



FERMAT

Storage-Attached Compute Accelerator

“People who are really serious about **software** should make their own **hardware**.”

–Alan Kay

Rationale for new hardware

STOIC built the fastest analytic engine on the market.

Customers love it, but they want it even faster.

And they want to use it on ever larger datasets.

Unfortunately, Moore's law came to an end.

FERMAT was built by the people who built STOIC.

Together, they will bring unmatched levels of performance.

Applications

BANKING

Sales, Risk, Finance, Compliance

INSURANCE

Customer Lifecycle, Asset Liability

ENERGY

Grid Optimization

GENOMICS

Gene Sequence Analysis

SECURITY

Text, Graph, and Geospatial Analysis

Design Principles

Designed for live interactions with **petascale** datasets.

Scale-up storage-attached compute system (vs. scale-out).

Bringing **compute to data** (vs. bringing data to compute).

Powered by next-generation **MPSoC** devices (CPU + FPGA).

Programmed with standard language and proprietary **DSL**.

Packaged as standard **PCIe** card (fits within any system).

Suitable for conventional data centers (forced **air cooling**).

Anatomy of an MPSoC

Xilinx Zynq UltraScale+ ZU9EG

Quad-core 64-bit ARM Cortex-A53 MPCore up to 1.5GHz

Dual-core 32-bit ARM Cortex-R5 MPCore up to 600MHz

Dual-core ARM Mali-400 MP2 GPU up to 667MHz

2,520 × 48-bit DSP Slices up to 500MHz

274,000 Look-Up Tables up to 500MHz

32.1Mb Block RAM


FERMAT x8

Hyperconverged System on a Board (HcSoB)

8 Modules with Map MPSoC



4 × 64-bit ARM Cores
2,520 × DSP Slices
16GB DDR4 SDRAM
8TB Flash Memory



32 × 64-bit ARM Cores
16 × 32-bit ARM RT Cores
20,928 × DSP Slices
2,729,600 Look-Up Tables
384GB DDR4 SDRAM
64TB Flash Memory
192GB/s RAM Bandwidth
83GB/s Flash Bandwidth
32GB/s PCIe Bandwidth
12 × 16GB/s Internal Links
21TeraFLOPS
350W

PCIe Gen3 ×16 Board with Reduce MPSoC

Design subject to changes

Proliant XL270



- 2U Server
- 2 × Xeon CPU
- 2 × PCIe Switch
- 1TB DDR4 SDRAM
- 2 × 2 × 100GbE Ports
- 8 × FERMAT x8 Boards

Apollo 6500



2 × HPE Proliant XL270
16 × FERMAT ×8 Boards
88 × 64-bit Xeon Cores
512 × 64-bit ARM Cores
334,848 × 48-bit DSPs
8TB DDR4 SDRAM
1PB Flash Memory
3.1TB/s SDRAM BW
1.3TB/s Flash BW

Suggested Rack

10 × HPE Apollo 6500

20 × HPE Proliant XL270

160 × FERMAT ×8 Cards

82TB DDR4 SDRAM

10PB Flash Memory

80 × 100GbE Ports

614TB Flash SSD

880 × 64-bit Intel Xeon Cores

5,120 × 64-bit ARM Cores


3,348,480 × 48-bit DSP Slices

31TB/s SDRAM Bandwidth

13TB/s Flash Bandwidth

8Tbps Network Bandwidth

3.4PetaFLOPS



Performance per Host

Device	Performance	Capacity	Bandwidth	Throughput	Celerity	Efficiency
XEON CPU Cores	1.5TeraFLOPS	1TB	50GB/s	32mB/FLOP	0.05TB ² /s	5GFLOPS/W
NVIDIA GPU Cores	84TeraFLOPS	128GB	6TB/s	68mB/FLOP	0.18TB ² /s	141GFLOPS/W
Map ARM Cores	1.5TeraFLOPS	1TB	1TB/s	667mB/FLOP	1TB ² /s	10GFLOPS/W
Map DSP Slices	161TeraFLOPS	512TB	666GB/s	4mB/FLOP	341TB ² /s	195GFLOPS/W

System Architecture

Running Linux instance on every MPSoC (8 Map Modules).

Using FERMAT SDK for data distribution and MapReduce.

From the host's OS, presents itself as 9 storage drives.

From the network, presents itself as 8 servers.

Board-level GlusterFS for module-level fault tolerance.

Cluster-level GlusterFS for board-level fault tolerance.

Polymorphic Data Indexing

Tabular index for **pivots**

Columnar index for **quantiles**

Relational index for **joins**

Hexastore index for **graphs**

Temporal index for **time series**

Inverted index for **full text search**

Finite-state transducer for **fuzzy text search**

Geospatial index for **locations, tracks, and geofences**

Multilevel MapReduce

Each board includes 8 Map modules with Flash Memory.

Data sharded across Map modules with polymorphic indexes.

Shard-level map phase handled by Map MPSoC devices.

Board-level reduce phase handled by Reduce MPSoC device.

Host-level reduce phase handled by local host.

Cluster-level reduce phase handled by any host in cluster.

Polysilicon Compiler with Optimizer

Device	Count	Precision	Media	Capacity	Bandwidth	Uplink
XEON CPU Cores	22	64-bit	SDRAM	512GB	50GB/s	100GbE
ARM Cores	4	64-bit	SDRAM	16GB	16GB/s	PCIe
DSP Slices	2,520	48-bit	Flash	8TB	10.4GB/s	Flip-Flops
Look-Up Tables	274,000	Any	Block RAM	31Mb	20PB/s	Flip-Flops

FERMAT SDK

FERMAT Polysilicon Compiler for CPU, GPU, DSP, and LUT.

FERMAT Data Indexing Engines.

FERMAT Multi-level MapReduce Engine.

Just-in-Time FERMAT to C++ Translation.

Just-in-Time C++ to Machine Code Compilation.

Incremental Machine Code Compiler.

FERMAT Language

Standard JavaScript 6 extended with operator overloading

Evaluated within browser or server

Compiled into C++ to run on CPU, GPU, or DSP (FPGA)

Directly linked to distributed in-memory data structures

Extended with the 450+ formula functions of Excel[®]

250+ functions for ETL, linear algebra, and time series

350+ functions for statistics and machine learning

FERMAT Optimizer

Compile-time and run-time optimizer.

Configured with actual data sharding and indexing plans.

Optimized with real-time resource utilization metrics.

Selecting compute devices for Map and Reduce phases.

Designing topology of multilevel Reduce phase.

Recommending data sharding and indexing optimizations.



info@fermat.co

Palo Alto · Paris · Basel