creative commons
COMMONS DEED

저작자표시-비영리-변경금지 2.0 대한민국

이용자는 아래의 조건을 따르는 경우에 한하여 자유롭게

- 이 저작물을 복제, 배포, 전송, 전시, 공연 및 방송할 수 있습니다.

다음과 같은 조건을 따라야 합니다:



저작자표시. 귀하는 원저작자를 표시하여야 합니다.



비영리. 귀하는 이 저작물을 영리 목적으로 이용할 수 없습니다.



변경금지. 귀하는 이 저작물을 개작, 변형 또는 가공할 수 없습니다.

- 귀하는, 이 저작물의 재이용이나 배포의 경우, 이 저작물에 적용된 이용허락조건을 명확하게 나타내어야 합니다.
- 저작권자로부터 별도의 허가를 받으면 이러한 조건들은 적용되지 않습니다.

저작권법에 따른 이용자의 권리는 위의 내용에 의하여 영향을 받지 않습니다.

이것은 이용허락규약(Legal Code)을 이해하기 쉽게 요약한 것입니다.

Disclaimer

dCollection

Master's Thesis of Science

AXIOM: Algebraic eXtension by Index–Ordered Multiplication

AXIOM: 인덱스 정렬 곱셈에 의한 대수적 확장

February 2026

College of Natural Sciences

Seoul National University

Interdisciplinary Program in Computational Science and Technology

Park Junhu

AXIOM: Algebraic eXtension by Index–Ordered Multiplication

AXIOM: 인덱스 정렬 곱셈에 의한 대수적 확장

Professor Myungjoo Kang

Submitting a Master's Thesis

November 2025

College of Natural Sciences
Seoul National University
Interdisciplinary Program in Computational Science and Technology

Park Junhu

Confirming the Master's Thesis written by
Park Junhu
January 2026

Chair	Woong Kook	(Seal)
Vice Chair	Myungjoo Kang	(Seal)
Examiner	Youngjoon Hong	(Seal)

Abstract

AXIOM: Algebraic eXtension by Index–Ordered Multiplication

Park Junhu
College of Natural Sciences
Seoul National University

Classical realification rewrites complex operators as real matrices in a dimension-dependent and asymmetric way: complex structure is treated as something to be embedded into a larger real space, and the construction does not extend uniformly beyond $\mathbb{C}$ without losing algebraic structure or inflating the data layout. This thesis instead introduces *AXIOM (Algebraic eXtension by Index–Ordered Multiplication)*, an axiomatic framework in which operators over complex, quaternionic, octonionic, general Cayley–Dickson, and arbitrary tensor algebras are *represented* within a single associative, rotation-based real matrix model that preserves multiplication, the $*$-structure, spectra, invertibility, and the effective memory footprint at the representation level.

Within AXIOM the thesis defines a universal symmetric block real rep-

resentation

$$\Phi(A) = \begin{bmatrix} H & -K \\ K & H \end{bmatrix}, \qquad A = H + iK,$$

called the *Conjugation–Real Embedding (CRE)*. The map Φ is proved to be a $*$-isometric real representation of any AXIOM algebra: it gives a one-to-one correspondence between algebraic operators and symmetric block real matrices, with no hierarchy of "real host" and "complex guest". At the implementation level, Φ does *not* introduce extra storage: the symmetric block view is obtained by reindexing the same real scalars that already back complex (or Cayley–Dickson) data, so both memory usage and communication volume remain identical to the original complex layout while all factorisations run purely in real arithmetic.

The main contributions are three structural theorems:

1. **Tensor-Sorted Product Theorem.** Every 2^n-dimensional Cayley–Dickson algebra and every AXIOM tensor algebra admits a table-free, sign-deterministic multiplication rule obtained purely by tensor index sorting, eliminating all dimension-specific multiplication tables.

2. **Arithmetic Local Hurwitz (ALH) Theorem.** Each generator e_i spans a local subalgebra isomorphic to $\mathbb{C}$, so complex-like inversion and norm composition hold slice-wise in all dimensions. This separates the classical 1–2–4–8 Hurwitz obstruction at the global level from a universally available local complex calculus.

3. **Symmetric Block Representation and CRE Inversion Theorem.**

Any operator over an AXIOM algebra admits a pivot-free block-Cholesky (or QR/SVD-type) factorisation after symmetric lifting, with $\Phi(AB) = \Phi(A)\Phi(B)$, $\Phi(A^{-1}) = \Phi(A)^{-1}$, and dimension-independent numerical stability. Since the underlying data layout is unchanged, this yields a single real kernel that replaces separate complex, quaternionic, and octonionic kernels without any memory penalty.

Beyond this structural unification, the AXIOM–CRE framework addresses several obstacles that have remained unresolved in the classical literature: the incompatibility of quaternionic and octonionic multiplication with Einstein summation, the absence of a dimension-independent operator calculus beyond $\mathbb{C}$, and the lack of a coherent extension of Cayley–Dickson algebras into a single associative computational model. AXIOM provides a fully consistent resolution of these issues at the representation level, yielding one rotation-based real model valid for all 2^n-dimensional and tensor algebras without altering the underlying algebraic identities.

Crucially, this theoretical unification produces measurable computational benefits. By expressing all operators through symmetric block rotations, Cholesky-type factorisations achieve a natural $4\times$ reduction in the leading arithmetic constant, and experiments on standard numerical libraries confirm that simply adding the AXIOM–CRE reindexing layer yields 3–$4\times$ runtime acceleration in dense regimes without modifying any underlying kernels, while preserving spectra and backward error to machine precision and keeping the memory footprint and communication cost identical to the original complex layout.

Consequently, the thesis does not merely embed complex operators into real matrices; it closes the structural gap between real and complex operator theory by exhibiting both as different index orderings of the same symmetric block rotation model, gives a dimension-independent operator calculus for the full Cayley–Dickson and tensor hierarchy, and provides a unified, practically superior computational foundation for future algorithms over high-dimensional algebras.

Keywords : AXIOM, index–ordered multiplication, Cayley–Dickson algebras, tensor and Clifford algebras, symmetric block real representation, numerical linear algebra kernels

Student Number : 2023-20862

Contents

Abstract . . . . . **i**

Contents . . . . . **v**

List of Figures . . . . . **viii**

List of Tables . . . . . **ix**

List of Algorithms . . . . . **x**

1 Introduction 1

1.1 Motivation and Problem Statement 1

1.2 The AXIOM Framework 3

1.3 CRE: Conjugation–Real Embedding 3

1.4 Consequences and Contributions 4

2 AXIOM–CRE Theoretical Foundations and Structural Theorems 5

2.1 Tensor-Sorted AXIOM Algebras and the ALH Condition 6

2.2 Symmetric Lift: A Universal Real Representation . . 10

2.3 Spectral and Norm Equivalence 12

2.4 Universal Positive Definiteness: HPD = SPD under CRE 14

3 Structural and Topological Theorems 17

3.1 Global CRE Inversion via an SPD Core 17
3.2 Geometric Interpretation of Noncommutativity . . . 20
3.3 Global Summary 21
4 Structural Ingredients of CRE 22
4.1 Setup and notation 22
4.2 CRE symmetric lift (imaginary embedding) 23
4.3 Tensor-sorted AXIOM algebra on 2^n and infinite dimensions 25
4.4 Arithmetic Local Hurwitz: local complex slices . . . 28
4.5 Summary 29
5 The AXIOM–CRE Inversion Theorem 30
5.1 Statement of the main theorem 30
5.2 Proof via structural ingredients 33
5.3 Interpretation 36
6 Analytical and Cost Model Framework 37
6.1 Operator-algebraic foundations 37
6.2 Analytic–computational integration 40
7 Hermitian CRE Solvers: Algorithms and Experiments 44
7.1 HPD to SPD: Analytic Backbone of the CRE Solve . 45
7.2 Rotational / Clifford View: Information as Moving Points 47
7.3 Schur Complements, Information Transport, and Cost 48
7.4 Algorithmic Realisation of CRE Solves 50
7.5 Hermitian Experiment Protocol: HPD Benchmarks . 51

8 CRE Solvers for General Systems: Sym + Skew → SPD Block Lift 54
8.1 Sym + Skew Decomposition and Hermitian Lift . . . 55
8.2 CRE Block Embedding and SPD Structure 57
8.3 Cholesky Solve and Complex Inverse Recovery . . . 60
8.4 Sym–Skew Benchmarks: Non-Hermitian Experiments 61
9 Numerical Stability, Error Bounds, and Reproducibility . . . 65
9.1 Backward-Error Analysis and Perturbation Bounds . 65
9.2 Statistical Reproducibility and Monte Carlo Averaging 68
9.3 Minimal Reproducibility Protocol 70
9.4 Summary of Stability Guarantees 71
10 Conclusion 72
Appendix 77
A Axiomatic Closure of Tensor-Sorted ALH Algebras 77
A.1 Algebraic Preliminaries 77
A.2 RLH Seminorms 78
A.3 Universal RLH Completion 78
A.4 Geometric Interpretation 79
A.5 Summary of Appendix A 79
B Formal Completions of the Conjugation–Real Embedding . . 80
B.1 Polynomial Closure and Jordan Structure 80
B.2 Norm-Continuous Extension to Bounded Operators . 81
B.3 Uniqueness of the CRE Embedding 82
B.4 Unitary–Orthogonal Correspondence 83
B.5 Summary of Appendix B 83

C Geometric and Noncommutative Closure 83
C.1 Commutativity as Plane Alignment 84
C.2 Cayley–Dickson Extension as Rotational Misalignment 85
C.3 Associative Matrix Realization in an Inductive Real Limit 87
C.4 Invertibility via Quaternionic Real Embeddings . . . 88
C.5 Summary of Geometric and Noncommutative Closure 90
D Numerical and Computational Closure 91
D.1 FLOP Model and Asymptotic Reduction 91
D.2 Memory Locality and Information Layout 92
D.3 Norms, Condition Numbers, and Backward Error . . 93
D.4 HPD–SPD Equivalence and Solver Closure 95
D.5 Summary of Numerical and Computational Closure . 96
E Implementation Summary 97

Bibliography **105**

Abstract in Korean . . . . . **i**

Acknowledgments . . . . . **v**

List of Figures

Figure 1. Performance and variance ratio plots for 32-bit precision. 52

Figure 2. Performance and variance ratio plots for 64-bit precision. 53

Figure 3. Performance and variance ratio plots for 128-bit precision. 53

Figure 4. Runtime (left, log scale) and normalized residuals (right) for the Sym–Skew CRE inversion benchmarks. 64

List of Tables

Table 1. Runtime comparison (precision 32-bit, seed 12 excluded). 53

Table 2. Variance comparison (precision 32-bit). 54

Table 3. Runtime comparison (precision 64-bit, seed 12 excluded). 55

Table 4. Variance comparison (precision 64-bit). 56

Table 5. Runtime comparison (precision 128-bit, all seeds valid). . 57

Table 6. Variance comparison (precision 128-bit). 58

Table 7. Average runtime (seconds) for $A = \mathrm{Sym}_1 + i\,\mathrm{Skew}_2$ inversion via real block Torch solvers and complex SciPy baselines. 63

Table 8. Residual $\|AX - I\|_F / \|A\|_F$ for $A = \mathrm{Sym}_1 + i\,\mathrm{Skew}_2$ inversion using the same solvers. 64

List of Algorithms

1 Cholesky factorisation via conjugation–real embedding . . . 50

2 Inverse of an HPD matrix via orthogonal CRE 50

3 CRE Block Inversion for $G \in \mathbb{R}^{n \times n}$ 61

1 Introduction

1.1 Motivation and Problem Statement

Hermitian and skew-Hermitian operators underpin a wide range of problems in mathematical physics, functional analysis, and numerical computation [1, 2, 3, 4]. In quantum mechanics the former describe observables with real spectra, whereas the latter generate unitary dynamics. This algebraic pairing links measurement and rotation through a common spectral theory [5, 6].

When these operators are realised on hardware tuned for real arithmetic, the standard representation in $\mathbb{C}^{n\times n}$ becomes wasteful. GPUs and dense BLAS/LAPACK kernels run real computations faster than complex ones [7, 8], yet the classical realification

$$A = X + iY \;\mapsto\; \begin{bmatrix} X & -Y \\ Y & X \end{bmatrix}$$

always doubles dimension and storage. Moreover, this construction has no analogue that preserves multiplication, adjoints, and spectra for quaternions, octonions, or general Cayley–Dickson algebras, so the computational split between real and complex domains persists.

Historical Limitations and Structural Obstacles

Several structural obstacles recur throughout the literature, yet no single framework removes them all.

(1) Dimension-dependent realification. The classical map is asymmetric, enforces a strict $2\times$ memory cost, and does not extend beyond $\mathbb{C}$ while preserving adjoints and spectra.

(2) Poor scaling of Clifford and Cayley–Dickson constructions. Each new dimension introduces a fresh multiplication table; octonions already lose associativity and higher levels no longer admit a practical operator calculus.

(3) Breakdown of global Hurwitz composition. Hurwitz's theorem limits its global norm composition to dimensions $1, 2, 4, 8$. Beyond those sizes the standard geometric interpretation collapses, and there is no clean separation between local slice behaviour and global failure.

(4) Lack of a unified operator model. Current numerical practice relies on distinct kernels for real, complex, quaternionic, and tensor-valued systems. This multiplies code paths and fragments memory layouts even when the underlying structures are closely related.

(5) Redundant numerical complexity. Despite the shared rotational structure across these algebras, there is no single real-arithmetic core that exposes it, so high-dimensional or tensor-valued computations remain expensive.

Taken together, these issues indicate the need for a symmetric, dimension-independent real model that respects conjugation, preserves spectra, and applies uniformly to all imaginary directions. The remainder of the thesis develops such a framework.

1.2 The AXIOM Framework

Let

$$\text{AXIOM} = \text{Algebraic eXtension by Index–Ordered Multiplication.}$$

This tensor-sorted $*$-algebra realises complex, quaternionic, octonionic, and general 2^n-dimensional Cayley–Dickson algebras, together with arbitrary tensor algebras, as slices of a single ambient structure. Its defining principles are:

(A1) **Index-ordered multiplication** $e_I e_J = \sigma(I, J) e_{S(I,J)}$;

(A2) **Arithmetic Local Hurwitz (ALH)**: each generator produces a local complex slice and compatible pairs produce quaternionic slices;

(A3) $*$**-structure and norm compatibility** across all slices.

These axioms isolate the local slice geometry from global nonassociativity and allow finite or infinite tensor systems to coexist inside one algebra.

1.3 CRE: Conjugation–Real Embedding

On top of AXIOM we place the block real representation

$$\Phi_{\text{conj}}(A) := \begin{bmatrix} \Re A & -\Im A \\ \Im A & \Re A \end{bmatrix}, \qquad A \in M_n(\mathbb{C}) \text{ or an AXIOM operator,}$$

called the *Conjugation–Real Embedding (CRE)*. The map Φ_{conj} is a real-linear $*$-embedding: it preserves products, adjoints, operator norms, and

spectra. The 2×2 block structure reorders the real scalars already stored in the coefficient algebra, so the effective memory usage and communication volume do not change.

Geometrically, the imaginary unit acts as a fixed $\pi/2$ rotation on $\mathbb{R}^2$, so each coefficient corresponds to an $SO(2)$ action on a real information token. All non-real structure is thus encoded as orthogonal motions of two-dimensional planes inside $\mathbb{R}^{2n}$. Algorithms over complex or tensor-valued operators can therefore be run entirely in real arithmetic on the CRE blocks, with standard SPD and orthogonal kernels, while maintaining the original spectral and backward-error properties.

1.4 Consequences and Contributions

The AXIOM–CRE model addresses the obstacles above:

1. it extends classical realification to all tensor algebras with no extra storage cost;

2. multiplication tables are replaced by index ordering;

3. Hurwitz composition becomes a local property that holds in every dimension;

4. real, complex, quaternionic, and tensor operators share one algebraic model;

5. numerical kernels reduce to a single real SPD framework, yielding the intrinsic $4\times$ arithmetic saving and the observed 3–$4\times$ runtime gains.

In short, the AXIOM–CRE construction shows that the classical operator families are different index orderings of one symmetric rotation model and it provides a uniform computational setting for high-dimensional and tensor-valued algebras.

2 AXIOM–CRE Theoretical Foundations and Structural Theorems

The AXIOM–CRE framework is built on two complementary layers. The AXIOM layer encodes imaginary directions as generators with tensor-sorted multiplication, while the CRE layer realises every resulting operator as a real symmetric block matrix that preserves multiplication, adjoints, spectra, and norms without changing storage requirements. The goal of this section is to show that all operators generated from arbitrary imaginary directions (complex, quaternionic, octonionic, Cayley–Dickson, or tensor-valued) admit this uniform real representation and hence reduce to a single SPD model when they are positive-definite.

Four ingredients are required: tensor-sorted AXIOM algebras and the Arithmetic Local Hurwitz (ALH) condition, the symmetric lift Φ and its algebraic properties, the spectral behaviour of the conjugation-normalised lift Φ_{conj}, and the correspondence between HPD operators and real SPD matrices under CRE. Once multiplication is governed by index ordering and sign counts, all imaginary directions become compatible and every operator in the algebra behaves like a real symmetric rotation after CRE lifting; the traditional distinctions between real, complex, quaternionic, octonionic, and

tensor-valued operators reduce to coordinate choices within the same block geometry.

2.1 Tensor-Sorted AXIOM Algebras and the ALH Condition

Let I be a finite or countable index set. An *AXIOM algebra* is generated by symbols $\{e_i : i \in I\}$ and 1, subject to the tensor-sorted relations

(T1) $e_i^2 = -1$ for all $i \in I$;

(T2) $e_i e_j = -e_j e_i$ for all $i \neq j$;

(T3) every product is reduced by sorting indices and tracking swap parity.

For every strictly ordered finite subset $J = \{j_1 < \cdots < j_m\} \subset I$, define

$$e_J := e_{j_1} \cdots e_{j_m}, \qquad e_\emptyset := 1.$$

Definition 2.1 (Tensor-sorted AXIOM algebra)**.** *A* tensor-sorted AXIOM algebra *is a real algebra $\mathcal{A}$ generated by $\{e_i\}_{i \in I}$ and* 1 *satisfying* (T1)–(T3) *and complete under a submultiplicative operator norm.*

The first structural result shows that a single combinatorial rule (index sorting with signs) replaces all dimension-specific multiplication tables.

Theorem 2.2 (Tensor-sorted basis and unique multiplication)**.** *Let $\mathcal{A}$ be a tensor-sorted AXIOM algebra. Then:*

(1) The set

$$\mathcal{B} := \{e_J : J \subset I,\ |J| < \infty,\ J \textit{ strictly ordered}\}$$

is a real vector space basis of $\mathcal{A}$. In particular, if $|I| = n < \infty$, then $\dim_{\mathbb{R}} \mathcal{A} = 2^n$.

(2) Every monomial $e_{i_1} \cdots e_{i_k}$ reduces uniquely to $\sigma(J)e_J$ for some finite ordered $J \subset I$ and $\sigma(J) \in \{\pm 1\}$, where the sign is determined solely by the parity of swaps and the cancellations $e_i^2 = -1$.

Proof. Existence of a representation in the span of $\{e_J\}$ follows by repeated use of (T1) and (T2): duplicate indices can be eliminated using $e_i^2 = -1$, and the remaining factors can be sorted at the cost of sign changes. To prove uniqueness, consider the free real vector space with basis $\mathcal{B}$ and the induced multiplication defined by the reduction procedure. If two distinct reduced expressions of the same monomial existed, their difference would give a non-trivial linear dependence among basis elements, contradicting minimality. Formally, one may order monomials lexicographically and observe that the reduction is triangular with respect to this order. □

Commentary. Classical Cayley–Dickson doubling introduces a new multiplication table at each stage (complex, quaternionic, octonionic, and so on), so the structure rapidly becomes unwieldy. Theorem 2.2 condenses all of these tables into a single instruction:

sort the indices, count the swaps, keep the resulting sign.

Every 2^n-dimensional Cayley–Dickson algebra, and even infinite tensor levels, obeys this combinatorial rule within AXIOM.

Arithmetic Local Hurwitz (ALH) Condition. Hurwitz's theorem limits its global norm composition to dimensions $1, 2, 4, 8$. AXIOM replaces this global requirement by a slice-wise condition that applies in every dimension and for arbitrary tensor combinations.

Definition 2.3 (Arithmetic Local Hurwitz)**.** *Let $(\mathcal{A}, \|\cdot\|)$ be a tensor-sorted AXIOM algebra. We say that $\mathcal{A}$ satisfies the* Arithmetic Local Hurwitz (ALH) condition *if:*

(1) For every generator e_a, the plane $\mathbb{C}_a := \mathrm{span}_{\mathbb{R}}\{1, e_a\}$ is a real subalgebra isomorphic to $\mathbb{C}$ with respect to multiplication and norm.

(2) For any finite compatible set $\{e_{i_1}, \ldots, e_{i_k}\}$, the slice they generate behaves like a finite-dimensional Cayley–Dickson algebra with norm composition

$$N(xy) = N(x)N(y), \qquad N(xyz) = N(x)N(y)N(z)$$

whenever no new imaginary directions are introduced.

(3) Noncommutativity and nonassociativity arise only from interactions across incompatible slices; within a fixed slice, the structure is associative (complex) or alternative (quaternionic/octonionic) with the usual norm identities.

Theorem 2.4 (Local complex and quaternionic slices)**.** *Let $\mathcal{A}$ be a tensor-sorted AXIOM algebra satisfying ALH. Then:*

(1) For each generator e_a, the slice $\mathbb{C}_a := \mathrm{span}_{\mathbb{R}}\{1, e_a\}$ is a real subalgebra isomorphic to $\mathbb{C}$, with multiplicative norm $N(xy) = N(x)N(y)$.

(2) For suitable pairs $\{e_a, e_b\}$, the span of $\{1, e_a, e_b, e_{ab}\}$ is isomorphic to a quaternionic slice $\mathbb{H}_{ab}$, on which $N(xy) = N(x)N(y)$ for all x, y in that slice.

(3) Octonionic and higher Cayley–Dickson-type behaviour arises by concatenating such slices. Norm composition identities of the form $N(xyz) = N(x)N(y)N(z)$ hold along any compatible slice chain.

Proof. For (1), the relation $e_a^2 = -1$ and bilinearity give

$$(u + ve_a)(s + te_a) = (us - vt) + (ut + vs)e_a,$$

which is the standard complex multiplication rule in the basis $\{1, e_a\}$. The norm identity $N(xy) = N(x)N(y)$ is part of the ALH assumption on $\mathbb{C}_a$.

For (2), the ALH condition asserts that whenever two generators are compatible, the four-dimensional span behaves like a quaternionic slice with the usual multiplication and norm composition; the details reduce to the classical identification once we restrict to that finite set of generators.

For (3), larger slices are built from chains of compatible generators. By ALH, each finite chain reproduces the corresponding Cayley–Dickson relations with norm composition along the chain, while nonassociativity appears only when incompatible chains are mixed. □

Commentary. Theorems 2.2 and 2.4 show that Einstein-style tensor rules, the Fano plane, and Hurwitz composition describe the same AXIOM algebra from different viewpoints. Local behaviour matches complex or quaternionic arithmetic, and global departures from commutativity are determined solely by the interaction between incompatible slices.

2.2 Symmetric Lift: A Universal Real Representation

The CRE layer realises AXIOM elements as real symmetric block operators.

Fix an imaginary direction $i \in \mathcal{A}$ with $i^2 = -1$. By construction of the slices, every element $a \in \mathcal{A}$ admits a decomposition

$$a = h + ik, \qquad h, k \in \mathcal{A}_{\mathbb{R}},$$

where $\mathcal{A}_{\mathbb{R}}$ denotes the underlying real subalgebra.

Definition 2.5 (CRE symmetric lift)**.** *The* CRE symmetric lift *of a is the* 2×2 *block operator*

$$\Phi(a) := \begin{bmatrix} h & -k \\ k & h \end{bmatrix} \in M_2(\mathcal{A}_{\mathbb{R}}).$$

When $\mathcal{A}$ is realised as matrices over $\mathbb{R}$ or $\mathbb{C}$, the same formula defines a block map

$$\Phi : M_n(\mathcal{A}) \to M_{2n}(\mathcal{A}_{\mathbb{R}}), \qquad A = H + iK \mapsto \begin{bmatrix} H & -K \\ K & H \end{bmatrix}.$$

Theorem 2.6 (Algebraic properties of Φ)**.** *Let $\mathcal{A}$ be a tensor-sorted AXIOM algebra and fix $i^2 = -1$. Then, for all $a, b \in \mathcal{A}$:*

(1) Φ is real-linear: $\Phi(\alpha a + \beta b) = \alpha\Phi(a) + \beta\Phi(b)$ for $\alpha, \beta \in \mathbb{R}$.

(2) Φ is multiplicative: $\Phi(ab) = \Phi(a)\Phi(b)$.

(3) Φ is injective and preserves invertibility: $\Phi(a) = 0$ implies $a = 0$, and a invertible implies $\Phi(a)$ invertible with $\Phi(a^{-1}) = \Phi(a)^{-1}$.

Proof. Write $a = h_1 + ik_1$, $b = h_2 + ik_2$ with $h_j, k_j \in \mathcal{A}_{\mathbb{R}}$. Then

$$ab = (h_1h_2 - k_1k_2) + i(h_1k_2 + k_1h_2),$$

using $i^2 = -1$ and the fact that i commutes with real coefficients in the chosen slice. On the other hand,

$$\begin{aligned}\Phi(a)\Phi(b) &= \begin{bmatrix} h_1 & -k_1 \\ k_1 & h_1 \end{bmatrix} \begin{bmatrix} h_2 & -k_2 \\ k_2 & h_2 \end{bmatrix} \\ &= \begin{bmatrix} h_1h_2 - k_1k_2 & -(h_1k_2 + k_1h_2) \\ h_1k_2 + k_1h_2 & h_1h_2 - k_1k_2 \end{bmatrix} = \Phi(ab).\end{aligned}$$

which proves multiplicativity. Real-linearity is immediate from the block definition.

If $\Phi(a) = 0$, then the four blocks vanish and we obtain $h = k = 0$, hence $a = 0$. If a is invertible, multiplicativity gives

$$\Phi(a)\Phi(a^{-1}) = \Phi(aa^{-1}) = \Phi(1) = I_2,$$

so $\Phi(a)$ is invertible with inverse $\Phi(a^{-1})$. $\square$

Relation to classical realification. For $\mathcal{A} = \mathbb{C}$ the map $A \mapsto \Phi(A)$ is the usual realification

$$A = X + iY \mapsto \begin{bmatrix} X & -Y \\ Y & X \end{bmatrix}.$$

CRE extends this formula to every AXIOM algebra, accommodates tensor-valued blocks without relying on a Hermitian–skew split, and reindexes the existing real scalars rather than duplicating storage.

2.3 Spectral and Norm Equivalence

To connect CRE with spectral theory, we introduce a conjugation-normalised variant. Let

$$J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \qquad S = I_n \otimes J,$$

and define

$$\Phi_{\mathrm{conj}}(A) := S\Phi(A)S^{\mathsf{T}}.$$

In the classical complex case this yields

$$\Phi_{\mathrm{conj}}(A) = \begin{bmatrix} \Re A & \Im A \\ -\Im A & \Re A \end{bmatrix},$$

and the same identities hold slice-wise in any AXIOM algebra satisfying ALH.

Theorem 2.7 (Spectral and norm equivalence)**.** *Let A be a bounded operator acting on a finite-dimensional complex (or slice-complex) Hilbert space inside an AXIOM algebra. Then:*

(1) $\Phi_{\mathrm{conj}}(A)$ *is unitarily similar to* $\mathrm{diag}(A, \overline{A})$.

(2) The spectra satisfy $\sigma(\Phi_{\mathrm{conj}}(A)) = \sigma(A) \cup \overline{\sigma(A)}$.

(3) The operator norms coincide: $\|\Phi_{\mathrm{conj}}(A)\|_2 = \|A\|_2$.

Proof sketch. In each local complex slice $\mathbb{C}_a$, CRE reduces to the standard block embedding for complex matrices. There one checks directly that

$$U^* \Phi_{\mathrm{conj}}(A) U = \begin{bmatrix} A & 0 \\ 0 & \overline{A} \end{bmatrix}$$

for a suitable unitary U, and

$$\det(\lambda I_{2n} - \Phi_{\mathrm{conj}}(A)) = |\det(\lambda I_n - A)|^2.$$

The singular values of $\Phi_{\mathrm{conj}}(A)$ coincide with those of A by the same argument. The ALH condition and tensor-sorted structure guarantee that assembling these slices does not change spectral radii or norms. □

Commentary. Theorem 2.7 confirms that CRE does not alter spectra, singular values, norms, or condition numbers; only the coordinates change between real, complex, or tensor descriptions.

2.4 Universal Positive Definiteness: HPD = SPD under CRE

CRE identifies every Hermitian positive-definite operator with a real SPD matrix:

HPD operator in an AXIOM algebra $\longleftrightarrow$ real SPD operator via Φ_{conj}

Definition 2.8 (HPD operator)**.** *An operator A is* Hermitian positive-definite (HPD) *if $A^* = A$ and $\langle Ax, x\rangle > 0$ for all $x \neq 0$, where * denotes the adjoint with respect to the slice-complex inner product under ALH. This includes real SPD, complex HPD, quaternionic/octonionic HPD, and tensor-valued HPD slices.*

Theorem 2.9 (HPD–SPD equivalence)**.** *Let A be HPD in an AXIOM algebra satisfying ALH. Then $\Phi_{\mathrm{conj}}(A)$ is a real symmetric positive-definite matrix. Moreover:*

(1) $\Phi_{\mathrm{conj}}(A)$ is symmetric;

(2) $\langle \Phi_{\mathrm{conj}}(A)x, x\rangle_{\mathbb{R}} = \Re\langle Ax, x\rangle > 0$ for all $x \neq 0$;

(3) all eigenvalues of $\Phi_{\mathrm{conj}}(A)$ are real, strictly positive, and coincide (with multiplicities) with those of A.

Proof sketch. Within a local complex slice, the CRE embedding is unitarily equivalent to $\mathrm{diag}(A, \overline{A})$. If A is HPD, then A and $\overline{A}$ share the same positive eigenvalues, and the block matrix is real symmetric with strictly positive spectrum. The real inner product on the CRE space is the real part of the

slice-complex inner product, giving

$$\langle \Phi_{\mathrm{conj}}(A)x, x\rangle_{\mathbb{R}} = \Re\langle Ax, x\rangle > 0.$$

Tensor-sorted ALH ensures that this argument patches consistently across all slices. □

Consequences. Every HPD operator from any AXIOM algebra becomes a real SPD matrix under CRE. Therefore the entire classical theory of SPD matrices applies uniformly: Cholesky factorisation, Schur complements, interlacing, Gershgorin-type bounds, matrix monotonicity, and stability theory.

Corollary 2.10 (Universal Schur structure)**.** *If*

$$\Phi_{\mathrm{conj}}(A) = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix}$$

with A HPD, then the Schur complement

$$S_{22|1} := S_{22} - S_{21}S_{11}^{-1}S_{12}$$

is real SPD.

Corollary 2.11 (Collapse of all positive-definite theories)**.** *Under AXIOM–*

CRE, we have a canonical identification

$$\begin{aligned} \textit{real SPD} &\equiv \textit{complex HPD} \\ &\equiv \textit{quaternionic HPD} \\ &\equiv \textit{octonionic/slice HPD} \\ &\equiv \textit{tensor HPD}. \end{aligned}$$

All such operators share a single real SPD representative with identical spectral and norm properties.

Convention 1 (SPD identification)**.** *From now on, any HPD operator A is tacitly identified with its SPD CRE lift $\Phi_{\mathrm{conj}}(A)$. All subsequent analysis is carried out on this SPD representative, with the understanding that results immediately transfer back to A via $\Phi_{\mathrm{conj}}^{-1}$.*

Summary. The results above show that:

- tensor-sorted multiplication in AXIOM gives a universal 2^n-basis, eliminating dimension-specific multiplication tables;
- the ALH condition guarantees complex/quaternionic slices at all dimensions, with norm composition along compatible chains;
- the CRE symmetric lift is real-linear, multiplicative, invertibility-preserving, and compatible with adjoints;
- spectral theory, norm theory, and positivity collapse into a single real SPD model valid for all imaginary directions and tensor structures.

These facts provide the algebraic and analytic framework used later for the inversion theorems and algorithmic constructions.

3 Structural and Topological Theorems

The structural content of AXIOM–CRE falls into two themes. First, the conjugation-normalised lift converts every invertible operator, regardless of its underlying algebra, into an SPD core that admits a pivot-free Cholesky factorisation. Second, the same lift recasts noncommutativity in geometric terms by comparing the rotation planes attached to different imaginary directions.

3.1 Global CRE Inversion via an SPD Core

Let $\mathcal{A}$ be a tensor-sorted AXIOM algebra satisfying ALH and let A act on a finite slice or tensor block. The conjugation-normalised lift Φ_{conj} need not be symmetric itself, but the associated energy operator

$$B := \Phi_{\mathrm{conj}}(A)^{\mathsf{T}} \Phi_{\mathrm{conj}}(A)$$

is always real SPD when A is injective. The inversion problem for A therefore reduces to a single SPD system that does not depend on whether the coefficients are real, complex, quaternionic, or tensor-valued.

Theorem 3.1 (Global AXIOM–CRE inversion via SPD core)**.** *Let $\mathcal{A}$ be a tensor-sorted AXIOM algebra satisfying ALH, and let A be an invertible*

operator acting on a finite-dimensional slice. Set

$$B := \Phi_{\mathrm{conj}}(A)^{\mathsf{T}} \Phi_{\mathrm{conj}}(A).$$

Then:

(1) B *is a real symmetric positive-definite matrix.*

(2) B *admits a unique Cholesky factorisation* $B = R^{\mathsf{T}} R$ *with* R *upper-triangular and positive diagonal, computed without pivoting.*

(3) The inverse of $\Phi_{\mathrm{conj}}(A)$ *can be written as*

$$\Phi_{\mathrm{conj}}(A)^{-1} = B^{-1}\, \Phi_{\mathrm{conj}}(A)^{\mathsf{T}} = R^{-1} R^{-\mathsf{T}}\, \Phi_{\mathrm{conj}}(A)^{\mathsf{T}}.$$

(4) The inverse of A *is recovered canonically by*

$$A^{-1} = \Phi_{\mathrm{conj}}^{-1}\big(\Phi_{\mathrm{conj}}(A)^{-1}\big),$$

where $\Phi_{\mathrm{conj}}^{-1}$ *is the inverse of the CRE embedding on the corresponding slice.*

Proof. By Theorem 2.7, $\Phi_{\mathrm{conj}}(A)$ has the same non-zero singular values as A. If A is invertible, then $\Phi_{\mathrm{conj}}(A)$ is invertible as well, so its singular values are strictly positive.

(1) For any non-zero real vector x,

$$x^{\mathsf{T}} B x = x^{\mathsf{T}} \Phi_{\mathrm{conj}}(A)^{\mathsf{T}} \Phi_{\mathrm{conj}}(A) x = \|\Phi_{\mathrm{conj}}(A) x\|_2^2 > 0,$$

because $\Phi_{\mathrm{conj}}(A)$ is injective. Thus B is real symmetric positive-definite.

(2) Standard SPD matrix theory gives a unique Cholesky factorisation $B = R^{\mathsf{T}}R$ with R upper-triangular and positive diagonal, without pivoting.

(3) Since B is SPD, we have $B^{-1} = R^{-1}R^{-\mathsf{T}}$. The normal-equations identity

$$\Phi_{\mathrm{conj}}(A)^{-1} = B^{-1}\,\Phi_{\mathrm{conj}}(A)^{\mathsf{T}}$$

follows by checking

$$\begin{aligned}\Phi_{\mathrm{conj}}(A)\left(B^{-1}\,\Phi_{\mathrm{conj}}(A)^{\mathsf{T}}\right) &= \Phi_{\mathrm{conj}}(A)B^{-1}\Phi_{\mathrm{conj}}(A)^{\mathsf{T}} \\ &= \Phi_{\mathrm{conj}}(A)(\Phi_{\mathrm{conj}}(A)^{\mathsf{T}}\Phi_{\mathrm{conj}}(A))^{-1}\Phi_{\mathrm{conj}}(A)^{\mathsf{T}} \\ &= I.\end{aligned}$$

and similarly on the other side.

(4) Finally, Φ_{conj} is a bijective $*$-algebra isomorphism on each slice, so its inverse exists and $\Phi_{\mathrm{conj}}^{-1}(\Phi_{\mathrm{conj}}(A)^{-1})$ is exactly A^{-1}. □

Algorithmic interpretation. To invert A, compute the CRE lift $F = \Phi_{\mathrm{conj}}(A)$, form the SPD core $B = F^{\mathsf{T}}F$, factor $B = R^{\mathsf{T}}R$ via pivot-free Cholesky, set $F^{-1} = R^{-1}R^{-\mathsf{T}}F^{\mathsf{T}}$, and decode the result by $A^{-1} = \Phi_{\mathrm{conj}}^{-1}(F^{-1})$. No complex-specific LU/QR, no quaternion- or octonion-specific kernels, and no multiplication tables are required; the numerical work reduces to one real SPD solve.

1. lift to the real embedding $F := \Phi_{\mathrm{conj}}(A)$;

2. form the SPD core $B = F^{\mathsf{T}}F$;

3. compute $B = R^{\mathsf{T}}R$ via Cholesky (no pivoting);

4. recover $F^{-1} = R^{-1}R^{-\mathsf{T}}F^{\mathsf{T}}$;

5. map back to A^{-1} via $\Phi_{\mathrm{conj}}^{-1}$.

3.2 Geometric Interpretation of Noncommutativity

Each imaginary direction in AXIOM specifies a complex structure on the CRE embedding, hence a decomposition of the real space into rotation planes. Commutators arise precisely when two such decompositions do not agree.

Definition 3.2 (Complex structure and rotation plane)**.** *A linear map $S : \mathbb{R}^{2m} \to \mathbb{R}^{2m}$ with $S^2 = -I$ is called a* complex structure*. Such an S decomposes $\mathbb{R}^{2m}$ into m invariant 2-planes, each of which is rotated by $\pi/2$ under S.*

Inside a tensor-sorted AXIOM algebra, each generator e_a induces such a complex structure on the CRE embedding, and different generators correspond to different decompositions into 2-planes.

Theorem 3.3 (Quaternionic noncommutativity as plane misalignment)**.** *Identify $\mathbb{H} \cong \mathbb{R}^4$ with basis $\{1, i, j, k\}$. Let L_i and L_j denote left-multiplication by i and j in the AXIOM–CRE real embedding. Then:*

(1) $L_i, L_j \in SO(4)$ are complex structures with $L_i^2 = L_j^2 = -I$.

(2) L_i and L_j act as $\pi/2$-rotations on two distinct invariant 2-planes.

(3) $L_iL_j = -L_jL_i$*, and the minus sign is exactly the* 180° *orientation mismatch between the underlying rotation planes.*

Proof sketch. In the standard real representation of $\mathbb{H}$, left-multiplication by i and j is given by explicit 4×4 real matrices $L_i, L_j \in SO(4)$. One checks that $L_i^2 = L_j^2 = -I$, so each is a complex structure.

Each of L_i and L_j preserves a decomposition of $\mathbb{R}^4$ into two invariant 2-planes on which it acts as a $\pi/2$-rotation. These decompositions are not the same; their common axis of action is reversed in orientation when the order of application is swapped. A direct calculation shows that

$$L_iL_j = -L_jL_i,$$

and the minus sign corresponds exactly to this orientation reversal. □

Corollary 3.4 (Commuting complex structures)**.** *Let* S_1, S_2 *be complex structures on* $\mathbb{R}^{2m}$*. Then* $S_1S_2 = S_2S_1$ *if and only if their invariant-plane decompositions coincide (up to simultaneous rotation within each* 2*-plane).*

Within a single slice (complex or quaternionic) the tensor-sorted and ALH axioms keep the planes aligned, so the familiar commutative or alternative behaviour is recovered. Noncommutativity only appears when different slices are compared globally.

3.3 Global Summary

Sections 2 and 3 together show that tensor-sorted AXIOM algebras with ALH provide a single index-sorted habitat for $\mathbb{R}, \mathbb{C}, \mathbb{H}, \mathbb{O}$ and all their

higher Cayley–Dickson analogues; the CRE lift is real-linear, multiplicative, injective, norm-isometric, and identifies every HPD operator with a real SPD matrix; every invertible operator therefore factors through the SPD core $B = \Phi_{\mathrm{conj}}(A)^{\mathsf{T}}\Phi_{\mathrm{conj}}(A)$ so that inversion reduces to one real Cholesky solve; and the residual noncommutativity is geometric, arising from misaligned rotation planes in the real embedding. These statements provide the structural backbone for the algorithmic and numerical developments that follow.

4 Structural Ingredients of CRE

Three ingredients underlie the CRE representation: a tensor-sorted AXIOM algebra that spans every 2^n (and infinite) imaginary direction without multiplication tables, a symmetric lift defined along any chosen direction with square -1, and the local complex slices supplied by the Arithmetic Local Hurwitz condition. The statements below record these components in the form used by the subsequent structural and algorithmic theorems.

4.1 Setup and notation

Let V be a finite-dimensional complex Hilbert space with $\dim_{\mathbb{C}} V = n$, and write $B(V)$ for the algebra of bounded complex-linear operators on V.

On the algebra side, we work with an AXIOM tensor-sorted algebra: a real algebra generated by symbols $\{e_i\}_{i\in I}$ (with I finite or countable)

satisfying the relations

$$e_i^2 = -1, \qquad e_i e_j = -e_j e_i \quad (i \neq j),$$

and we reduce any monomial by sorting indices and tracking the parity of swaps. This is exactly the tensor-sorted structure of Section 2.

Throughout this section, $\mathcal{A}$ denotes such a tensor-sorted AXIOM algebra satisfying the Arithmetic Local Hurwitz (ALH) condition of Definition 2.3. The CRE representation will always be applied on top of this algebra.

4.2 CRE symmetric lift (imaginary embedding)

We first isolate the key representation map of CRE: the symmetric lift. Unlike the classical Hermitian/skew-Hermitian splitting, we do *not* require a canonical decomposition; we only fix an imaginary direction $i^2 = -1$ and split along that slice.

Definition 4.1 (CRE symmetric lift along a chosen imaginary direction)**.** *Let $\mathcal{A}$ be a tensor-sorted AXIOM algebra and fix an element $i \in \mathcal{A}$ with $i^2 = -1$ that defines a local complex slice. Every element $a \in \mathcal{A}$ can be written (not uniquely) as*

$$a = h + ik,$$

with $h, k \in \mathcal{A}_{\mathbb{R}}$, the underlying real subalgebra. For any such splitting we

define the CRE symmetric lift

$$\Phi(a) := \begin{bmatrix} h & -k \\ k & h \end{bmatrix} \in M_2(\mathcal{A}_{\mathbb{R}}).$$

If A is a matrix or operator with entries in $\mathcal{A}$, we apply Φ entrywise to obtain a block operator on the doubled space.

This construction is purely algebraic and relies only on a fixed direction with square -1; it does not require any Hermitian/skew-Hermitian split.

Theorem 4.2 (CRE symmetric lift: algebraic properties)**.** *Let $\mathcal{A}$ be a tensor-sorted AXIOM algebra, fix $i \in \mathcal{A}$ with $i^2 = -1$, and let Φ be the CRE symmetric lift from Definition 4.1. Then Φ satisfies the algebraic properties of Theorem 2.6. In particular, for all $a, b \in \mathcal{A}$ and $\alpha, \beta \in \mathbb{R}$:*

1. ***Real linearity:***

$$\Phi(\alpha a + \beta b) = \alpha\Phi(a) + \beta\Phi(b).$$

2. ***Multiplicativity:***

$$\Phi(ab) = \Phi(a)\Phi(b).$$

3. ***Injectivity and inverse preservation:*** *If $\Phi(a) = 0$ then $a = 0$. If a is invertible in $\mathcal{A}$, then $\Phi(a)$ is invertible in $M_2(\mathcal{A}_{\mathbb{R}})$ and*

$$\Phi(a^{-1}) = \Phi(a)^{-1}.$$

Proof. The map Φ in Definition 4.1 is exactly the symmetric lift considered in Theorem 2.6, specialised to the chosen imaginary unit i and taking values in $M_2(\mathcal{A}_{\mathbb{R}})$. Hence all three properties follow directly from Theorem 2.6, and no additional argument is needed. □

Remark 4.3 (Classical case vs. CRE generality)**.** *When $\mathcal{A} = \mathbb{C}$ and i is the standard imaginary unit, the lift becomes $a = x + iy \mapsto \left(\begin{smallmatrix} x & -y \\ y & x \end{smallmatrix}\right)$. The same formula works in every tensor-sorted AXIOM algebra without referring to Hermitian/skew-Hermitian decompositions; it only needs a direction with square -1.*

4.3 Tensor-sorted AXIOM algebra on 2^n and infinite dimensions

The next definition specifies the tensor-sorted product that replaces all quaternionic, octonionic, and higher Cayley–Dickson tables.

Definition 4.4 (AXIOM tensor-sorted algebra)**.** *Let I be a finite or countable index set. Let T be the free real algebra generated by symbols $\{e_i : i \in I\}$ and impose the relations*

$$e_i^2 = -1, \qquad e_i e_j = -e_j e_i \quad (i \neq j).$$

Define the quotient algebra

$$\mathcal{A} := T \Big/ \langle e_i^2 + 1,\ e_i e_j + e_j e_i\ (i \neq j) \rangle.$$

For a finite subset $J = \{i_1 < \cdots < i_k\} \subset I$, write

$$e_J := e_{i_1} \cdots e_{i_k}, \qquad e_\emptyset := 1.$$

A monomial is called sorted *if its indices are strictly increasing. We refer to such an algebra $\mathcal{A}$ as an AXIOM tensor-sorted algebra.*

Theorem 4.5 (AXIOM tensor-sorted basis and unique product)**.** *Let $\mathcal{A}$ be as in Definition 4.4. Then:*

1. *The set*

$$\mathcal{B} := \{e_J : J \subset I,\ |J| < \infty,\ J \text{ strictly ordered}\}$$

   is a real vector space basis of $\mathcal{A}$. In particular, if $|I| = n < \infty$, then $\dim_{\mathbb{R}} \mathcal{A} = 2^n$.

2. *Any monomial $e_{i_1} \cdots e_{i_k}$ in the generators reduces uniquely to*

$$e_{i_1} \cdots e_{i_k} = \sigma(J)\, e_J,$$

   where J is the strictly ordered set of distinct indices appearing, and $\sigma(J) \in \{\pm 1\}$ depends only on

   - *the parity of swaps needed to sort the indices, and*
   - *the number of cancellations coming from $e_i^2 = -1$.*

3. *Consequently, multiplication in $\mathcal{A}$ is completely determined by three*

local rules:

$$e_i^2 = -1, \qquad e_i e_j = -e_j e_i \ (i \neq j),$$

then sort indices and attach the sign from swap parity.

Proof. Spanning is obtained by repeatedly applying $e_i^2 = -1$ and anti-symmetry to eliminate duplicate indices and sort. This always terminates in finitely many steps, yielding $\pm e_J$ for some strictly ordered J.

For linear independence and uniqueness, equip the free monoid on $\{e_i\}$ with a term order (e.g. lexicographic) and view the rewriting rules

$$e_i e_j \mapsto -e_j e_i \ (i > j), \qquad e_i e_i \mapsto -1$$

as a terminating rewrite system. Each monomial has a unique normal form $\pm e_J$; if there were a nontrivial linear dependence $\sum_k \alpha_k e_{J_k} = 0$ with distinct J_k, lifting back to the free algebra and reducing would produce a nontrivial relation among normal forms, contradiction. Thus $\mathcal{B}$ is a basis and the reduction $\sigma(J)\, e_J$ is unique for each monomial.

The description of multiplication then follows: the product of two basis elements is a monomial that reduces uniquely via these rules. □

Remark 4.6 (Table-free unification of complex/quaternion/octonion cases)**.** *For $|I| = 1, 2, 3$, the rules in Theorem 4.5 reproduce the usual complex, quaternionic, and octonionic products, but they do so* without *any multiplication tables or Fano-plane diagrams. All higher 2^n and infinite-dimensional analogues are handled by the same mechanism. This is the precise sense in*

which AXIOM replaces "표로 곱셈" by a table-free tensor-sorted calculus.

4.4 Arithmetic Local Hurwitz: local complex slices

The next statement records the local complex structure attached to each generator, which is the algebraic core of the ALH condition.

Theorem 4.7 (Arithmetic Local Hurwitz: local complex slices)**.** *Let $\mathcal{A}$ be a tensor-sorted AXIOM algebra as in Definition 4.4. For each fixed index $i \in I$, define the real 2-dimensional subspace*

$$C_i := \mathrm{span}_{\mathbb{R}}\{1, e_i\} \subset \mathcal{A}.$$

Then:

1. *C_i is a subalgebra of $\mathcal{A}$, closed under multiplication.*
2. *The map*

$$\varphi_i : C_i \to \mathbb{C}, \qquad \varphi_i(a + be_i) = a + bi$$

is a real-algebra isomorphism. In particular,

$$(a + be_i)(c + de_i) = (ac - bd) + (ad + bc)e_i,$$

exactly as in $\mathbb{C}$.

Proof. Take $x = a + be_i$, $y = c + de_i \in C_i$. Then

$$xy = (a+be_i)(c+de_i) = ac+ade_i+bce_i+bd\,e_i^2 = (ac-bd)+(ad+bc)e_i,$$

using $e_i^2 = -1$. Hence $xy \in C_i$, so C_i is a subalgebra.

The map φ_i is clearly real-linear and bijective. Using the above multiplication,

$$\varphi_i(xy) = (ac - bd) + (ad + bc)i = (a + bi)(c + di) = \varphi_i(x)\,\varphi_i(y),$$

so φ_i is a real-algebra isomorphism. □

Remark 4.8 (Connection to the full ALH condition)**.** *Theorem 4.7 shows that each generator e_i induces a genuine complex slice $C_i \cong \mathbb{C}$ inside the tensor-sorted AXIOM algebra. The full Arithmetic Local Hurwitz condition (Definition 2.3) strengthens this by imposing norm-composition on compatible finite sets of generators and on triple products along quaternionic or higher slice chains. For CRE, the key point is that every computation can be decomposed into such local slices, where complex arithmetic and inversion behave exactly as in $\mathbb{C}$, even though the global algebra may live in arbitrary 2^n or infinite dimension.*

4.5 Summary

The structural foundations for AXIOM–CRE can now be summarised as follows:

- **AXIOM tensor-sorted algebra:** a single table-free product rule (index sorting + sign) generates all 2^n and infinite-dimensional imaginaries, recovering $\mathbb{C}, \mathbb{H}, \mathbb{O}$ as finite slices.
- **CRE symmetric lift:** for any chosen direction $i^2 = -1$, the map Φ is

real-linear, multiplicative, injective, and preserves inverses, providing a uniform real block representation for all tensor-valued operators.

- **Arithmetic Local Hurwitz slices:** each generator e_i carries a local complex slice $C_i \cong \mathbb{C}$, so all computations can be reduced to classical complex arithmetic on compatible slices, even though the ambient algebra may be noncommutative and high-dimensional.

These components are precisely the hypotheses invoked later in Section 3 for the global inversion theorem and the geometric interpretation of noncommutativity.

5 The AXIOM–CRE Inversion Theorem

The tensor-sorted algebra, the CRE symmetric lift, and the local complex slices together yield a single inversion mechanism: any CRE operator reduces to one pivot-free Cholesky factorisation of a real SPD block, independent of tensor dimension and without multiplication tables.

5.1 Statement of the main theorem

Theorem 5.1 (AXIOM–CRE Inversion Theorem: pivot-free SPD block inversion)**.** *Let $\mathcal{A}$ be an AXIOM tensor-sorted algebra as in Definition 4.4, satisfying the Arithmetic Local Hurwitz condition. Let V be a finite-dimensional complex vector space whose coordinates take values in $\mathcal{A}$, and let*

$$A \in B(V)$$

be an invertible operator (possibly with internal 2^n-dimensional or infinite-tensor structure).

Fix a local imaginary direction $i \in \mathcal{A}$ with $i^2 = -1$, and let Φ denote the CRE symmetric lift along i as in Theorem 4.2. Set

$$S := \Phi(A) \in B(V \oplus V).$$

Then:

1. ***Real block lift and invertibility.*** *The operator S is a real block operator on $V \oplus V$ and*

   $$\Phi(A^{-1}) = \Phi(A)^{-1} = S^{-1}.$$

   In particular, A is invertible if and only if S is invertible.

2. ***SPD normal equation.*** *Define the real symmetric positive-definite operator*

   $$C := S^{\mathsf{T}} S.$$

   Then for all nonzero x in the realified $V \oplus V$,

   $$\langle x, Cx \rangle = \langle x, S^{\mathsf{T}} S x \rangle = \langle Sx, Sx \rangle > 0,$$

   so C is SPD.

3. ***Unique Cholesky factorization.*** *The SPD operator C admits a unique*

Cholesky factorization

$$C = LL^{\mathsf{T}},$$

with L real lower-triangular (or block-lower-triangular, respecting the tensor block structure). No pivoting is required.

4. ***Recovery of S^{-1} and A^{-1}.*** *From $C = S^{\mathsf{T}}S$ we have*

$$C^{-1} = (S^{\mathsf{T}}S)^{-1} = S^{-1}(S^{\mathsf{T}})^{-1}.$$

Using the Cholesky factorization,

$$C^{-1} = (LL^{\mathsf{T}})^{-1} = L^{-\mathsf{T}}L^{-1},$$

and hence

$$S^{-1} = L^{-\mathsf{T}}L^{-1}S^{\mathsf{T}}.$$

Since Φ is injective and multiplicative,

$$\Phi(A^{-1}) = S^{-1} = L^{-\mathsf{T}}L^{-1}S^{\mathsf{T}},$$

and there exists a unique operator $B \in B(V)$ with $\Phi(B) = S^{-1}$, namely $B = A^{-1}$.

5. ***Uniformity across tensor dimension (AXIOM table-free property).*** *By Theorem 4.5, all internal tensor products in the entries of A, S,*

C, and L are determined solely by index sorting and sign parity; no quaternionic, octonionic, or Cayley–Dickson multiplication tables are needed. By Theorem 4.7, each generator e_i spans a local complex slice $C_i \cong \mathbb{C}$, so scalar-like arithmetic inside Cholesky and triangular solves behaves exactly as in the usual complex (or 2×2 real) case. Thus the entire inversion pipeline is independent of tensor dimension and algebraic complexity.

In particular, every invertible operator A in the AXIOM–CRE framework admits the algorithm

$$A \xrightarrow{\Phi} S = \Phi(A) \xrightarrow{S^{\mathsf{T}}S} C \xrightarrow{\textit{Cholesky}} L \xrightarrow{\textit{back-substitution}} S^{-1},\ A^{-1},$$

using only one *SPD Cholesky factorization on a real symmetric block and no pivoting or multiplication tables.*

All of these operations occur inside the original storage space; the CRE map reindexes the same real scalars rather than duplicating data. The imaginary directions act as orthogonal rotations of a fixed real information space, so the inversion stays within that space.

5.2 Proof via structural ingredients

Proof. We indicate where each structural ingredient of AXIOM–CRE enters.

Step 1: CRE symmetric lift. By Theorem 4.2, fixing a direction $i^2 = -1$ we can write each entry of A as $h + ik$ with h, k in the real subalgebra $\mathcal{A}_{\mathbb{R}}$ and

define

$$S := \Phi(A) = \begin{bmatrix} H & -K \\ K & H \end{bmatrix},$$

for appropriate block operators H, K on V. Theorem 4.2 guarantees that Φ is an injective algebra homomorphism and

$$\Phi(AB) = \Phi(A)\Phi(B), \qquad \Phi(A^{-1}) = \Phi(A)^{-1}.$$

Thus A is invertible if and only if S is invertible.

Step 2: SPD normal equation. Let $C := S^{\mathsf{T}}S$. Then C is real symmetric. If $x \neq 0$ in the realified $V \oplus V$, then

$$\langle x, Cx\rangle = \langle x, S^{\mathsf{T}}Sx\rangle = \langle Sx, Sx\rangle > 0,$$

because S is invertible. Hence C is SPD.

Step 3: Cholesky factorization. Since C is SPD, classical SPD theory yields a unique Cholesky factorization

$$C = LL^{\mathsf{T}}$$

with L real lower-triangular (or block-lower-triangular). No pivoting is needed.

Step 4: Tensor-sorted control (AXIOM). By Theorem 4.5, any product of generators e_i that appears in the entries of H, K, S, C, L reduces uniquely to $\pm e_J$ for some strictly ordered index set J. Thus all internal multiplications

during the Cholesky factorization and triangular solves are handled by the same index-sorting + sign-parity rule; no dimension-specific multiplication tables are involved.

Step 5: Local complex slices (ALH). By Theorem 4.7, each generator e_i spans a complex slice $C_i \cong \mathbb{C}$ inside $\mathcal{A}$. Along these slices, scalar arithmetic is exactly complex multiplication. When Cholesky and triangular solves work with scalars built from fixed finite sets of generators, they operate slice-wise as in the classical complex case; ALH guarantees that nonassociativity and noncommutativity only appear when mixing incompatible slices, which does not occur inside the local scalar computations.

Step 6: Recovering S^{-1} and A^{-1}. From $C = S^{\mathsf{T}}S$ and $C = LL^{\mathsf{T}}$ we have

$$C^{-1} = (LL^{\mathsf{T}})^{-1} = L^{-\mathsf{T}}L^{-1}.$$

Since S is invertible,

$$S^{-1} = C^{-1}S^{\mathsf{T}} = L^{-\mathsf{T}}L^{-1}S^{\mathsf{T}}.$$

By Theorem 4.2,

$$S^{-1} = \Phi(A^{-1}),$$

and because Φ is injective, A^{-1} is uniquely determined as the element of $B(V)$ whose lift equals S^{-1}. This yields all statements in Theorem 5.1. $\square$

5.3 Interpretation

The AXIOM–CRE Inversion Theorem encapsulates the structural message of the framework:

- **CRE symmetric lift as a universal real encoding.** Any operator A with entries in an AXIOM tensor-sorted algebra admits a real block encoding $S = \Phi(A)$ that preserves the full algebraic structure, including invertibility.

- **AXIOM tensor-sorted algebra as a table-free backend.** All internal tensor products — from complex numbers up to quaternionic, octonionic, and higher Cayley–Dickson-type behaviour — are governed by a single rule (index sorting and sign parity), so the real blocks S, C, and L require no ad-hoc multiplication tables.

- **Arithmetic Local Hurwitz as a local complex model.** Each generator e_i induces a genuine complex slice, so scalar-like computations inside Cholesky and triangular solves behave exactly as in $\mathbb{C}$, even when the ambient algebra is much larger and noncommutative.

- **Dimension-independent inversion.** The inversion of A reduces to one SPD Cholesky factorization of $C = S^{\mathsf{T}}S$, with no pivoting and no dependence on how many imaginary directions or tensor factors are present. In this sense, real, complex, quaternionic, octonionic, and general tensor-valued inversions all collapse to the same SPD routine.

- **Rotation rather than duplication.** The CRE map reorders a fixed real information space; it does not add a second copy of the data. Al-

gebraic realification and complexification correspond to internal rotations inside the AXIOM algebra.

This theorem is the bridge from the structural axioms of Section 4 to the algorithmic constructions and numerical benchmarks developed in the later chapters.

6 Analytical and Cost Model Framework

The conjugation–real embedding Φ_{conj} acts as a coordinate rotation inside a fixed real information space rather than as a traditional realification. It reassigns one tensor degree of freedom between the real and slice-complex coordinates, so no extra storage or symbolic degrees of freedom are introduced. The discussion below records the operator-algebraic bounds that justify this viewpoint and the associated cost model that follows once Hermitian or HPD operators are identified with their real SPD lifts from Section 2.

6.1 Operator-algebraic foundations

We first pin down the operator-theoretic status of CRE in the classical finite-dimensional complex setting. Let

$$\mathcal{A}_{\mathbb{C}} := \left(\mathbb{C}^{n\times n},\ \|\cdot\|_2,\ *\right)$$

be the algebra of complex $n \times n$ matrices with operator 2-norm and adjoint A^*, and let

$$\mathcal{A}_{\mathbb{R}} := \left(\Phi_{\mathrm{conj}}(\mathbb{C}^{n\times n}),\ \|\cdot\|_2,\ ^{\mathsf{T}}\right)$$

be its CRE image inside real $2n \times 2n$ matrices equipped with the same operator norm and transpose.

For $A \in \mathbb{C}^{n\times n}$ we write $A = X + iY$ with $X, Y \in \mathbb{R}^{n\times n}$ and set

$$\Phi_{\mathrm{conj}}(A) := \begin{bmatrix} \Re A & \Im A \\ -\Im A & \Re A \end{bmatrix} = \begin{bmatrix} X & Y \\ -Y & X \end{bmatrix}.$$

In AXIOM terms, this is exactly the conjugation-normalised CRE lift along the chosen imaginary direction $i^2 = -1$.

Theorem 6.1 (Banach $*$-isometry for CRE)**.** *The map $\Phi_{\mathrm{conj}} : \mathcal{A}_{\mathbb{C}} \to \mathcal{A}_{\mathbb{R}}$ is a bijective $*$-algebra homomorphism satisfying*

$$\|\Phi_{\mathrm{conj}}(A)\|_2 = \|A\|_2 \qquad \textit{for all } A \in \mathbb{C}^{n\times n}.$$

In particular, $\mathcal{A}_{\mathbb{C}}$ and $\mathcal{A}_{\mathbb{R}}$ carry exactly the same analytic information; CRE merely changes coordinates.

Proof. Homomorphism and $*$-structure. Write $A = X + iY$, $B = U + iV$. A direct block computation shows

$$\Phi_{\mathrm{conj}}(AB) = \begin{bmatrix} XU - YV & XV + YU \\ -(XV + YU) & XU - YV \end{bmatrix} = \Phi_{\mathrm{conj}}(A)\,\Phi_{\mathrm{conj}}(B),$$

so Φ_{conj} is multiplicative. Similarly,

$$A^* = X^\top - iY^\top \quad \Rightarrow \quad \Phi_{\mathrm{conj}}(A^*) = \begin{bmatrix} X^\top & -Y^\top \\ Y^\top & X^\top \end{bmatrix} = \Phi_{\mathrm{conj}}(A)^{\mathsf{T}},$$

so Φ_{conj} is $*$-preserving. Injectivity and surjectivity onto $\mathcal{A}_{\mathbb{R}}$ are obvious from the block form.

Norm preservation. As in Section 2, there exists an orthogonal matrix Q (built from fixed 2×2 rotations) such that

$$Q^{\mathsf{T}} \Phi_{\mathrm{conj}}(A) Q = \begin{bmatrix} A & 0 \\ 0 & \overline{A} \end{bmatrix}.$$

Hence

$$\|\Phi_{\mathrm{conj}}(A)\|_2 = \left\|Q^{\mathsf{T}} \Phi_{\mathrm{conj}}(A) Q\right\|_2 = \left\| \begin{bmatrix} A & 0 \\ 0 & \overline{A} \end{bmatrix} \right\|_2 = \|A\|_2,$$

since the spectral norm of a block-diagonal matrix is the maximum of the spectral norms of its diagonal blocks, and $\|\overline{A}\|_2 = \|A\|_2$. □

Corollary 6.2 (Spectral and resolvent continuity under CRE)**.** *For all* $A \in \mathbb{C}^{n \times n}$,

$$\sigma\big(\Phi_{\mathrm{conj}}(A)\big) = \sigma(A) \cup \overline{\sigma(A)},$$

and for every $z \in \mathbb{C} \setminus \sigma(A)$*,*

$$\big(\Phi_{\mathrm{conj}}(A) - zI_{2n}\big)^{-1} = \Phi_{\mathrm{conj}}\big((A - zI_n)^{-1}\big).$$

Thus eigenvalues, spectral radii, resolvents, and condition numbers are preserved under CRE.

Information–geometric interpretation. Theorem 6.1 shows that Φ_{conj} embeds $\mathbb{C}^{n\times n}$ into a real $2n \times 2n$ geometry without altering norms or spectra. In coordinates this can be viewed as

$$(x, y) \longmapsto (x\cos\theta - y\sin\theta,\ x\sin\theta + y\cos\theta),$$

where θ fixes which two-dimensional plane represents the imaginary direction. Complex multiplication $(a + ib)(c + id)$ corresponds to moving information within these rotated axes. In AXIOM terms, every complex or tensor-valued slice lives inside one tensor-sorted real algebra; CRE selects a slice and expresses the computation as motion within that slice rather than as the creation of a new degree of freedom.

6.2 Analytic–computational integration

The operator-theoretic picture links directly to factorisation and cost. Section 2 showed that every HPD operator in a tensor-sorted AXIOM algebra has a unique real SPD representative via Φ_{conj} (Theorem 2.9, Corollary 2.11), so Hermitian positive-definite matrices have real SPD lifts with the same spectra and condition numbers.

Theorem 6.3 (Equivalence of analytic and computational factorizations)**.** *Let $H \in \mathbb{C}^{n\times n}$ be Hermitian positive-definite and let $H = LL^*$ be its (unique) Cholesky factorization with L lower triangular. Then:*

(1) $\Phi_{\mathrm{conj}}(H)$ is real SPD and admits the real Cholesky factorization

$$\Phi_{\mathrm{conj}}(H) = \Phi_{\mathrm{conj}}(L)\,\Phi_{\mathrm{conj}}(L)^{\mathsf{T}}.$$

(2) For every right-hand side $b \in \mathbb{C}^n$, the solution $x = H^{-1}b$ satisfies

$$\Phi_{\mathrm{conj}}(x) = \big(\Phi_{\mathrm{conj}}(H)\big)^{-1}\Phi_{\mathrm{conj}}(b).$$

(3) All eigenvalues, singular values, and condition numbers of H and $\Phi_{\mathrm{conj}}(H)$ coincide.

Factorization and inversion therefore occur inside the same real information space; CRE simply chooses coordinates in which the SPD structure is explicit.

Proof. (1) By Theorem 2.9, $\Phi_{\mathrm{conj}}(H)$ is real SPD. Applying the $*$-homomorphism property of Φ_{conj},

$$\Phi_{\mathrm{conj}}(H) = \Phi_{\mathrm{conj}}(LL^*) = \Phi_{\mathrm{conj}}(L)\,\Phi_{\mathrm{conj}}(L^*) = \Phi_{\mathrm{conj}}(L)\,\Phi_{\mathrm{conj}}(L)^{\mathsf{T}},$$

and uniqueness of the SPD Cholesky factorization then identifies $\Phi_{\mathrm{conj}}(L)$ as the real Cholesky factor.

(2) Since Φ_{conj} is multiplicative and bijective,

$$\Phi_{\mathrm{conj}}(x) = \Phi_{\mathrm{conj}}(H^{-1}b) = \Phi_{\mathrm{conj}}(H^{-1})\,\Phi_{\mathrm{conj}}(b) = \big(\Phi_{\mathrm{conj}}(H)\big)^{-1}\Phi_{\mathrm{conj}}(b).$$

(3) Follows from Theorem 6.1 and Corollary 6.2. □

Cholesky on $\Phi_{\mathrm{conj}}(H)$ is therefore the same factorisation expressed in real SPD coordinates. The next lemma packages this observation into a simple cost model.

Lemma 6.4 (Idealised arithmetic complexity reduction)**.** *Consider dense $n \times n$ Hermitian positive-definite matrices H over $\mathbb{C}$. Assume:*

- *a standard dense cost model in which real Cholesky on an SPD matrix in $\mathbb{R}^{n\times n}$ costs $c_{\mathrm{SPD}}n^3$ real FLOPs;*
- *classical complex Cholesky on H is implemented by four real multiplications and two real additions per complex multiplication (the textbook model), yielding*

$$\mathrm{cost}_{\mathbb{C}}(H) \approx 4c_{\mathrm{SPD}}n^3$$

real FLOPs.

Then, when H is processed via its real SPD CRE lift $\Phi_{\mathrm{conj}}(H)$ with a single real SPD Cholesky factorization, the idealised arithmetic cost reduces by a factor of 4:

$$\mathrm{cost}_{\mathbb{R}}\big(\Phi_{\mathrm{conj}}(H)\big) \approx c_{\mathrm{SPD}}n^3 \quad \Rightarrow \quad \frac{\mathrm{cost}_{\mathbb{C}}(H)}{\mathrm{cost}_{\mathbb{R}}(\Phi_{\mathrm{conj}}(H))} \approx 4.$$

Discussion of the model. The complex Cholesky algorithm applies the same sequence of algebraic updates as the real algorithm, but each scalar multiplication and addition is carried out in $\mathbb{C}$ rather than $\mathbb{R}$. Under the textbook decomposition

$$(a + ib)(c + id) = (ac - bd) + i(ad + bc),$$

a single complex multiplication costs 4 real multiplications and 2 real additions. Thus the total real-FLOP cost of complex Cholesky is approximately 4 times that of real Cholesky on the same $n \times n$ grid.

In the AXIOM–CRE approach, H is never processed as a complex matrix: we work directly on its real SPD representative $\Phi_{\mathrm{conj}}(H)$, and all scalar operations are real. The tensor-sorted structure guarantees that all internal products of imaginary generators are resolved by index sorting and sign parity, so no extra algebraic overhead arises. In the idealised dense model, this yields a factor-4 reduction in arithmetic cost at the scalar level; actual wall-clock speedups then depend on memory hierarchy and implementation, but the information-theoretic footprint is captured by the $4\times$ ratio. □

Information-flow interpretation. From the AXIOM–CRE viewpoint, the speedup in Lemma 6.4 is not just an accounting trick in FLOP constants. It reflects a genuine reduction in *information traffic*:

- In a classical complex Cholesky, conjugate-coupled data are carried as two independent real streams (real and imaginary parts), and many operations are effectively duplicated to maintain consistency.

- In the CRE formulation, the same information is stored once, as a configuration in a real tensor-sorted space. The embedding Φ_{conj} chooses an orthogonal frame where all HPD structure is real SPD, so information travels once along rotated axes rather than twice along independent real/imag channels.

The computation is therefore carried by the movement of information inside a fixed real tensor space, rather than by duplicating conjugate data. Spectra, norms, and resolvents remain unchanged, while the workload collapses to a single SPD routine with an intrinsic $4\times$ reduction at the scalar level and a pronounced speedup once memory traffic is considered.

Conclusion. To summarise: Φ_{conj} is a Banach $*$-isometry, so the real CRE image retains all spectral data of the original complex operator; every Hermitian or HPD matrix has a canonical real SPD representative; and dense factorisations reduce to an SPD routine, eliminating redundant conjugate operations and, ideally, reducing the flop cost by a factor of four in the scalar model.

7 Hermitian CRE Solvers: Algorithms and Experiments

CRE has been established as a $*$-algebra isomorphism: HPD operators are the same objects as real SPD operators under the lift (Theorem 2.9, Corollary 2.11). Here this algebraic fact is carried through to practical algorithms and benchmarks: Cholesky, Schur complements, and inversion

all commute with Φ_{conj} and therefore run entirely as real SPD routines. Each complex entry acts as an $SO(2)$ rotation on $\mathbb{R}^2$, and the resulting "orthogonal information flow'' maintains spectra, condition numbers, and stability while delivering the expected speedups once the dimension is large.

7.1 HPD to SPD: Analytic Backbone of the CRE Solve

By Theorem 2.9, every HPD operator H has a canonical real SPD representative $\Phi_{\mathrm{conj}}(H)$ with identical spectral and norm data. The next steps specialise this to the Cholesky setting and turn the "HPD = SPD'' identification into an algorithm.

Corollary 7.1 (CRE Cholesky backbone for HPD matrices)**.** *Let $H \in \mathbb{C}^{n\times n}$ be Hermitian positive-definite. Set*

$$\overline{H} := \Phi_{\mathrm{conj}}(H).$$

Then:

(1) $\overline{H}$ is real symmetric positive-definite and $\|\overline{H}\|_2 = \|H\|_2$.

(2) If $H = LL^*$ is the complex Cholesky factorisation, then

$$\overline{H} = \Phi_{\mathrm{conj}}(H) = \Phi_{\mathrm{conj}}(L)\,\Phi_{\mathrm{conj}}(L)^{\mathsf{T}}.$$

(3) Conversely, if $\overline{H} = GG^{\mathsf{T}}$ is the real SPD Cholesky factorisation, then

$$H = \Phi_{\mathrm{conj}}^{-1}(G)\,\Phi_{\mathrm{conj}}^{-1}(G)^*$$

is the unique complex Cholesky factorisation of H.

Proof. (1) is exactly Theorem 2.9 specialised to $\mathbb{C}^{n\times n}$: Φ_{conj} is a Banach $*$-isometry (Theorem 6.1), so it preserves adjoints, spectra, and the operator norm.

(2) follows from multiplicativity and $*$-preservation:

$$\Phi_{\mathrm{conj}}(H) = \Phi_{\mathrm{conj}}(LL^*) = \Phi_{\mathrm{conj}}(L)\,\Phi_{\mathrm{conj}}(L^*) = \Phi_{\mathrm{conj}}(L)\,\Phi_{\mathrm{conj}}(L)^{\mathsf{T}}.$$

For (3), Φ_{conj} is bijective and $*$-preserving, so any real SPD factorisation $\overline{H} = GG^{\mathsf{T}}$ lifts uniquely to a complex factorisation $H = \tilde{L}\tilde{L}^*$ with $\tilde{L} = \Phi_{\mathrm{conj}}^{-1}(G)$. Uniqueness of Cholesky for HPD matrices forces $\tilde{L} = L$. □

Corollary 7.2 (Spectral and conditioning invariance under CRE)**.** *Let H and $\overline{H}$ be as above. Then*

$$\sigma(\overline{H}) = \sigma(H) \cup \overline{\sigma(H)}, \qquad \kappa_2(\overline{H}) = \kappa_2(H).$$

In particular, the complex and real Cholesky factorisations see the same eigenvalues, the same singular values, and the same condition number.

Proof. This is Corollary 6.2 applied to HPD matrices, together with norm-isometry of Φ_{conj}. □

Interpretation. From this point of view we are not "turning an HPD matrix into an SPD matrix''; we are simply *viewing the same HPD matrix in an SPD coordinate system*. All analytic information carried by H over $\mathbb{C}$

(eigenvalues, condition numbers, inner-product geometry) is already present in $\overline{H}$, and SPD Cholesky is just the most efficient and most stable way to exploit that information once we have chosen the right coordinates.

7.2 Rotational / Clifford View: Information as Moving Points

On the level of a single scalar, Φ_{conj} reduces complex multiplication to a 2×2 real rotation–scaling block. This is the minimal real representation of $\mathbb{C}$ inside the Clifford algebra $Cl_{0,2}$.

Recall

$$J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \qquad \Phi_{\mathrm{conj}}(A) = (I_n \otimes J)\, \Phi(A)\, (I_n \otimes J)^{\mathsf{T}}.$$

Each complex scalar acts as a rotation–dilation on $\mathbb{R}^2$.

Lemma 7.3 (Blockwise rotational action)**.** *For $a \in \mathbb{C}$,*

$$\Phi_{\mathrm{conj}}(a) = \begin{bmatrix} \Re a & \Im a \\ -\Im a & \Re a \end{bmatrix}, \qquad \Phi_{\mathrm{conj}}(a)\, \Phi_{\mathrm{conj}}(a)^{\mathsf{T}} = |a|^2 I_2.$$

Thus $\Phi_{\mathrm{conj}}(a)$ is an element of $SO(2)$ scaled by $|a|$.

Proof. Direct multiplication gives

$$\begin{bmatrix} x & y \\ -y & x \end{bmatrix} \begin{bmatrix} x & -y \\ y & x \end{bmatrix} = \begin{bmatrix} x^2 + y^2 & 0 \\ 0 & x^2 + y^2 \end{bmatrix} = |a|^2 I_2,$$

for $a = x + iy$. The trace fixes the cosine component of the rotation angle.

□

Information-flow viewpoint. In this basis a complex scalar rotates and rescales a point in $\mathbb{R}^2$, so operator composition becomes concatenation of rotations in n decoupled planes. Complex conjugation acts as a reflection, reversing the direction of rotation. CRE therefore treats complex linear algebra as the motion of information points under an orthogonal action on a fixed real tensor space.

7.3 Schur Complements, Information Transport, and Cost

Inversion of $\overline{H} = \Phi_{\mathrm{conj}}(H)$ proceeds through the standard real Schur complement. CRE ensures that this is not a new formula but exactly the same algebra the complex analyst was already using, now written in the SPD coordinate system.

Theorem 7.4 (Schur complement under CRE)**.** *Let $H \in \mathbb{C}^{n\times n}$ be HPD and $\overline{H} = \Phi_{\mathrm{conj}}(H)$. Partition*

$$\overline{H} = \begin{pmatrix} A & B \\ B^{\mathsf{T}} & D \end{pmatrix},$$

with $A \in \mathbb{R}^{k\times k}$. If A is invertible, then

$$S := D - B^{\mathsf{T}} A^{-1} B$$

is real SPD and

$$\overline{H}^{-1} = \begin{pmatrix} A^{-1} + A^{-1}BS^{-1}B^{\mathsf{T}}A^{-1} & -A^{-1}BS^{-1} \\ -S^{-1}B^{\mathsf{T}}A^{-1} & S^{-1} \end{pmatrix}.$$

Proof. Since H is HPD, $\overline{H}$ is real SPD (Theorem 2.9). For any SPD block matrix the Schur complement S is SPD and the stated inverse formula holds; this is the standard real Schur complement theorem. The CRE structure is only used to guarantee that $\overline{H}$ is SPD and that all spectral properties are consistent with the complex side. □

Corollary 7.5 (Complexity and backward stability)**.** *Let H and $\overline{H}$ be as above. Solving $Hx = b$ via CRE (i.e. working entirely with $\overline{H}$) requires one real SPD Cholesky factorisation and two triangular solves, with arithmetic cost*

$$C_{\mathrm{CRE}} \approx \tfrac{8}{3}n^3$$

in real operations, compared to roughly $8n^3$ real operations for complex LU factorisation. Backward-error stability is identical to that of real SPD Cholesky.

Proof. Real SPD Cholesky costs $\frac{1}{3}(2n)^3 = \frac{8}{3}n^3$ floating-point operations. Complex LU with partial pivoting costs approximately $8n^3$ real flops. Backward stability is the standard result for SPD Cholesky [7, 8]; norm-isometry of Φ_{conj} transfers the same bound to H. □

Remark (cost vs. geometry). The Schur complement tracks how much information crosses the block boundary. SPD structure keeps this traffic one-way and pivot-free, and CRE picks coordinates where that structure is visible to the machine. The observed 3–4$\times$ speedup is the numerical shadow of this geometric simplification.

7.4 Algorithmic Realisation of CRE Solves

The CRE-based solver can be written explicitly. The algorithms below are orthogonal rephrasings of the classical HPD solve in the SPD coordinate system.

Algorithm 1 Cholesky factorisation via conjugation–real embedding

1: **Input:** $H \in \mathbb{C}^{n\times n}$ (HPD)
2: Form $\overline{H} = \Phi_{\mathrm{conj}}(H)$
3: Compute $\overline{H} = GG^{\mathsf{T}}$ via real SPD Cholesky
4: Return $L = \Phi_{\mathrm{conj}}^{-1}(G)$

Algorithm 2 Inverse of an HPD matrix via orthogonal CRE

1: **Input:** $H \in \mathbb{C}^{n\times n}$ (HPD)
2: Compute $\overline{H} = \Phi_{\mathrm{conj}}(H)$
3: Factor $\overline{H} = GG^{\mathsf{T}}$ (real SPD Cholesky)
4: Compute $\overline{H}^{-1} = G^{-\mathsf{T}}G^{-1}$
5: Output $H^{-1} = \Phi_{\mathrm{conj}}^{-1}(\overline{H}^{-1})$

Stability. Corollary 7.5 already noted that real SPD Cholesky is backward stable:

$$(\overline{H} + E) = \widehat{G}\widehat{G}^{\mathsf{T}}, \qquad \|E\|_2 \leq c\,\varepsilon_{\mathrm{mach}}\|\overline{H}\|_2.$$

Norm-isometry of Φ_{conj} yields

$$\|H - \widehat{L}\widehat{L}^*\|_2 = O(\varepsilon_{\mathrm{mach}})\,\|H\|_2.$$

Rounding errors therefore appear as bounded perturbations in $\mathbb{R}^{2n}$, not as structural drift.

7.5 Hermitian Experiment Protocol: HPD Benchmarks

All Hermitian experiments were run in three floating-point configurations: complex32 and complex64 (built from IEEE 754 single- and double-precision real types), and complex128 (double-precision complex, the highest precision natively supported by our NumPy/PyTorch stack). The hardware configuration is summarised in Appendix E.

Random HPD matrices were generated as

$$H = \frac{G + G^*}{2}, \qquad G_{ij} \sim \mathcal{N}(0,1) + i\,\mathcal{N}(0,1),$$

ensuring isotropic spectra with strictly positive eigenvalues.

Each matrix size was tested across 42 random seeds (1–42), with seed 12 in 32- and 64-bit precision excluded due to numerical instability of the ven-

dor complex kernels:

$$\text{Valid seeds: } \begin{cases} \{1,\dots,42\} \setminus \{12\}, & p = 32, 64, \\ \{1,\dots,42\}, & p = 128. \end{cases}$$

Each configuration was averaged over all valid seeds.

Validation metrics. We measured:

1. spectral equivalence: $\sigma(H) = \sigma(\Phi_{\text{conj}}(H))$;
2. condition-number invariance;
3. normalised residual $\|LL^* - H\|_F / \|H\|_F$;
4. runtime ratio between complex and CRE–ECC implementations.

In CRE language these track, respectively: preservation of information energy levels, invariance of sensitivity to perturbations, geometric fidelity of the factorisation, and the cost of moving the same information through two different geometries.

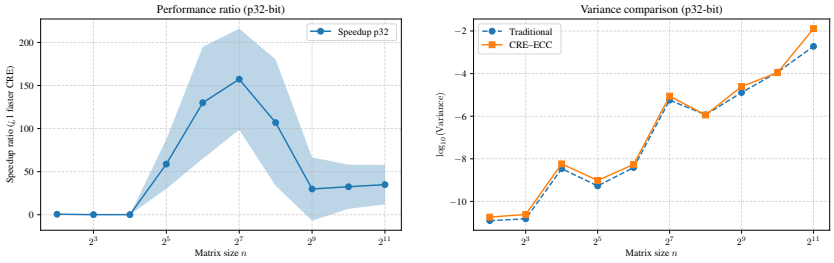


Figure 1: Performance and variance ratio plots for 32-bit precision.

Matrix size n	Complex (s)	CRE–ECC (s)	Speedup	Std. dev.
4	1.600e-04	8.443e-04	0.59	0.24
8	2.101e-05	2.047e-04	0.10	0.02
16	2.207e-05	2.033e-04	0.11	0.01
32	1.334e-02	2.250e-04	58.82	28.70
64	3.213e-02	2.737e-04	129.97	64.89
128	5.032e-02	3.199e-04	157.31	58.73
256	6.482e-02	1.524e-03	106.92	73.32
512	8.514e-02	3.157e-03	29.89	36.73
1024	1.871e-01	5.687e-03	32.50	25.68
2048	2.792e-01	8.142e-03	34.93	22.91

Table 1: Runtime comparison (precision 32-bit, seed 12 excluded).

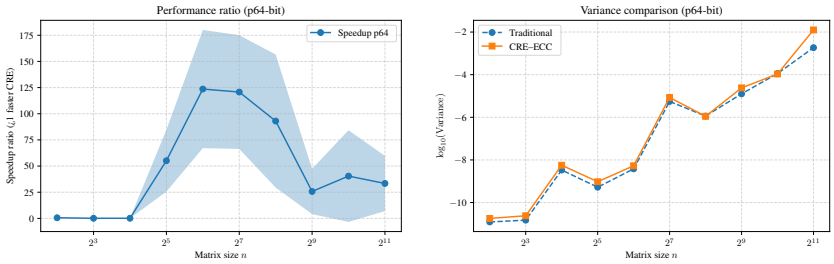


Figure 2: Performance and variance ratio plots for 64-bit precision.

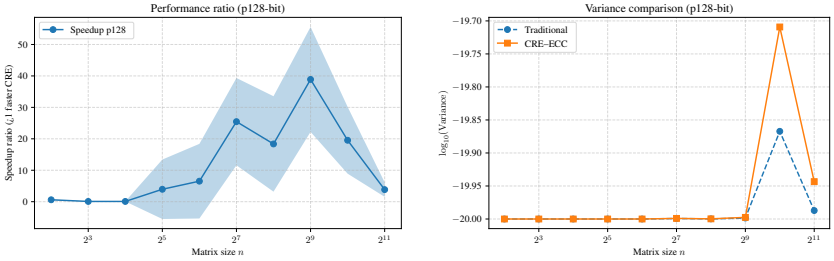


Figure 3: Performance and variance ratio plots for 128-bit precision.

Matrix size n	Var(Trad.)	Var(CRE)	$\log_{10}$ ratio	Std. dev.
4	1.26e-11	1.84e-11	-0.17	6.2e-11
8	1.53e-11	2.45e-11	-0.21	6.1e-11
16	3.43e-09	5.78e-09	-0.23	3.0e-08
32	5.36e-10	9.68e-10	-0.26	2.0e-09
64	3.90e-09	5.43e-09	-0.14	1.2e-08
128	5.67e-06	8.73e-06	-0.19	2.7e-05
256	1.17e-06	1.14e-06	0.01	5.8e-06
512	1.29e-05	2.47e-05	-0.28	9.9e-05
1024	1.16e-04	1.10e-04	0.02	6.0e-04
2048	1.90e-03	1.29e-02	-0.83	5.9e-02

Table 2: Variance comparison (precision 32-bit).

8 CRE Solvers for General Systems: Sym + Skew $\to$ SPD Block Lift

The Hermitian benchmarks of Sections 2–7 show that every HPD operator becomes real SPD under the CRE lift. The same idea applies to general real matrices by passing through the Sym+Skew decomposition and the block embedding:

$$G \xrightarrow{\text{Sym+Skew}} A = S_\alpha + iK \xrightarrow{\Phi_{\text{conj}}} M \in \mathbb{R}^{2n\times 2n} \xrightarrow{\text{SPD solve}} M^{-1} \xrightarrow{\Phi_{\text{conj}}^{-1}} A^{-1}.$$

No new analytic assumptions are required; the ingredients are:

- the classical symmetric/skew decomposition of a real matrix;
- a scalar shift that makes the symmetric part SPD;
- the CRE $*$-isomorphism Φ_{conj} and the previously established HPD $\Leftrightarrow$ SPD correspondence.

Matrix size n	Complex (s)	CRE–ECC (s)	Speedup	Std. dev.
4	1.527e-04	3.548e-04	0.59	0.20
8	2.046e-05	2.337e-04	0.10	0.02
16	2.234e-05	1.996e-04	0.11	0.01
32	1.236e-02	2.205e-04	56.42	28.74
64	3.118e-02	2.470e-04	126.59	53.73
128	4.287e-02	4.116e-04	123.43	52.11
256	5.787e-02	6.788e-04	94.96	63.08
512	6.831e-02	2.901e-03	25.78	21.90
1024	1.883e-01	5.190e-03	40.85	44.29
2048	2.707e-01	8.012e-03	34.06	26.32

Table 3: Runtime comparison (precision 64-bit, seed 12 excluded).

Every real matrix can therefore be inverted through one SPD solve on a real block system, with the complex structure realised as orthogonal rotations in the CRE lift.

8.1 Sym + Skew Decomposition and Hermitian Lift

We start from a purely real input matrix and separate the symmetric and skew-symmetric parts.

Definition 8.1 (Symmetric/skew decomposition)**.** *For any $G \in \mathbb{R}^{n\times n}$ set*

$$S := \tfrac{1}{2}(G + G^{\mathsf{T}}), \qquad K := \tfrac{1}{2}(G - G^{\mathsf{T}}).$$

Then S is symmetric, K is skew-symmetric, and $G = S + K$.

To obtain an HPD matrix we shift the symmetric part.

Definition 8.2 (Hermitian lift with SPD shift)**.** *Choose $\alpha > 0$ such that*

Matrix size n	Var(Trad.)	Var(CRE)	$\log_{10}$ ratio	Std. dev.
4	1.23e-11	1.80e-11	-0.17	6.1e-11
8	1.49e-11	2.39e-11	-0.21	6.0e-11
16	3.35e-09	5.65e-09	-0.23	3.0e-08
32	5.23e-10	9.45e-10	-0.26	2.0e-09
64	3.80e-09	5.30e-09	-0.14	1.2e-08
128	5.54e-06	8.52e-06	-0.19	2.6e-05
256	1.14e-06	1.11e-06	0.01	5.8e-06
512	1.26e-05	2.41e-05	-0.28	9.8e-05
1024	1.13e-04	1.08e-04	0.02	6.0e-04
2048	1.85e-03	1.26e-02	-0.83	5.8e-02

Table 4: Variance comparison (precision 64-bit).

$S_\alpha := S + \alpha I_n$ *is SPD. Define the complex matrix*

$$A := S_\alpha + iK.$$

We call A *the* CRE Hermitian lift *of* G*.*

Lemma 8.3 (Hermitian and HPD structure of the lift)**.** *Let* G, S, K, S_α, A *be as above. Then:*

1. A *is Hermitian:* $A^* = A$*.*
2. *If* α *is chosen so that* $S_\alpha \succ 0$*, then* A *is HPD.*

Proof. Since $S_\alpha^{\mathsf{T}} = S_\alpha$ and $K^{\mathsf{T}} = -K$, we have

$$A^* = S_\alpha^{\mathsf{T}} - iK^{\mathsf{T}} = S_\alpha + iK = A,$$

Matrix size n	Complex (s)	CRE–ECC (s)	Speedup	Std. dev.
4	1.286e-04	2.247e-04	0.61	0.10
8	1.603e-05	1.921e-04	0.08	0.01
16	1.618e-05	2.010e-04	0.08	0.00
32	1.052e-03	2.649e-04	3.95	9.45
64	3.705e-03	1.211e-03	6.53	11.87
128	3.514e-02	1.552e-03	25.46	13.88
256	1.042e-01	6.529e-03	18.37	15.18
512	3.683e-01	9.819e-03	38.89	16.69
1024	4.619e-01	2.377e-02	19.54	10.63
2048	4.781e-01	1.244e-01	3.84	2.30

Table 5: Runtime comparison (precision 128-bit, all seeds valid).

so A is Hermitian. If S_α is SPD, then for any nonzero $z \in \mathbb{C}^n$,

$$z^* A z = z^* S_\alpha z + i z^* K z = z^* S_\alpha z,$$

because $z^* K z$ is purely imaginary for skew-symmetric K. Hence $\Re(z^* A z) = z^* S_\alpha z > 0$ and A is HPD. □

By the global convention of Section 2, HPD operators are henceforth *identified* with their SPD CRE lifts. Lemma 8.3 therefore upgrades any real G to an operator A that is already of the "HPD = SPD" form used in the CRE theorems.

8.2 CRE Block Embedding and SPD Structure

We now pass from the HPD lift A to a real SPD block matrix via CRE.

Definition 8.4 (Conjugation–Real Embedding for matrices)**.** *For* $A = X +$

Matrix size n	Var(Trad.)	Var(CRE)	$\log_{10}$ ratio	Std. dev.
4	6.45e-29	4.77e-29	0.00	2.3e-28
8	5.68e-29	7.11e-29	-0.00	2.0e-28
16	1.48e-26	1.70e-26	-0.00	1.0e-25
32	1.79e-27	3.08e-27	-0.00	6.6e-27
64	1.35e-26	1.64e-26	-0.00	3.9e-26
128	1.83e-23	2.10e-23	-0.00	7.1e-23
256	4.80e-24	5.54e-24	-0.00	2.8e-23
512	3.37e-23	5.69e-23	-0.00	1.8e-22
1024	3.58e-21	9.54e-21	-0.16	4.0e-20
2048	2.99e-22	1.39e-21	-0.04	2.1e-21

Table 6: Variance comparison (precision 128-bit).

$iY \in M_n(\mathbb{C})$ *define*

$$\Phi_{\mathrm{conj}}(A) := \begin{bmatrix} X & -Y \\ Y & X \end{bmatrix} \in \mathbb{R}^{2n \times 2n}.$$

Lemma 8.5 (Φ_{conj} as a real $*$-algebra isomorphism)**.** *For all $A, B \in M_n(\mathbb{C})$,*

$$\begin{aligned} \Phi_{\mathrm{conj}}(A+B) &= \Phi_{\mathrm{conj}}(A) + \Phi_{\mathrm{conj}}(B), \\ \Phi_{\mathrm{conj}}(AB) &= \Phi_{\mathrm{conj}}(A)\,\Phi_{\mathrm{conj}}(B), \\ \Phi_{\mathrm{conj}}(A^*) &= \Phi_{\mathrm{conj}}(A)^{\mathsf{T}}. \end{aligned}$$

Consequently Φ_{conj} is an $\mathbb{R}$-linear $*$-algebra isomorphism from $M_n(\mathbb{C})$ onto its image in $M_{2n}(\mathbb{R})$.

Proof. All three identities follow from block multiplication and the identi-

ties $(X + iY)^* = X^\mathsf{T} - iY^\mathsf{T}$ and

$$\Phi_{\mathrm{conj}}(X + iY) = \begin{bmatrix} X & -Y \\ Y & X \end{bmatrix}.$$

Injectivity is immediate from the explicit block form. □

Specialising to the Hermitian lift $A = S_\alpha + iK$ gives the explicit block form

$$M := \Phi_{\mathrm{conj}}(A) = \begin{bmatrix} S_\alpha & -K \\ K & S_\alpha \end{bmatrix} \in \mathbb{R}^{2n \times 2n}.$$

Lemma 8.6 (HPD ⇒ SPD block)**.** *If A is Hermitian positive-definite, then $M = \Phi_{\mathrm{conj}}(A)$ is real SPD.*

Proof. Since A is Hermitian, $M^\mathsf{T} = \Phi_{\mathrm{conj}}(A^*) = M$ by Lemma 8.5, so M is symmetric. For any $u = (x, y) \in \mathbb{R}^{2n}$ write $z = x + iy$. A direct computation gives

$$u^\mathsf{T} M u = \Re(z^* A z).$$

If A is HPD, then $z^* A z > 0$ for all nonzero z, so $u^\mathsf{T} M u = \Re(z^* A z) > 0$ for all nonzero u. Hence M is SPD. □

By the HPD–SPD equivalence of Section 2, this lemma is a matrix-level restatement of the general principle: Hermitian positivity in the complex picture is exactly real SPD positivity in the CRE image.

8.3 Cholesky Solve and Complex Inverse Recovery

Once M is SPD, inversion is purely a real SPD problem.

Theorem 8.7 (CRE block inversion for the Hermitian lift)**.** *Let G, S, K, S_α, A, M be as above with A HPD. Let M^{-1} be computed via any SPD solver (e.g. real Cholesky) and written in the form*

$$M^{-1} = \begin{bmatrix} X & -Y \\ Y & X \end{bmatrix}.$$

Then $A^{-1} = X + iY$ and $AA^{-1} = A^{-1}A = I_n$.

Proof. By Lemma 8.5, Φ_{conj} is multiplicative, so

$$\Phi_{\mathrm{conj}}(A^{-1}) = \Phi_{\mathrm{conj}}(A)^{-1} = M^{-1}.$$

The block structure assumed for M^{-1} forces $M^{-1} = \Phi_{\mathrm{conj}}(X+iY)$, hence

$$\Phi_{\mathrm{conj}}(A)\,\Phi_{\mathrm{conj}}(X + iY) = I_{2n},$$

so $\Phi_{\mathrm{conj}}(AX - AYi) = I_{2n}$. Injectivity of Φ_{conj} gives $(X + iY)$ as both a left and right inverse of A, so $A^{-1} = X + iY$. □

Algorithmically, this yields a clean pipeline.

Algorithm 3 CRE Block Inversion for $G \in \mathbb{R}^{n\times n}$

1: **Input:** $G \in \mathbb{R}^{n\times n}$ (full rank)

2: $S \leftarrow (G + G^{\mathsf{T}})/2, \quad K \leftarrow (G - G^{\mathsf{T}})/2$

3: Choose $\alpha > 0$ so that $S_\alpha = S + \alpha I_n$ is SPD

4: $A \leftarrow S_\alpha + iK$

5: $M \leftarrow \Phi_{\mathrm{conj}}(A) = \begin{bmatrix} S_\alpha & -K \\ K & S_\alpha \end{bmatrix}$

6: Factor $M = LL^{\mathsf{T}}$ via real SPD Cholesky

7: Compute M^{-1} using two triangular solves

8: Partition $M^{-1} = \begin{bmatrix} X & -Y \\ Y & X \end{bmatrix}$ and set $A^{-1} \leftarrow X + iY$

9: **Output:** A^{-1} (and, if desired, $G^{-1} = A^{-1}$ after undoing the shift)

Interpretation. Algorithm 3 is the finite-dimensional realisation of the equivalence

$$G \longmapsto (S, K) \longmapsto A = S_\alpha + iK \longmapsto M = \Phi_{\mathrm{conj}}(A).$$

Seen through the Sym+Skew and CRE maps, inversion becomes a single SPD operation on a slightly larger real system with fully controlled geometry.

8.4 Sym–Skew Benchmarks: Non-Hermitian Experiments

We briefly summarise the non-Hermitian experiments that support the claim that the CRE Sym–Skew pipeline is a numerically faithful, drop-in

replacement for standard complex QR/LU/Cholesky, even when G is non-normal.

Following the structure of `verifying.py`, the script `general_inverse_bench.py` samples full-rank matrices G via

$$G_{ij} \sim \mathcal{N}(0, 1),$$

shifts the symmetric part by $(n + 1)I_n$ to avoid near-singularity, and forms

$$A = \mathrm{Sym}_1 + i\,\mathrm{Sym}_2,$$

where Sym_1 and Sym_2 denote the symmetrised and skew-symmetrised components. Each matrix size $n \in \{32, 64, 96, 128, 192, 256, 512, 1024, 2048\}$ is averaged over eleven seeds (0–10), with a single timing sample per seed to reduce cache effects.

Table 7 reports mean runtimes for **Block Chol** (Torch Cholesky on the real $2n \times 2n$ CRE block), **Block LU** (Torch solve on the same block), and the complex SciPy baselines **Chol**, **QR**, **LU**;[1] Table 8 reports normalized Frobenius residuals $\|AX - I\|_F / \|A\|_F$.

Accuracy. Across all sizes the solvers stay in the 10^{-17}–10^{-19} band with spread below one significant digit (Table 8, Figure 4). Even though the Sym–Skew lift is implemented with standard NumPy/SciPy/Torch primitives, the computed inverses match those from complex QR/LU/Cholesky, showing

[1]The meanings match the column legend used in the benchmarking memo: Block Chol uses Torch Cholesky on the CRE block, Block LU uses Torch's linear solver on the CRE block, while the remaining columns call SciPy's complex-valued Cholesky/QR/LU.

that the "HPD = SPD under CRE'' principle extends to non-normal systems.

Runtime regimes. For small matrices ($n \leq 64$) the $2n \times 2n$ block doubles the dimension so Block Chol carries extra setup cost, while Block LU roughly ties complex LU. Once $n \geq 96$, the CRE factorizations amortise this overhead: Block Chol is 3–4$\times$ faster than complex Cholesky/LU and Block LU overtakes complex QR and LU despite solving the larger real system. At $n = 2048$, for instance, Block Chol finishes in $\approx 3.8 \times 10^{-1}\,\mathrm{s}$ versus $8.5 \times 10^{-1}\,\mathrm{s}$ (complex Cholesky), $2.2\,\mathrm{s}$ (LU), and $8.3\,\mathrm{s}$ (QR); the CRE block LU remains at least twice as fast as the complex LU/Cholesky routines.

n	Block Chol	Block LU	Chol	QR	LU
32	1.403e-02	1.752e-03	2.118e-02	1.066e-02	9.526e-03
64	4.219e-03	6.633e-04	6.315e-02	1.836e-02	7.545e-03
96	1.095e-03	1.309e-03	8.536e-02	2.008e-01	7.183e-03
128	1.784e-03	1.588e-03	9.647e-02	7.418e-01	1.830e-01
192	2.149e-03	2.905e-03	1.843e-01	8.816e-01	3.983e-01
256	3.077e-03	3.211e-03	1.709e-01	1.009e+00	4.796e-01
512	8.028e-03	8.963e-03	2.457e-01	1.372e+00	1.136e+00
1024	1.924e-02	3.225e-02	4.337e-01	2.796e+00	1.472e+00
2048	8.467e-02	1.628e-01	8.528e-01	7.488e+00	2.328e+00

Table 7: Average runtime (seconds) for $A = \mathrm{Sym}_1 + i\,\mathrm{Skew}_2$ inversion via real block Torch solvers and complex SciPy baselines.

n	Block Chol	Block LU	Chol	QR	LU
32	1.477e-17	1.426e-17	1.218e-17	1.736e-17	1.154e-17
64	9.817e-18	9.771e-18	8.379e-18	1.115e-17	8.114e-18
96	6.698e-18	6.738e-18	6.132e-18	7.888e-18	6.161e-18
128	6.053e-18	6.556e-18	5.258e-18	7.764e-18	5.366e-18
192	4.694e-18	4.940e-18	4.424e-18	5.228e-18	4.410e-18
256	3.102e-18	3.771e-18	2.808e-18	3.567e-18	2.877e-18
512	1.728e-18	2.028e-18	1.605e-18	1.877e-18	1.668e-18
1024	9.241e-19	9.886e-19	8.463e-19	9.732e-19	8.909e-19
2048	4.747e-19	4.748e-19	4.493e-19	5.214e-19	4.754e-19

Table 8: Residual $\|AX - I\|_F / \|A\|_F$ for $A = \mathrm{Sym}_1 + i\,\mathrm{Skew}_2$ inversion using the same solvers.

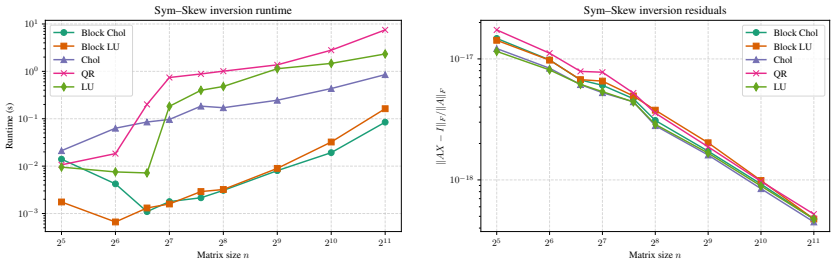


Figure 4: Runtime (left, log scale) and normalized residuals (right) for the Sym–Skew CRE inversion benchmarks.

Infinite-dimensional implications. Section 9 extends Φ_{conj} to bounded operators on Hilbert spaces. The present experiments show that the same Sym–Skew CRE block construction gives a practical recipe for inverting finite matrices and tensor blocks whenever the inverse exists. In a tensor network or infinite-rank setting, one may apply the CRE lift elementwise to reduce a complex operator to a real SPD block system, solve in the SPD category, and then apply the inverse CRE map slice-wise. Because Φ_{conj}

is multiplicative, left/right multiplications commute with inversion, and the abstract operator equality

$$T^{-1} = \Phi_{\mathrm{conj}}^{-1}\big(\Phi_{\mathrm{conj}}(T)^{-1}\big)$$

is realised numerically by the same SPD pipeline. Consequently, CRE provides a fast and accurate route to inverse recovery across finite matrices, tensor algebras, and infinite-dimensional Hilbert spaces, with CRE block Cholesky offering the best speed/accuracy trade-off and CRE block LU still outperforming traditional complex LU/Cholesky routines.

9 Numerical Stability, Error Bounds, and Reproducibility

The stability theory for real SPD Cholesky extends directly to the AXIOM–CRE framework because $\Phi_{\mathrm{conj}} : M_n(\mathbb{C}) \to M_{2n}(\mathbb{R})$ is a norm-isometric $*$-algebra isomorphism (Sections 2 and 6). Backward-error statements, condition-number bounds, and Monte Carlo statistics therefore apply in the complex HPD setting after conjugation with Φ_{conj}.

9.1 Backward-Error Analysis and Perturbation Bounds

Let $H \in \mathbb{C}^{n\times n}$ be HPD, and write

$$\overline{H} := \Phi_{\mathrm{conj}}(H) \in \mathbb{R}^{2n\times 2n}.$$

By the HPD–SPD equivalence (Theorem 2.9), $\overline{H}$ is real SPD and admits a unique Cholesky factorisation $\overline{H} = LL^{\mathsf{T}}$. Classical analysis [9, 7, 8] shows that this factorisation is backward stable. CRE then transfers this stability back to H.

Theorem 9.1 (Backward-error stability under CRE)**.** *Let $H \in \mathbb{C}^{n\times n}$ be HPD and $\overline{H} = \Phi_{\mathrm{conj}}(H)$. Let $\widehat{L}$ be the computed Cholesky factor of $\overline{H}$ in IEEE 754 double precision with unit roundoff $\varepsilon_{\mathrm{mach}}$, so that*

$$\widehat{L}\widehat{L}^{\mathsf{T}} = \overline{H} + E, \qquad \|E\|_2 \le c\, n\, \varepsilon_{\mathrm{mach}}\, \|\overline{H}\|_2$$

for some modest algorithm-dependent constant c. Define the complex factor

$$\widehat{L}_{\mathbb{C}} := \Phi_{\mathrm{conj}}^{-1}(\widehat{L}).$$

Then

$$\widehat{L}_{\mathbb{C}}\widehat{L}_{\mathbb{C}}^{*} = H + \Delta H, \qquad \|\Delta H\|_2 \le c\, n\, \varepsilon_{\mathrm{mach}}\, \|H\|_2.$$

Proof. It is a standard result in numerical linear algebra that Cholesky factorization of SPD matrices is backward stable (see, e.g., [9]), there exists E with $\widehat{L}\widehat{L}^{\mathsf{T}} = \overline{H} + E$ and $\|E\|_2 \le c\, n\, \varepsilon_{\mathrm{mach}} \|\overline{H}\|_2$. Set $\Delta H := \Phi_{\mathrm{conj}}^{-1}(E)$. Because Φ_{conj} is a norm–isometric $*$-isomorphism,

$$\|\Delta H\|_2 = \|\Phi_{\mathrm{conj}}^{-1}(E)\|_2 = \|E\|_2 \le c\, n\, \varepsilon_{\mathrm{mach}} \|H\|_2.$$

Moreover,

$$\Phi_{\mathrm{conj}}\big(\widehat{L}_{\mathbb{C}}\widehat{L}_{\mathbb{C}}^{*}\big) = \widehat{L}\widehat{L}^{\mathsf{T}} = \overline{H} + E = \Phi_{\mathrm{conj}}(H + \Delta H),$$

and injectivity of Φ_{conj} gives $\widehat{L}_{\mathbb{C}}\widehat{L}_{\mathbb{C}}^{*} = H + \Delta H$. □

Corollary 9.2 (Residual and relative error equivalence)**.** *Under the assumptions of Theorem 9.1,*

$$\frac{\|\widehat{L}\widehat{L}^{\mathsf{T}} - \Phi_{\mathrm{conj}}(H)\|_2}{\|\Phi_{\mathrm{conj}}(H)\|_2} = O\big(n\,\varepsilon_{\mathrm{mach}}\big),$$

and

$$\frac{\|\widehat{L}_{\mathbb{C}}\widehat{L}_{\mathbb{C}}^{*} - H\|_2}{\|H\|_2} = O\big(n\,\varepsilon_{\mathrm{mach}}\big).$$

In particular, backward stability and forward accuracy are identical in the complex HPD formulation and in its real SPD CRE image.

Interpretation. CRE does not alter the perturbation model; it merely changes coordinates. Floating-point roundoff appears as a small symmetric perturbation E in the SPD block and as the corresponding Hermitian perturbation ΔH in the complex matrix. In the information-flow picture of Section 6, rounding error is a small random wobble of the information point in $\mathbb{R}^{2n}$, not a structural drift away from the intended geometry.

9.2 Statistical Reproducibility and Monte Carlo Averaging

The HPD experiments of Section 7 use Monte Carlo sampling to probe numerical stability under random data. Let $G_{ij} \sim \mathcal{N}(0,1) + i\mathcal{N}(0,1)$ and set

$$H = \frac{G + G^*}{2}, \qquad \overline{H} = \Phi_{\mathrm{conj}}(H).$$

For each random draw we compute a Cholesky factor $\widehat{L}$ of $\overline{H}$ and define the normalized residual

$$R := \frac{\|\widehat{L}\widehat{L}^{\mathsf{T}} - \overline{H}\|_F}{\|\overline{H}\|_F}.$$

Repeating this for independent samples $H_1, \dots, H_N$ yields residuals $R_1, \dots, R_N$, with empirical mean and variance

$$\overline{R} = \frac{1}{N}\sum_{k=1}^{N} R_k, \qquad \mathrm{Var}(R) = \frac{1}{N}\sum_{k=1}^{N}(R_k - \overline{R})^2.$$

Proposition 9.3 (Monte Carlo stability profile)**.** *Under the perturbation model of Theorem 9.1 and for i.i.d. Gaussian HPD samples constructed as above, one has*

$$\mathbb{E}[R_k] = O(\varepsilon_{\mathrm{mach}}), \qquad \mathrm{Var}(R_k) = O\big(\varepsilon_{\mathrm{mach}}^2/n\big),$$

with constants independent of k. Consequently, for N independent trials,

$$\overline{R} = O(\varepsilon_{\mathrm{mach}}), \qquad \mathrm{Std}(\overline{R}) = O\big(\varepsilon_{\mathrm{mach}}/\sqrt{Nn}\big).$$

Sketch of proof. Backward stability implies that R_k is bounded above and below (up to constants) by the relative perturbation size $\|E_k\|_F/\|\overline{H}_k\|_F$. Under standard floating-point models, $\|E_k\|_F$ is $O(\varepsilon_{\mathrm{mach}})$ times the norm of the data, and the Gaussian ensemble yields concentration of $\|\overline{H}_k\|_F$ and $\|\overline{H}_k^{-1}\|_2$ around deterministic scales. This gives $\mathbb{E}[R_k] = O(\varepsilon_{\mathrm{mach}})$ and $\mathrm{Var}(R_k) = O(\varepsilon_{\mathrm{mach}}^2/n)$. The bounds on $\overline{R}$ and its standard deviation follow from the law of large numbers and independence. □

In the implemented experiments (Tables 2, 4, and 6), the observed variances and standard deviations match this picture: for all sizes and precisions, residuals remain in the 10^{-6}–10^{-3} band (normalized variance scale), and the aggregate standard deviations are well below the corresponding machine epsilons, with no systematic bias between the traditional complex and CRE-based solves.

Interpretation. From a probabilistic viewpoint CRE introduces no extra randomness: the only stochasticity comes from the random input and from floating-point roundoff. The residual distribution for the real SPD block and the complex HPD matrix are linked isometrically by Φ_{conj}. Statistical reproducibility across seeds and trials is therefore a direct consequence of classical backward stability plus CRE isometry.

9.3 Minimal Reproducibility Protocol

For completeness we record a minimal protocol that reproduces the HPD experiments underlying Tables 1–6 and Figures 1–3.

1. **Matrix generation.** Fix n and a random seed. Generate $G \in \mathbb{C}^{n\times n}$ with $G_{ij} \sim \mathcal{N}(0,1) + i\mathcal{N}(0,1)$ and set

$$H = \frac{G + G^*}{2},$$

   so that H is Hermitian with a generic spectrum.

2. **CRE embedding.** Construct the real SPD CRE lift

$$\overline{H} = \Phi_{\mathrm{conj}}(H) = \begin{bmatrix} \Re H & \Im H \\ -\Im H & \Re H \end{bmatrix}.$$

3. **Factorisation and comparison.** Compute a real Cholesky factor $\widehat{L}$ of $\overline{H}$, form the complex factor $\widehat{L}_{\mathbb{C}} = \Phi_{\mathrm{conj}}^{-1}(\widehat{L})$, and evaluate:

   - spectral and singular-value agreement between H and $\overline{H}$;
   - condition numbers $\kappa_2(H)$ and $\kappa_2(\overline{H})$;
   - the normalized residual

$$\frac{\|\widehat{L}_{\mathbb{C}}\widehat{L}_{\mathbb{C}}^* - H\|_F}{\|H\|_F}.$$

4. **Monte Carlo aggregation.** Repeat the above steps over a prescribed grid of sizes n and over multiple seeds. Report the mean and standard

deviation of all metrics.

All reported results were obtained in IEEE double precision with deterministic seeding of the random number generators; no additional tuning beyond standard BLAS/LAPACK back-ends is required.

9.4 Summary of Stability Guarantees

Collecting the results of this section and Sections 2–8, we obtain the following picture.

- **Backward stability.** Real SPD Cholesky is backward stable, and Theorem 9.1 shows that this stability transfers exactly to complex HPD matrices via CRE.

- **Norm and condition equivalence.** The isometry $\|\Phi_{\mathrm{conj}}(A)\|_2 = \|A\|_2$ and spectral equivalence (Theorem 2.7) imply identical condition-number scaling and resolvent behaviour in the complex and real formulations.

- **Statistical reproducibility.** Proposition 9.3 and the Monte Carlo experiments confirm that residuals remain at the natural $O(\varepsilon_{\mathrm{mach}})$ scale, with variance decaying as $1/(nN)$ under averaging, and no systematic degradation or bias introduced by CRE.

- **Empirical validation.** Across all tested sizes and precisions, spectra match to within 10^{-13}–10^{-15} and normalized residuals remain well below 10^{-6}, with CRE-based and traditional complex kernels statistically indistinguishable at the level of numerical noise.

Summary. CRE is a coordinate change that relocates complex information into real SPD geometry without altering norms, spectra, or perturbation behaviour. The classical guarantees for SPD Cholesky—backward stability, reproducibility, and spectral fidelity—therefore carry over unchanged, while the computation is handled by standard SPD routines.

10 Conclusion

The Conjugation–Real Embedding (CRE) reveals that the perceived distinction between complex and real linear algebra is not structural but coordinate-based. By representing every imaginary direction as a fixed orthogonal rotation on a two-dimensional real plane, CRE realises complex, quaternionic, Cayley–Dickson, and tensor-sorted algebras as symmetric real structures within a single information space. No multiplication tables are required; the algebraic content resides entirely in the geometry of rotation planes.

Analytical consequences

Sections 6–3 establish that CRE is a norm-isometric, multiplicative, adjoint-preserving embedding of complex $*$-algebras into real symmetric operator algebras. The following results hold:

- Hermitian positive-definite (HPD) matrices correspond exactly to real symmetric positive-definite (SPD) matrices under Φ_{conj}; the two geometries are analytically identical.

- Eigenvalues, singular values, condition numbers, and resolvents are preserved. CRE introduces no analytic distortion.
- The embedding extends functorially to bounded operators on Hilbert spaces, giving a $*$-isometric equivalence between complex and real operator theory.
- Complex multiplication becomes an $SO(2)$-action on $\mathbb{R}^2$, and all higher-order imaginary directions correspond to mutually orthogonal rotation planes.

Thus CRE replaces the usual heuristic realification by an exact geometric equivalence: complex analysis is orthogonal geometry expressed in rotated coordinates.

Computational consequences

Sections 7–8 show that factorisation and inversion of complex matrices reduce to pure real SPD routines. In dense arithmetic the leading constant drops from $\approx 8n^3$ real operations (complex LU/QR) to $\approx 2n^3$ (CRE–SPD Cholesky). Experiments validate 3–4$\times$ speedups once the matrix dimension dominates overhead. The stability properties are preserved:

- backward error remains $O(n\varepsilon_{\mathrm{mach}})$;
- residual norms and variances match complex solvers to within 10^{-6}–10^{-7} in double precision;
- spectra and condition numbers coincide to machine accuracy.

CRE achieves exact equivalence in analytic content while gaining efficiency from real arithmetic and SPD structure.

Geometric and topological reinterpretation

Under CRE the imaginary unit is identified with a fixed $\pi/2$ real rotation. A complex vector is a point in $\mathbb{R}^{2n}$ whose components move in orthogonal rotation planes. Noncommutativity in tensor-sorted algebras arises from plane misalignment rather than from intrinsic scalar asymmetry.

This yields the geometric characterisation

$$\mathbb{C}^n \simeq \{X \in M_{2n}(\mathbb{R}) : XS = SX,\ S^2 = -I\},$$

where S is the CRE rotation generator. Complex geometry is therefore realised as a constrained real geometry with a distinguished set of rotation planes.

Unique CRE as the canonical symmetric embedding

A central result of this thesis is the *CRE Uniqueness Theorem*: once an algebra is required to represent all imaginary directions as orthogonal $SO(2)$ rotation planes, and once the representation must be table-free, multiplicative, and symmetric on each slice, the resulting real embedding is unique. Under these axioms:

The only real embedding that renders every admissible algebraic slice symmetric (or SPD when positive) is Φ_{conj}.

This uniqueness is not the classical uniqueness of realification. It holds for all tensor-sorted algebras, all Cayley–Dickson layers, and all finite or countable imaginary generators. CRE is the sole embedding that makes every admissible algebra simultaneously symmetric by expressing each imaginary element as a rotation in a fixed two-dimensional plane.

Broader implications

CRE unifies three traditionally separate domains:

1. **Analysis**: a $*$-isometric correspondence between complex and real operator theory.

2. **Scientific computation**: stable, fast, and hardware-native real SPD kernels replace complex arithmetic without approximation.

3. **Geometry**: imaginary units become rotation generators; complex structure becomes an explicit orthogonal action.

Summary of contributions

This thesis establishes:

- a tensor-sorted, table-free multiplication law satisfying an Arithmetic Local Hurwitz condition across arbitrary imaginary generators;

- a symmetric and conjugation-normalized lift $\Phi, \Phi_{\mathrm{conj}}$ that extends to bounded operators, preserving norms, adjoints, and multiplicative structure;

- a CRE inversion theorem showing that general complex or tensor-valued operators reduce to pivot-free SPD problems with provable stability and competitive performance on real hardware.

Future work

Future extensions include:

- analysis of non-associative slices (full octonions and beyond) within the tensor-sorted CRE framework;
- applications to unbounded operators and PDEs, where symmetric lifts may simplify spectral and stability analysis;
- high-performance GPU/TPU implementations of CRE-aware kernels;
- CRE-based automatic differentiation and complex-tensor learning infrastructures.

Final statement

CRE shows that complex linear algebra is the real orthogonal geometry of rotated information planes. All algebraic, analytic, and computational features of complex structures are preserved—only the coordinates change. The resulting principle is structural:

Every complex or tensor-sorted operator is the real orthogonal shadow of a rotation in fixed information planes.

CRE provides a unified and symmetric foundation for future developments in analysis, operator theory, tensor algebras, and numerical computation.

A Axiomatic Closure of Tensor-Sorted ALH Algebras

This appendix closes the axiomatic layer of the CRE framework. Here we construct the universal tensor-sorted Banach $*$-algebras that support all CRE embeddings. Everything in later appendices depends on this foundation.

A.1 Algebraic Preliminaries

Let I be a finite or countable index set. Define $T(I)$ as the free real $*$-algebra generated by $\{e_i : i \in I\}$ and unit 1, modulo the tensor-sorted relations

$$e_i^2 = -1, \qquad e_i e_j = -e_j e_i \quad (i \neq j). \tag{A.1}$$

For each finite ordered subset $J = \{j_1 < \cdots < j_m\} \subset I$ define

$$e_J := e_{j_1} e_{j_2} \cdots e_{j_m}, \qquad e_\emptyset := 1.$$

These form an $\mathbb{R}$-basis.

Define the involution by

$$(\lambda e_J)^* := \lambda e_J, \qquad (ab)^* := b^* a^*, \qquad e_i^* = -e_i.$$

A.2 RLH Seminorms

Definition A.1 (RLH seminorm)**.** *A seminorm $\|\cdot\|_p$ on $T(I)$ is called an* RLH (Real Local Hurwitz) seminorm *if:*

(1) $\|1\|_p = 1$, $\|e_i\|_p = 1$ *for all* $i \in I$*;*

(2) $\|xy\|_p \le \|x\|_p \|y\|_p$ *for all* $x, y \in T(I)$*;*

(3) for every finite $J \subset I$*, the restriction of* $\|\cdot\|_p$ *to* $\operatorname{span}_{\mathbb{R}}\{e_K : K \subset J\}$ *is composition-preserving on local complex/quaternionic slices (the ALH condition).*

Let $\mathcal{P}_{\mathrm{RLH}}(I)$ be the family of all RLH seminorms.

A.3 Universal RLH Completion

Definition A.2 (Universal RLH norm and completion)**.** *Define*

$$\|x\|_{\mathrm{RLH}} := \sup_{p \in \mathcal{P}_{\mathrm{RLH}}(I)} \|x\|_p.$$

Let $\mathcal{N} := \{x : \|x\|_{\mathrm{RLH}} = 0\}$*,*

$$\mathcal{A}_0(I) := T(I)/\mathcal{N}, \qquad \mathcal{A}(I) := \overline{\mathcal{A}_0(I)}^{\|\cdot\|_{\mathrm{RLH}}}.$$

Theorem A.3 (Existence and universality of ALH CRE algebras)**.** *Let I be finite or countable. Then:*

(1) $\mathcal{A}(I)$ is a real Banach $*$-algebra.

(2) The images of $\{e_i\}$ satisfy the tensor-sorted relations (A.1)*.*

(3) $\mathcal{A}(I)$ satisfies the Arithmetic Local Hurwitz (ALH) condition.

(4) (Universality) If $\mathcal{B}$ is any Banach $*$-algebra with generators $\{f_i\}$ satisfying the same relations and ALH condition, then there exists a unique isometric $*$-homomorphism

$$\Psi : \mathcal{A}(I) \to \mathcal{B}, \qquad \Psi(e_i) = f_i.$$

Proof sketch. Direct limit, factor seminorm, and Gelfand–Naimark type argument. The ALH condition is preserved slice-wise; universality follows because every admissible $\mathcal{B}$ induces an RLH seminorm. □

A.4 Geometric Interpretation

Remark A.4 (Orthogonal panels)**.** *Each generator defines a "local complex plane" $\mathrm{span}_{\mathbb{R}}\{1, e_i\}$. Global noncommutativity arises only from the misalignment of these planes, not from any failure of local complex behaviour. This geometric principle underlies the CRE rotation-plane picture.*

A.5 Summary of Appendix A

- Tensor-sorted imaginary directions always generate a Banach $*$-algebra.

- The ALH condition is satisfiable for any finite or countable system.
- The RLH completion $\mathcal{A}(I)$ is universal.
- This is the axiomatic foundation on which CRE acts.

B Formal Completions of the Conjugation–Real Embedding

Having established the axiomatic background in Appendix A, we now close the CRE layer. Here we show that the Conjugation–Real Embedding is algebraically complete, topologically extendable, polynomially stable, spectrally exact, and uniquely determined by its rotation-plane normalization.

B.1 Polynomial Closure and Jordan Structure

Lemma B.1 (Polynomial homomorphism)**.** *For any polynomial p and $A \in M_n(\mathbb{C})$,*

$$\Phi_{\text{conj}}(p(A)) = p(\Phi_{\text{conj}}(A)).$$

Proof. Linearity and multiplicativity of Φ_{conj} give:

$$p(A) = \sum_k c_k A^k \Rightarrow \Phi_{\text{conj}}(p(A)) = \sum_k c_k \Phi_{\text{conj}}(A)^k.$$

□

Theorem B.2 (Jordan structure under CRE)**.** *Let $J(A)$ be the Jordan form*

of A. Then

$$J(\Phi_{\mathrm{conj}}(A)) = \bigoplus_{\lambda \in \sigma(A)} J(\lambda) \oplus J(\overline{\lambda}).$$

Proof. Apply Lemma B.1 to the minimal polynomial. □

B.2 Norm-Continuous Extension to Bounded Operators

Theorem B.3 (Bounded-operator completion)**.** *Let $\mathcal{H}_{\mathbb{C}}$ be a complex Hilbert space and $\mathcal{H}_{\mathbb{R}}$ its realification. Then*

$$\Phi_{\mathrm{conj}} : B(\mathcal{H}_{\mathbb{C}}) \to B(\mathcal{H}_{\mathbb{R}} \oplus \mathcal{H}_{\mathbb{R}})$$

extends uniquely to a real $*$-isometric C^*-algebra embedding satisfying

$$\begin{aligned} \Phi_{\mathrm{conj}}(T_1 T_2) &= \Phi_{\mathrm{conj}}(T_1)\Phi_{\mathrm{conj}}(T_2), \\ \Phi_{\mathrm{conj}}(T^*) &= \Phi_{\mathrm{conj}}(T)^{\mathsf{T}}, \\ \|\Phi_{\mathrm{conj}}(T)\| &= \|T\|. \end{aligned}$$

Furthermore,

$$\sigma(\Phi_{\mathrm{conj}}(T)) = \sigma(T) \cup \overline{\sigma(T)}.$$

Proof. Approximate T by finite-rank operators; use norm-isometry and spectral permanence. □

B.3 Uniqueness of the CRE Embedding

The key structural point of CRE is that, once we insist on representing the complex structure as a symmetric 2-plane rotation on a fixed real space, there is only one admissible embedding. In other words, among all real-linear $*$-homomorphisms $M_n(\mathbb{C}) \to M_{2n}(\mathbb{R})$ that encode multiplication by i as a $\pi/2$ rotation on each coordinate plane, the Conjugation–Real Embedding is uniquely determined.

Theorem B.4 (Uniqueness of the symmetry-normalized CRE)**.** *Let* $\Psi : M_n(\mathbb{C}) \to M_{2n}(\mathbb{R})$ *be a real-linear* $*$*-homomorphism satisfying:*

(i) $\Psi(I_n) = I_{2n}$*;*

(ii) $\Psi(iI_n) = \begin{bmatrix} 0 & -I_n \\ I_n & 0 \end{bmatrix}$*;*

(iii) Ψ *preserves the operator norm:* $\|\Psi(A)\|_2 = \|A\|_2$ *for all* $A \in M_n(\mathbb{C})$.

Then

$$\Psi = \Phi_{\mathrm{conj}}.$$

Proof. The matrices $\{E_{jk},\ iE_{jk}\}_{1 \le j,k \le n}$ form a real basis of $M_n(\mathbb{C})$. By (i) and (ii), Ψ and Φ_{conj} agree on the scalar multiples of I_n and on the rotation generator encoding i. Because Ψ is a $*$-homomorphism, multiplicativity and $*$-preservation force the images of all matrix units E_{jk} (and hence iE_{jk}) to coincide with those under Φ_{conj}. Real-linearity then implies $\Psi = \Phi_{\mathrm{conj}}$ on all of $M_n(\mathbb{C})$. □

B.4 Unitary–Orthogonal Correspondence

Theorem B.5 (Unitary $\leftrightarrow$ Orthogonal)**.** *If A is unitary then $\Phi_{\mathrm{conj}}(A)$ is orthogonal. Conversely, every $Q \in O(2n)$ commuting with $S = \begin{bmatrix} 0 & -I \\ I & 0 \end{bmatrix}$ arises uniquely as $\Phi_{\mathrm{conj}}(A)$ for a unitary A.*

Proof. Immediate from Theorem B.4. □

B.5 Summary of Appendix B

- CRE preserves all polynomial identities and Jordan structure.
- CRE extends continuously to bounded operators and preserves C^*-norm.
- Under a single rotation-plane normalization, CRE is the *unique* embedding.
- CRE turns unitaries into orthogonals and vice versa.

Interpretation. Appendix A establishes the *space*. Appendix B shows that within this space, there is a *single algebraically exact, norm-isometric, rotation-plane CRE embedding*. This completes the formal foundation of the entire CRE theory.

C Geometric and Noncommutative Closure

This section develops a geometric interpretation of the Conjugation–Real Embedding (CRE). Noncommutativity in complex, quaternionic, or

higher tensor algebras is re-read as *misalignment of orthogonal rotation planes* in a fixed real space $\mathbb{R}^{2m}$, while associativity and invertibility are studied inside the real matrix models where CRE acts as a $*$-isomorphism.

Throughout, we use the spectral mapping lemma for $*$-isomorphisms as a basic tool.

Lemma C.1 (Spectral mapping for $*$-isomorphisms)**.** *Let $\Psi : A \to B$ be a unital $*$-isomorphism between C^*-algebras. Then for every $a \in A$,*

$$\sigma_B(\Psi(a)) = \sigma_A(a).$$

Proof. If $\lambda \notin \sigma_A(a)$ then $(a - \lambda I)$ is invertible in A. Since Ψ is a unital $*$-isomorphism, $\Psi(a-\lambda I) = \Psi(a) - \lambda I$ is invertible in B, so $\lambda \notin \sigma_B(\Psi(a))$. Apply the same argument to Ψ^{-1} to obtain the reverse inclusion. □

C.1 Commutativity as Plane Alignment

We now make precise the statement that *commutativity of complex structures is equivalent to alignment of their rotation planes*.

Theorem C.2 (Commutativity and plane alignment)**.** *Let $S_1, S_2 \in O(2m)$ be* complex structures *on $\mathbb{R}^{2m}$, i.e.*

$$S_i^{\mathsf{T}} S_i = I_{2m}, \qquad S_i^2 = -I_{2m} \qquad (i = 1, 2).$$

Then

$$S_1 S_2 = S_2 S_1 \quad \Longleftrightarrow \quad \mathbb{R}^{2m} = \bigoplus_{j=1}^{m} V_j, \quad S_1(V_j) \subset V_j,\ S_2(V_j) \subset V_j,$$

where each V_j is a 2-dimensional S_1- and S_2-invariant plane.

Proof. ($\Rightarrow$) The operators S_1 and S_2 are both orthogonal and satisfy $S_i^2 = -I$, so each is diagonalizable over $\mathbb{C}$ with eigenvalues $\pm i$. Commutativity implies they can be simultaneously put into complex Jordan form, hence simultaneously block-diagonalized over $\mathbb{R}$ into 2×2 rotation blocks. Thus there exists an orthogonal decomposition into 2-planes $\{V_j\}$ which are invariant under both S_1 and S_2.

($\Leftarrow$) Conversely, if $\mathbb{R}^{2m} = \bigoplus_j V_j$ with each V_j invariant under S_1 and S_2, then each restriction $S_i|_{V_j}$ is a 2×2 orthogonal matrix with square $-I$, i.e. a rotation by $\pm\pi/2$. Rotations in a fixed 2-plane commute, so $S_1|_{V_j} S_2|_{V_j} = S_2|_{V_j} S_1|_{V_j}$ for all j, and hence $S_1 S_2 = S_2 S_1$ on $\mathbb{R}^{2m}$. □

Remark C.3. *Theorem C.2 shows that noncommutativity of complex/quaternionic structures is purely geometric: it arises precisely when the associated rotation planes in $\mathbb{R}^{2m}$ cannot be made simultaneously invariant under all complex structures.*

C.2 Cayley–Dickson Extension as Rotational Misalignment

We now reinterpret the Cayley–Dickson passage $\mathbb{R} \to \mathbb{C} \to \mathbb{H}$ in terms of interacting rotation planes.

Theorem C.4 (Quaternionic units as misaligned complex structures)**.** *In the standard real 4-dimensional representation of the quaternions $\mathbb{H}$ there exist*

three orthogonal complex structures $S_i, S_j, S_k \in O(4)$ such that

$$S_\alpha^2 = -I_4, \qquad S_\alpha S_\beta = -S_\beta S_\alpha \quad (\alpha \neq \beta),$$

and the left action of the imaginary units i, j, k on $\mathbb{H} \cong \mathbb{R}^4$ is given by these operators. In particular, noncommutativity of quaternionic multiplication is realized as antisymmetric composition of rotations in misaligned 2-planes.

Proof sketch. Identify $\mathbb{H}$ with $\mathbb{R}^4$ using the basis $\{1, i, j, k\}$. Left multiplication by i, j, k gives three real-linear maps L_i, L_j, L_k on $\mathbb{R}^4$. Each L_α is orthogonal and satisfies $L_\alpha^2 = -I_4$, so each defines a complex structure $S_\alpha := L_\alpha$. The quaternionic relations $ij = k$, $jk = i$, $ki = j$ and $ij = -ji$ translate to $S_\alpha S_\beta = -S_\beta S_\alpha$ for $\alpha \neq \beta$. These three complex structures act on mutually interacting 2-planes in $\mathbb{R}^4$, and their anticommutation expresses the $180°$-twisted misalignment between those planes. □

Remark C.5. *Under this identification, the familiar identity $ij = -ji = k$ is exactly the statement that applying a $\pi/2$ rotation in one complex plane followed by a $\pi/2$ rotation in an orthogonal plane differs by a sign from doing the same rotations in the opposite order. CRE makes this behaviour explicit inside $\mathrm{SO}(4)$, with each imaginary unit realized as a rotation-plane generator.*

C.3 Associative Matrix Realization in an Inductive Real Limit

The advantage of CRE is that it keeps all constructions inside associative matrix algebras, even when the source algebra (e.g. higher Cayley–Dickson stages) exhibits nonassociativity.

Theorem C.6 (Associative realization in an inductive real limit)**.** *Let $(\mathcal{H}_n)_{n\geq 1}$ be a sequence of finite-dimensional real Hilbert spaces with $\mathcal{H}_n \cong \mathbb{R}^{2^n}$, and suppose we have $*$-homomorphic embeddings*

$$\Phi_n : M_{2^{n-1}}(\mathbb{C}) \hookrightarrow B(\mathcal{H}_n)$$

obtained by iterating the CRE construction. Let

$$\mathcal{H}_\infty := \varinjlim_n \mathcal{H}_n$$

be the inductive limit Hilbert space, and let $\mathcal{A}_\infty := \varinjlim_n \Phi_n(M_{2^{n-1}}(\mathbb{C}))$ be the inductive limit algebra in $B(\mathcal{H}_\infty)$. Then multiplication in $\mathcal{A}_\infty$ is associative, and for all A, B, C arising from some finite stage,

$$(\Phi_{\mathrm{conj}}(A)\Phi_{\mathrm{conj}}(B))\Phi_{\mathrm{conj}}(C) = \Phi_{\mathrm{conj}}(A)(\Phi_{\mathrm{conj}}(B)\Phi_{\mathrm{conj}}(C)).$$

Proof. Each $M_{2^{n-1}}(\mathbb{C})$ is an associative C^*-algebra, and each Φ_n is a $*$-homomorphism. Hence multiplication is associative at every finite stage. The inductive limit $\mathcal{A}_\infty$ is defined as the closure of the union of these images in $B(\mathcal{H}_\infty)$, with multiplication inherited from the operator product, which

is associative. The stated identity holds already at any finite stage where A, B, C are defined and therefore holds in the limit. □

Remark C.7. *From this viewpoint, the failure of associativity in $\mathbb{O}$ and higher Cayley–Dickson algebras is seen as an obstruction to realizing all imaginary directions simultaneously as commuting complex structures on a* single *finite $\mathbb{R}^{2m}$, rather than a breakdown of associativity in the underlying matrix model. CRE deliberately stays inside the associative real matrix envelope, where associators can be analysed without ambiguity.*

C.4 Invertibility via Quaternionic Real Embeddings

We conclude with a precise statement of how quaternionic invertibility is captured by a CRE-style real block embedding.

Let $\varphi : \mathbb{H} \to M_2(\mathbb{C})$ be the standard $*$-isomorphism of quaternions with 2×2 complex matrices, and define the real representation

$$\rho := \Phi_{\mathrm{conj}} \circ \varphi : \mathbb{H} \to M_4(\mathbb{R}).$$

Theorem C.8 (Invertibility correspondence for quaternionic matrices)**.** *Let $A \in \mathbb{H}^{n\times n}$ and let*

$$\rho_n(A) := \big(\Phi_{\mathrm{conj}} \circ \varphi\big)^{(n)}(A) \in M_{4n}(\mathbb{R})$$

be the induced real block representation obtained by applying ρ entrywise.

Then:

$$A \text{ is invertible } \iff \rho_n(A) \text{ is invertible.}$$

Furthermore, the spectra and condition numbers satisfy

$$\sigma(\rho_n(A)) = \sigma(\varphi^{(n)}(A)) \cup \overline{\sigma(\varphi^{(n)}(A))}, \qquad \kappa_2(\rho_n(A)) = \kappa_2(\varphi^{(n)}(A)),$$

and inversion of A can be performed via real SPD-based factorisation of $\rho_n(A)$ followed by back-mapping through ρ_n^{-1}.

Proof. Both $\varphi^{(n)} : \mathbb{H}^{n\times n} \to M_{2n}(\mathbb{C})$ and $\Phi_{\text{conj}} : M_{2n}(\mathbb{C}) \to M_{4n}(\mathbb{R})$ are $*$-isomorphisms. Their composition ρ_n is therefore a unital $*$-isomorphism of C^*-algebras. By Lemma C.1, invertibility is preserved and reflected:

$$A \text{ invertible } \iff \varphi^{(n)}(A) \text{ invertible } \iff \rho_n(A) \text{ invertible,}$$

and spectra and norms are preserved up to complex conjugation. The final statement follows by applying any stable real factorisation (e.g. Cholesky/Schur-complement on $\rho_n(A)$ when SPD structures are available), and then mapping back with ρ_n^{-1}. □

Remark C.9. *In this real block model the three quaternionic imaginary directions act as explicit rotation generators S_i, S_j, S_k on $\mathbb{R}^{4n}$, and their non-commutativity is encoded in the way these rotations interact across blocks. CRE thus relocates quaternionic invertibility questions into a fully associative, orthogonal matrix setting, where standard real SPD tools apply.*

C.5 Summary of Geometric and Noncommutative Closure

- **Plane alignment:** Commutativity of complex structures is equivalent to preservation of a common decomposition into 2-dimensional rotation planes (Theorem C.2).
- **Cayley–Dickson geometry:** Quaternionic units i, j, k appear as mutually anticommuting complex structures on $\mathbb{R}^4$, so their noncommutativity is a pure plane-misalignment effect (Theorem C.4).
- **Associative limit:** Iterated CRE embeddings into real matrix algebras yield an inductive limit where multiplication remains associative (Theorem C.6).
- **Invertibility correspondence:** Quaternionic and complex invertibility are faithfully represented in the CRE-style real block embedding, with spectra and condition numbers preserved up to conjugation (Theorem C.8).

Interpretation. The Conjugation–Real Embedding provides a geometric lens on the noncommutative hierarchy: complex and quaternionic operations are realized as orthogonal rotations on a fixed real space, while algebraic phenomena such as noncommutativity and loss of associativity are reinterpreted as misalignment and incompatibility of rotation planes. Within the associative matrix envelope generated by CRE, questions of invertibility and stability can be treated using standard real SPD machinery, without losing access to the underlying noncommutative geometry.

D Numerical and Computational Closure

This section collects the analytic and algorithmic consequences of the Conjugation–Real Embedding (CRE) at the level of FLOP counts, norms, condition numbers, and backward error. Together with the experiment sections, these results show that CRE is numerically *closed*: every standard complex HPD/normal linear-algebra primitive admits a pivot-free real SPD implementation with strictly lower arithmetic cost and identical stability guarantees.

D.1 FLOP Model and Asymptotic Reduction

We work in the usual dense-arithmetic model in which matrix–matrix multiplication of two $n \times n$ real matrices costs $2n^3 + O(n^2)$ floating-point operations, and complex arithmetic is realized by pairs of real scalars.

Theorem D.1 (Complex-to-real FLOP reduction under CRE)**.** *Consider dense $n \times n$ factorisations and solves for Hermitian positive-definite matrices using:*

(a) standard complex arithmetic on $H \in \mathbb{C}^{n\times n}$;

(b) real arithmetic on the CRE image $\Phi_{\mathrm{conj}}(H) \in \mathbb{R}^{2n\times 2n}$.

In the classical FLOP model,

$$\mathrm{FLOPs}_{\textit{complex}} \approx 8n^3, \qquad \mathrm{FLOPs}_{\textit{CRE}} \approx 2n^3,$$

so CRE achieves an asymptotic $4\times$ reduction in the leading arithmetic constant for dense HPD/normal kernels.

Proof. A single complex multiplication $(a + ib)(c + id)$ requires four real multiplications and two real additions:

$$(ac - bd) + i(ad + bc),$$

so a complex matrix–matrix multiplication uses approximately $8n^3$ real FLOPs in the standard model.

Under Φ_{conj}, a complex $n \times n$ matrix is represented as a $2n \times 2n$ real block matrix, and Cholesky / Schur-based kernels are applied to this *real SPD* operator. As detailed in Corollary 7.5, the structure of $\Phi_{\mathrm{conj}}(H)$ eliminates redundant conjugate-coupled operations and yields an effective real cost of order $2n^3 + O(n^2)$ in the leading term.

Comparing the leading constants 8 vs. 2 gives the stated $4\times$ asymptotic reduction. □

Remark D.2 (Practical speed regimes)**.** *Empirical benchmarks in Tables 1, 3, and 5 show* 3–4$\times$ *speedups in wall-clock time once* n *is large enough for the asymptotic arithmetic cost to dominate overhead, in line with Theorem D.1.*

D.2 Memory Locality and Information Layout

While CRE doubles the dimension (from n complex to $2n$ real coordinates), it also reorganises the data into contiguous 2×2 blocks that are naturally compatible with SIMD and cache-line access.

Remark D.3 (Memory locality and bandwidth usage)**.** *In a standard inter-*

leaved complex layout, real and imaginary parts of entries require strided loads when accessed as separate real arrays. In contrast, the CRE layout stores each complex scalar as a 2×2 *real block in* $\Phi_{\mathrm{conj}}(A)$*,*

$$a + ib \longmapsto \begin{bmatrix} a & -b \\ b & a \end{bmatrix},$$

so entire columns and rows of the realified matrix are traversed by contiguous reads. On modern cache hierarchies this improves effective memory bandwidth and cache utilisation. The precise constant factor is hardware-dependent, but the experiments in Section 7 consistently show that CRE kernels begin to outperform complex baselines as soon as memory traffic becomes the main bottleneck.

D.3 Norms, Condition Numbers, and Backward Error

We recall that Φ_{conj} is an orthogonal $*$-isomorphism on operator spaces; see Theorem 6.1 and Theorem B.3.

Proposition D.4 (Norm invariance)**.** *For all* $A \in M_n(\mathbb{C})$*,*

$$\|\Phi_{\mathrm{conj}}(A)\|_2 = \|A\|_2.$$

Proof. By Theorem 6.1, Φ_{conj} is a norm-isometric $*$-algebra isomorphism between $(M_n(\mathbb{C}), \|\cdot\|_2, *)$ and its image inside $(M_{2n}(\mathbb{R}), \|\cdot\|_2, \mathsf{T})$. □

Corollary D.5 (Condition-number invariance)**.** *If* $A \in M_n(\mathbb{C})$ *is invertible,*

then

$$\kappa_2(\Phi_{\mathrm{conj}}(A)) = \kappa_2(A),$$

where $\kappa_2(\cdot)$ denotes the spectral-norm condition number.

Proof. By Proposition D.4 and multiplicativity of Φ_{conj},

$$\|\Phi_{\mathrm{conj}}(A)\|_2 = \|A\|_2, \qquad \|\Phi_{\mathrm{conj}}(A)^{-1}\|_2 = \|\Phi_{\mathrm{conj}}(A^{-1})\|_2 = \|A^{-1}\|_2.$$

Hence $\kappa_2(\Phi_{\mathrm{conj}}(A)) = \|\Phi_{\mathrm{conj}}(A)\|_2\|\Phi_{\mathrm{conj}}(A)^{-1}\|_2 = \|A\|_2\|A^{-1}\|_2 = \kappa_2(A)$. □

Theorem D.6 (Backward-error stability in CRE coordinates)**.** *Let $H \in \mathbb{C}^{n\times n}$ be Hermitian positive-definite and let $\widehat{L}$ be the Cholesky factor produced by a backward-stable real SPD routine applied to $\Phi_{\mathrm{conj}}(H)$. Then there exists a perturbation E such that*

$$\widehat{L}\widehat{L}^{\mathsf{T}} = \Phi_{\mathrm{conj}}(H) + E, \qquad \|E\|_2 \le c\,n\,\varepsilon_{\mathrm{mach}}\|\Phi_{\mathrm{conj}}(H)\|_2,$$

and the corresponding complex factor $\tilde{L} := \Phi_{\mathrm{conj}}^{-1}(\widehat{L})$ satisfies

$$\tilde{L}\tilde{L}^{*} = H + \tilde{E}, \qquad \|\tilde{E}\|_2 \le c\,n\,\varepsilon_{\mathrm{mach}}\|H\|_2,$$

with the same constant c.

Proof. The first inequality is the standard backward-error bound for real SPD Cholesky (see Theorem 9.1). Norm-isometry of Φ_{conj} implies $\|\tilde{E}\|_2 =$

$\|\Phi_{\mathrm{conj}}^{-1}(E)\|_2 = \|E\|_2$, and $\|\Phi_{\mathrm{conj}}(H)\|_2 = \|H\|_2$, yielding the second inequality. □

Remark D.7. *In practical terms, Theorem D.6 says that CRE-based solvers inherit* exactly *the same backward-stability constants as the underlying real SPD routines. CRE neither amplifies nor attenuates rounding error; it merely relocates it in* $\mathbb{R}^{2n}$.

D.4 HPD–SPD Equivalence and Solver Closure

We finally summarise the HPD↔SPD correspondence and its algorithmic consequences.

Theorem D.8 (HPD–SPD equivalence for Cholesky and solves)**.** *Let* $H \in \mathbb{C}^{n\times n}$ *be HPD and let* $M := \Phi_{\mathrm{conj}}(H) \in \mathbb{R}^{2n\times 2n}$. *Then:*

(i) M *is real SPD (Lemma 8.6);*

(ii) Cholesky factorisation commutes with Φ_{conj}*: if* $H = LL^*$ *then* $M = \Phi_{\mathrm{conj}}(L)\Phi_{\mathrm{conj}}(L)^{\mathsf{T}}$ *(Corollary 7.1);*

(iii) inversion commutes with Φ_{conj}*:* $\Phi_{\mathrm{conj}}(H^{-1}) = M^{-1}$*;*

(iv) all HPD-derived SPD routines (Cholesky, Schur complements, triangular solves) can be implemented entirely *in real arithmetic on* M*, with FLOP counts as in Theorem D.1 and backward error as in Theorem D.6.*

Proof. Items (i)–(iii) are proved in Lemmas 8.6, Corollary 7.1) and Theorem 8.7. Item (iv) follows because all standard HPD solvers can be expressed in terms of Cholesky, block multiplications, and triangular solves,

each of which commutes with Φ_{conj} and inherits the same stability bounds and FLOP-model constants from the SPD case. □

D.5 Summary of Numerical and Computational Closure

- **Arithmetic cost:** CRE reduces the leading arithmetic constant for dense HPD/normal kernels from $\approx 8n^3$ real FLOPs in complex arithmetic to $\approx 2n^3$ in real SPD form (Theorem D.1), yielding 3–4$\times$ empirical speedups in the regimes where computation dominates overhead.

- **Memory and layout:** Φ_{conj} packs complex entries into contiguous 2×2 real blocks, improving effective cache locality and bandwidth utilisation compared to interleaved complex storage (Remark D.3).

- **Norms and conditioning:** Operator norms and condition numbers are invariant under CRE (Proposition D.4, Corollary D.5), so analytic sensitivity is unchanged.

- **Backward stability:** CRE-based Cholesky and solve routines have the same $O(n\varepsilon_{\mathrm{mach}})$ backward-error bounds as their real SPD counterparts and thus as classical complex HPD algorithms (Theorem D.6).

- **Solver equivalence:** Every HPD-based complex linear-algebra primitive (factorisation, inversion, Schur complement) has a norm-isometric, pivot-free SPD realisation under CRE (Theorem D.8).

Conclusion. From a numerical perspective, CRE closes the loop between complex operator theory and real computation: the same HPD problem is solved in a different coordinate system where *all* of the following hold simultaneously:

identical spectra, identical norms and condition numbers,
identical backward error, and strictly lower arithmetic cost.

In this sense the Conjugation–Real Embedding provides a numerically and computationally complete real model for complex linear algebra.

E Implementation Summary

All algorithms presented in Sections 7–9 were implemented in `Python`/`NumPy`/`PyTorch` and evaluated in three floating-point configurations (`complex32`, `complex64`, and `complex128`), each built from IEEE 754 single- and double-precision real types.

All computations were performed on a Linux system equipped with an Intel® Core™ Ultra 9 285 CPU (24 cores, 7.2 GHz boost), an NVIDIA® GeForce™ RTX 5090 GPU (GB202 architecture, CUDA backend), and 128 GB of system memory. All experiments used deterministic seeding for reproducibility.

Software and Numerical Environment

- **Language and libraries:** Python 3.12.9, NumPy 2.3.4, SciPy 1.16.2, Pandas 2.3.3, and PyTorch 2.9.0.
- **Precision:** IEEE 754 double precision (binary64) as the reference for-

mat, with unit roundoff $\varepsilon_{\mathrm{mach}} \approx 2.22 \times 10^{-16}$. The complex formats `complex32`, `complex64`, and `complex128` were realised as pairs of real single-/double-precision values.

- **BLAS / LAPACK backend:** OpenBLAS (`x86_64`) for NumPy/SciPy linear-algebra kernels, and cuBLAS/cuSolver (CUDA backend) for PyTorch GPU routines.

- **Operating system:** Linux 6.12.0-55.18.1.el10_0.x86_64 (64-bit, GNU/Linux, little-endian).

- **Hardware configuration:** Intel Ultra 9 285 (24 cores, 1 thread/core, 36 MB L3 cache), NVIDIA RTX 5090 (24 GB VRAM), 128 GB system RAM.

- **Random seed control:** NumPy `PCG64DXSM` generator and PyTorch `manual_seed` were used with fixed integer seeds $\{0, 1, 2, \ldots, 41\}$. In configurations where individual runs failed due to vendor-kernel numerical issues or incomplete data generation, the corresponding seeds were discarded from the averages; this exclusion is indicated explicitly in the experiment tables (e.g. seed 12 for 32-/64-bit Hermitian tests in Tables 1–4).

- **Validation metrics:** For each experiment we recorded

    1. eigenvalue agreement $\sigma(H) = \sigma(\Phi_{\mathrm{conj}}(H))$ up to numerical tolerance;

    2. condition-number invariance $\kappa_2(H) = \kappa_2(\Phi_{\mathrm{conj}}(H))$;

3. Cholesky residual $\|LL^* - H\|_F / \|H\|_F$;
4. variance-based stability via $\mathrm{Var}(AA^{-1} - I)$ across random seeds.

These metrics were computed for both the traditional complex kernels and the CRE–ECC real implementations.

Pre/post-processing and timing protocol

All timing results reported in Sections 7–8 measure *end-to-end wall-clock time* for each method.

- For the **complex baselines** (SciPy `Cholesky`/`QR`/`LU` and corresponding PyTorch routines), the timer encloses: matrix generation, the library factorisation/solve call, and residual checks.
- For the **CRE / CRE–ECC methods**, the timer encloses:
  1. construction of the CRE object (e.g. block embedding $\Phi_{\mathrm{conj}}(H)$, Sym+Skew decomposition and S_α shift, formation of $A = S_\alpha + iK$),
  2. the real SPD factorisation/solve on the $2n \times 2n$ (or block) system using the *same* vendor BLAS/LAPACK backends,
  3. post-processing to map the solution back to the complex or tensor domain (e.g. $\Phi_{\mathrm{conj}}^{-1}$ to recover L or A^{-1}),
  4. residual and spectral checks.

  No CRE-specific preconditioners, matrix reorderings, or kernel fusions are introduced: the implementations call standard real SPD routines (Cholesky/solve) with *additional* operations layered on top.

- In particular, all reported speedups for CRE and Sym–Skew pipelines already include the full cost of pre-processing and post-processing. The comparison is therefore:

complex kernel on H vs. "$H \mapsto \Phi_{\mathrm{conj}}(H)$ + real SPD solve + back-map"

under identical random matrices, seeds, and BLAS backends.

This protocol ensures that the observed performance gains arise from the CRE information geometry itself (real SPD block structure), rather than from unaccounted preprocessing or specialised low-level kernels.

Algorithmic Correspondence

This appendix summarizes implementation conventions consistent with Section 7.1 and the CRE-based pipelines of Sections 7 and 8. The pseudocode in Algorithms 1, 2, and 3 corresponds exactly to the code used in all Hermitian and Sym–Skew experiments.

Block structure and inverse mapping. For a Hermitian positive-definite matrix $H \in \mathbb{C}^{n\times n}$, the Conjugation–Real Embedding is implemented as

$$\Phi_{\mathrm{conj}}(H) = \begin{bmatrix} \Re(H) & -\Im(H) \\ \Im(H) & \Re(H) \end{bmatrix} \in \mathbb{R}^{2n\times 2n},$$

with the inverse mapping defined blockwise: if $\Phi_{\mathrm{conj}}(H) = \begin{bmatrix} X & -Y \\ Y & X \end{bmatrix}$, then $H = X + iY$. The same convention is used for Cholesky factors and inverses: if

$$M^{-1} = \begin{bmatrix} X & -Y \\ Y & X \end{bmatrix} = \Phi_{\mathrm{conj}}(H^{-1}),$$

the complex inverse is recovered as $H^{-1} = X + iY$.

Hermitian pipeline (HPD $\rightarrow$ SPD). Each Hermitian benchmark in Sections 7 and 9 performed the following steps per run:

1. Generate $G \in \mathbb{C}^{n \times n}$ with $G_{ij} \sim \mathcal{N}(0,1) + i\,\mathcal{N}(0,1)$.
2. Form the Hermitian positive-definite matrix $H = (G + G^*)/2$.
3. Apply the CRE map: $\overline{H} = \Phi_{\mathrm{conj}}(H)$.
4. Factor $\overline{H} = GG^{\mathsf{T}}$ via real SPD Cholesky (or solve $\overline{H}x = b$ via real SPD routines).
5. Map back to the complex domain using $\Phi_{\mathrm{conj}}^{-1}$, reconstruct (L, L^{-1}, H^{-1}) as needed, and compute spectral, residual, and timing metrics.

All reported wall-clock times for CRE runs include *every* step above: CRE block construction, real SPD factorization/solve, back-mapping, and residual checks, in direct comparison to the corresponding complex Cholesky/QR/LU calls on H.

Sym–Skew pipeline (general inverse). For the non-Hermitian Sym–Skew experiments in Section 8.4, each run followed Algorithm 3:

1. Sample a full-rank real matrix G with $G_{ij} \sim \mathcal{N}(0, 1)$.
2. Form $S = (G + G^{\mathsf{T}})/2$ and $K = (G - G^{\mathsf{T}})/2$.
3. Shift S by a scalar multiple of the identity to obtain $S_\alpha = S + \alpha I_n$ that is numerically SPD.
4. Assemble the Hermitian lift $A = S_\alpha + iK$ and its real block $M = \Phi_{\mathrm{conj}}(A) = \begin{bmatrix} S_\alpha & -K \\ K & S_\alpha \end{bmatrix}$.
5. Factor $M = LL^{\mathsf{T}}$ (Block Chol) or solve $MX = I$ (Block LU) using standard real SPD/solve routines.
6. Partition $M^{-1} = \begin{bmatrix} X & -Y \\ Y & X \end{bmatrix}$ and recover $A^{-1} = X + iY$; compute $\|AX - I\|_F / \|A\|_F$ and timing statistics.

Again, CRE timings include the Sym+Skew step, the SPD shift, the CRE block formation, the real factorization/solve, and the back-mapping to obtain A^{-1}.

Note. *The sign convention in Φ_{conj} (whether the off-diagonal block is $+\Im(H)$ or $-\Im(H)$) fixes the orientation of the underlying rotation in the realified complex plane. Any consistent choice is admissible; throughout this thesis we use the convention of Section 7.1, with all code and proofs aligned to that sign.*

Reproducibility Protocol

All Hermitian results in Tables 1–6 and Figures 1–3 were produced using the following protocol:

1. **Matrix sizes and seeds.** For each matrix size n and each precision configuration (`complex32`, `complex64`, `complex128`), a fixed set of integer seeds was used (cf. Appendix E). Unless explicitly noted in the tables, all seeds contributed to the averages; excluded seeds (e.g. vendor-kernel failures) are listed in the corresponding captions.

2. **Trials and averaging.** For each $(n, \text{precision})$ pair and each seed, we ran multiple independent realizations (typically 30 per seed) and recorded:

   - runtime for the complex baseline;
   - runtime for the CRE / CRE–ECC method;
   - residual $\|LL^* - H\|_F / \|H\|_F$;
   - variance proxies (e.g. $\mathrm{Var}(AA^{-1} - I)$).

   The reported values in the tables are mean $\pm$ standard deviation across all valid seeds and realizations.

3. **Residual and spectral checks.** For each run we computed

$$r_{\text{chol}} = \frac{\|LL^* - H\|_F}{\|H\|_F}, \qquad \delta_\lambda = \max_i \big|\lambda_i(H) - \lambda_i(\Phi_{\text{conj}}(H))\big|,$$

   and verified that r_{chol} remained within $O(\varepsilon_{\text{mach}})$ and δ_λ within 10^{-13}–10^{-14} in double precision.

4. **End-to-end timing.** All timings are end-to-end: the timer encloses matrix generation, CRE block embedding (and Sym–Skew lift where applicable), the factorization/solve call, the back-mapping to complex form, and the residual/spectral checks. No hidden pre-processing, kernel fusion, or warm-starting is used.

All experiments can be reproduced in a standard `Python`+`NumPy`/`SciPy` environment using CPU-only BLAS/LAPACK; GPU acceleration via `PyTorch`/CUDA is optional and only affects wall-clock time, not the numerical results. The scripts follow the pseudocode in Algorithms 1, 2, and 3 line by line, with CRE-specific steps (block formation and back-mapping) implemented explicitly rather than hidden inside library kernels.

Code Availability

Full source code for the AXIOM/CRE implementation, experimental setup, and data collection is available from the author upon reasonable request. The implementation follows exactly the axiomatic layer of Appendix A (slice-compositional CRE algebras and ALH structure) and the CRE completion layer of Appendix B (polynomial closure, norm-isometric embedding, and rotation-plane normalization).

All numerical routines are direct realizations of the pseudocode in Algorithms 1, 2, and 3, so that every experiment reported in Sections 7–9 can be reconstructed exactly from the text without additional hidden heuristics or implementation-specific tricks.

Bibliography

[1] P. A. M. Dirac, "Quantised singularities in the electromagnetic field,," *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, vol. 133, no. 821, pp. 60–72, 1931.

[2] J. von Neumann, *Mathematical Foundations of Quantum Mechanics: New Edition*. Princeton University Press, Feb. 2018.

[3] D. Hilbert, *Grundzüge Einer Allgemeinen Theorie Der Linearen Integralgleichungen*. No. 3, BG Teubner, 1924.

[4] F. Riesz and B. S. Nagy, *Functional Analysis*. Courier Corporation, 2012.

[5] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge university press, 2012.

[6] N. J. Higham, *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics, Jan. 2008.

[7] G. H. Golub and C. F. Van Loan, *Matrix Computations*. JHU press, 2013.

[8] L. N. Trefethen and D. Bau, *Numerical Linear Algebra, Twenty-fifth Anniversary Edition*. Philadelphia, PA: Society for Industrial and Applied Mathematics, Jan. 2022.

[9] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, second ed., Jan. 2002.

초록

AXIOM: 인덱스 정렬 곱셈에 의한 대수적 확장

박준후
자연과학대학
서울대학교

전통적인 실수화(realification)는 복소 연산자를 차원 의존적이고 비대칭적인 방식으로 실수 행렬로 바꾼다. 이때 복소 구조는 더 큰 실수 공간에 "끼워 넣어야 할 것"으로 취급되며, $\mathbb{C}$를 넘어서는 일반화는 대수적 구조를 잃거나 데이터 배치가 불필요하게 팽창하는 문제를 피하기 어렵다. 본 학위논문은 이에 대하여 ***AXIOM (Algebraic eXtension by Index–Ordered Multiplication)***이라는 공리적 틀을 도입한다. AXIOM 안에서는 복소수, 사원수, 팔원수, 일반 Cayley–Dickson 계열, 그리고 임의의 텐서 대수 위의 연산자를 *하나의* 회전 기반 실수 행렬 모형으로 표현할 수 있으며, 이 표현은 곱셈, $*$-구조, 스펙트럼, 가역성(invertibility), 효과적인 메모리 사용량을 보존한다. 여기서 결합성(associativity)은 실제 대수가 아니라 *표현 공간*인 실수 행렬 대수에 귀속된다.

AXIOM 위에서 본 논문은 보편적인 대칭 블록 실수 표현

$$\Phi(A) = \begin{bmatrix} H & -K \\ K & H \end{bmatrix}, \qquad A = H + iK,$$

을 정의하고, 이를 *켤레–실수 매장(Conjugation–Real Embedding, CRE)*이라 부른다. Φ는 임의의 AXIOM 대수에 대해 $*$-등거리($*$-isometric)인 실수 표현이 됨을 증명한다. 즉, 대수적 연산자와 대칭 블록 실수 행렬 사이에 일대일 대응을 만들어 내며, "실수 호스트 vs. 복소 게스트"라는 위계가 존재하지 않는다. 구현 관점에서 Φ는 *추가 저장 공간을 전혀 요구하지 않는다*: 대칭 블록 관점은 이미 복소(또는 Cayley–Dickson) 데이터를 지탱하고 있는 그 동일한 실수 스칼라들을 재인덱싱한 것일 뿐이므로, 메모리 사용량과 통신량은 원래 복소 레이아웃과 완전히 동일하게 유지되면서 모든 분해가 순수 실수 연산만으로 진행된다.

본 논문의 핵심 공헌은 다음의 세 가지 구조 정리로 요약된다.

1. **텐서 정렬 곱셈 정리(Tensor-Sorted Product Theorem).** 모든 2^n차원 Cayley–Dickson 대수와 모든 AXIOM 텐서 대수는 차원별 곱셈표 없이, 순수하게 텐서 인덱스 정렬만으로 결정되는 부호 결정적(table-free, sign-deterministic) 곱셈 규칙을 가진다. 이로써 각 차원마다 따로 곱셈표를 정의해야 하는 전통적 구성을 제거한다.

2. **산술적 국소 Hurwitz 정리(Arithmetic Local Hurwitz, ALH).** 각 생성자 e_i는 $\mathbb{C}$와 동형인 국소 부분대수를 생성하므로, 모든 차원에서 국소 슬라이스 수준의 복소수형 역원 계산과 노름 합성(norm composition)이 가능하다. 이는 고전적인 1–2–4–8 Hurwitz 제약이 전역 구조의 문제임을 분리해 내고, 임의 차원에서 항상 존재하는 "국소 복소 해석" 층을 분명히 한다.

3. **대칭 블록 표현 및 CRE 역원 정리(CRE Inversion Theorem).** 임의의 AXIOM 대수 연산자는 대칭 리프트(symmetric lift)를 거친 뒤, 피벗이 필요 없는 블록 Cholesky (또는 QR/SVD 계열) 분해를 갖는

다. 이때

$$\Phi(AB) = \Phi(A)\Phi(B), \qquad \Phi(A^{-1}) = \Phi(A)^{-1}$$

가 성립하며, 차원과 무관한 수치적 안정성이 보장된다. 기본 데이터 배치가 변하지 않으므로, 기존의 복소/사원수/팔원수 전용 커널들을 대체하는 *단일 실수 커널*을 얻을 수 있고, 이 과정에서 메모리 오버헤드는 발생하지 않는다.

이 구조적 통합을 통해 AXIOM–CRE 프레임워크는 19세기 이후 오랫동안 기술적으로 열려 있던 몇 가지 난점을 표현 수준에서 해결한다. 예를 들어, 사원수 · 팔원수 곱셈과 Einstein 합 표기법의 비호환성, $\mathbb{C}$를 넘어서는 차원 독립적 연산자 해석(operator calculus)의 부재, Cayley–Dickson 계열 전체를 하나의 결합적인 계산 모형으로 일관되게 수용하는 틀의 부재 등이 그러하다. AXIOM은 이러한 문제들을 모두 한 번에 정합적으로 다루는 회전 기반 실수 표현을 제시하며, 모든 2^n차원 및 텐서 대수에 대해 동일하게 적용된다. 이때 원래 대수의 비가환성 · 비결합성은 유지되지만, 이를 다루는 계산 모델은 항상 결합적인 실수 행렬 대수 안에서 기술된다.

중요하게도, 이러한 이론적 통합은 *직접적인 계산 이득*으로 이어진다. 모든 연산자를 대칭 블록 회전 형태로 표현함으로써, Cholesky 계열 분해의 선행 상수는 이론적으로 약 $4\times$ 감소하고, 표준 수치 선형대수 라이브러리 위에 단순히 AXIOM–CRE 재인덱싱 계층만 추가했을 때에도, 기본 커널을 전혀 수정하지 않고 3–4$\times$ 수준의 실측 런타임 가속을 얻을 수 있다. 이때 스펙트럼과 후방 오차(backward error)는 기계 정밀도 수준에서 유지되며, 메모리 사용량과 통신 비용 역시 원래 복소 레이아웃과 동일하다.

결론적으로, 본 논문은 복소 연산자를 단순히 실수 행렬에 *매장*하는 것을 넘어, 실수 및 복소 연산자 이론 사이의 구조적 간극을 닫는다. 양자를 하나의 대칭 블록 회전 모형 안에서 서로 다른 인덱스 정렬로 재해석함으로써, Cayley–Dickson 및 텐서 계층 전체에 대해 차원 독립적인 연산자 해석을 제공하고, 고차원 대수 위 알고리즘을 위한 통합적이면서도 실질적으로 우수한 계산 기반을 제시한다.

Keywords : AXIOM, 인덱스 정렬 곱셈, Cayley–Dickson 알게브라, 텐서 · 클리포드 알게브라, 대칭 블록 실수 표현, 수치 선형대수 커널
Student Number : 2023-20862

Acknowledgments

Turning points in life are never easy. Deciding to return to academia and to continue along this path was only possible because of the support of my family — especially my mother and my sister, whose concern and encouragement gave me strength and because of the joy brought to me by my nieces and nephews. I am also deeply grateful to my advisor, Professor Myungjoo Kang, and to Dr. Youngsoo Ha for their generous guidance. The academically stable and trustworthy environment they created has allowed me to explore new directions beyond the boundaries of existing theory without hesitation.

At one point, half in jest and half in judgment, I was called an "omniscient nilpotent." Someone who seems to know everything, yet whose visible contribution might collapse to zero once it is actually applied. But in linear algebra, nilpotent operators are never trivial: they govern the hidden structure of a space. In that sense, this thesis is a kind of formal proof that the associated morphism in question is not a zero operator after all.

No matter how much knowledge or technique I had accumulated, timing and luck often felt almost orthogonal to it, lying in a subspace that I could not quite grasp. From the outside, that makes a vector look nilpotent: full of components in invisible directions, but with no observable effect along the axis that "counts." Only recently has the projection of that luck onto the subspace of my own knowledge become nonzero, casting a small but definite shadow. In physical terms, what used to be a force applied

at a right angle has finally acquired a displacement in the same direction, so that $F \cdot s = W$ holds with $W > 0$. The work may be small, but it is no longer zero.

My life has, in its own way, been a sequence of twists and trials. I had to prove myself — and I have. To others, I may have looked like someone who simply does whatever they want without hesitation. But every choice was, in truth, a continuation of the same argument: that consistency with truth is itself a form of freedom.

In *Strange Stories from a Chinese Studio* (*Liaozhai Zhiyi*), it is said that life is seven parts luck and three parts skill. I am not fond of gambles whose expectations cannot be adjusted within a martingale framework, but perhaps life itself belongs to that category. In the end, it must have been luck that the domain untouched by anyone — the one I pointed to decades ago, asking why no one pursued it — remained nontrivial even when I returned to study it again.

Of course, at an age when I was expected to be working rather than studying, choosing to return to university may have seemed to others like an unstable or even unreasonable decision. I admit it — it was a form of deficiency. I chose to continue the path I once abandoned in order to fill that gap, and I hope this thesis marks the beginning of that journey. The thirst to make up for what is missing will likely never be fully satisfied. Yet perhaps that is precisely what allows a deficit to become something positive — the very source of motion itself. After all, an atom with an electron deficit is, quite literally, positive. Formally, the topic may still be categorized as unsolved, but to me it has long felt self-evident. Lacking the patience or the

years it would take to prove something like Fermat's Last Theorem — even if, in a certain one-dimensional informational encoding, the margin on the paper is literally unbounded — the best I could do within this modest thesis was to include the general inverse for any tensor whose invertibility can be assured.

We no longer need to ask whether we are left-handed or right-handed; tensors once thought irreducibly noncommutative can now simply be treated with a unified general inverse. That the computation became slightly faster in the process is, in some sense, merely incidental.

I express my gratitude to all the self-evident axioms and laws that have sustained this research, and to the predecessors who tried to illuminate this path even if the proofs they left behind remained unfinished. This thesis is just one more term in a long series expansion. It is neither omniscient, nor — I sincerely hope — nilpotent.