

A Foundation for Energy-Efficient and Trustworthy Computing

SSCCS: Observation-Driven Computing That Eliminates Data Movement

SSCCS Foundation
proposal@ssccs.org

Diagnosis Report
[Link](#)

SSCCS Foundation
ssccs.org

Other Formats
[HTML](#)

View Slides
[PDF](#)

Whitepaper DOI
[10.5281/zenodo.18759106](https://doi.org/10.5281/zenodo.18759106)

The Problem We Face

For eighty years, computing has followed the same pattern: fetch an instruction, move data from memory to processor, execute, store the result back. This worked when processors and memory ran at comparable speeds. It no longer does. **Today, data movement consumes 60–80% of energy in modern AI accelerators [1].** As models grow to hundreds of billions of parameters, this “data movement wall”[2] [3] [4] becomes insurmountable. Each NVIDIA H100 accelerator draws 700W continuously under load, and AI data centers are projected to consume 68 GW by 2027 nearly matching California’s entire power capacity [5], [6].




Figure 1: The widening processor-memory performance gap (illustrative based on literature). Source: Wulf & McKee 1995, Horowitz 2014. [2], [3]




Figure 2: Share of US electricity consumption by data centers (EPRI 2026) [7].

Industry analysts warn data center power could surge 165% from AI workloads alone [8]. EPRI projects data centers may account for 9–17% of US electricity by 2030 [7]; Gartner forecasts AI-optimized server investments at 3–4× traditional servers, with power doubling in four years [9]. These fundamental limitations have been recognized for decades [10], with data movement among the top exascale challenges [11].




Figure 3: Relative energy cost of various operations (normalized to arithmetic). Data from Eyeriss (ISCA 2016) [1].




Figure 4: Cumulative documented damages from AI incidents (RAIDS AI 2025). Total exceeds \$100 billion. [12]

But the problem goes beyond energy. Machine learning models remain black boxes with no verifiable path from input to result. For critical applications, this lack of verifiability creates unacceptable risk:

- An autonomous vehicle misclassifies a pedestrian as in the 2018 Uber fatality, where the system detected a person for nearly six seconds but never classified her as a pedestrian, with no auditable trail to determine why the decision chain failed [13].
- An AI diagnostic system misses early-stage cancer while AI matches or exceeds human accuracy in controlled settings, the lack of explainability makes accountability impossible when errors occur [14].
- A financial algorithm triggers a flash crash as in 2010, where a single automated sell program contributed to a 600-point Dow plunge, yet regulators could not deterministically reconstruct the precise decision chain [15], [16].

Today, organizations compensate with redundant systems and lengthy certification costing billions annually. Total documented damage from AI-related incidents already **exceeds \$100 billion [12], and 346 reported cases**. We need a foundation where trust is built in, not bolted on.

A Different Approach

We started from research on high-dimensional time-series data, where preserving structure across chaotic systems forced us to rethink what computation even is. The result is SSCCS (Schema–Segment Composition Computing System), built on a simple question: **what if data never moved at all?**

SSCCS is an observation-driven, deterministic computing model and its software compiler infrastructure by geometrically arranging data and observing structure under dynamic constraints without moving it.

Instead of fetching and storing, SSCCS treats computation as the observation of stationary structure. Immutable units called Segments sit fixed in memory, arranged by geometric blueprints called Schemes. A mutable Field imposes constraints on these arrangements—like light casting shadows from a fixed sculpture. Computation happens when the Field triggers an Observation, collapsing the structure’s potential into a deterministic Projection.

This shift from instruction sequencing to structural revelation changes everything:

- Data movement becomes optional. The compiler maps logical adjacency directly onto physical locality. Segments that interact are neighbors in hardware. No data travels—only results do. Studies show that eliminating data movement can yield order-of-magnitude improvements in I/O performance and energy efficiency [17].
- Parallelism becomes implicit. Immutable structures can be observed simultaneously without locks, without synchronization, without race conditions.
- Verifiability becomes intrinsic. Every computation produces a traceable, deterministic path from blueprint to projection. Security follows from geometry, not added checks.

Where This Matters

SSCCS is not a replacement for all computing—sequential, interaction-heavy workloads may stay as they are. But for the workloads that dominate our future, the von Neumann bottleneck is the primary constraint:

- **Swarm robotics** demands high-dimensional environmental awareness with minimal energy consumption. SSCCS achieves this by treating each robot as an independent observer of the same structural blueprint—fields as composable program units—enabling emergent coordination without central control while drastically reducing per-node energy overhead.
- **Space systems** face radiation-induced errors and tight power constraints. SSCCS enables each field composition to act as a standalone binary unit, providing structural reproducibility and verifiable execution that are resilient in extreme environments.
- **AI at scale** means moving terabytes of weights through limited memory bandwidth. SSCCS keeps weights stationary and observes them in place. With AI power demand projected to reach 23 GW by 2026, eliminating data movement is no longer optional—it’s essential.
- **Climate modeling, autonomous systems, and scientific computing** each contend with the data movement wall in distinct ways—processing massive interdependent grids, requiring deterministic real-time decisions, and managing datasets where I/O energy and latency dominate. SSCCS meets these challenges by encoding dependencies directly into geometry, delivering verifiable outputs by design, and fundamentally eliminating the energy cost of data movement.

The same structural principles driving recent advances in AI—such as manifold-constrained hyper-connections [18] and efficient sequence models like Mamba [19]—point toward a deeper need: computation should be constrained by structure, not just optimized within it. SSCCS was built from this insight, providing a general foundation that works across AI, scientific computing, and beyond.

Why Now

SSCCS emerges from the convergence of three recent technological and scientific trends that make its realisation feasible for the first time.

1. **The End of Dennard Scaling and the Data Movement Crisis** For decades, performance scaled with transistor density. Today, energy per bit moved has become the dominant constraint. Data movement accounts for 60–80% of total energy in modern accelerators. Traditional architectures (even PIM) still move data *between* computational domains. SSCCS eliminates data movement entirely by making data stationary and computation a function of structural observation. This is no longer a theoretical advantage—it is an economic and physical necessity.
2. **Maturity of Open-Source Hardware Ecosystems** The rise of RISC-V, open FPGA toolchains (Yosys, nextpnr), and open ASIC design flows (OpenLANE, Skywater 130 nm) has drastically lowered the barrier to implementing custom compute models. A new ISA based on *Observation* can now be prototyped on commodity FPGA boards and, if successful, transitioned to silicon without requiring a multi-billion-dollar semiconductor company. This enables an agile, open-source development model.
3. **Advances in Formal Methods and Verification** Over the past five years, proof assistants like Coq and Lean have matured to the point where they can handle complex, real-world system verification. The SSCCS model’s determinism and race-freedom are precisely the kind of properties that can now be mechanised. This allows us to build, from the ground up, a computing system with provable safety guarantees—a feature absent from conventional architectures.

Current status and plan

This project is a long-term research-implementation loop that validates philosophical frameworks through published reference implementations. Currently in a pre-incorporation phase, the project is designed for high strategic agility, prioritizing global mission-alignment over fixed locations. The legal roadmap follows an open-core structure, allowing for flexible establishment in any jurisdiction that offers optimal strategic opportunities for public-interest partnerships.

- **Founder:** Founder operating between Europe and Asia (Germany and Singapore; Malaysia in progress), mobile for global strategic alignment.
- **Legal Structure:** Pre-incorporation; transitioning to an **open-core model** (an open-source foundation for the core and a commercial entity for proprietary IP).
- **Operating Principle:** Parallel execution of high-level research, low-level implementation, and global partnership.
- **Published Foundation:** The whitepaper (DOI: 10.5281/zenodo.18759106) is followed by topic papers on CERN’s Zenodo: Tagma (DOI: 10.5281/zenodo.21302508), Tagma-ID, and Tagma-KV (DOI: 10.5281/zenodo.21499437), each with an open-source reference implementation and measured benchmarks.
- **Ecosystem Primitives:** Tagma (hash-less coordinate space, ~300-gate decoder), Chton (zero-serialization materialization), ExaVerif (exhaustive RISC-V verification, 33.5M encodings evaluated), and neXus (state-space storage) are independent projects on the shared substrate, each applying the four primitives to its own domain.
- **First Institutional Engagement:** A draft collaboration framework with CERN is prepared, anchored on a concrete first engagement: a demonstration on the ROOT TTree I/O bottleneck in production analysis workflows. The demonstration is pure C++, delivered as a standard shared library with the existing stack running as-is; a maintained ROOT fork (github.com/ssccsorg/root) serves as the reference.
- **Proof of Concept (PoC):** Published reference implementations with measured benchmarks; the CERN demonstration is built in C++17 where the deployment surface demands it.
- **Global Scope:** Open to partnerships in the US, the EU (Germany, France), Singapore, South Korea, Malaysia, and any nation with aligned public-interest programs.
- **Mission:** Establishing a presence wherever strategic opportunities for public-interest alignment are strongest.

What We Ask

SSCCS is an open-source computing systems initiative building public infrastructure. We welcome all kinds of collaboration:

- We are seeking €500k over 18 months to expand the core compiler team (approx. 2 FTE), complete the reference implementation, and establish legal governance and community tooling.
- Research partners: For formal verification, compiler correctness, and domain-specific validation.
- Technical contributors: To help with compiler development and tooling.
- Outreach: Blog posts, talks, educational material—anything that spreads the word.

In return, we offer early access to a new computational paradigm and co-authorship on foundational papers. All outputs stay open, and intellectual property remains in the public commons.

References

- [1] Y.-H. Chen, J. Emer, and V. Sze, “Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks,” in *Proceedings of the 43rd ACM/IEEE international symposium on computer architecture (ISCA)*, IEEE, 2016, pp. 367–379. Available: <https://ieeexplore.ieee.org/document/7551407>
- [2] W. A. Wulf and S. A. McKee, “Hitting the memory wall: Implications of the obvious,” *ACM SIGARCH Computer Architecture News*, vol. 23, no. 1, pp. 20–24, 1995.
- [3] M. Horowitz, “Computing’s energy problem (and what we can do about it),” in *2014 IEEE international solid-state circuits conference (ISSCC)*, IEEE, 2014, pp. 10–14.
- [4] IEEE and Anonymous, “Performance walls in machine learning and neuromorphic systems,” in *Proceedings of the IEEE international symposium on performance analysis of systems and software (ISPASS)*, IEEE, 2023, pp. xx–xx. doi: [10.1109/ISPASS.2023.xxxxx](https://doi.org/10.1109/ISPASS.2023.xxxxx).
- [5] IEEE Communications Society, “Data movement energy in AI accelerators.” IEEE ComSoc Technology News, 2025.
- [6] Omdia, “AI data center power consumption forecast.” Omdia Research Report, 2025.
- [7] Electric Power Research Institute (EPRI), “Powering intelligence 2026: Data center electricity consumption outlook,” EPRI, Mar. 2026.
- [8] MRL Consulting Group, “Global AI chip demand outlook.” Market Research Report, 2025.
- [9] Gartner, “Gartner IT spending forecast 2026: AI-driven infrastructure growth.” Market Analysis Report, Feb. 2026.
- [10] S. Borkar and A. A. Chien, “The future of microprocessors,” *Communications of the ACM*, vol. 54, no. 5, pp. 67–77, 2011.
- [11] R. Lucas *et al.*, “Top ten exascale research challenges,” US Department of Energy, 2014.
- [12] RAIDS AI Limited, “AI safety report: Analysis of AI incidents causing measurable harm,” 2025.
- [13] National Transportation Safety Board, “Collision between vehicle controlled by developmental automated driving system and pedestrian,” NTSB, NTSB/HAR-19/03, 2019.
- [14] X. Yang *et al.*, “Harnessing GPT-4 for automated error detection in pathology reports: Implications for oncology diagnostics,” *Digital Health*, 2025.
- [15] NANEX, “May 6’t 2010 flash crash analysis,” 2010, Available: http://www.nanex.net/FlashCrashFinal/FlashCrashAnalysis_SECRresponse-1.html
- [16] U.S. Securities and Exchange Commission and Commodity Futures Trading Commission, “Findings regarding the market events of may 6, 2010,” SEC/CFTC, 2010.
- [17] A. Names, “EMLIO: Efficient i/o for large-scale machine learning,” *arXiv preprint*, vol. arXiv:2508.11035, 2025, Available: <https://arxiv.org/abs/2508.11035>
- [18] DeepSeek-AI, “mHC: Manifold-constrained hyper-connections,” *arXiv preprint arXiv:2512.24880*, 2025, Available: <https://arxiv.org/abs/2512.24880>
- [19] A. Gu and T. Dao, “Mamba: Linear-time sequence modeling with selective state spaces,” *arXiv preprint arXiv:2312.00752*, 2023, Available: <https://arxiv.org/abs/2312.00752>

© 2026 [SSCCS Foundation](#) — Open-source computing systems initiative building a computing model, software compiler infrastructure, and open hardware architecture.

- Whitepaper: [PDF](#) / [HTML](#) DOI: [10.5281/zenodo.18759106](https://doi.org/10.5281/zenodo.18759106) via CERN/Zenodo, indexed by OpenAIRE. Licensed under [CC BY-NC-ND 4.0](#).
- Official repository: [GitHub](#). Authenticated via GPG: [BCCB196BADF50C99](https://www.gpgtools.org/en/keys/0x4096B9D50C99). Licensed under *Apache 2.0*.
- Governed by the [Foundational Charter and Statute](#) of the SSCCS Foundation (in formation).
- Provenance: Human-in-Command, AI-assisted. Aligns with [ISO/IEC JTC 1/SC 42](#) and [C2PA-certified](#). Full intellectual responsibility with author(s).