



On the Status, Requirements, and Expectations of Neuro-Inspired High-Performance Computing
Johannes Gebert, Qifeng Pan, Lukas Stockmann, Hartwig Anzt, Christian Mayr

NICE 2026, Georgia Institute of Technology, Atlanta, USA

The HPC Landscape in 2026

Traditional HPC:

- Von Neumann architecture
- CPUs + GPU accelerators
- Mature software stacks
- Proven over decades

Neuro-Inspired Computing:

- Hardware-based AI
- Emerging technology
- Medium to large scale systems
- Ready for industrialization?

Three Central Criteria

1. Extreme Scaling Considerations

- Distributed memory + computation
- Local processors with local memory
- Highly performant interconnect in case of communication

2. Distributed & Localized Algorithms

- Fine-grained parallelism
- Avoid global synchronization
- Avoid global communication

3. Event-Based & Asynchronous Communication

- Dynamic sparsity
- Thresholding and information gating

1. AI/LLMs

- Transformer operations
- Attention mechanisms
- Low-power inference
- Edge deployment
- ...

2. Sparse Matrix

- CG/GMRES/... solvers
- FEMs
- CFDs
- ...

3. Optimization

- QUBO problems
- Hopfield networks
- Job scheduling
- Graph partitioning
- ...

A recurring challenge: Toy problems vs. real-world software stacks (OpenFOAM, GROMACS, etc.).
We need toy problems, but we have to estimate if they potentially translate into production.

Traditional HPC:

- Compilers (architecture-specific)
- Libraries (BLAS, FFT, etc.)
- MPI (distributed memory)
- OpenMP (shared memory)
- CUDA/OpenACC (GPU)

Neuro-Inspired Needs:

- Event-driven computation
- Asynchronous memory access
- SNN frameworks
- Massive parallelism support
- Low-power sparse patterns

Neuromorphic Intermediate Representation (NIR):

Hardware-agnostic abstraction enabling portability across platforms.

1. Software Ecosystem

- Compatibility with existing HPC stacks
- Mature drivers and toolchains
- Long-term vendor support
- People & documentation

2. Service & Maintainability

- Guaranteed spare parts supply
- Responsive technical support
- Service-level agreements (SLAs)

3. Security & Multi-Tenancy

- Hardware-level security features
- Resource isolation
- Data protection standards

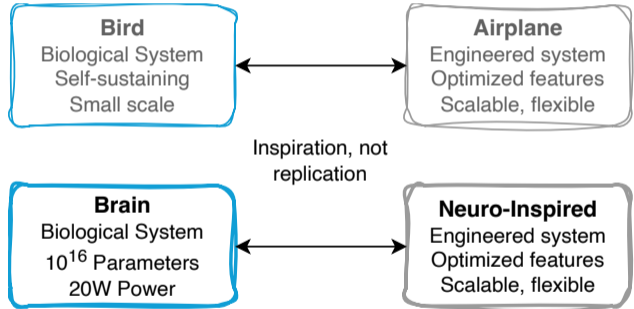
4. Thermal Management

- Air vs. liquid cooling
- Energy efficiency
- Compute density

Bird Analogy

Traditional HPC:

- Cherry-pick best neuro-inspired features
- Combine with traditional HPC components
- Pragmatic approach, not dogmatic



Call to Action

1. Collaboration

- Algorithm developers and hardware designers
- Academia and industry partnerships

2. Pragmatism

- Include non-neuro-inspired components where needed
- Fast interconnects, programmability

3. Software Development

- Treat neuromorphic software as first-class citizen
- Invest in algorithm adaptation

4. Continuing Investment

- Map simulation workloads to neuro-hardware
- Develop AI applications
- Create viable large-scale systems