

The logo for ARRAY, featuring the word "ARRAY" in a bold, white, sans-serif font centered within a dark blue rectangular background.

ARRAY

EXECUTIVE DETAIL

FOR THE IMPLEMENTATION OF:

*A COMPUTE-IN-MEMORY ARITHMETIC CORE WITH
+85% EFFICIENCY AND THROUGHPUT GAINS
USING STANDARD CMOS TECHNOLOGY*

www.fastarithmeticunits-dataroom.com

[ARRAY ARCHITECTURE, CORP.]

Computing Infrastructure, from hyperscale data centers down to edge devices, is one of the global energy market drivers, and in developed countries it has the fastest growing energy demand of any industry. **Most of a processor's time delay and power overhead is spent moving data between processing cores and memory arrays.** Furthermore, every arithmetic operation in every processor depends on the decades old carry-over paradigm of addition. and **state-of-the-art Parallel Adders are not fast enough for modern applications.** They are complex, unscalable and power-hungry circuits.

This document constitutes an executive brief of our proprietary Simple and Linear Fast Adder (SLFA) which **solves the energy-wall problem** at the mathematical level, not the physical (material) level. **By colocalizing logic and memory in a single grid, we eliminate the bottleneck responsible for 60-90% of time and energy consumption.** The SLFA enables area-optimized Fast Arithmetic Units capable of executing matrix multiplication (the mathematical operation behind **AI training, graphics processing, cryptography, and Digital Signal Processing**), and Hash compression functions (**algorithm for mining Bitcoin**) with minimal silicon and energy foot-print. A patent application is in the national phase in key jurisdictions: US, China, Japan, Korea, India, Singapore, UK, and Canada. A third-party valuation from InTraCoM Group (Germany-based global patent valuation experts) **places current value at +\$8 Million USD.** Additionally, we have made available a comprehensive technology and market report from Anuation Research & Consulting LLP.

Unlike other CIM proposals that require exotic materials and billions in R&D, our architecture works with today's technology. **The innovation is in the mathematical algorithm which redefines addition,** and enables operations to be executed in a regular array of standard CMOS memory cells and logic gates. Any leading foundry can manufacture without costly and risky R&D, giving us **a direct path to high-volume, low-cost, next-generation processors.** While highly applicable to Bitcoin Mining and AI data center processors, **the FAU is broadly adaptable across multiple compute-intensive applications.**

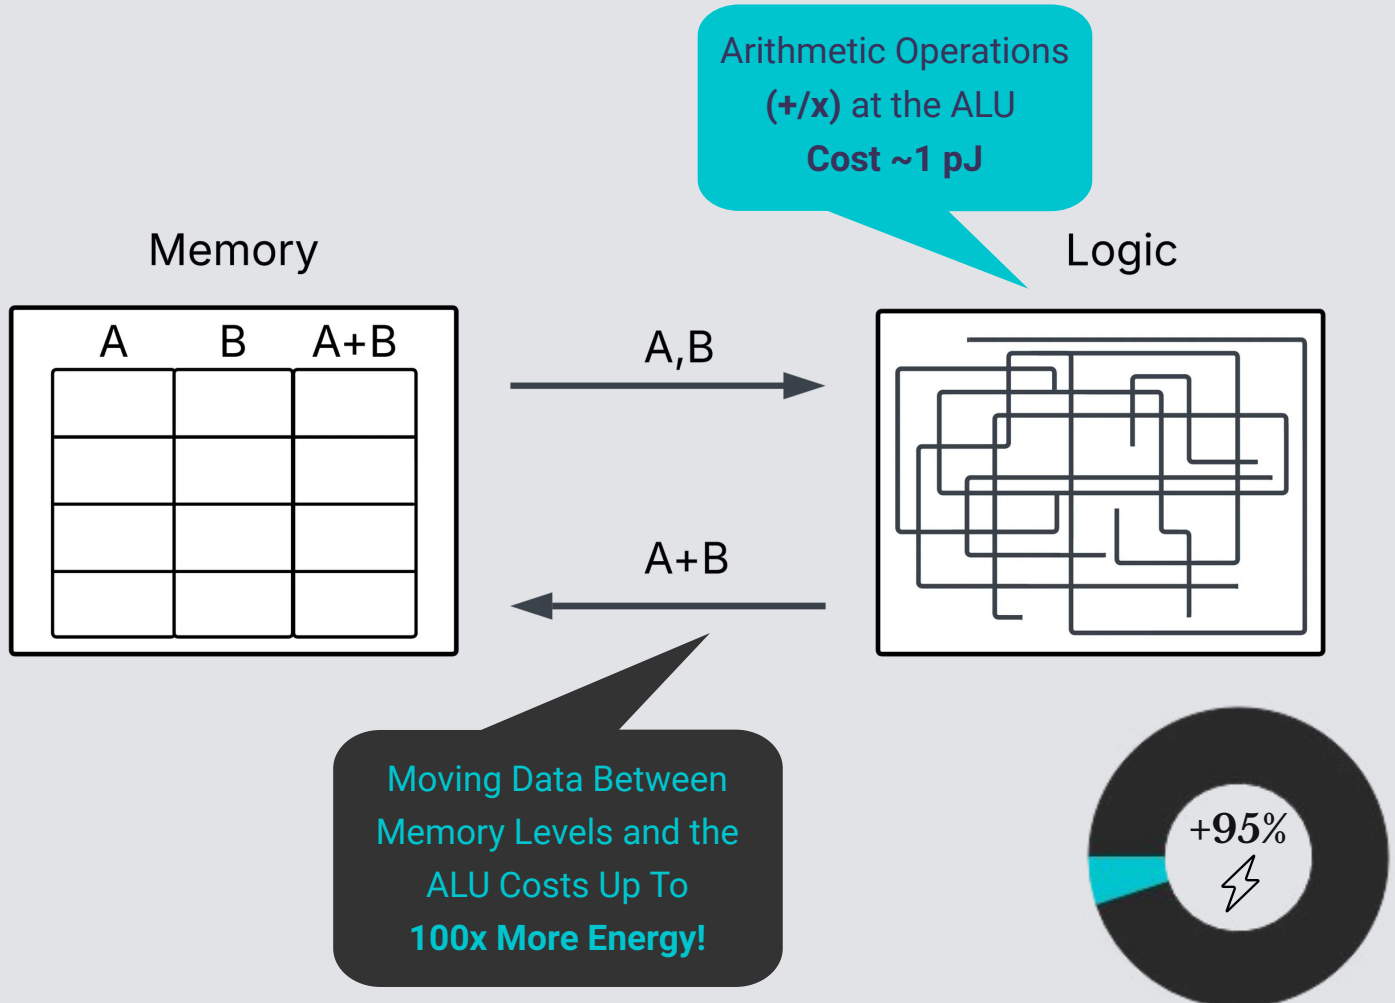
Index

Problem.....	4
Solution.....	7
Core IP: Simple and Linear Fast Adder.....	8
Product and Technology.....	9
Innovation and Differentiator.....	10
Comparison Table.....	11
Market Impact.....	12
Potential Market Sizes (TAM, SAM, SOM).....	13
Business Model and Competition (Potential Partners).....	15
Annual Projections and Base Data.....	16
Traction.....	17
What's Next? (Our Ask).....	18
Technology Readiness Level (TRL-3).....	19
Technical Road Map.....	20
Use of Funds.....	22
Table Summary for Investors and Partners.....	23
Let's Partner.....	24

PROBLEM

High-speed & reliable arithmetic cores are the main objective for high-performance processors such as those specialized for graphics processing, and AI. **The single most important subunit determining performance in any processor is the ALU (Arithmetic Logic Unit), responsible for ~40% of all operations** executed in an average CPU and can reach +80% of operations executed in GPUs or other ASICs. ALUs house adders and other arithmetic cores. Adders are a type of digital circuits necessary to perform more complicated operations such as matrix operations.


Current **computer architectures separate the Memory (which acts like a warehouse) and the ALU (which acts like a production line) into two different subunits.** Therefore, information has to be relayed back and forth between these two subunits. **The bandwidth between the Memory and the ALU is like a road connecting a Warehouse and a Production Line.** If we increase the capacity of the warehouse and the production line too much, then the road is going to suffer traffic jams and schedules will collapse.




PROBLEM

In a conventional Von Neumann Architecture, data must travel significant distances because the processing and storage is done in separate units. The Von Neumann bottleneck makes up **60-90% of the time latency and energy consumption** in modern processors due to four main reasons:

Peripheral distance travelled between units is huge compared to the distances travelled inside the subunits.



“Bandwidth problem” when trying to increase bit capacity.



Energy costs are disproportionate.

~1 pJ to Add/Multiply
vs.
~100 pJ to Move Data



A single mathematical operation requires data to move **back and forth several times** between memory and compute units.



PROBLEM

- **One category of solutions is to design algorithm-specific accelerators through an extreme ASIC approach** (custom built designs that execute a single workload, very efficiently). These niche designs are hardwired to accelerate very specific processes and any modification to the executed algorithm requires restructuring the entire design. The result is an expensive architecture that you can only sell a few units of, before the algorithms or models are modified/updated. Example of this include **Bitcoin Mining ASICs, Etched's AI accelerators, Neuroblade's SQL accelerators, etc.** These solutions work on macro-level functionality.

- **Another class of solutions physically relocates memory arrays and processing cores closer to each other (Near-Memory Computing)** by chopping them up into smaller sections and placing these closer to each other. This involves complex designs, high control overhead and does not solve the real problem. Non-traditional Packaging, and thermal issues (due to huge wafer-size dies) are also a downside. Companies like **NeuroBlade, Etched and Cerebras** have solved the off-chip memory-wall, eliminating up to 95% of energy consumption at the system and macro system-level. However, **the remaining 5% we target still represents an enormous amount of resources.**

- Compute-In-Memory architectures that perform **analog computation inside a memory cell involve transistor-level physics.** This is the most expensive alternative because of the R&D costs for developing experimental transistor types, based on analog effects. Trade-offs include error prone results, short life spans, and manufacturing difficulty. Examples include **FeFET, MRAM, ReRAM, CRAM, and IBMs Northpole chip.** These memory types eliminate the Von Neumann Bottleneck offering transformative improvements (**up to 90% Time & Energy efficiency**), but are difficult and expensive to manufacture.

R&D Budgets into CIM (Last Decade)	
Based on Public Information	
IBM	\$400 Million USD
intel	\$1 Billion USD
SAMSUNG	\$500 Million USD

SOLUTION

While competitors chase incremental optimizations, our SLFA-based grid bypasses the bottleneck altogether—delivering Fast Arithmetic Units for ultra-efficient computation across many types of processing units. **Our architecture is not a complicated niche design, nor does it require R&D investment into new transistors.** The Simple Linear Fast Adder embeds computational logic with memory elements at the bit level.

- **Low-Power, Low-Cost Architecture**

- **Dense Packing** (More cores per wafer!)
- **Regular Grid Design** (Better thermal profile and distributed heat).
- **Scalable Bit Capacity**

- **Wide Range of Applications**

- Is a universal building block for high-performance arithmetic cores, benefitting most massively parallel, arithmetic-heavy, and power-sensitive processors.
- Generalizes to a **Manhattan-style grid for multi-input addition** (multiplication) and **area-efficient matrix multiplication.**

- **In-Situ Processing (Super-Fast and Energy-Efficient)**

- **Eliminates the Von Neumann bottleneck** by colocalizing logic and memory at the bit level. The distance between memory and its associated logic is reduced to the distance between a Flip-Flop and its neighbor gate—essentially zero.
- Projects 30-50% improvements in latency and power efficiency.
- Higher possible clock speeds (shorter critical paths and gate depth).

- **Based on Standard CMOS Transistor Technology** (Faster time-to-market).

- No exotic materials or new transistor types.
- No new manufacturing requirements (such as new fabs or new machines).

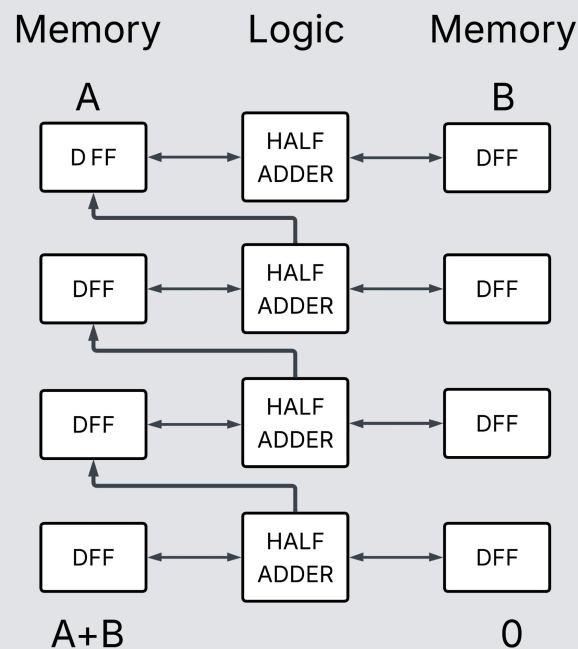
The memory element we require to implement the proposed architecture does not require research for "special compute cells". Our design uses a mature, standard memory cell known as edge-triggered D-type flip-flops (DFF). **The innovation is in the mathematical algorithm which redefines addition, and that can be executed with an array of ordinary components (DFFs+AND/XOR gates).** Some of the biggest advantages of using DFF's are:

- They are part of every **standard cell library.**
- The **fundamental building block** of all digital design.
- DFFs are the **easiest on-chip memory elements to manufacture.**
- Perfectly regular, dense, and have **decades of speed, power, and area optimization.**

This Adder architecture is not an incremental improvement, **it's a complete rethinking of how mathematical operations are executed.** The result is an arithmetic core that redefines processing power. If applied to Bitcoin mining ASICs, **this architecture enables more hash engines per wafer, significantly lower J/TH,** and greater operational resilience through market cycles. Our IP design delivers not just a better mining chip, but a new architectural paradigm for ultra-efficient, deterministic computation applicable across a wide range of high-performance use cases.

CORE IP: SIMPLE AND LINEAR FAST ADDER (SLFA)

Our SLFA serves as the fundamental building block for optimizing various processing cores. For example, the 64 rounds of the SHA-256 compression function constitute close to 90% of Bitcoin mining's computational workload. The SLFA can be generalized to **a rectangular grid of standard memory cells, capable of executing the 64-round compression function In-Memory, restricting data movement between columns.** The regular, Manhattan-routed grid of streets and avenues yields predictable wire lengths and minimal routing congestion enabling higher clock frequencies and lower operating voltages, saving energy and reducing thermal signature.



The theoretical foundations, mathematical algorithms, circuit logic, and RTL design have been detailed in numerous international conferences and peer-reviewed publications. An **international PCT patent application is currently in national phase in eight key jurisdictions: US, China, Japan, Korea, India, Singapore, UK, and Canada.** The patent has received a favorable written opinion from the International Searching Authority (USPTO) with all 15 claims approved without amendment—a rare outcome indicating strong novelty and industrial applicability. **A third-party valuation from InTraCoM Group (Germany-based global patent valuation experts) places current value at +\$8 Million USD.** InTraCoM has extensive experience working with international blue-chip corporations, financial institutions, and technology-driven organizations, reinforcing the credibility and institutional rigor of this valuation. Additionally, we have **a comprehensive technology and market report from Anuation Research & Consulting LLP, including a detailed market study that illustrates the practical usability of the product and demonstrates real-world applications** of the underlying technology, which can be shared to support further evaluation and commercial understanding.

PRODUCT & TECHNOLOGY

We don't just place a computing engine next to a memory subarray. We have solved the memory-wall problem at the algorithmic (mathematical) level, not the physical (material) level. **By redefining the very algorithm for addition, proximity of memory and logic is possible at the bit level.** Our patented Simple and Linear Fast Adder (SLFA) core is directly embedded within on-chip Flip Flop arrays, dramatically reducing data movement between memory and arithmetic cores. **Switching and routing activity is minimized, lowering energy per operation while increasing throughput**—particularly in hashing-dominant workloads (mining), vector/matrix operations (AI, graphics), and signal processing. This linear adder can be generalized to a **rectangular grid of identical nodes connected with Manhattan-style routing of streets and avenues, that executes addition of multiple inputs.** Rather than delivering incremental process/node gains, **this is an architectural shift that meaningfully expands the performance-per-watt achievable** in ASIC designs.

- **30-40% Improvement in Bitcoin Mining** Throughput and Efficiency.
- **+75% Projected Increase in Mining Profitability.**

Yearly Profits per 1,000 ASICs

Scenario	Bitcoins Generated	Energy Cost (MUSD)	Profit in Bitcoins	Profit Increase
Current ASICs	50	2.34	27	--
+30% Time & Energy Efficiency	65	1.7	48	+75%
+40% Time & Energy Efficiency	70	1.4	56	+100%

INNOVATION AND DIFFERENTIATORS

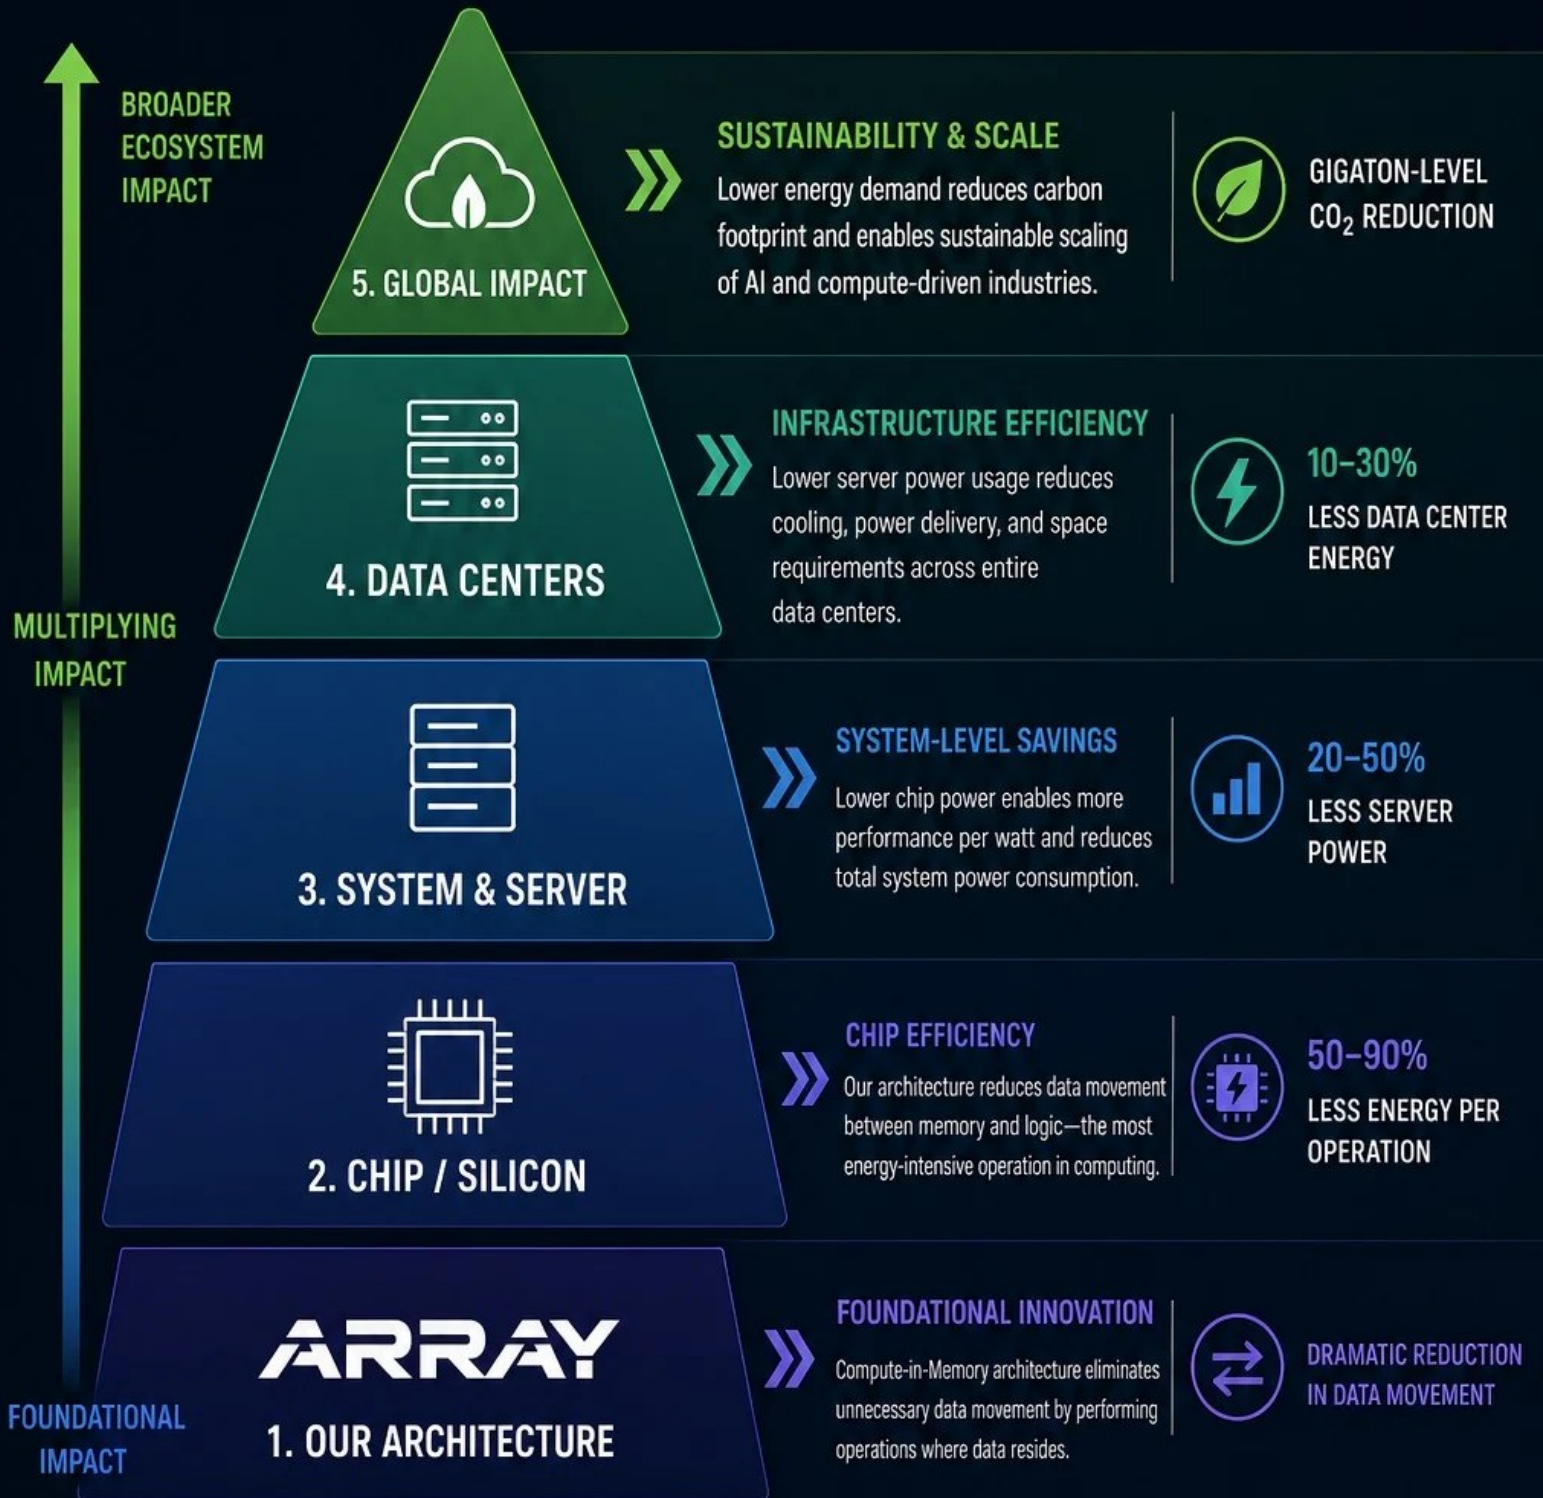
Our IP's value proposition is digital, precise, scalable, and immediately fabricable CIM arithmetic cores using the industry's most mature and cost-effective memory element: the CMOS Flip-Flop. **Once the GDSII prototype design files are complete, it can be manufactured in a CMOS logic process (TSMC, Intel, Samsung) immediately.** This general-purpose adder can be easily scaled and adapted to execute complex operations. Unlike other non-Von Neumann architectures, **our IP is fully CMOS-compatible and manufacturable using standard foundry flows,** no high-risk R&D cycles. This is achieved through our new addition algorithm that is fundamentally different from the carry-over paradigm of addition. **We did not improve the old algorithm. We replaced it.**

- **Manufacturability:** Our design is manufacturable with standard CMOS technology (TSMC, Samsung, Intel) with dramatically simplified place-and-route, reduced layers, higher yield, and faster time-to-market—requiring no exotic materials or processes due to its predictable, regular layout from design through production.
- **Performance:** The architecture delivers logarithmic average time, a constant power profile, and linear scalability of bit width.
- **Targets a Foundational Problem:** While most solutions address the off-chip memory-wall through optimized hierarchies and near-memory computing, we provide the fundamental solution to the on-chip memory-wall—the separation between logic and registers—which is the only real memory-wall remaining.
- **Wide Applicability:** Unlike niche CIM solutions that accelerate specific processes, our SLFA is a universal building block for high-performance arithmetic cores across massively parallel, arithmetic-heavy domains—from Bitcoin mining to AI edge devices and DSP—because we solved addition of multiple inputs at the mathematical level, the most fundamental operation in almost any processor.
- **Market Resilience:** In the Bitcoin Mining ASIC industry, this means our chips will remain profitable during price dips, our customers will survive bear markets competitors won't, and we will control the mining efficiency curve for years.

Comparison Table

	Traditional Adder Architecture		Non Von Neumann Architectures			
	Ripple Carry-Over (1st Gen)	Parallel (Carry-Look Ahead, Carry-Save, Kogge-Stone, etc.)	Near-Memory Computing	Compute-In-Memory Transistor Arrays (FeFETs, MRAM, ReRAM, etc.)	ASIC (On-Chip Pipelines)	Simple and Linear Fast Adder (Bit-Level Embedded Logic & Memory)
Competitors	Low-Power Systems Texas Instruments	Traditional Processors Intel, Nvidia, IBM, AMD, etc.	NeuroBlade, Nvidia, MythicAI	IBM, Intel, Samsung, SK Hynix, MythicAI	BitMain, Etched, Nvidia, Cerebras, NeuroBlade	Our IP
Energy (Efficiency)	◐	○	◑	●	●	●
Time (Latency)	○	◐	◑	●	●	●
Manufacturing (Costs and Requirements)	●	◐	◐	○	◐	●
Technology (R&D)	●	◑	◐	○	◐	●
Applicability (Use Cases)	◐	●	◐	◐	◐	●
Scalability (Bit Length)	●	◐	◐	◐	◐	●

Our architecture addresses a foundational inefficiency in the computing stack by **reducing data movement**, unlocking significant energy savings across data centers and compute-intensive applications.



THE COMPOUNDING BENEFIT

Small savings at the silicon level compound into massive impact at global scale.

=



TRANSFORMING THE ENERGY ECONOMICS OF COMPUTING

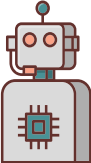
LESS DATA MOVEMENT. **LOWER ENERGY.** GREATER IMPACT.

POTENTIAL MARKET SIZES

The potential market size for our CIM arithmetic core is massive (**arithmetic-heavy processor cores are the main engine behind: AI & Machine Learning, Bitcoin Mining, Digital Signal Processing, GPUs, Cryptography**).


Artificial Intelligence **\$200USD/year in MP**

Matrix multiplication is the core operation of neural networks. It directly affects the training speed and scalability of AI models.




Graphics and Gaming **\$100USD/year in MP**

Transformations like rotation, scaling, and translation in 3D graphics rely on matrix operations.




Blockchain and Bitcoin Mining **\$60USD/year in MP**

Matrix multiplication is fundamental to many cryptographic algorithms used for secure key exchange, encryption, and decryption.




Digital Signal Processing **\$40USD/year in MP**

DSP for audio, image, and video data relies on matrix multiplication for tasks like imaging, filtering, coordinate transformations, and data compression.




Data Analysis and Big Data **\$60USD/year in MP**

Principal Component Analysis (PCA) and ML models leverage matrix multiplication to analyze correlations and patterns in massive datasets.




Cryptography and Security **\$30USD/year in MP**

Matrix multiplication is fundamental to cryptographic algorithms used for securing sensitive data, especially in real-time applications.




Scientific Simulations **\$30USD/year in MP**

In fields like physics, chemistry, and weather modelling, matrix operations are crucial for solving simulations and numerical methods.



Optimization Problems **\$20USD/year in MP**

Fast matrix multiplication accelerates decision-making and problem-solving in real-time systems, from logistics to robotics.



***Market Sizes and Annual Spending on Micro Processors per Year** are derived from studies and industry reports from Gartner, IDC, McKinsey, MarketsandMarkets, Precedence Research, and chip procurement trends.

TOTAL ADDRESSABLE MARKET

Biggest Industries:

- Blockchain and Bitcoin Mining
- Artificial Intelligence
- Computer Graphics and Video Gaming
- Scientific Computing and Simulation
- Digital Signal Processing
- Data Analysis and Big Data
- Optimization Problems
- Cryptography and Security

Global Annual Spending on Microprocessors:

\$540 Billion USD

SERVICEABLE AVAILABLE MARKET

Example Industry:

- Blockchain and Bitcoin Mining

Annual Spending on Mining Processors:

\$60 Billion USD

SERVICEABLE OBTAINABLE MARKET

Specific Product:

CIM Arithmetic-Logic Core for
Blockchain and Bitcoin Mining

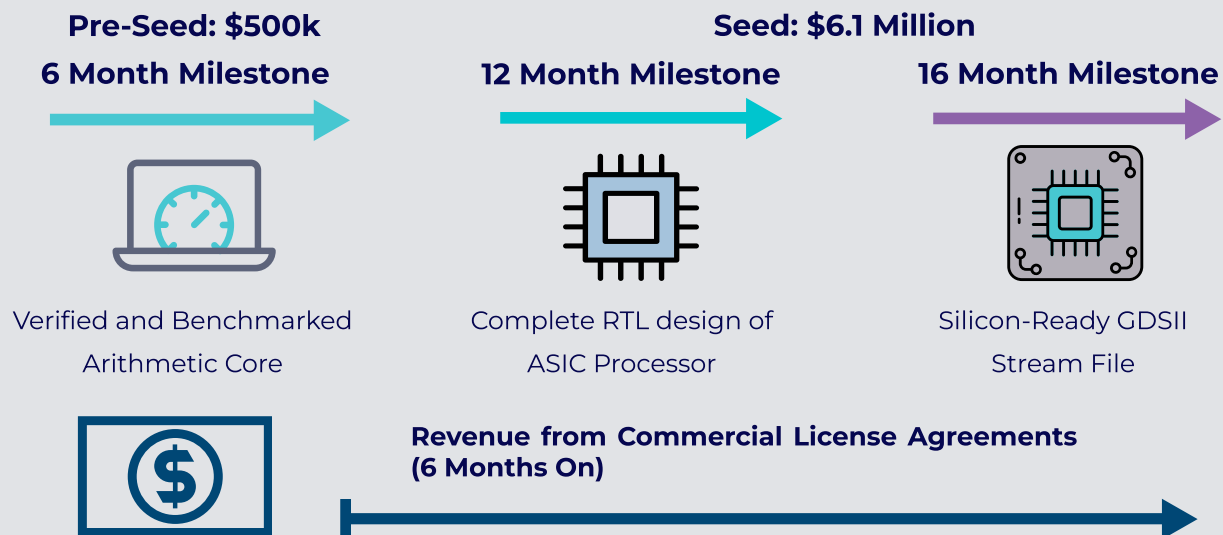
1% of Annual Spending on Mining Processors:

\$600 Million USD

BUSINESS MODEL AND PROCESS TO MARKET

Although we are ready to start licensing as of now, every de-risking step will significantly increase the IP valuation. A verified and patented prototype design will be licensed to ASIC design firms and foundries, generating revenue in the form of:

- **Upfront Licensing Fees** (+\$5 Million USD per Licensee, per year)
- **Royalties**
 - % Share in Revenue
 - % Share in Manufacturing Yields (in the case of the Bitcoin mining ASIC industry we will look for in-kind payment).
- **Mining operations** using our Share in Manufacturing Yields (Bitcoin mining yields from our ASICs).
 - Our chips **remain profitable during price dips**.
 - Our customers **survive bear markets** competitors don't.
 - We **control the mining efficiency curve** for years.



COMPETITION (POTENTIAL PARTNERS)

Processor IP & design firms, foundries and other industry elements are not our direct competitors, they are our future partners and licensees. They are working within the Von Neumann architecture's fundamental limitations, and we are selling the IP for a new arithmetic core that makes their entire product lines more efficient. We aim to partner with them, not compete with them. That is why **we see other design firms and ASIC manufacturer's more as potential clients and partners** rather than competitors. Our key differentiator is that we solved the problem with a new mathematical algorithm and architecture. This makes us faster to market and massively more cost-effective.

ANNUAL PROJECTION


Expected Profit/Year for our Partnership

The top three Global ASIC manufacturers each have average annual productivity rates varying between **100-300k ASICs/year**. Below we project three different case scenarios of yearly profits from mining with a **1-3% share of ASIC manufacturing yields** operating with **30 to 40% efficiency and performance increase**.



BASE DATA

Complementary Data for the Projections



TRACTION

2020

Theoretical background.
Mathematical Proofs.

2022

International (PCT)
patent application filed
in key global
jurisdictions.

2024

- IEEE Conference, Publication, and Peer-Review
- Independent IP Valuation
- Third-Party Market Report

2025

Logical Synthesis and
Analysis for In-Memory
Matrix Multiplication
(mathematical operation
behind AI and Computer
Graphics) and Bitcoin
Mining Engine.

2026

Ready for Silicon
Prototyping.

WHAT'S NEXT? (OUR ASK)

We are positioned to initiate early-stage licensing discussions, while recognizing that **validating and benchmarking a prototype will significantly enhance the value and commercial readiness of the IP.** The next steps include emulation, prototyping, and benchmarking. **Each milestone meaningfully reduces technical risk and strengthens validation,** unlocking substantial upside in licensing potential and long-term revenue through recurring fees and royalties.

The next two years are about further **de-risking the technology at a physical level and creating the licensable asset** (GDSII tape-out ready Stream File), with an **investment of \$6.6 Million USD.** Our cash burn-rate will be primarily driven by the engineering team's salaries, signing bonuses, tools, and resources. **This investment is divided into two stages:**

- 1. \$500k Pre-seed funding (6-8 months for Adder Core):** The first stage will take **6 months to complete and require a capital investment of \$500k USD.** We will validate and benchmark key parameters of our CIM Adder Core, **emulated and benchmarked on Virtex 7 FPGAs, completely de-risking and increasing the value of our underlying IP portfolio.** This package—verified RTL, simulation models, and performance data—is an **immediately licensable asset for semiconductor companies.** It proves our **30-40% efficiency gains** and positions us for a larger Series A.
- 2. \$6.6M Seed funding (16 months for SHA-256 Compression Engine):** The second stage, **lasting 10 months, requires a capital investment of the remaining \$6.1 Million USD to prototype a Compute-In-Memory Hash Engine.** Our goal is to deliver GDSII files within 16 months. We will build-up a world-class nine-person team to deliver a complete, tapeout-ready SoC accelerator for AI or Bitcoin Mining. The technical team for executing the development plan is divided into **two teams, one for RTL to Physical Design, and another for Validation/Benchmarking.** This development plan includes physical design (RTL to GDSII flow), verification, and lab validation—a fully manufacturable prototype ready for **licensing to ASIC design firms, manufacturers, and foundries by the 22nd month.**

Technology Readiness Level (TRL-3): From Math Theory, to Licensable IP, to SoC

	TRL	Definition	Exit Criteria
Completed (Founder Funded)	1	Basic principles observed and reported.	Peer reviewed publication of research underlying the proposed concept/application. (Peer-reviewed Publications, International Conferences, etc).
	2	Technology concept and/or application formulated.	Documented description of the application/concept that addresses feasibility and benefit. (Patent Application, Valuation, and Market Report).
First Stage Investment- 6 Months \$500K USD	3	Analytical and experimental critical function and/or characteristic proof of concept.	Documented analytical/experimental results validating predictions of key parameters.
	4	Component and/or breadboard validation in laboratory environment.	Documented test performance demonstrating agreement with analytical predictions. Documented definition of relevant environment.
Second Stage Investment- 10 Months \$6.1 Million USD	5	Component and/or breadboard validation in relevant environment.	Documented test performance demonstrating agreement with analytical predictions. Documented definition of scaling requirements.
	6	System/sub-system model or prototype demonstration in an operational environment.	Documented test performance demonstrating agreement with analytical predictions.
Hand-Over/Joint Development Licensing Revenue	7	System prototype demonstration in an operational environment.	Documented test performance demonstrating agreement with analytical predictions.
	8	Actual system completed and "flight qualified" through test and demonstration.	Documented test performance verifying analytical predictions.
	9	Actual system flight proven through successful mission operations.	Documented operational results.

ROAD MAP

Every time we further derisk, our IP will have a revaluation. That is why we will not be delivering just a paper concept or a partial prototype. **Our goal is to design a novel hardware architecture, prove correctness of its functionality, characterize its performance and generate Pre-Silicon files**, in a minimal timeline. The next steps include emulating, prototyping, and benchmarking in collaboration with the University of Guadalajara. We have laid out a **16-month development plan to prototype a CIM Arithmetic Core and a fully functional block that executes the 64-round SHA-256 compression function In-Memory**.

A three phase approach separates the Proof-of-Concept of the arithmetic core, from the complex task of building a full system around it, and then taking it to Pre-Silicon files. These are three distinct challenges. We can fast track a 16-month path to achieve a fully licensable and benchmarked product (GDSII tapeout-ready files), given the right funding and strategic partnering.

- 1. Prototype Arithmetic Core (6 Month Mark)
- 2. Integration to ASIC Design Standards (12 Month Mark)
- 3. Integration to ASIC Manufacturing Standards (16 Month Mark)

1. Prototype Arithmetic Core

A Fully Functional and Verified Core arithmetic unit (a 32-bit adder/multiplier based on FAU) is emulated and completely verified for functional correctness. We will design, verify, and benchmark the FAU's "local operability", functional correctness, performance, and power characteristics in isolation, proving its fundamental advantage.

3. Integration to ASIC Manufacturing Standards

A complete "prototype" as an industry-grade pre-silicon GDSII files, ready to be sent to a foundry for fabrication. This asset is incredibly valuable for licensing.



2. Integration to ASIC Design Standards

Integrate the arithmetic core into a complete System-on-Chip for a SHA-256 hashing pipeline. This includes building the surrounding infrastructure of the core for control logic, data movement, etc. to create a fully functional system.

License & Trade Agreements

Integration to client's products.

ROAD MAP

We will deliver a complete, proven IP package built around the GDSII files, conforming a licensable IP asset, as it provides the manufacturer with everything they need to confidently integrate it into their own SoC designs. It is our final objective to provide manufacturers with the finished, tested blueprint and all its documentation.

I. Prototype Arithmetic Core: Develop a time and energy-efficient Arithmetic Core.

- i.* Onboarding
- ii.* Implement Registers in Verilog
- iii.* RTL Design of CIM Adder
- iv.* Verify and Test
- v.* Emulate and Benchmark
- iv.* Publish Results

II. Integrate Arithmetic Core into a Mining ASIC Architecture: We integrate our Arithmetic Core to a Bitcoin mining architecture with existing ASIC design standards.

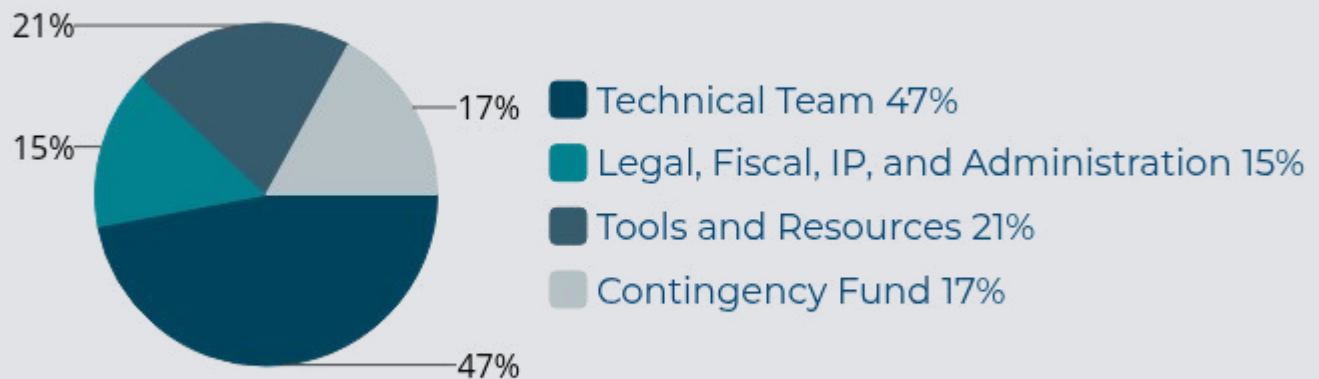
III. Tapeout-Ready GDSII File Stream: Integrate our prototype to existing manufacturing standards for a fully functional SoC mining architecture, and adapt our technology to manufacturers' needs, in order to secure contracts. This is the final, signed-off database that describes the complete geometric layout of the entire chip in polygon form. It is the universal format sent to the foundry (TSMC, Samsung, etc.) to create the photomasks used in fabrication. It represents a design that has passed all functional verification, timing checks, and physical verification rules. It consists of the following elements:

- **Gate-Level Net List:** Post-synthesis, technology-mapped description of the circuit. A list of standard cells and their connections. The licensee uses it for their own final verification.
- **Timing Sign-off:** The final Static Timing Analysis (STA) reports showing all timing paths are closed (i.e., the chip will run at the target speed).
- **Power Sign-off:** Power integrity analysis (IR Drop, Electromigration) reports proving the power network is robust.
- **Physical Verification Reports:** Documentation proving the layout passed Design Rule Check) and LVS (Layout vs. Schematic) checks for the target foundry process.
- **System Verilog/UVM Testbench and Test Cases:** The complete environment used to verify the design. This allows the licensee to run their own simulations and regressions to validate the IP within their specific system.
- **SDC (Synopsys Design Constraints) File:** The file that defines the timing constraints (clocks, input/output delays) for the entire design. This is critical for the licensee's integration and verification flow.
- **Comprehensive Documentation:**
 - Architecture Spec: How the block works.
 - Integration Guide: How to connect it to a system.
 - Verification Plan and Report: What was tested and the results.

IV. License and Commercial Agreements: collaborate with mining ASIC designing teams.

USE OF FUNDS

Our detailed budget shows that **\$6.6 million** is the capital required to de-risk this technology completely and deliver a finished, tapeout-ready design to our manufacturing partners. This funding will allow us to **1) Build and retain a world-class team for a period of 24 months** to see this through from architecture to sign-off and support our clients during their final implementation, **2) Procure the essential industry-grade software and hardware** necessary to design and test to industry standards, and **3) Execute our three-phase 16-Month Timeline**. With this investment, we will not be delivering just a paper concept or a partial prototype; we will be delivering the next generation of processor in a final GDSII Stream File. This asset is what we will license, and its value is confirmed by a third-party valuation of the underlying IP.



TIMELINE

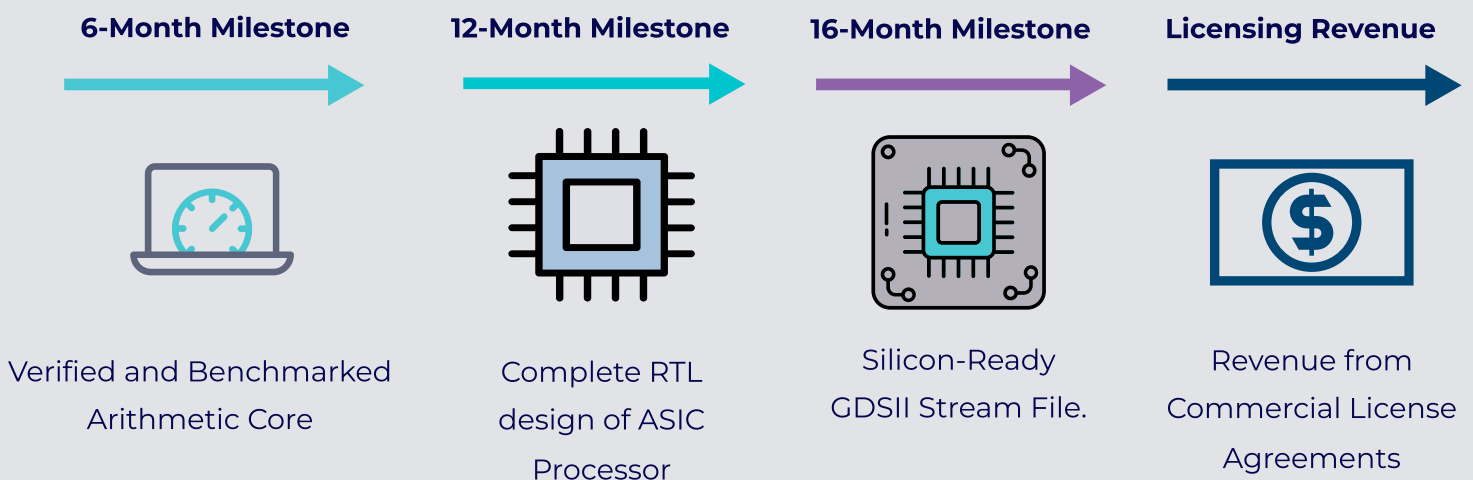


Table Summary for Investors and Partners

	Arithmetic Processing Core	SHA-256 Compression Engine
Investment Required	\$500, 000 USD	\$6' 100,000.00 USD
Primary Deliverable	Verified RTL core (Prototype) + FPGA benchmarks + simulation models to demonstrate performance and commercial impact.	Full integration of our Arithmetic Core into a Tapeout-ready GDSII File Stream for a Bitcoin mining System-on-Chip.
Delivery time	6 months	10-12 months
Industrial Applicability	<ul style="list-style-type: none"> • AI accelerators • Digital Signal Processing • Cryptography and Security • Computer Graphics 	Blockchain and Bitcoin Mining
Benefit, broad overview	This package—verified RTL, simulation models, and performance data—is an immediately licensable asset for semiconductor companies. It proves our 30-40% efficiency gains and positions us for a larger Series A.	Our architecture is expected to deliver a meaningful performance and efficiency advantage over existing market solutions, positioning us to capture a significant share of a rapidly expanding market opportunity estimated at +\$60 Billion USD.
Business model	<p><i>Licensing Arithmetic Core:</i></p> <ul style="list-style-type: none"> • Upfront Licensing Fees • Royalties 	<p><i>Mining Bitcoin:</i></p> <ul style="list-style-type: none"> • Upfront Licensing Fees • Percentage share of manufacturing yields (in-kind payment) to be employed in Mining
Shareholding	7% – 10% (Preferential Stock)	15% – 20% (Series A)
Implied Post-Money Valuation	\$5-7 Million USD	\$30-40 Million USD

LET'S PARTNER

Our primary market strategy is to leverage our foundational IP to co-design, manufacture, and operate **a new generation of highly efficient Bitcoin mining ASICs**, creating a dominant position in the market through strategic partnerships. As a built-in, passive fallback, we have a clear path to monetization through licensing our proprietary arithmetic core to established processor manufacturers, ensuring value capture with minimal additional investment.

We are building strategic collaborations and partnerships with academic, and research experts, as well as industry leading institutions to advance this unique focus into the computational throughput of Mining ASICs to **establish a new standard reference in global mining output**.

- We are seeking **\$6.6 Million USD** of investment to develop a licensable asset and co-design efficient mining ASICs.
- Our patent has received a Third Party valuation by Intra-Com Group for **\$8M but this valuation will increase in the next 12 months (as of January 2026) as countries grant the patent**.
- We are offering a **stake in Commercial and Licensing Agreements and other forms of revenue** stemming from our ASIC Prototype for the duration of the patent (through 2042).


THANK YOU!

 www.fastarithmeticunits-dataroom.com



Juan Pablo Ramirez

Architect, Project Manager & CEO
jramirez@binaryprojx.com



Pablo César Vázquez Estrella

Chief Financial Officer
pablo.vazquez@binaryprojx.com

www.fastarithmeticunits-dataroom.com