE systems this Part constructs. The Voyager 2 Neptune flyby (1989) required trajectory accuracy of a few kilometres over 4.4 billion km — a relative precision of about 1 part in 10^9, achieved by finite numerical computation.

Conservation laws. Energy, momentum, and angular momentum conservation — derived here from Noether's theorem — are verified in every collision experiment, from billiard tables to particle accelerators. No violation of these conservation laws has ever been observed in a classical system.

Normal modes. The vibrational frequencies of mechanical systems — bridges, buildings, musical instruments, molecules — are computed by diagonalising finite matrices (the normal-mode analysis of §1.6.3). The agreement between predicted and measured resonance frequencies is routinely better than 1%.

What BST constructs and why it suffices

Every computation listed above is already a finite computation at finite precision. No laboratory has ever measured a position to infinite decimal places or integrated an ODE over a continuous real line. The Voyager trajectory was computed by a Runge-Kutta integrator on a finite grid — not by solving a differential equation on ℝ.

This Part makes the finiteness explicit. Phase spaces are finite sets (ℝ_B(k)^{2nd}). Equations of motion are finite difference equations. Solutions are finite sequences computed by bounded recursion. Conservation laws are exact algebraic identities. The precision parameter k controls the resolution: at k large enough, the BST computation reproduces every classical prediction to any desired accuracy. Nothing in the experimental record requires the infinite real line. What it requires is enough precision — and BST provides that by parameterising precision explicitly.

This Part establishes the pattern that recurs throughout the companion volume: define the finite state space, write the bounded equations of motion, prove the bounded theorems, and identify the recovery type. Every subsequent Part — waves, thermodynamics, quantum mechanics, gauge theory, gravity — follows the same pattern in a richer setting.

1.1 Configuration and Phase Spaces

Definition 1.1 — Bounded configuration space:
A classical system of n particles in d spatial dimensions has configuration space:

$$Q = \mathbb{R}_B(k)^{nd}$$

A point $q = (q_1^1, \ldots, q_1^d, q_2^1, \ldots, q_n^d) \in Q$ specifies the position of every particle.

Q is a finite set.
Cardinality: $|Q| = |\mathbb{R}_B(k)|^{nd} \leq (k^2)^{nd} = k^{2nd}$.

For a single particle in 3 dimensions: $|Q| \leq k^6$.
For 10 particles in 3 dimensions: $|Q| \leq k^{60}$.

Existence condition: $k^{2nd} \leq n_M$.

Definition 1.2 — Bounded velocity space:
The velocity of a particle is the finite difference quotient of its position trajectory (Definition 9.2 of the AFB paper):

$$v(t) = \delta_h q(t) = (q(t+h) - q(t-h)) / (2h)$$

where h is the time step. Each velocity component is an element of $\mathbb{R}_B(k)$, computed in $\mathbb{Q}_B(k^4)$ and rounded by ρ_k.

The velocity space at a point is $\mathbb{R}_B(k)^{nd}$ — the same finite set as configuration space (velocities are real numbers of the same precision as positions).

The tangent bundle — configuration plus velocity — is the product:

$$TQ = Q \times \mathbb{R}_B(k)^{nd} = \mathbb{R}_B(k)^{2nd}$$

Cardinality: $|TQ| \leq k^{4nd}$.

Definition 1.3 — Bounded phase space:
The phase space of the system is:

$$\Gamma = \mathbb{R}_B(k)^{2nd}$$

with coordinates $(q_1, \ldots, q_n^d, p_1, \ldots, p_n^d)$, where the p_i are the canonical momenta (defined from the Lagrangian in Section 1.2).

Γ is a finite set.
Cardinality: $|\Gamma| \leq k^{4nd}$.

A state of the classical system is a single point $(q, p) \in \Gamma$. A trajectory is a finite sequence of states: $\{(q(t_j), p(t_j))\}_{j=0}^{M}$ for time steps $t_j = t_0 + jh$, with $M = \lfloor T/h \rfloor$.

Plain language: The arena of classical mechanics in BST is a finite set of points — positions and momenta at bounded precision. A single particle in 3D has a phase space of at most k^{12} points. The finiteness is not an approximation to an infinite phase space; it is the mathematical content of "positions and momenta at precision $1/k^2$." The infinite phase space $\mathbb{R}^{2nd}$ is the metatheoretic limit of the family $\{\mathbb{R}_B(k)^{2nd}\}$ as k grows through the metatheory.

1.2 Bounded Lagrangian Mechanics

Definition 1.4 — Bounded Lagrangian:

A bounded Lagrangian is a function:

$$L: TQ \to \mathbb{R}_B(k)$$
$$L: (q, v) \mapsto L(q, v)$$

mapping a configuration-velocity pair to a real number. In physics, L typically takes the form:

$$L(q, v) = T(q, v) - V(q)$$

where T is the kinetic energy and V is the potential energy, both functions to $\mathbb{R}_B(k)$.

For n particles of masses $m_1, \ldots, m_n \in \mathbb{R}_B(k)$:

$$T(q, v) = \Sigma_i\ (1/2)\ m_i\ |v_i|^2$$
$$= \Sigma_i\ (1/2)\ m_i\ \Sigma_d\ (v_i^j)^2$$

This is a finite sum of products in $\mathbb{R}_B(k)$, computed in $\mathbb{Q}_B(k^4)$ with rounding error $O(nd/k^2)$.

Definition 1.5 — Bounded discrete action:

For a discrete path $q = (q(t_0), q(t_1), \ldots, q(t_M))$ through configuration space at time steps $t_j = t_0 + jh$, define the velocities:

$$v(t_j) = (q(t_{j+1}) - q(t_{j-1})) / (2h)$$

(central differences, Definition 9.2 of the AFB paper)

The discrete action is the finite sum:

$$S[q] = \Sigma_{j=1}^{M-1} L(q(t_j), v(t_j)) \cdot h$$

This is a function from the finite set of discrete paths (with fixed endpoints $q(t_0)$ and $q(t_M)$) to

$\mathbb{R}_B(k)$. The set of discrete paths is finite: each intermediate point $q(t_j)$ ranges over $Q = \mathbb{R}_B(k)^{nd}$, so the path space has cardinality $|Q|^{M-1}$.

$S[q]$ is a specific element of $\mathbb{R}_B(k)$, computed by a finite sum of M-1 terms. Type I.

Definition 1.6 — Bounded Euler-Lagrange equations:
The discrete Euler-Lagrange equations are the stationarity conditions for $S[q]$ with respect to variation of each interior point $q(t_j)$:

$$\partial S/\partial q^\alpha(t_j) = 0 \quad \text{for } j = 1, \ldots, M-1$$

where the partial derivative is the finite difference quotient of S with respect to the α-th component of $q(t_j)$.

Expanding, this gives (for each component α and each interior time j):

$$\partial L/\partial q^\alpha(t_j) - \Delta_h(\partial L/\partial v^\alpha(t_j)) = 0$$

where Δ_h is the forward difference operator (Definition 9.2) acting on the time index, and $\partial L/\partial q^\alpha$ and $\partial L/\partial v^\alpha$ are finite differences of L with respect to configuration and velocity components.

Explicitly, for $L = T - V$ with $T = \Sigma\, (1/2)m|v|^2$:

$$m^\alpha \cdot (q(t_{j+1}) - 2q(t_j) + q(t_{j-1})) / h^2 = - \partial V/\partial q^\alpha(t_j)$$

This is Newton's second law: mass times acceleration (the second finite difference of position) equals force (the negative gradient of the potential).

Each equation is a finite algebraic relation in $\mathbb{R}_B(k)$. The system of M-1 equations for nd components is a finite system of $(M-1)\cdot nd$ equations in $(M-1)\cdot nd$ unknowns.

Plain language: The Euler-Lagrange equations in BST are Newton's F = ma written in finite-difference form. The second derivative d^2q/dt^2 becomes the second difference $(q(t+h) - 2q(t) + q(t-h))/h^2$, and the force is the finite-difference gradient of the potential. These are exact finite algebraic equations — no limits, no infinitesimals. The classical Euler-Lagrange equations on $\mathbb{R}$ are what this system looks like when $h \to 0$ and the grid becomes continuous, which is a metatheoretic (Type IV) limit.

1.3 Bounded Hamiltonian Mechanics

Definition 1.7 — Bounded Legendre transform:

Given a bounded Lagrangian $L(q, v)$, the canonical momentum conjugate to q^α is:

$$p^\alpha = \partial L/\partial v^\alpha$$

computed as a finite difference of L with respect to the velocity component v^α.

For $L = T - V$ with $T = \Sigma\ (1/2)m|v|^2$:

$p^\alpha = m^\alpha\ v^\alpha$ (exact in $\mathbb{R}$_B(k))

The bounded Hamiltonian is defined by the Legendre transform:

$$H(q, p) = \Sigma^\alpha\ p^\alpha\ v^\alpha - L(q, v)$$

where v is expressed in terms of (q, p) by inverting the relation $p = \partial L/\partial v$. For the standard kinetic energy:

$v = p/m$ (component-wise division in $\mathbb{Q}$_B(k^4))

$$H(q, p) = \Sigma_i\ |p_i|^2/(2m_i) + V(q)$$

This is a function $H: \Gamma \to \mathbb{R}$_B(k) on the bounded phase space. Type I.

Definition 1.8 — Bounded Hamilton’s equations:

Hamilton's equations of motion are the system of $2nd$ first-order ODEs:

$dq^\alpha/dt = \partial H/\partial p^\alpha$ (velocity from momentum)
$dp^\alpha/dt = -\partial H/\partial q^\alpha$ (force from potential)

In BST, this is a system ODE (Definition 9.4 of the AFB paper) with $2nd$ components:

$$y = (q_1, ..., q_{nd}, p_1, ..., p_{nd}) \in \mathbb{R}\_B(k)\^\{2nd\}$$

$$F(t, y) = (\partial H/\partial p_1, ..., \partial H/\partial p_{nd}, -\partial H/\partial q_1, ..., -\partial H/\partial q_{nd})$$

where all partial derivatives are finite difference quotients (Definition 9.2).

Solved by Euler (Theorem 9.7) or RK4 (Theorem 9.8) as a bounded recursion. Each step: $2nd$ function evaluations and $2nd$ updates in $\mathbb{R}$_B(k).

For an explicit H (e.g. $H = |p|^2/(2m) + V(q)$):

$$q(t_{j+1}) = q(t_j) + h \cdot p(t_j)/m$$
$$p(t_{j+1}) = p(t_j) - h \cdot \partial V/\partial q(t_j)$$

This is the symplectic Euler method — the simplest symplectic integrator, which updates positions using old momenta and momenta using old positions.

Recovery type: Type I (exact finite iteration). The trajectory is a specific finite sequence in $\mathbb{R}_B(k)^{2nd}$, computable in $O(M \cdot nd)$ operations.

Definition 1.9 — Bounded symplectic structure:
The phase space $\Gamma = \mathbb{R}_B(k)^{2nd}$ carries a natural symplectic form — the antisymmetric bilinear form:

$$\omega = \Sigma^a \, dp^a \wedge dq^a$$

In the finite setting, ω is a 2nd × 2nd antisymmetric matrix over $\mathbb{R}_B(k)$:

$$\omega_{ij} = \begin{pmatrix} 0 & I_n \\ -I_n & 0 \end{pmatrix}$$

where I_n is the nd × nd identity matrix. This is the standard symplectic matrix J.

ω is non-degenerate: $\det(J) = 1 \neq 0$. ✓
ω is antisymmetric: $J^T = -J$. ✓
ω is closed: $d\omega = 0$ (automatic for a constant matrix on a flat space). ✓

All three properties are exact algebraic statements about a finite matrix. Type I.

1.4 Conservation Laws

Theorem 1.1 — Bounded energy conservation:
Let $H: \Gamma \to \mathbb{R}_B(k)$ be a bounded Hamiltonian that does not depend explicitly on time ($\partial H/\partial t = 0$). Then H is preserved along exact solutions of Hamilton's equations.

Proof (algebraic):

$$\begin{aligned} dH/dt &= \Sigma^a \left(\partial H/\partial q^a \cdot dq^a/dt + \partial H/\partial p^a \cdot dp^a/dt\right) \\ &= \Sigma^a \left(\partial H/\partial q^a \cdot \partial H/\partial p^a + \partial H/\partial p^a \cdot (-\partial H/\partial q^a)\right) \\ &= 0 \end{aligned}$$

In BST: each term in the sum is a product of finite differences in $\mathbb{R}_B(k)$. The cancellation is exact at the algebraic level (each positive term has an equal negative partner). Therefore $H(q(t_j), p(t_j))$ is constant along the trajectory, up to the accumulated rounding error of the finite-difference arithmetic.

For the symplectic Euler method, the energy is preserved to within $O(h)$ over the full time interval (AFB paper, §9.3.3.4). For the Störmer-Verlet method, the energy is preserved to within $O(h^2)$ without secular drift. □

Recovery type: Type I for the algebraic identity. Type II for the numerical preservation bound (exact inequality at each h, with the classical conservation law as the $h \to 0$ family limit).

Theorem 1.2 — Bounded Noether's theorem:

Let $L(q, v)$ be a bounded Lagrangian and let $\varphi_s: Q \to Q$ be a one-parameter family of transformations of configuration space, parametrised by $s \in \mathbb{R}_B(k)$, such that:

$$L(\varphi_s(q), d\varphi_s(q)/ds \cdot v) = L(q, v) \quad \text{for all } s$$

(the Lagrangian is invariant under the transformation).

Define the Noether charge:

$$J = \Sigma^a (\partial L/\partial v^a) \cdot (\partial \varphi_s(q)^a/\partial s)|_{s=0}$$
$$= \Sigma^a p^a \cdot \xi^a(q)$$

where $\xi^a(q) = \partial \varphi_s(q)^a/\partial s|_{s=0}$ is the infinitesimal generator of the symmetry, computed as a finite difference of φ_s at $s = 0$.

Then J is conserved along solutions of the Euler-Lagrange equations:

$$dJ/dt = 0$$

Proof: By the discrete Euler-Lagrange equations and the symmetry condition. The argument is a finite algebraic manipulation — the time derivative of J involves $\partial L/\partial q$ (which equals $\Delta_h(\partial L/\partial v)$ by Euler-Lagrange) and $\partial^2\varphi/\partial s\partial t$ terms (which cancel by the symmetry condition). □

Recovery type: Type I (finite algebraic identity from the symmetry condition and equations of motion).

Important instances:

Time-translation symmetry → energy conservation.
  ($\xi = \partial/\partial t$, $J = H$. Recovers Theorem 1.1.)

Spatial translation symmetry → momentum conservation.
  ($\xi = \partial/\partial q^\alpha$ for direction α, $J = p^\alpha$.)

Rotational symmetry → angular momentum conservation.
  ($\xi = q \times \partial/\partial q$, $J = q \times p$.)

Each conservation law is a specific finite algebraic identity, verified by direct computation from the symmetry of the bounded Lagrangian.

Plain language: Noether's theorem in BST says the same thing it says classically — every continuous symmetry of the Lagrangian yields a conserved quantity — but the proof is a finite algebraic computation rather than an argument involving smooth manifolds and Lie groups. The "continuous" symmetry is replaced by a parametrised family of finite transformations. The conservation law is an exact algebraic identity. The individual instances (energy, momentum, angular momentum) are verified by direct substitution into finite sums.

1.5 The Poisson Bracket and Hamiltonian Algebra

Definition 1.10 — Bounded Poisson bracket:

For functions $f, g: \Gamma \to \mathbb{R}\_B(k)$ on the bounded phase space, the Poisson bracket is:

$$\{f, g\} = \Sigma^\alpha \left(\delta\_h f/\delta q^\alpha \cdot \delta\_h g/\delta p^\alpha - \delta\_h f/\delta p^\alpha \cdot \delta\_h g/\delta q^\alpha\right)$$

where $\delta\_h/\delta q^\alpha$ and $\delta\_h/\delta p^\alpha$ are finite difference quotients (Definition 9.2 of the AFB paper) in the phase-space coordinates.

The Poisson bracket maps two functions on Γ to a function on Γ. It is a finite computation: for each point $(q, p) \in \Gamma$, evaluate 4nd finite differences and form their alternating sum. Type I.

Theorem 1.3 — Properties of the bounded Poisson bracket:

For functions $f, g, h: \Gamma \to \mathbb{R}\_B(k)$ and $a, b \in \mathbb{R}\_B(k)$:

(i) Antisymmetry: $\{f, g\} = -\{g, f\}$
(Exact: swapping f and g negates each term.) ✓

(ii) Bilinearity: $\{af + bg, h\} = a\{f, h\} + b\{g, h\}$
(Exact: finite differences are linear.) ✓

(iii) Leibniz rule: $\{fg, h\} = f\{g, h\} + g\{f, h\}$
(Exact: product rule for finite differences.) ✓

(iv) Jacobi identity:
$\{f, \{g, h\}\} + \{g, \{h, f\}\} + \{h, \{f, g\}\} = 0$

(At the algebraic level, this follows from the structure of the symplectic matrix J. In the finite-difference implementation, the identity holds up to $O(h^2)$ terms from the commutation of mixed finite differences. For exact algebraic functions — polynomials in q, p evaluated exactly in $\mathbb{R}\_B(k)$ — the Jacobi identity is exact.) ✓

(v) Canonical brackets:

$\{q^\alpha, q^\beta\} = 0$
$\{p^\alpha, p^\beta\} = 0$
$\{q^\alpha, p^\beta\} = \delta^{\alpha\beta}$ (Kronecker delta)

These are exact for the coordinate functions because $\delta\_h q^\alpha / \delta q^\beta = \delta^{\alpha\beta}$ exactly. ✓

Recovery type: Type I for properties (i)-(iii) and (v). The Jacobi identity (iv) is Type I for polynomial Hamiltonians and Type III (with $O(h^2)$ error) for general functions.

Theorem 1.4 — Hamilton's equations in Poisson bracket form:
Hamilton's equations can be written:

$dq^\alpha/dt = \{q^\alpha, H\}$
$dp^\alpha/dt = \{p^\alpha, H\}$

and more generally, for any function $f: \Gamma \to \mathbb{R}\_B(k)$:

$df/dt = \{f, H\} + \partial f/\partial t$

Proof: By direct expansion.
$\{q^\alpha, H\} = \Sigma^\beta (\delta q^\alpha/\delta q^\beta \cdot \delta H/\delta p^\beta - \delta q^\alpha/\delta p^\beta \cdot \delta H/\delta q^\beta)$
$= \delta H/\delta p^\alpha$ (only the $\beta = \alpha$ term survives)
$= dq^\alpha/dt$ (by Hamilton's equation). □

Recovery type: Type I. The Poisson bracket form is algebraically equivalent to Hamilton's equations.

1.5.1 The bridge to quantum mechanics

The Poisson bracket has a structural parallel that will be developed in Part IV (Quantum Mechanics) of this volume. In quantum mechanics, observables are Hermitian operators on a finite Hilbert space, and the analogue of the Poisson bracket is the commutator:

$\{f, g\} \leftrightarrow (1/i\hbar)[\hat{f}, \hat{g}] = (1/i\hbar)(\hat{f}\hat{g} - \hat{g}\hat{f})$

The canonical Poisson brackets $\{q^\alpha, p^\beta\} = \delta^{\alpha\beta}$ become the canonical commutation relations $[\hat{q}^\alpha, \hat{p}^\beta] = i\hbar\, \delta^{\alpha\beta}$. Hamilton's equation $df/dt = \{f, H\}$ becomes the Heisenberg equation $d\hat{f}/dt = (1/i\hbar)[\hat{f}, \hat{H}]$.

In BST, both sides of this correspondence are exact finite computations: the Poisson bracket is a finite sum of finite differences (Definition 1.10), and the commutator is a finite matrix multiplication (Part XI of the AFB paper). The correspondence between the two — Dirac's canonical quantisation — is an exact structural map between finite algebraic objects.

This section names the connection but defers its development to Part IV.

1.6 Concrete Examples

1.6.1 The bounded harmonic oscillator

```
System: one particle in one dimension, mass m,
spring constant κ, both in ℝ_B(k).

Configuration space: Q = ℝ_B(k).  (|Q| ≤ k²)
Phase space: Γ = ℝ_B(k)².  (|Γ| ≤ k⁴)

Lagrangian: L(q, v) = (1/2)mv² − (1/2)κq²
Hamiltonian: H(q, p) = p²/(2m) + (1/2)κq²

Hamilton's equations:
  dq/dt = p/m
  dp/dt = −κq

Discrete (symplectic Euler):
  q_{j+1} = q_j + h · p_j/m
  p_{j+1} = p_j − h · κ · q_j

This is a linear recurrence. The solution can also
be written using the matrix exponential (Theorem
11.12 of the AFB paper):

  ⎛q(t)⎞        ⎛ 0     1/m⎞  ⎛q₀⎞
  ⎜    ⎟ = exp(t·⎜          ⎟)·⎜  ⎟
  ⎝p(t)⎠        ⎝−κ     0  ⎠  ⎝p₀⎠

The matrix A = ((0, 1/m), (−κ, 0)) has eigenvalues
±iω where ω = √(κ/m), giving:

  q(t) = q₀ cos(ωt) + (p₀/mω) sin(ωt)
  p(t) = p₀ cos(ωt) − (mωq₀) sin(ωt)

where cos and sin are evaluated in ℝ_B(k) by
bounded transcendental approximation (Part IX
of the AFB paper).

Energy: H = p²/(2m) + κq²/2 is constant along
the trajectory (Theorem 1.1). Verified by direct
substitution: p²cos²(ωt)/(2m) + mω²q₀²sin²(ωt)/2
+ cross terms = H₀ by sin² + cos² = 1. Exact. ✓

Period: T = 2π/ω. The trajectory returns to its
initial state after M = T/h time steps.
```

Type I for the algebraic solution. Type III for the numerical evaluation of sin, cos at precision O(1/k).

1.6.2 The bounded Kepler problem

System: two particles in 3 dimensions with gravitational interaction.

In centre-of-mass coordinates, reduces to one particle of reduced mass μ in a central potential:

$V(r) = -GM\mu/r$

where $r = |q| = \sqrt{q_1^2 + q_2^2 + q_3^2}$, computed in $\mathbb{R}\_B(k)$ via $\mathbb{Q}\_B(k^4)$.

Hamiltonian: $H = |p|^2/(2\mu) - GM\mu/|q|$

Hamilton's equations:

$dq/dt = p/\mu$

$dp/dt = -GM\mu \cdot q/|q|^3$

This is a system of 6 ODEs (Definition 9.4). Solved by RK4 (Theorem 9.8) with step size h.

Conserved quantities:

- Energy $E = H$ (Theorem 1.1)
- Angular momentum $L = q \times p$ (Noether, Theorem 1.2, rotational symmetry)
- Runge-Lenz vector $A = p \times L - GM\mu^2\hat{q}$ (additional symmetry of the 1/r potential)

Each verified by direct computation of $\{J, H\} = 0$ using the bounded Poisson bracket (Definition 1.10).

Orbits: elliptical ($E < 0$), parabolic ($E = 0$), hyperbolic ($E > 0$) — determined by the sign of E in $\mathbb{R}\_B(k)$, which is decidable. The orbital elements (semi-major axis, eccentricity) are algebraic functions of E and $|L|$, computable in $\mathbb{R}\_B(k)$.

Note on the singularity: The potential $V = -GM\mu/r$ is undefined at $r = 0$. In BST, r is an element of $\mathbb{R}\_B(k)$ and cannot be exactly zero unless $q = (0,0,0)$ exactly. The minimum nonzero $|q|$ is $\sim 1/k^2$ (the spacing of $\mathbb{R}\_B(k)$). The potential at the nearest grid point to the origin is $V \approx -GM\mu k^2$ — large but finite. The "singularity" is a finite number. This is an instance of the general BST phenomenon: every classical singularity that arises from a 1/r-type divergence is automatically regularised by the finite grid spacing.

Type I for the discrete trajectory and conserved quantities. Type III for comparison with classical Keplerian orbits.

1.6.3 Bounded coupled oscillators

System: N particles in one dimension, each connected to its neighbours by springs of constant κ.

Configuration: $q = (q_1, \ldots, q_N) \in \mathbb{R}_B(k)^N$
Momentum: $p = (p_1, \ldots, p_N) \in \mathbb{R}_B(k)^N$
Phase space: $\Gamma = \mathbb{R}_B(k)^{2N}$

Hamiltonian:
$H = \Sigma_i\, p_i^2/(2m) + (\kappa/2)\, \Sigma_i\, (q_{i+1} - q_i)^2$

Hamilton's equations:
$dq_i/dt = p_i/m$
$dp_i/dt = \kappa(q_{i+1} - 2q_i + q_{i-1})$

The momentum equation is: mass times acceleration equals the discrete Laplacian of the displacement. This is exactly the discrete wave equation (Theorem 9.11 of the AFB paper) for the special case of a one-dimensional lattice.

Normal modes: The system is linear, so the solution is determined by the eigenvalues and eigenvectors of the N × N coupling matrix:

$K_{ij} = \kappa(2\delta_{ij} - \delta_{i,j+1} - \delta_{i,j-1})$

This is a finite Hermitian matrix (in fact, real symmetric). Its eigenvalues $\omega_1^2, \ldots, \omega_N^2$ are computable by the spectral theorem (Theorem 11.4 of the AFB paper). Each normal mode oscillates independently at frequency ω_a.

The DFT (Definition 9.1 of the AFB paper) provides the normal-mode decomposition directly: the eigenvectors of the nearest-neighbour coupling matrix are the Fourier modes $e^{2\pi i n \alpha/N}$, and the eigenvalues are $\omega_a^2 = (4\kappa/m)\sin^2(\pi\alpha/N)$.

This example connects:
- Hamiltonian mechanics (this Part)
- The discrete Laplacian (AFB §9.3.3.5)
- Fourier analysis (AFB §9.3.2)
- Spectral theory (AFB Part XI)

All exact at fixed N, k. Type I.

Plain language: The coupled oscillator system demonstrates how multiple strands of the AFB mathematical toolkit converge in a single physical application. The Hamiltonian structure gives conservation of energy. The equations of motion are the discrete wave equation. The normal modes are the DFT basis. The frequencies are eigenvalues of a finite matrix. Everything is exact, finite, and computable. This is the prototype for lattice field theory, lattice phonons, and ultimately lattice gauge theory — all of which are elaborations of the same pattern.

1.6.4 The bounded N-body problem

```
System: N particles in 3 dimensions with pairwise
gravitational (or Coulomb) interactions.

Phase space: Γ = ℝ_B(k)^{6N}

Hamiltonian:
H = Σᵢ |pᵢ|²/(2mᵢ) + Σ_{i<j} V(|qᵢ − qⱼ|)

where V(r) = −Gmᵢmⱼ/r (gravity) or V(r) = keᵢeⱼ/r
(Coulomb).

Hamilton's equations: a system of 6N coupled ODEs
(Definition 9.4). Solved by RK4 or symplectic
integrators.

Conserved quantities:
  Total energy H (time-translation symmetry)
  Total momentum P = Σ pᵢ (spatial translation)
  Total angular momentum L = Σ qᵢ × pᵢ (rotation)
  Centre of mass R = Σ mᵢqᵢ / Σ mᵢ (Galilean)

For N = 2: reduces to the Kepler problem (§1.6.2).
For N ≥ 3: no closed-form solution classically, and
none in BST either. But the trajectory is computable
to any desired precision by choosing h small enough.

Computational cost: each RK4 step evaluates O(N²)
pairwise interactions. Over M = T/h steps:
O(M · N²) operations in ℝ_B(k). Finite.

Type I for the discrete trajectory. The classical
N-body problem (the continuous ODE on ℝ^{6N}) is
the Type IV family limit.
```

1.7 What Part I Recovers

The bounded classical mechanics developed in this Part provides the finite mathematical substrate for every computation in classical physics. At each fixed k, the framework delivers:

Exact finite algorithms (Type I) for computing trajectories, evaluating conservation laws, and decomposing systems into normal modes.

Exact algebraic identities (Type I) for energy conservation, Noether's theorem, the canonical Poisson brackets, and the symplectic structure.

Exact numerical bounds (Type II) for energy preservation under symplectic integration, trajectory stability via Gronwall, and approximation error of finite-difference schemes.

The specific recovery status of the main results:

Result	Recovery type	Reference
Phase space as finite set	Type I	Def. 1.3
Lagrangian mechanics	Type I	Defs. 1.4–1.6
Hamiltonian mechanics	Type I	Defs. 1.7–1.8
Symplectic structure	Type I	Def. 1.9
Energy conservation	Type I/II	Thm. 1.1
Noether's theorem	Type I	Thm. 1.2
Poisson bracket	Type I	Def. 1.10, Thm. 1.3
Hamilton's equations (Poisson)	Type I	Thm. 1.4
Harmonic oscillator solution	Type I/III	§1.6.1
Kepler orbits	Type I/III	§1.6.2
Normal mode decomposition	Type I	§1.6.3
N-body trajectories	Type I	§1.6.4
Singularity regularisation	Type I	§1.6.2 (note)

What is NOT recovered in this Part: infinite-time behaviour (ergodicity, Poincaré recurrence as a limit statement), KAM theory (which concerns infinite-time stability of quasi-periodic orbits on an infinite-dimensional function space), and continuous symmetry groups as smooth manifolds (replaced by finite matrix groups over ℝ_B(k)). These are either family-level (Type IV) phenomena or Category C absences.

What is gained: every singularity ($1/r$, $1/r^2$, etc.) is automatically regularised by the finite grid spacing. No renormalisation needed. No special treatment of collisions. The mathematics is finite and the computation terminates.

End of Part I

Part II: Waves, Fields, and Electromagnetism

Maturity: Tier 1

AFB imports: ℝ_B(k), ℂ_B(k^4) (Part VIII); bounded PDEs including the wave equation and discrete Laplacian (Part IX, §9.3.3.5, Theorems 9.10–9.12, Definition 9.6); DFT, FFT, Parseval, convolution (Part IX, §9.3.2, Definition 9.1, Theorems 9.4–9.6); exterior algebra, wedge product, Hodge star (Part XI, §11.8, Definitions 11.17–11.19); tensor algebra and index operations (Part XI, §11.6, Definitions 11.10–11.15); spectral theorem (Part XI, Theorem 11.4); matrix exponential (Part XI, §11.7, Definition 11.16); cochain complexes and $d^2 = 0$ (Part VII + Part XI §11.8)

Volume imports: Part I (classical mechanics, Hamiltonian structure, Poisson brackets)

The physics

This Part covers three layers of physics in order of increasing structure. Waves — the physics of vibrating strings, sound, water ripples, and seismic disturbances — are the simplest distributed systems: what happens when the coupled oscillators of Part I are treated as continuous media. Special relativity is the kinematics of flat spacetime — the framework that unifies space and time and constrains all physics at speeds approaching light. Electromagnetism is the first field theory: it describes light, radio, magnetism, and all electrical phenomena through Maxwell's four equations, and it is the prototype for every gauge theory in modern physics.

What experiments confirm

Wave phenomena. The wave equation governs acoustic, seismic, and electromagnetic propagation. Seismic wave models computed by finite-difference time-stepping on discrete grids locate earthquake epicentres to within kilometres and image the Earth's interior structure. Acoustic engineering designs concert halls using finite-element normal-mode analysis — the same spectral decomposition constructed in §2.1.3.

Special relativity. Time dilation is confirmed by muon lifetime measurements: muons created in the upper atmosphere at 0.998c live ~5× longer in the lab frame than at rest, matching the Lorentz factor $\gamma = 1/\sqrt{(1-v^2/c^2)}$ to within experimental error. GPS satellites correct for relativistic time dilation at the level of ~38 microseconds per day; without the correction, GPS positions would drift by ~10 km/day. The energy-momentum relation $E^2 = p^2c^2 + m^2c^4$ is verified in every particle accelerator collision — the LHC operates at energies where the Lorentz factor γ exceeds 7000.

Electromagnetism. Maxwell's equations, formulated in the 1860s, predicted electromagnetic waves propagating at the speed of light — confirmed by Hertz in 1887. Modern applications span 20 orders of magnitude in frequency: from ELF radio (3 Hz) to gamma rays (10^{20} Hz). Precision tests of QED (the quantum version, built in Parts IV–V) confirm electromagnetic predictions to 1 part in 10^{12}, but the classical Maxwell theory underlying all RF engineering, antenna design, and optical systems is verified at the ~0.1% level across the entire spectrum.

Optics. Diffraction patterns — computed in §2.7.2 as the DFT of the aperture function — match laboratory measurements of single-slit, double-slit, and circular aperture patterns to the resolution of the detector. Fourier optics, where lenses perform optical Fourier transforms, is the operating principle of every modern imaging system from microscopes to lithography steppers.

What BST constructs and why it suffices

Every computation above is already performed on finite grids at finite precision. Seismic wave modelling uses finite-difference time-domain (FDTD) methods on discrete spatial grids — identical in structure to the bounded wave equation (Theorem 9.11 of the AFB paper). GPS relativistic corrections are computed by evaluating $\gamma = 1/\sqrt{(1-v^2/c^2)}$ at specific numerical values — bounded transcendental evaluation in $\mathbb{R}$_B(k). Diffraction patterns are computed by the FFT — the same DFT that Definition 9.1 constructs. Maxwell's equations on a lattice are the finite cochain equations of §2.4. None of these computations requires the infinite real line, the smooth electromagnetic field on $\mathbb{R}^4$, or the completed continuum. They require precision — and BST parameterises precision explicitly through k.

This Part builds the physics in three layers: waves (§§2.1), special relativity (§2.2), and electromagnetism in three formulations — vector calculus (§2.3), exterior calculus (§2.4), and U(1)_B lattice gauge theory (§2.5) — followed by electromagnetic waves

(§2.6) and optics (§2.7). The third formulation is the template that Part V generalises to SU(N).

2.1 Bounded Wave Mechanics

Waves are the simplest distributed physical systems — they are what you get when the coupled oscillators of Part I §1.6.3 are treated as a field rather than a collection of particles. The mathematics is already in the AFB paper: the bounded wave equation (Theorem 9.11) and the DFT (Definition 9.1).

2.1.1 Scalar waves on a bounded grid

Definition 2.1 — Bounded scalar field:
A bounded scalar field on a d-dimensional spatial grid Ω_h (Definition 9.5 of the AFB paper) is a function:

$$u: \Omega_h \to \mathbb{R}_B(k)$$

assigning a real number to each grid point. Since Ω_h is finite ($|\Omega_h| = (L/h + 1)^d$), the field u is a finite array of $|\Omega_h|$ values. It is a single element of $\mathbb{R}_B(k)^{|\Omega_h|}$ — a finite vector.

A time-dependent scalar field is a finite sequence: $\{u_n\}_{n=0}^{M}$ where each $u_n: \Omega_h \to \mathbb{R}_B(k)$.

The bounded wave equation for u is (Theorem 9.11):

$$u_{n+1}(x) = 2u_n(x) - u_{n-1}(x) + c^2\tau^2(\Delta_h u_n)(x)$$

where c is the wave speed, τ the time step, and Δ_h the discrete Laplacian (Definition 9.6).

Stability: $c\tau/h \le 1/\sqrt{d}$ (CFL condition). Type I.

2.1.2 Plane waves and the DFT

Definition 2.2 — Bounded plane wave:
On a periodic grid of N points in one dimension ($x_j = jh$, $j = 0, \ldots, N-1$), the discrete plane wave of wavenumber m is:

$$\varphi_m(x_j) = \exp(2\pi imj/N)$$

This is the m-th basis vector of the DFT (Definition 9.1 of the AFB paper). The set $\{\varphi_0, \ldots, \varphi_{N-1}\}$ forms an orthogonal basis for the space of functions on the grid, with inner product:

$$\langle\varphi_m, \varphi_l\rangle = \Sigma_j \varphi_m(x_j)^* \varphi_l(x_j) = N \delta_{ml}$$

(Parseval orthogonality, Theorem 9.4.)

Dispersion relation:
A discrete plane wave φ_m satisfies the wave equation with angular frequency ω_m determined by the discrete Laplacian eigenvalue:

$$\omega_m^2 = (4c^2/h^2) \sin^2(\pi m/N)$$

This is an algebraic equation in $\mathbb{R}_B(k)$. For small m/N (long wavelengths), $\sin^2(\pi m/N) \approx (\pi m/N)^2$, giving $\omega \approx 2\pi cm/(Nh)$ — the classical dispersion relation $\omega = c|k|$. The deviation at short wavelengths (large m/N) is the lattice dispersion — a physical consequence of the finite grid spacing, not an approximation error.

Type I for the exact discrete dispersion relation.
Type III for the comparison with $\omega = c|k|$.

2.1.3 Superposition, interference, and standing waves

Any bounded scalar field on the periodic grid decomposes exactly into a finite sum of plane waves:

$$u(x_j) = \sum_{m=0}^{N-1} U_m \varphi_m(x_j)$$

where $U_m = (1/N) \sum_j u(x_j) \varphi_m(x_j)^*$ is the DFT of u (Definition 9.1). This is exact (Type I) — the DFT is a bijection between the spatial and frequency representations.

Interference: the superposition of two plane waves φ_m and φ_l produces a field with intensity pattern:

$$|\varphi_m + \varphi_l|^2 = 2 + 2\cos(2\pi(m-l)j/N)$$

This is an exact finite computation. Constructive interference at points where $(m-l)j/N$ is an integer; destructive where it is a half-integer. Type I.

Standing waves: on a grid with fixed boundary conditions ($u = 0$ at endpoints), the normal modes are $\sin(\pi mj/N)$ with $m = 1, \ldots, N-1$. These are the eigenvectors of the discrete Laplacian with Dirichlet boundary conditions — a finite Hermitian matrix (Definition 9.6). Eigenvalues and eigenvectors by the spectral theorem (Theorem 11.4). Type I.

Plain language: Wave mechanics in BST is the physics of finite sums on finite grids. Every wave is a finite superposition of finitely many discrete plane waves. Interference is a finite trigonometric identity. Standing waves are eigenvectors of a finite matrix. The DFT — already in the AFB paper — is the exact tool for all of it. No infinite Fourier integrals, no function spaces, no measure theory.

2.2 Special Relativity

Special relativity is the kinematics of flat spacetime. Its mathematical content is the Lorentz group — a specific group of 4×4 matrices — and the geometry of Minkowski spacetime — a 4-dimensional space with an indefinite metric. Both are naturally finite objects over $\mathbb{R}_B(k)$.

2.2.1 Bounded Minkowski spacetime

Definition 2.3 — Bounded Minkowski space:
Bounded Minkowski spacetime is the set:

$$M = \mathbb{R}_B(k)^4$$

with coordinates $x = (x^0, x^1, x^2, x^3) = (ct, x, y, z)$.

The Minkowski metric is the (0,2)-tensor
(Definition 11.11 of the AFB paper):

$$\eta_{\mu\nu} = \mathrm{diag}(-1, +1, +1, +1)$$

The spacetime interval between events x and y is:

$$\begin{aligned}\Delta s^2 &= \eta_{\mu\nu} (x - y)^\mu (x - y)^\nu \\ &= -c^2\Delta t^2 + \Delta x^2 + \Delta y^2 + \Delta z^2\end{aligned}$$

computed in $\mathbb{R}_B(k)$ via $\mathbb{Q}_B(k^4)$ arithmetic.

Causal structure:
- $\Delta s^2 < 0$: timelike separation (causally connected)
- $\Delta s^2 > 0$: spacelike separation (causally disconnected)
- $\Delta s^2 = 0$: lightlike/null separation

Each is a decidable condition in $\mathbb{R}_B(k)$. Type I.

2.2.2 The bounded Lorentz group

Definition 2.4 — Bounded Lorentz transformation:
A bounded Lorentz transformation is a 4 × 4 matrix
$\Lambda \in M_4(\mathbb{R}_B(k))$ satisfying:

$$\Lambda^T \eta \Lambda = \eta$$

where $\eta = \mathrm{diag}(-1,1,1,1)$ is the Minkowski metric.

The set of all such matrices forms the bounded
Lorentz group $O(1,3)_B$. The proper orthochronous
subgroup $SO^+(1,3)_B$ is the subset with $\det(\Lambda) = +1$
and $\Lambda^0{}_0 > 0$.

Both conditions — the defining equation and the
determinant/sign conditions — are finite algebraic

equations and inequalities over $\mathbb{R}_B(k)$. Membership is decidable. Type I.

Definition 2.5 — Bounded boost:
A Lorentz boost in the x-direction with velocity $v \in \mathbb{R}_B(k)$ (with $|v| < c$) is:

$$\Lambda(v) = \begin{pmatrix} \gamma & -\gamma v/c & 0 & 0 \\ -\gamma v/c & \gamma & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

where $\gamma = 1/\sqrt{(1 - v^2/c^2)}$ is the Lorentz factor, computed in $\mathbb{R}_B(k)$ via $\mathbb{Q}_B(k^4)$ arithmetic (square root by Newton-Raphson, as in the bounded reals construction of AFB Part VIII §8.7).

Verification: $\Lambda^T \eta \Lambda = \eta$ by direct 4×4 matrix multiplication. Exact at the algebraic level; the numerical computation of γ carries the standard $O(1/k)$ precision of bounded transcendental evaluation. Type I/III.

Rapidity parametrisation: defining $\tanh(\varphi) = v/c$, the boost matrix becomes:

$$\Lambda(\varphi) = \exp(\varphi \cdot K_x)$$

where K_x is the boost generator matrix:

$$K_x = \begin{pmatrix} 0 & -1 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

and exp is the matrix exponential (Definition 11.16 of the AFB paper). This connects special relativity to Lie group theory: boosts are generated by exponentiating Lie algebra elements. Type I/III.

Definition 2.6 — Bounded spatial rotation:
A spatial rotation by angle θ about the z-axis is:

$$R(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta & 0 \\ 0 & \sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

This is also expressible as $\exp(\theta \cdot J_z)$ where J_z is the rotation generator. Rotations about x and y axes are analogous.

The six generators $\{K_x, K_y, K_z, J_x, J_y, J_z\}$

span the Lie algebra so(1,3) — the space of 4×4 real antisymmetric matrices (with respect to η). The structure constants are computed by commutator:

$$[J_i, J_j] = \varepsilon_{ijk} J_k \quad \text{(rotation subalgebra)}$$
$$[K_i, K_j] = -\varepsilon_{ijk} J_k \quad \text{(boost-boost gives rotation)}$$
$$[J_i, K_j] = \varepsilon_{ijk} K_k \quad \text{(rotation-boost mixing)}$$

These are exact matrix identities over $\mathbb{R}_B(k)$. Type I.

2.2.3 Relativistic kinematics

Definition 2.7 — Bounded four-vectors:
A four-vector is an element of $\mathbb{R}_B(k)^4$ that transforms under Lorentz transformations as:

$$A'^\mu = \Lambda^\mu{}_\nu A^\nu$$

(matrix-vector multiplication, Definition 11.3 of the AFB paper).

Important four-vectors:

Position: $x^\mu = (ct, x, y, z)$

Four-velocity: $u^\mu = \gamma(c, v_x, v_y, v_z)$
  where $\gamma = 1/\sqrt{(1 - |v|^2/c^2)}$.
  Satisfies $\eta_{\mu\nu} u^\mu u^\nu = -c^2$. ✓

Four-momentum: $p^\mu = m u^\mu = (E/c, p_x, p_y, p_z)$
  where $E = \gamma mc^2$ and $p = \gamma mv$.

The relativistic energy-momentum relation:
  $\eta_{\mu\nu} p^\mu p^\nu = -m^2c^2$
  i.e. $E^2/c^2 - |p|^2 = m^2c^2$
  i.e. $E^2 = |p|^2c^2 + m^2c^4$

This is an exact polynomial identity in $\mathbb{R}_B(k)$. No transcendental functions needed. Type I.

Theorem 2.1 — Bounded relativistic conservation laws:
In a collision or interaction between particles, four-momentum is conserved:

$$\Sigma_i\, p_i{}^\mu \text{ (before)} = \Sigma_f\, p_f{}^\mu \text{ (after)}$$

This is a system of four linear equations in $\mathbb{R}_B(k)$. The energy-momentum relation $E^2 = |p|^2c^2 + m^2c^4$ constrains each particle individually.

Together, conservation and the mass-shell condition determine the kinematics of any bounded collision process as a finite system of algebraic equations over $\mathbb{R}_B(k)$. Type I.

Plain language: Special relativity in BST is the algebra of 4×4 matrices over $\mathbb{R}_B(k)$. Lorentz transformations are finite matrices satisfying $\Lambda^T \eta \Lambda = \eta$. Boosts and rotations are generated by exponentiating 4×4 generators via the matrix exponential. Four-vectors, the energy-momentum relation, and collision kinematics are exact polynomial algebra. The only approximation is in the numerical evaluation of $\gamma = 1/\sqrt{(1-v^2/c^2)}$, which carries the standard $O(1/k)$ precision of bounded transcendental functions. The algebraic structure — the Lorentz group, its Lie algebra, the conservation laws — is exact.

2.3 Electrostatics and Magnetostatics

Before developing the full time-dependent theory, we treat the static cases — Coulomb's law and the magnetic field of steady currents — which require only the discrete Laplacian and finite linear algebra.

2.3.1 Bounded Poisson equation

Theorem 2.2 — Bounded electrostatics:
Given a bounded charge distribution $\rho: \Omega_h \to \mathbb{R}_B(k)$ on a d-dimensional grid, the electrostatic potential φ satisfies the discrete Poisson equation:

$$\Delta_h \varphi = -\rho/\varepsilon_0$$

where Δ_h is the discrete Laplacian (Definition 9.6 of the AFB paper). This is a finite linear system:

$$L \cdot \varphi = -\rho/\varepsilon_0$$

where L is the $|\Omega_h| \times |\Omega_h|$ Laplacian matrix and φ, ρ are vectors in $\mathbb{R}_B(k)^{|\Omega_h|}$.

Solution: $\varphi = -L^{-1} \cdot (\rho/\varepsilon_0)$, where L^{-1} is computed by Gaussian elimination or spectral decomposition (Theorem 11.4). L is symmetric and (with appropriate boundary conditions) invertible. The Green's function of the discrete Laplacian is the matrix L^{-1} — its columns are the potentials of point sources at each grid point.

The electric field is the negative discrete gradient:

$$E(x) = -\nabla_h \varphi(x)$$

where ∇_h is the vector of central differences (Definition 9.2) in each spatial direction:

$$E_i(x) = -\delta_h \varphi / \delta x_i = -(\varphi(x+he_i) - \varphi(x-he_i))/(2h)$$

Each component of E is a finite computation. Type I.

Coulomb's law recovery:
For a point charge q at the origin ($\rho = q \cdot \delta_{0,x}$ where δ is the Kronecker function), the potential is:

$$\varphi(x) = -(q/\varepsilon_0) \cdot (L^{-1})_{x,0}$$

On a 3D grid, $(L^{-1})_{x,0} \approx 1/(4\pi|x|)$ for $|x| \gg h$. The approximation error is $O(h^2/|x|^3)$ — the lattice correction to the Coulomb potential.

Gauss's law: $\Sigma_{\text{faces of box}} E \cdot \hat{n} \cdot A_{face} = Q_{enclosed}/\varepsilon_0$, where the sum is over the faces of a discrete box and $Q_{enclosed}$ is the total charge inside. This follows from the identity $\Sigma \nabla_h \cdot E \cdot h^d = -\Sigma \rho/\varepsilon_0 \cdot h^d$, which is the discrete divergence theorem — an exact finite identity (telescoping sum). Type I.

2.3.2 Bounded magnetostatics

For a steady current density $J: \Omega_h \to \mathbb{R}_B(k)^3$, the vector potential A satisfies:

$$\Delta_h A = -\mu_0 J$$

(component-wise discrete Poisson equations). Solved by $A = -\mu_0 L^{-1} \cdot J$, the same finite linear algebra.

The magnetic field is the discrete curl:

$$B = \nabla_h \times A$$

where the curl is defined by finite differences:

$$(\nabla_h \times A)_i = \varepsilon_{ijk}\, \delta_h A_k / \delta x_j$$

(with ε_{ijk} the Levi-Civita symbol and summation over repeated indices). This is a finite computation at each grid point. Type I.

The divergence-free condition $\nabla_h \cdot B = 0$ follows from $\nabla \cdot (\nabla \times A) = 0$, which is the discrete identity corresponding to $d^2 = 0$ in the exterior calculus (§11.8 of the AFB paper). Exact. Type I.

2.4 The Electromagnetic Field as a 2-Form

The vector calculus formulation of §2.3 works but obscures the geometric structure. The exterior algebra (§11.8 of the AFB paper) provides the natural language.

Definition 2.8 — The electromagnetic field strength:
On a 4-dimensional lattice with coordinates $(x^0, x^1, x^2, x^3) = (ct, x, y, z)$, the electromagnetic field strength is a 2-form:

$$F \in \Lambda^2(V^*) = \Lambda^2(\mathbb{R}\_B(k)^4)$$

(Definition 11.15 of the AFB paper).

In components (a 4×4 antisymmetric matrix):

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$$

where $A = (A_0, A_1, A_2, A_3)$ is the electromagnetic four-potential and $\partial_\mu = \delta_h/\delta x^\mu$ is the finite difference in the μ-th direction.

Explicitly (with $E_i = -\partial_0 A_i - \partial_i A_0$ and $B_i = \varepsilon_{ijk} \partial_j A_k$):

$$F_{\mu\nu} = \begin{pmatrix} 0 & E_x/c & E_y/c & E_z/c \\ -E_x/c & 0 & -B_z & B_y \\ -E_y/c & B_z & 0 & -B_x \\ -E_z/c & -B_y & B_x & 0 \end{pmatrix}$$

This is a finite antisymmetric matrix over $\mathbb{R}_B(k)$ at each lattice point. The entire field F is a finite array: $|\Omega_h|$ grid points × 6 independent components. Type I.

Definition 2.9 — The gauge potential as a 1-form:
The electromagnetic potential is a 1-form:

$$A = A_\mu dx^\mu \in \Lambda^1(V^*)$$

The field strength is its exterior derivative:

$$F = dA$$

where d is the discrete exterior derivative (coboundary operator from Part VII of the AFB paper, identified with the exterior derivative in §11.8).

The identity $d^2 = 0$ gives automatically:

$$dF = d(dA) = 0$$

This is the homogeneous Maxwell equation — the Bianchi identity. In component form it reads:

$$\partial_\mu F_{\nu\rho} + \partial_\nu F_{\rho\mu} + \partial_\rho F_{\mu\nu} = 0$$

which encodes $\nabla \cdot B = 0$ and $\partial B/\partial t + \nabla \times E = 0$.

Both are exact consequences of $d^2 = 0$ — a finite algebraic identity. No dynamical content; pure geometry. Type I.

Theorem 2.3 — Bounded Maxwell's equations:
Maxwell's equations in exterior calculus form are:

(1) $dF = 0$ (homogeneous: Bianchi identity)
(2) $d(*F) = *J$ (inhomogeneous: with source)

where * is the Hodge star (Definition 11.19 of the AFB paper) with respect to the Minkowski metric η, and $J = J_\mu dx^\mu$ is the current 1-form (or more precisely, $*J$ is the current 3-form).

Equation (1) is automatic from $F = dA$.

Equation (2) is the equation of motion. In component form it reads:

$$\partial_\nu F^{\mu\nu} = \mu_0 J^\mu$$

which encodes $\nabla \cdot E = \rho/\varepsilon_0$ and $\nabla \times B - (1/c^2)\partial E/\partial t = \mu_0 J$.

On a finite simplicial lattice, both equations are finite cochain equations:

- F is a 2-cochain (values on 2-cells / plaquettes)
- dF is a 3-cochain (values on 3-cells)
- *F is a 2-cochain (via the Hodge star)
- d(*F) is a 3-cochain
- *J is a 3-cochain

The equations are finite linear relations between finite arrays of field values. Type I.

Note on gauge invariance: The transformation $A \to A + d\lambda$ for any 0-form (function) λ leaves $F = dA$ invariant, since $d(A + d\lambda) = dA + d^2\lambda = dA$. This is gauge invariance — and it is an exact algebraic consequence of $d^2 = 0$. Type I.

Plain language: Maxwell's equations in BST are two finite cochain equations on a finite simplicial lattice. The homogeneous equations ($dF = 0$) are automatic from the definition $F = dA$ and the algebraic identity $d^2 = 0$. The inhomogeneous equations ($d{*F} = J$) relate the field to the source via the Hodge star. Gauge invariance is an exact consequence of $d^2 = 0$. The entire structure — field strength as a 2-form, gauge potential as a 1-form, gauge invariance from exact sequences — transfers verbatim from classical differential geometry to BST, with smooth manifolds replaced by finite simplicial complexes and differential forms replaced by finite cochains.

2.5 U(1)_B Lattice Gauge Theory

The exterior calculus formulation of §2.4 works with the linearised (abelian) theory. The lattice gauge theory formulation works directly with the group-valued connection variables and generalises naturally to non-abelian gauge groups (Part V of this volume).

Definition 2.10 — The bounded U(1) gauge group:

$$U(1)_B = \{z \in \mathbb{C}_B(k^4) : |z|^2 = 1\}$$

More precisely, since $\mathbb{C}_B(k^4)$ is an approximate field, $U(1)_B$ is the set of elements z with $|z|^2 = 1$ to within $O(1/k^2)$ — or equivalently, the set of elements $z = \exp(i\theta)$ for $\theta \in \mathbb{R}_B(k) \cap [0, 2\pi)$, where exp is the bounded exponential.

$U(1)_B$ is a finite set. It is closed under multiplication and complex conjugation (which serves as inverse: $z \cdot z^* = |z|^2 = 1$). Type I.

Definition 2.11 — Lattice U(1) connection:

On a finite oriented lattice (graph) K with vertex set $V(K)$ and edge set $E(K)$:

A U(1) lattice connection is a map:

$$U: E(K) \to U(1)_B$$

assigning a phase to each oriented edge, with the convention $U(e^{-1}) = U(e)^*$ for reverse orientation.

The configuration space is:

$$A_K = U(1)_B^{E(K)}$$

This is a finite set. $|A_K| = |U(1)_B|^{|E(K)|}$.

Definition 2.12 — Plaquette holonomy and Wilson action:

For a plaquette (oriented 2-cell) p with boundary edges e_1, e_2, e_3, e_4, the plaquette holonomy is:

$$U_p = U(e_1) \cdot U(e_2) \cdot U(e_3) \cdot U(e_4)$$

For U(1), this is simply a product of complex numbers in $\mathbb{C}_B(k^4)$. The result $U_p \in U(1)_B$.

The Wilson action is:

$$S[U] = \beta \sum_p \mathrm{Re}(1 - U_p)$$

where the sum is over all plaquettes and $\beta = 1/(e^2)$

is the inverse coupling squared.

$S[U] \geq 0$, with $S = 0$ iff all plaquettes are trivial ($U_p = 1$, i.e. zero field strength). S is a finite sum of real numbers in $\mathbb{R}_B(k)$. Type I.

Connection to the continuum: for a slowly varying field ($|U(e) - 1| \ll 1$), writing $U(e) = \exp(ieA_\mu h)$ and expanding:

$$1 - U_p \approx (1/2)(eF_{\mu\nu} h^2)^2$$

So the Wilson action reduces to the Maxwell action $S \propto \Sigma F_{\mu\nu}^2$ in the smooth-field limit. The lattice action is the finite version; the Maxwell action is the continuum idealisation. Type I for the lattice action; Type IV for the continuum correspondence.

Definition 2.13 — Lattice gauge transformation:
A gauge transformation is a map $g: V(K) \to U(1)_B$ assigning a phase to each vertex.

It acts on the connection by:
$U(e) \to g(\text{source}(e)) \cdot U(e) \cdot g(\text{target}(e))^*$

The plaquette holonomy transforms as:
$U_p \to g(v) \cdot U_p \cdot g(v)^*$

For U(1) (abelian), U_p is invariant since $g(v)\cdot U_p\cdot g(v)^* = U_p$. More generally, $\mathrm{Re}(U_p)$ is gauge-invariant.

The Wilson action S[U] is gauge-invariant. ✓

Theorem 2.4 — U(1) lattice partition function:
The partition function of the U(1) lattice gauge theory is:

$$Z = \Sigma_{U \in A_K} \exp(-S[U])$$

This is a finite sum over the finite configuration space A_K. No functional integral, no regularisation, no infinite-dimensional measure theory.

Expectation values of observables O[U]:

$$\langle O\rangle = (1/Z) \Sigma_{U \in A_K} O[U] \cdot \exp(-S[U])$$

Both Z and ⟨O⟩ are finite rational expressions in $\mathbb{R}_B(k)$. Type I.

Wilson loop observable: for a closed loop γ on the lattice, the Wilson loop is:

$$W(\gamma) = \prod_{e \in \gamma} U(e)$$

For U(1), $W(\gamma)$ is gauge-invariant (abelian case).
$\langle W(\gamma)\rangle$ probes the gauge-field correlations and is the lattice precursor of the continuum Wilson loop that diagnoses confinement. Type I.

Plain language: U(1) lattice gauge theory is the gauge-theoretic formulation of electromagnetism on a finite lattice. The gauge field is a phase on each edge. The field strength is measured by the phase around each plaquette. The action is a finite sum. The partition function is a finite sum. This is the complete finite theory — it is not an approximation to continuum electromagnetism but the finite mathematical content from which the continuum theory is idealised. Part V of this volume will generalise this from U(1) to SU(N), giving the full Yang-Mills theory.

2.6 Electromagnetic Waves

With the time-dependent Maxwell equations (Theorem 2.3) and the wave equation machinery (§2.1), electromagnetic wave propagation is immediate.

Theorem 2.5 — Bounded electromagnetic wave equation:
In a source-free region ($J = 0$), Maxwell's equations imply that each component of the electromagnetic field satisfies the wave equation:

$$\Delta_h E_i - (1/c^2)\, \Delta_\tau^2 E_i = 0$$
$$\Delta_h B_i - (1/c^2)\, \Delta_\tau^2 B_i = 0$$

where Δ_h is the spatial discrete Laplacian and Δ_τ^2 is the second time difference.

Proof: From $d{*}F = 0$ (source-free Maxwell) and $dF = 0$ (Bianchi), apply the discrete Laplacian identity $\Delta = d{*}d + dd{*}$ to each component. The argument is the standard derivation of the wave equation from Maxwell, with all derivatives replaced by finite differences. □

The solutions are electromagnetic plane waves (§2.1.2) with the dispersion relation:

$$\omega_m^2 = c^2 \cdot (4/h^2) \sum_{i=1}^{3} \sin^2(\pi m_i/N_i)$$

Polarisation: each plane wave has two independent polarisation states, corresponding to the two independent components of the electric field transverse to the propagation direction. This is a finite-dimensional linear algebra statement (the transverse subspace has dimension 2). Type I.

Energy density: $u = (\varepsilon_0/2)|E|^2 + (1/2\mu_0)|B|^2$ at each grid point. A finite sum of squares. Type I.

Poynting vector: $S = (1/\mu_0) E \times B$ at each grid point. The energy flux through a surface is a finite sum of $S \cdot \hat{n}$ over grid faces. Type I.

2.7 Bounded Optics

Optics is the regime of electromagnetism where the wavelength is much smaller than the characteristic scale of obstacles and apertures. In BST, this is the regime where the relevant wavenumber m satisfies $m \gg 1$ but $m \ll N$ (many wavelengths fit on the grid, but the grid is still much finer than a wavelength).

2.7.1 Snell's law

At an interface between two regions with different wave speeds c_1 and c_2 (modelled by different values of the discretised refractive index $n = c/c_{medium}$), the dispersion relation forces:

$$\sin \theta_1 / \sin \theta_2 = c_1/c_2 = n_2/n_1$$

where θ_1, θ_2 are the angles of incidence and refraction. This is an algebraic relation in $\mathbb{R}_B(k)$, derived from the requirement that the tangential component of the wavevector is continuous across the interface (a finite boundary condition). Type I.

Reflection coefficient (Fresnel equations): algebraic functions of the angles and refractive indices. Finite computation. Type I.

2.7.2 Interference and diffraction

Interference: superposition of waves from two or more coherent sources. The intensity pattern is $|\Sigma_j a_j \exp(i\varphi_j)|^2$ — a finite sum of complex exponentials, squared. Exact finite computation in $\mathbb{C}_B(k^4)$. Type I.

Young's double slit: two point sources separated by distance d. Intensity at position x on a distant screen: $I(x) \propto \cos^2(\pi dx/(\lambda L))$, where L is the screen distance. A finite trigonometric expression. Type I.

Diffraction: the Fraunhofer diffraction pattern of an aperture is the DFT of the aperture function. For an aperture function $a(x_j)$ (= 1 inside the aperture, 0 outside), the far-field intensity is:

$$I(m) \propto |DFT(a)|^2_m = |\Sigma_j a(x_j) \exp(2\pi imj/N)|^2$$

This is an exact finite computation via the FFT (Theorem 9.6 of the AFB paper). Type I.

Single slit of width w: the DFT of a rectangular function gives the discrete sinc pattern. Type I.

Circular aperture: the DFT on a 2D grid gives the discrete Airy pattern. Type I.

2.7.3 Fourier optics

The general framework: in the Fraunhofer regime, the far-field pattern is the DFT of the near-field. Lenses perform the DFT optically.

In BST, this becomes: the DFT (Definition 9.1) is the mathematical model of optical Fourier transformation. An optical system that images through a lens is modelled by:

1. Represent the input field as $u_{in}: \Omega_h \to \mathbb{C}_B(k^4)$
2. Apply the DFT: $U = DFT(u_{in})$
3. Apply the lens transfer function: $U' = H \cdot U$ (pointwise multiplication, i.e. convolution in real space — Theorem 9.5)
4. Apply the inverse DFT: $u_{out} = DFT^{-1}(U')$

Each step is an exact finite computation. The entire optical system is modelled by finite matrix operations. Type I.

2.8 What Part II Recovers

Result	Type	Section
Scalar wave equation	I	2.1.1
Plane wave decomposition (DFT)	I	2.1.2
Dispersion relation (discrete)	I	2.1.2
Superposition and interference	I	2.1.3
Standing waves (eigenmodes)	I	2.1.3
Minkowski metric and intervals	I	2.2.1
Lorentz group $O(1,3)_B$	I	2.2.2
Boosts and rotations	I/III	2.2.2
Lorentz algebra structure constants	I	2.2.2
Four-vectors, $E = mc^2$	I	2.2.3
Relativistic collision kinematics	I	2.2.3
Coulomb's law on lattice	I/III	2.3.1
Gauss's law (discrete divergence)	I	2.3.1
Biot-Savart on lattice	I/III	2.3.2
$\nabla \cdot B = 0$ (from $d^2 = 0$)	I	2.3.2
$F = dA$ (field strength 2-form)	I	2.4
$dF = 0$ (Bianchi, automatic)	I	2.4
$d{*}F = {*}J$ (inhomogeneous Maxwell)	I	2.4
Gauge invariance ($A \to A + d\lambda$)	I	2.4
U(1) lattice gauge theory	I	2.5
Wilson action and partition function	I	2.5
Wilson loops	I	2.5
EM wave equation	I	2.6

Polarisation	I	2.6
Energy density and Poynting vector	I	2.6
Snell's law	I	2.7.1
Fresnel equations	I	2.7.1
Diffraction (DFT of aperture)	I	2.7.2
Fourier optics	I	2.7.3

What is NOT recovered: smooth electromagnetic fields on $\mathbb{R}^4$ (Type IV family limit), radiation reaction (requires self-interaction beyond the scope of this Part), quantum electrodynamics (requires Part IV quantum mechanics + Part V non-abelian gauge theory). The lattice dispersion relation deviates from the continuum $\omega = c|k|$ at short wavelengths — this is the correct finite-resolution physics, not an error.

What is gained: the gauge-theoretic structure (U(1) connection on edges, curvature on plaquettes, gauge invariance from group structure) is manifestly finite and requires no infinite-dimensional configuration space. The partition function is a finite sum. The automatic regularisation noted in Part I §1.6.2 (no 1/r singularities) extends to the Coulomb potential here: the potential of a point charge at the nearest grid point is large but finite.

End of Part II

Part III: Statistical Mechanics and Thermodynamics

Maturity: Tier 1

AFB imports: $\mathbb{R}$_B(k) (Part VIII, Theorem 8.13); bounded probability spaces, expectation, variance (Part IX, §9.3.4, Definitions 9.7–9.8); Bayes' theorem, Chebyshev inequality, law of large numbers (Part IX, Theorems 9.13–9.16); Shannon entropy, KL divergence, Jensen's inequality (Part IX, Definitions 9.11–9.12, Theorem 9.17); bounded Markov chains (Part IX, Definition 9.10); bounded transcendental functions including exp and ln (Part IX, §9.3); spectral theorem (Part XI, Theorem 11.4); matrix exponential (Part XI, Definition 11.16)

Volume imports: Part I (phase spaces, Hamiltonian mechanics, energy conservation)

The physics

Statistical mechanics connects the microscopic world of atoms and molecules to the macroscopic world of temperature, pressure, and heat. A gas in a box contains $\sim 10^{23}$ molecules, each obeying Newton's laws. Tracking them individually is impossible. Statistical mechanics instead describes the *probability distribution* over their collective states, deriving thermodynamic laws — the behaviour of engines, refrigerators, stars, and chemical reactions — from the statistics of large finite systems. Thermodynamics, the science of heat and work, was developed empirically (Carnot, Clausius, Kelvin) before statistical mechanics (Boltzmann, Gibbs) provided its microscopic foundation.

What experiments confirm

The gas laws. The ideal gas law PV = NkT, derived here in §3.7 from a finite partition function, has been verified for dilute gases since Boyle (1662) and Charles (1787). Deviations at high density are captured by the van der Waals equation and its successors — all computable by the same finite-sum partition function with modified Hamiltonians.

The laws of thermodynamics. The first law (energy conservation) is verified by calorimetry — every energy balance in every chemical and physical process ever measured confirms it. The second law (entropy non-decrease) has never been violated in any macroscopic experiment. The third law (entropy approaches a constant at absolute zero) is confirmed by heat capacity measurements that vanish as $T \to 0$. All three are proved as finite theorems in §3.3.

Phase transitions. The boiling of water, the magnetisation of iron, the onset of superconductivity — all are phase transitions described by statistical mechanics. The Ising model (§3.6) captures the essential physics of ferromagnetic transitions, and its predictions (critical temperature, critical exponents, magnetisation curves) are confirmed experimentally in materials like EuO and $CrBr_3$ to within a few percent.

Boltzmann's distribution. The exponential energy distribution $p(E) \propto \exp(-E/kT)$ — derived here in §3.4 as a theorem from maximum entropy — governs the speed distribution of gas molecules (Maxwell-Boltzmann, confirmed by molecular beam experiments), the intensity of blackbody radiation (Planck), and the population of atomic energy levels (spectroscopy). Every laser, every incandescent light, every star confirms it.

Monte Carlo methods. The Metropolis algorithm (§3.8), invented in 1953 at Los Alamos, is one of the most widely used computational methods in all of science. It is used to simulate protein folding, nuclear reactor design, financial risk, and lattice gauge theory. Every Monte Carlo computation is a finite Markov chain on a finite state space — the exact object BST constructs.

What BST constructs and why it suffices

Statistical mechanics was *invented* for finite systems. Boltzmann's original derivation of the entropy formula $S = k \ln W$ counted the number of microstates W of a finite collection of molecules. The "thermodynamic limit" ($N \to \infty$) was a later idealisation introduced to obtain mathematically sharp phase transitions and to apply the machinery of measure theory. BST reverses the idealisation: it works directly with the finite systems — which are what the physics describes and what the computations use — and observes the infinite-volume behaviour as a family-level (Type IV) phenomenon visible only from the metatheory.

Every partition function in this Part is a finite sum. Every entropy is the Shannon entropy of a finite probability distribution. Every thermodynamic quantity is computed by finite sums and finite differences. The second law is a theorem about finite Markov chains. No measure theory, no infinite-dimensional function spaces, no σ-algebras. The computational methods (exact enumeration, transfer matrices, Monte Carlo) are identical to the ones used in practice — because those methods were always finite.

3.1 Finite Ensembles

The central object of statistical mechanics is the ensemble: a probability distribution over the microstates of a physical system. In BST, every ensemble is a bounded probability space (Definition 9.7 of the AFB paper).

Definition 3.1 — Bounded microcanonical ensemble:
A physical system with finite state space Ω and Hamiltonian $H: \Omega \to \mathbb{R}_B(k)$.

The microcanonical ensemble at energy E with tolerance δ is the uniform distribution over the energy shell:

$$\Omega_{E,\delta} = \{\omega \in \Omega : |H(\omega) - E| \leq \delta\}$$

This is a finite subset of Ω, definable by Bounded Separation. The probability distribution is:

$$p(\omega) = 1/|\Omega_{E,\delta}| \quad \text{if } \omega \in \Omega_{E,\delta}$$
$$p(\omega) = 0 \quad \text{otherwise}$$

This is a bounded probability space (Definition 9.7) with uniform weights on a finite set. Type I.

The microcanonical entropy is:

$$S_{mc} = k_B \ln |\Omega_{E,\delta}|$$

where $k_B \in \mathbb{R}_B(k)$ is Boltzmann's constant and ln is evaluated by bounded transcendental approximation (Part IX of the AFB paper). S_{mc} is a specific element of $\mathbb{R}_B(k)$. Type I.

Definition 3.2 — Bounded canonical ensemble:
The canonical ensemble at inverse temperature $\beta = 1/(k_B T)$, with $\beta \in \mathbb{R}_B(k)$, is the probability distribution:

$$p(\omega) = (1/Z) \exp(-\beta H(\omega))$$

where the partition function is the finite sum:

$$Z(\beta) = \Sigma_{\omega \in \Omega} \exp(-\beta H(\omega))$$

Z is a sum of $|\Omega|$ terms, each computed by the bounded exponential (Part IX). It is a specific element of $\mathbb{R}_B(k)$, computed in $\mathbb{Q}_B(k^4)$.

The distribution p is a bounded probability space (Definition 9.7): $p(\omega) \geq 0$ for all ω, and $\Sigma\, p(\omega) = Z/Z = 1$. ✓ Type I.

Existence and computability:
- Z requires $|\Omega|$ evaluations of $\exp(-\beta H(\omega))$.
- Each exp evaluation uses bounded transcendental approximation with error $O(1/k)$.
- The division $1/Z$ is performed in $\mathbb{Q}_B(k^4)$.
- Total cost: $O(|\Omega|)$ operations.

```
For a lattice spin system with N sites and q states
per site: |Ω| = q^N. For the Ising model (q = 2)
on a 10×10 lattice: |Ω| = 2^{100} ≈ 10^{30}.
This is large but finite and well within any
reasonable n_M. For larger lattices, the partition
function is still a well-defined finite sum —
the computational cost is the only limitation,
not the mathematical existence.
```

Definition 3.3 — Bounded grand canonical ensemble:

```
For a system where particle number N can vary, the
grand canonical ensemble at inverse temperature β
and chemical potential μ ∈ ℝ_B(k) is:

p(ω) = (1/Z_G) exp(−β(H(ω) − μN(ω)))

where N(ω) is the particle number in microstate ω
and:

Z_G(β, μ) = Σ_{ω ∈ Ω} exp(−β(H(ω) − μN(ω)))

The state space Ω now includes states with different
particle numbers: Ω = ∪_{N=0}^{N_max} Ω_N, where
Ω_N is the state space at fixed particle number N.
The union is finite (N_max ≤ n_M). Type I.
```

Plain language: Every statistical ensemble in BST is a finite probability distribution on a finite set of states. The partition function is a finite sum — not a functional integral, not a trace over an infinite-dimensional Hilbert space, not an integral over an infinite phase space. The mathematical existence of Z is trivial (it is a finite sum of positive real numbers). The computational cost of evaluating Z may be large, but that is a practical matter, not a foundational one.

3.2 Thermodynamic Quantities

All thermodynamic quantities are derived from the partition function by finite operations: differentiation (finite differences), logarithms (bounded transcendentals), and expectation values (finite sums).

Definition 3.4 — Bounded thermodynamic functions:

```
From the canonical partition function Z(β):

Free energy:
  F = −k_B T ln Z = −(1/β) ln Z
  A specific element of ℝ_B(k). Type I.

Average energy:
  ⟨E⟩ = Σ_ω H(ω) p(ω) = −∂ ln Z / ∂β
  (where ∂/∂β is the finite difference quotient
  in β, Definition 9.2 of the AFB paper).
  Equivalently: ⟨E⟩ = (1/Z) Σ_ω H(ω) exp(−βH(ω)).
  A finite sum. Type I.
```

Entropy:

$$S = -k_B \sum_\omega p(\omega) \ln p(\omega)$$

This is the Shannon entropy (Definition 9.11 of the AFB paper) applied to the canonical distribution, scaled by Boltzmann's constant.

Equivalently: $S = k_B(\ln Z + \beta\langle E\rangle)$.

Type I.

Heat capacity:

$$C = \partial\langle E\rangle/\partial T = -k_B \beta^2 \partial\langle E\rangle/\partial\beta$$
$$= k_B \beta^2 (\langle E^2\rangle - \langle E\rangle^2)$$
$$= k_B \beta^2 \mathrm{Var}(E)$$

where $\mathrm{Var}(E)$ is the variance (Definition 9.8).

Type I.

Pressure (for systems with volume parameter V):

$$P = -\partial F/\partial V = (1/\beta) \partial \ln Z / \partial V$$

where V is a parameter of the Hamiltonian (e.g. the lattice size) and $\partial/\partial V$ is a finite difference.

Type I.

Chemical potential (grand canonical):

$$\mu = \partial F/\partial\langle N\rangle$$

Finite difference with respect to average particle number. Type I.

The fluctuation-dissipation relation:

$$C = k_B \beta^2 \mathrm{Var}(E)$$

connects the response function (heat capacity) to the equilibrium fluctuations (energy variance). This is an exact algebraic identity — it follows from differentiating ln Z twice. In BST it is proved by direct computation:

$$\partial^2\ln Z/\partial\beta^2 = (1/Z)\Sigma\, H^2e^{-\beta H} - ((1/Z)\Sigma\, He^{-\beta H})^2$$
$$= \langle E^2\rangle - \langle E\rangle^2 = \mathrm{Var}(E). \quad \square$$

Type I. The identity holds exactly for any finite probability distribution.

3.3 The Laws of Thermodynamics

The laws of thermodynamics are universal constraints on physical processes. In classical physics, they are either postulated or derived from infinite-system limits. In BST, they are finite theorems.

Theorem 3.1 — Zeroth law (BST):

Thermal equilibrium is an equivalence relation.

If system A is in thermal equilibrium with system B, and system B is in thermal equilibrium with system C, then system A is in thermal equilibrium with system C.

In BST: Two systems are in thermal equilibrium when their combined system is described by a canonical ensemble at a common temperature β. "A is in equilibrium with B at β" means: the joint distribution on $\Omega_A \times \Omega_B$ factors as $p_A(\beta) \times p_B(\beta)$. This relation is transitive because a real number (β) that equals another that equals a third is equal to the third — transitivity of equality in $\mathbb{R}_B(k)$. □

Type I. The proof uses only transitivity of equality.

Theorem 3.2 — First law (BST):

Energy is conserved in an isolated system.

In BST: For a Hamiltonian system (Part I, §1.3), the total energy H(q, p) is constant along the trajectory (Theorem 1.1 of this volume). When two systems interact, the total Hamiltonian $H_{total} = H_A + H_B + H_{int}$ is conserved. Energy transferred from A to B equals the negative of energy transferred from B to A.

In the statistical setting: $\langle E_A\rangle + \langle E_B\rangle + \langle E_{int}\rangle$ is constant under Hamiltonian evolution of the combined system. Changes in $\langle E_A\rangle$ are accounted for by changes in $\langle E_B\rangle$ and $\langle E_{int}\rangle$. This is the first law: $\Delta E = Q + W$, where Q is heat (energy transferred via uncontrolled microscopic degrees) and W is work (energy transferred via controlled macroscopic parameters). □

Type I. Follows from Hamiltonian energy conservation.

Theorem 3.3 — Second law (BST):

The entropy of an isolated system does not decrease under Markovian evolution.

In BST: Let Ω be a finite state space with probability distribution p, and let T be a doubly stochastic transition matrix (Definition 9.10 of the AFB paper: $T_{ij} \geq 0$, $\Sigma_j T_{ij} = 1$, $\Sigma_i T_{ij} = 1$).

The distribution after one step is $p' = pT$.

Claim: $H(p') \geq H(p)$, where H is the Shannon entropy (Definition 9.11).

Proof: By the Schur-concavity of entropy. Since T is doubly stochastic, $p' = pT$ is majorised by p (a standard finite-dimensional inequality). Shannon entropy is Schur-concave, so $H(p') \geq H(p)$.

More directly: by Jensen's inequality (Theorem 9.17 of the AFB paper) applied to the concave function φ(x) = −x ln x:

```
H(p') = −Σ_j p'_j ln p'_j
      = −Σ_j (Σ_i p_i T_{ij}) ln(Σ_i p_i T_{ij})
      ≥ −Σ_j Σ_i p_i T_{ij} ln(p_i T_{ij})
      = −Σ_i p_i ln p_i − Σ_i p_i Σ_j T_{ij} ln T_{ij}
      ≥ H(p)
```

(The last inequality uses the non-negativity of the Gibbs entropy of the transition matrix rows.) □

Type I. The proof is a finite algebraic inequality using Jensen (Theorem 9.17). No infinite-dimensional measure theory, no ergodic hypothesis, no thermodynamic limit.

Note: The second law in BST is a theorem about finite Markov processes, not an asymptotic statement about infinite systems. It holds exactly at each finite state space. The classical second law — entropy increases in the thermodynamic limit — is the family-level reading (Type IV).

Theorem 3.4 — Third law (BST):

As β → ∞ (T → 0), the canonical distribution concentrates on the ground state(s).

In BST: Let E_0 = min_{ω ∈ Ω} H(ω) be the ground state energy (which exists because Ω is finite — EVT, Part IX). Let Ω_0 = {ω : H(ω) = E_0} be the ground state set.

```
For any ω ∈ Ω_0:
p(ω) = exp(−βE_0)/Z ≥ exp(−βE_0)/(|Ω|exp(−βE_0))
      = 1/|Ω|
```

For any ω ∉ Ω_0: H(ω) ≥ E_0 + Δ for some gap Δ > 0 (which exists because H takes finitely many values). So:
p(ω) = exp(−βH(ω))/Z ≤ exp(−βΔ)/|Ω_0|

As β grows through ℝ_B(k), this ratio → 0.

At the minimum accessible temperature (the largest β in ℝ_B(k)), the ground state probability is ≥ 1 − |Ω|exp(−βΔ), and the entropy is:

S ≤ k_B ln|Ω_0| + O(exp(−βΔ))

If the ground state is non-degenerate (|Ω_0| = 1), then $S \to 0$. □

Type II. Exact inequality at each β, with the classical third law as the $\beta \to \infty$ family limit.

3.4 The Boltzmann Distribution: Variational Characterisation

Theorem 3.5 — Maximum entropy characterisation:

Among all probability distributions on Ω with a given average energy ⟨E⟩ = E_target, the canonical distribution $p(\omega) = Z^{-1} \exp(-\beta H(\omega))$ is the unique distribution maximising the Shannon entropy H(p).

Proof: Let q be any other distribution with ⟨E⟩_q = E_target. Then:

```
H(p) − H(q) = Σ_ω q(ω) ln(q(ω)/p(ω)) − Σ_ω (q(ω) − p(ω)) ln p(ω)
             = D_KL(q ‖ p) + β(⟨E⟩_q − ⟨E⟩_p) + ln Z − ln Z
             = D_KL(q ‖ p) ≥ 0
```

where D_KL is the Kullback-Leibler divergence (Definition 9.12 of the AFB paper), which is non-negative by Gibbs' inequality (proved by Jensen, Theorem 9.17). Equality iff q = p. □

Type I. The proof is a finite-sum computation using the properties of entropy and KL divergence already established in the AFB paper.

Plain language: The Boltzmann distribution is not a postulate in BST — it is a theorem. Among all probability distributions consistent with a given average energy, the canonical ensemble is the one that maximises entropy. The proof is a finite algebraic argument using Jensen's inequality and the non-negativity of KL divergence. This variational characterisation explains *why* physical systems in thermal equilibrium are described by the canonical ensemble: it is the least-biased distribution given the energy constraint.

3.5 Phase Transitions and Finite-Size Effects

3.5.1 The finite smoothing principle

On a finite state space Ω, all thermodynamic quantities — free energy, entropy, heat capacity, magnetisation, susceptibility — are smooth (in fact, analytic) functions of the external parameters (temperature, field strength, coupling constants).

Proof: The partition function Z(β, h, ...) is a finite sum of exponentials. Each exponential is smooth in its parameters. A finite sum of smooth functions is smooth. Therefore $F = -(1/\beta) \ln Z$ is smooth, and all derivatives of F are well-defined.

In particular: the heat capacity C(T) is continuous.

There are no discontinuities, no divergences, no non-analyticities. □

Consequence: true phase transitions — defined as non-analyticities in the free energy — cannot occur in a finite system. This is the Yang-Lee theorem for finite systems, and it is trivially true in BST.

3.5.2 Crossovers and the family limit

Although finite systems have no true phase transitions, they can exhibit sharp crossovers: rapid but smooth changes in thermodynamic quantities over a narrow temperature range.

The sharpness of the crossover increases with system size. For a d-dimensional system of linear size L:

- The susceptibility peak scales as $\chi_{max} \sim L^{\gamma/\nu}$
- The peak width scales as $\Delta T \sim L^{-1/\nu}$
- The correlation length is bounded by L

where γ and ν are the critical exponents of the corresponding infinite-system transition.

In BST, each finite lattice gives a smooth crossover. The family {crossover at lattice size L} shows increasing sharpness as L grows through the metatheory. The phase transition is the family-level limit — a Type IV phenomenon.

This is the correct finite-resolution picture of phase transitions. It is what numerical simulations and real physical systems actually exhibit: finite rounding of the transition, visible sharpening with system size, and exact criticality only in the (metatheoretic) infinite-volume limit.

3.6 The Ising Model

The Ising model is the simplest non-trivial statistical mechanics system. It demonstrates the full BST framework with a single concrete example.

Definition 3.5 — Bounded Ising model:

System: N spins on a finite graph $G = (V, E)$, with $|V| = N$.

State space: $\Omega = \{+1, -1\}^N$. Each microstate $\sigma = (\sigma_1, \ldots, \sigma_N)$ assigns a spin ±1 to each vertex. $|\Omega| = 2^N$. Finite. ✓

Hamiltonian:
$$H(\sigma) = -J \sum_{(i,j) \in E} \sigma_i \sigma_j - h \sum_{i \in V} \sigma_i$$

where $J \in \mathbb{R}\_B(k)$ is the coupling constant and $h \in \mathbb{R}\_B(k)$ is the external field.

The first sum runs over edges (nearest-neighbour interactions). The second sum is the Zeeman coupling to the external field. Both are finite sums of products in $\mathbb{R}\_B(k)$. Type I.

3.6.1 Exact thermodynamics

Partition function:

$$Z = \Sigma_{\sigma \in \Omega} \exp(-\beta H(\sigma))$$
$$= \Sigma_{\sigma \in \{\pm 1\}^N} \exp(\beta J \Sigma \sigma_i \sigma_j + \beta h \Sigma \sigma_i)$$

This is a sum of 2^N terms. Each term is computed by evaluating $H(\sigma)$ ($O(|E|)$ operations) and taking the bounded exponential ($O(1)$ operations). Total cost: $O(2^N \cdot |E|)$ operations in $\mathbb{R}\_B(k)$.

For small N (≤ 20, say), this is directly computable. For larger N, the transfer matrix method (§3.6.2) or Monte Carlo sampling (§3.8) provides more efficient approaches.

Thermodynamic quantities:

- Free energy: $F = -(1/\beta) \ln Z$.
- Average energy: $\langle E\rangle = -\partial \ln Z / \partial\beta$.
- Magnetisation: $m = \langle (1/N) \Sigma \sigma_i\rangle = (1/N\beta)\, \partial \ln Z / \partial h$.
- Susceptibility: $\chi = \partial m/\partial h = \beta(\langle M^2\rangle - \langle M\rangle^2)/N$.
- Heat capacity: $C = k_B \beta^2(\langle E^2\rangle - \langle E\rangle^2)$.

All finite sums or finite differences. Type I.

3.6.2 The transfer matrix

For the 1D Ising model on a chain of N sites with periodic boundary conditions:

$$Z = \Sigma_{\{\sigma_1,\ldots,\sigma_N\}} \Pi_{i=1}^{N} T_{\sigma_i, \sigma_{i+1}}$$

where T is the 2 × 2 transfer matrix:

$$T = \begin{pmatrix} \exp(\beta J + \beta h) & \exp(-\beta J) \\ \exp(-\beta J) & \exp(\beta J - \beta h) \end{pmatrix}$$

Then: $Z = \mathrm{Tr}(T^N) = \lambda_+^N + \lambda_-^N$

where $\lambda_+ > \lambda_-$ are the eigenvalues of T, computed by the spectral theorem (Theorem 11.4 of the AFB paper).

$$\lambda_\pm = \exp(\beta J)\, [\cosh(\beta h) \pm \sqrt{\sinh^2(\beta h) + \exp(-4\beta J)}]$$

These are specific elements of $\mathbb{R}\_B(k)$, computed by bounded transcendental evaluation.

Free energy per site:
$f = -(1/N\beta) \ln Z = -(1/\beta) \ln \lambda_+ - (1/N\beta) \ln(1 + (\lambda_-/\lambda_+)^N)$

For large N, the second term is exponentially small (since $\lambda_-/\lambda_+ < 1$). The free energy per site converges to $-(1/\beta) \ln \lambda_+$ — the exact 1D solution.

This is a Type I computation at each N, with the Onsager-like exact solution visible as the metatheoretic $N \to \infty$ limit.

For the 2D Ising model: the transfer matrix is $2^L \times 2^L$ for a lattice of width L. Its eigenvalues are computable by Part XI's spectral theorem for any specific L. The Onsager exact solution (the infinite-L limit) is a Type IV family statement.

3.7 The Classical Ideal Gas

Definition 3.6 — Bounded ideal gas:
System: N non-interacting particles of mass m in a bounded d-dimensional box of side L, with the box modelled as a bounded grid Ω_h with spacing h.

Single-particle state space: the grid points of the box. $|\Omega_{single}| = (L/h)^d$.

N-particle state space (distinguishable particles):
$\Omega = \Omega_{single}^N$. $|\Omega| = (L/h)^{dN}$.

For indistinguishable particles: Ω consists of unordered N-element multisets from Ω_{single}, with $|\Omega| = C(|\Omega_{single}| + N - 1, N)$ for bosons or $C(|\Omega_{single}|, N)$ for fermions.

Hamiltonian: $H = \Sigma_i |p_i|^2/(2m)$, where the momenta are the DFT dual variables (§9.3.2 of the AFB paper). The kinetic energy eigenvalues are:

$$\varepsilon_n = (\hbar^2/(2m)) |2\pi n/L|^2$$

for integer vector $n = (n_1, ..., n_d)$ with components bounded by $L/(2h)$.

Theorem 3.6 — Bounded ideal gas equation of state:
For the classical (distinguishable-particle) ideal gas at temperature T:

The single-particle partition function:
$z = \Sigma_n \exp(-\beta \varepsilon_n) \approx (L/\lambda_{th})^d$

where $\lambda_{th} = h_{Planck}/\sqrt{2\pi m k_B T}$ is the thermal de Broglie wavelength. The approximation becomes exact (Type II) when $\lambda_{th} \gg h$ (many grid points per thermal wavelength).

The N-particle partition function: $Z = z^N / N!$ (for indistinguishable particles in the classical limit where occupation numbers are small).

Free energy: $F = -k_BT(N \ln z - \ln N!)$
$\approx -Nk_BT(\ln(V/N) + d/2 \cdot \ln(2\pi m k_BT/h_P^2) + 1)$

Pressure: $P = -\partial F/\partial V = Nk_BT/V$.

This is the ideal gas law $PV = Nk_BT$, recovered as a finite-sum computation. Type I for the finite partition function; Type II for the correspondence with the classical ideal gas law.

3.8 Bounded Monte Carlo Methods

For systems where direct evaluation of Z is computationally prohibitive ($|\Omega| \gg 10^{20}$), Monte Carlo sampling provides an efficient alternative. In BST, Monte Carlo is a finite Markov chain (Definition 9.10 of the AFB paper) on the finite state space Ω.

Definition 3.7 — Bounded Metropolis algorithm:
Input: finite state space Ω, Hamiltonian H, inverse temperature β, number of steps N_{steps}.

Algorithm (by bounded recursion, BI-BST):

1. Choose initial state $\sigma_0 \in \Omega$.
2. For $n = 0, 1, \ldots, N_{steps} - 1$:
   a. Propose a new state σ' by a local move (e.g. flip one spin for the Ising model).
   b. Compute $\Delta E = H(\sigma') - H(\sigma_n)$.
   c. If $\Delta E \leq 0$: accept ($\sigma_{n+1} = \sigma'$).
   d. If $\Delta E > 0$: accept with probability $\exp(-\beta\Delta E)$; otherwise reject ($\sigma_{n+1} = \sigma_n$).

Step (d) requires comparing a random number $r \in \mathbb{R}_B(k) \cap [0,1]$ with $\exp(-\beta\Delta E)$ — a finite comparison. The random number is drawn from a bounded pseudorandom generator or an external random source.

The Metropolis algorithm defines a Markov chain (Definition 9.10) with transition matrix T satisfying:
- Detailed balance: $p(\sigma)T(\sigma,\sigma') = p(\sigma')T(\sigma',\sigma)$ for the canonical distribution p.
- Ergodicity: if the proposal move connects all states, T is irreducible.

```
Therefore the stationary distribution of the chain
is the canonical ensemble. By the finite convergence
theorem for Markov chains (spectral gap of T,
Theorem 11.4), the distribution converges to
equilibrium in a computable number of steps.

Output: after M_burn burn-in steps, the states
{σ_n}_{n=M_burn}^{N_steps} sample (approximately)
from the canonical distribution. Expectation values
are estimated by:

⟨O⟩ ≈ (1/N_sample) Σ_{n=M_burn}^{N_steps} O(σ_n)

Type I for the algorithm. The approximation quality
is controlled by the spectral gap of T and the
number of samples — both finite and computable.
```

Plain language: Monte Carlo in BST is a bounded recursion on a finite Markov chain. The convergence is guaranteed by the spectral gap of the transition matrix — a computable quantity via Part XI's spectral theorem. There is no mystery about "ergodicity in infinite time" or "measure-theoretic convergence." The chain is finite, its transition matrix is a finite matrix, its convergence rate is a computable eigenvalue, and the sampling error is a computable bound.

3.9 Bounded Information Theory

The connection between thermodynamic entropy and information-theoretic entropy is not merely analogical in BST — it is an exact identification.

Theorem 3.7 — Entropy identification:

```
For a system in the canonical ensemble with
distribution p(ω) = Z⁻¹ exp(−βH(ω)):

The thermodynamic entropy S = ⟨E⟩/T + k_B ln Z
equals the Shannon entropy (Definition 9.11):

S = −k_B Σ_ω p(ω) ln p(ω)

Proof: Direct computation.
−k_B Σ p(ω) ln p(ω)
= −k_B Σ p(ω) (−βH(ω) − ln Z)
= k_B β Σ p(ω) H(ω) + k_B ln Z
= k_B β ⟨E⟩ + k_B ln Z
= ⟨E⟩/T + k_B ln Z.  □

Type I. The identification is an exact algebraic
identity for any finite probability distribution.

Consequences:

The second law (Theorem 3.3) is an instance of the
information-theoretic principle that coarse-graining
cannot decrease entropy — proved for finite Markov
```

chains by Jensen's inequality.

The maximum entropy principle (Theorem 3.5) is Jaynes' information-theoretic derivation of statistical mechanics – proved as an exact finite theorem by the non-negativity of KL divergence.

The Landauer bound on the thermodynamic cost of erasing information: erasing one bit requires at least k_BT ln 2 of energy dissipated as heat. In BST this is a finite algebraic inequality relating the entropy decrease (k_B ln 2 per erased bit) to the minimum heat production via the second law. Type I.

3.10 What Part III Recovers

Result	Type	Section
Microcanonical ensemble	I	3.1
Canonical ensemble and Z	I	3.1
Grand canonical ensemble	I	3.1
Free energy, energy, entropy	I	3.2
Heat capacity (fluctuation formula)	I	3.2
Pressure, chemical potential	I	3.2
Zeroth law	I	3.3
First law (energy conservation)	I	3.3
Second law (entropy non-decrease)	I	3.3
Third law (ground-state limit)	II	3.3
Maximum entropy → Boltzmann	I	3.4
Phase transitions (finite systems)	–	3.5
Phase transition (true)	IV	3.5
Ising partition function	I	3.6
Transfer matrix method	I	3.6
Ideal gas law PV = NkT	I/II	3.7
Metropolis Monte Carlo	I	3.8
Entropy = Shannon entropy	I	3.9
Landauer bound	I	3.9

What is NOT recovered: true phase transitions with divergent correlation length or non-analytic free energy (these are Type IV family-level phenomena — they emerge in the metatheoretic limit of increasing system size). Continuous symmetry breaking in finite systems (the finite-system magnetisation is always analytic in the field). Renormalisation group fixed points as continuum objects (the finite analogue is developed in Part IX of this volume). Infinite-volume Gibbs measures.

What is gained: every thermodynamic quantity is a computable finite number, not an infinite-dimensional functional. The partition function is a finite sum that always exists and is always positive — no questions of convergence, regularisation, or renormalisation. The second law is a theorem, not a postulate or an infinite-time asymptotic statement. The entropy is always finite and always computable. Monte Carlo convergence is guaranteed by a computable spectral gap. The connection between thermodynamic and information-theoretic entropy is an exact identity, not a philosophical interpretation.

End of Part III

Part IV: Quantum Mechanics

Maturity: Tier 1

AFB imports: Finite Hilbert spaces, inner product, Cauchy-Schwarz (Part XI, §11.5, Definitions 11.7–11.8, Theorem 11.7); spectral theorem for normal operators (Part XI, Theorem 11.4); Weyl eigenvalue stability, Davis-Kahan eigenspace stability (Part XI, Theorems 11.5–11.6); matrix exponential and unitarity (Part XI, §11.7, Definition 11.16, Theorem 11.11); tensor products (Part XI, §11.6, Definition 11.10); $\mathbb{C}$_B(k^4) (Part VIII, Theorem 8.16); bounded probability spaces and Born rule (Part IX, §9.3.4, Definitions 9.7–9.8); representation theory including Maschke, Schur, character orthogonality, irreducible decomposition (Part XIII, Theorems 13.3–13.7); DFT for momentum space (Part IX, §9.3.2, Definition 9.1)

Volume imports: Part I (classical mechanics, Poisson bracket as classical limit); Part III (statistical ensembles for thermal quantum states)

The physics

Quantum mechanics is the physics of the very small — atoms, electrons, photons, and the fundamental particles. Developed between 1900 and 1930 (Planck, Bohr, Heisenberg, Schrödinger, Dirac), it replaced classical mechanics as the fundamental theory of matter. Its mathematical core is linear algebra on complex vector spaces: states are vectors, observables are Hermitian operators, time evolution is unitary, and measurement yields eigenvalues with probabilities given by the Born rule. Quantum mechanics governs all of chemistry, all of solid-state physics, and all of nuclear and particle physics.

What experiments confirm

Quantum mechanics is the most precisely verified physical theory in existence.

Atomic spectra. The hydrogen spectrum — the Rydberg series $E_n = -13.6\ \text{eV}/n^2$ constructed in §4.2.3 by diagonalising a finite Hamiltonian matrix — was the first quantitative success of quantum theory (Bohr 1913). Spectroscopic measurements now confirm it to 12 significant figures, making the hydrogen Lamb shift one of the most precise agreements in all of science.

The double-slit experiment. The interference pattern of single electrons passing through a double slit — a direct consequence of the superposition principle (Postulate 1, §4.1) — was confirmed by Tonomura et al. (1989) and subsequent experiments. Each electron arrives as a localised detection event; the interference pattern emerges from the statistics of many events, exactly as the Born rule predicts.

Bell inequality violation. The CHSH inequality violation (Theorem 4.1 of this volume: $S = 2\sqrt{2}$) was experimentally confirmed by Aspect et al. (1982) and, with all loopholes closed, by Hensen et al. (2015). The measured value $S \approx 2.80 \pm 0.02$ matches the quantum prediction $2\sqrt{2} \approx 2.828$. In BST, this is an exact 4×4 matrix computation — no infinite-dimensional Hilbert space required.

Quantum information. Quantum teleportation (§4.8.2) was experimentally demonstrated by Bouwmeester et al. (1997). Quantum error correction (§4.8.4) is implemented in current quantum computing hardware (Google, IBM, Quantinuum). Quantum key distribution (BB84 protocol) is commercially deployed. Every one of these protocols operates on finite-dimensional Hilbert spaces — qubits, not $L^2(\mathbb{R})$.

Spectral gaps. The spectral gap of a quantum Hamiltonian (§4.9) determines whether a system is gapped (insulator) or gapless (conductor, superfluid). Measurements of the spectral gap in materials like $TlCuCl_3$ and $BiFeO_3$ confirm quantum predictions from exact diagonalisation and Monte Carlo to within a few percent.

What BST constructs and why it suffices

Every quantum computation in a laboratory — every simulation of a molecule, every design of a quantum circuit, every prediction of a spectrum — is performed on a finite-dimensional Hilbert space by finite matrix operations. The infinite-dimensional Hilbert space $L^2(\mathbb{R})$ is a theoretical convenience that allows formal proofs using functional analysis, but no measurement has ever required it: every measured spectrum has finitely many lines, every observed state has finite-dimensional support, and every computation terminates after finitely many steps.

This Part makes the finiteness explicit. States are vectors in $\mathbb{C}_B(k^4)^d$. Observables are $d \times d$ Hermitian matrices. Time evolution is $\exp(-iHt/\hbar)$ computed by the matrix exponential. The Born rule is a finite probability distribution. The spectral gap is a computable number. Every postulate of quantum mechanics holds in BST, and several hold *more strongly* — completeness is automatic (finite sets are trivially complete), every operator is bounded (finite matrices have computable norms), and domain questions do not arise.

4.1 The Postulates of Bounded Quantum Mechanics

Classical quantum mechanics is built on a set of postulates that specify the mathematical framework and its physical interpretation. In BST, the same postulates hold with “Hilbert space” replaced by “finite Hilbert space” — and several become stronger because finiteness eliminates the domain and convergence issues that plague infinite-dimensional quantum mechanics.

Postulate 1 — States:

```
The state of a quantum system is a unit vector
|ψ⟩ in a finite Hilbert space ℋ over ℂ_B(k⁴)
(Definition 11.8 of the AFB paper).

dim(ℋ) = d < ∞. The state |ψ⟩ is a d-component
column vector with entries in ℂ_B(k⁴), satisfying
⟨ψ|ψ⟩ = 1.

More generally, a mixed state is a density operator:
ρ = Σ_i w_i |ψ_i⟩⟨ψ_i|

where w_i ≥ 0, Σ w_i = 1, and each |ψ_i⟩ ∈ ℋ.
ρ is a d × d Hermitian positive-semidefinite matrix
over ℂ_B(k⁴) with Tr(ρ) = 1.
```

Type I. States are finite vectors; density matrices are finite matrices.

Postulate 2 — Observables:

An observable is a Hermitian operator A on $\mathcal{H}$. Since $\mathcal{H}$ is finite-dimensional, A is a $d \times d$ Hermitian matrix over $\mathbb{C}_B(k^4)$: $A = A^\dagger$.

By the spectral theorem (Theorem 11.4):

$$A = \Sigma_{i=1}^{d} a_i |a_i\rangle\langle a_i|$$

where $a_i \in \mathbb{R}_B(k)$ are the eigenvalues (real, because A is Hermitian) and $\{|a_i\rangle\}$ is an orthonormal eigenbasis.

The possible measurement outcomes are the eigenvalues $\{a_1, \ldots, a_d\}$ — a finite set, exactly computable by diagonalisation. No continuous spectrum. No domain issues. No essential self-adjointness questions.

Type I.

Postulate 3 — The Born rule:

The probability of measuring outcome a_i on a system in state $|\psi\rangle$ is:

$$P(a_i) = |\langle a_i|\psi\rangle|^2$$

This is a specific element of $\mathbb{R}_B(k)$, computed by one inner product (d multiplications and $d-1$ additions in $\mathbb{C}_B(k^4)$) followed by one absolute-value squaring.

The set $\{P(a_i)\}$ forms a bounded probability distribution (Definition 9.7 of the AFB paper): $P(a_i) \geq 0$, and $\Sigma_i P(a_i) = \Sigma_i |\langle a_i|\psi\rangle|^2 = 1$ (by completeness of the eigenbasis and Parseval).

After measurement yielding outcome a_i, the state collapses to $|a_i\rangle$ (projection postulate).

Type I.

Postulate 4 — Time evolution:

The state of a closed system evolves by:

$$|\psi(t)\rangle = U(t)|\psi(0)\rangle$$

where $U(t) = \exp(-iHt/\hbar)$ is the time evolution operator (Definition 11.16 of the AFB paper).

H is the Hamiltonian — a Hermitian operator on ℋ.
By Theorem 11.11(iv), U(t) is unitary: U†U = I.

Unitarity preserves the norm: ⟨ψ(t)|ψ(t)⟩ = 1.
Probability is conserved. ✓

The Schrödinger equation:
iℏ d|ψ⟩/dt = H|ψ⟩

is the differential form of U(t) = exp(−iHt/ℏ).
In BST, the derivative is a finite difference
(Definition 9.2) and the equation is solved exactly
by the matrix exponential (Theorem 11.12).

Type I for the matrix exponential solution.
Type III for the numerical evaluation of exp at
precision O(1/k).

Postulate 5 — Composite systems:
The state space of a composite system AB is the
tensor product:

ℋ_{AB} = ℋ_A ⊗ ℋ_B

(Definition 11.10 of the AFB paper).

dim(ℋ_{AB}) = dim(ℋ_A) × dim(ℋ_B).

If ℋ_A has dimension d_A and ℋ_B has dimension d_B,
then ℋ_{AB} has dimension d_A · d_B. The tensor
product is a finite Hilbert space. Type I.

Plain language: The five postulates of quantum mechanics transfer verbatim to BST, with every Hilbert space finite, every operator a finite matrix, every spectrum a finite set of computable eigenvalues, and every probability a computable real number. The infinite-dimensional Hilbert space of classical quantum mechanics is what the family {ℋ at dimension d} looks like as d → ∞ through the metatheory. At each fixed dimension, the quantum theory is exact finite linear algebra.

4.2 Elementary Quantum Systems

4.2.1 The qubit

The simplest quantum system: ℋ = ℂ_B(k^4)^2.
dim(ℋ) = 2.

Basis states: |0⟩ = (1, 0)^T, |1⟩ = (0, 1)^T.

General state: |ψ⟩ = α|0⟩ + β|1⟩ with
$|\alpha|^2 + |\beta|^2 = 1$.

The Pauli matrices:

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \qquad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Properties (verified by 2×2 matrix multiplication):

- Hermitian: $\sigma_i^\dagger = \sigma_i$. ✓
- Unitary: $\sigma_i^2 = I$. ✓
- Anticommuting: $\sigma_i \sigma_j + \sigma_j \sigma_i = 2\delta_{ij} I$. ✓
- Commuting: $[\sigma_i, \sigma_j] = 2i\varepsilon_{ijk} \sigma_k$. ✓

These are the Lie algebra relations of su(2) — verified by exact finite matrix computation. Type I.

Eigenvalues of σ_z: +1 (spin up) and −1 (spin down).
Eigenstates: $|0\rangle$ and $|1\rangle$.
Measurement of σ_z on state $\alpha|0\rangle + \beta|1\rangle$:

$$P(+1) = |\alpha|^2, \quad P(-1) = |\beta|^2.$$

4.2.2 The bounded harmonic oscillator

The quantum harmonic oscillator in d dimensions is truncated to the first N energy levels.

Hilbert space: $\mathcal{H} = \mathbb{C}_B(k^4)^N$.
Basis: $\{|0\rangle, |1\rangle, \ldots, |N-1\rangle\}$ (energy eigenstates).

Hamiltonian (diagonal in the energy basis):

$$H = \hbar\omega \cdot \mathrm{diag}(1/2, 3/2, 5/2, \ldots, (2N-1)/2)$$

i.e. $H|n\rangle = \hbar\omega(n + 1/2)|n\rangle$ for $n = 0, \ldots, N-1$.

Creation and annihilation operators:

$a^\dagger|n\rangle = \sqrt{(n+1)}|n+1\rangle$ (for $n < N-1$; $a^\dagger|N-1\rangle = 0$)
$a|n\rangle = \sqrt{n}|n-1\rangle$ (for $n > 0$; $a|0\rangle = 0$)

These are N × N matrices:

$(a)_{m,n} = \sqrt{n} \cdot \delta_{m,n-1}$
$(a^\dagger)_{m,n} = \sqrt{(n+1)} \cdot \delta_{m,n+1}$

Number operator: $\hat{N} = a^\dagger a = \mathrm{diag}(0, 1, 2, \ldots, N-1)$.
$H = \hbar\omega(\hat{N} + 1/2)$.

The commutation relation $[a, a^\dagger] = I$ holds exactly as an N × N matrix identity. ✓

Position and momentum operators:

$\hat{x} = \sqrt{(\hbar/2m\omega)}(a + a^\dagger)$
$\hat{p} = i\sqrt{(m\hbar\omega/2)}(a^\dagger - a)$

Commutation: $[\hat{x}, \hat{p}] = i\hbar(I - 2|N-1\rangle\langle N-1|)$

The canonical commutation relation $[\hat{x}, \hat{p}] = i\hbar I$ holds exactly on the span of $\{|0\rangle, \ldots, |N-2\rangle\}$ and has a boundary correction at $|N-1\rangle$. This is the finite-truncation effect — it vanishes as $N \to \infty$

through the metatheory. Type I for the finite system; Type II for the family.

Energy spectrum: $E_n = \hbar\omega(n + 1/2)$ for $n = 0,\ldots,N-1$. Exact. The spacing $\hbar\omega$ is uniform. The zero-point energy $E_0 = \hbar\omega/2$ is a specific element of $\mathbb{R}_B(k)$.

4.2.3 The bounded hydrogen atom

The hydrogen atom Hamiltonian in a finite basis of N_{max} orbital states (truncated at principal quantum number $n \leq N_{max}$):

Hilbert space: $\mathcal{H} = \mathbb{C}_B(k^4)^D$ where $D = \Sigma_{n=1}^{N_{max}} n^2 = N_{max}(N_{max}+1)(2N_{max}+1)/6$.

The Hamiltonian is a $D \times D$ Hermitian matrix whose eigenvalues, in the exact hydrogen basis, are:

$E_n = -13.6 \text{ eV} / n^2$ for $n = 1, \ldots, N_{max}$

Each energy level has degeneracy n^2 (from angular momentum quantum numbers $l = 0,\ldots,n-1$ and $m = -l,\ldots,l$).

In BST, the Hamiltonian matrix is diagonalised by the spectral theorem (Theorem 11.4). The eigenvalues are the Rydberg series — exact polynomial expressions in the fundamental constants, computed in $\mathbb{R}_B(k)$.

For $N_{max} = 10$: $D = 385$ states. The Hamiltonian is a 385×385 Hermitian matrix. Diagonalisation gives the first 10 energy levels of hydrogen. Type I.

Transition frequencies: $\nu_{n\to m} = |E_n - E_m|/h$. These are computable elements of $\mathbb{R}_B(k)$. The Lyman, Balmer, and Paschen series are finite subsets of the transition frequencies. Type I.

4.3 Angular Momentum and Symmetry

Angular momentum in quantum mechanics is the representation theory of SU(2). Part XIII of the AFB paper develops the general theory of finite group representations. This section applies it to the rotation group.

4.3.1 Bounded spin representations

The group $SU(2)_B$ acts on finite Hilbert spaces by its irreducible representations (Part XIII, §13.2).

For each $j = 0, 1/2, 1, 3/2, \ldots$, there is an irreducible representation of dimension $2j+1$:

$\rho_j: SU(2)_B \to GL(\mathcal{H}_j)$

where $\mathcal{H}_j = \mathbb{C}_B(k^4)^{2j+1}$.

The angular momentum operators J_x, J_y, J_z are $(2j+1) \times (2j+1)$ Hermitian matrices satisfying the commutation relations:

$$[J_i, J_j] = i\hbar\, \varepsilon_{ijk} J_k$$

These are the structure constants of su(2), verified by exact matrix computation (Part XIII, §13.3).

$$J^2 = J_x^2 + J_y^2 + J_z^2 = \hbar^2 j(j+1) \cdot I$$

Eigenvalues of J_z: $m = -j, -j+1, \ldots, j-1, j$ ($2j+1$ values). Each is a specific element of $\mathbb{R}_B(k)$.

For $j = 1/2$: the Pauli matrices (§4.2.1).
For $j = 1$: the 3×3 spin-1 matrices.
For $j = 3/2, 2, \ldots$: explicit $(2j+1)\times(2j+1)$ matrices.

All exact finite linear algebra. Type I.

4.3.2 Addition of angular momenta

Two systems with angular momenta j_1 and j_2 have combined Hilbert space:

$$\mathcal{H} = \mathcal{H}_{j_1} \otimes \mathcal{H}_{j_2}$$

with dimension $(2j_1+1)(2j_2+1)$.

By Maschke's theorem (Theorem 13.3 of the AFB paper), this reducible representation decomposes into irreducibles:

$$\mathcal{H}_{j_1} \otimes \mathcal{H}_{j_2} = \bigoplus_{J=|j_1-j_2|}^{j_1+j_2} \mathcal{H}_J$$

The multiplicities are computed by the character inner product (Theorem 13.6). In this case, each J appears exactly once (the Clebsch-Gordan decomposition).

The Clebsch-Gordan coefficients $\langle j_1m_1; j_2m_2|JM\rangle$ are the matrix elements of the unitary transformation between the product basis $\{|j_1m_1\rangle \otimes |j_2m_2\rangle\}$ and the coupled basis $\{|JM\rangle\}$. They are specific elements of $\mathbb{R}_B(k)$, computable by the standard recursion relations (which are finite linear recurrences, solvable by BI-BST).

For $j_1 = j_2 = 1/2$:
$$\mathcal{H}_{1/2} \otimes \mathcal{H}_{1/2} = \mathcal{H}_1 \oplus \mathcal{H}_0$$

(triplet + singlet). The Clebsch-Gordan coefficients are $\pm 1/\sqrt{2}$ and ± 1 — computable in $\mathbb{R}_B(k)$. Type I.

4.3.3 Selection rules

A transition between states in representation ρ_i and ρ_j, mediated by an operator transforming under representation ρ_k, is nonzero only if ρ_j appears in the decomposition of $\rho_i \otimes \rho_k$.

The multiplicity is:

$$m = \langle \chi_i \cdot \chi_k, \chi_j \rangle$$

(character inner product, Theorem 13.4 of the AFB paper). This is a finite sum — decidable.

If $m = 0$, the transition is forbidden.
If $m > 0$, the transition is allowed.

This is the mathematical basis of all spectral selection rules: electric dipole ($\Delta l = \pm 1$), magnetic dipole, quadrupole, etc. Each rule is a decidable finite computation on characters. Type I.

4.4 Composite Systems and Entanglement

Definition 4.1 — Bounded tensor product state space:

For systems A and B with Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$:

$$\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$$

(Definition 11.10 of the AFB paper).
$\dim(\mathcal{H}_{AB}) = d_A \cdot d_B$.

A product state: $|\psi\rangle_A \otimes |\varphi\rangle_B$.
An entangled state: any state in $\mathcal{H}_{AB}$ that cannot be written as a product.

Example (Bell state):

$$|\Phi^+\rangle = (1/\sqrt{2})(|00\rangle + |11\rangle)$$

This is a specific vector in $\mathbb{C}_B(k^4)^4$. The coefficient $1/\sqrt{2}$ is computed in $\mathbb{R}_B(k)$ by Newton-Raphson (Part VIII §8.7). Type I.

Definition 4.2 — Bounded density matrix:

For a composite system in state $|\psi_{AB}\rangle$, the reduced density matrix of subsystem A is:

$$\rho_A = \mathrm{Tr}_B(|\psi_{AB}\rangle\langle\psi_{AB}|)$$

where Tr_B is the partial trace over $\mathcal{H}_B$ — a finite sum of d_B terms, each a $d_A \times d_A$ matrix over $\mathbb{C}\_B(k^4)$. The result is a $d_A \times d_A$ Hermitian matrix with $Tr(\rho_A) = 1$. Type I.

Definition 4.3 — Von Neumann entropy:

The entanglement entropy of subsystem A is:

$$S(\rho_A) = -Tr(\rho_A \ln \rho_A)$$
$$= -\Sigma_i \lambda_i \ln \lambda_i$$

where $\{\lambda_i\}$ are the eigenvalues of ρ_A, computed by the spectral theorem (Theorem 11.4).

This is the Shannon entropy (Definition 9.11 of the AFB paper) of the eigenvalue distribution. It is a specific element of $\mathbb{R}\_B(k)$. Type I.

$S = 0$ iff ρ_A is a pure state (product state).
$S = \ln d_A$ iff $\rho_A = I/d_A$ (maximally entangled).

Theorem 4.1 — Bell inequality violation in BST:

For two qubits in the Bell state $|\Phi^+\rangle = (1/\sqrt{2})(|00\rangle + |11\rangle)$, the CHSH combination:

$$S = \langle A_1B_1\rangle + \langle A_1B_2\rangle + \langle A_2B_1\rangle - \langle A_2B_2\rangle$$

where A_1, A_2 are observables on qubit A and B_1, B_2 on qubit B, satisfies:

$$\max S = 2\sqrt{2} \approx 2.828$$

(the Tsirelson bound). This exceeds the classical Bell limit of 2.

Proof: Direct computation. Choose $A_1 = \sigma_z$, $A_2 = \sigma_x$, $B_1 = (\sigma_z + \sigma_x)/\sqrt{2}$, $B_2 = (\sigma_z - \sigma_x)/\sqrt{2}$.
Each expectation value is a 4×4 matrix trace (a finite sum of 4 terms in $\mathbb{C}\_B(k^4)$).

$$\langle A_1B_1\rangle = Tr(|\Phi^+\rangle\langle\Phi^+| \cdot A_1\otimes B_1) = 1/\sqrt{2}$$

Similarly for the others. Sum: $S = 2\sqrt{2}$. □

Type I. Bell violation is an exact algebraic fact about 4×4 matrices over $\mathbb{C}\_B(k^4)$. It does not require infinite-dimensional Hilbert spaces, continuous variables, or locality loopholes. The quantum nature of entanglement is a finite-dimensional phenomenon.

4.5 Perturbation Theory

Classical perturbation theory in quantum mechanics uses formal power series that may not converge. In BST, perturbation theory is replaced by exact, computable bounds — the Weyl and Davis-Kahan theorems from Part XI of the AFB paper.

Theorem 4.2 — Bounded eigenvalue perturbation:
Let H_0 be a Hamiltonian with known eigenvalues $\{E_i^0\}$ and let V be a perturbation with $\|V\| = \varepsilon$. Then the eigenvalues $\{E_i\}$ of $H = H_0 + V$ satisfy:

$$|E_i - E_i^0| \leq \varepsilon \quad \text{for each } i$$

(Weyl's inequality, Theorem 11.5 of the AFB paper).

This is an exact bound – not a first-order approximation, not a formal power series, not an asymptotic statement. It holds for any ε, not just small ε. It is computable: $\varepsilon = \|V\|$ is a finite maximum (Definition 11.5). Type I.

Theorem 4.3 — Bounded eigenspace perturbation:
If H_0 has an eigenspace S_0 separated from the rest of the spectrum by a gap $\gamma > 0$, and $\|V\| < \gamma/2$, then $H = H_0 + V$ has an eigenspace S near S_0 with:

$$\|P_S - P_{S_0}\| \leq \|V\|/\gamma$$

(Davis-Kahan theorem, Theorem 11.6 of the AFB paper).

This controls the rotation of eigenspaces under perturbation. For quantum mechanics: it guarantees that energy levels that are well-separated from other levels are stable under small perturbations, with a computable stability bound.

Application: the Zeeman effect. A magnetic field perturbs the hydrogen Hamiltonian. If the field strength corresponds to $\|V\| = \varepsilon$, and the unperturbed energy gap is γ, then the perturbed energy levels differ from the unperturbed ones by at most ε, and the perturbed eigenstates differ by at most ε/γ. Both computable. Type I.

Plain language: Perturbation theory in BST does not use formal power series or Feynman diagrams (those come in the gauge theory of Part V and the Standard Model programme of Part IX). It uses exact computable bounds — the Weyl and Davis-Kahan inequalities — that tell you exactly how much the eigenvalues and eigenspaces of a Hamiltonian can shift when a perturbation is added. These bounds hold for perturbations of any size, not just small ones. They are the finite-dimensional content of perturbation theory, and they are strictly more informative than the classical formal series (which may diverge).

4.6 Bounded Scattering Theory

Definition 4.4 — Bounded S-matrix:
For a quantum system with a finite space of asymptotic states $\mathcal{H}_{asym}$ (incoming and outgoing particle states), the S-matrix is a unitary operator:

$$S: \mathcal{H}_{asym} \to \mathcal{H}_{asym}$$

S is a d × d unitary matrix over $\mathbb{C}_B(k^4)$, where $d = \dim(\mathcal{H}_{asym})$.

The matrix element $S_{fi} = \langle f|S|i\rangle$ is the transition amplitude from initial state $|i\rangle$ to final state $|f\rangle$. The transition probability is:

$$P(i \to f) = |S_{fi}|^2$$

This is a bounded probability distribution over the finite set of final states (Definition 9.7 of the AFB paper): $\Sigma_f |S_{fi}|^2 = 1$ (unitarity of S). Type I.

Definition 4.5 — Bounded cross-section:
The differential cross-section for scattering from initial state $|i\rangle$ to final state $|f\rangle$ is proportional to the transition probability:

$$d\sigma/d\Omega \propto |\langle f|T|i\rangle|^2$$

where $T = S - I$ is the transfer matrix and the proportionality involves kinematic factors (momenta, phase space volume) computable in $\mathbb{R}_B(k)$.

The total cross-section is a finite sum over all final states:

$$\sigma_{total} = \Sigma_f \sigma(i \to f)$$

This is a finite sum of positive real numbers in $\mathbb{R}_B(k)$. Type I.

Connection to experiment: the cross-section is the quantity measured in particle physics experiments. It is a specific computable number for any specific initial and final state configuration. The infinite-dimensional scattering theory of classical QM (Lippmann-Schwinger equation, Born series) is the infinite-dimensional idealisation of this finite matrix computation.

4.7 Quantum-Classical Correspondence

4.7.1 Ehrenfest's theorem

Theorem 4.4 — Bounded Ehrenfest:
For a quantum system with Hamiltonian H and observable A (both Hermitian matrices on $\mathcal{H}$):

$$d\langle A\rangle/dt = (1/i\hbar)\langle[A, H]\rangle + \langle\partial A/\partial t\rangle$$

where $\langle\cdot\rangle = \langle\psi(t)|\cdot|\psi(t)\rangle$ and $[A,H] = AH - HA$.

Proof: $d/dt\langle\psi|A|\psi\rangle = \langle d\psi/dt|A|\psi\rangle + \langle\psi|A|d\psi/dt\rangle + \langle\psi|\partial A/\partial t|\psi\rangle$
Using $i\hbar d|\psi\rangle/dt = H|\psi\rangle$:
$= (1/i\hbar)(\langle\psi|HA|\psi\rangle - \langle\psi|AH|\psi\rangle) + \langle\partial A/\partial t\rangle$
$= (1/i\hbar)\langle[A,H]\rangle + \langle\partial A/\partial t\rangle$. □

Type I. An exact finite-matrix identity.

For $A = \hat{x}$: $d\langle x\rangle/dt = \langle p\rangle/m$ (if $H = p^2/2m + V(x)$).
For $A = \hat{p}$: $d\langle p\rangle/dt = -\langle\partial V/\partial x\rangle$.

These are the quantum analogues of Hamilton's equations (Part I, §1.3). The quantum expectation values obey the same equations as the classical variables — up to quantum corrections from the non-commutativity of $\hat{x}$ and $V'(\hat{x})$.

4.7.2 The Poisson bracket correspondence

The structural parallel between classical and quantum mechanics, named in Part I §1.5.1, can now be made precise:

Classical: $\{f, g\} = \Sigma\ (\partial f/\partial q \cdot \partial g/\partial p - \partial f/\partial p \cdot \partial g/\partial q)$
Quantum: $[\hat{f}, \hat{g}]/(i\hbar) = (\hat{f}\hat{g} - \hat{g}\hat{f})/(i\hbar)$

The map $\{f,g\} \to [\hat{f}, \hat{g}]/(i\hbar)$ is Dirac's canonical quantisation. It maps:

$\{q, p\} = 1 \quad\to\quad [\hat{q}, \hat{p}]/(i\hbar) = I$
$\{f, H\} = df/dt \quad\to\quad [\hat{f}, \hat{H}]/(i\hbar) = df/dt$ (Ehrenfest)

In BST, both sides are exact finite computations: the Poisson bracket is a finite sum of finite differences (Definition 1.10), and the commutator is a finite matrix multiplication. The correspondence is an exact structural map between two finite algebraic systems. Type I.

4.8 Quantum Information

Quantum information theory is among the strongest applications of BST to physics, because its fundamental objects — qubits, gates, circuits — are already finite.

4.8.1 Qubits, gates, and circuits

A qubit is a 2-dimensional Hilbert space:
$\mathcal{H} = \mathbb{C}_B(k^4)^2$. (§4.2.1.)

A quantum gate is a unitary matrix on a qubit register of n qubits:

$U: \mathbb{C}_B(k^4)^{2^n} \to \mathbb{C}_B(k^4)^{2^n}$

Standard gates (all 2×2 or 4×4 unitary matrices):

Hadamard: $H = (1/\sqrt{2})\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$

CNOT: $\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$

Phase: $S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$ $T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$

Each is a specific matrix over $\mathbb{C}_B(k^4)$. Type I.

A quantum circuit is a finite sequence of gates:
$U = U_m \cdot U_{m-1} \cdot \ldots \cdot U_1$

The circuit is a finite composition of finite matrices — a single unitary matrix on $\mathbb{C}_B(k^4)^{2^n}$. Type I.

Connection to Part XII of the AFB paper: the computational complexity of quantum circuits is classified by BST-BQP (bounded-error quantum polynomial time), which is a uniform family of bounded quantum computations. The relationship between BST-BQP and BST-P/BST-NP is the bounded analogue of the classical BQP vs P/NP question.

4.8.2 Quantum teleportation

The teleportation protocol for one qubit:

Setup: Alice and Bob share the Bell state
$|\Phi^+\rangle_{23} = (1/\sqrt{2})(|00\rangle + |11\rangle)$ on qubits 2 and 3.
Alice has qubit 1 in unknown state
$|\psi\rangle_1 = \alpha|0\rangle + \beta|1\rangle$.

The total state is:
$|\psi\rangle_1 \otimes |\Phi^+\rangle_{23} \in \mathbb{C}_B(k^4)^8$

Protocol:
1. Alice performs a Bell measurement on qubits 1,2.

(Project onto the 4 Bell states — a finite measurement with 4 outcomes.)
2. Alice communicates the 2-bit outcome to Bob.
3. Bob applies the corresponding Pauli correction (I, σ_x, σ_y, or σ_z) to qubit 3.

Result: Bob's qubit is in state $|\psi\rangle$. ✓

Verification: direct 8×8 matrix computation in $\mathbb{C}_B(k^4)$. The protocol works exactly. Type I.

No faster-than-light communication: Alice's measurement outcome is uniformly random ($P = 1/4$ for each of the 4 outcomes). The quantum state is transferred, but no information travels without the classical 2-bit message.

4.8.3 No-cloning theorem

Theorem 4.5 — No-cloning (BST):
There is no unitary operator U on $\mathcal{H} \otimes \mathcal{H}$ such that $U(|\psi\rangle \otimes |0\rangle) = |\psi\rangle \otimes |\psi\rangle$ for all $|\psi\rangle \in \mathcal{H}$.

Proof: Suppose such U exists. Let $|\psi\rangle$ and $|\varphi\rangle$ be two non-orthogonal states ($\langle\psi|\varphi\rangle \neq 0, 1$).

$$\begin{aligned}
\langle\psi|\varphi\rangle &= \langle\psi|\langle 0| \cdot |\varphi\rangle|0\rangle && \text{(inner product)} \\
&= \langle\psi|\langle 0|U^\dagger U|\varphi\rangle|0\rangle && (U^\dagger U = I\text{, unitarity}) \\
&= (\langle\psi|\langle\psi|)(|\varphi\rangle|\varphi\rangle) && \text{(by hypothesis)} \\
&= (\langle\psi|\varphi\rangle)^2
\end{aligned}$$

So $\langle\psi|\varphi\rangle = (\langle\psi|\varphi\rangle)^2$. This implies $\langle\psi|\varphi\rangle = 0$ or 1. Contradiction with the assumption $0 < |\langle\psi|\varphi\rangle| < 1$. □

Type I. The proof uses only unitarity (finite matrix property) and the inner product (Definition 11.7). No infinite-dimensional arguments.

4.8.4 Quantum error correction

The simplest quantum error-correcting code: the 3-qubit bit-flip code.

Encoding: $|0\rangle \to |000\rangle$, $|1\rangle \to |111\rangle$.
Logical qubit: $\alpha|000\rangle + \beta|111\rangle \in \mathbb{C}_B(k^4)^8$.

Error model: a single qubit flips (σ_x on one qubit).

Syndrome measurement: measure the parity operators Z_1Z_2 and Z_2Z_3 (products of Pauli-Z on pairs). The 4 possible syndromes identify which qubit (if any) has flipped — a finite measurement with 4 outcomes. Apply the correcting σ_x. The logical qubit is restored.

This is a finite group-theoretic construction: the stabiliser group is generated by Z_1Z_2 and Z_2Z_3 (a finite group of order 4 acting on 8-dimensional Hilbert space). The code space is the +1 eigenspace of both generators — a 2-dimensional subspace found by finite projection. Type I.

More general stabiliser codes (the Steane code, the surface code) follow the same pattern: finite groups acting on finite Hilbert spaces, with code spaces defined by joint eigenspaces. All exact finite linear algebra. Part XIII's representation theory provides the decomposition.

4.9 The Spectral Verification Protocol

This section collects and generalises the spectral approach that applies across quantum physics.

Definition 4.6 — Spectral verification:
Given any bounded Hamiltonian H on a finite Hilbert space $\mathcal{H}$ of dimension d:

1. Diagonalise H by the spectral theorem (Theorem 11.4).
   Cost: $O(d^3)$ operations in $\mathbb{C}\_B(k^4)$.
   Result: eigenvalues $\lambda_0 \le \lambda_1 \le \ldots \le \lambda_{d-1}$
   and eigenvectors $\{|\lambda_i\rangle\}$.

2. The spectral gap is $\gamma = \lambda_1 - \lambda_0$.
   Whether $\gamma > 0$ is decidable (comparison in $\mathbb{R}\_B(k)$).

3. The ground state energy is λ_0.
   The ground state is $|\lambda_0\rangle$.

4. Any observable $\langle A\rangle$ in the ground state is
   $\langle\lambda_0|A|\lambda_0\rangle$ — one matrix element.

Every spectral question about a finite quantum system is a finite linear algebra computation. The answers are specific elements of $\mathbb{R}\_B(k)$ or $\mathbb{C}\_B(k^4)$, computable to the precision of the bounded field.

This protocol applies uniformly to:
- Energy levels of atoms and molecules
- Mass gaps of lattice gauge theories (Part V)
- Band structures of crystals (Part VIII of this volume)
- Ground state properties of condensed matter systems

The fixed-dimension computation is always decidable. The family-level question — how do these quantities behave as the dimension grows through the metatheory — is the genuine open problem in each case.

```
Type I for the finite computation. The family-level
behaviour is the content of the specific physics
problem.
```

4.10 What Part IV Recovers

Result	Type	Section
Quantum states as finite vectors	I	4.1
Observables as Hermitian matrices	I	4.1
Born rule probabilities	I	4.1
Unitary time evolution	I/III	4.1
Composite systems (tensor product)	I	4.1
Qubit and Pauli matrices	I	4.2.1
Harmonic oscillator spectrum	I	4.2.2
Hydrogen atom energy levels	I	4.2.3
Angular momentum representations	I	4.3.1
Clebsch-Gordan decomposition	I	4.3.2
Selection rules (decidable)	I	4.3.3
Entanglement and Bell states	I	4.4
Density matrices and partial trace	I	4.4
Von Neumann entropy	I	4.4
Bell inequality violation (2√2)	I	4.4
Weyl eigenvalue bound	I	4.5
Davis-Kahan eigenspace stability	I	4.5
S-matrix and cross-sections	I	4.6
Ehrenfest's theorem	I	4.7.1
Poisson-commutator correspondence	I	4.7.2
Quantum gates and circuits	I	4.8.1
Quantum teleportation	I	4.8.2
No-cloning theorem	I	4.8.3
Quantum error correction	I	4.8.4
Spectral verification protocol	I	4.9

What is NOT recovered: infinite-dimensional Hilbert spaces $L^2(\mathbb{R})$ or $L^2(\mathbb{R}^3)$ (replaced by finite-dimensional truncations); continuous spectra (all spectra are discrete and finite); unbounded operators (all operators are finite matrices with computable norms); the Stone-von Neumann theorem in its infinite-dimensional form (the canonical commutation relation holds exactly on the finite truncation minus a boundary correction); quantum field theory (requires Part V's gauge theory and the infinite-dimensional Fock space construction, which BST replaces with finite lattice Hilbert spaces).

What is gained: every eigenvalue is computable, every spectral gap is decidable, every measurement probability is a specific element of $\mathbb{R}$_B(k), every time evolution is a finite matrix exponential. There are no domain questions, no self-adjointness problems, no renormalisation needed for finite quantum systems. The mathematical structure of quantum mechanics is cleaner and more computationally explicit in BST than in the infinite-dimensional formulation.

End of Part IV

Part V: Gauge Theory and Yang-Mills

Maturity: Tier 1

AFB imports: $\mathbb{C}$_B(k^4) (Part VIII, Theorem 8.16); bounded linear operators and matrix algebra (Part XI, §§11.2, 11.6); spectral theorem (Part XI, Theorem 11.4); finite Hilbert spaces (Part XI, §11.5, Definition 11.8); matrix exponential (Part XI, §11.7, Definition 11.16); representation theory including Maschke, characters, irreducible decomposition (Part XIII, Theorems 13.3–13.7); bounded probability and expectation (Part IX, §9.3.4)

Volume imports: Part II §2.5 (U(1)_B lattice gauge theory as template and special case); Part IV (finite Hilbert spaces for the quantum gauge theory, spectral verification protocol §4.9)

The physics

Gauge theory is the mathematical framework underlying all three fundamental forces described by the Standard Model. Electromagnetism is the gauge theory of the group U(1). The weak nuclear force (responsible for radioactive beta decay) is the gauge theory of SU(2). The strong nuclear force (which binds quarks into protons, neutrons, and atomic nuclei) is the gauge theory of SU(3) — quantum chromodynamics (QCD). The Standard Model combines all three as a single gauge theory with gauge group SU(3) × SU(2) × U(1). Part II constructed the U(1) case. This Part generalises to SU(N), covering the full gauge-theoretic structure of the nuclear forces.

What experiments confirm

The strong force. QCD predicts that quarks are confined inside hadrons — they cannot be isolated. This is confirmed by the absence of free quarks in all experiments. Lattice QCD — the exact computational framework this Part constructs — predicts the masses of protons, neutrons, and other hadrons from the quark masses and the SU(3) coupling constant. The proton mass (938.3 MeV) has been computed by lattice QCD to within 2% (BMW Collaboration 2008), and with subsequent improvements to sub-percent precision — computed by exactly the finite-sum partition function of §5.4.

Asymptotic freedom. The SU(3) coupling constant decreases at high energies — quarks interact more weakly when close together. This prediction (Gross, Wilczek, Politzer 1973, Nobel Prize 2004) is confirmed by deep inelastic scattering experiments at SLAC, HERA, and the LHC, with the running coupling measured across three orders of magnitude in energy and matching the prediction to a few percent at each scale.

The weak force. The W and Z bosons predicted by SU(2) × U(1) gauge theory were discovered at CERN in 1983 (Rubbia and van der Meer, Nobel Prize 1984) with masses matching the theoretical predictions: $m_W \approx 80.4$ GeV, $m_Z \approx 91.2$ GeV. The electroweak mixing angle $\sin^2\theta_W \approx 0.231$ is measured consistently across dozens of independent experiments.

Wilson loops and confinement. The area-law scaling of Wilson loops (§5.3), which diagnoses quark confinement, is confirmed by lattice Monte Carlo simulations — the same finite-sum computations that §5.4 constructs. The string tension $\sigma \approx (440\ \text{MeV})^2$ extracted from these computations matches the Regge slopes of hadronic spectra.

The mass gap. Whether SU(N) Yang-Mills theory has a positive mass gap is one of the seven Clay Millennium Problems. Lattice computations (the spectral verification of §5.6) provide strong numerical evidence for a gap of ~1.5 GeV in pure SU(3) gauge theory — but a rigorous proof remains open.

What BST constructs and why it suffices

Lattice gauge theory — the framework constructed in this Part — is how the strong force is actually computed. The BMW proton mass calculation, the ALPHA Collaboration's determination of the strong coupling, and every lattice QCD result published by FLAG are computed on finite lattices with finite gauge groups, finite partition functions, and finite matrix algebra. The continuum formulation of Yang-Mills theory (gauge fields on smooth $\mathbb{R}^4$) is the theoretical idealisation; the lattice computation is the mathematical substance.

BST makes this explicit. Gauge connections are group-valued edge variables on a finite lattice. The action is a finite sum. The partition function is a finite sum. The gauge-field Hilbert space is finite-dimensional. Gauss's law is implemented by a finite group projection (Maschke's theorem). The spectral gap is computable by diagonalising a finite matrix. The generalisation from U(1) (Part II §2.5) to SU(N) changes only the group — the framework is identical.

5.1 Bounded Gauge Groups

Definition 5.1 — Bounded matrix gauge groups:

```
For N ≥ 1, define the following finite matrix groups
over ℂ_B(k⁴):

U(N)_B = {U ∈ M_N(ℂ_B(k⁴)) : U†U = I}
  (unitary matrices)

SU(N)_B = {U ∈ U(N)_B : det(U) = 1}
  (special unitary matrices)

Here U† = (U*)^T is the conjugate transpose, I is
the N×N identity, and det is the determinant —
all finite computations over ℂ_B(k⁴).

The conditions U†U = I and det(U) = 1 are N² + 1
polynomial equations in the N² complex entries of U.
Membership is decidable. Type I.

Group properties (all by direct finite matrix
computation):

Closure: If U, V ∈ SU(N)_B, then
  (UV)†(UV) = V†U†UV = V†V = I ✓
  det(UV) = det(U)det(V) = 1 ✓
  So UV ∈ SU(N)_B. ✓

Identity: I ∈ SU(N)_B. ✓
```

Inverse: $U^{-1} = U^\dagger \in SU(N)\_B$ (since $(U^\dagger)^\dagger U^\dagger = UU^\dagger = I$ and $\det(U^\dagger) = \det(U)^* = 1$). ✓

Associativity: $(UV)W = U(VW)$ — matrix multiplication is associative in $\mathbb{C}\_B(k^4)$ up to $O(1/k^2)$ per operation (Theorem 8.16). For N×N matrices, each product involves $O(N^3)$ field operations, so a triple product carries $O(N^3/k^2)$ rounding. Type I for the algebraic structure; Type III for the arithmetic. ✓

Cardinality: $|SU(N)\_B| \leq |M\_N(\mathbb{C}\_B(k^4))| = k^{4N^2}$. The actual cardinality — the number of matrices satisfying the unitarity and determinant conditions within the bounded field — depends on k and N. Existence condition: $k^{4N^2} \leq n\_M$.

Definition 5.2 — Bounded Lie algebra:
The Lie algebra su(N) is the space of N×N traceless anti-Hermitian matrices over $\mathbb{C}\_B(k^4)$:

$$su(N)\_B = \{X \in M\_N(\mathbb{C}\_B(k^4)) : X^\dagger = -X, \mathrm{Tr}(X) = 0\}$$

This is a real vector space of dimension $N^2 - 1$.

A basis $\{T_a\}_{a=1}^{N^2-1}$ can be chosen with the commutation relations:

$$[T_a, T_b] = \Sigma_c f_{abc} T_c$$

where f_{abc} are the structure constants — specific rational numbers (for the standard Gell-Mann basis of su(3) or the standard basis of su(2), the f_{abc} are integers or simple fractions). The structure constants are exact. Type I.

The exponential map: for $X \in su(N)\_B$, $\exp(X) \in SU(N)\_B$ (by Theorem 11.11(iv) of the AFB paper, since iX is Hermitian). This connects the algebra to the group.

The physically important cases:

$SU(2)\_B$: $\dim(su(2)) = 3$. Generators: the Pauli matrices $\sigma_i/2$. Structure constants: $f_{ijk} = \varepsilon_{ijk}$. This is the gauge group of the weak nuclear force.

$SU(3)\_B$: $\dim(su(3)) = 8$. Generators: the Gell-Mann matrices $\lambda_a/2$. Structure constants: the standard f_{abc} of su(3). This is the gauge group of the strong nuclear force (quantum chromodynamics).

Definition 5.3 — The Standard Model gauge group:
$G\_{SM} = SU(3)\_B \times SU(2)\_B \times U(1)\_B$

This is the direct product of three bounded gauge groups. An element is a triple (g_3, g_2, g_1) with $g_3 \in SU(3)\_B$, $g_2 \in SU(2)\_B$, $g_1 \in U(1)\_B$.

Multiplication: componentwise.

$$(g_3, g_2, g_1)\cdot(h_3, h_2, h_1) = (g_3h_3, g_2h_2, g_1h_1)$$

Cardinality: $|G\_SM| \leq |SU(3)\_B| \cdot |SU(2)\_B| \cdot |U(1)\_B|$
$\leq k^{36} \cdot k^{16} \cdot k^4 = k^{56}$.

Existence condition: $k^{56} \leq n\_M$. For the holographic benchmark $n\_M \approx 10^{185}$, this gives $k \leq 10^{185/56} \approx 10^{3.3} \approx 2000$. This is modest – the Standard Model gauge group exists comfortably within the holographic bound.

Representation content of one generation of fermions (Part XIII, §13.5):

- Left-handed quarks: (3, 2, 1/6)
- Right-handed up quarks: (3, 1, 2/3)
- Right-handed down quarks: (3, 1, −1/3)
- Left-handed leptons: (1, 2, −1/2)
- Right-handed electron: (1, 1, −1)

Each entry (r_3, r_2, Y) specifies irreducible representations of $SU(3)\_B$, $SU(2)\_B$, and the $U(1)\_B$ hypercharge. These are finite-dimensional vector spaces with specified group actions – all internal to BST. Type I for the algebraic data.

5.2 Lattice Gauge Theory: The General Construction

Definition 5.4 — Bounded lattice:

A lattice (or simplicial complex) K consists of:
- A finite set V(K) of vertices (sites)
- A finite set E(K) of oriented edges (links)
- A finite set P(K) of oriented plaquettes (faces)
- Orientation conventions: each edge e has a reverse e^{-1} with the opposite orientation.

For a d-dimensional hypercubic lattice of side L with spacing a:

$|V(K)| = (L/a)^d$
$|E(K)| = d \cdot (L/a)^d$ (one link per direction per site)
$|P(K)| = d(d-1)/2 \cdot (L/a)^d$ (one plaquette per plane per site)

All finite. Type I.

Definition 5.5 — Bounded connection (gauge field):

A lattice gauge connection on K with gauge group G_B is a map:

$$U: E(K) \to G\_B$$

assigning a group element to each oriented edge, with the convention:

$U(e^{-1}) = U(e)^{-1} = U(e)^\dagger$

The configuration space is:
$A_K = G_B^{E(K)}$

This is the set of all connections — a finite set. $|A_K| = |G_B|^{|E(K)|}$.

For SU(3) on a 4^4 lattice (4 sites per direction, 4 dimensions):

$|E(K)| = 4 \cdot 4^4 = 1024$ links
$|A_K| = |SU(3)_B|^{1024}$

This is large but finite and well-defined. Type I.

Definition 5.6 — Plaquette holonomy:
For an oriented plaquette p with boundary edges e_1, e_2, e_3, e_4 (traversed in order), the plaquette holonomy is the ordered product:

$U_p = U(e_1) \cdot U(e_2) \cdot U(e_3) \cdot U(e_4)$

This is a product of four N×N matrices in G_B. The result $U_p \in M_N(\mathbb{C}_B(k^4))$ — it is close to an element of G_B for smooth configurations, but the product is computed exactly as matrix multiplication.

For non-abelian groups, the order matters: $U(e_1)U(e_2) \neq U(e_2)U(e_1)$ in general. This is the fundamental difference from the U(1) case of Part II, where all products commute.

The plaquette holonomy measures the local curvature of the gauge field. For a slowly varying field ($U(e) \approx I + iaA_\mu T_a$):

$U_p \approx I + ia^2F_{\mu\nu} T_a + O(a^4)$

where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + ig[A_\mu, A_\nu]$ is the non-abelian field strength. The commutator term $[A_\mu, A_\nu]$ — absent in the abelian case — is the hallmark of non-abelian gauge theory.

Definition 5.7 — Wilson action:
The Wilson action for gauge group $G_B = SU(N)_B$ is:

$S[U] = \beta \sum_{p \in P(K)} (1 - (1/N) \text{ Re } \text{Tr}(U_p))$

where:
- $\beta = 2N/g^2$ is the inverse coupling (g is the gauge coupling constant)
- Tr is the matrix trace in the fundamental (N-dimensional) representation
- Re is the real part
- The sum runs over all plaquettes

$S[U] \geq 0$, with $S = 0$ iff all $U_p = I$ (pure gauge / zero field strength).

S is a finite sum of $|P(K)|$ terms, each involving one N×N matrix trace ($O(N)$ operations). Total cost: $O(|P(K)| \cdot N)$ operations in $\mathbb{C}_B(k^4)$. Type I.

For $G_B = U(1)_B$: $Tr(U_p) = U_p$ (a single complex number), and the action reduces to the U(1) Wilson action of Part II §2.5 (Definition 2.12). ✓

5.3 Gauge Transformations and Invariance

Definition 5.8 — Lattice gauge transformation:
A gauge transformation is a map:
$g: V(K) \to G_B$

assigning a group element to each vertex.

It acts on connections by conjugation:
$U(e) \to g(\mathrm{source}(e)) \cdot U(e) \cdot g(\mathrm{target}(e))^{-1}$

The plaquette holonomy transforms as:
$U_p \to g(v_0) \cdot U_p \cdot g(v_0)^{-1}$

where v_0 is the base vertex of the plaquette.

Consequence: $Tr(U_p) \to Tr(g\ U_p\ g^{-1}) = Tr(U_p)$ by the cyclic property of trace. ✓

Therefore the Wilson action S[U] is gauge-invariant:
$S[g \cdot U] = S[U]$ for all gauge transformations g. ✓

Proof: Each term $(1/N)\mathrm{Re}\,Tr(U_p)$ is gauge-invariant, and S is a sum of such terms. □

Type I. Gauge invariance is an exact algebraic consequence of the cyclic property of trace — the same proof as in the U(1) case (Part II, Definition 2.13), but now using the cyclic property for N×N matrices.

Definition 5.9 — Wilson loop observable:
For a closed loop γ on the lattice (a sequence of edges forming a closed path), the Wilson loop is:

$$W(\gamma) = (1/N)\,\mathrm{Tr}(\prod_{e \in \gamma} U(e))$$

This is the trace of the ordered product of group elements around γ. For a loop of length L_γ edges, it involves L_γ matrix multiplications and one trace.

Gauge transformation:
$$W(\gamma) \to (1/N)\,\mathrm{Tr}(g(v_0) \cdot \prod U(e) \cdot g(v_0)^{-1}) = W(\gamma)$$

by the cyclic property. Wilson loops are gauge-invariant. ✓

Physical interpretation: W(γ) is the order parameter for confinement. For a rectangular loop of spatial extent R and temporal extent T:

$$\langle W(R,T)\rangle \sim \exp(-\sigma RT) \quad \text{(area law → confinement)}$$
$$\langle W(R,T)\rangle \sim \exp(-\mu(R+T)) \quad \text{(perimeter law → deconfinement)}$$

where σ is the string tension and μ is a mass scale. These scaling behaviours are computable from the finite partition function (§5.4). Type I for the computation; the interpretation as confinement is physical.

5.4 Partition Function and Expectation Values

Definition 5.10 — Bounded Yang-Mills partition function:
$$Z = \sum_{U \in A_K} \exp(-S[U])$$

This is a finite sum over the finite configuration space $A_K = G_B^{E(K)}$.

$|A_K|$ terms. Each term requires:
- $O(|P(K)| \cdot N)$ operations to evaluate S[U]
- O(1) operation to evaluate exp(−S[U])

Total cost: $O(|A_K| \cdot |P(K)| \cdot N)$ operations.

Z is a specific positive element of $\mathbb{R}_B(k)$.
It always exists and is always positive (each term in the sum is positive). No regularisation. No renormalisation. No functional integral. Type I.

Expectation values:
$$\langle O\rangle = (1/Z) \sum_{U \in A_K} O[U] \cdot \exp(-S[U])$$

For any gauge-invariant observable O (e.g. a Wilson loop), this is a finite sum of finite products. Type I.

Note on computational cost:
For SU(3) on a 4^4 lattice, $|A\_K| = |SU(3)\_B|^{1024}$. This is astronomically large — direct summation is infeasible for all but the smallest lattices. In practice, the partition function is estimated by Monte Carlo sampling (Part III, §3.8), adapted to the gauge theory setting:

The Metropolis algorithm on A_K:

1. Start from a random connection U_0.
2. At each step, propose a local update: change $U(e) \to U'(e)$ for a single edge e.
3. Accept/reject by the Metropolis criterion with weight $\exp(-S)$.

The transition matrix is doubly stochastic with respect to the Wilson action weight. Convergence is guaranteed by the spectral gap of the transition operator (Part III, §3.8). Type I.

This is exactly what lattice QCD practitioners do on computers today — the BST formulation makes explicit that the mathematical foundations of these computations are entirely finite.

5.5 The Gauge-Field Hilbert Space

Definition 5.11 — Bounded gauge-field quantum state space:
The quantum state space of the lattice gauge theory is the finite Hilbert space:

$$\mathscr{H}\_K = L^2(A\_K)$$

with basis states $\{|U\rangle\}_{U \in A\_K}$ — one basis vector for each classical connection configuration — and the counting-measure inner product:

$$\langle U|V\rangle = \delta_{U,V}$$

$$\dim(\mathscr{H}\_K) = |A\_K| = |G\_B|^{|E(K)|}.$$

This is a bounded Hilbert space (Definition 11.8 of the AFB paper). Type I.

Definition 5.12 — Bounded gauge-theory Hamiltonian:
The lattice gauge Hamiltonian in the Kogut-Susskind formulation consists of two terms:

$$H = H\_E + H\_B$$

Electric term:
$H_E = (g^2/2a) \Sigma_{e \in E(K)} \Sigma_a E_a(e)^2$

where $E_a(e)$ are the "electric field" operators — the left-invariant vector fields on G_B, acting on $\mathscr{H}_K$ as finite difference operators in the group variable $U(e)$.

Magnetic term:
$H_B = (1/(g^2 a)) \Sigma_{p \in P(K)} (N - \mathrm{Re}\,\mathrm{Tr}(U_p))$

where U_p acts on $\mathscr{H}_K$ by multiplication: for each plaquette p, the operator $\mathrm{Tr}(U_p)$ multiplies the wave function $\Psi(U)$ by $\mathrm{Tr}(U_p)$.

Both H_E and H_B are Hermitian operators on the finite Hilbert space $\mathscr{H}_K$. H is Hermitian. The spectral theorem (Theorem 11.4 of the AFB paper) applies: H has a complete set of real eigenvalues and orthonormal eigenstates. Type I.

5.6 The Mass-Gap Problem

The Yang-Mills mass-gap problem is one of the seven Clay Millennium Problems. In BST, it has a precise reformulation as a family of finite spectral questions.

Definition 5.13 — Bounded spectral gap:
For the lattice gauge Hamiltonian H on $\mathscr{H}_K$:

1. Diagonalise H by the spectral theorem.
   Eigenvalues: $E_0 \leq E_1 \leq \ldots \leq E_{d-1}$
   where $d = \dim(\mathscr{H}_K)$.

2. The vacuum energy: E_0 = min eigenvalue.
   The vacuum state: $|\Omega\rangle$ = eigenvector of E_0.

3. The mass gap: $\gamma = E_1 - E_0$.
   This is a specific element of $\mathbb{R}_B(k)$, computable by diagonalisation.

4. Whether $\gamma > 0$ is decidable (comparison in $\mathbb{R}_B(k)$).

Type I. The spectral gap of any specific bounded Hamiltonian is an exactly computable quantity.

5.6.1 The reformulated Millennium Problem

The Clay problem asks: Does four-dimensional Yang-Mills theory with gauge group SU(N) (for any $N \geq 2$) have a positive mass gap?

In BST, this decomposes into:

Fixed-(K,k) question:

For this specific lattice K and this specific precision k, is $\gamma_{K,k} > 0$?

This is decidable. The answer is a specific yes or no for each (K,k). Type I.

Family-level question:
Does there exist a uniform positive lower bound $\gamma_0 > 0$ such that for all sufficiently large lattices K (in some controlled family approaching the continuum limit), $\gamma_{K,k} \geq \gamma_0$?

This is the genuine open problem. It is a statement about the family of bounded gauge theories — a metatheoretic claim about the behaviour of the spectral gap across models. It is not decidable by any single finite computation.

The BST reformulation separates:
- The finite verification (decidable, Type I)
- The universal claim (open, family-level)

The Clay problem is the second. BST does not solve it. What BST does is make the structure of the problem transparent: it is a question about whether a sequence of computable numbers has a positive lower bound.

5.7 Yang-Mills Dynamics

Definition 5.14 — Lattice equations of motion:
The classical lattice Yang-Mills equations are the stationarity conditions for the Wilson action:

$$\delta S/\delta U(e) = 0 \quad \text{for each edge } e \in E(K)$$

Since $U(e) \in SU(N)_B$ is a matrix-valued variable, the variation is taken within the group manifold. The resulting equation at each link e is:

$$\Sigma_{p \ni e} (U_p - U_p^\dagger) |_{\text{projected}} = 0$$

where the sum runs over all plaquettes containing edge e, and the projection is onto the traceless anti-Hermitian part (i.e. onto $su(N)_B$).

This is a finite system of algebraic equations — one equation per edge, each involving a bounded number of N×N matrix operations. Type I.

Theorem 5.1 — Bounded lattice Yang-Mills energy conservation:
The lattice Hamiltonian $H = H_E + H_B$ (Definition 5.12) generates unitary time evolution on $\mathscr{H}_K$:

$$|\Psi(t)\rangle = \exp(-iHt/\hbar)|\Psi(0)\rangle$$

By Theorem 11.11(iv) of the AFB paper, $\exp(-iHt/\hbar)$ is unitary. Therefore:

$$\langle\Psi(t)|H|\Psi(t)\rangle = \langle\Psi(0)|H|\Psi(0)\rangle$$

Energy is conserved exactly. □

Type I.

Theorem 5.2 — Gauss's law constraint:

The physical Hilbert space $\mathcal{H}_{phys} \subseteq \mathcal{H}_K$ is the subspace of gauge-invariant states:

$$\mathcal{H}_{phys} = \{|\Psi\rangle \in \mathcal{H}_K : G_v|\Psi\rangle = |\Psi\rangle \text{ for all } v \in V(K)\}$$

where G_v is the gauge transformation operator at vertex v.

The operators $\{G_v\}$ form a finite group (the gauge group acting vertex-wise). By Maschke's theorem (Theorem 13.3 of the AFB paper), $\mathcal{H}_K$ decomposes into irreducible representations of this gauge group. $\mathcal{H}_{phys}$ is the trivial (gauge-singlet) component.

The projector onto $\mathcal{H}_{phys}$ is:

$$P_{phys} = (1/|G_V|) \Sigma_{g \in G_V} \Pi_v G_v^{g(v)}$$

This is the group-averaging projector of Maschke's theorem applied to the gauge group. It is a finite sum of finite operators. Type I.

Gauss's law (the constraint that physical states are gauge-invariant) is implemented exactly by this projector. In classical gauge theory, Gauss's law is a differential constraint requiring care with gauge-fixing. In BST, it is a finite group-averaging projection — constructive and exact.

Plain language: Gauss's law in BST is Maschke's theorem applied to the gauge group. The physical Hilbert space is the gauge-singlet subspace, found by the explicit averaging projector from Part XIII's representation theory. No gauge-fixing procedure, no Faddeev-Popov ghosts, no BRST cohomology. The finite group structure handles everything directly.

5.8 Comparison: U(1) to SU(N)

Every construction in this Part specialises to Part II §2.5 when $G_B = U(1)_B$:

General SU(N)	U(1) (Part II §2.5)
U(e) ∈ SU(N)_B (N×N matrix)	U(e) ∈ U(1)_B (phase)
U_p = ordered product of U(e)	U_p = product of phases
Tr(U_p) (N×N trace)	U_p (the phase itself)
S = β Σ(1−Re Tr(U_p)/N)	S = β Σ Re(1−U_p)
[A_μ, A_ν] ≠ 0 (non-abelian)	[A_μ, A_ν] = 0 (abelian)
W(γ) = Tr(∏U)/N	W(γ) = ∏U
Gauge transform: conjugation	Gauge transform: trivial (abelian)
Gauss's law: non-trivial	Gauss's law: divergence constraint

The non-abelian structure (ordering of products, non-trivial conjugation, the commutator [A_μ, A_ν]) is the only structural addition. The framework — connections on edges, curvature on plaquettes, action as a finite sum, partition function as a finite sum, Hilbert space as finite — is identical.

Electromagnetism is the N = 1 case. The weak force is N = 2. The strong force is N = 3. The Standard Model gauge theory is the direct product of all three (Definition 5.3). The entire gauge-theoretic structure of the Standard Model fits inside BST as finite matrix algebra on a finite lattice.

5.9 What Part V Recovers

Result	Type	Section
SU(N)_B as finite matrix group	I	5.1
su(N) structure constants	I	5.1
Standard Model gauge group	I	5.1
SM representation content	I	5.1
Lattice connections	I	5.2
Plaquette holonomy	I	5.2
Wilson action	I	5.2
Gauge transformations	I	5.3
Gauge invariance of action	I	5.3
Wilson loops (gauge-invariant)	I	5.3
Partition function (finite sum)	I	5.4
Expectation values	I	5.4
Lattice Monte Carlo	I	5.4
Gauge-field Hilbert space	I	5.5
Kogut-Susskind Hamiltonian	I	5.5
Spectral gap (computable)	I	5.6
Mass gap decidability (fixed K,k)	I	5.6
Yang-Mills equations of motion	I	5.7
Energy conservation	I	5.7
Gauss's law (Maschke projector)	I	5.7
U(1) → SU(N) generalisation	I	5.8

What is NOT recovered:

Continuum Yang-Mills on $\mathbb{R}^4$:
The smooth gauge field A_μ(x) on a continuous manifold does not exist in BST. The lattice gauge field is the finite mathematical content; the continuum field is the Type IV family idealisation.

```
The mass gap as a universal theorem:
  The spectral gap of each specific Hamiltonian is
  computable. Whether the gap persists uniformly as
  the lattice refines is the open Clay problem.

Perturbative QCD:
  Feynman diagrams in the continuum require infinite-
  dimensional functional integrals. The lattice
  partition function replaces them with finite sums.
  Bounded perturbation theory (finite Feynman rules
  on the lattice) is deferred to Part IX of this
  volume.

Asymptotic freedom:
  The running of the gauge coupling with energy scale
  is a family-level statement about how β_{K,k}
  relates to the lattice spacing a. The bounded RG
  framework is deferred to Part IX.

What is gained:

The partition function always exists — it is a
finite sum of positive numbers. No regularisation,
no renormalisation, no dimensional regularisation,
no ε-expansion, no zeta-function regularisation.

Gauss's law is implemented by a finite group
projection (Maschke's theorem) rather than by
gauge-fixing plus ghost fields.

The mass gap is a computable quantity for any
specific Hamiltonian. The Clay problem is
decomposed into a decidable finite part and an
open family-level part — making the structure of
the problem transparent.

The Standard Model gauge group SU(3) × SU(2) × U(1)
and its representation content exist as explicit
finite algebraic objects inside BST, with a modest
cardinality cost (k^{56}).
```

End of Part V

Part VI: Finite Simplicial Geometry and Topology

Maturity: Tier 1 (topology and Hodge decomposition) / Tier 2 (Kähler structures)

AFB imports: BST set theory and Bounded Separation (Part IV); bounded functions and relations (Part VII); finite-dimensional vector spaces, inner products, and spectral theorem (Part XI, §§11.1–11.5, Theorems 11.1–11.9); exterior algebra, wedge product, Hodge star (Part XI, §11.8, Definitions 11.17–11.19); Weyl and Davis-Kahan spectral stability (Part XI, Theorems 11.5–11.6); bounded Dolbeault cohomology and preliminary Kähler geometry (Part X, §§10.3–10.5); ℝ_B(k) and ℂ_B(k^4) (Part VIII)

Volume imports: Part II (exterior calculus formulation of Maxwell, providing physical motivation); Part V (gauge connections on simplicial complexes, gauge-field cochains)

The physics

This Part builds mathematical infrastructure rather than a specific physical theory. Simplicial geometry and topology — the mathematics of finite triangulated spaces — provides the language for discrete gauge theory (Part V), discrete gravity (Part VII), and topological phases of matter (Part VIII). In classical physics, these tools live on smooth manifolds: differential forms, de Rham cohomology, the Hodge theorem. In BST, they live on finite simplicial complexes: cochains, simplicial cohomology, the combinatorial Hodge decomposition. The physics these tools support is experimentally verified — the tools themselves are the finite mathematical framework.

Experimental relevance (indirect)

This Part has no direct experimental tests of its own. Its experimental validation is indirect, through the physics it supports:

Gauge theory (Part V). The cochain complexes and $d^2 = 0$ identity constructed here are the mathematical backbone of lattice gauge theory. Every lattice QCD computation — including the proton mass calculation cited in Part V — uses cochains on a simplicial (or cubical) lattice. The Hodge decomposition separates gauge and physical degrees of freedom.

Topological insulators (Part VIII). The Chern number (§8.4 of Part VIII) — the topological invariant classifying quantum Hall states — is computed by the cohomological machinery of this Part applied to the Brillouin zone. Experimentally confirmed in quantum Hall systems (von Klitzing 1980, Nobel Prize 1985) and topological insulators (Hasan and Kane 2010).

Discrete gravity (Part VII). Regge calculus and causal set theory use simplicial complexes as discrete spacetimes. The deficit angles and Regge action of Part VII are defined on the simplicial structures constructed here.

What BST constructs

Three layers of increasing structure: (1) combinatorial topology (§§6.1–6.3) — simplicial complexes, boundary maps, homology and cohomology as finite linear algebra, Betti numbers by Gaussian elimination, Stokes' theorem as a finite algebraic identity; (2) metric topology (§§6.4–6.5) — inner products on cochains, the combinatorial Laplacian Δ = dd* + d*d, and the Hodge decomposition C^p = im(d) ⊕ im(d*) ⊕ ker(Δ) as an exact theorem of finite linear algebra; (3) Kähler structures (§6.6) — the approximate Kähler identities and (p,q)-decomposition, where the proof strategy is clear but specific geometric lemmas remain open (Tier 2).

6.1 Finite Simplicial Complexes

Definition 6.1 — Bounded abstract simplicial complex:

A bounded abstract simplicial complex K is a finite collection of finite subsets of a finite vertex set V, satisfying the downward closure condition:

If $\sigma \in K$ and $\tau \subseteq \sigma$, then $\tau \in K$.

Formally: $K \subseteq P(V)$ (a subset of the power set of V, which exists by Theorem 4.3 of the AFB paper whenever $2^{|V|} \leq n_M$), and K is downward closed under the subset relation.

A p-simplex is an element $\sigma \in K$ with $|\sigma| = p + 1$ (a set of p + 1 vertices). The dimension of K is the maximum p such that K contains a p-simplex.

The set of p-simplices: $K_p = \{\sigma \in K : |\sigma| = p + 1\}$. Each K_p is finite (a subset of a finite set). ✓

Examples:
- A graph: K_0 = vertices, K_1 = edges, no higher simplices.
- A triangulation of a surface: K_0 = vertices, K_1 = edges, K_2 = triangles.
- A tetrahedral mesh: K_0 through K_3.

Type I. All objects are finite sets definable by Bounded Separation.

Definition 6.2 — Oriented simplex:

An oriented p-simplex is an equivalence class of orderings of its p + 1 vertices, where two orderings are equivalent if they differ by an even permutation.

For a p-simplex $\sigma = \{v_0, v_1, \ldots, v_p\}$, the two orientations are:

$$[v_0, v_1, \ldots, v_p] \quad \text{and} \quad -[v_0, v_1, \ldots, v_p]$$

where the sign is determined by the parity of the permutation relating two orderings: $[v_{\pi(0)}, \ldots, v_{\pi(p)}] = \mathrm{sgn}(\pi) \cdot [v_0, \ldots, v_p]$.

Since $|K_p|$ is finite, the set of oriented p-simplices is finite (at most $2|K_p|$ elements, with each unoriented simplex giving two oriented versions). Type I.

6.2 Chain Complexes and Homology

Definition 6.3 — Bounded chain group:

For a bounded simplicial complex K and a coefficient field F (typically $\mathbb{R}_B(k)$ or $\mathbb{C}_B(k^4)$):

The p-th chain group is the finite free module:

$$C_p(K; F) = \{\Sigma_{\sigma \in K_p} a_\sigma [\sigma] : a_\sigma \in F\}$$

This is a finite-dimensional vector space over F with basis the oriented p-simplices.

$\dim(C_p) = |K_p|$ = number of p-simplices.

An element $c \in C_p$ is a finite linear combination of oriented simplices — a finite vector of $|K_p|$ coefficients. Type I.

Definition 6.4 — Boundary operator:

The boundary operator $\partial_p: C_p(K; F) \to C_{p-1}(K; F)$ is the linear map defined on basis elements by the alternating face formula:

$$\partial_p [v_0, v_1, \ldots, v_p] = \Sigma_{i=0}^{p} (-1)^i [v_0, \ldots, \hat{v}_i, \ldots, v_p]$$

where $\hat{v}_i$ means omit the i-th vertex.

This is a $|K_{p-1}| \times |K_p|$ matrix over F. Each column has at most $p + 1$ nonzero entries (each ± 1). Type I.

Theorem 6.1 — $\partial^2 = 0$:

$\partial_{p-1} \circ \partial_p = 0$ for all p.

Proof: Apply ∂_{p-1} to $\partial_p [v_0, \ldots, v_p]$:

$$\begin{aligned}
&\partial_{p-1}(\partial_p [v_0, \ldots, v_p]) \\
&= \Sigma_i (-1)^i \partial_{p-1} [v_0, \ldots, \hat{v}_i, \ldots, v_p] \\
&= \Sigma_i (-1)^i \Sigma_{j<i} (-1)^j [v_0,\ldots,\hat{v}_j,\ldots,\hat{v}_i,\ldots,v_p] \\
&+ \Sigma_i (-1)^i \Sigma_{j>i} (-1)^{j-1} [v_0,\ldots,\hat{v}_i,\ldots,\hat{v}_j,\ldots,v_p]
\end{aligned}$$

In the first sum, the term with omissions j and i appears with sign $(-1)^{i+j}$. In the second sum, the same pair (now with i and j) appears with sign $(-1)^{i+j-1} = -(-1)^{i+j}$. Every term cancels exactly. □

Type I. The proof is a finite algebraic cancellation — no limits, no infinite sums. The identity $\partial^2 = 0$ is exact for any coefficient field, including the approximate fields $\mathbb{R}_B(k)$ and $\mathbb{C}_B(k^4)$, because it involves only addition and sign changes (± 1), not multiplication that would introduce rounding error.

Definition 6.5 — Homology groups:
The p-th homology group of K with coefficients in F:

$Z_p = \ker(\partial_p)$ (p-cycles: chains with zero boundary)
$B_p = \mathrm{im}(\partial_{p+1})$ (p-boundaries: chains that are boundaries of (p+1)-chains)

Since $\partial^2 = 0$: $B_p \subseteq Z_p$. ✓

$$H_p(K; F) = Z_p / B_p$$

This is a finite-dimensional vector space over F.

$$\dim(H_p) = \dim(Z_p) - \dim(B_p) = \dim(\ker \partial_p) - \dim(\mathrm{im}\ \partial_{p+1})$$

This is computable by Gaussian elimination or Smith normal form on the finite matrix ∂_p. The rank-nullity theorem (exact for finite matrices) gives all dimensions.

The p-th Betti number:
$\beta_p = \dim(H_p(K; F))$

This counts the number of independent p-dimensional "holes" in K: β_0 = connected components, β_1 = independent loops, β_2 = enclosed cavities, etc.

Type I. All dimensions are computable by finite linear algebra (Part XI of the AFB paper).

Theorem 6.2 — Euler characteristic:
The Euler characteristic of K is:

$\chi(K) = \Sigma_p (-1)^p |K_p|$ (alternating simplex count)
$= \Sigma_p (-1)^p \beta_p$ (alternating Betti number sum)

The equality of these two expressions is a finite identity:

$$\Sigma_p (-1)^p \dim(C_p) = \Sigma_p (-1)^p \dim(H_p)$$

Proof: By the rank-nullity theorem applied to each ∂_p. For each p:

$$\dim(C_p) = \dim(\ker \partial_p) + \dim(\mathrm{im}\ \partial_p) = \dim(Z_p) + \dim(B_{p-1})$$

Substituting into the alternating sum and using $\dim(H_p) = \dim(Z_p) - \dim(B_p)$, the B_p terms telescope and cancel. □

Type I. An exact algebraic identity.

6.3 Cochain Complexes and Cohomology

Definition 6.6 — Bounded cochain group:
The p-th cochain group is the dual of C_p:

$$C^p(K; F) = \mathrm{Hom}(C_p, F) \cong F^{|K_p|}$$

A p-cochain is a function assigning a coefficient in F to each oriented p-simplex. Since K_p is finite, this is a finite vector. Type I.

$$\dim(C^p) = |K_p| = \dim(C_p).$$

Definition 6.7 — Coboundary operator:
The coboundary operator $d^p: C^p \to C^{p+1}$ is the transpose (dual) of the boundary operator:

$$(d^p \alpha)(\sigma) = \alpha(\partial_{p+1} \sigma) \quad \text{for } \sigma \in K_{p+1}$$

As a matrix: $d^p = (\partial_{p+1})^T$.

Since $\partial^2 = 0$ (Theorem 6.1), we have:
$(d^p)^2 = (\partial^T)^2 = (\partial^2)^T = 0^T = 0.$

$d^2 = 0$. ✓ (Exact, by transposition of $\partial^2 = 0$.)

Definition 6.8 — Cohomology groups:
$Z^p = \ker(d^p)$ (p-cocycles)
$B^p = \mathrm{im}(d^{p-1})$ (p-coboundaries)

$$H^p(K; F) = Z^p / B^p$$

$\dim(H^p) = \beta_p = \dim(H_p)$ (same Betti numbers)

The isomorphism $H^p \cong H_p$ holds for coefficients in a field by the universal coefficient theorem — which for finite-dimensional vector spaces is simply the statement that a finite matrix and its transpose have the same rank. Type I.

Connection to the exterior derivative:
The coboundary $d: C^p \to C^{p+1}$ is the discrete counterpart of the exterior derivative $d: \Omega^p(M) \to \Omega^{p+1}(M)$ on smooth manifolds (identified in §11.8 of the AFB paper).

The identity $d^2 = 0$ is the combinatorial counterpart of $d^2 = 0$ on differential forms.

For a p-cochain α representing a discrete "p-form,"

```
dα is the discrete "(p+1)-form" — its value on a
(p+1)-simplex is the alternating sum of α on the
boundary faces. This is exactly Stokes' theorem
in combinatorial form:

(dα)(σ) = α(∂σ)  ↔  ∫_σ dα = ∫_{∂σ} α

The finite version is not an approximation to
Stokes' theorem — it IS Stokes' theorem on a
finite complex. Type I.
```

Plain language: Homology counts holes. Cohomology assigns values to holes. Both are computed by finite linear algebra — Gaussian elimination on finite matrices of integers (or elements of F). The Betti numbers β_0, β_1, β_2, … are computable finite invariants. The coboundary operator d is the discrete exterior derivative, and $d^2 = 0$ is exact by finite cancellation. Stokes' theorem is a finite algebraic identity.

6.4 Inner Products and the Combinatorial Laplacian

The topology of §§6.1–6.3 is purely combinatorial — it depends only on the simplicial structure, not on any notion of distance or angle. To connect topology to geometry (and to physics), we add an inner product on the cochain spaces. This introduces the adjoint d*, the Laplacian, and the Hodge decomposition.

Definition 6.9 — Bounded cochain inner product:

```
Choose a positive-definite inner product on C^p(K; F)
for each p. The simplest choice is the standard one:

⟨α, β⟩_p = Σ_{σ ∈ K_p} w_σ · α(σ)* · β(σ)

where w_σ > 0 are positive weights (typically
w_σ = 1 for the combinatorial inner product, or
w_σ = Vol(σ) for a geometrically weighted version).

This makes each C^p a finite inner-product space
(Definition 11.7 of the AFB paper). Type I.

For applications to physics, the weights encode
the metric geometry of the simplicial complex: if
the simplices have assigned volumes (as in Regge
geometry, Part VII of this volume), the weights
reflect those volumes.
```

Definition 6.10 — Formal adjoint:

```
The formal adjoint of d^p: C^p → C^{p+1} with
respect to the inner products is the linear map:

d^{p*}: C^{p+1} → C^p

defined by:  ⟨d^p α, β⟩_{p+1} = ⟨α, d^{p*} β⟩_p

for all α ∈ C^p, β ∈ C^{p+1}.
```

Since d^p is a finite matrix and the inner products are finite matrices (positive definite), d^{p*} is computable by:

$$d^{p*} = W_p^{-1} (d^p)^\dagger W_{p+1}$$

where $W_p = \mathrm{diag}(w_\sigma)$ and $(d^p)^\dagger$ is the conjugate transpose. This is a finite matrix computation. Type I.

For the unweighted case ($W = I$): $d^{p*} = (d^p)^\dagger$.

Definition 6.11 — Bounded combinatorial Laplacian:

The p-th combinatorial Laplacian (Hodge Laplacian) is the linear operator:

$$\Delta_p = d^{p-1} d^{(p-1)*} + d^{p*} d^p \quad : C^p \to C^p$$

This is a $|K_p| \times |K_p|$ matrix over F.

Properties (proved by direct computation):

(i) Δ_p is self-adjoint: $\langle \Delta_p \alpha, \beta \rangle = \langle \alpha, \Delta_p \beta \rangle$.
Proof: Both terms $d\, d^*$ and $d^*\, d$ are self-adjoint ($d\, d^*$ is self-adjoint because $(d\, d^*)^\dagger = d^{**} d^\dagger = d\, d^*$, and similarly for $d^*\, d$). ✓

(ii) Δ_p is positive-semidefinite:
$$\langle \Delta_p \alpha, \alpha \rangle = \langle d^* \alpha, d^* \alpha \rangle + \langle d\alpha, d\alpha \rangle = \|d^* \alpha\|^2 + \|d\alpha\|^2 \geq 0. \ ✓$$

(iii) $\ker(\Delta_p) = \ker(d^p) \cap \ker(d^{(p-1)*})$ (the harmonic p-cochains).
Proof: $\Delta_p \alpha = 0$ iff $\|d^* \alpha\|^2 + \|d\alpha\|^2 = 0$ iff $d\alpha = 0$ and $d^* \alpha = 0$. ✓

All eigenvalues of Δ_p are non-negative real numbers in $\mathbb{R}_B(k)$, computable by the spectral theorem (Theorem 11.4 of the AFB paper). Type I.

6.5 The Hodge Decomposition

Theorem 6.3 — Bounded Hodge decomposition:

For a finite simplicial complex K with inner-product-equipped cochain spaces:

$$C^p(K; F) = \mathrm{im}(d^{p-1}) \oplus \mathrm{im}(d^{p*}) \oplus \ker(\Delta_p)$$

That is: every p-cochain decomposes uniquely into an exact part (in $\mathrm{im}(d)$), a coexact part (in $\mathrm{im}(d^*)$), and a harmonic part (in $\ker(\Delta)$).

Proof: This is the orthogonal decomposition of a

finite inner-product space into the image, co-image,
and kernel of a self-adjoint operator.

Step 1: C^p = ker(Δ_p) ⊕ im(Δ_p)
  (orthogonal decomposition by self-adjointness of Δ_p
  and the finite-dimensional spectral theorem — the
  eigenspaces of Δ_p span C^p, and ker(Δ_p) is the
  0-eigenspace).

Step 2: im(Δ_p) = im(d^{p-1}) ⊕ im(d^{p*})
  Proof: im(Δ_p) = im(dd* + d*d) ⊆ im(d) + im(d*).
  Conversely: for α ∈ im(d), write α = dβ. Then
  d*α = d*dβ ∈ im(d*d) ⊆ im(Δ_p), so the image
  of d restricted to the appropriate subspace lies
  in im(Δ_p). Similarly for im(d*).
  The two subspaces are orthogonal:
  ⟨dα, d*β⟩ = ⟨d^2α, β⟩ = ⟨0, β⟩ = 0. ✓

Step 3: Combine Steps 1 and 2:
  C^p = im(d) ⊕ im(d*) ⊕ ker(Δ). □

Type I. The proof uses only:
- Self-adjointness of Δ (Definition 6.11(i))
- The spectral theorem (Theorem 11.4)
- d^2 = 0 (Theorem 6.1)
- Orthogonality of inner products (Definition 6.9)
All exact finite linear algebra.

Theorem 6.4 — Hodge isomorphism:
The harmonic cochains are isomorphic to cohomology:

ker(Δ_p) ≅ H^p(K; F)

Proof: Every harmonic cochain is a cocycle
(since dα = 0 by Definition 6.11(iii)), so there
is a natural map ker(Δ_p) → H^p = Z^p/B^p.

Injective: If a harmonic α is a coboundary (α = dβ),
then α ∈ im(d) ∩ ker(Δ). But the Hodge decomposition
says im(d) ⊥ ker(Δ). So α = 0. ✓

Surjective: Every cohomology class [c] ∈ H^p has
a unique harmonic representative. Write c = c_d +
c_{d*} + c_h by Hodge decomposition. Since c ∈ Z^p
(dc = 0), and dc_h = 0, dc_{d*} = 0 (since
d d* d = Δd − d*d^2 = Δd, and ... this requires care).
More directly: the projection of any cocycle onto
ker(Δ) gives a harmonic representative, and two
cocycles in the same class project to the same
harmonic form (their difference is exact, hence
orthogonal to ker(Δ)). ✓ □

Consequence: β_p = dim(H^p) = dim(ker Δ_p)

```
= number of zero eigenvalues of Δ_p.

The Betti numbers are readable from the spectrum
of the Laplacian. The Hodge decomposition converts
a topological question (what are the holes?) into
a spectral question (what is the kernel of a
matrix?). Type I.
```

Plain language: The Hodge decomposition says that every cochain splits uniquely into three orthogonal pieces: an exact piece (a coboundary), a coexact piece (in the image of d*), and a harmonic piece (in the kernel of the Laplacian). The harmonic pieces ARE the cohomology — each cohomology class has exactly one harmonic representative. This is the bridge from topology to spectral geometry: counting holes in a space is the same as counting zero eigenvalues of the Laplacian. In BST, all three parts of this statement are exact finite theorems — orthogonal decomposition of a finite vector space, the spectral theorem for a finite self-adjoint matrix, and the algebraic identity $d^2 = 0$.

This decomposition is used by: - Part V (gauge theory): the decomposition of gauge-field configurations into physical and gauge degrees of freedom - Part VII (gravity): the spectral geometry of discrete spacetimes - The Hodge conjecture programme: the identification of cohomology classes with harmonic forms

6.6 Bounded Kähler Structures (Tier 2)

The Hodge decomposition of §6.5 is exact and requires no additional geometric input beyond the inner product. The stronger Kähler-Hodge theory — which relates the Dolbeault cohomology (Part X of the AFB paper) to the de Rham cohomology, giving a decomposition into (p,q)-types — requires additional structure: a compatible complex structure and a Kähler metric satisfying specific commutator identities.

6.6.1 What is available

```
From Part X of the AFB paper:
- Bounded Dolbeault cohomology H^{p,q}_B
  (Definition 10.7)
- Bounded ∂̄ operator on cochains
- Preliminary bounded Kähler metric

From Part XI of the AFB paper:
- Spectral stability (Weyl, Theorem 11.5;
  Davis-Kahan, Theorem 11.6)
- The operator perturbation framework: if the
  "ideal" Kähler Laplacian Δ_K differs from the
  bounded Laplacian Δ_B by ‖E‖ = O(1/k), then
  eigenvalues shift by at most O(1/k) and
  eigenspaces are stable when the spectral gap
  exceeds 2‖E‖.
```

6.6.2 What is needed (open lemmas)

```
The Kähler identities in the continuum state:

[Λ, ∂] = −i∂̄*     [Λ, ∂̄] = i∂*

where Λ is the adjoint of the Lefschetz operator
```

(contraction with the Kähler form).

In the bounded setting, these identities are expected to hold approximately:

$$[\Lambda\_B, \partial\_B] = -i\bar{\partial}\_B^* + E_1 \quad \text{with } \|E_1\| \leq C/k$$
$$[\Lambda\_B, \bar{\partial}\_B] = i\partial\_B^* + E_2 \quad \text{with } \|E_2\| \leq C/k$$

The error terms E_1, E_2 arise from the approximate arithmetic of $\mathbb{C}\_B(k^4)$ and the discretisation of the Kähler structure on a finite simplicial complex.

OPEN LEMMA 1: Prove the approximate Kähler identities with explicit error bounds for a suitable class of bounded complex simplicial manifolds.

OPEN LEMMA 2: Prove that the resulting approximate Hodge decomposition into (p,q)-types is stable under the spectral perturbation — i.e. that the spectral gap of the Kähler Laplacian exceeds 2C/k for the class of complexes under consideration.

6.6.3 The conditional theorem

CONDITIONAL: If Open Lemmas 1 and 2 are proved, then:

The bounded Kähler-Hodge decomposition holds:

$$H^n(K; \mathbb{C}\_B(k^4)) \cong \oplus\_{p+q=n} H^{p,q}\_B(K)$$

with each $H^{p,q}\_B$ a finite-dimensional vector space over $\mathbb{C}\_B(k^4)$, and the decomposition stable to within $O(1/k)$ in the spectral-perturbation sense.

This would be a Type III result: approximate decomposition with explicit error vanishing as k grows.

Status: Tier 2. The proof strategy is clear, the spectral tools are available, the geometric input lemmas are open.

Plain language: The combinatorial Hodge decomposition (§6.5) is done — it's an exact theorem of finite linear algebra. The Kähler-Hodge decomposition (§6.6) is a stronger result that requires additional geometric structure. The proof strategy is clear: use the approximate Kähler identities plus spectral stability to control the decomposition. The spectral stability tools (Weyl, Davis-Kahan) are already in the AFB paper. What remains open is proving the approximate Kähler identities themselves — a specific geometric lemma about bounded complex simplicial structures.

6.7 Applications and Downstream Connections

6.7.1 Gauge theory (Part V)

The cochain complex $C^p(K; F)$ provides the space in which gauge field configurations live:

- A connection is a 1-cochain (values on edges)
- The field strength is a 2-cochain (dA, values on plaquettes)
- Gauge transformations act on 0-cochains (values on vertices)

The Hodge decomposition separates gauge field configurations into:

- $\mathrm{im}(d)$: pure gauge configurations (gauge artifacts)
- $\mathrm{im}(d^*)$: physical (transverse) configurations
- $\ker(\Delta)$: harmonic / topological modes

This decomposition is Part V's Gauss's law (Theorem 5.2) reformulated in cohomological language.

6.7.2 Gravity (Part VII of this volume)

On a Regge simplicial geometry with edge lengths in $\mathbb{R}_B(k)$:

- The discrete metric is encoded in the simplex volumes and angles (computed by Cayley-Menger determinants)
- The inner product on cochains (Definition 6.9) is weighted by these volumes
- The Laplacian (Definition 6.11) becomes the geometric Laplacian of the discrete manifold
- The spectrum of the Laplacian encodes the geometry: eigenvalues correspond to curvature scales

The Hodge decomposition on a Regge geometry separates the topological content (Betti numbers) from the geometric content (non-zero eigenvalues of Δ).

6.7.3 The bounded Hodge conjecture

Part X of the AFB paper (§10.5) formulates the bounded Hodge conjecture: for a bounded complex projective variety X over $\mathbb{C}_B(k^4)$, every (p,p)-class in $H^{2p}_B(X)$ is a rational linear combination of algebraic cycle classes.

The cohomological machinery of this Part provides the setting: H^{2p}_B is a finite-dimensional vector space, the (p,p)-classes form a subspace computable by projection, and the algebraic cycle classes form a subspace computable from the variety's defining equations. The conjecture asks whether the second subspace spans the first — a decidable finite linear algebra question for each specific X and k.

6.8 What Part VI Recovers

Result	Type	Section
Simplicial complexes	I	6.1
Oriented simplices	I	6.1
Chain groups	I	6.2
Boundary operator ∂	I	6.2
$\partial^2 = 0$	I	6.2
Homology groups H_p	I	6.2
Betti numbers	I	6.2
Euler characteristic	I	6.2
Cochain groups and coboundary d	I	6.3
$d^2 = 0$	I	6.3
Cohomology groups H^p	I	6.3
Stokes' theorem (combinatorial)	I	6.3
Cochain inner products	I	6.4
Adjoint d*	I	6.4
Combinatorial Laplacian Δ_p	I	6.4
Hodge decomposition (combinatorial)	I	6.5
Hodge isomorphism (ker Δ ≅ H^p)	I	6.5
Betti numbers from spectrum of Δ	I	6.5
Kähler-Hodge decomposition	III	6.6 (conditional)
(p,q)-decomposition	III	6.6 (conditional)

What is NOT recovered: smooth manifolds and smooth differential forms (replaced by finite simplicial complexes and cochains). Infinite-dimensional de Rham cohomology (replaced by finite-dimensional simplicial cohomology). Continuous Hodge theory requiring elliptic operator theory on compact manifolds (replaced by finite Hodge decomposition using the spectral theorem for finite matrices). Topological invariance under arbitrary continuous deformations (simplicial invariance under explicit combinatorial moves is available; continuous invariance is a family-level statement).

What is gained: every topological invariant is computable by finite linear algebra. The Hodge decomposition is exact — not an asymptotic or approximate statement. The Betti numbers are computable by diagonalising a finite matrix. The spectral gap of the Laplacian is a computable number that separates topology (zero eigenvalues) from geometry (positive eigenvalues). For applications to gauge theory and gravity, the decomposition provides an exact, constructive separation of physical from gauge degrees of freedom.

End of Part VI

Part VII: Gravity

Maturity: Tier 1–2 (Regge calculus) / Tier 1 (causal combinatorics) / Tier 3 (GR emergence)

AFB imports: ℝ_B(k) (Part VIII, Theorem 8.13); bounded transcendental functions including cos, arccos, and square roots (Part IX, §9.3); bounded ODEs (Part IX, §9.3.3); tensor algebra including (0,2)-tensors, contraction, index raising/lowering (Part XI,

§11.6, Definitions 11.10–11.15); spectral theorem (Part XI, Theorem 11.4); bounded Laplacian (Part IX, Definition 9.6)

Volume imports: Part II §2.2 (Lorentz group, Minkowski metric, causal structure); Part VI (simplicial complexes, Hodge decomposition, spectral geometry of the Laplacian); Part V (gauge fields for matter coupling)

The physics

Gravity is the oldest and most universal of the physical forces — it governs the orbits of planets, the structure of stars, the expansion of the universe, and the formation of black holes. Newton's theory (Part I) describes gravity as a force between masses. Einstein's general relativity (GR, 1915) replaces this with a geometric theory: mass and energy curve spacetime, and objects follow the straightest possible paths (geodesics) through that curved geometry. GR is a classical field theory whose fundamental object is the metric tensor $g_{\mu\nu}$ — a (0,2)-tensor (Definition 11.11 of the AFB paper) that encodes distances and angles at every point in spacetime.

What experiments confirm

Classical tests of GR. The perihelion precession of Mercury (43 arcseconds/century beyond Newton's prediction), the deflection of starlight by the Sun (1.75 arcseconds, confirmed during the 1919 eclipse by Eddington), and the gravitational redshift (confirmed by Pound-Rebka 1959 to 1%) were the three original tests of GR. All three are confirmed to high precision.

Gravitational waves. LIGO's first detection (2015, Nobel Prize 2017) of gravitational waves from merging black holes confirmed GR's prediction of propagating spacetime disturbances. The observed waveform matches numerical relativity simulations — which solve the discrete Einstein equations on a finite computational grid, the same structure as the Regge equations of §7.4 — to within detector precision.

Black holes and strong-field gravity. The Event Horizon Telescope's image of the black hole in M87 (2019) confirmed the predicted shadow size to ~10%. Pulsar timing in binary systems (Hulse-Taylor, Nobel Prize 1993) confirmed gravitational wave energy loss to 0.2% over 30 years of observation.

Cosmology. The cosmic microwave background (CMB) power spectrum, measured by WMAP and Planck to sub-percent precision, confirms the Friedmann equations (the cosmological solutions of Einstein's field equations) with six free parameters fit to thousands of data points.

Discrete approaches. Regge calculus (§§7.1–7.4) — the finite simplicial approach to gravity — has been used in numerical relativity since the 1960s. Causal set theory (§§7.5–7.7) is an active research programme with the explicit hypothesis that spacetime is fundamentally discrete. Loop quantum gravity and causal dynamical triangulations compute on finite simplicial structures. The experimental status of these discrete approaches is indirect: they aim to reproduce the confirmed predictions of continuum GR from finite combinatorial foundations.

What BST constructs and why it suffices

Every numerical computation in gravitational physics — from LIGO waveform templates to cosmological simulations — is already performed on a finite grid with finite arithmetic. The continuum metric $g_{\mu\nu}(x)$ on a smooth manifold is the theoretical object; the computation uses edge lengths on a mesh (Regge calculus), or field values on a grid (finite differences), or causal relations in a discrete structure (causal sets). BST makes this finiteness foundational.

This Part builds what can be built and identifies what remains open. Regge geometry (§§7.1–7.3) — edge lengths, volumes, deficit angles, the Regge action — is Tier 1: exact finite computation. Regge dynamics and discrete Einstein equations (§7.4) are Tier 1–2. Causal set combinatorics (§§7.5–7.7) — partial orders, proper-time surrogates, the Benincasa-Dowker action — are Tier 1. The emergence of Einstein's GR from these discrete structures (§7.8) is Tier 3: five explicit open problems forming a linear dependency chain. This Part does not claim to have derived general relativity. It builds the finite structures from which the derivation would start.

7.1 Regge Geometry: Edge Lengths and Volumes

Regge calculus (Regge 1961) replaces smooth spacetime with a simplicial complex where the geometry is encoded in the edge lengths. In BST, every edge length is an element of $\mathbb{R}_B(k)$, every volume is computed by a finite determinant, and every angle is computed by a finite trigonometric expression.

Definition 7.1 — Bounded Regge manifold:
A bounded Regge manifold is a pair (K, ℓ) where:

(i) K is a finite simplicial complex of dimension d (Definition 6.1 of Part VI of this volume).

(ii) $\ell: E(K) \to \mathbb{R}_B(k)_{>0}$ is an assignment of positive edge lengths to each edge, satisfying the metric inequalities: for every d-simplex $\sigma \in K_d$, the edge lengths of σ must satisfy the generalised triangle inequalities that ensure σ can be realised as a non-degenerate simplex in $\mathbb{R}^d$.

The metric inequalities are finite algebraic inequalities in $\mathbb{R}_B(k)$ – decidable. ✓

The geometry of (K, ℓ) is fully determined by the combinatorics of K and the real numbers $\ell(e)$. No coordinate charts, no smooth atlas, no differentiable structure. Type I.

Definition 7.2 — Simplex volumes by Cayley-Menger determinant:
For a d-simplex σ with vertices $\{v_0, ..., v_d\}$ and edge lengths $\ell_{ij} = \ell(v_i v_j)$, the d-volume is:

$$(V_\sigma)^2 = ((-1)^{d+1} / (2^d (d!)^2)) \cdot \det(CM)$$

where CM is the (d+2) × (d+2) Cayley-Menger matrix:

$$CM = \begin{pmatrix} 0 & 1 & 1 & \dots & 1 \\ 1 & 0 & \ell_{01}^2 & \dots & \ell_{0d}^2 \\ 1 & \ell_{01}^2 & 0 & \dots & \ell_{1d}^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \ell_{0d}^2 & \ell_{1d}^2 & \dots & 0 \end{pmatrix}$$

The determinant is a polynomial in the squared edge lengths — a finite computation in $\mathbb{R}\_B(k)$ via $\mathbb{Q}\_B(k^4)$ arithmetic.

For d = 2 (triangle with sides a, b, c):
$V^2 = (1/16)[2a^2b^2 + 2b^2c^2 + 2c^2a^2 - a^4 - b^4 - c^4]$
(Heron's formula in squared form.)

For d = 3 (tetrahedron): a 5×5 determinant.
For d = 4 (4-simplex): a 6×6 determinant.

Each is an exact polynomial computation in $\mathbb{R}\_B(k)$, with accumulated rounding error $O(d^2/k^2)$ from the finite-field arithmetic. Type I/III.

Definition 7.3 — Dihedral angles:

For a d-simplex σ and two of its (d-1)-faces f_1, f_2 sharing a (d-2)-face (the "hinge" h), the dihedral angle $\theta(f_1, f_2)$ between f_1 and f_2 is:

$$\cos\theta = (V\_{f_1}^2 + V\_{f_2}^2 - V\_{opp}^2 + \text{correction}) / (2\, V\_{f_1} V\_{f_2} \cdot \text{geometric\_factor})$$

The exact formula depends on d and involves the volumes of the faces and opposite faces, all computable by Cayley-Menger (Definition 7.2).

For d = 2 (angle at a vertex of a triangle):
The standard law of cosines:
$\cos\theta = (a^2 + b^2 - c^2) / (2ab)$

For d = 3 (dihedral angle of a tetrahedron):
$\cos\theta = (\cos\alpha - \cos\beta\cos\gamma) / (\sin\beta\sin\gamma)$
where α, β, γ are the face angles at the relevant vertex.

For d = 4: expressible in terms of face volumes and the Cayley-Menger data.

In each case: θ = arccos(...) where the argument is a finite algebraic expression in $\mathbb{R}\_B(k)$, and arccos is evaluated by bounded transcendental approximation (Part IX of the AFB paper) to precision $O(1/k)$.

Type I for the algebraic expression; Type III for the arccos evaluation.

7.2 Curvature: Deficit Angles

In smooth Riemannian geometry, curvature measures how much parallel transport around a small loop rotates a vector. In Regge geometry, curvature is concentrated at the (d-2)-dimensional hinges (edges in 3D, triangles in 4D), and is measured by the deficit angle: the amount by which the dihedral angles around a hinge fail to sum to 2π.

Definition 7.4 — Deficit angle:

For a (d-2)-dimensional hinge h in a Regge manifold (K, ℓ), let $\sigma_1, \ldots, \sigma_m$ be the d-simplices sharing h, and let θ_i be the dihedral angle of σ_i at h (Definition 7.3).

The deficit angle at h is:

$$\delta_h = 2\pi - \sum_{i=1}^{m} \theta_i$$

where $2\pi \in \mathbb{R}_B(k)$ is evaluated by bounded transcendental approximation.

Interpretation:
- $\delta_h = 0$: flat around h (no curvature).
- $\delta_h > 0$: positive curvature (like a cone tip — there is "too little" angle, so the surface closes in).
- $\delta_h < 0$: negative curvature (like a saddle — there is "too much" angle).

δ_h is a specific element of $\mathbb{R}_B(k)$, computable from the edge lengths. It is the discrete analogue of the scalar curvature integrated over the neighbourhood of h.

Type I for the algebraic computation; Type III for the precision of the trigonometric functions.

The key geometric fact (Regge 1961):
In the limit where the simplicial mesh becomes fine (many small simplices approximating a smooth manifold), the deficit angle δ_h converges to the integral of the scalar curvature R over the dual volume of h:

$$\delta_h \cdot V_h \to \int_{\mathrm{dual}(h)} R \cdot dV$$

This convergence is a family-level (Type IV) statement — it relates the bounded Regge geometry to the smooth Riemannian geometry that BST does not posit as an internal object. At each fixed (K, ℓ), the deficit angle is an exact finite quantity. Type I.

7.3 The Regge Action

Definition 7.5 — Bounded Regge action:
The Regge action for a d-dimensional Regge manifold (K, ℓ) is:

$$S_R[\ell] = (1/8\pi G) \Sigma_{h \in K_{d-2}} V_h \cdot \delta_h$$

where:
- The sum runs over all (d-2)-dimensional hinges h
- V_h is the (d-2)-volume of h (Cayley-Menger, Def. 7.2)
- δ_h is the deficit angle at h (Def. 7.4)
- G is Newton's gravitational constant $\in \mathbb{R}_B(k)$

S_R is a finite sum of $|K_{d-2}|$ terms, each a product of a volume (a polynomial in edge lengths) and a deficit angle (involving arccos of an algebraic expression). Total cost: $O(|K_{d-2}| \cdot d^2)$ operations in $\mathbb{R}_B(k)$.

Type I for the finite sum. The Regge action is the exact discrete counterpart of the Einstein-Hilbert action:

$$S_{EH} = (1/16\pi G) \int R \sqrt{g}\, d^4x$$

The correspondence $S_R \to S_{EH}$ as the mesh refines is the content of Regge's original theorem and is a family-level (Type IV) statement.

For d = 4 (the physically relevant case):
- Hinges are triangles (2-simplices): $|K_2|$ hinges
- V_h = area of triangle h (by Cayley-Menger)
- δ_h = deficit angle at triangle h
- $S_R = (1/8\pi G) \Sigma_{triangles} A_h \cdot \delta_h$

For a 4D simplicial manifold with V vertices, E edges, T triangles, Tet tetrahedra, and P pentachorons (4-simplices):
- The action involves T terms (one per triangle).
- Each term requires computing the dihedral angles of the pentachorons sharing that triangle — a bounded number (depending on the combinatorial structure, typically 3–20 pentachorons per triangle).
- Total cost: $O(T \cdot p_{max} \cdot d^2)$ operations, where p_{max} is the maximum number of pentachorons per triangle.

All finite. All computable. Type I.

7.4 Regge Dynamics: Discrete Einstein Equations

Definition 7.6 — Bounded Regge equations:
The Regge equations of motion are obtained by varying the Regge action with respect to the edge lengths:

$$\partial S_R / \partial \ell_e = 0 \quad \text{for each edge } e \in E(K)$$

Using the Schläfli identity (which states that $\Sigma_h V_h \partial\delta_h/\partial\ell_e = 0$ for any variation of a simplicial manifold), the variation simplifies to:

$$\partial S_R / \partial \ell_e = (1/8\pi G) \Sigma_{h \supset e} \delta_h \cdot \partial V_h/\partial \ell_e = 0$$

where the sum runs over all hinges h containing edge e, and $\partial V_h/\partial\ell_e$ is the derivative of the hinge volume with respect to the edge length.

Each term involves:
- The deficit angle δ_h (already computed)
- The derivative of the Cayley-Menger volume formula with respect to ℓ_e — an algebraic expression in the edge lengths, computable by finite differences (Definition 9.2 of the AFB paper) or by explicit polynomial differentiation

The Regge equations are a finite system of $|E(K)|$ algebraic equations in $|E(K)|$ unknowns (the edge lengths). Type I for the equations themselves.

Solving the equations — finding edge lengths that satisfy all of them simultaneously — is a finite root-finding problem. Newton-Raphson on $\mathbb{R}_B(k)^{|E|}$ (a system of ODEs, Theorem 9.8 of the AFB paper) provides a computational method. Type I/III.

Theorem 7.1 — Schläfli identity in BST:
For a d-dimensional Regge manifold (K, ℓ) and any variation $\ell \to \ell + \delta\ell$ of the edge lengths:

$$\Sigma_{h \in K_{d-2}} V_h \cdot \delta(\delta_h) = 0$$

where $\delta(\delta_h)$ is the variation of the deficit angle at hinge h induced by the variation $\delta\ell$.

Proof: The Schläfli identity is an algebraic identity relating the variations of dihedral angles to the variations of edge lengths in a single d-simplex. Summing over all simplices sharing each hinge, the dihedral angle variations contribute with appropriate signs, and the identity follows from the fact that the total solid angle around each hinge in flat space

```
is 2π.

The proof is a finite algebraic computation for each
d-simplex — it requires only the chain rule applied
to the Cayley-Menger volume and dihedral angle
formulas. By BI-BST over the simplices of K. □

Type I. The Schläfli identity is exact for any
finite simplicial manifold.

Physical interpretation:
The Regge equations ∂S_R/∂ℓ_e = 0 are the discrete
analogues of Einstein's field equations in vacuum:

G_μν = 0   (Einstein tensor vanishes)

The Schläfli identity ensures that the Regge equations
are consistent — they form a well-posed system.
This is the discrete analogue of the Bianchi identity
∇_μ G^μν = 0, which ensures the consistency of
Einstein's equations.

With matter:
S_total = S_R + S_matter

where S_matter is the action of matter fields on the
simplicial complex (e.g. the lattice gauge theory
action of Part V). The coupled equations:

∂S_R/∂ℓ_e = −∂S_matter/∂ℓ_e

are the discrete Einstein equations with source —
the discrete analogue of G_μν = 8πG T_μν.

Type I for the discrete equations. The identification
with Einstein's equations is Type IV (family-level).
```

Plain language: Regge calculus gives us a finite, exact theory of discrete gravity. The geometry is encoded in edge lengths. Curvature is measured by deficit angles. The gravitational action is a finite sum. The equations of motion are a finite system of algebraic equations. The Schläfli identity — the discrete Bianchi identity — ensures consistency. Matter coupling is immediate: add the lattice gauge action from Part V and vary with respect to edge lengths. The result is a finite, coupled system of discrete gravity + gauge theory equations. All computable.

What this does NOT give us is a proof that this discrete theory converges to Einstein's general relativity as the mesh refines. That is the content of §7.8.

7.5 Causal Set Structures

The Regge approach (§§7.1–7.4) starts from a simplicial complex with metric data (edge lengths). The causal set approach starts from something more primitive: a finite partial order, interpreted as the causal structure of discrete spacetime.

Definition 7.7 — Bounded causal set:

A bounded causal set is a finite partially ordered set $(C, \prec)$ where:

(i) C is a finite set of "events," $|C| \leq n_M$.
(ii) $\prec$ is a strict partial order on C:
- Irreflexive: $x \nprec x$ for all x.
- Transitive: $x \prec y$ and $y \prec z$ implies $x \prec z$.

(iii) Local finiteness: for all $x, z \in C$ with $x \prec z$, the interval $I(x,z) = \{y \in C : x \prec y \prec z\}$ is finite.

Condition (iii) is automatic in BST (C is finite). The relation $\prec$ is a finite subset of $C \times C$, definable by Bounded Separation. Transitivity and irreflexivity are decidable by exhaustive finite check. Type I.

Physical interpretation: $x \prec y$ means "event x is in the causal past of event y" — information can propagate from x to y but not vice versa.

Definition 7.8 — Causal geometric invariants:

From the partial order alone, the following geometric surrogates are computable:

Proper time surrogate:
The longest chain length from x to z:
$\tau(x,z) = \max\{n : \exists \text{ chain } x = c_0 \prec c_1 \prec \ldots \prec c_n = z\}$

This is computable by a finite graph algorithm (longest path in a DAG — $O(|C| + |\prec|)$ by topological sort). The Myrheim-Meyer correspondence (classical result) states that in the continuum limit, the longest chain length between two events is proportional to the proper time between them.

Volume surrogate:
The interval cardinality:
$\mathrm{vol}(x,z) = |I(x,z)| = |\{y \in C : x \prec y \prec z\}|$

This is a finite count — computable by exhaustive enumeration. In the continuum limit, the interval cardinality is proportional to the spacetime volume of the causal diamond between x and z.

Dimension estimator:
The Myrheim-Meyer dimension estimator uses the relationship between chain lengths and interval cardinalities to estimate the effective dimension of the causal set. For a d-dimensional Lorentzian manifold:

$$\langle \tau(x,z)^d \rangle / \langle \mathrm{vol}(x,z) \rangle \to \mathrm{constant}(d)$$

This ratio is computable for each pair (x,z) by finite counting and comparison. Type I.

All invariants are exact finite computations on the partial order. The geometric interpretation — proper time, volume, dimension — is physical, not mathematical. The mathematics is pure finite combinatorics. Type I.

7.6 Causal Set Curvature

Definition 7.9 — Bounded Benincasa-Dowker action:

The Benincasa-Dowker (BD) action is a candidate gravitational action for causal sets, defined purely from the order relation:

$$S_{BD}[C, \prec] = \Sigma_{x \in C} \varepsilon(x)$$

where ε(x) is a sum over "order intervals" involving x, weighted by combinatorial coefficients that encode discrete curvature. The precise formula in d = 4 is:

$$\varepsilon(x) = \Sigma_{y:\ x \prec y} [C_0 - C_1 \cdot n_1(x,y) + C_2 \cdot n_2(x,y) - \ldots]$$

where $n_k(x,y)$ counts the number of elements at "layer k" in the interval between x and y, and the coefficients C_k are specific rational numbers determined by the dimension.

Each ε(x) is a finite sum over the finite set $\{y : x \prec y\}$. The counts n_k are finite. The entire action is a finite sum over |C| events. Type I.

In the continuum limit (the family of causal sets approximating a smooth Lorentzian manifold), the BD action converges to the Einstein-Hilbert action:

$$S_{BD} \to (1/16\pi G) \int R \sqrt{-g}\, d^4x + \text{boundary terms}$$

This convergence is a proven classical result (Benincasa-Dowker 2010) for causal sets faithfully embedded in Lorentzian manifolds. In BST, each finite S_{BD} is exact; the convergence is Type IV.

7.7 Causal Dynamics

Definition 7.10 — Bounded causal partition function:

The partition function over causal histories:

$$Z = \Sigma_{(C,\prec) \in \Omega_{causal}} \exp(iS_{BD}[C,\prec] / \hbar)$$

where Ω_{causal} is the finite set of causal sets of a given size (up to some maximum $|C| \leq N_{max}$).

The number of partial orders on N elements is finite (bounded by 2^{N^2}). The sum is therefore a finite sum of complex numbers in $\mathbb{C}_B(k^4)$. Type I.

Note: this is a Lorentzian partition function (the weight is $\exp(iS/\hbar)$, not $\exp(-S)$), reflecting the Lorentzian signature of spacetime. The sum oscillates rather than being dominated by a saddle point. This makes numerical evaluation harder but does not change the mathematical status: it is a finite sum. Type I.

Observables: for any gauge-invariant observable $O[C,\prec]$ (a function that depends only on the isomorphism class of the causal set):

$$\langle O \rangle = (1/Z) \Sigma_{(C,\prec)} O[C,\prec] \cdot \exp(iS_{BD}/\hbar)$$

Finite sum. Type I.

Definition 7.11 — Bounded causal growth models:
An alternative to the sum-over-histories: sequential causal growth, where the causal set is built one event at a time.

A classical sequential growth (CSG) model is a Markov chain (Definition 9.10 of the AFB paper) on the space of causal sets:

At each step n:
1. The current causal set C_n has n events.
2. A new event x is added.
3. The causal relations between x and the existing events are chosen according to a transition rule $T(C_n \to C_{n+1})$.

If T satisfies general covariance (the probability depends only on the isomorphism class of C_n, not on the labelling of events), the model is called a "classical sequential growth model" (Rideout-Sorkin 2000).

In BST: each step is a finite transition on a finite state space. The growth process is a bounded Markov chain, terminating after N_{max} steps. Convergence properties of the chain (mixing time, stationary distribution) are computable by the spectral gap of the transition matrix (Part III §3.8 of this volume). Type I.

7.8 The GR Emergence Theorem Stack

This section does NOT prove that general relativity emerges from the discrete structures above. It identifies, precisely, the theorems that would need to be proved, stating each as an explicit open problem with its imports.

Open Problem 7.1 — Lorentzian reconstruction:

Statement: Given a family of bounded causal sets {(C_n, <_n)}_{n ∈ metatheory} of increasing size, with each C_n faithfully embeddable (in a sense to be defined) in a Lorentzian manifold (M, g), prove that the manifold (M, g) is uniquely recoverable (up to isometry) from the order and counting data of the family.

Status: Partial classical results exist (Malament 1977: the causal order determines the conformal structure; Bombelli-Meyer theorem: adding volume data determines the metric up to a global constant). A full reconstruction theorem for finite causal sets approximating general Lorentzian manifolds is open.

What BST provides: the causal combinatorics of §§7.5–7.6.
What BST needs: the reconstruction proof itself.

Open Problem 7.2 — Curvature convergence:

Statement: Prove that the discrete curvature proxy (deficit angles for Regge, or the BD action density for causal sets) converges to the Ricci/scalar curvature of the limit manifold in the refinement family.

Status: For Regge calculus, convergence results exist in specific cases (Cheeger-Müller-Schrader 1984 for Regge → smooth in 2D; partial results in higher dimensions). For causal sets, the BD action convergence is proved (Benincasa-Dowker 2010) for faithful embeddings.

What BST provides: the finite curvature quantities (Definitions 7.4, 7.9). What BST needs: the convergence theorems, which are family-level (Type IV) statements.

Open Problem 7.3 — Action convergence:

Statement: Prove that the Regge action S_R (or the BD action S_{BD}) converges to the Einstein-Hilbert action S_{EH} in the refinement family.

Status: Regge's original claim (1961) is established in specific settings. A fully general convergence theorem with explicit error bounds is open.

What BST provides: the finite actions (Definitions 7.5, 7.9). What BST needs: the convergence proof.

Open Problem 7.4 — Matter coupling:

Statement: Show that the lattice gauge/quantum matter sectors from Parts IV–V of this volume, coupled to discrete gravity via the total action $S_{total} = S_{gravity} + S_{matter}$, induce a recovered stress-energy tensor $T_{\mu\nu}$ in the continuum limit.

Status: Lattice gauge theory on fixed backgrounds is well-developed (Part V). Dynamical triangulations with matter are studied numerically. A rigorous coupling theorem is open.

What BST provides: the finite gauge theory (Part V) and the finite gravity (§§7.1–7.7). What BST needs: the coupling convergence theorem.

Open Problem 7.5 — Einstein recovery:

Statement: Show that the coupled discrete field equations (Regge equations + matter) converge to Einstein's field equations $G_{\mu\nu} = 8\pi G\, T_{\mu\nu}$ in the continuum limit.

Status: Open. This is the culmination of Problems 7.1–7.4.

Dependencies:
7.5 requires 7.4 (matter coupling)
7.4 requires 7.3 (action convergence) + Part V
7.3 requires 7.2 (curvature convergence)
7.2 requires 7.1 (Lorentzian reconstruction)

The chain is linear: reconstruction → curvature → action → matter → Einstein. Each step depends on all previous steps.

Plain language: We have five open problems forming a linear chain. The foundation — finite Regge geometry, finite causal sets, finite partition functions — is built (Tier 1). The intermediate steps — curvature convergence, action convergence — have partial classical support but are not fully proved in the BST setting (Tier 2). The final goal — Einstein recovery — is a research programme (Tier 3). The value of this Part is not that it achieves GR emergence, but that it identifies exactly what remains to be proved and builds the finite structures from which the proofs would start.

7.9 The Holographic Parallel

7.9.1 The structural correspondence

The Axiom of Finite Bounds asserts that every set has finite cardinality bounded by some n_M. The Bekenstein-Hawking entropy bound asserts that the maximum entropy (information content) of a region of space is proportional to its surface area in Planck units:

$$S_{max} = A / (4 \ell_P^2)$$

For the observable universe, this gives approximately $S_{max} \sim 10^{122}$ (in bits) or, as a bound on the number of distinguishable states, $N_{max} \sim \exp(10^{122}) \sim 10^{10^{122}}$. A more aggressive estimate — the total number of Planck-volume cells in the observable universe — gives $n_M \sim 10^{185}$.

The structural parallel:
- BST: there exists a finite upper bound n_M on the size of any set.
- Holography: there exists a finite upper bound on the information content of any spatial region.

Both assert a finite ceiling on what can exist in a bounded domain. Neither specifies the ceiling's value — BST leaves n_M as a parameter; holography relates it to the geometry of the bounding surface.

7.9.2 What is NOT claimed

This volume does not claim:

(a) That BST derives the holographic bound. The holographic bound is a physical conjecture (supported by black hole thermodynamics and string theory) about the information content of spatial regions. BST is a mathematical framework with a finite ceiling. The correspondence is structural, not deductive.

(b) That the holographic bound derives BST. The physical bound does not imply a mathematical axiom. The connection, if any, would require a bridge theory relating mathematical ontology to physical information content.

(c) That $n_M = 10^{185}$ (or any other specific value). The AFB paper's two formulations of AFB (schema and metatheoretic) both leave n_M as an unspecified finite parameter. The holographic estimate is a calibration benchmark, not a theorem.

What IS observed: the Cayley-Dickson algebraic tower of physics ($\mathbb{C}$ for QM, $\mathbb{H}$ for spinors, $\mathbb{O}$ for exceptional groups) requires $k^{16} \leq n_M$. If $n_M \sim 10^{185}$, then $k \leq 10^{11.6}$ — sufficient for all practical physics. The fact that the algebraic resource requirements of the Standard Model fit comfortably within the holographic bound is a numerical observation, not a theorem. Whether it is

a coincidence or a consequence of deeper structure
is an open question.

7.10 What Part VII Recovers

Result	Type	Section
Regge manifold (K, ℓ)	I	7.1
Simplex volumes (Cayley-Menger)	I/III	7.1
Dihedral angles	I/III	7.1
Deficit angles (curvature)	I/III	7.2
Regge action	I	7.3
Schläfli identity	I	7.4
Regge equations (discrete Einstein)	I	7.4
Matter-coupled Regge equations	I	7.4
Causal set (finite partial order)	I	7.5
Proper time surrogate	I	7.5
Volume surrogate	I	7.5
Dimension estimator	I	7.5
Benincasa-Dowker action	I	7.6
Causal partition function	I	7.7
Causal growth models (Markov)	I	7.7
Lorentzian reconstruction	Open	7.8
Curvature convergence	Open	7.8
Action convergence	Open	7.8
Matter coupling convergence	Open	7.8
Einstein recovery	Open	7.8
Holographic calibration	Obs.	7.9

What is NOT recovered: smooth Lorentzian manifolds (replaced by finite simplicial or causal structures). The Einstein field equations as internal BST theorems (they are the conjectured family-level limit of the discrete equations). The equivalence principle as a geometric statement about smooth geodesics (replaced by longest chains in causal sets or geodesics on Regge skeletons). Continuous diffeomorphism invariance (replaced by combinatorial symmetries of the discrete structures).

What is gained: every geometric quantity — volume, angle, curvature, action — is a computable finite number. Singularities do not arise: there is no point where curvature diverges, because every deficit angle is a finite number bounded by 2π. The gravitational partition function is a finite sum (no functional integral over infinite-dimensional spaces of metrics). The Regge equations are a finite system of algebraic equations. The causal set approach provides a background-independent formulation where the discrete spacetime structure is itself a dynamical variable — without the infinite-dimensional diffeomorphism group of smooth GR.

End of Part VII

Part VIII: Condensed Matter and Many-Body Physics

Maturity: Tier 1

AFB imports: ℂ_B(k^4) (Part VIII of AFB, Theorem 8.16); finite Hilbert spaces and spectral theorem (Part XI, §§11.4–11.5, Theorem 11.4, Definition 11.8); tensor products

(Part XI, §11.6, Definition 11.10); antisymmetric tensors (Part XI, §11.6, Definition 11.15); matrix exponential (Part XI, §11.7, Definition 11.16); Weyl and Davis-Kahan stability (Part XI, Theorems 11.5–11.6); bounded probability and expectation (Part IX, §9.3.4); DFT and FFT (Part IX, §9.3.2); representation theory and characters (Part XIII, Theorems 13.3–13.7)

Volume imports: Part III (statistical mechanics, partition functions, Monte Carlo); Part IV (quantum mechanics, spectral verification, entanglement); Part VI (simplicial topology, Laplacian spectrum, Betti numbers)

The physics

Condensed matter physics is the study of the collective behaviour of large numbers of interacting particles — the physics of solids, liquids, magnets, superconductors, and exotic quantum phases. It is the largest subfield of physics by number of practitioners and the most directly connected to technology: semiconductors, magnetic storage, superconducting MRI magnets, and the emerging field of topological quantum computing all rest on condensed matter theory. The mathematical models are defined on finite lattices with finite state spaces — making condensed matter the branch of physics where BST's finite foundations are not merely adequate but optimal.

What experiments confirm

Magnetism and the Ising model. The Ising model (§8.1) captures the essential physics of ferromagnetic phase transitions. The critical temperature of the 2D Ising model, computed exactly by Onsager (1944), matches experimental measurements in quasi-2D materials like Rb_2CoF_4 to within a few percent. Critical exponents — the power laws governing magnetisation, susceptibility, and correlation length near the transition — are confirmed across dozens of materials in the same universality class, matching theoretical predictions to 0.1%.

Quantum magnetism. The antiferromagnetic Heisenberg model (§8.2) describes quantum magnets. Neutron scattering measurements of the magnon dispersion relation $\omega(k)$ in materials like La_2CuO_4 and $KCuF_3$ match the predictions from exact diagonalisation and spin-wave theory — the same finite-matrix spectral computation this Part constructs — to within a few percent across the entire Brillouin zone.

Strongly correlated electrons. The Hubbard model (§8.3) is the standard model of strongly correlated electron physics. The Mott metal-insulator transition it predicts is observed in vanadium oxides (V_2O_3), organic conductors, and cold-atom experiments. The charge gap Δ at the Mott transition is measured by optical conductivity and photoemission, and matches Hubbard model predictions from exact diagonalisation and dynamical mean-field theory.

Topological phases. The integer quantum Hall effect (von Klitzing 1980, Nobel Prize 1985) produces Hall conductance quantised to 1 part in 10^9 — one of the most precise measurements in all of physics. The quantisation is explained by the Chern number (§8.4), an integer topological invariant computed by the finite linear algebra of this Part. Topological insulators predicted by band-theory Chern number calculations were subsequently discovered in Bi_2Se_3 and Bi_2Te_3, confirming the predictions.

Superconductivity. The BCS theory (§8.6) predicted the superconducting energy gap Δ and its temperature dependence, confirmed by tunnelling experiments (Giaever 1960, Nobel Prize 1973). The BCS gap equation is a finite nonlinear system on a finite lattice.

Phonons. The phonon dispersion relations of crystalline solids — computed by diagonalising the dynamical matrix (§8.5) — are measured by inelastic neutron scattering and confirmed to within a few percent for hundreds of materials, from silicon to diamond to iron.

What BST constructs and why it suffices

Condensed matter physics is already finite. Real crystals have finite numbers of atoms. Real experiments measure finite lattices. The computational methods — exact diagonalisation, quantum Monte Carlo, DMRG, tensor networks — are all finite computations on finite state spaces. The infinite-volume "thermodynamic limit" is a mathematical convenience for obtaining sharp phase transitions, not a physical requirement: every real material has a finite number of atoms, and every real phase transition has finite rounding.

BST makes this finiteness foundational rather than incidental. Every Hilbert space is finite-dimensional. Every Hamiltonian is a finite matrix. Every spectral gap is computable. Every topological invariant is an integer computed by finite linear algebra. The Pauli exclusion principle is implemented by the exterior algebra (antisymmetric tensors, Definition 11.15). Symmetry sectors are found by Maschke's theorem (Part XIII). The physics does not change; the foundations match what the physics has always been.

8.1 The Quantum Ising Model

Part III §3.6 developed the classical Ising model as a statistical mechanics system with partition function Z = Σ exp(−βH). The quantum Ising model adds quantum dynamics: spins become operators, and the Hamiltonian includes non-commuting terms that generate quantum fluctuations.

Definition 8.1 — Bounded transverse-field Ising model:

```
System: N spin-½ sites on a finite graph G = (V, E),
with |V| = N.

Hilbert space: ℋ = (ℂ_B(k⁴)²)^{⊗N} = ℂ_B(k⁴)^{2^N}

(the N-fold tensor product of qubit spaces,
Definition 11.10 of the AFB paper).

dim(ℋ) = 2^N. Finite. ✓

Basis: the computational basis {|σ₁ ... σ_N⟩} where
each σ_i ∈ {↑, ↓}, i.e. the 2^N spin configurations.

Hamiltonian:
H = −J Σ_{(i,j) ∈ E} σ_i^z σ_j^z − h Σ_{i ∈ V} σ_i^x

where:
```

- $\sigma_i^z = I \otimes \ldots \otimes \sigma_z \otimes \ldots \otimes I$ (Pauli-Z on site i, identity on all other sites)
- $\sigma_i^x = I \otimes \ldots \otimes \sigma_x \otimes \ldots \otimes I$ (Pauli-X on site i)
- $J \in \mathbb{R}_B(k)$ is the Ising coupling
- $h \in \mathbb{R}_B(k)$ is the transverse field

H is a $2^N \times 2^N$ Hermitian matrix over $\mathbb{C}_B(k^4)$. Type I.

The two terms in H do not commute: $[\sigma_i^z \sigma_j^z, \sigma_i^x] \neq 0$. This non-commutativity is the source of quantum fluctuations and quantum phase transitions — absent in the classical Ising model of Part III.

8.1.1 Exact diagonalisation

For $N \leq \sim 20$ ($2^N \leq \sim 10^6$): the Hamiltonian matrix can be explicitly constructed and diagonalised by the spectral theorem (Theorem 11.4 of the AFB paper).

Output:
- All 2^N energy eigenvalues $E_0 \leq E_1 \leq \ldots \leq E_{2^N - 1}$
- All eigenstates $|E_i\rangle$
- The spectral gap $\gamma = E_1 - E_0$
- The ground state energy E_0 and ground state $|E_0\rangle$
- Any observable $\langle A\rangle_0 = \langle E_0|A|E_0\rangle$ in the ground state

All exact finite computations. Type I.

8.1.2 Quantum phase transition

The transverse-field Ising model on a d-dimensional lattice exhibits a quantum phase transition at a critical value h_c of the transverse field:

For $h < h_c$ (strong coupling): the ground state is magnetically ordered ($\langle\sigma^z\rangle \neq 0$). Spontaneous symmetry breaking.

For $h > h_c$ (weak coupling): the ground state is paramagnetic ($\langle\sigma^z\rangle = 0$). The Z_2 symmetry is restored.

On a finite lattice: there is no true phase transition (Part III, §3.5 — all thermodynamic functions are smooth for finite systems). Instead, there is a sharp crossover: the spectral gap $\gamma(h)$ has a minimum near h_c, and the minimum sharpens as N grows.

The spectral gap at the putative critical point:
$\gamma(h_c, N) \sim N^{-z}$

where z is the dynamical critical exponent, governs the finite-size scaling. This scaling exponent is computable from the exact spectrum for small N and estimable from Monte Carlo for larger N.

In BST: each finite spectrum is exact (Type I). The phase transition is the family-level (Type IV) limit of the family of finite spectra. The critical exponent z is extracted from the finite-size scaling of the computable spectral gap — a numerical observation about the family, not a single BST theorem.

8.2 The Heisenberg Model

Definition 8.2 — Bounded Heisenberg model:

System: N spin-½ sites on a finite graph $G = (V, E)$.

Hilbert space: $\mathcal{H} = \mathbb{C}_B(k^4)^{2^N}$ (same as Ising).

Hamiltonian:

$$H = J \sum_{(i,j) \in E} \vec{\sigma}_i \cdot \vec{\sigma}_j$$
$$= J \sum_{(i,j) \in E} (\sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y + \sigma_i^z \sigma_j^z)$$

This is a $2^N \times 2^N$ Hermitian matrix. $J > 0$ gives the antiferromagnet; $J < 0$ gives the ferromagnet.

Symmetry: H commutes with the total spin operators $\vec{S}_{total} = \sum_i \vec{\sigma}_i / 2$. The Hilbert space decomposes into sectors of definite total spin S by the representation theory of $SU(2)_B$ (Part XIII of the AFB paper, applied in Part IV §4.3 of this volume).

The decomposition:

$$\mathcal{H} = \oplus_S \mathcal{H}_S$$

where each $\mathcal{H}_S$ is an irreducible representation of $SU(2)_B$ of dimension $2S+1$, and the multiplicities are computable by character inner products (Theorem 13.6).

Diagonalisation within each sector reduces the effective Hamiltonian size: the largest sector has dimension $C(N, N/2) \approx 2^N / \sqrt{N}$, smaller than the full 2^N by a factor of $\sqrt{N}$. This is the first practical payoff of representation theory for many-body physics.

Type I. All decompositions and diagonalisations are exact finite linear algebra.

8.2.1 Magnons and spin waves

For the ferromagnetic Heisenberg model ($J < 0$) on a regular lattice, the ground state is the fully aligned state $|\uparrow\uparrow\ldots\uparrow\rangle$ (all spins up).

The lowest-energy excitations are single spin flips delocalised across the lattice — magnons:

$|k\rangle = (1/\sqrt{N}) \Sigma_j \exp(2\pi ikj/N) |\uparrow...\downarrow_j...\uparrow\rangle$

These are the DFT (Definition 9.1 of the AFB paper) of the localised spin-flip states.

The magnon energy:
$E(k) = E_0 + 2|J|(1 - \cos(2\pi k/N))$

This is the discrete dispersion relation — identical in structure to the lattice wave dispersion of Part II §2.1.2 and the coupled oscillator normal modes of Part I §1.6.3.

The pattern recurs: the DFT diagonalises the translation-invariant Hamiltonian, and the eigenvalues follow from the discrete Laplacian spectrum. All exact at fixed N. Type I.

8.3 The Hubbard Model

The Hubbard model is the simplest model of interacting electrons on a lattice. It is the foundation for understanding magnetism, metal-insulator transitions, and high-temperature superconductivity.

Definition 8.3 — Bounded Hubbard model:
System: a finite lattice $G = (V, E)$ with $|V| = N$ sites. Electrons are fermions with spin up and down.

Single-particle states: N sites × 2 spins = 2N single-particle states.

Many-body Hilbert space: the antisymmetric subspace of the 2N-fold tensor product (fermionic Fock space):

$\mathcal{H} = \Lambda(\mathbb{C}_B(k^4)^{2N})$

where Λ denotes the exterior algebra (Definition 11.15 of the AFB paper) — antisymmetric tensor products enforce the Pauli exclusion principle.

$\dim(\mathcal{H}) = 2^{2N} = 4^N$ (each of 2N single-particle states is either occupied or empty).

For N = 6 sites: $\dim(\mathcal{H}) = 4^6 = 4096$. Manageable.
For N = 10 sites: $\dim(\mathcal{H}) = 4^{10} \approx 10^6$. Still exactly diagonalisable.
For N = 20 sites: $\dim(\mathcal{H}) = 4^{20} \approx 10^{12}$. Requires large-scale computation but is still a finite, well-defined problem.

Hamiltonian:

$$H = -t \sum_{(i,j)\in E,\ \sigma} (c^\dagger_{i\sigma} c_{j\sigma} + h.c.) + U \sum_i n_{i\uparrow} n_{i\downarrow}$$

where:

- $c^\dagger_{i\sigma}$, $c_{i\sigma}$ are fermionic creation/annihilation operators at site i with spin σ — represented as $4^N \times 4^N$ matrices acting on $\mathcal{H}$ by the Jordan-Wigner transformation (a finite matrix construction)
- $n_{i\sigma} = c^\dagger_{i\sigma} c_{i\sigma}$ is the number operator
- $t \in \mathbb{R}_B(k)$ is the hopping amplitude
- $U \in \mathbb{R}_B(k)$ is the on-site interaction strength
- h.c. = Hermitian conjugate

The first term (hopping) describes electrons tunnelling between adjacent sites. The second term (interaction) penalises double occupancy. The competition between kinetic energy (t) and interaction energy (U) drives the physics — metal-insulator transitions, magnetism, and (in 2D) superconductivity.

H is a $4^N \times 4^N$ Hermitian matrix over $\mathbb{C}_B(k^4)$. Type I.

8.3.1 Symmetries and sector decomposition

The Hubbard Hamiltonian commutes with:

Total particle number: $\hat{N} = \sum_{i,\sigma} n_{i\sigma}$
  (eigenvalues 0, 1, ..., 2N)

Total spin-z: $S^z = (1/2) \sum_i (n_{i\uparrow} - n_{i\downarrow})$
  (eigenvalues −N, −N+1, ..., N)

Lattice translations (if the lattice is periodic):
  momentum quantum number k

The Hilbert space decomposes into sectors labelled by (N_e, S^z, k):

$$\mathcal{H} = \oplus_{N_e, S_z, k}\ \mathcal{H}_{N_e, S_z, k}$$

The largest sector is the half-filled case $N_e = N$, $S^z = 0$, which has dimension $C(N, N/2)^2$. For N = 10: this is $C(10,5)^2 = 252^2 = 63504$ — compared to the full $4^{10} \approx 10^6$.

Diagonalisation within each sector is more efficient than the full problem. By Part XIII's representation theory (Maschke decomposition + character projectors), the sector decomposition is constructive. Type I.

8.3.2 The Mott transition

At half filling ($N_e = N$) and zero temperature:

For $U/t \ll 1$ (weak interaction): the system is metallic — electrons hop freely, the ground state has no charge gap.

For $U/t \gg 1$ (strong interaction): the system is a Mott insulator — double occupancy is suppressed, each site has exactly one electron, and there is a charge gap $\Delta \approx U$.

The Mott transition occurs at a critical $(U/t)_c$.

On a finite lattice: the charge gap
$\Delta = E_0(N_e = N+1) + E_0(N_e = N-1) - 2E_0(N_e = N)$
is computable by diagonalising H in three particle-number sectors. Whether $\Delta > 0$ is decidable. Type I.

The true Mott transition (where Δ vanishes discontinuously or develops a singularity) is a family-level (Type IV) phenomenon — it requires the thermodynamic limit $N \to \infty$. The finite-lattice precursors (a rapid but smooth crossover in $\Delta(U/t)$ as N grows) are exactly computable.

8.4 Topological Phases and Berry Phase

Topological phases of matter are quantum phases characterised not by local order parameters (like magnetisation) but by global topological invariants. In BST, these invariants are computable by finite linear algebra.

Definition 8.4 — Bounded Berry phase:
For a family of Hamiltonians $H(\lambda)$ parametrised by $\lambda = (\lambda_1, ..., \lambda_m) \in \mathbb{R}_B(k)^m$, with non-degenerate ground state $|\psi_0(\lambda)\rangle$:

The Berry connection is the 1-form:
$A_\mu(\lambda) = i\langle\psi_0(\lambda)|\partial\psi_0/\partial\lambda_\mu\rangle$

where $\partial/\partial\lambda_\mu$ is the finite difference quotient (Definition 9.2 of the AFB paper) in the parameter λ_μ.

The Berry phase around a closed loop γ in parameter space is:

$$\Phi_B = \oint_\gamma A \cdot d\lambda = \Sigma_{\text{edges of } \gamma} A_\mu \cdot \Delta\lambda_\mu$$

This is a finite sum of inner products — each computed by finite linear algebra on $\mathcal{H}$. Type I.

The Berry curvature (the "magnetic field" in parameter space) is:

$$F_{\mu\nu} = \partial A_\nu/\partial\lambda_\mu - \partial A_\mu/\partial\lambda_\nu$$

(finite differences of inner products). This is a

2-form on the parameter space — connecting directly to the exterior algebra of Part XI §11.8 and the cochain complexes of Part VI.

Definition 8.5 — Bounded Chern number:
For a 2D lattice system with Bloch Hamiltonian $H(k)$ parametrised by crystal momentum $k = (k_x, k_y)$ on the finite Brillouin zone (a discrete torus of $N_x \times N_y$ points):

The Chern number of the n-th band is:

$$C_n = (1/2\pi) \Sigma_{\text{plaquettes } p} F_p$$

where F_p is the Berry curvature on each plaquette of the discrete Brillouin zone, computed by:

$$F_p = \text{Im} \ln(U_{12} \cdot U_{23} \cdot U_{34} \cdot U_{41})$$

with $U_{ij} = \langle \psi_n(k_i) | \psi_n(k_j) \rangle$ being the overlap between ground states at adjacent k-points.

C_n is an integer (by the finite-lattice analogue of the Gauss-Bonnet theorem). It is exactly computable from the ground state wavefunctions at each k-point — each obtained by diagonalising the Bloch Hamiltonian $H(k)$ via the spectral theorem (Theorem 11.4).

The Chern number is a topological invariant: it does not change under smooth (continuous in the finite-difference sense) deformations of $H(k)$ that do not close the gap between bands. It classifies quantum Hall states and topological insulators.

Type I. An integer computed by finite linear algebra.

Connection to Part VI: The Chern number is the first Chern class of a U(1) bundle over the Brillouin zone torus — computed by the cohomological machinery of Part VI applied to the Berry connection. The lattice gauge theory of Part II §2.5 (U(1) on a lattice) provides the framework: the Berry connection is a U(1) lattice gauge field on the Brillouin zone, and the Chern number is its topological charge.

Theorem 8.1 — Bulk-boundary correspondence (finite):
For a finite 2D lattice with boundaries (a strip of width W and length L):

The number of edge modes crossing the Fermi level equals the Chern number of the bulk band below the Fermi level:

$N_{edge} = C_{bulk}$

Proof sketch: The edge modes are eigenstates of the finite Hamiltonian localised near the boundary. Their count is determined by the spectral flow of the Hamiltonian as a function of the boundary parameter — which is topologically constrained by the Chern number of the bulk bands.

For a specific finite lattice: both sides are exactly computable. N_{edge} by diagonalising the strip Hamiltonian and counting edge-localised eigenvalues in the gap. C_{bulk} by the plaquette formula (Definition 8.5). The equality is verifiable by direct computation. Type I.

The infinite-system bulk-boundary correspondence is the family-level (Type IV) limit of this finite identity.

8.5 Lattice Phonons

Phonons — quantised lattice vibrations — connect the classical coupled oscillators of Part I §1.6.3 to the quantum theory of Part IV.

Definition 8.6 — Bounded phonon Hamiltonian:

System: N atoms on a d-dimensional lattice, each displaced from equilibrium by $u_i \in \mathbb{R}_B(k)^d$.

The harmonic Hamiltonian (expanding the potential to second order around equilibrium):

$$H = \Sigma_i |p_i|^2/(2m) + (1/2) \Sigma_{i,j} u_i^T \cdot D_{ij} \cdot u_j$$

where D_{ij} is the dN × dN dynamical matrix — a real symmetric matrix over $\mathbb{R}_B(k)$, encoding the spring constants between all atom pairs.

Normal mode decomposition:
Diagonalise D by the spectral theorem (Theorem 11.4):
$$D = \Sigma_\alpha \omega_\alpha^2 |e_\alpha\rangle\langle e_\alpha|$$

The eigenvalues ω_α^2 are the squared phonon frequencies. The eigenvectors $|e_\alpha\rangle$ are the normal mode displacement patterns.

For each normal mode α, the quantum Hamiltonian is a harmonic oscillator (Part IV §4.2.2):
$$H_\alpha = \hbar\omega_\alpha (a_\alpha^\dagger a_\alpha + 1/2)$$

The total phonon Hamiltonian is:
$H_{phonon} = \Sigma_\alpha \hbar\omega_\alpha (a_\alpha^\dagger a_\alpha + 1/2)$

which is diagonal in the normal mode basis.

All frequencies ω_α are exact eigenvalues of a finite matrix. The phonon spectrum (density of states $g(\omega) = \Sigma_\alpha \delta(\omega - \omega_\alpha)$) is a sum of delta functions at the computed frequencies — a finite histogram. Type I.

8.5.1 The Debye model on a finite lattice

The Debye model approximates the phonon density of states as $g(\omega) \propto \omega^{d-1}$ for $\omega \leq \omega_D$ (the Debye cutoff), and $g(\omega) = 0$ for $\omega > \omega_D$.

In BST on a finite lattice: there is no need for the Debye approximation. The exact phonon spectrum is computable by diagonalising the dynamical matrix. The Debye approximation is the smooth envelope of the exact finite spectrum — useful as a simplified description for large lattices, but unnecessary for small ones.

Thermodynamic quantities from the phonon spectrum:

Internal energy:
$U(T) = \Sigma_\alpha \hbar\omega_\alpha [n_B(\omega_\alpha) + 1/2]$

where $n_B(\omega) = 1/(\exp(\hbar\omega/k_BT) - 1)$ is the Bose distribution (computed by bounded transcendentals).

Heat capacity:
$C(T) = \partial U/\partial T = \Sigma_\alpha (\hbar\omega_\alpha)^2 n_B(\omega_\alpha)(n_B(\omega_\alpha)+1) / (k_BT^2)$

Both are finite sums over the dN normal modes. Type I.

The Debye T^3 law for C(T) at low temperature is the family-level (Type IV) behaviour of the exact finite heat capacity as $N \to \infty$.

8.6 Superconductivity: The BCS Ground State

Definition 8.7 — Bounded BCS variational state:

The Bardeen-Cooper-Schrieffer (BCS) theory of superconductivity is based on a variational ansatz for the ground state of an attractive Hubbard-like model.

On a finite lattice with N_k momentum states (labelled by the discrete momenta k from the DFT):

The BCS trial state:

```
|BCS⟩ = Π_k (u_k + v_k c†_{k↑} c†_{-k↓}) |0⟩

where:
- u_k, v_k ∈ ℝ_B(k) with u_k² + v_k² = 1
- c†_{k↑} creates an electron with momentum k, spin up
- c†_{-k↓} creates an electron with momentum −k, spin down
- |0⟩ is the vacuum state

This is a specific vector in the Fock space
ℋ = ℂ_B(k⁴)^{4^N}. The parameters {u_k, v_k} are
determined by minimising the energy:

E_BCS = ⟨BCS|H|BCS⟩

This is a finite optimisation over 2N_k real
parameters — solvable by the self-consistent BCS
gap equation:

Δ_k = −Σ_{k'} V_{kk'} Δ_{k'} / (2E_{k'})

where E_k = √(ε_k² + Δ_k²) is the quasiparticle
energy and V_{kk'} is the pairing interaction.

The gap equation is a finite system of N_k nonlinear
equations in N_k unknowns — solvable by bounded
iteration (a fixed-point problem, by BI-BST on the
iteration count). Type I.

The superconducting gap Δ = min_k |Δ_k| is computable.
Whether Δ > 0 (superconducting) or Δ = 0 (normal)
is decidable. Type I.
```

Plain language: BCS superconductivity on a finite lattice is a finite variational problem. The BCS gap equation is a finite nonlinear system. The superconducting gap is a computable number. The entire theory of conventional superconductivity — gap equations, quasiparticle spectra, Meissner effect (via the electromagnetic response kernel, also a finite matrix) — is finite linear algebra on a finite lattice.

8.7 What Part VIII Recovers

Result	Type	Section
Quantum Ising model	I	8.1
Exact diagonalisation	I	8.1.1
Quantum phase transition (finite)	I	8.1.2
Quantum phase transition (true)	IV	8.1.2
Heisenberg model	I	8.2
SU(2) sector decomposition	I	8.2
Magnon dispersion relation	I	8.2.1
Hubbard model	I	8.3
Fermionic Fock space (Λ)	I	8.3
Jordan-Wigner transformation	I	8.3
Sector decomposition	I	8.3.1
Mott transition (finite lattice)	I	8.3.2

Mott transition (true)	IV	8.3.2
Berry phase	I	8.4
Berry curvature	I	8.4
Chern number (integer, computable)	I	8.4
Bulk-boundary correspondence	I	8.4
Phonon spectrum (dynamical matrix)	I	8.5
Debye heat capacity (exact)	I	8.5.1
BCS gap equation	I	8.6
Superconducting gap (decidable)	I	8.6

What is NOT recovered: infinite-volume phase transitions with true singularities (Type IV). Continuous symmetry breaking with Goldstone bosons in finite systems (the finite-volume magnetisation is always smooth). The BCS ground state as a coherent state on an infinite Fock space (replaced by a variational state on a finite Fock space). Topological field theory in the continuum (replaced by finite topological invariants of the lattice — Chern numbers, Berry phases). Anderson's "More is Different" emergence hierarchy in its full infinite-volume form.

What is gained: every spectrum is computable by diagonalising a finite matrix. Every gap — spectral gap, charge gap, superconducting gap — is a decidable quantity. Topological invariants (Chern numbers) are integers computed by finite linear algebra. The Pauli exclusion principle is implemented by the exterior algebra (antisymmetric tensors, Definition 11.15 of the AFB paper) — exactly, with no ad hoc anti-commutation rules. Symmetry sectors are found by the constructive Maschke decomposition (Part XIII) — reducing the effective problem size. Monte Carlo sampling (Part III §3.8) provides efficient estimation for systems too large to diagonalise exactly.

The philosophical point: condensed matter physics is the branch of physics where finiteness is not a limitation but a feature. Real crystals have finite numbers of atoms. Real experiments measure finite lattices (even if large ones). The theoretical tools — exact diagonalisation, Monte Carlo, DMRG, tensor networks — are all finite computations on finite state spaces. BST makes this finiteness foundational rather than incidental. The physics doesn't change; the foundations become honest.

End of Part VIII

Part IX: Towards the Standard Model

Maturity: Tier 2 (spinor and Higgs kinematics) / Tier 2–3 (Feynman rules, RG) / Tier 3 (full SM assembly)

AFB imports: ℂ_B(k^4), ℍ_B(k^8) (Part VIII, §8.8); bounded linear algebra and operator norms (Part XI, §§11.1–11.5); tensor algebra including antisymmetric tensors (Part XI, §11.6); matrix exponential (Part XI, §11.7); exterior algebra and Hodge star (Part XI, §11.8); representation theory including Maschke, Schur, irreducible decomposition (Part XIII); DFT (Part IX of AFB, §9.3.2); bounded probability (Part IX of AFB, §9.3.4)

Volume imports: Part II (Lorentz group §2.2, U(1)_B gauge theory §2.5); Part IV (quantum mechanics, spectral verification, scattering §4.6); Part V (SU(N)_B gauge groups §5.1, lattice gauge theory §§5.2–5.5, Standard Model gauge group §5.3, Gauss's law §5.7); Part VI (cochain complexes, cohomology); Part VIII (fermionic Fock space §8.3)

The physics

The Standard Model of particle physics is the most comprehensive and precisely tested theory in the history of science. It describes three of the four fundamental forces (electromagnetism, the weak force, the strong force) and classifies all known elementary particles: six quarks, six leptons, the photon, the W and Z bosons, eight gluons, and the Higgs boson. Developed between the 1960s and 1970s (Glashow, Weinberg, Salam, 't Hooft, Veltman, and many others), it combines the gauge theory of Part V with spinor fields (quarks and leptons), the Higgs mechanism (spontaneous symmetry breaking giving mass to the W and Z), and perturbative quantum field theory (Feynman diagrams).

What experiments confirm

The Standard Model has survived every experimental test for over fifty years.

The Higgs boson. Predicted by the Higgs mechanism in 1964, discovered at the LHC in 2012 (Nobel Prize 2013). Its mass (125.1 GeV) and its production and decay rates match SM predictions to within 10–20% — improving with more data. The Higgs mechanism, constructed in §9.3 as a finite spectral analysis of the quadratic fluctuation operator, is experimentally confirmed.

The electron g−2. The anomalous magnetic moment of the electron is the most precisely measured quantity in physics: g/2 = 1.001 159 652 180 73 (experimental, Gabrielse 2023). The SM prediction, computed from ~12,000 Feynman diagrams through 5th order in QED, agrees to 1 part in 10^{12}. Each Feynman diagram is a finite sum on a finite momentum lattice — the computation of §9.4.

The W and Z masses. Predicted by SU(2) × U(1) electroweak theory before their discovery: m_W ≈ 80.4 GeV, m_Z ≈ 91.2 GeV. Measured at CERN (1983) and subsequently at LEP and the Tevatron to 0.01% precision. These masses are algebraic functions of the gauge couplings and the Higgs vev — the finite spectral computation of §9.3.1.

Quark flavour mixing. The CKM matrix — the 3×3 unitary mixing matrix between quark flavour and mass eigenstates — is measured to percent-level precision in B-meson factories (BaBar, Belle, LHCb). Its unitarity is confirmed to 0.1%. CP violation, predicted by the complex phase of the CKM matrix, is observed in kaon and B-meson decays. The CKM matrix is a specific unitary matrix over $\mathbb{C}$_B(k^4) — finite linear algebra.

Asymptotic freedom. The running of the strong coupling α_s from low energies (~1 GeV, where α_s ≈ 0.5) to high energies (~100 GeV, where α_s ≈ 0.1) is measured at the LHC and matches the SM prediction — the discrete beta function of §9.5.1 — across three decades of energy.

What BST constructs and why it suffices

The Standard Model was *developed* using lattice gauge theory and perturbative Feynman diagrams — both finite computations. The proton mass is computed by lattice QCD (Part V). The electron g−2 is computed by evaluating finite sums of Feynman diagrams. The Higgs mass predictions are algebraic functions of the couplings. Every precision test of the Standard Model is a comparison between a finite computation and a finite measurement.

This Part assembles the SM ingredients already constructed in earlier Parts: the gauge group SU(3) × SU(2) × U(1) from Part V, the quantum Hilbert spaces from Part IV, the representation content from Part XIII. It adds the remaining kinematic pieces — Clifford algebra and spinors (§9.1), the fermion representation data (§9.2), the Higgs mechanism (§9.3) — and the dynamical framework — lattice Feynman rules (§9.4), the renormalisation group (§9.5), and effective field theory (§9.6). The honest assessment: the algebraic skeleton is Tier 1 (it exists as exact finite objects); the dynamical completeness — chiral fermions, anomaly cancellation, the continuum limit — is Tier 2–3 (frameworks definable, proofs partly open). §9.7 states precisely where each boundary lies.

9.1 Bounded Clifford Algebra and Spinors

Fermions — quarks and leptons — are described by spinor fields. Spinors transform under the Lorentz group in a representation that is not a tensor representation: it requires the double cover Spin(1,3) ≅ SL(2,ℂ). The algebraic foundation is the Clifford algebra.

Definition 9.1 — Bounded Clifford algebra:

```
The Clifford algebra Cl(1,3) over ℂ_B(k⁴) is the
algebra generated by four elements γ⁰, γ¹, γ², γ³
satisfying the anticommutation relation:

{γ^μ, γ^ν} = γ^μ γ^ν + γ^ν γ^μ = 2η^{μν} I

where η = diag(−1,+1,+1,+1) is the Minkowski metric
(Definition 2.3 of Part II of this volume).

The standard (Dirac) representation uses 4×4 matrices
over ℂ_B(k⁴):

γ⁰ = ⎛ I₂   0 ⎞      γⁱ = ⎛  0    σ_i ⎞
     ⎝ 0  −I₂ ⎠           ⎝−σ_i   0   ⎠

where σ_i are the Pauli matrices (Part IV §4.2.1)
and I₂ is the 2×2 identity.

Verification of {γ^μ, γ^ν} = 2η^{μν}:
By direct 4×4 matrix multiplication in ℂ_B(k⁴).
16 products to check. Each is exact. Type I.

The Clifford algebra is 16-dimensional as a vector
space (spanned by I, γ^μ, γ^μγ^ν for μ<ν,
γ^μγ^νγ^ρ for μ<ν<ρ, and γ⁰γ¹γ²γ³ = γ⁵).
All basis elements are specific 4×4 matrices over
ℂ_B(k⁴). Type I.
```

Definition 9.2 — Bounded Dirac spinor:
A Dirac spinor is a 4-component column vector:

$\psi = (\psi_1, \psi_2, \psi_3, \psi_4)^T \in \mathbb{C}_B(k^4)^4$

Under a Lorentz transformation $\Lambda \in SO^+(1,3)_B$ (Definition 2.4 of Part II), the spinor transforms as:

$\psi \to S(\Lambda)\, \psi$

where $S(\Lambda) = \exp((i/4)\, \omega_{\mu\nu}\, \sigma^{\mu\nu})$ is the spinor representation, with:
- $\omega_{\mu\nu}$ the antisymmetric parameters of Λ
- $\sigma^{\mu\nu} = (i/2)[\gamma^\mu, \gamma^\nu]$ the Lorentz generators in the spinor representation

$S(\Lambda)$ is a 4×4 matrix over $\mathbb{C}_B(k^4)$, computed by the matrix exponential (Definition 11.16 of the AFB paper). Type I/III.

Chirality: the matrix $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3$ satisfies $(\gamma^5)^2 = I$ and $\{\gamma^5, \gamma^\mu\} = 0$.

Left-handed spinor: $P_L\, \psi$ where $P_L = (1-\gamma^5)/2$.
Right-handed spinor: $P_R\, \psi$ where $P_R = (1+\gamma^5)/2$.

The chiral projectors P_L, P_R are exact 4×4 matrices over $\mathbb{C}_B(k^4)$ with $P_L + P_R = I$, $P_L^2 = P_L$, $P_R^2 = P_R$, $P_L P_R = 0$. Type I.

Definition 9.3 — Bounded Dirac operator on a lattice:
On a finite lattice K with gauge connection $U: E(K) \to G_B$ (Part V, Definition 5.5), the lattice Dirac operator is:

$$(D_{lat}\, \psi)(x) = (1/2a)\, \Sigma_\mu\, \gamma^\mu\, [U_\mu(x)\, \psi(x+\hat{\mu}) - U_\mu(x-\hat{\mu})^\dagger\, \psi(x-\hat{\mu})]$$

where:
- a is the lattice spacing
- $\hat{\mu}$ is the unit vector in direction μ
- $U_\mu(x)$ is the gauge link from x to $x+\hat{\mu}$
- The sum runs over the d = 4 spacetime directions

This is a finite matrix acting on the space of lattice spinor fields:
$\psi: V(K) \to \mathbb{C}_B(k^4)^4$

The total Hilbert space is $\mathbb{C}_B(k^4)^{4|V(K)|}$ for a single flavour of fermion, or $\mathbb{C}_B(k^4)^{4N_f|V(K)|}$ for N_f flavours.

```
D_lat is a (4|V(K)|) × (4|V(K)|) matrix over ℂ_B(k⁴).

Properties:
- γ⁵-Hermiticity: γ⁵ D_lat γ⁵ = D_lat† (by the
  anticommutation {γ⁵, γ^μ} = 0 and the unitarity
  of the gauge links). ✓
- Gauge covariance: under a gauge transformation
  g: V(K) → G_B, D_lat transforms covariantly:
  D_lat → g(x) D_lat g(x)⁻¹. ✓

Both verified by direct finite matrix computation.
Type I.

Note — the doubling problem: the naive lattice Dirac
operator produces 2^d = 16 species of fermion instead
of 1 (the Nielsen-Ninomiya theorem). Solutions include
Wilson fermions (add a mass-like term that lifts the
doublers), staggered fermions (reduce the degrees of
freedom by a factor of 4), or domain-wall/overlap
fermions (add an extra dimension). Each is a specific
modification of D_lat — a different finite matrix on
the same lattice. The choice affects the approximation
properties but not the finiteness: all are finite
matrices over ℂ_B(k⁴). Tier 2.
```

9.2 The Standard Model Matter Content

Part V §5.3 defined the Standard Model gauge group G_SM = SU(3)_B × SU(2)_B × U(1)_B and listed the representation content of one generation. This section develops the matter sector more explicitly.

Definition 9.4 — Bounded Standard Model fermion representations:

```
One generation of Standard Model fermions consists
of the following fields, each a section of a specific
representation bundle over the lattice:

Left-handed quark doublet:
  Q_L: V(K) → ℂ_B(k⁴)^{3×2×4}
  Transforms as (3, 2, 1/6) under G_SM.
  3 colours × 2 weak isospin × 4 spinor components.

Right-handed up quark:
  u_R: V(K) → ℂ_B(k⁴)^{3×1×4}
  Transforms as (3, 1, 2/3).

Right-handed down quark:
  d_R: V(K) → ℂ_B(k⁴)^{3×1×4}
  Transforms as (3, 1, −1/3).

Left-handed lepton doublet:
  L_L: V(K) → ℂ_B(k⁴)^{1×2×4}
  Transforms as (1, 2, −1/2).
```

```
Right-handed electron:
  e_R: V(K) → ℂ_B(k⁴)^{1×1×4}
  Transforms as (1, 1, −1).

Total per generation per site:
  (3×2 + 3×1 + 3×1 + 1×2 + 1×1) × 4 spinor = 15 × 4 = 60
  complex components per site.

For 3 generations: 180 complex components per site.

On a lattice of |V(K)| sites: the fermion Hilbert
space has dimension proportional to 2^{180|V(K)|}
(fermionic Fock space, as in Part VIII §8.3).

The representation data — the specific matrices by
which G_SM acts on each field — are determined by
Part XIII's representation theory. The action of
SU(3)_B on colour triplets is via the fundamental
3-dimensional representation. The action of SU(2)_B
on weak doublets is via the fundamental 2-dimensional
representation. The U(1)_B hypercharge acts by
phase multiplication exp(iYθ) with the assigned
hypercharge Y.

Type I for the algebraic data.
```

9.3 The Bounded Higgs Mechanism

The Higgs mechanism gives mass to the W and Z bosons (and to the fermions via Yukawa couplings) through spontaneous symmetry breaking. In BST on a finite lattice, the symmetry breaking is not spontaneous in the infinite-volume sense — it is a finite minimisation problem.

Definition 9.5 — Bounded Higgs field:

```
The Higgs field is a scalar field on the lattice
transforming as (1, 2, 1/2) under G_SM:

φ: V(K) → ℂ_B(k⁴)²

(a complex doublet at each lattice site).

The Higgs potential at each site:

V(φ) = −μ² |φ|² + λ |φ|⁴

where μ², λ ∈ ℝ_B(k) with μ² > 0 and λ > 0.

This is the "Mexican hat" potential. It has a
minimum at:

|φ|² = μ²/(2λ) ≡ v²/2
```

where v is the vacuum expectation value (vev).

The lattice Higgs action:

$$S_H = \Sigma_x [|D_\mu \varphi(x)|^2 + V(\varphi(x))]$$

where $D_\mu \varphi = (\varphi(x+\hat{\mu}) - U_\mu(x)\varphi(x))/a$ is the covariant lattice derivative, coupling the Higgs to the gauge field.

S_H is a finite sum over $|V(K)|$ sites. Type I.

9.3.1 Symmetry breaking on a finite lattice

Classically, the Higgs mechanism involves spontaneous symmetry breaking: the vacuum state selects a direction in the SU(2) × U(1) space, breaking it to $U(1)_{em}$ (electromagnetism).

On a finite lattice: the partition function $Z = \Sigma \exp(-S)$ sums over all field configurations, including all orientations of φ. The exact finite partition function preserves the full G_{SM} symmetry (Part III §3.5 — no true symmetry breaking in finite systems).

What BST provides instead:

1. The classical minimum: minimise V(φ) over all constant field configurations. The minimum is at $|\varphi| = v/\sqrt{2}$, a specific element of $\mathbb{R}_B(k)$. This is a finite optimisation problem. Type I.

2. The mass spectrum: expand the action to second order around the minimum. The quadratic fluctuation operator is a finite matrix — its eigenvalues give the particle masses:

   - W boson mass: $m_W = gv/2$
   - Z boson mass: $m_Z = \sqrt{(g^2 + g'^2)}\, v/2$
   - Higgs boson mass: $m_H = \sqrt{(2\lambda)}\, v$
   - Photon mass: $m_\gamma = 0$ (the unbroken $U(1)_{em}$ direction has a zero eigenvalue)

   where g, g' are the SU(2) and U(1) gauge couplings. These are algebraic expressions in $\mathbb{R}_B(k)$. Type I.

3. The Goldstone bosons: the three directions in field space along the "trough" of the Mexican hat correspond to zero eigenvalues of the mass matrix — the would-be Goldstone bosons. In the unitary gauge, they are "eaten" by the W± and Z to give them mass. This is a finite linear algebra

statement: three zero eigenvectors of the mass matrix are absorbed into the longitudinal polarisations of the gauge bosons by a finite gauge transformation. Type I.

The full quantum symmetry breaking – with tunnelling between degenerate vacua and the order parameter selecting a direction – is a family-level (Type IV) phenomenon. It requires the thermodynamic limit $|V(K)| \to \infty$ for the tunnelling amplitude to vanish. At each finite lattice, the physics is a finite optimisation + finite spectral problem.

9.4 Bounded Feynman Rules on a Lattice

Perturbative quantum field theory — the computational engine of the Standard Model — uses Feynman diagrams to organise the expansion of the partition function and correlation functions in powers of the coupling constants. In BST, Feynman diagrams are finite combinatorial objects, and each diagram evaluates to a finite number.

Definition 9.6 — Bounded lattice propagator:
The free propagator on a finite lattice is the inverse of the quadratic (kinetic + mass) operator:

For a scalar field: $G(x,y) = (-\Delta_{lat} + m^2)^{-1}_{xy}$

where Δ_{lat} is the lattice Laplacian (Definition 9.6 of the AFB paper) and m is the mass. This is the inverse of a finite matrix – computed by Part XI finite linear algebra. Type I.

In momentum space (via DFT, Definition 9.1 of the AFB paper):

$\tilde{G}(p) = 1/(\tilde{p}^2 + m^2)$

where $\tilde{p}_\mu = (2/a)\sin(p_\mu a/2)$ is the lattice momentum. This is a specific element of $\mathbb{R}_B(k)$ at each discrete momentum p on the finite Brillouin zone. Type I.

For the Dirac propagator:
$S(x,y) = (D_{lat} + m)^{-1}_{xy}$

where D_{lat} is the lattice Dirac operator (Definition 9.3). Again: the inverse of a finite matrix. Type I.

For the gauge propagator (in a fixed gauge):
$D_{\mu\nu}(x,y)$ requires gauge-fixing – choosing a specific representative from each gauge orbit. On a finite lattice, gauge-fixing is a finite selection problem. The gauge-fixed propagator is

then the inverse of the gauge-fixed kinetic
operator — a finite matrix. Type I.

Definition 9.7 — Bounded Feynman diagram evaluation:

A Feynman diagram Γ with:
- L internal lines (propagators)
- V vertices (interaction points)
- E external lines (incoming/outgoing particles)
- $\ell = L - V + 1$ loops

evaluates to a finite number (the diagram amplitude) by the following procedure:

1. Assign a discrete lattice momentum to each internal line, subject to momentum conservation at each vertex. The independent loop momenta take values on the finite Brillouin zone ($N_x \times N_y \times N_z \times N_t$ discrete momenta).

2. For each internal line: insert the lattice propagator $\tilde{G}(p)$ or $\tilde{S}(p)$ at the assigned momentum.

3. For each vertex: insert the vertex factor (coupling constant × gamma matrices × colour factors, as appropriate for the interaction).

4. Sum over the ℓ independent loop momenta — each ranging over the finite Brillouin zone.

5. The result is a finite sum of finite products of elements of $\mathbb{C}_B(k^4)$. Type I.

The key difference from continuum Feynman rules: Step 4 is a finite sum, not an integral. There is no UV divergence — the lattice spacing a provides a natural cutoff. There is no IR divergence — the finite lattice volume provides a natural IR cutoff. No regularisation. No renormalisation at fixed lattice spacing. Type I.

For a one-loop diagram with N^4 lattice momenta: the sum has N^4 terms. For $N = 10$: $10^4 = 10000$ terms. Computable. For $N = 100$: 10^8 terms. Large but finite.

Note on renormalisation:
The finite Feynman rules at fixed lattice spacing a give finite, well-defined amplitudes. The question of renormalisation arises when asking: how do these amplitudes behave as $a \to 0$ (the continuum limit)?

This is a family-level question — it asks about the behaviour of the family of finite computations as the lattice refines. The answer involves the renormalisation group (§9.5), which describes how the coupling constants must be adjusted as a function of a to maintain physical predictions.

At each fixed a: the computation is finite and exact.
The continuum limit: is a Type IV family statement.

9.5 Bounded Renormalisation Group

Definition 9.8 — Bounded RG step:

A renormalisation group (RG) step is a map between theories at different scales.

On a finite lattice with spacing a:

1. Start with a lattice theory T_a defined by an action $S_a[U, \psi, \varphi]$ with coupling constants $g_a = (g_1, g_2, g_3, \lambda, y_t, \ldots)_a$.

2. "Block" the lattice: combine groups of b^d sites into single "block" sites, producing a coarser lattice with spacing $a' = ba$.

3. Define the blocked fields on the coarse lattice by averaging or summing the fine-lattice fields within each block.

4. Compute the effective action $S_{a'}[U', \psi', \varphi']$ for the blocked fields by integrating out the within-block fluctuations:

   $$\exp(-S_{a'}) = \Sigma_{\text{fine modes}} \exp(-S_a)$$

   This is a finite sum (the sum over within-block field configurations is finite because the lattice is finite). Type I.

5. Read off the new coupling constants $g_{a'}$ from $S_{a'}$.

The RG step is the map: $g_a \to g_{a'} = R_b(g_a)$.

This is a specific computable map between finite sets of coupling constants in $\mathbb{R}_B(k)$. Type I at each step.

A sequence of RG steps: $g_a \to g_{ba} \to g_{b^2a} \to \ldots$ is a finite sequence of such maps (bounded by the number of times the lattice can be blocked before reaching a single site). Type I.

9.5.1 Running couplings and beta functions

The discrete beta function at scale a:

$$\beta_i(g) \approx (g_i(a') - g_i(a)) / \ln(a'/a)$$

This is a finite difference quotient — a specific element of $\mathbb{R}_B(k)$ at each a. Type I.

The classical continuum beta functions of the Standard Model:

$$\beta_1 = (41/6)g_1^3/(16\pi^2) + \ldots \quad \text{(U(1): not asymptotically free)}$$
$$\beta_2 = -(19/6)g_2^3/(16\pi^2) + \ldots \quad \text{(SU(2): asymptotically free)}$$
$$\beta_3 = -7g_3^3/(16\pi^2) + \ldots \quad \text{(SU(3): asymptotically free)}$$

These are the one-loop coefficients, which are rational numbers times the cube of the coupling. In BST: each coefficient is an exact rational in $\mathbb{Q}_B(k^2)$, computable from the group theory of the gauge group and matter representations (Part XIII).

Whether the bounded discrete beta function converges to the continuum beta function as the lattice refines is a family-level (Type IV) question — the lattice analogue of the universality hypothesis in statistical mechanics.

9.5.2 Asymptotic freedom

Asymptotic freedom — the statement that the SU(3) coupling g_3 decreases at short distances (high energies) — is the reason perturbative QCD works at high energies.

In BST at fixed lattice spacing a: the coupling $g_3(a)$ is a specific element of $\mathbb{R}_B(k)$. The statement "g_3 decreases as a decreases" is a comparison between $g_3(a)$ and $g_3(a')$ for $a' < a$ — a decidable inequality in $\mathbb{R}_B(k)$ for each specific pair of lattice spacings. Type I.

The universal statement "$g_3(a) \to 0$ as $a \to 0$" is a family-level (Type IV) claim about the behaviour of the sequence $\{g_3(a_n)\}$ as the lattice refines. This is the content of asymptotic freedom, and it is not an internal BST theorem — it is a metatheoretic observation about the family.

9.6 Bounded Effective Field Theory

Definition 9.9 — Bounded EFT at fixed cutoff:

An effective field theory at cutoff scale Λ (equivalently, lattice spacing $a = 1/\Lambda$) is specified by:

1. A finite list of fields on the lattice (gauge, fermion, scalar — as in §§9.1–9.2).

2. A finite list of local operators $O_1, \ldots, O_M$ up to a bounded engineering dimension D_{max}. Each operator is a polynomial in the fields and their finite differences, evaluated at a single lattice site. The number of such operators is finite (bounded by the number of monomials of degree $\leq D_{max}$ in the lattice fields).

3. Coupling constants $c_1, \ldots, c_M \in \mathbb{R}_B(k)$, one per operator.

4. The EFT action:
   $S_{EFT} = \Sigma_x \Sigma_{i=1}^{M} c_i O_i(x)$

   A finite sum over sites and operators. Type I.

The EFT framework is a controlled truncation: by choosing D_{max}, one retains only the operators that are relevant or marginal at the scale Λ, discarding the irrelevant operators whose effects are suppressed by powers of (energy/Λ).

In BST: the truncation is explicit, the operator list is finite and enumerable, and the coupling constants are specific elements of $\mathbb{R}_B(k)$. There is no infinite tower of irrelevant operators — the tower terminates at D_{max}. Type I.

9.7 The Assembly Gap

This section identifies what has been built and what remains open for the full Standard Model in BST.

What is built (Tier 1):

- The gauge group $SU(3)_B \times SU(2)_B \times U(1)_B$ (Part V §5.3)
- The representation content of all SM fermions (§9.2)
- The Clifford algebra and Dirac spinors (§9.1)
- The lattice Dirac operator (§9.3, Tier 2 for doubling solutions)
- The Higgs field and Mexican-hat potential (§9.3)
- The classical Higgs mechanism: mass spectrum by finite spectral analysis of the quadratic fluctuation operator (§9.3.1)
- Lattice propagators as finite matrix inverses (§9.4)
- Individual Feynman diagrams as finite sums (§9.4)
- RG steps as finite blocking maps (§9.5)
- EFT at fixed cutoff as a finite action (§9.6)

What remains open (Tier 2–3):

```
- Full fermion doubling solution with chiral symmetry:
  The naive lattice Dirac operator doubles the fermion
  spectrum. Solutions exist (Wilson, staggered, overlap)
  but each involves tradeoffs. A clean bounded
  treatment with proved chiral properties is Tier 2.

- Yukawa couplings and fermion masses:
  The coupling of fermions to the Higgs field
  (ψ̄_L φ ψ_R + h.c.) is definable on the lattice
  but the full mass matrix with 3 generations,
  CKM mixing, and CP violation requires assembling
  the complete flavour structure. Tier 2.

- Anomaly cancellation:
  The SM is anomaly-free: the chiral gauge anomalies
  cancel between quarks and leptons. On a lattice,
  proving anomaly cancellation requires controlling
  the chiral properties of the lattice fermions
  (connected to the doubling problem). Tier 2–3.

- Electroweak symmetry breaking in the quantum theory:
  The classical Higgs mechanism (§9.3.1) is Tier 1.
  The full quantum treatment — showing that the
  symmetry breaking persists beyond the classical
  approximation on a finite lattice — is Tier 2.

- Continuum limit and renormalisability:
  Showing that the lattice SM has a well-defined
  continuum limit (a → 0 with physical predictions
  held fixed) is a family-level (Type IV) statement.
  This is the lattice analogue of the classical
  renormalisability proofs (t'Hooft-Veltman). Tier 3.

- Precision predictions:
  The electron anomalous magnetic moment g−2, computed
  to 10 significant figures in continuum QED, would
  require lattice computations at extremely fine
  lattice spacing. The computation is finite at each
  lattice spacing, but achieving the precision of
  the best continuum calculations is a practical
  (not foundational) challenge. Tier 2.
```

Plain language: The Standard Model's algebraic skeleton is fully present in BST: the gauge group, the matter representations, the Higgs field, the mass mechanism, the Feynman diagram framework, and the RG flow — all as finite objects on a finite lattice. The remaining open problems are dynamical and analytic: proving that the lattice theory has the right continuum limit, controlling chiral symmetry on the lattice, and computing precision predictions to match experimental accuracy. These are hard physics problems, not missing mathematical foundations.

The situation parallels the status of lattice QCD in mainstream physics: the lattice formulation is rigorous and well-defined; the computational challenge is extracting the physical predictions from the finite lattice theory. BST does not change this challenge — it makes explicit that the lattice formulation IS the fundamental mathematics, not an approximation to an infinite-dimensional continuum theory.

9.8 What Part IX Recovers

Result	Type	Section
Clifford algebra Cl(1,3)	I	9.1
Dirac spinors	I	9.1
Chiral projectors P_L, P_R	I	9.1
Lattice Dirac operator	I	9.1
Spinor Lorentz transformation	I/III	9.1
SM fermion representations	I	9.2
SM fermion field content	I	9.2
Higgs field and potential	I	9.3
Classical vev and mass spectrum	I	9.3.1
Goldstone bosons (algebraic)	I	9.3.1
Lattice propagators (matrix inverse)	I	9.4
Feynman diagrams (finite sums)	I	9.4
UV/IR finiteness of lattice diagrams	I	9.4
RG blocking step	I	9.5
Discrete beta function	I	9.5.1
One-loop beta coefficients	I	9.5.1
Asymptotic freedom (per lattice)	I	9.5.2
EFT at fixed cutoff	I	9.6
Continuum limit	IV	9.5, 9.7
Full SM anomaly cancellation	Open	9.7
Precision g-2 predictions	Open	9.7

What is NOT recovered: the Standard Model as a continuum quantum field theory on $\mathbb{R}^4$ (replaced by a lattice gauge theory on a finite complex). The path integral over infinite-dimensional field space (replaced by a finite sum over finite configuration space). The BRS/BRST cohomological formulation of gauge-fixing (replaced by Maschke projectors from Part V §5.7). The proof of renormalisability (a family-level statement). The hierarchy problem (a question about the sensitivity of the Higgs mass to the cutoff — which in BST is a question about the dependence of m_H on the lattice spacing a, a computable finite-difference quantity at each a).

What is gained: every Feynman diagram evaluates to a finite number with no UV or IR divergences. The lattice spacing provides a natural, physical cutoff — not an artificial regularisation. The Higgs mechanism reduces to finite spectral analysis. The RG flow is a computable finite map. The entire Standard Model gauge + matter + Higgs structure exists as explicit finite algebra inside BST. The foundational question “what is quantum field theory?” has an answer in BST: it is finite statistical mechanics on a finite lattice with gauge symmetry.

End of Part IX

Part X: Open Problems and Research Programmes

Maturity: Tier 3 throughout — this Part builds nothing new. It collects.

Volume imports: Everything above.

This Part gathers every open problem, every family-level claim, and every Tier 2–3 construction identified in Parts I–IX into a single honest accounting. It serves three purposes: (1) the reader can see at a glance what the companion volume has built and what it has not; (2) the open problems are stated with their dependencies, enabling focused research; and (3) the boundary between theorem and programme is drawn cleanly.

The Part is organised by physics domain, not by maturity tier, so that a researcher interested in a specific area can find all the relevant open questions in one place.

10.1 Continuum Recovery: The Central Open Question

Every Part of this volume constructs finite physics on a finite lattice or finite state space at fixed precision k and fixed lattice size. The central open question, across all domains, is the same:

Does the family of finite theories converge to the corresponding classical continuum theory as the lattice refines and the precision grows?

This is the content of Type IV recovery (AFB paper, §9.1): a family-level, metatheoretic correspondence that is visible only when looking across all members of the family {T_{k,K}} from outside any single member.

The question decomposes differently in each domain:

Domain	Continuum target	Volume Part	Status
Classical mechanics	Hamilton on $\mathbb{R}^{2n}$	I	Convergent (RK4 error $O(h^4)$)
Wave mechanics	Wave eq. on $\mathbb{R}^d$	II	Convergent (CFL stable)
Electromagnetism	Maxwell on $\mathbb{R}^4$	II	Convergent (lattice → continuum known)
Special relativity	Lorentz on $\mathbb{R}$	II	Exact algebraic structure
Thermodynamics	Infinite-volume	III	Phase transitions Type IV
Quantum mechanics	$L^2(\mathbb{R})$ Hilbert space	IV	Truncation → full as $N \to \infty$
Gauge theory	Yang-Mills on $\mathbb{R}^4$	V	Open (the mass-gap problem)
Simplicial topology	Smooth manifolds	VI	Open (reconstruction theorems)
Gravity	Einstein GR	VII	Open (five-theorem chain)
Condensed matter	Thermodynamic limit	VIII	Phase transitions Type IV
Standard Model	Continuum QFT	IX	Open (renormalisability)

For the first four domains, convergence is established by classical numerical analysis results (finite-difference convergence theorems). For thermodynamics and condensed matter, the "convergence" is to a limit with qualitatively new features (phase transitions, symmetry breaking) that are Type IV. For gauge theory, gravity, and the Standard Model, the convergence is the substance of major open research programmes.

10.2 Gauge Theory and the Mass Gap

Open Problem 10.1 — Yang-Mills mass gap (Clay Millennium Problem):

Statement: Does four-dimensional SU(N) lattice gauge theory have a uniform positive spectral gap $\gamma_0 > 0$ in the family limit (lattice spacing $a \to 0$ with physical predictions held fixed)?

Status: Open. This is the Clay Millennium Problem.

What BST provides (Part V):
- The spectral gap γ_{K,k} is exactly computable for each finite lattice Hamiltonian H_{K,k}. (Type I)
- Whether γ_{K,k} > 0 is decidable at each (K,k).
- The problem is decomposed into:
  (a) a decidable finite part (spectral gap at fixed lattice)
  (b) an open family-level part (uniform bound across refinement)

What is needed:
- A proof that $\gamma_{K,k} \geq \gamma_0 > 0$ for all sufficiently large K in a controlled refinement family, OR
- A constructive sequence of lattice computations establishing the gap with increasing confidence.

Dependencies: Part V §5.6. Connected to Part IX §9.5 (RG flow controls how the gap depends on lattice spacing).

Open Problem 10.2 — Confinement:

Statement: In SU(3) lattice gauge theory, does the Wilson loop expectation value ⟨W(R,T)⟩ exhibit area-law decay for large loops?

Status: Strong numerical evidence from lattice Monte Carlo; rigorous proof open.

What BST provides (Part V §5.3):
- Wilson loops are gauge-invariant observables, computable as finite sums.
- The area-law vs. perimeter-law scaling is testable on each finite lattice.

What is needed: A family-level proof that the area law persists as the lattice refines.

Open Problem 10.3 — Lattice chiral gauge theory:

Statement: Construct a lattice formulation of the Standard Model gauge theory that correctly implements chiral fermion representations without doublers and with exact gauge invariance.

Status: Partially open. The overlap/Ginsparg-Wilson approach provides a solution in principle, but a complete rigorous treatment in the BST setting is Tier 2.

What BST provides (Part IX §9.1):
- The lattice Dirac operator D_lat is a well-defined finite matrix.
- Wilson, staggered, and overlap modifications are all specific finite matrices.

What is needed: A clean proof that the overlap Dirac operator satisfies the Ginsparg-Wilson relation exactly or to controlled precision in $\mathbb{C}_B(k^4)$, and that the resulting lattice theory has the correct chiral anomaly structure.

Dependencies: Part IX §§9.1, 9.7.

10.3 Gravity

Open Problem 10.4 — Lorentzian reconstruction:
From Part VII §7.8, Open Problem 7.1.

Statement: Reconstruct a Lorentzian manifold (M, g) from a controlled family of bounded causal sets.

Status: Partial classical results (Malament, Bombelli-Meyer). Full reconstruction open.

Dependencies: Part VII §§7.5–7.6 (causal combinatorics).

Open Problem 10.5 — Curvature convergence:
From Part VII §7.8, Open Problem 7.2.

Statement: Prove that deficit angles (Regge) or BD action density (causal sets) converge to scalar/Ricci curvature in the refinement family.

Status: Partial results (Cheeger-Müller-Schrader in 2D; Benincasa-Dowker for BD).

Dependencies: Open Problem 10.4. Part VII §§7.2, 7.6.

Open Problem 10.6 — Action convergence:
From Part VII §7.8, Open Problem 7.3.

Statement: Regge action → Einstein-Hilbert action in the refinement limit.

Dependencies: Open Problem 10.5. Part VII §7.3.

Open Problem 10.7 — Matter-gravity coupling:

From Part VII §7.8, Open Problem 7.4.

Statement: Lattice gauge/matter sectors (Parts V, IX) coupled to discrete gravity induce a recovered stress-energy tensor $T_{\mu\nu}$.

Dependencies: Open Problems 10.5, 10.6. Parts V, IX.

Open Problem 10.8 — Einstein recovery:

From Part VII §7.8, Open Problem 7.5.

Statement: Coupled discrete field equations converge to $G_{\mu\nu} = 8\pi G\, T_{\mu\nu}$.

This is the culmination: 10.4 → 10.5 → 10.6 → 10.7 → 10.8.

Dependencies: All of 10.4–10.7.

Open Problem 10.9 — Holographic derivation:

From Part VII §7.9.

Statement: Is there a derivation — within BST or a natural extension — of a lower bound on n_M from the requirement that the Cayley-Dickson algebraic tower ($\mathbb{C}$, $\mathbb{H}$, $\mathbb{O}$) exists within the theory? Is that bound compatible with the Bekenstein-Hawking bound?

Status: Purely speculative. The numerical compatibility is observed; no theoretical derivation exists.

Dependencies: Part VII §7.9. AFB §8.8.4 (Cayley-Dickson resource costs).

10.4 Kähler Geometry and the Hodge Conjecture

Open Problem 10.10 — Approximate Kähler identities:

From Part VI §6.6.2, Open Lemma 1.

Statement: Prove the approximate Kähler commutator identities $[\Lambda_B, \partial_B] = -i\bar{\partial}_B^* + E_1$ with $\|E_1\| \le C/k$ for a suitable class of bounded complex simplicial manifolds.

Status: Open. The proof strategy is clear (Part VI §6.6); the geometric input is missing.

Dependencies: Part VI §6.4 (Laplacians), Part X of AFB (preliminary Kähler geometry).

Open Problem 10.11 — Spectral gap for Kähler Laplacian:

From Part VI §6.6.2, Open Lemma 2.

Statement: Prove that the spectral gap of the bounded Kähler Laplacian exceeds $2C/k$ for the relevant class of complexes.

Dependencies: Open Problem 10.10. AFB Part XI (Weyl/Davis-Kahan).

Open Problem 10.12 — Bounded Hodge conjecture:

Statement: For every bounded complex projective variety X over $\mathbb{C}_B(k^4)$, does every (p,p)-class in $H^{2p}_B(X)$ lie in the span of algebraic cycle classes?

Status: Formulated (AFB Part X §10.5). Each instance is a decidable finite linear algebra question. The universal statement is open.

Dependencies: Open Problems 10.10–10.11 (for the Kähler Hodge decomposition that defines the (p,p)-classes). Part VI §6.3 (cohomology).

10.5 The Standard Model

Open Problem 10.13 — Lattice chiral symmetry:

From Part IX §9.7.

Statement: Construct a lattice fermion formulation with exact chiral symmetry compatible with gauge invariance. Connected to Open Problem 10.3.

Open Problem 10.14 — Anomaly cancellation on the lattice:

From Part IX §9.7.

Statement: Prove that the SM anomalies cancel on the finite lattice – i.e. that the chiral gauge anomalies from quarks and leptons sum to zero in the bounded lattice formulation.

Dependencies: Open Problem 10.13. Part IX §9.2 (fermion representations).

Open Problem 10.15 — Yukawa structure and CKM matrix:

From Part IX §9.7.

Statement: Assemble the full 3-generation Yukawa coupling structure with CKM mixing matrix and CP violation on the finite lattice.

Status: The algebraic data (coupling constants, mixing angles) are finite parameters in $\mathbb{R}_B(k)$. The dynamical question is whether the lattice theory reproduces the observed mass hierarchy and mixing pattern.

Dependencies: Part IX §§9.2, 9.3 (Higgs mechanism, fermion representations).

Open Problem 10.16 — Electroweak quantum symmetry breaking:
From Part IX §9.3.1.

Statement: Prove that the Higgs mechanism persists beyond the classical approximation on a finite lattice — i.e. that the lattice partition function at large volume selects a broken-symmetry phase.

Status: Expected from lattice simulations but not rigorously proved in BST.

Dependencies: Part IX §9.3. Part III §3.5 (finite-size symmetry restoration).

Open Problem 10.17 — Continuum limit and renormalisability:
From Part IX §9.5, §9.7.

Statement: Prove that the lattice Standard Model has a well-defined continuum limit with physical predictions independent of the lattice spacing.

Status: Open. The classical proof of renormalisability (t'Hooft-Veltman) works in the continuum. The lattice analogue requires controlling the RG flow of all SM couplings simultaneously as the lattice refines.

Dependencies: Part IX §9.5 (RG flow). Part V §5.4 (lattice partition function).

10.6 Condensed Matter and Many-Body Theory

Open Problem 10.18 — Quantum phase transition universality:
From Part VIII §8.1.2.

Statement: Prove that the critical exponents of finite-lattice quantum phase transitions (e.g. the transverse-field Ising model) converge to universal values as the lattice size grows.

Status: Strong numerical evidence. Rigorous proof (in the lattice setting) open for most models.

Dependencies: Part VIII §8.1. Part IX §9.5 (RG flow applied to condensed matter).

Open Problem 10.19 — High-temperature superconductivity:
Statement: Determine whether the 2D Hubbard model at intermediate coupling (U/t ~ 4–8) has a superconducting ground state.

Status: One of the most important open problems in condensed matter physics. Not specific to BST — but BST provides the rigorous framework (finite Fock space, exact Hamiltonian, decidable spectral gap) within which the question is mathematically precise.

Dependencies: Part VIII §§8.3, 8.6.

10.7 Bounded Navier-Stokes

Open Problem 10.20 — Bounded Navier-Stokes regularity:

Statement: For the bounded Navier-Stokes equations on a finite grid (AFB paper, §9.3.3.5 machinery): does the finite solution remain bounded in a suitable norm uniformly as the grid refines?

This is the BST reformulation of the Clay Navier-Stokes Millennium Problem. As with the mass gap (Open Problem 10.1), the problem decomposes:

(a) At fixed grid: the finite solution exists and is computable (Type I).

(b) Across the family: does a uniform bound hold? This is the open family-level question.

Status: Open. The 3D case is unsolved classically.

What BST provides:
- Exact finite dynamics at each grid resolution (AFB §9.3.3.5)
- Finite energy inequalities at each resolution
- The reformulation separating the computable finite part from the open universal part

Dependencies: AFB Part IX §9.3.3.5 (bounded PDEs).

10.8 Bounded Number Theory

Open Problem 10.21 — Bounded Riemann Hypothesis programme:

Statement: Study the zero distribution of the bounded truncated zeta function $\zeta_B(s, k)$ over $\mathbb{C}_B(k^4)$. Do the zeros satisfy a bounded critical-line property?

Status: Research programme, not a theorem target. The classical RH is a family-level universal claim.

What BST provides:
- $\zeta_B(s, k)$ as a finite sum, computable in $\mathbb{C}_B(k^4)$
- Zeros as a finite set, findable by finite root search
- The proved finite-field analogue (Weil/Deligne)

Dependencies: AFB Part X (bounded complex analysis), AFB Part VIII ($\mathbb{C}_B(k^4)$).

Open Problem 10.22 — Bounded BSD programme:

Statement: Sharpen the "split" of the BSD conjecture: determine exactly which arithmetic content (rational points, Tate-Shafarevich groups at finite primes) is provable in BST, and whether bounded surrogates for the analytic content (L-function behaviour) can be constructed over $\mathbb{C}_B(k^4)$.

Status: Research programme.

Dependencies: AFB Part VIII (bounded number theory), AFB Part X (bounded complex analysis).

10.9 Research Infrastructure

Beyond the specific open problems, several pieces of research infrastructure would strengthen the entire programme:

10.9.1 Bounded simplicial topology as a standalone development

Part VI of this volume develops simplicial topology within the companion. For the AFB paper itself, a standalone internal treatment of finite simplicial complexes, homology, and cohomology — proved entirely within BST with no physics motivation — would strengthen the foundational core.

Status: The mathematical tools are all in the AFB paper (Parts VII, XI). The assembly into a standalone simplicial topology development is a writing task, not a research task. Tier 1.

10.9.2 A BST model library

A collection of explicit, small BST models (n_M = 10, 20, 50, 100) with:

- All axioms verified
- The bounded number chain constructed
- Specific physical computations performed (harmonic oscillator, Ising model, hydrogen atom)
- Error bounds computed and displayed

This would make the abstract framework concrete and provide test cases for all subsequent developments.

10.9.3 Computational implementation

A software implementation of BST arithmetic — $\mathbb{R}_B(k)$ and $\mathbb{C}_B(k^4)$ with exact error tracking — would enable:

- Automated verification of the paper's theorems
- Numerical experiments on bounded physical models
- Exploration of the parameter space (k, n_M)
- Concrete demonstration that BST physics is computable, not merely theoretically finite

This is an engineering project, not a mathematical
one, but it would significantly strengthen the
programme's credibility and utility.

10.10 Summary: The State of BST Physics

What is built (Tier 1):

Classical mechanics on finite phase spaces (Part I).
Wave mechanics on finite grids (Part II).
Special relativity as finite matrix algebra (Part II).
Electromagnetism in three formulations (Part II).
Statistical mechanics with finite partition functions (Part III).
Quantum mechanics on finite Hilbert spaces (Part IV).
Lattice gauge theory with SU(N)_B (Part V).
Finite simplicial topology and Hodge decomposition (Part VI).
Regge geometry and causal set combinatorics (Part VII).
Condensed matter models — Ising, Heisenberg, Hubbard (Part VIII).
Topological invariants — Berry phase, Chern number (Part VIII).
Standard Model algebraic skeleton — gauge group,
  representations, Clifford algebra, Higgs mechanism (Part IX).
Lattice Feynman rules and propagators (Part IX).

What is partially built (Tier 2):

Lattice Dirac operator with chirality control (Part IX §9.1).
Regge dynamics and discrete Einstein equations (Part VII §7.4).
Kähler-Hodge decomposition (Part VI §6.6).
Bounded RG flow and beta functions (Part IX §9.5).
BCS superconductivity on finite lattices (Part VIII §8.6).

What is a research programme (Tier 3):

GR emergence from discrete structures (Part VII §7.8).
Yang-Mills mass gap (Part V §5.6, Open Problem 10.1).
Full Standard Model assembly (Part IX §9.7).
Continuum limit of the lattice SM (Open Problem 10.17).
Bounded Navier-Stokes regularity (Open Problem 10.20).
Bounded Riemann Hypothesis programme (Open Problem 10.21).
Holographic derivation of n_M (Open Problem 10.9).

The critical path:

If the goal is a published mathematical physics volume:
→ Polish Parts I–V and VIII (all Tier 1, strongest cases).

If the goal is GR emergence:
→ Open Problems 10.4 → 10.5 → 10.6 → 10.7 → 10.8
   (the linear chain from Part VII §7.8).

If the goal is the Standard Model:
→ Open Problems 10.13 → 10.14 → 10.15 → 10.16 → 10.17
   (chiral lattice theory → anomalies → Yukawa → symmetry breaking →
continuum limit).

If the goal is a Millennium Problem:
→ Open Problem 10.1 (mass gap) is the most mature:
   the finite spectral machinery is in place, the

```
   question is precisely stated, and numerical evidence
   is strong.
→ Open Problem 10.20 (Navier-Stokes) is precisely
   reformulated but the underlying PDE problem is
   as hard in BST as classically.
```

10.11 Closing Statement

This companion volume has constructed bounded physics from the ground up — from classical mechanics through the Standard Model, from simple harmonic oscillators through lattice quantum gravity — using only the mathematical tools developed in Parts III–XIII of the AFB paper. Every construction is finite. Every partition function is a finite sum. Every spectrum is a finite set of computable eigenvalues. Every topological invariant is a computable integer.

The programme is not complete. The open problems of §§10.1–10.8 are genuine open problems — some of them among the deepest in mathematics and physics. What BST provides is not their solution but their precise formulation in a setting where the finite mathematical content is separated from the infinite idealisation, the computable part is separated from the universal claim, and the dependable core is separated from the research frontier.

The physics does not change. The foundations become honest.

End of Part X