Corporate & tech

April. 2023

Scale Big Data & AI

Boost Big Data & AI 20x, being 10x more energy & cost efficient, with **Processing In Memory (PIM)**

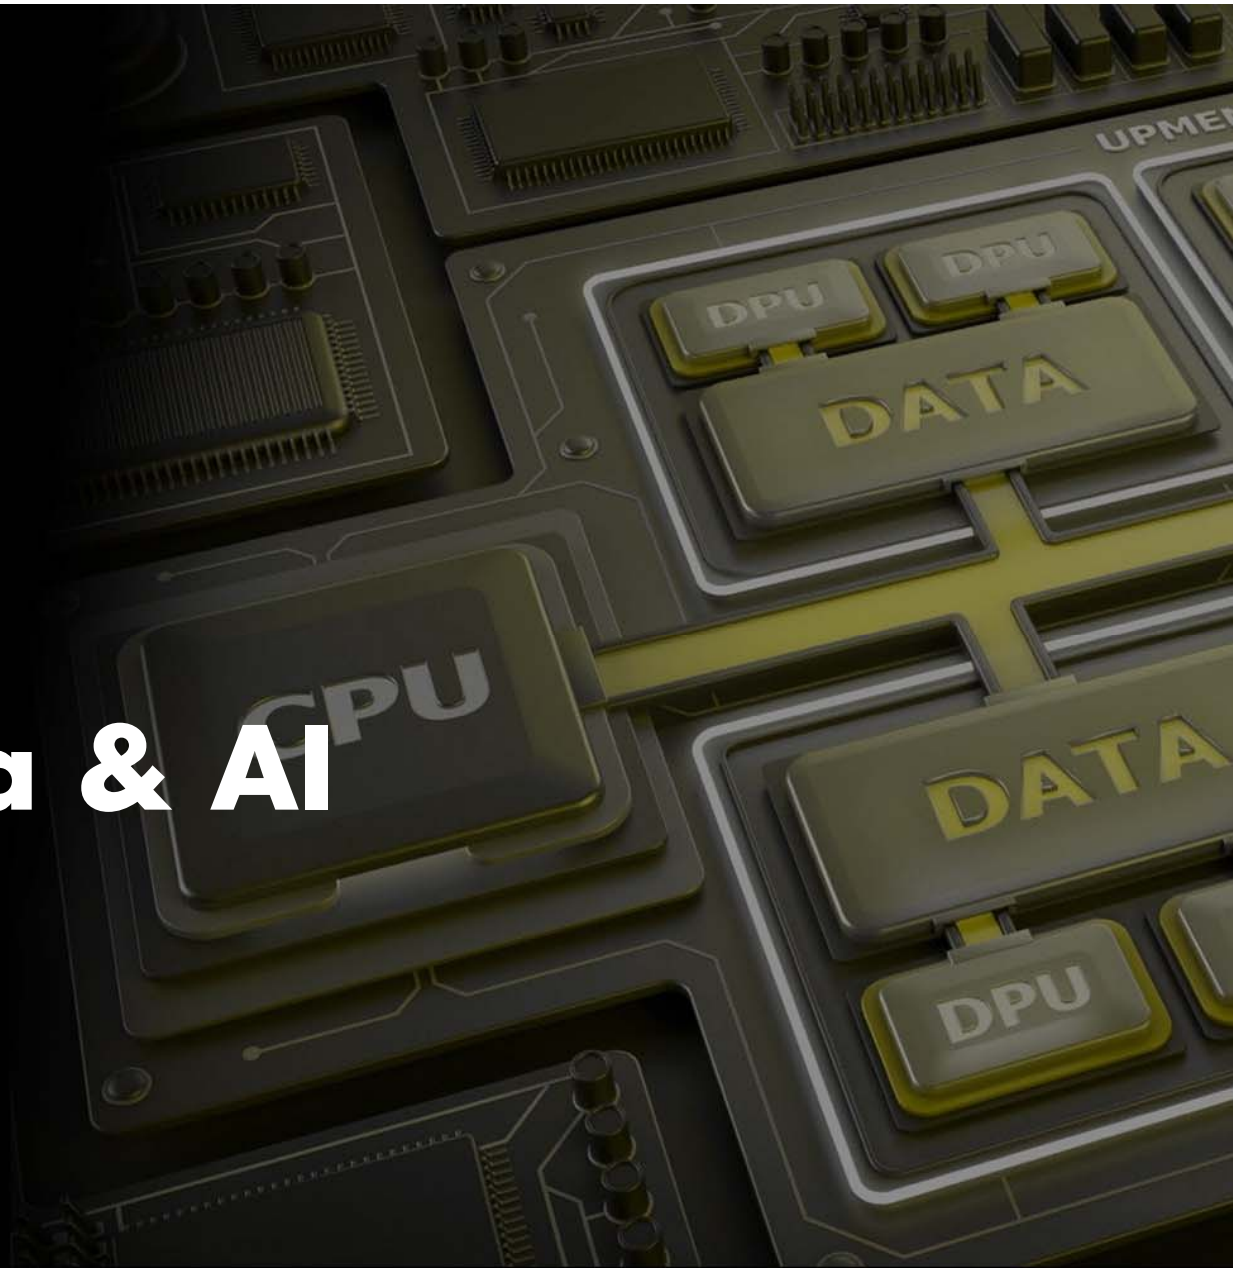

Corporate introduction

-

UPMEM in a nutshell

Leader & Pioneer in Processing In-Memory (PIM) on DRAM

INNOVATION
- The award-winning **leader in revolutionary data-centric computing**
- Pioneering Processing-In-Memory (PIM) **solving CPU-Memory bottleneck**

IMPACT
- **20x faster** processing
- **10x less energy** consumption, 10x lower TCO
- Marginal additional hardware & implementation cost
- 20% reduction in datacenter energy & hardware footprint

BUSINESS MODEL
- **Fabless semiconductor**, with strong foundry partnership
- Server vendors will assemble and ensure distribution towards data centers

APPLICATIONS
- **Big Data** (analytics, genomics) & **AI**: ~50% of data center applications
- PIM to later target edge

CUSTOMERS
- **Large Cloud Service Providers** will buy and deploy UPMEM PIM accelerators
- Large data heavy application owners start creating market pull

illumina · SAP · intel · ANSYS OPTISWORKS 安诺优达 · HUAWEI · THALES

UPMEM PIM-DRAM wafer

PIM Chip

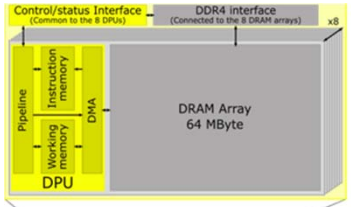

PIM Module

PIM Server

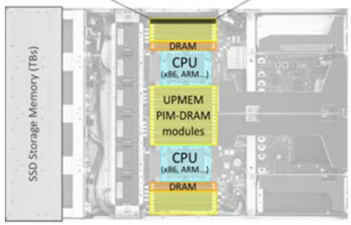

Total Funding: $17M

Founded: 2015

2022 POC revenues
2025E Revenue: $30M+

Total Patents: 11

Employee Count: ~23

Headquarters:
Grenoble, France

Supported by

Western Digital Capital · PSMC · C4Ventures · SUPER NOVA INVEST · bpifrance · partech · Co-funded by the European Union

PROBLEM

CPU computing nodes cannot cope with data-intensive apps

Global datacenter energy consumption booming

- 2% of global electricity, growing fast
  - 40% of which is compute
- Datacenter capacity is a main bottleneck for cloud players

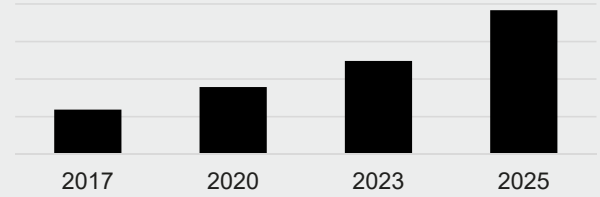

Source : [Lean ICT Materials] Forecast Model. The Shift Project Oct 2018 (data from Andrae & Edler, 2015)]

Data movement is the dominant compute system energy cost

- ~80% of energy consumption in the computing node occurs moving data between the main memory & processor (DRAM and CPU)
- 63% of server system energy consumption

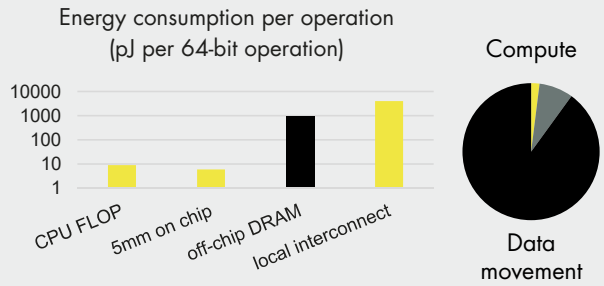

Source: Lawrence Berkeley National Laboratory; SK Hynix 2021

Current compute architectures have inherent data bottlenecks

- Data-intensive applications, becoming dominant, are memory bound
- Existing solutions don't meet market needs
  - Big Data Analytics, genomics…
- Memory access is the main bottleneck, not compute

Data-related architectural bottlenecks are the main obstacle to improving computing performance

up mem

SOLUTION

The Next Generation data center compute architecture

PIM inhibits movement of dana
by performing calculations in memory

PIM bypasses structural bottlenecks
inherent within existing server architecture

UPMEM PIM unique patented technology

- Leveraging existing industry processes, protocols, server architecture and memory slots, allowing smooth integration while maintaining familiar programmability

Standard compute node

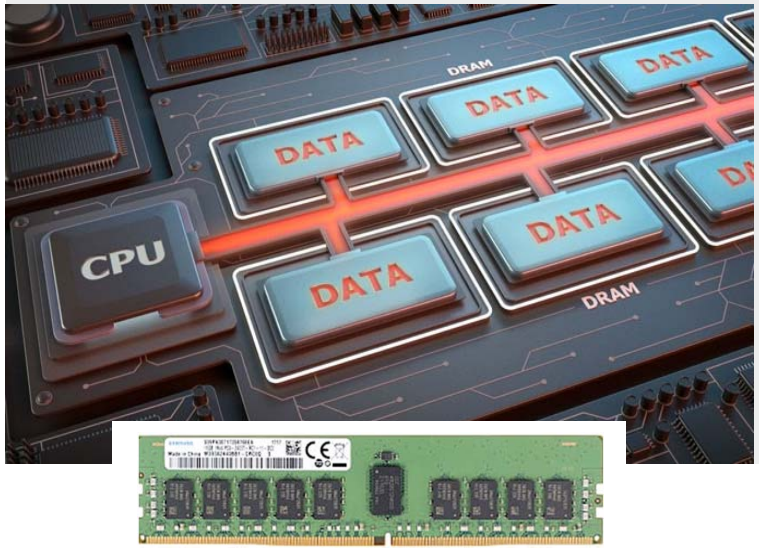

PIM system

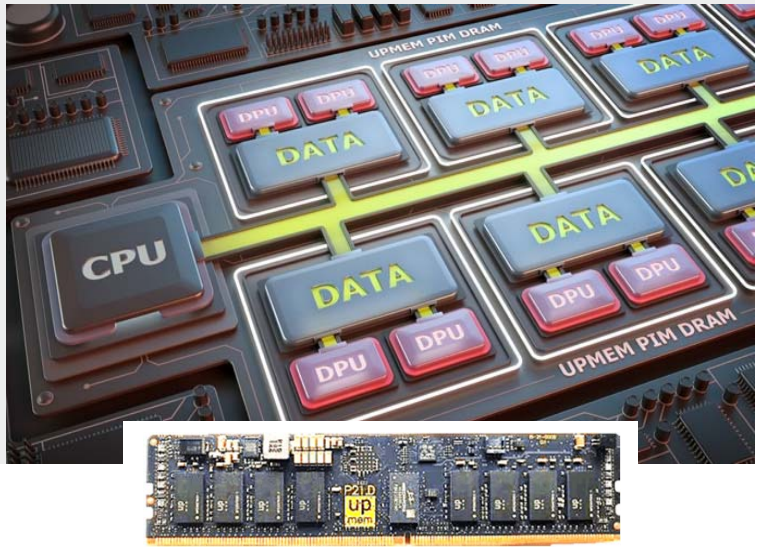

HIGHLIGHTS

- Up to 40 times faster
- Up to 10 times less energy at system level for the same compute power
- Up to 10 times lower TCO at system level
- Massive scalability

SOLUTION

UPMEM ultra-efficient scalable PIM-DRAM accelerator solution

- Hardware acceleration of data analytics
- Massive local and global data bandwidth
- Predictable performance
- Offloads data analytics while freeing CPUs bottlenecks and charge
- Scalable system integration
- Hardware
- Software

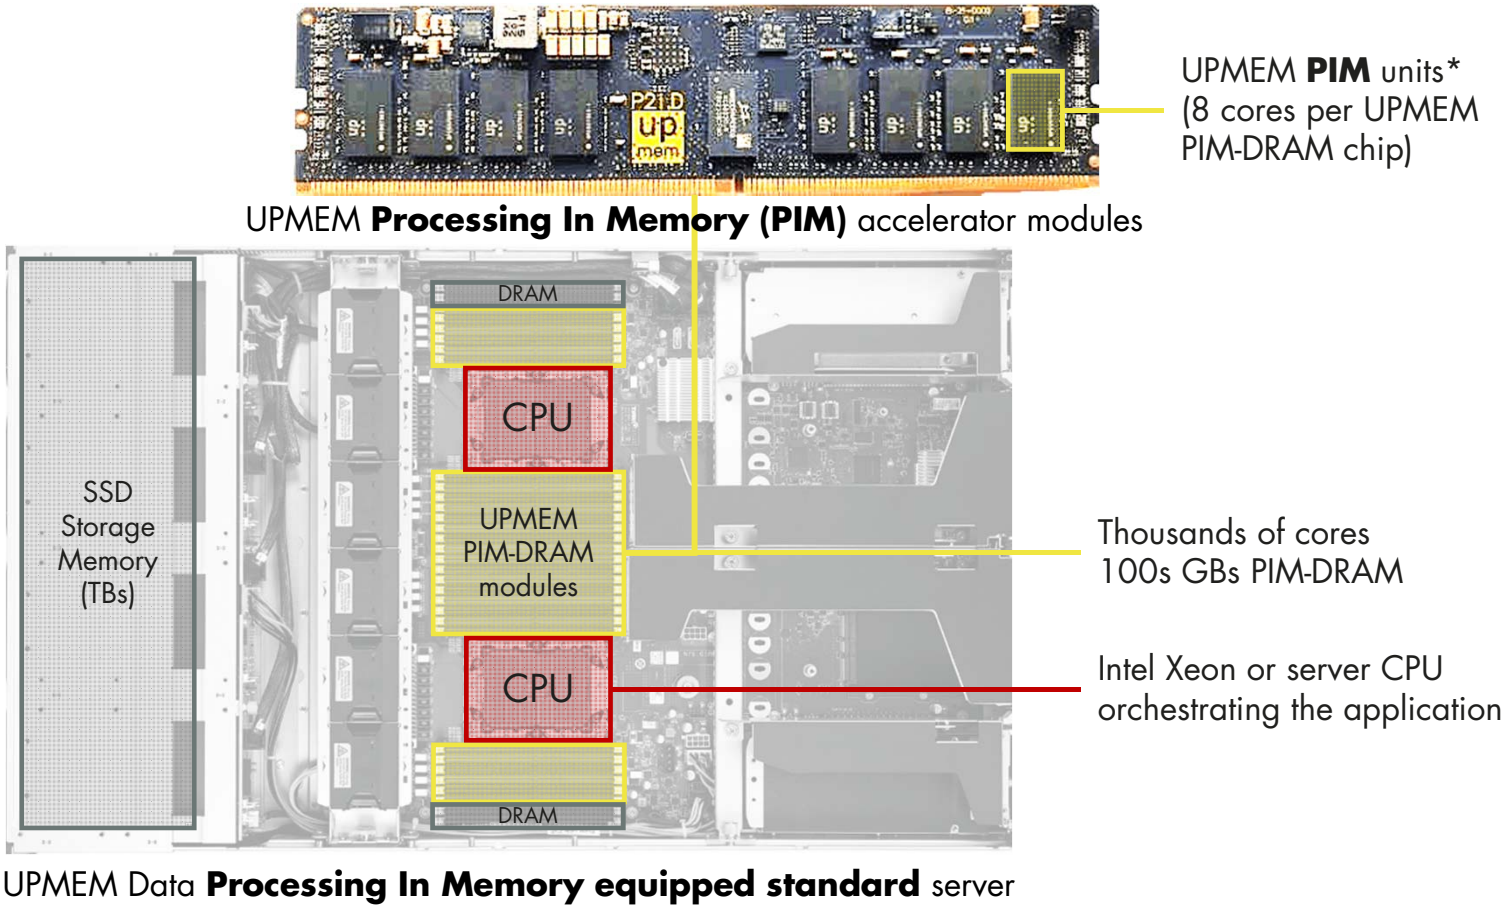

* A **PIM** unit is an app core (DPU) associated with its 64MB DRAM

SOLUTION

Easily deployable with widely accessible programming effort on applications

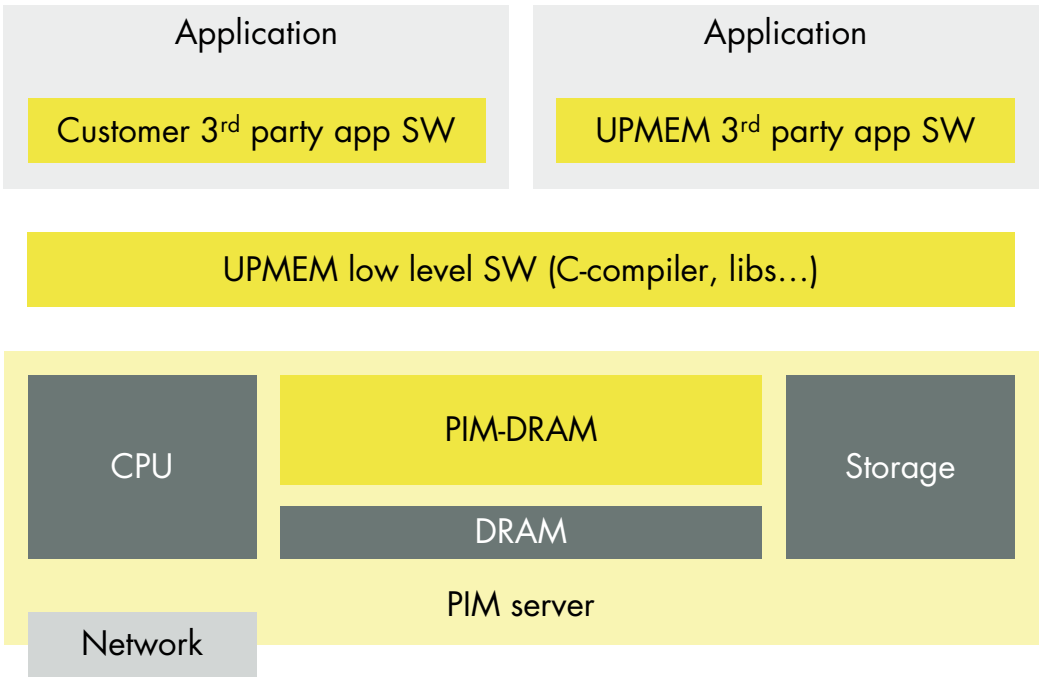

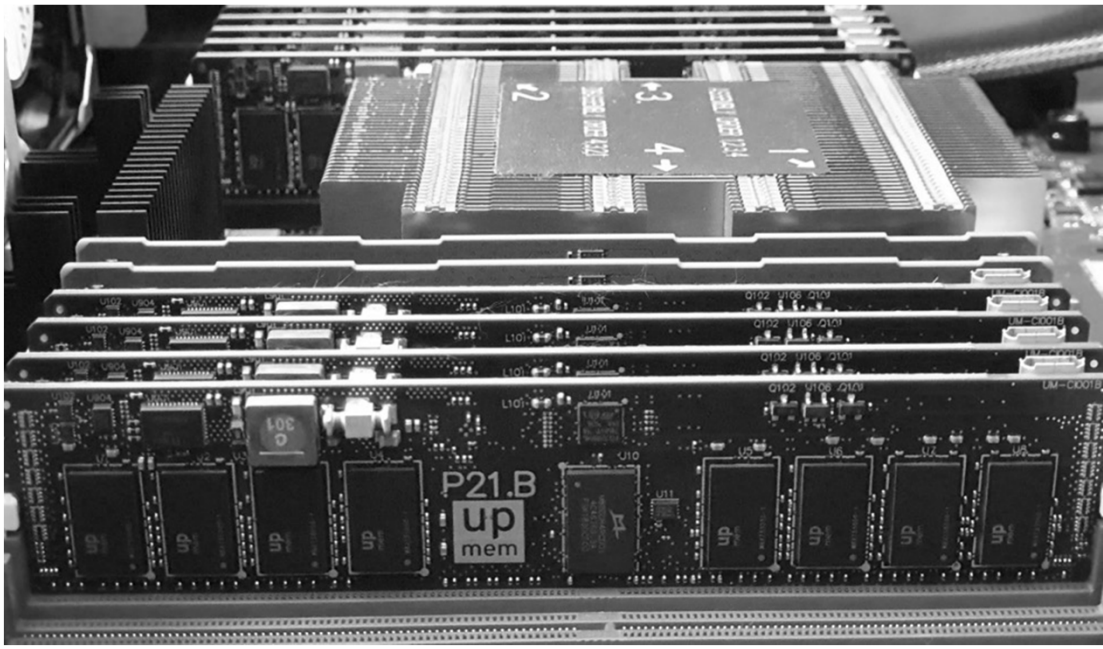

up
mem

PRODUCT MARKET FIT

UPMEM PIM: game changer for analytics & genomics

Targeting data-intensive compute-intensive apps

Data Analytics

Customer

Leading Telecom Equipment Provider

Use Case

- Index Search of string text in big data sets

Customer Relevance

- Execution time is a top pain point with TBs of data to move and compute
- No appropriate server accelerator (GPU, FPGA)

Results

- **26x Increase** In Throughput
- **25x Decrease In Total Cost of Ownership**
- **15x Energy Efficient per Request**
- Accessible Programing
- Scalable Deployment

Speed-up vs Lucene*: x26 vs best CPU to x50 vs same CPU

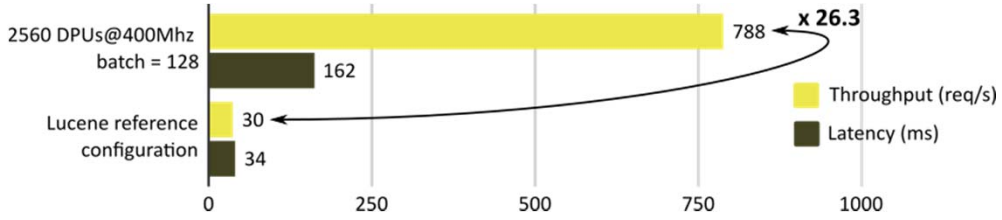

Energy x15 better and TCO x25 better vs same CPU

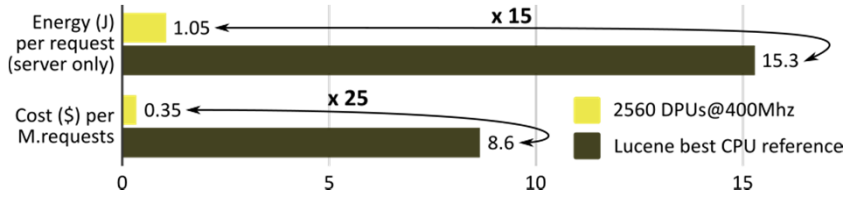

* On Wikipedia index

TECHNOLOGY

UPMEM 1st PIM-DRAM based MVP available & shipping

Tech breakthroughs

1st PIM-DRAM

Design & implement processor (DPU) into DRAM process

- with proprietary ISA

Silicon proven design

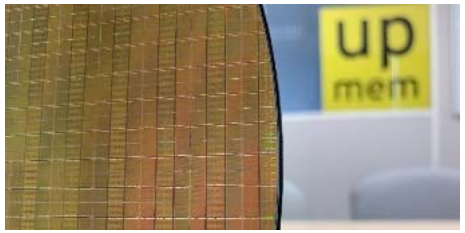

UPMEM PIM wafer

Fabbed with DRAM process

Mass-production grade in DRAM foundry

DDR4 compatible

DDR4 memory protocol compliant
Hacking DDR4 for control

JEDEC SSO **PIM standardization**

UPMEM PIM chip

Standard DIMM for unchanged server hardware architecture

- developed internally

128 processors per DIMM, 8GB

System proven

UPMEM module with UPMEM PIM chips

System Integration

Allow PIM DRAM memory sharing (CPU & PIM-DPU) with appropriate server BIOS/Firmware settings

Thousands of PIM units, i.e. PIM DPUs per PIM server

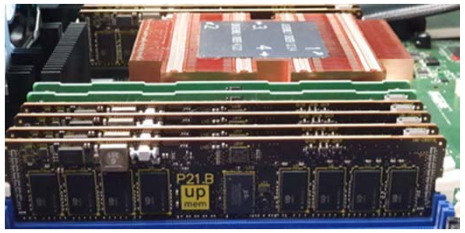

UPMEM PIM qualified server

Running in operational conditions on x86 servers 24/7

Programmability

Allow programs to share tasks and data between CPU and PIM DPUs

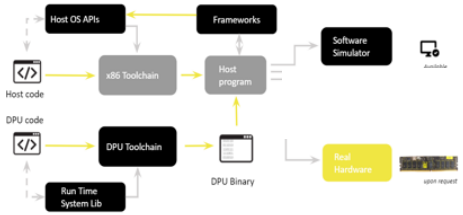

UPMEM proprietary SDK

Allow familiar C, Rust programming of DPUs

Uses standard LLVM compilation toolchain

Benchmarks & use cases

Hundreds of PIM developers

up
mem

SOLUTION

UPMEM ultra-efficient scalable PIM-DRAM accelerator solution

PIM Chip

Ultra Efficient PIM-DRAM Processor

- Implemented on unchanged DRAM process
- High & efficient data-bandwidth

Frequency	Up to 500 MHz
Register width	32-bit (64-bit support)
Bandwidth to DRAM	1 GB/s
DPUs per PIM Chip	8 DPU/PIM unit per chip
PIM Chip DRAM	8x64 MB = 4Gb
PIM Unit WRAM/IRAM	64 kB/24kB

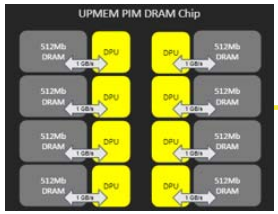

PIM Module

DDR4 Compliant PIM Module

- Readily inserted in the memory slots
- DDR4 speed

JEDEC Compliance	Yes
Memory	DDR4-2400 8 GB
PIM Chips per Module	16
DPUs per Module	128
Operating tempelature	0 to 95C

UPMEM **PIM** accelerator module

PIM Server

DDR Server Integration

- 160GB PIM-DRAM & 2560 DPUs @ 400MHz
- Bios Patching only

Chassis	E.g. INTEL Server System
CPUs	E.g. INTEL Xeon
Memory	4x32/64 GB DDR4 Dual Rank DRAM per socket
PIM Modules per Server	20
Memory Bandwidth	2.56 TB/s

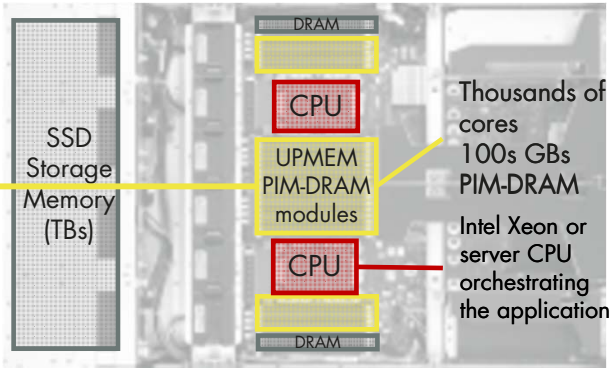

up
mem

BUSINESS MODEL

Selling & distributing UPMEM PIM

Cloud Services Provider

AWS, MS Azure, Google Cloud…OVH…

Own use or cloud services (search, analytics, DB…)

Fabless semiconductor model, selling accelerator solution

Main customers are Cloud Services Providers

Initially sell to large application owners

Sales

Cloud service Provider → Solution/ App Provider → End user

(Analytics, genomics)

UPMEM PIM Accelerator

HW & Software Dev Kit

Strategic partners

Delivery

Key components design → Chip/IC Manufacturing → Server vendors (OEM/ODM)

IP

Technology well protected & recognized by IT & SC community

Robust Patent Portfolio

Granted & Published

UPM1	DRAM circuit provided with a built-In processor
UPM2	Computer device provided with Processing-In-Memory and narrow access ports
UPM3	Row hammer correction logic for DRAM with integrated processor
UPM4	Combined instruction for addition and checking of terminals
UPM5	Methods and devise for bypassing the internal cache of an advanced DRAM controller
UPM6	Die Stacking-Heat Sink for protecting a DRAM device from the row hammer effect
UPM7	Silver Bullet method and circuit for protecting a DRAM device from the row hammer effect

Unpublished

UPM 8 to 11	

Collaboration with >20 Leading IT Labs

IT Labs are exploring and building use cases with benchmarks in EU, US, Israel, and Asia

Award Winning & Well Referenced

- Selected as Top 100 emerging startups to follow in 2020
- Products are highlighted by mainstream chip community with increasing frequency
- Leading standardization for DDR-PIM with JEDEC

Impact

Reduce by 20% the datacenter energy & hardware footprint

Scale Big data & AI, enabling new critical digital services

Genomics

Analytics, search, DB

AI

Security

EIC Accelerator contribution

-

EIC contribution

Support UPMEM tech scale up with EIC Accelerator program

Scaling from POCs with MVP to deployable product with Cloud Service Providers (CSP)

Showcase PIM Use
Accelerating key workloads of high potential apps with pilot customers & IT Labs

Grow Market Use
Co-design & nurture deployable apps on scalable platforms

Scale & Deploy through CSPs
Cloud Service Providers (CSPs) deploy with server & software ecosystem

2022

Large App PoC Customers

- **Tier 1 Paying Customers with successful POCs** in US, China, Korea, Europe
- Big Data Analytics, search, Genomics…

With MVP, while designing 2nd generation product

Software Ecosystem Ignition

- 25 Labs globally with 3 EU-backed projects
- Hundreds of PIM developers

Prepare Early Strategic Partners

- Scalable DRAM foundry & ecosystem
- 3 top server partners

2023

Large App Deployed Customers

- **Codesign deployable PIM apps**
- Prepare large sales with SW effort

While completing & fabbing 2nd generation product

- Performance, security, scalability

Mature Software & Add Use Cases

- Apps & Libraries
- More IT labs & customer use cases

Expand Strategic Partner Network

- Improve tech & manufacturing partners
- With dedicated 3rd party PIM qualifies servers

2024/2025

Sell, leverage 1st mover advantage

- **Large deployment with CSPs at scale**
- Nurture large application owners' adoption

With mature and qualified solution

While continuing R&D

Full Scale Software Support

- Full-suite of software tools
- Comprehensive application libraries

Leverage Strategic Partners

- Upgrade technology & systems
- Add software partners

2.5M€ EIC grant

6M€ EIC Fund equity

up mem

EIC Accelerator program

Work packages of EIC Accelerator program EnergiaPIM

Tech

- Chip design and architecture
- Chip manufacturing
- Board and system design, manufacturing
- Productization and qualification
- Software developments

- Use cases
  - Genomics
  - Search
  - Analytics
  - AI
    - Analytics
    - LLMs

Pre-commercial activities

- Dissemination
- Engaging developer's community
  - Open-source
  - Academics
- Tech marketing

EIC Accelerator program

Takeaways from our journey with EIC

PROs

- Provide deep tech financial support hard to find in Europe to scale
- Dedicated and supporting teams
- Networking
- Visibility

CONs

- Bureaucratic long process and schedules
  - Hopefully solved

Thank you

Gilles HAMOU, CEO & co-founder
ghamou@upmem.com