On Recursively Reducible Structures in High-Dimensional Riemann Theta Functions

Masaru Moriyama

April 2026

Technical Note / Conceptual Supplement

1. Background and Motivation

The numerical computation of Riemann theta functions is known to require computational effort that grows exponentially with the dimension g (O((2N+1)^g)). This "curse of dimensionality" has limited practical exact computation to roughly $g \le 10$ in general settings.

It is also known that simplifications occur when the period matrix has special structure, such as complete block diagonalization or degenerate limits. However, these have generally been treated as trivial factorizations or limiting cases, and have not been systematically developed as a framework for practical computability in high dimensions.

2. Observed Phenomenon

In the course of implementing structure-dependent algorithms based on S(k,k)-type decompositions (including k = 2, 3, 5, …), the following phenomenon was observed:

For certain classes of period matrices, the evaluation of theta functions becomes tractable through a recursive reduction of the effective dimension.

More specifically, for period matrices Ω possessing a certain structure, the evaluation problem decomposes hierarchically into lower-dimensional subproblems, and this process can be applied repeatedly.

As a result, the computational scaling departs significantly from the previously assumed exponential growth, exhibiting behavior that is effectively polynomial or nearly linear in practice. This has been confirmed to hold even in the extremely high-dimensional regime $g > 10^4$.

3. Relationship to Existing Theory

This phenomenon is not unrelated to existing theory, but does not appear to have been explicitly identified or systematically organized.

- In Riemann Theta Function Theory, simplifications under special structure (factorization, degeneration, etc.) are known.
- In Integrable Systems, period matrices with structure arise in finite-gap solutions and related contexts.
- In Siegel moduli spaces, studies of special subvarieties (loci) exist.

The present observation differs in the following respect: a substantial reduction in computational cost occurs without assuming complete block diagonalization or degenerate limits.

Numerical approaches for high-dimensional theta functions include tensor-train and hyperbolic cross methods (Claeys–Seiner et al., 2023), which achieve approximate evaluation up to $g \sim 60$ with relative error ~0.01, but face fundamental limits at ultra-high dimensions. The FLINT implementation (Elkies–Kieffer, 2025) provides certified precision for general period matrices, but assumes Siegel reduction and its practicality at ultra-high genus remains open.

4. Characteristic Properties

The observed structure is believed to possess the following properties:

- Not a complete block diagonalization in the standard sense
- Not a degenerate limit in the usual sense
- Recursive reduction of effective dimension is possible during the evaluation process
- Exact (non-approximate) computation is possible under structural constraints

These properties place the observed structure intermediate between classical reducible systems and general high-dimensional systems.

5. Terminology (Provisional)

To describe this phenomenon, the following provisional terminology is introduced:

Recursively Reducible Theta Structures

This refers to a subclass of period matrices that admit recursive reduction of the effective dimension in theta function evaluation. S(2,2) is the first and most computationally efficient example; the concept generalizes to the family S(k,k) ($k = 3, 5, \ldots$) and to mixed-base decompositions. We emphasize that at this stage, this is a descriptive term based on observed computational properties, not a rigorous mathematical definition.

6. Computational Implications

For structures belonging to this class, the computational cost of theta function evaluation changes from exponential to effectively polynomial scaling.

This enables:

- Exact computation in ultra-high dimensions
- Construction of reliable benchmarks for structured problems

Furthermore, computing across the full S(k,k) family dramatically expands the range of g that can be covered exactly:

- S(2,2): $g = 2^n$ series
- S(3,3): $g = 3^n$ series
- S(5,5): $g = 5^n$ series

Since these series are nearly disjoint, combining the exact values from each yields a **structured benchmark dataset** covering a wide range of g. This dataset serves as a reference for verifying the approximation accuracy of general methods and for constructing perturbative extensions.

In the implementation, stable computation has been confirmed even beyond g = 20,000.

Use as exact reference values for approximate methods on general structures — in particular, where approaches such as those of Claeys–Seiner et al. (2023) are limited to g ~ 60 with relative error ~0.01, exact values under S(k,k) structure provide direct verification benchmarks in the ultra-high-dimensional regime.

7. Open Problems

This observation raises the following questions:

- Rigorous mathematical formulation of recursively reducible structures
- Clarification of their position within Siegel moduli space
- Criteria for distinguishing non-trivial reducibility from hidden factorization
- Treatment of approximate structure (perturbation)
- Applications to high-dimensional finite-gap solutions and related topics
- Implementation extension to general S(k,k) and preparation of benchmark datasets for each series
- The problem of determining the optimal decomposition base k for a given g (prime factorization-type optimization)

8. Scope and Limitations

This note does not claim:

- A general solution method for arbitrary period matrices
- A complete classification of such structures
- A proof of the complete novelty of the phenomenon

The purpose of this note is to delineate and describe a “computationally accessible high-dimensional regime” that has not previously been systematically organized.

While this note is written primarily with S(2,2) as the leading example, the same logic applies to the entire S(k,k) family, extending the class of computable g to arbitrary prime powers and their products. S(2,2) is positioned as the implementation with the highest padding efficiency within this family.

9. Implementation

An implementation based on S(2,2) decomposition is publicly available at:

https://github.com/Moriyamax/s22-theta-acceleration

This implementation constitutes the experimental basis for the observations described in this note.

10. Moduli Space of Computational Structures (Conceptual Note)

The period matrices Ω considered in this note are not derived from algebraic curves and thus do not belong to the Jacobian locus. Instead, they exhibit a *computational structure* that enables recursive dimensional reduction in the evaluation of Riemann theta functions.

The equivalence classes of such Ω may form a moduli space that is independent of classical geometric moduli. While a rigorous definition of this **computational moduli space** is beyond the scope of the present note, the S(2,2) construction—and its generalizations to the S(k,k) families—can be regarded as concrete examples of this new type of moduli.

Summary

This note reports a phenomenon in high-dimensional Riemann theta functions whereby, for period matrices with certain structure, the evaluation problem recursively reduces to lower dimensions, enabling exact computation in a regime previously considered intractable. This suggests a new aspect of computability in high-dimensional theta theory.

The S(k,k) family (k = 2, 3, 5, …) provides a systematic and nearly disjoint collection of exactly computable g-series, forming a structured benchmark dataset of broad utility for high-dimensional theta function research.