

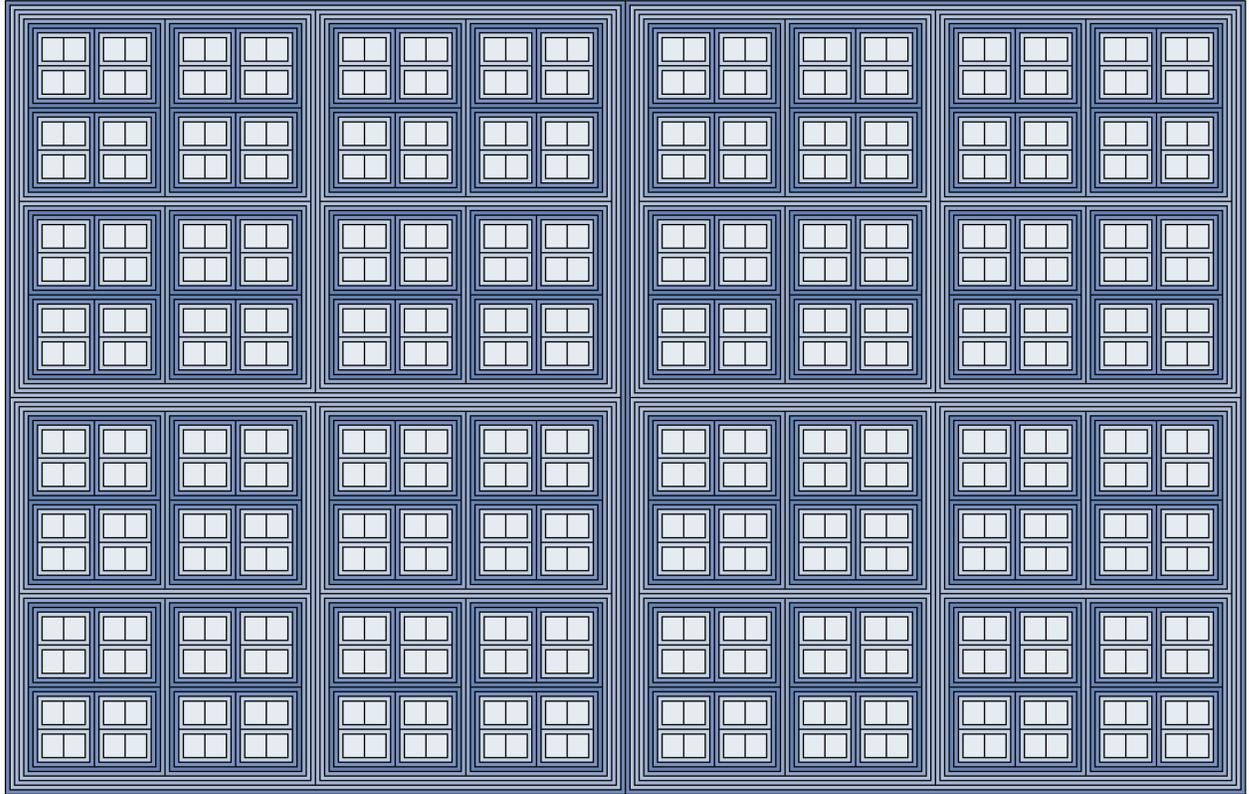
SciSci Research

サイサイ・リサーチ

Inventions No. 1

sci-sci.org/hensel-cpu

DOI: 10.5281/zenodo.17288854



## HENSEL CPU (ヘンゼル CPU):

A 2-Adic Computing Architecture for  
Exact Arithmetic

James Douglas Boyd

**SciSci Research, Inc.**

Boulder, Colorado, United States

[www.sci-sci.org](http://www.sci-sci.org)

**Copyright** © 2025 by SciSci Research, Inc. All Rights Reserved.

**Citation Format:**

Boyd, J.D. (2025). HENSEL CPU (ヘンゼル): A 2-Adic Computing Architecture for Exact Arithmetic. SciSci Inventions, 1(1). DOI: 10.5281/zenodo.17288854

# Contents

<b>1</b>	<b>Introducing Hensel</b>	<b>2</b>
1.1	A CPU Architecture for $\mathbb{Q}_2$ . . . . .	2
1.1.1	Defining Exactness . . . . .	2
1.1.2	The Question of Technical Risk . . . . .	3
1.1.3	The Prize of Exact Arithmetic . . . . .	3
1.2	Report Scope . . . . .	4
<b>2</b>	<b>Novel Architectural and Coding Features</b>	<b>5</b>
2.1	Encoding Standards . . . . .	5
2.1.1	FC 1-2025 Operand Encoding . . . . .	5
2.1.2	FC 2-2025 Instruction Encoding . . . . .	5
2.2	The Hensel Processor . . . . .	5
2.2.1	2AALUs . . . . .	6
2.2.2	PR Cluster . . . . .	6
<b>3</b>	<b>FC 1-2025</b>	<b>8</b>
3.1	A Quick Review of 2-adic Expansions and Coefficients . . . . .	8
3.2	FC 1-2025 Encoding . . . . .	8
3.3	Warm-Up: FC-1-2025 Numbers . . . . .	10
<b>4</b>	<b>The Processor</b>	<b>12</b>
4.1	The Processor Register Cluster . . . . .	12
4.2	Nested-Clustered Design . . . . .	13
4.3	Cluster Load-Store and Addressing . . . . .	13
<b>5</b>	<b>Parallelization</b>	<b>16</b>
5.1	Parallelization . . . . .	16
5.2	Parallelized Load-Store . . . . .	16
5.3	FC 2-2025 Encoding . . . . .	16
5.4	Parallelization and Non-Archimedean Distance . . . . .	17
5.5	Max $(\Delta_2^{h_\ell})$ Parallelization . . . . .	19
5.6	Constraints on Optimization . . . . .	22
<b>6</b>	<b>Further Prospects</b>	<b>23</b>
6.1	Error Correction and Fault Tolerance . . . . .	23
6.1.1	Error-Checking with $(\chi, \pi)$ -Payloads . . . . .	23
6.1.2	$\Delta_2^\mu$ Trajectory-Defect Checks . . . . .	23
6.1.3	LSU Consensus . . . . .	23
6.2	Supercomputing . . . . .	24

# 1 Introducing Hensel

## 1.1 A CPU Architecture for $\mathbb{Q}_2$

Introduced preliminarily in this report is the Hensel CPU (ヘンゼル CPU), designed according to a novel computing architecture with the aspiration of replacing floating-point arithmetic with exact arithmetic performed in  $\mathbb{Q}_2$  (i.e., 2-adic arithmetic). Here, exact arithmetic denotes arithmetic which, although subject to computational bounds (e.g., constrained by a given bit-width), is nonetheless performed on operands whose representation in bits is unique. The Hensel architecture is designed for arithmetic in  $\mathbb{Q}_2$ , rather than  $\mathbb{R}$ . In floating-point arithmetic, "floating-point numbers" (i.e., elements of  $\mathbb{Q}$ ) approximate elements of  $\mathbb{R}$  with finite precision. It has been known among some computer scientists since the 1970's that  $p$ -adic numbers (e.g., elements of  $\mathbb{Q}_2$ ) – where, from Ostrowski's theorem, we know that  $\mathbb{Q}_p$  and  $\mathbb{R}$  are the only completions of  $\mathbb{Q}$  – admit exact, unique representations with finite encodings. (Examples of precedents include the so-called "Hensel codes" of Krishnamurthy *et al.*, the work of Horspool-Hehner, and Doris' system for exact  $p$ -adic arithmetic in Magma.)

Unlike the above precedents, Hensel is an architecture, rather than an algorithm or software package. Thus, it is designed to perform arithmetic in  $\mathbb{Q}_2$  at the machine level, where  $\mathbb{Q}_2$  is distinct from all other  $\mathbb{Q}_{p>2}$  in that the coefficients of 2-adic expansions are  $a_i \in \{0, 1\}$  and thus can be written in bits. Descending to the architectural level carries several prospective advantages to CPU users. First, users can utilize a CPU performing arithmetic in  $\mathbb{Q}_2$  at the machine level without any familiarity with 2-adic arithmetic; operands in  $\mathbb{Q}_2$  can be expressed – for instance, symbolically – in user-recognizable form in higher-level programming languages. Thus, unlike software such as Magma, which offers exact arithmetic to number theorists familiar with  $\mathbb{Q}_p$ , the Hensel CPU is designed to bring the accu-

racy and performance of exactness to general users. Moreover, by developing a design to realize exact arithmetic at an architectural level, SciSci Research and its Future Computing group endeavor to confront the barriers posed by floating-point to high-performance computing (HPC), and help to realize an exact HPC capability unhindered by tradeoffs between performance, accuracy, and cost.

### 1.1.1 Defining Exactness

Although no computing system can overcome the limitation of computing with finite resources over fields with cardinalities of the continuum, computers can be designed such that the operands over which they compute are unique; such is the aspiration of exactness pursued by SciSci Research and Future Computing in designing the Hensel CPU. Thus, the criterion of exactness is not to be mistaken for the promise of computability. One will ineluctably encounter examples of arithmetic over particular numbers that the Hensel CPU cannot perform within its bit-width (in which case one will receive an error message, rather than an approximation), but the architecture is designed such that, when it can perform arithmetic, it does so exactly. Furthermore, the architecture, which is designed to favor scaling towards supercomputing applications, is advanced with the aspiration of extending exactness to the largest collection of operands possible (e.g., with a large bit-width).

It should be noted that the Hensel CPU is not entirely untethered from the question of approximation insofar as its instructions and operand encodings are developed with the assistance of software such as Sage and Magma, which return so-called "lazy" representations of 2-adic numbers (i.e., up to a specified precision). Nonetheless, arithmetic performed by a Hensel CPU can nonetheless remain exact so long as its accepted operands and instructions are restricted to those with unique representations that can be encoded within the CPU bit-width.

### 1.1.2 The Question of Technical Risk

It should be emphasized that the novelty of the Hensel architecture resides not in an argument regarding the prospect of exact arithmetic, a critique of floating-point, or an observation that  $\mathbb{Q}_p$  can be used advantageously for exact arithmetic; such arguments can already be found in the literature. The novel value proposition of the Hensel CPU is found in its realization of 2-adic arithmetic at the hardware level. Although a rough design for the architecture is presented here, the technical risks of realizing novel hardware (e.g., arithmetic logical units and processor registers) remain outstanding. Nonetheless, it should be emphasized that an architecture for operands in  $\mathbb{Q}_2$ , whose expansion coefficients can be encoded in bits, should be compatible with extant MOSFET technology. Such compatibility significantly de-risks the Hensel CPU prospect relative to other computing paradigms such as quantum computing, in which case one must develop wholly new electronics, such as transistors, for computing with qubits. Thus, although the Hensel CPU will involve new hardware, it doesn't require a paradigmatic alternative to current electronics and semiconductor technology; it merely requires a new CPU that performs arithmetic in  $\mathbb{Q}_2$  with such technology.

### 1.1.3 The Prize of Exact Arithmetic

In floating-point arithmetic, real numbers (i.e., elements of  $\mathbb{R}$ ) are approximated by "floating-point numbers", which are rationals (i.e., elements of  $\mathbb{Q}$ ). Reals are given decimal representations, which are Cauchy sequences of rationals; in the case of floating-point, these are truncated to be of finite precision. Of course, given finite resources, representations must be finite. The coding-theoretic issue, in the case of  $\mathbb{R}$ , pertains to an analytic issue: real numbers don't have unique Cauchy sequences; they can only be given up to equivalence. Moreover,  $\mathbb{R}$  is in

fact a field of equivalence classes of Cauchy sequences. As a consequence, the accuracy limitations from which floating-point arithmetic suffers can be seen as a consequence of giving finite-precision representations of non-unique approximations of elements of  $\mathbb{R}$ . For instance, suppose, given a bit-width allowing 8 decimal places, one tries to approximate  $\frac{1}{3}$ . One cannot distinguish the approximation from  $\frac{33333333}{100000000}$ ; the representation of  $\frac{1}{3}$  is non-unique.

In the case of  $\mathbb{Q}_2$ , every 2-adic number has a unique 2-adic expansion, which is an infinite series  $\sum_{i=z}^{\infty} a_i 2^i$  (where  $z \in \mathbb{Z}$  and  $a_i \in \{0, 1\}$ ). Writing the coefficients of these series, we obtain unique binary representations. One can compute the exact values of 2-adic expansions of  $n \in \mathbb{Q}_2$  using the convergent properties of infinite series with respect to the 2-adic norm,

$$|n|_2 = 2^{-v_2(n)} \quad (1)$$

where  $v_2$  is the the 2-adic valuation  $v_2 : \mathbb{Q} \rightarrow \mathbb{Z} \cup \infty$  (i.e.,  $\infty$  in the case of  $n = 0$ ). For instance, the following series for  $\frac{1}{3} \in \mathbb{Q}_2$  convergences with respect to  $|\cdot|_2$ :

$$\begin{aligned} \frac{1}{3} &= \\ &1 + 2(1 + 2^2 + 2^4 + \dots) \\ &= 1 + \frac{2}{1 - 2^2} \end{aligned} \quad (2)$$

as is the case for  $\frac{1}{5} \in \mathbb{Q}_2$ :

$$\begin{aligned} \frac{1}{5} &= \\ &1 + 2^2(1 + 2^4 + 2^8 + \dots) + \\ &2^3(1 + 2^4 + 2^8 + \dots) \\ &= 1 + \frac{2^2}{1 - 2^4} \end{aligned} \quad (3)$$

As discussed in this report, in the case of the Hensel CPU, the crux of its value proposition is computation with finite, efficient encodings of 2-adic expansions in a manner that preserves uniqueness and hence, exactness.

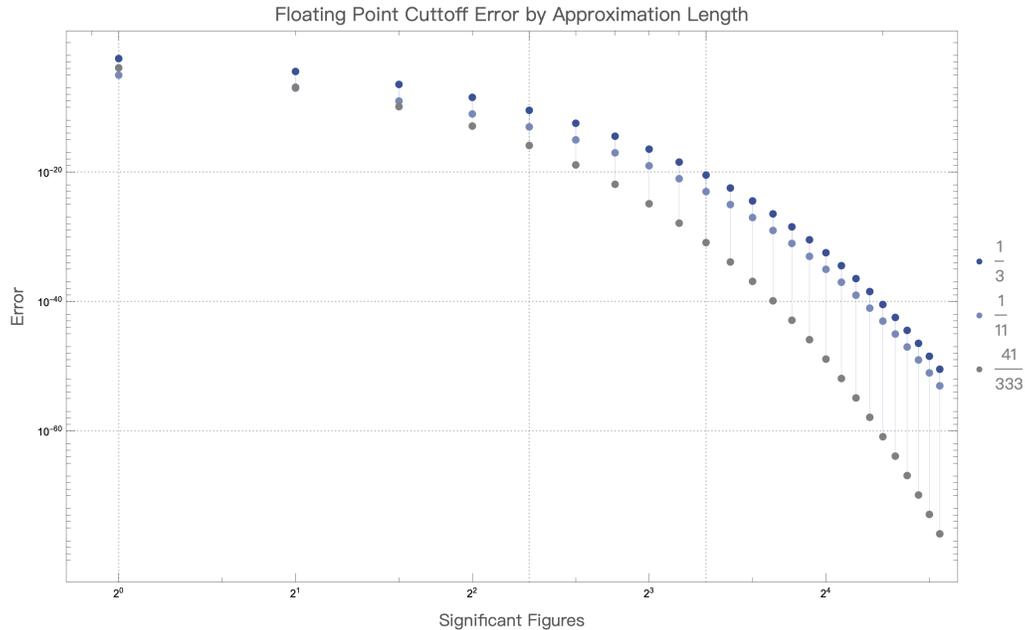


Figure 1: A plot of approximation errors for  $\frac{1}{3}$  (i.e.,  $\frac{1}{3} - \frac{3.3 \times 10^i}{10^{i+1}}$ ,  $i \in \{1, 3, \dots, 49\}$ ),  $\frac{1}{11}$  (i.e.,  $\frac{1}{11} - \frac{\sum_{k=0}^i 9 \times 10^{k+1}}{10^{i+3}}$ ,  $i \in \{2, 4, \dots, 50\}$ ), and  $\frac{41}{333}$  (etc.) as one increases  $i$ .

## 1.2 Report Scope

With reports on the [Virtual Hensel](#), [load-store architecture](#), and [2AALUs](#) now published, the role of this original report on the Hensel CPU can now be lent greater context. It is not a reference on how the computing with the Hensel architecture works. The Virtual Hensel report was written to provide and explain the Virtual Hensel as an elementary proof-of-principle demonstration of exact computation with the Hensel architecture. As a complement to the Virtual Hensel report, the 2AALU report and load-store report offer a (relatively short) overview of how arithmetic and load-store are performed at the level of circuit design and combinational logic. Nonetheless, these two reports, with great regularity, refer back conceptually to this report. Moreover, both the Virtual Hensel, load-store, and 2AALU reports, very much pre-occupied with the practical details of general exact computing capability, seldom discuss matters to do with  $p$ -adic analysis or  $\mathbb{Q}_2$ , which are very much central to this re-

port. Thus, this original report serves, and is expected to serve for some time, as a prerequisite text on the Hensel architecture in a manner that guides the reader from topics in pure mathematics such as 2-adic expansions, Cauchy sequences, and non-archimedean distance to features unique to the Hensel architecture such as FC encodings, 2AALUs, nested carrier packaging,  $\chi$ -addresses, and distributed load-store. In many instances, this report stops short of explaining how architectural features – imbricating computer engineering strategies with number-theoretic affordances – are realized in practice as computing capabilities, instead remarking that the reader can refer to the Virtual Hensel, load-store, and 2AALU reports.

The most rudimentary objective of this report is to introduce how 2-adic numbers are to be stored and operated upon in the Hensel architecture. This requires a coding-theoretic description of 2-adic operands, a data-structural description of how they are

stored, and a logical-functional description of how they are subject to arithmetic operations. In light of the novelty of a CPU architecture designed for  $\mathbb{Q}_2$ , it is necessary to proffer this description in a manner that begins with p-adic analysis and proceeds to computer engineering. With the intersection of these disciplines relatively empty, it is necessary to begin at a fundamental level, describing operands and operations at the level of lists and functions. Doing so will involve relatively simple mathematics but a moderately sophisticated regime of notation for assigning mathematical descriptions to the architectural components of the CPU.

The Hensel CPU architecture is designed for a load-store instruction-set architecture (ISA). This first report is intended to preliminarily introduce the architectural features, especially Hensel processor components, whose load-store roles are most integral. Thus, the report will dedicate particular attention to process registers and arithmetic logical units, as well as the architectural properties with which their design is endowed for 2-adic computing. This report intends to describe their design and function, as well as characterize (and provide mathematical proofs of) properties anticipated to be of utility in HPC.

Looking ahead, beyond the rudimentary presentation given here, more rarefied architectural descriptions – including ISA, microarchitecture, and implementation – will each be given their own subsequent reports. This report is the first of what will be a Hensel series by SciSci Research and its Future Computing group.

## 2 Novel Architectural and Coding Features

This report introduces, in addition to several components of the Hensel CPU architecture, standards used by SciSci and Future Computing for representing 2-adic numbers and instructions (comparable in purpose to standards given for floating-point arithmetic,

such as IEEE 754-1985). These standards are necessarily internal, rather than industrial; that is to say, they are used internally by SciSci and Future Computing as a coding-theoretic basis for architectural design specifications.

### 2.1 Encoding Standards

#### 2.1.1 FC 1-2025 Operand Encoding

This report introduces FC 1-2025 (Future Computing Standard 1-2025 for 2-adic Arithmetic), a coding standard for 2-adic numbers. FC 1-2025 gives an efficient coding of the coefficients of 2-adic expansions, and thus provides an exact, non-approximate encoding of 2-adic numbers (*pace* the  $\mathbb{R}$ -approximations found indelibly in floating-point arithmetic). That is to say, all  $n \in \mathbb{Q}_2$  whose FC 1-2025 encoding is within the bit-width of the CPU can be subject to exact arithmetic. (As discussed in Section 6.2, this bit-width can be very large.)

#### 2.1.2 FC 2-2025 Instruction Encoding

The second standard introduced in this report is FC 2-2025, a standard for encoding parallelizable arithmetic instructions for executing operations on FC-1-2025-encoded 2-adic numbers.

### 2.2 The Hensel Processor

Hensel's arithmetic logical units (ALUs) and processor registers (PRs) are designed according to a novel, nested framework, according to which a moderate number of low-cost ALUs and PRs are assembled into a cluster, with two key prospective advantages (one afforded by nesting and the other by clustering). First, the nested structure is utilized for optimizing storage of FC-1-2025-encoded operands in the PRs and execution of FC-2-2025-encoded instructions by ALUs due to the correspondence between this nested structure and the relationships between 2-adic numbers. (For instance, note that the distance between 2-adic numbers

is non-archimedean; rather than forming a "number line" like the case of  $\mathbb{R}$ , it forms a nested structure.) Second, clustering allows for parallelization of the optimal storage and computing capabilities made available by nested design.

### **2.2.1 2AALUs**

2-adic arithmetic logical units (2AALUs) are arithmetic logical units that perform arithmetic in  $\mathbb{Q}_2$  in parallelized fashion.

### **2.2.2 PR Cluster**

Processor registers are also designed in a cluster,  $\Xi_{PR}$ . Operands are stored in the processor (and main memory) in FC 1-2025 format according to a triple-tree convention discussed in this report. The individual PRs in  $\Xi_{PR}$ , coordinating with the load-store unit (LSU), are instructed by the control unit (CU) to supply operands to  $\Xi_{PR}$  for parallelized execution and to receive outputs.

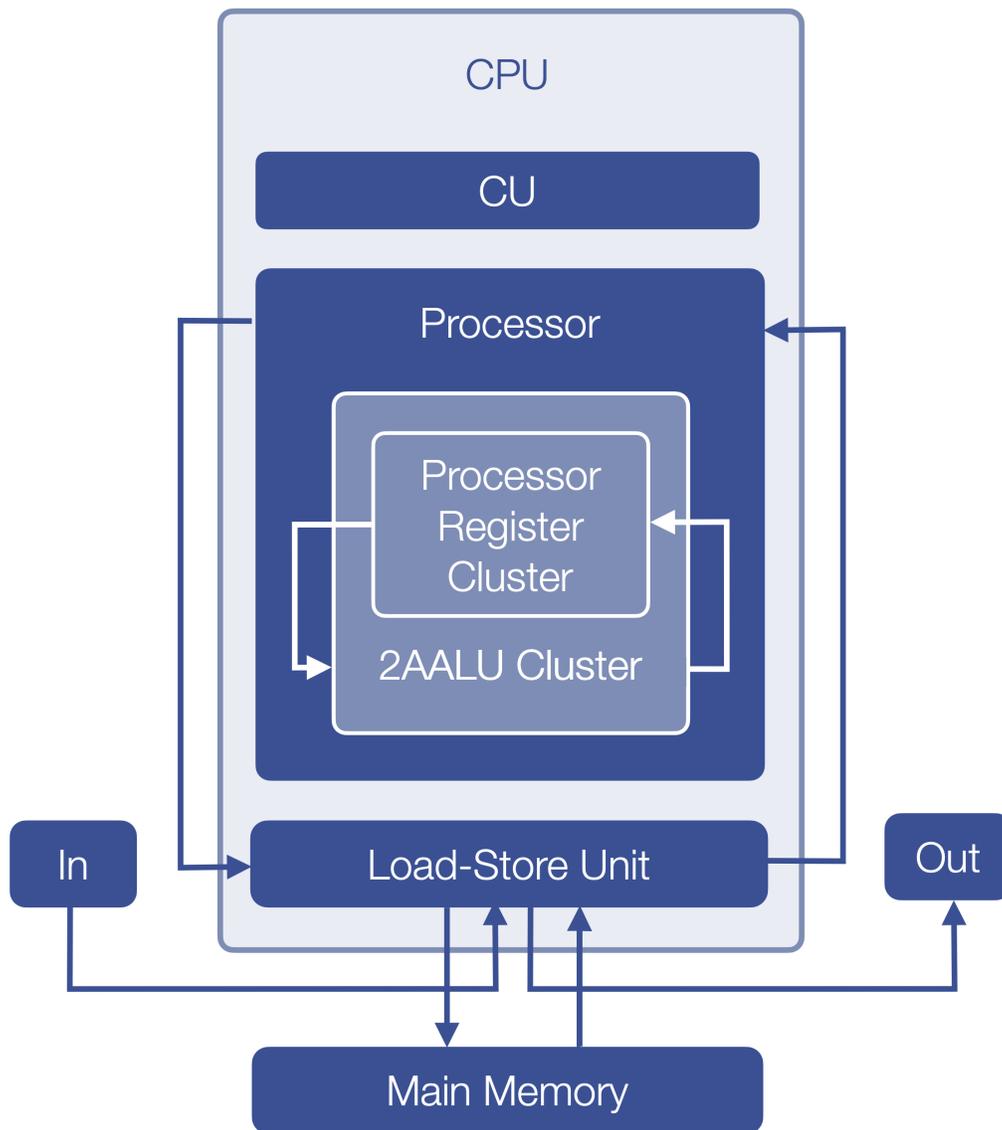


Figure 2: Block diagram of the Hensel CPU, processor, control unit (CU), load-store unit (LSU), and main memory, highlighting 2AALU and processor register cluster relations. (For simplicity of presentation, in- and outgoing arrows with respect to the the CU are omitted.)

### 3 FC 1-2025

#### 3.1 A Quick Review of 2-adic Expansions and Coefficients

Recall from  $p$ -adic analysis that a 2-adic number ( $n \in \mathbb{Q}_2$ ) admits the following expansion

$$n := \sum_{i=z}^{\infty} a_i 2^i, \quad a_i \in \{0, 1\}, \quad z \in \mathbb{Z} \quad (4)$$

These unique expansions give exact sums owing to the convergent properties of their series with respect the 2-adic norm. For instance, in the case of the following,

$$a(2^i + 2^{i+k} + 2^{i+2k} + \dots) = \frac{a}{1-2^k} \quad (5)$$

whereas, in the case of  $\mathbb{R}$ , geometric series  $\sum_{k=0}^{\infty} a_k r^k$  only converge when  $|r| < 1$ , it is the case that 2-adic expansions converge with respect to  $|r|_2$  even when  $|r| \geq 1$ . Properties such as this, as well as the uniqueness of 2-adic expansions, provide a basis for exact arithmetic.

For purposes of introducing the FC 1-2025, we turn our attention to the coefficients  $a_i$ . Let's revisit the example of  $\frac{1}{3} \in \mathbb{Q}_2$ . The first 20 summands (including 0 summands) in its 2-adic expansion are

$$\frac{1}{3} = 1 + 2 + 2^3 + 2^5 + 2^7 + 2^9 + 2^{11} + 2^{13} + 2^{15} + 2^{17} + 2^{19} + \dots \quad (6)$$

(Here, the coefficient for  $2^3$ , for instance, is 1, whereas the coefficient for  $2^4$ , for instance, is 0.) We then write the coefficients from right to left, with an overbar over repeating coefficients. In this case,

$$\frac{1}{3} = \overline{011}1.2 \quad (7)$$

The first 20 summands in the expansion of  $\frac{1}{5}$  are

$$\frac{1}{5} = 1 + 2^2 + 2^3 + 2^6 + 2^7 + 2^{10} + 2^{11} + 2^{14} + 2^{15} + 2^{18} + 2^{19} + \dots \quad (8)$$

which can be abbreviated as

$$\frac{1}{5} = \overline{0011}01.2 \quad (9)$$

Abbreviations of this kind are encoded according to the FC 1-2025 standard as follows.

#### 3.2 FC 1-2025 Encoding

We begin by taking the coefficients  $a_i$  of a 2-adic expansion sequence of a given  $n \in \mathbb{Q}_2$ , which will be ordered from right to left. The coefficient-list-form  $S$  of the sequence for the 2-adic expansion of  $n \in \mathbb{Q}_2$  is as follows:

$$S(n) = (a_{\infty}, \dots, a_1, a_0, \dots, a_{z+1}, a_z) \quad (10)$$

The FC 1-2025 encoding of a given  $n \in \mathbb{Q}_2$  consists of finite sublists of  $S(n)$  which still give a unique encoding:

$$FC(S(n)) := (\perp, R_{FC}, \perp, L_{FC}, \perp, T_{FC}) \quad (11)$$

$R_{FC}$  is the sublist of repeating coefficients in the 2-adic expansion (e.g.,  $\overline{0011}$  for  $\frac{1}{3}$ ). Non-empty  $R_{FC}$  encodings always begin with a 1, and are of length  $\geq 2$ . (For instance,  $R_{FC}(-1) := (1, 1)$ .) Generally, for all sublists, sequences of 0's – e.g., for integers – are treated as empty  $R_{FC}$  encodings.)  $L_{FC}$  is the sublist of non-repeating coefficients to the left of the "decimal" point (e.g.,  $(0, 1)$  for the case of  $\frac{1}{5}$ ), and are minimized so as to not include repeated entries in  $R_{FC}$ , but are written so that  $R_{FC}$  can begin with a 1.  $T_{FC}$  is the sublist of coefficients for summands with negative exponents, typically written to the right of the "decimal point" (e.g.,  $1$  in the case of  $\frac{1}{2}$ , which is written as  $.1_2$ ).  $\perp$  is the "no operation" symbol separating R, L, and T. For instance,

$$FC\left(S\left(\frac{1}{3}\right)\right) := (\perp, (0, 1), \perp, (1), \perp, ()) \quad (12)$$

To take another example,

$$FC\left(S\left(\frac{1}{5}\right)\right) := (\perp, (0, 0, 1, 1), \perp, (0, 1), \perp, ()) \quad (13)$$

Another important example is as follows: by convention,  $FC(S(0)) := (\perp, \perp, (0), \perp)$ . Visual illustrations of several FC 1-2025 encodings are given in Figure 3 and Figure 4.

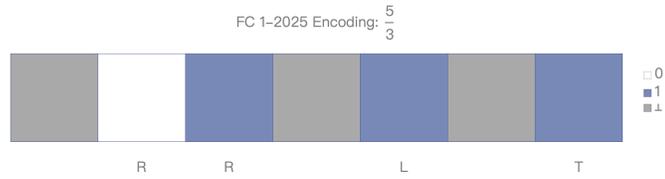


Figure 3: Array plot visualization (with color legend) of the FC 1-2025 Encoding of  $\frac{5}{3}$ .

Rational	2-adic Expansion	2-adic Number	FC2 1-2025 Encoding
$\frac{1}{2}$	$2^{-1} = \frac{1}{2}$	$0.1_2$	
$\frac{1}{5}$	$1 + 2^2(1 + 2^4 + 2^8 + \dots) + 2^3(1 + 2^4 + 2^8 + \dots) = 1 + \frac{2^2}{1-2^4} + \frac{2^3}{1-2^4} = \frac{1}{5}$	$\overline{0011}01_2$	
$\frac{7}{10}$	$2^{-1} + 1 + 2^2(1 + 2^4 + 2^8 + \dots) + 2^3(1 + 2^4 + 2^8 + \dots) = 2^{-1} + 1 + \frac{2^2}{1-2^4} + \frac{2^3}{1-2^4} = \frac{7}{10}$	$\overline{0011}01.1_2$	
$\frac{1}{3}$	$1 + 2(1 + 2^2 + 2^4 + \dots) = 1 + \frac{2}{1-2^2} = \frac{1}{3}$	$\overline{01}1_2$	
$\frac{3}{7}$	$1 + 2^2(1 + 2^3 + 2^6 + 2^9 + \dots) = 1 + \frac{2^2}{1-2^3} = \frac{3}{7}$	$\overline{001}01_2$	
$\frac{22}{7}$	$2 + 2^3 + 2^4(1 + 2^3 + 2^6 + \dots) + 2^5(1 + 2^3 + 2^6 + \dots) = 2 + 2^3 + \frac{2^4}{1-2^3} + \frac{2^5}{1-2^3} = \frac{22}{7}$	$\overline{011}1010_2$	

Figure 4: Table of six 2-adic numbers, their 2-adic expansions, and FC 1-2025 encodings.

### 3.3 Warm-Up: FC-1-2025 Numbers

As a kind of warm-up tutorial for studying finite 2-adic encodings, we'll consider the case of FC-1-2025-encoded 2-adic numbers, a mock standard preceding the FC-3-2025 standard used by the Hensel CPU. One can think of FC-1-2025 encoded 2-adic numbers as given by a "triple-tree" data structure, i.e., three binary trees  $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$  whose parent nodes are glued to a common vertex. With respect to the FC 1-2025 format,  $R_{FC(-)}$  data are stored as  $\mathcal{B}_1$ ,  $L_{FC(-)}$  as  $\mathcal{B}_2$ , and  $T_{FC(-)}$  as  $\mathcal{B}_3$ . ( $\mathcal{T}$ -form examples for FC ( $S(\frac{7}{12})$ ), FC ( $S(\frac{9}{20})$ ), and FC ( $S(\frac{47}{60})$ ) are shown in Figure 6.) Why is it useful to start out thinking this way? Binary tree encoding assigns each  $a_i$  in  $R_{FC}$ ,  $L_{FC}$ , or  $T_{FC}$  to a vertex  $v_i \in \mathcal{B}_k^*$ , beginning with  $v_0^{\mathcal{B}_k}$ ,

with the coefficient assigned thereto indexed as  $a_0$ . Thus, one can think of  $R_{FC}$ ,  $L_{FC}$ , or  $T_{FC}$  as each stored in a  $\mathcal{B}_k^*$  as its own 2-adic number (beginning to the left of the "decimal point" at  $a_0$ ), though nevertheless together giving an FC-1-2025 encoding as a triple-tree data-structure. (Beginning each  $\mathcal{B}_k^*$  encoding with  $a_0$  eliminates the need to encode  $T_{FC}$  using negative indices, which is of consequence for  $\max(\Delta_2^\mu)$  parallelization, as discussed subsequently.) This tree-structure warmup also conveys the notion that the individual  $R_{FC}$ ,  $L_{FC}$ , or  $T_{FC}$  blocks can be regarded as their own number, and indeed, this is how FC-3-2025 numbers are loaded to the processor cluster (via a  $\chi$ -ID system discussed extensively in other reports).

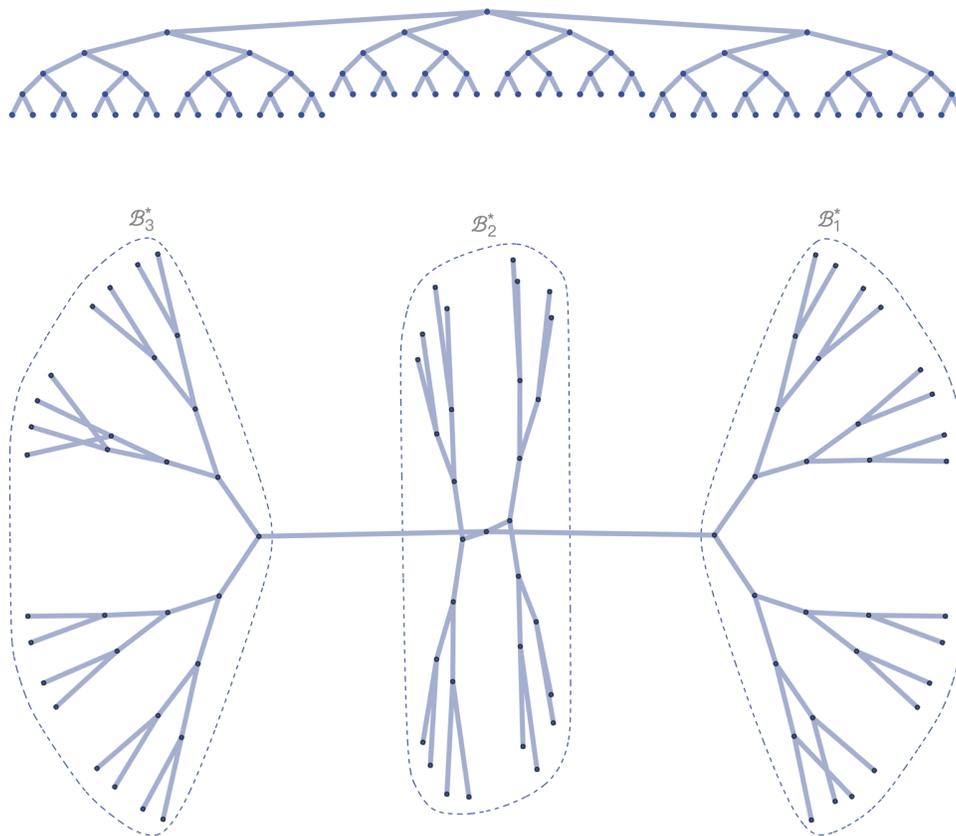


Figure 5: Top: an example illustration of  $\mathcal{T}$ . Bottom: the same illustration of  $\mathcal{T}$  highlighting  $\mathcal{B}_1^*$ ,  $\mathcal{B}_2^*$ , and  $\mathcal{B}_3^*$ . Note that these graph embeddings, for simplicity, collapse  $v_\perp$  onto  $v_0^{\mathcal{B}_2}$ .

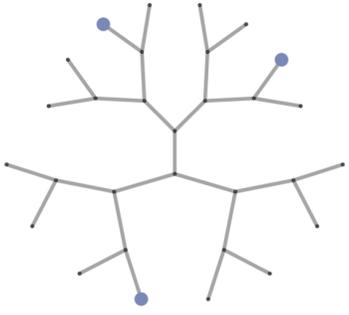
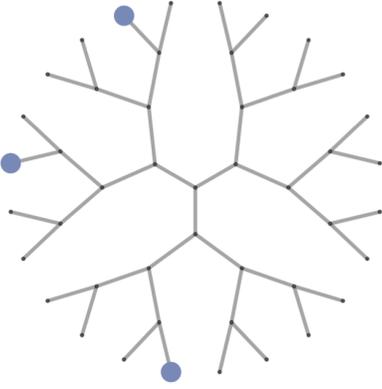
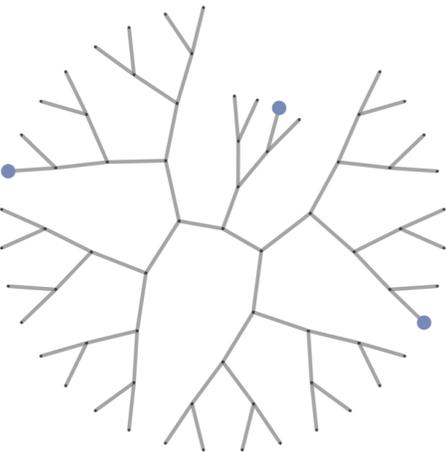
Rational	Triple-Tree Form	FC 1-2025 Encoding
$\frac{7}{12}$		
$\frac{9}{20}$		
$\frac{47}{60}$		

Figure 6: Triple-tree plots for three FC-1-2025-encoded 2-adic numbers. (Here, the embeddings separate  $v_{\perp}$  from the  $v_0^{B_i}$ .)

## 4 The Processor

### 4.1 The Processor Register Cluster

The Hensel CPU architecture features a novel processor design, which is that of a cluster composed of smaller units (i.e., small, low-cost PRs), belonging to a processor register cluster  $\Xi_{PR}$  with a nested structure.

Processor registers, being distributed throughout the cluster, are loaded in distributed fashion by "loaders", or small units, collectively comprising part of the load-store-unit (LSU), that load operands to the appropriate registers in the cluster. The cluster is compartmentalized, with each compartment housing a loader. Loads to (or reloads between) registers handled by loaders operating concert in the LSU circuit. Each compartment is packaged in a carrier, with carriers packaged in nested fashion (whilst still being directly surface-mounted) in correspondence with the topology of the LSU circuit (as described in the load-store report). Loaders in the circuit feed their outputs to one another, with the carriers of the loaders that they feed packaged within their carriers, recurrently, such that the carrier of a given loader at level  $\ell$  feeds its output to an loader at level  $\ell - 1$  whose carrier is packaged within its own carrier, as visualized in Figure 8. The term "level" refers, technically, to the level in the processor circuit tree in which the loader is situated, with a carrier for an loader at level  $\ell - 1$  packaged within the carrier of the  $\ell - 1$  loader that is its parent node in the circuit tree. (See the [load-store report](#) for further details.) The  $\Xi_{PR}$  is a collection of clustered processor registers,  $\mathcal{C}_{PR}$ , whose containers are packaged in the lowest-level cluster carriers, as shown in Figure 7. Thus, the overall cluster consists of nested-packaged loader carriers, with processor register carriers packaged at the innermost level. Loaders positioned at different levels can modify different  $a_i$ , and, doing so simultaneously at different levels, parallelize loader operations. This is achieved via execution of FC-2-2025-encoded instruc-

tions, as discussed subsequently.

Operands are loaded to  $\Xi_{PR}$  in distributed fashion. Given the  $\mathcal{T}$ -form of a given  $FC(S(n))$ , the individual  $\mathcal{B}_k^*$  are loaded to distinct clustered processor registers in  $\Xi_{PR}$ . The PR cluster consists of a master PR,  $\mathcal{M}_{PR}$ , as well as clustered PRs, written as  $\mathcal{C}_{PR}$ . Together, they comprise the  $\Xi_{PR}$ , which we write in notation as a set of PRs  $\Xi_{PR} := \left( \bigcup_{i=1}^{\mathcal{N}} \mathcal{C}_{PR_i} \right) \cup \mathcal{M}_{PR}$  (where  $\mathcal{N} < 2^B$ , with B being the architecture bit-width). The individual  $\mathcal{B}_k^*$  are loaded to specific  $\mathcal{C}_{PR}$  according to individual encoding blocks called  $\chi$ -IDs, discussed later in this report (and to greater effect in the [Virtual Hensel report](#)). A given  $\chi$ -ID is loaded to the  $\mathcal{C}_{PR}$  whose address matches the  $\chi$ -ID entries.  $\mathcal{M}_{PR}$  is responsible for reassembling ("recompounding") these individual blocks back to a whole operand for storage purposes. Each  $\mathcal{C}_{PR}$  is loaded via receipt of an "activation input", which we'll term a  $\pi$ -sequence, issued by  $\mathcal{M}_{PR}$  as prompted by the LSU. A  $\pi$ -sequence both activates a given  $\mathcal{C}_{PR}$  and encodes information about the kind of  $\chi$ -ID to be loaded (e.g., for a  $R_{FC}$ ,  $L_{FC}$ , or  $T_{FC}$  block). Thus, to load an operand, the  $\mathcal{M}_{PR}$  loads the  $\chi$ -IDs of its constituent blocks to the address-matching  $\mathcal{C}_{PR}$ , and tells the  $\mathcal{C}_{PR}$  what kind of block is being loaded. Operands can be reassembled ("recompounded") with this information. The operands can also be subject to arithmetic operations, resulting in new  $\chi$ -IDs, which are loaded to new  $\mathcal{C}_{PR}$  with matching addresses, and recompounded by the  $\mathcal{M}_{PR}$  to obtain the output.

The LSU, via its loaders, performs load-store on operands stored according to parallelizable instructions. Given some  $FC(S(q))$  subject to an arithmetic operation  $\mu$  yielding output  $\mu(FC(S(q)))$  these instructions amount to modifying the  $FC(S(q))$  address-matching  $\chi$ -ID in order to obtain the address-matching  $\chi$ -ID for the output  $FC(S(\mu(q)))$ . The instructions guide the loaders at nest levels  $2 \leq \ell \leq \mathcal{L} - 1$  in modifying the  $\chi$ -ID entries in parallel, where modifications of coefficients at lower  $i$  are performed by  $\mathcal{A}_{\ell,j}$  at lower  $\ell$  (i.e., innermost in the nest-structure) and greater  $i$  by

$\mathcal{A}_{\ell,j}$  at greater  $\ell$  (i.e., outermost in the nest-structure).

## 4.2 Nested-Clustered Design

In the Hensel architecture, each cluster is to be physically designed in nested form. From an engineering perspective, one can build carriers of differing sizes and mount them to the printed circuit board with carriers at lower levels packaged within carriers at levels, as shown in Figure 7. The process register carriers are, in turn, to be packaged inside the lowest-level cluster carriers, and thus the most deeply nested within the carrier packaging structure, as shown in Figure 8. It is this choice of inter-carrier packaging that gives the nested structure of the processor. (See the [load-store report](#) for further discussion.)

A cluster of nest depth  $\mathcal{L}$  is designed so that, proceeding from the innermost loader level  $\ell = 2$  to level  $\mathcal{L} - 1$  (with level 1 storing the  $\Xi_{PR}$  and level  $\mathcal{L}$  storing the  $\mathcal{M}_{PR}/\mathcal{M}_{LSU}$ ), there are  $2^{\mathcal{L}-\ell}$  loaders  $\{\xi_{\ell,2^{\mathcal{L}-\ell}}, \dots, \xi_{\ell,1}\}$  at each level  $\ell$ . Each carrier at level  $\ell$  will contain  $2^{\ell-1}$  loader carriers (with the exception of the  $\xi_{2,j}$ , whose carrier packages the the  $\mathcal{C}_{PR}$  carriers) and will be packaged alongside another carrier by a carrier at level  $\ell + 1$  (with the exception of  $\{\xi_{\mathcal{L}-1,1}, \xi_{\mathcal{L}-1,2}\}$ , which are packaged within the  $\mathcal{M}_{PR}$  carrier). We can describe the nest structure of the register cluster in terms set membership, where, as a shorthand,  $\xi_{\ell,j} \equiv \{\}\_{\ell,j}$ . With this shorthand, we can write the nest structure as follows, beginning at  $\ell = 2$  and moving outward by  $\ell + 1$ :

$$\{\}\_{\ell+1,j} := \begin{cases} \{\{\}\_{\ell,j}, \{\}\_{\ell,j}\} & (\ell - 1)2 \\ \{\{\}\_{\ell,j}, \{\}\_{\ell,j}\} & (\ell - 1) \uparrow 2 \end{cases}, \quad \{\}\_{2,j} = \{\{\}, \{\}\} \quad (14)$$

## 4.3 Cluster Load-Store and Addressing

Operands are subject to load-store in distributed fashion, with the  $R_{FC(S(-))}$ ,  $L_{FC(S(-))}$ , or  $T_{FC(S(-))}$  blocks of an operand's encoding each loaded to a different  $\mathcal{C}_{PR}$ . As described in greater detail in the [Virtual Hensel report](#), IDs, called  $\chi$ -IDs, are generated from these individual blocks, which are similar to the blocks themselves, but subject to a few coding tricks that permit a sizeable quantity of blocks to be subject to load-store by a relatively small number of processor registers. The address  $\mathfrak{A}$  assigned to a given  $\mathcal{C}_{PR}$  is simply the  $\chi$ -ID of the  $R_{FC(S(-))}$ ,  $L_{FC(S(-))}$ , or  $T_{FC(S(-))}$  encoding block load that it accepts. We denote a  $\mathcal{C}_{PR}$  with address  $\mathfrak{A}$  as  $\mathcal{C}_{PR}(\mathfrak{A})$ . Thus, distributed load-store is a matter of ID-address matching (i.e.,  $\chi$ - $\mathfrak{A}$  matching).

A given  $\mathcal{C}_{PR}$  is loaded by activating it with a  $\pi$ -sequence, which contains an encoding that distinguishes the  $\chi$ -ID by its block type (i.e.,  $R_{FC(S(-))}$ ,  $L_{FC(S(-))}$ , or  $T_{FC(S(-))}$ ), such that an operand, although being subject to load-store with its encoding broken down ("decompounded") into its constituent blocks, can nonetheless be recovered ("recompounded"). The LSU instructs the  $\mathcal{M}_{PR}$  to load a  $\mathcal{C}_{PR}$  as follows:

$$\lambda : (\chi, \pi) \rightarrow \mathcal{C}_{PR}(\mathfrak{A}_{\pi}^{\chi}) \quad (15)$$

The  $\lambda$  mapping, as written above, is in fact a simplified description of what is done in practice by the loaders. See the [Virtual Hensel report](#) for a finer description of loads, and the [load-store report](#) for a finer description of re-loads, which are applied to output operands.

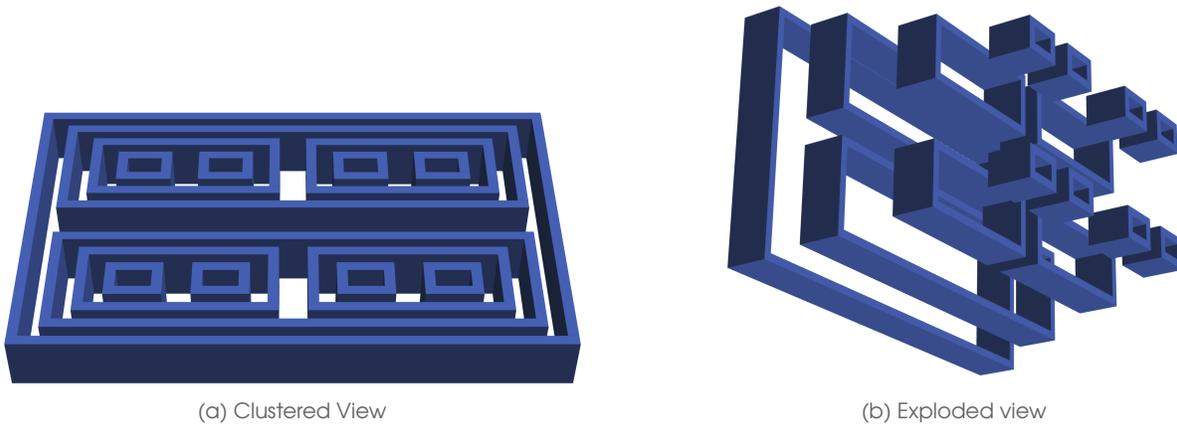


Figure 7: Illustration of the nested structure in the register cluster.

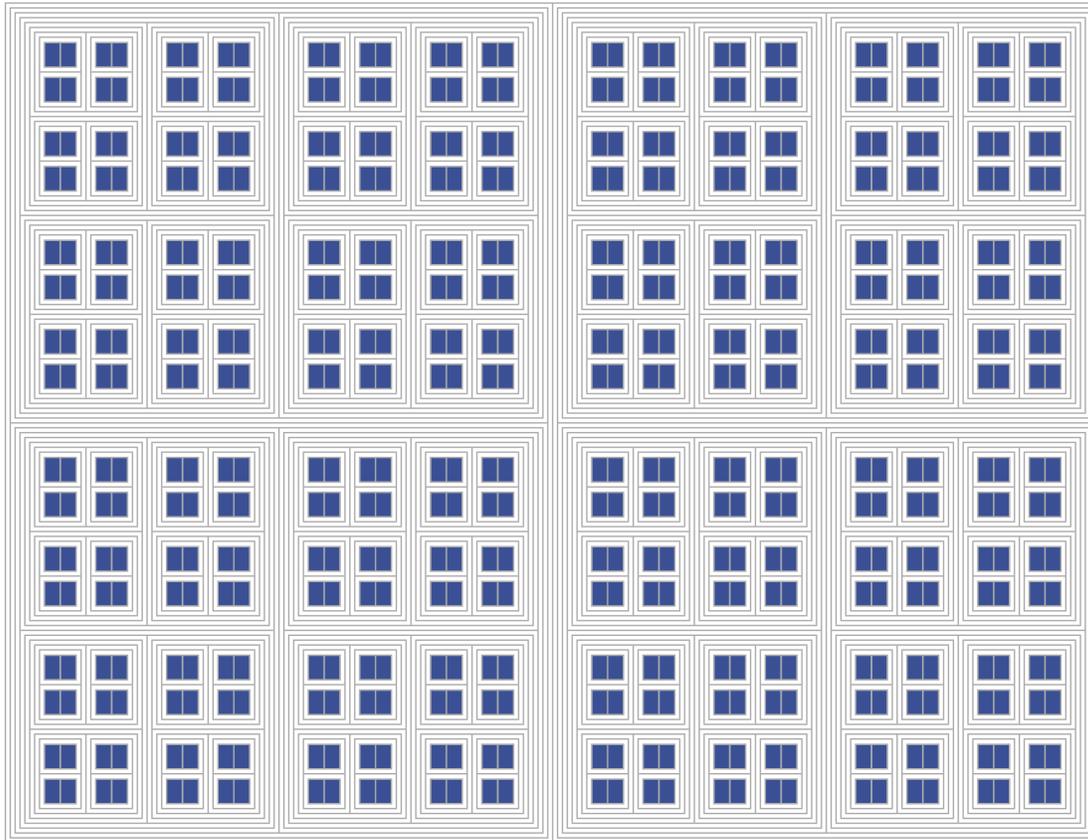


Figure 8: Cluster distribution of  $C_{PR}$  at the innermost level of the register cluster. (Note that the number of nested layers in this example is greater than in Figure 7, simply for purposes of visual variety.)

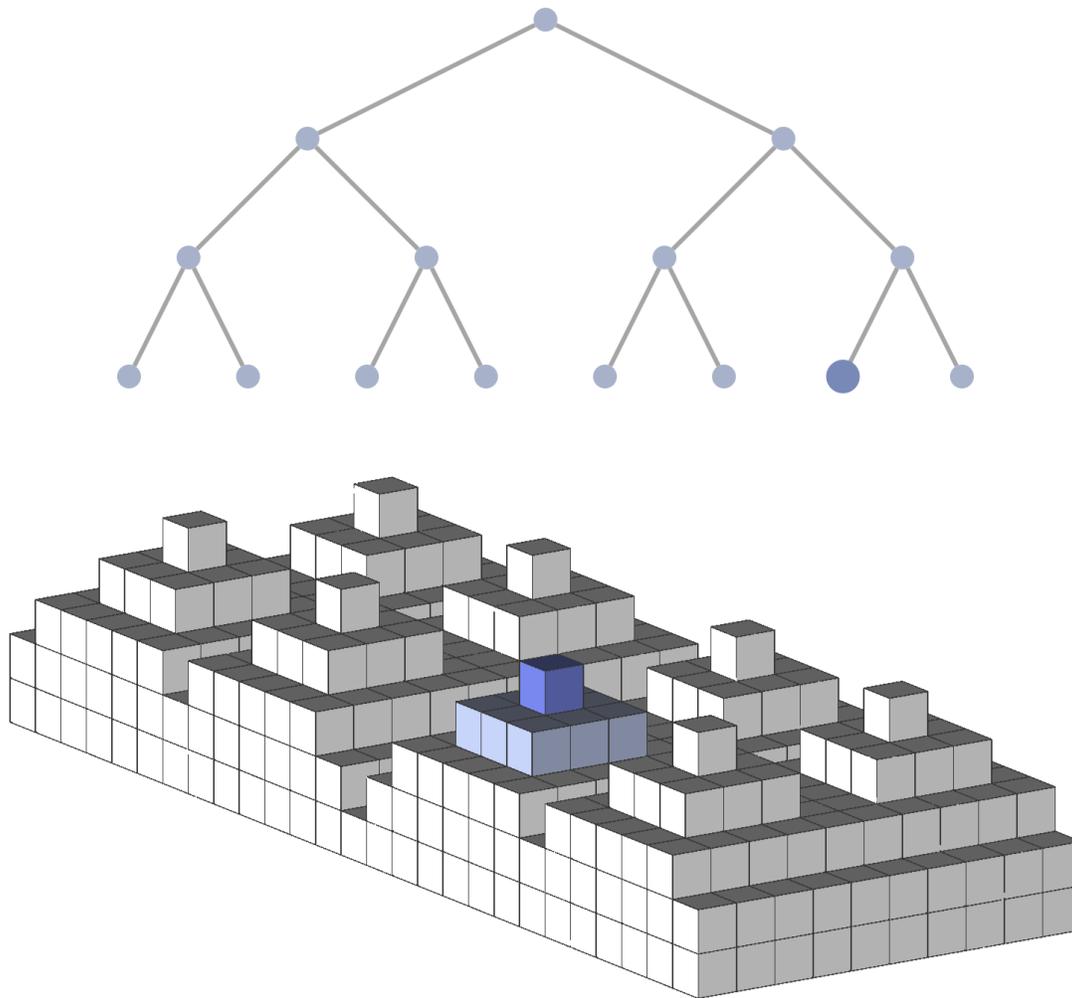


Figure 9: Illustration of the correspondence between  $\mathcal{B}_k^*$ -form tree structure and  $\mathcal{C}_{PR}$  cluster position. In this and other similar illustrations, the  $\mathcal{C}_{PR}$  are positioned at the "top", with layers "below" being loader layers. Here, the top-most layers is nest-innermost in the cluster (i.e., at level  $\ell = 1$ ) and the bottom-most layer is nest-outermost in the cluster (i.e., at level  $\ell = \mathcal{L} - 1$ ).

## 5 Parallelization

### 5.1 Parallelization

2AALUs perform arithmetic in parallelized fashion. For instance, in the case of dyadic addition, operations on each pair of entries in the two  $\chi$ -IDs of the operands are done in parallel. Arithmetic is thus reduced to a two step process: parallelized arithmetic without "extra digits" (i.e., from carrying), done in parallel, and a second step for adding extra "digits" if necessary, which is also done in parallel. Finally, a computation of an FC-2-2025 instruction is performed in order to guide the LSU in handling input and output operands in the register cluster. The 2AALU report describes the above arithmetic procedures in greater detail. We will now attend to load-store.

### 5.2 Parallelized Load-Store

FC 2-2025 instructions guide the loading and reloading of operands in the register cluster. For a load, the instructions direct the LSU loaders to load the relevant operand  $\chi$ -ID to the register with the matching address. In the case of a reload, the FC-2-2025 instructions tell the loaders, which entries in the address need to be changed.

### 5.3 FC 2-2025 Encoding

Modification of entries in  $\chi$ -IDs is nothing more than mere bit-flipping (e.g.,  $(1, 1, 1) \rightarrow (1, 1, 0)$ ). Indeed, as discussed in the load-store report, at the circuit level, modifications are executed in combinational logic by a given loader by changing the input value (i.e., 0 or 1) for the logic gates at a particular level. At times, it is expedient to elide discussion of combinatorial logic and instead denote loader operations in terms of "hop calculus", which, rather than explain the circuit-level combinational logic underpinning loader operations, describes the effect of loader modifications, due to their effect on a given  $\chi$ -ID, on corresponding changes

in  $\mathcal{C}_{PR}$  to which the  $\chi$ -ID will be loaded. Indeed, a single loader modification can alter a  $\chi$ -ID such that its new address-matching  $\mathcal{C}_{PR}$  is located in a notably different location in the processor cluster. However, because  $\mathcal{C}_{PR}$  are located at the ends of the circuit in a specified manner (as discussed in the [load-store report](#)), the change-in-address consequent of loader modifications is predictable, and very much related to the notion of 2-adic distance. A change-in-address can be thought of as a "hop" across the processor, and the loader modification a "hop operation".

Each loader in the LSU circuit performs an operation and passes its output  $\tau$  to the next loader downstream in the circuit tree. A loader can pass a  $\tau$  via two kinds of hops. Because each loader in the circuit tree (at level  $\ell > 2$ ) has two children, i.e.,  $\{\xi_{\ell,1}, \xi_{\ell,2}\}$  (and because, accordingly, each loader carrier at level  $\ell + 1$  packages two loader carriers at level  $\ell$ ), it is the case that  $\tau$  can be passed according to one of two hops:

$$\tau h_{\ell}^{\sigma} : \mathcal{A}_{\ell,1} \rightarrow \mathcal{A}_{\ell,2} \quad (16)$$

or

$$\tau h_{\ell}^{\bar{\sigma}} : \mathcal{A}_{\ell,2} \rightarrow \mathcal{A}_{\ell,1} \quad (17)$$

Hereafter, we will write them simply as  $h_{\ell}^{\sigma}$  and  $h_{\ell}^{\bar{\sigma}}$ , with  $\tau$  implicit. (Alternatively, the loader can perform no operation: a non-hop  $h_{\ell}^{\perp}$ .)

The standard encoding for hop instructions, FC 2-2025, is as follows:

$$(\perp, \mathfrak{P}_R, \perp, \mathfrak{P}_L, \perp, \mathfrak{P}_T) \quad (18)$$

where  $\mathfrak{P}_R$  is a sublist of instructions for  $R_{FC}$ ;  $\mathfrak{P}_L$ , for  $L_{FC}$ ; and  $\mathfrak{P}_T$ , for  $T_{FC}$ . The elements populating these  $\mathfrak{P}_{(-)}$  sublists are  $h_{\ell}^{\sigma}$  and  $h_{\ell}^{\bar{\sigma}}$  (as well as  $h_{\ell}^{\perp}$ ) where  $h_{\ell}^{\sigma}$  is encoded as  $(0, 1)$  and  $h_{\ell}^{\bar{\sigma}}$  is encoded as  $(1, 0)$ . (Additionally, a non-hop is encoded as  $(0, 0)$ .) Each instruction is in turn separated by a  $\perp$ .

FC 2-2025 encoding can be applied to instructions for addition, subtraction, multiplication, or division operations whose inputs and outputs are  $n \in \mathbb{Q}_2$  within the allowed bit-width, with the advantage of forgoing

operations such as carrying in carry arithmetic; the loader just perform modifications according to stored instructions. In practice, the Hensel performs arithmetic on  $\chi$ -IDs, which are encoded according to the FC-3-2025 standard, as introduced in the [Virtual Hensel report](#). In this report, for introductory purposes, we'll consider arithmetic on FC-1-2025 encoded operands. (We'll write "FC\* 2-2025" to describe instructions on FC-1-2025 encoded operands, since, in practice, they are encoded for  $\chi$ -IDs.) For instance, recall that the FC 1-2025 encoding for  $\frac{1}{3}$  is  $(\perp, 0, 1, \perp, 1, \perp)$ . Adding  $\frac{1}{5}$  and  $\frac{1}{3}$  is encoded in the following FC\* 2-2025 instruction:

$$(\perp, \perp, (1, 0), \perp, (0, 1), \perp, (0, 1), \perp, (0, 0), \perp, \perp, (0, 1), \perp, (1, 0), \perp, (1, 0), \perp, (0, 1), \perp, \perp) \quad (19)$$

To be more precise, this instruction performs a modification on  $FC(S(\frac{1}{3}))$  and returns  $FC(S(\mu^{(+, \frac{1}{5})}(\frac{1}{3}))) = FC(S(\frac{8}{15}))$ .

However, for purposes of evaluating parallelization performance, it is convenient to write FC(\*)-2-2025-encoded instructions in parallelized form, or  $\mathfrak{P}$ -form, where the  $\mathfrak{P}$ -form of a given  $\mathfrak{P}_{(-)}$  instruction list for a modification  $\mu(FC(S(q)))$  is given in terms of hop calculus. It is no more than a list  $\mathcal{H}$  of  $h_\ell^{(-)}$  instructions, which are written from right to left for loaders from levels  $\ell = 2$  to  $\ell = \mathcal{L} - 1$ :

$$\mathfrak{P}(\mu(FC(S(q)))) =: \mathcal{H} = (h_m^{(-)}, \dots, h_2^{(-)}) \quad (20)$$

where  $\mathcal{L} - 1 \geq m$ . Thus, a parallel computation performing a modification  $\mu$  (written  $\mathfrak{P}(\Xi(\mu(-)))$ ) of depth  $\mathcal{D} = \text{Length}(\mathcal{H}) = m - 1$ , will execute  $(m - 1)$ -many  $h_\ell^{(-)}$  operations. When  $\chi$ -modification is performed for  $R_{FC}$ ,  $L_{FC}$ , or  $T_{FC}$ , each is treated as a separate number and, each having its own  $\mathfrak{P}_{(-)}$ , has its own  $\mathfrak{P}$ -form, with  $\chi$  entries beginning at  $a_0$  and operations beginning at  $h_2^{(-)}$ .

## 5.4 Parallelization and Non-Archimedean Distance

Hensel's nested structure is designed to optimally store, and efficiently compute with, numbers in  $\mathbb{Q}_2$  in a manner commensurate with 2-adic arithmetic. For instance, the Hensel processor efficiently maximizes 2-adic output-operand distance over minimal operations because its nested structure is commensurate with 2-adic distance, which, unlike the archimedean case of  $\mathbb{R}$ , which is given with respect to a number line, is instead non-archimedean and gives a nested structure; the PR cluster is designed according to this structure.

Given the nested structure of the cluster, a hop operation affects, at level  $\ell = 1$ , the board location of the output-ID-matching PR. Change-in-location is more distal when hops are performed at higher  $\ell$  than at lower  $\ell$ . Whereas at lower  $\ell$ , hop operations can be performed all the while remaining nested within the same loader packaging at higher  $\ell$ , hops at higher  $\ell$  in turn also affect all lower  $\ell$ , due to nested packaging. Thus, one might suppose that the greater  $\ell$  is, the greater the effect of a hop operation on the output. This would be rather inefficient, as it would imply that clearing greater arithmetic distances require  $\pi$ -sequence passing over greater board distances. However, with respect to 2-adic distance, one finds the opposite to be the case: the lower- $\ell$  hops have the greatest affect on the output. This is a consequence of triple-tree data-structure and nested-clustered design.

Recall the 2-adic distance between  $p, q \in \mathbb{Q}_2$ :

$$|q - p|_2 = 2^{-v_2(q-p)} \quad (21)$$

where  $v_2$  is the 2-adic valuation. Given an operand  $q$  and output  $\mu(q)$ , we'll denote the 2-adic output-operand distance as follows:

$$\Delta_2^\mu(q) := |\mu(q) - q|_2 \quad (22)$$

Let's consider, in the parallelized case, the relationship between  $\Delta_2^{\mathcal{H}}(-) = \sum_{\ell}^{\mathcal{L}-1} \Delta_2^{h_\ell^{(-)}}(-)$

and  $\mathfrak{D}(\mathcal{H})$ , the parallelization depth. Notably, greater  $\mathfrak{D}(\mathcal{H})$ , requiring  $h_\ell^{(-)}$  at ever-greater  $\ell$ , probes smaller 2-adic distances per  $h_\ell^{(-)}$  operation, or to be more precise:

$$\max \left( \Delta_2^{h_\ell^{(-)}}(\text{FC}(S(q))) \right) \propto \frac{1}{\mathfrak{D}(\mathcal{H})} \quad (23)$$

Thus, for greater  $\mathfrak{D}(\mathcal{H})$ , the upper bound on  $\Delta_2^{h_\ell^{(-)}}$  decreases with each additional level.

For instance, modification of  $a_0$  has the greatest effect on  $\Delta_2^\mu$ . Consider the case of  $q = 1$  and  $\mu := (+, 1)$ . In this case,  $\Delta_2^{(+,1)}(1) = 2^0 = 1$ . This only involves modification of  $a_0$  (thus,  $\mathfrak{D} = 1$ ). Or, consider the case of  $a_1$ -modification: take the case where  $q = 1$  and  $\mu := (+, 2)$ . In this case,  $\Delta_2^{(+,2)}(1) = 2^{-1} = \frac{1}{2}$ , a smaller distance. Of course, modification of  $a_0$  can affect smaller

distances too. For instance, if  $q = 32$  and  $\mu := (+, 1)$ , then  $\Delta_2^{(+,32)}(1) = 2^{-5}$ . However, there is no modification of  $a_{i>1}$  that gives a distance  $\Delta_2^{h_\ell^{(-)}} \geq 2^0$ .

At first glance, the above statement may seem surprising, for, in the usual case of  $a_i$ -modification, one can indeed obtain values  $\Delta_2^\mu > 1$  when negative exponents appear in the 2-adic expansion of the operand or output, that is, coefficients to the right of the "decimal point". However,  $\mathfrak{A}$ -form encodings give instructions for  $\mathcal{T}$ -form FC-1-2025 operands, which are of  $\mathcal{B}_k^*$  structure.  $T_{FC}$  entries, which do correspond to negative exponents in a typical 2-adic expansion, are encoded in  $\mathcal{B}_3^*$  beginning with  $a_0$  and thus non-negative.  $\mathcal{T}$ -form abolishes index-negativity and thereby bounds output-operand distance to  $0 \leq \Delta_2^\mu \leq 1$ .

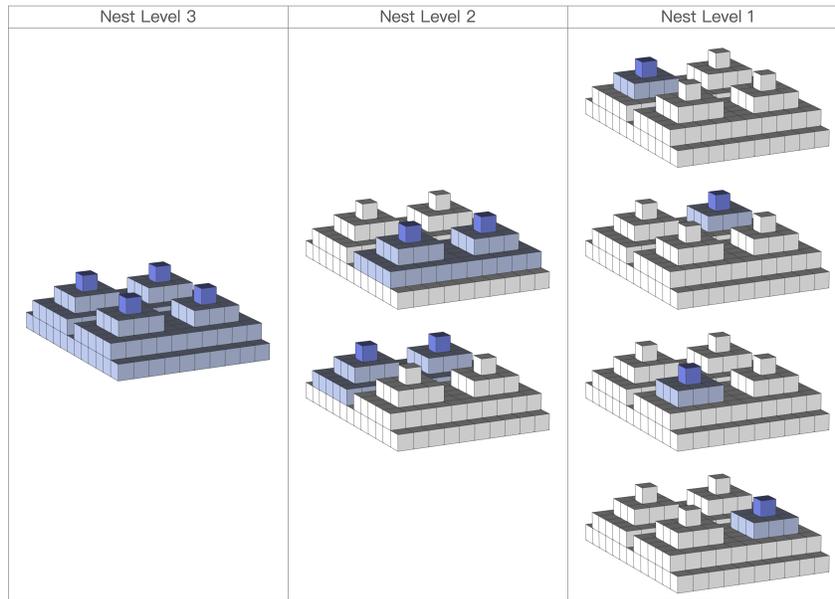


Figure 10: Table highlighting, per level, the loaders a hop can reach (in light blue) and the PRs they can reach (in dark blue).

### 5.5 Max $(\Delta_2^{h_\ell})$ Parallelization

If we look at  $\Delta_2^{h_\ell}$  per level  $\ell$ , i.e., per  $\Delta_2^{h_\ell^{(-)}}$ , we find that the inner-to-outer level execution of right-to-left-indexed  $\mathfrak{P}$ -form instructions necessarily maximizes, relative to depth  $\mathfrak{D}$ , the distance  $\Delta_2^{h_2^{(-)}}$  that can be cleared per level  $\ell$ , as shown in Figure 11. As a consequence, given some  $\mathcal{A}_{2,i}$  performing a  $h_2^{(-)}$  hop,  $\max(\Delta_2^{h_2^{(-)}}$ ) is greater than  $\max(\Delta_2^{h_{i>2}^{(-)}}$ ).

Such is advantageous for parallelization, as it means that  $\mathcal{A}_{\ell,j}$  are maximally level-efficient, per  $\ell$  (and with respect to  $\mathfrak{D}$ ), in bringing the operand to the output value. Thus, the  $\mathfrak{D}$  values for parallel computation can readily be minimized by following the principle of  $\mathfrak{D}$ -minimization through  $\Delta_2^{h_\ell^{(-)}}$ -maximization, where  $\mathfrak{D}$ -minimization necessarily minimizes the number of loaders involved, and is thus a benchmark for parallelization efficiency. The relationship between nest level execution of FC 2-2025-encoded  $\mathfrak{P}$  instructions is illustrated, for some rudimentary cases, in Figure 12.

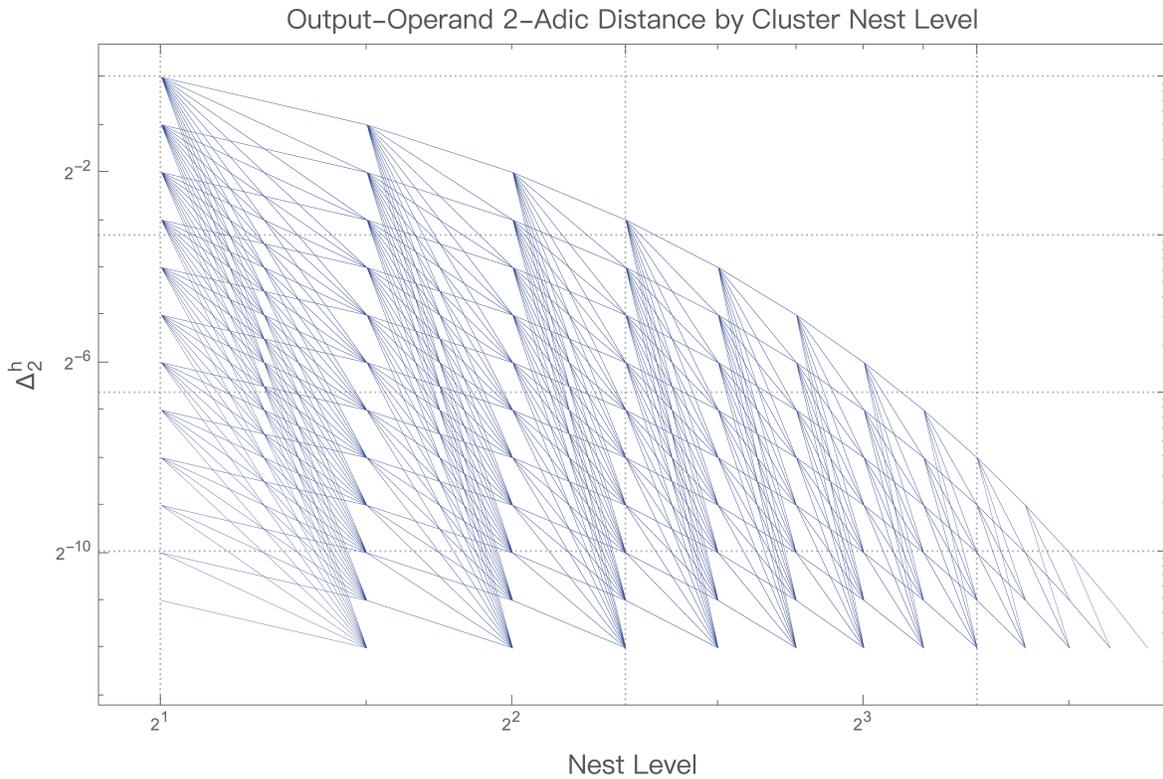


Figure 11: Plot showing possible  $\Delta_2^{h_\ell^{(-)}}$  values per  $\ell$  level, where  $2 \leq \ell \leq 14$ .

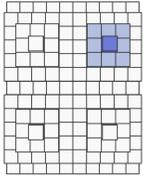
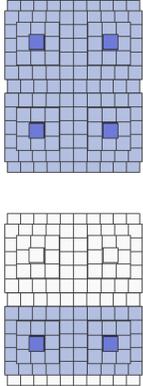
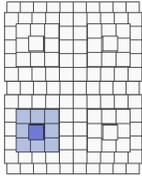
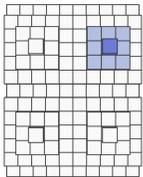
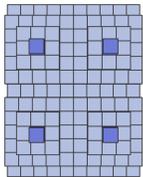
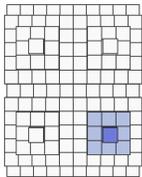
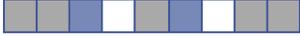
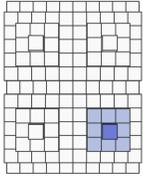
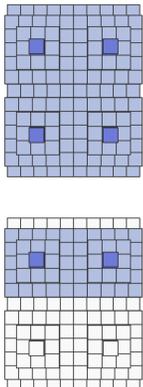
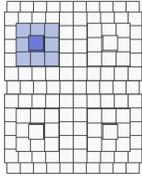
Sum	FC 2-2025 Instructions	Loader Operations		
1+1		Operand: 1 	$h_2^\sigma, h_3^{\bar{\sigma}}$ 	Output: 2 
1+2		Operand: 1 	$h_2^\sigma$ 	Output: 3 
3-3		Operand: 3 	$h_2^{\bar{\sigma}}, h_3^{\bar{\sigma}}$ 	Output: 0 

Figure 12: Elementary visualizations of FC\*-2-2025-encoded instructions.

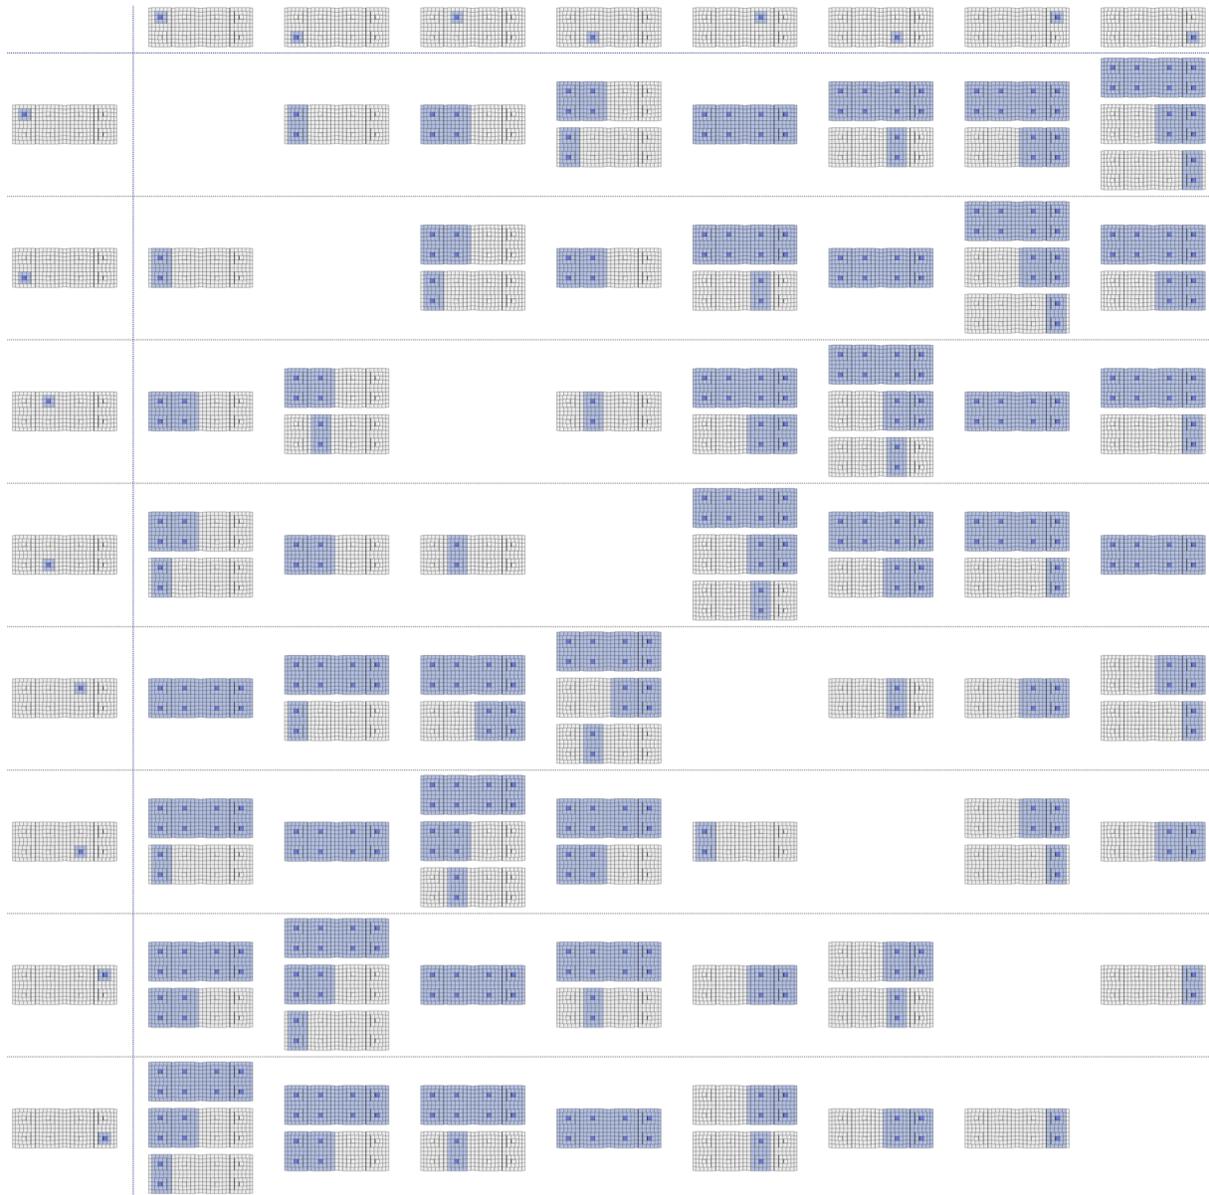


Figure 13: A  $\mathfrak{B}$ -form "instruction table", where the columns are  $C_{PR}$  whose  $\chi$ -addresses are the  $\mathcal{T}$ -form FC 1-2025 operand inputs  $\{0, 1, 2, 3, 4, 5, 6, 7\} \in \mathbb{Q}_2$ , and the rows are the same  $C_{PR}$  whose  $\chi$ -addresses are taken as outputs. Shown in the table are the  $h_\ell^{(-)}$  operations in the instructions that return each output from each input.

### 5.6 Constraints on Optimization

Although parallelization lends the efficiency of simultaneity to computation, and the  $\max(\Delta_2^H)$  property makes operations at each nest level economical, there must necessarily exist constraints on parallelization optimality, which should be articulated so as to measure performance relative thereto. The key constraint, of course, is a depth constraint: certain computations will require many parallel operations. Depth costs are incurred when operands and outputs are many hops removed from one another in  $\Xi_{PR}$ , which will ineluctably require parallelized computation involving several loaders to effectuate hops.

As an example, let's consider possible values of  $q$ , namely  $\{0, 1, 2, 3, 4, 5, 6, 7\} \in \mathbb{Q}_2$ . Let's also consider 8 values for  $\mu(q)$ , which will also be  $\{0, 1, 2, 3, 4, 5, 6, 7\}$ . Figure 13 provides a table showing the  $\mathcal{H}$  needed to obtain a given  $\mathcal{T}$ -form of  $\mathfrak{P}(\Xi(\text{FC}(S(\mu(q)))))$  from the  $\mathcal{T}$ -form of  $\text{FC}(S(q))$ . The following  $8 \times 8$  matrix

gives the  $\mathfrak{D}(\mathcal{H})$  values for each:

$$M = \begin{pmatrix} 0 & 1 & 1 & 2 & 1 & 2 & 2 & 3 \\ 1 & 0 & 2 & 1 & 2 & 1 & 3 & 2 \\ 1 & 2 & 0 & 1 & 2 & 3 & 1 & 2 \\ 2 & 1 & 1 & 0 & 3 & 2 & 2 & 1 \\ 1 & 2 & 2 & 3 & 0 & 1 & 1 & 2 \\ 2 & 1 & 3 & 2 & 1 & 0 & 2 & 1 \\ 2 & 3 & 1 & 2 & 1 & 2 & 0 & 1 \\ 3 & 2 & 2 & 1 & 2 & 1 & 1 & 0 \end{pmatrix} \quad (24)$$

Note that entries in  $M$  near the left-to-right diagonal are small, whereas the opposite is true for the right-to-left diagonal; in the latter case, the operands and outputs are stored in  $\mathcal{C}_{PR}$  that are hop-distal from one another in the cluster. The values near these diagonals are sufficiently disparate such that if one interpolates a surface  $\mathfrak{x}$  from the matrix (i.e., with coordinates  $(i = q, j = \mu(q), M_{i,j})$ ), the gap between high- and low- $\mathfrak{D}(\mathcal{H})$  along the diagonals gives a hole in the surface (see Figure 14). Heuristically, one can think of this "topological invariant" in the interpolated surface as exhibiting the difference in  $\mathfrak{D}(\mathcal{H})$  between hop-proximal and hop-distal  $\mathcal{C}_{PR}$ .

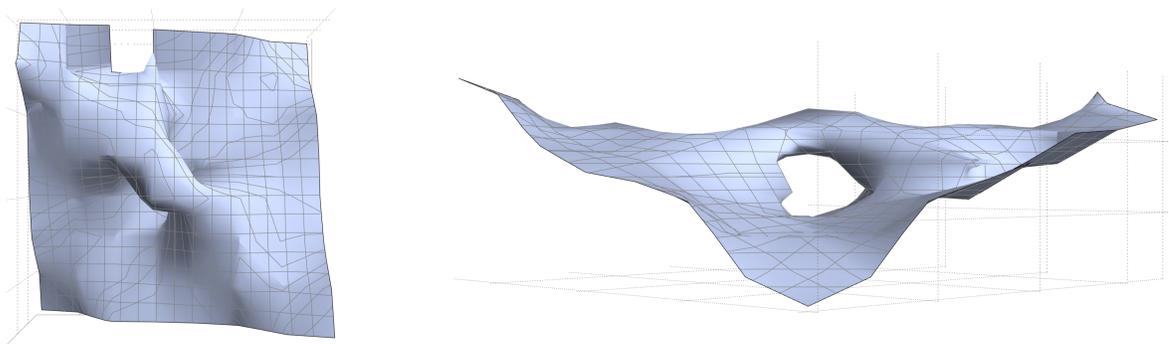


Figure 14: A 3D plot of  $\mathfrak{x}$  seen from two viewing angles.

## 6 Further Prospects

### 6.1 Error Correction and Fault Tolerance

With the prospective advantage of Hensel CPU architecture being exact arithmetic, a design that ensures checks against computational error, and is resilient to computational faults, is of paramount importance in delivering arithmetic performance. Sketched cursorily are some Hensel-architecture-compatible mechanisms for delivering error correction and fault tolerance. These sketches are all inefficient but give indications of the kind of mechanisms that can be developed.

#### 6.1.1 Error-Checking with $(\chi, \pi)$ -Payloads

Error detection for operand encoding can be performed by taking advantage of the  $\chi$ -ID system. Suppose a given encoding block is mis-transmitted (e.g., between a  $C_{PR}$  and an  $\mathcal{A}$ , between the  $\mathcal{A}$  and  $\mathcal{M}_{PR}$ , etc.). Because the process begins by sending the appropriate  $\chi$  from  $\mathcal{M}_{PR}$  to the appropriate  $C_{PR}$ , one could readily implement a check for operand-transmission consistency during

LSU operations by requiring that the loaders involved continue to transmit  $\chi$  in their payloads by requiring that  $\tau$  payloads become  $(\chi, \tau)$ -payloads, in which case FC-1-2025-form errors could be easily detected and located.

#### 6.1.2 $\Delta_2^\mu$ Trajectory-Defect Checks

Because loaders at a given nest level  $\ell$  can perform computations that affect  $\Delta_2^\mu$  by a certain 2-adic distance, an error in computation or transmission at any given level  $\ell$  will be evident if it effectuates a change in  $\mu(q)$  beyond the  $\Delta_2^\mu$  range permissible for that level  $\ell$ . Figure 15 gives an illustration.

#### 6.1.3 LSU Consensus

One could also, albeit at the cost of efficiency, implement a distributed-consensus framework within the processor itself, namely by employing an odd number of redundant clusters whose outputs are sent to  $\mathcal{M}_{PR}$  upon clearing a vote. Note that such a process does not jeopardize parallelization, for the redundant LSUs could compute in parallel with respect to one another.

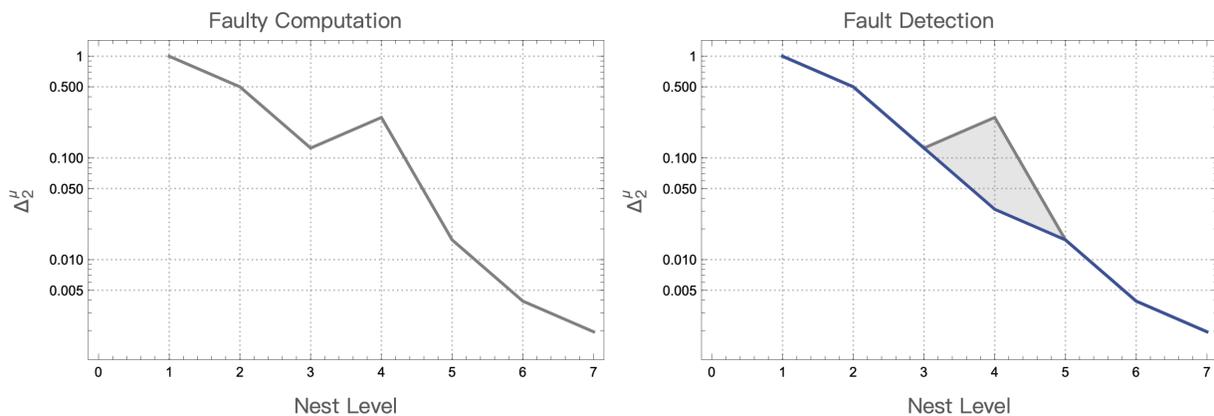


Figure 15: Illustration of  $\Delta_2^\mu$  trajectory-defect detection.

### 6.2 Supercomputing

The register cluster is designed so that it can implement multiple F-2-2025 encoded  $\mathcal{H}$  instruction lists simultaneously. That is to say, the  $\mathcal{H}$  operations are executed in parallel and, what is more, multiple  $\mathcal{H}$  instruction lists can themselves be executed in parallel. Given an instruction superlist

$$\Pi := (\mathcal{H}_i) \tag{25}$$

the upper bound on  $\text{Length}(\Pi)$ , the number of  $\mathcal{H}_i$  that can be executed simultaneously, is comparable to the number of loaders in the cluster (since, in the most efficient case, each loader is performing an operation at

any given time). Because in the case of FC-3-2025 encodings each operand is broken into blocks are loaded to four registers, the number of permissible parallel operations would be

$$\max(\Pi) = \frac{2^{\mathcal{L}-2}!}{4!(2^{\mathcal{L}-2} - 4)!} \tag{26}$$

Thus, a processor cluster with a depth of  $\mathcal{L} - 2 = 15$  would be able to handle as many as  $10^{19}$  operands at once. (Figure 16 shows a cluster of depth slightly lower.) One would then just need enough 2AALUs.  $\mathcal{L} - 2 = 17$  should cross the zetaXOPS performance threshold, as shown in Figure 17.

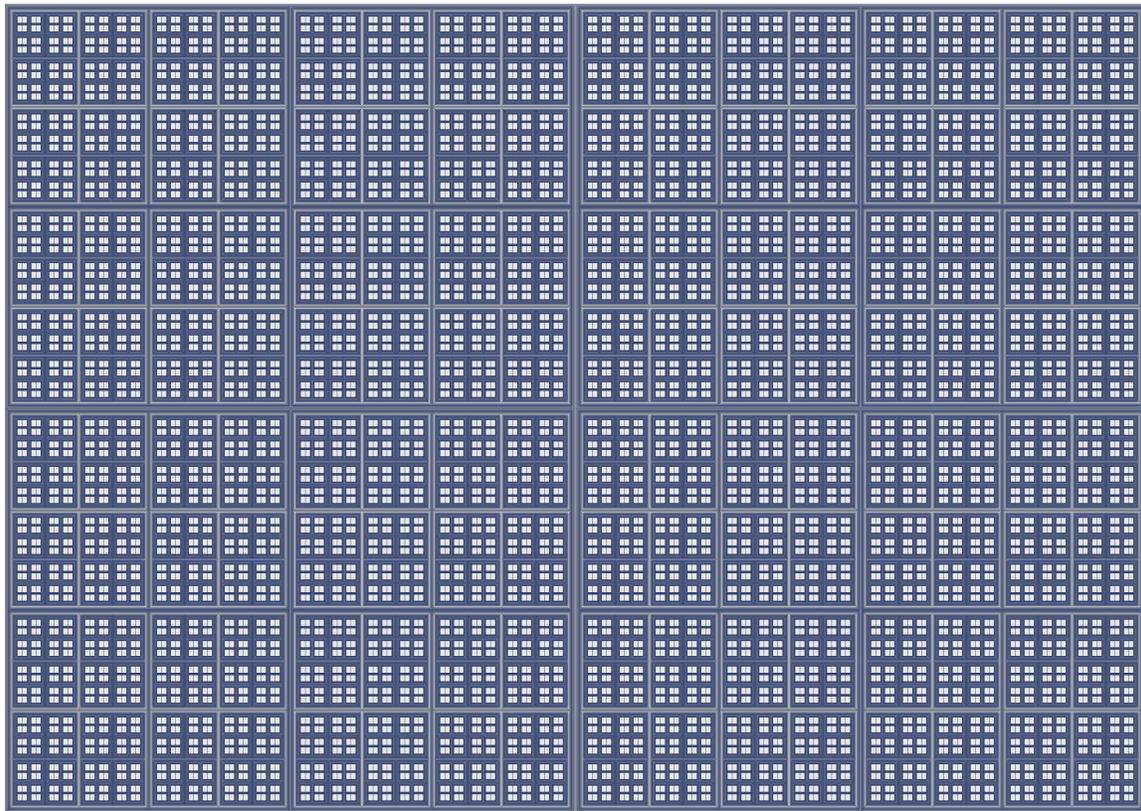


Figure 16: A register cluster of level depth  $\mathcal{L} = 13$ .

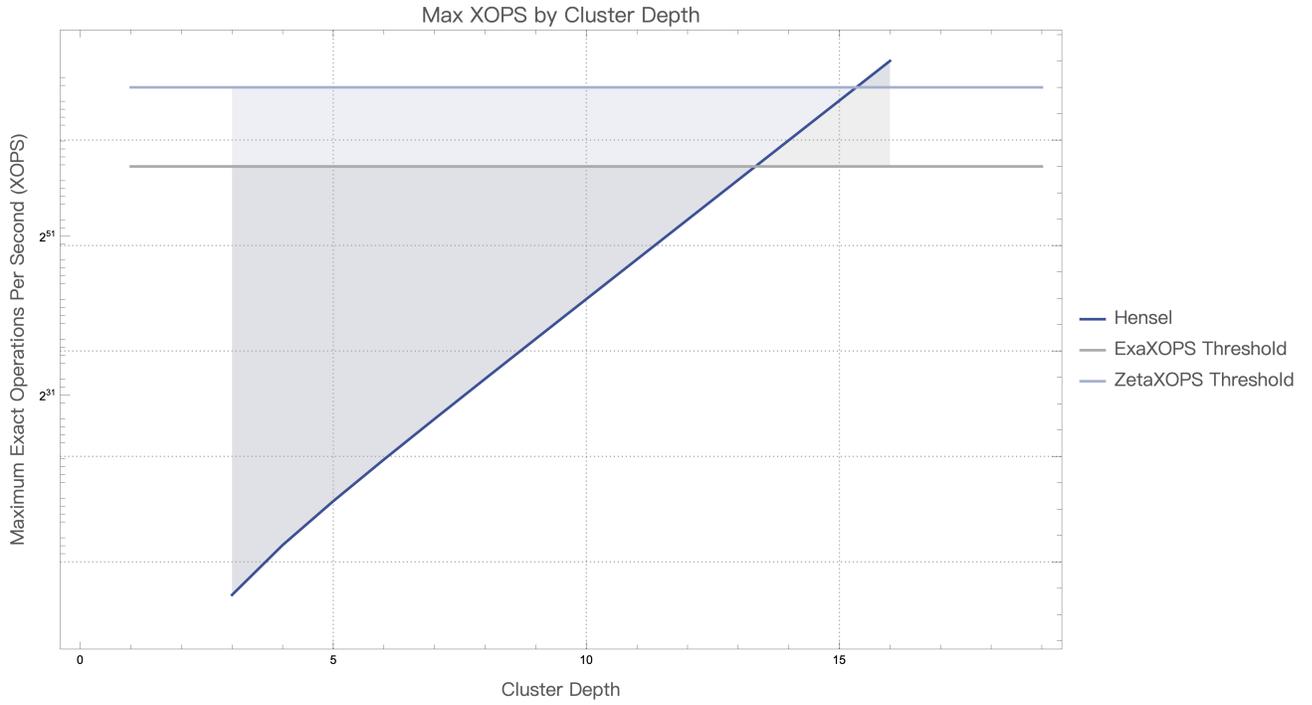
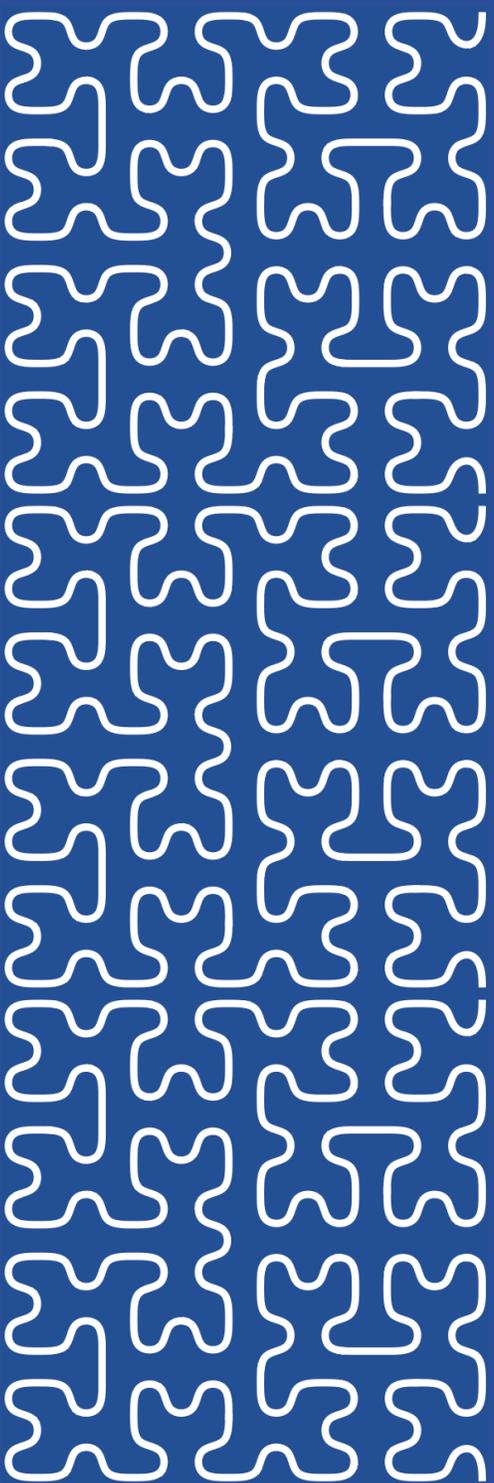


Figure 17: Best-case prediction of the cluster depth required for exaXOPS and zetaXOPS computing performance. (Here, "Depth" refers to the  $\mathcal{L}$  (minus 2 value).)



# SciSci Research

サイサイ・リサーチ