D E Shaw Research

The ANTON 3 ASIC: a Fire-Breathing Monster for Molecular Dynamics Simulations

Hot Chips 33

24 August 2021

The Anton 3 hardware team

Peter J. Adams[†], Brannon Batson, Alistair Bell, Jhanvi Bhatt[†], J. Adam Butts, Timothy Correia, Bruce Edwards, Peter Feldmann, Christopher H. Fenton, Anthony Forte, Joseph Gagliardo, Gennette Gill, Maria Gorlatova[†], Brian Greskamp, J.P. Grossman[†], Jeremy Hunt, Bryan L. Jackson, Mollie M. Kirk[†], Jeffrey S. Kuskin[†], Roy J. Mader, Richard McGowen[†], Adam McLaughlin, Mark A. Moraes, Mohamed Nasr[†], Lawrence J. Nociolo, Lief O'Donnell, Andrew Parker, Jon L. Peticolas, Terry Quan, T. Carl Schwink[†], Keun Sup Shim, Naseer Siddique[†], Jochen Spengler[†], Michael Theobald, Brian Towles[†], William Vick[†], Stanley C. Wang, Michael Wazlowski, Madeleine J. Weingarten, John M. Williams[†], David E. Shaw

† Work conducted while at D. E. Shaw Research; author's affiliation has subsequently changed.

Molecular dynamics (MD) simulation

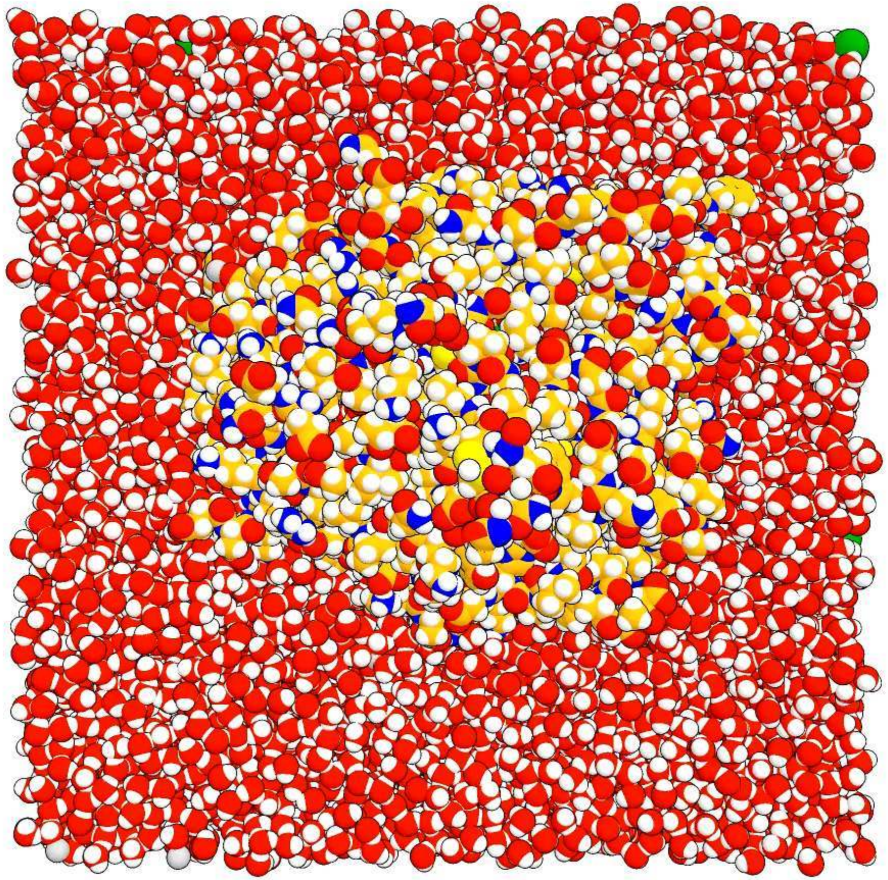

- Understand biomolecular systems through their motions
- Numerical integration of Newton's laws of motion
  - Model atoms as point masses
  - Compute forces on every atom based on current positions
  - Update atom velocities and positions in discrete time steps of a few femtoseconds
- Force computation described by a model: the force field

Biomolecular force fields

$F =$ 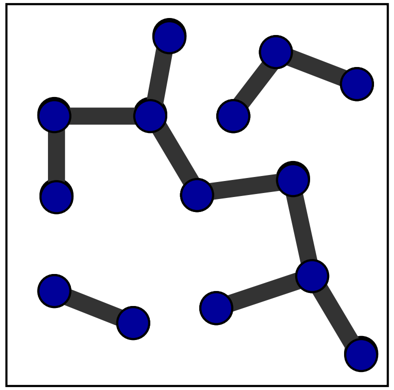 $+$ 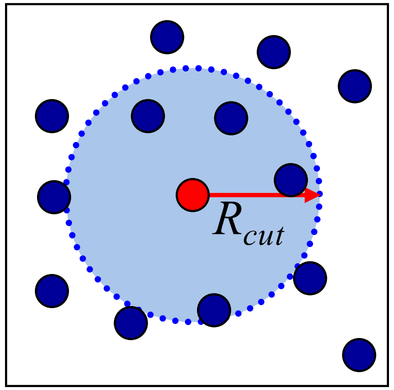 $+$ 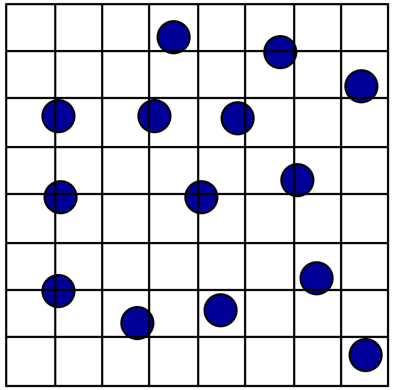

Bonded

Near electrostatic, Van der Waals

Distant electrostatic

Meet ANTON

ANTON 2: The defending champion

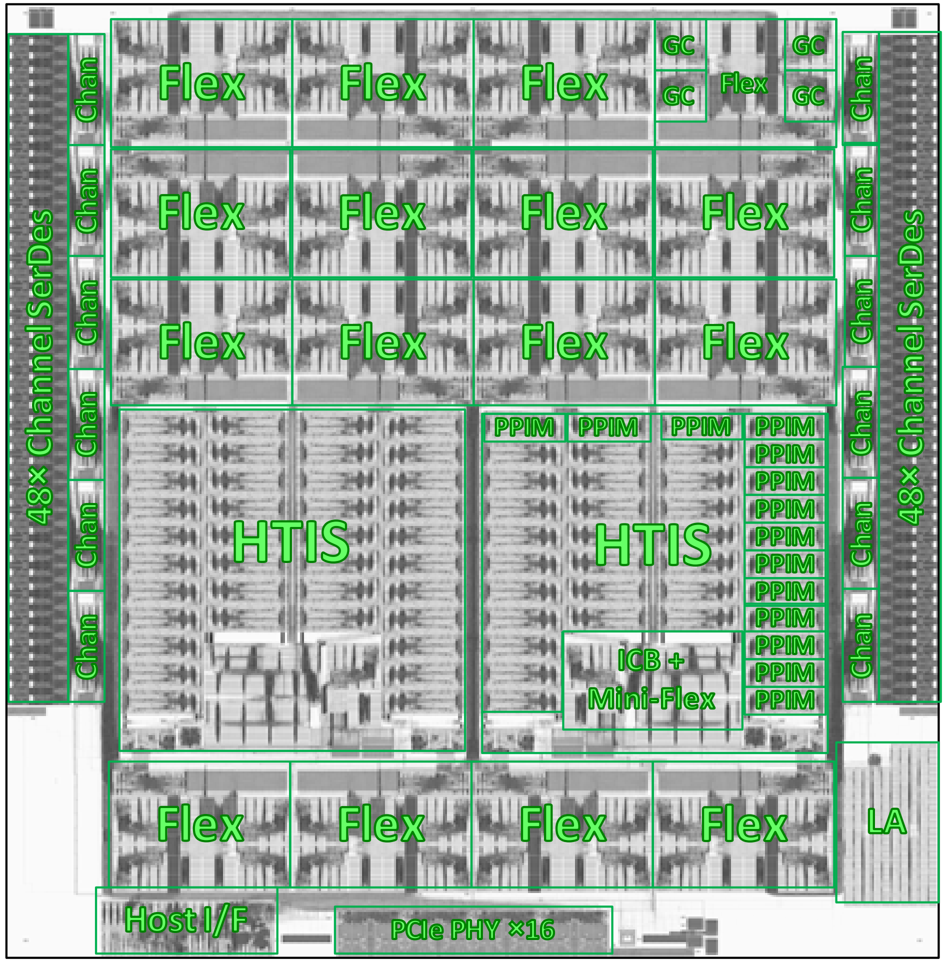

Flex Tile

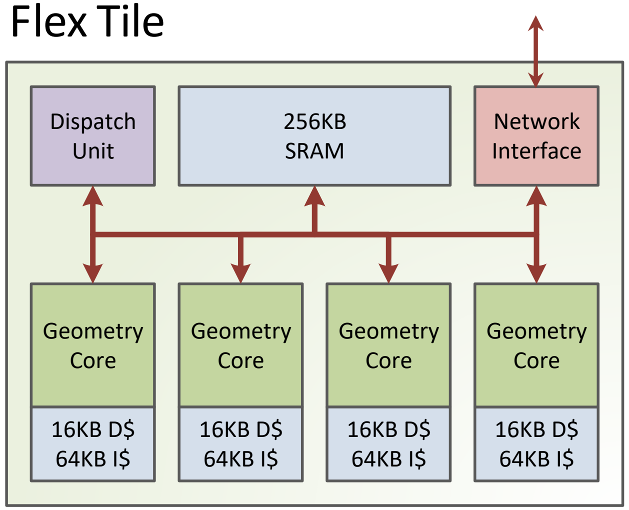

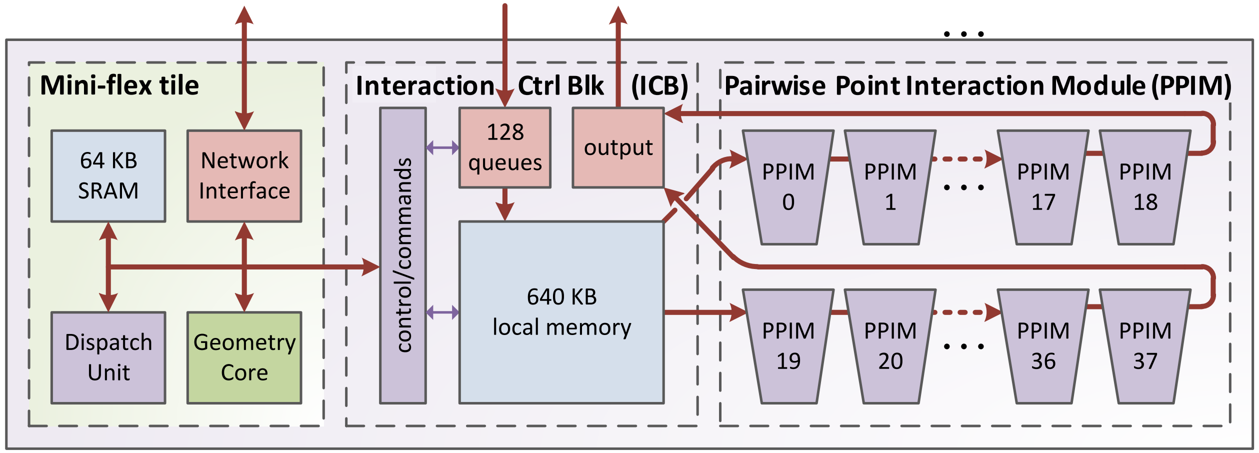

High-Throughput Interaction Subsystem (HTIS)

How do we make it better?

- Increase computational throughput
  - Pairwise force computation
  - Programmable cores
- Address exposed bottlenecks
  - Bond computation
  - Communication bandwidth
- Improve utility
  - Maximum simulation capacity
  - Programmability
  - Flexibility
- Manage design complexity

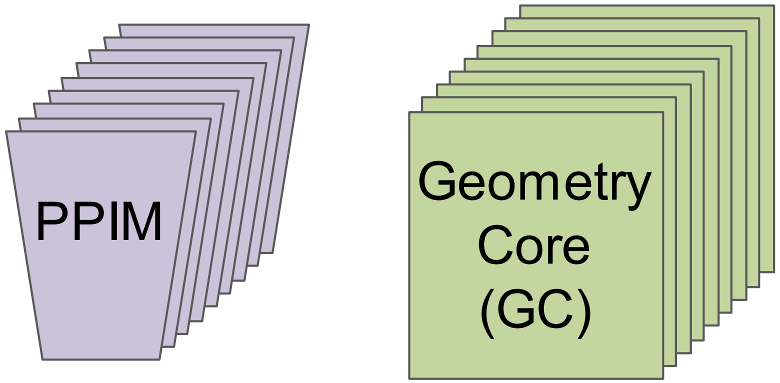

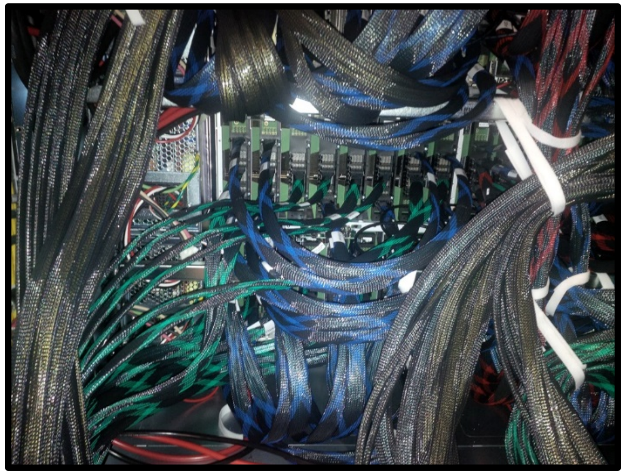

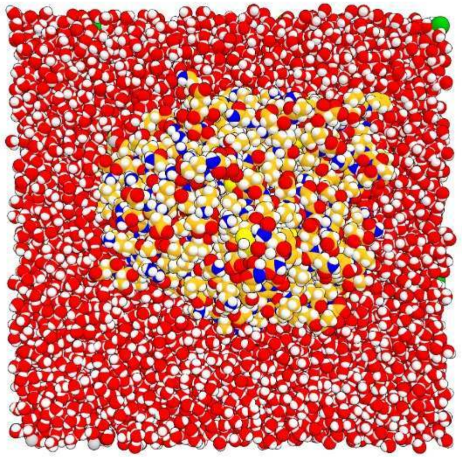

Concentrated compute: The Core Tile

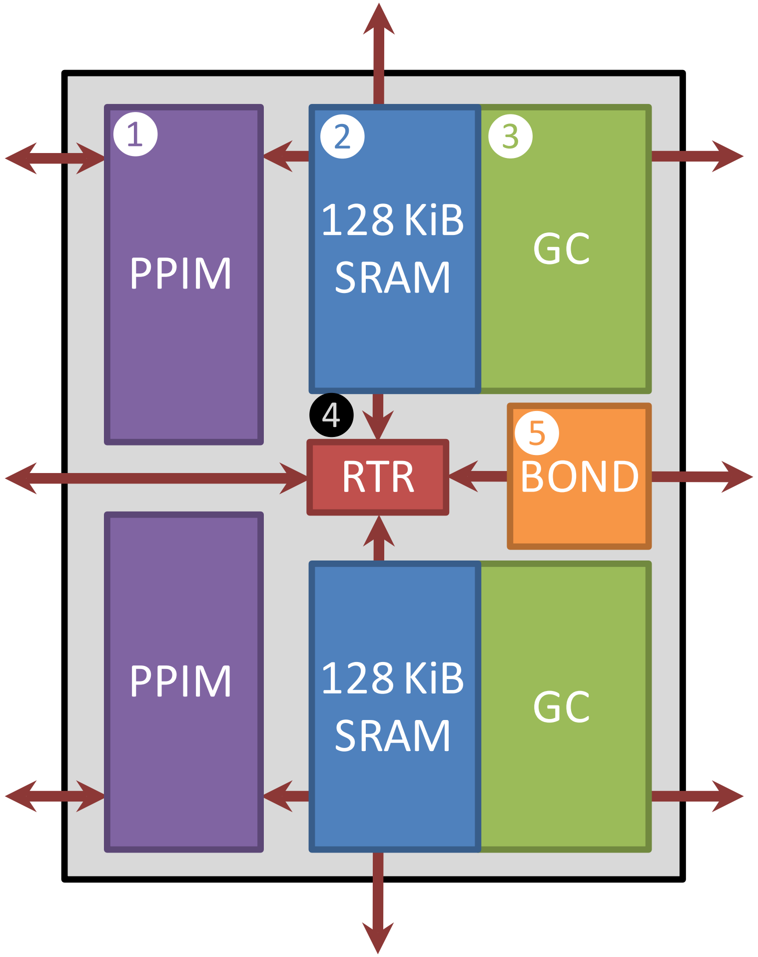

- Evolutionary changes
  - ① Support additional functional forms
  - ② Increase memory capacity
  - ③ Tune instruction set for MD application
  - ③ Increase code density

- **Revolutionary changes**
  - ④ Co-locate compute resources
  - ⑤ Specialize bonded force computation
  - ① Double effective density of pairwise interaction calculation
  - ② ④ Implement fine-grained synchronization within memory and network

D E Shaw Research

Bond calculator

Term	Stretch	Angle	Dihedral / Torsion
Atoms	2	3	4
	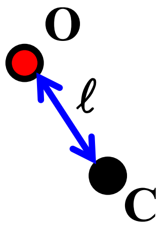	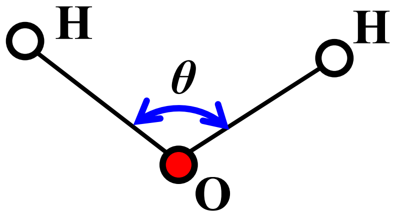	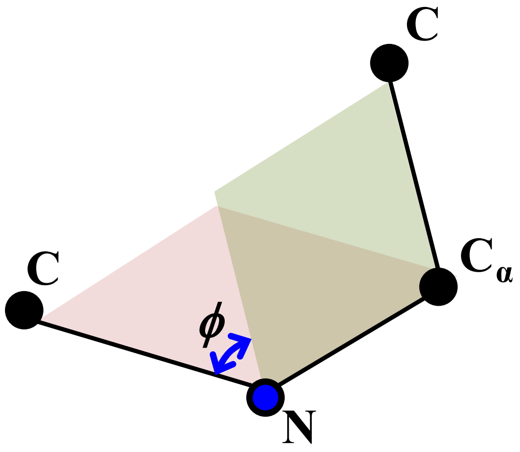
Coordinate (q)	ℓ	θ	ϕ
Potential (V)	$k\,(\ell - \ell_0)^2$	$k\,(\theta - \theta_0)^2$	$\sum_{n \le 6} k_n \cos\,(n\phi - \phi_0)$

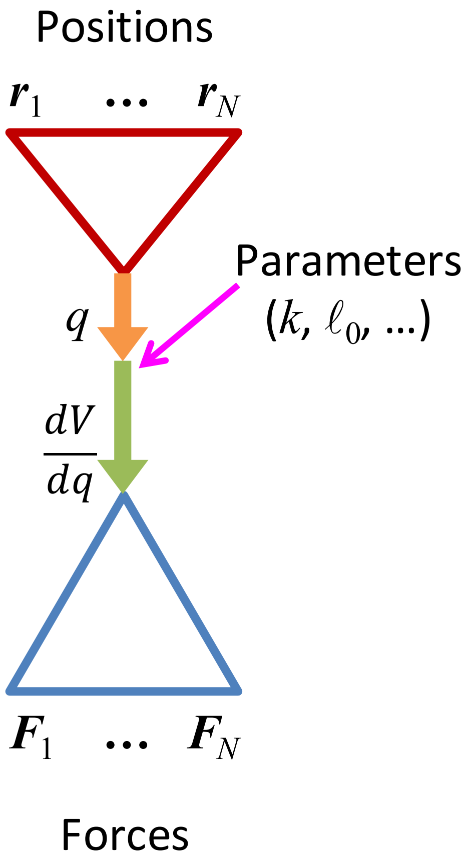

Near versus far

Volume within radius r

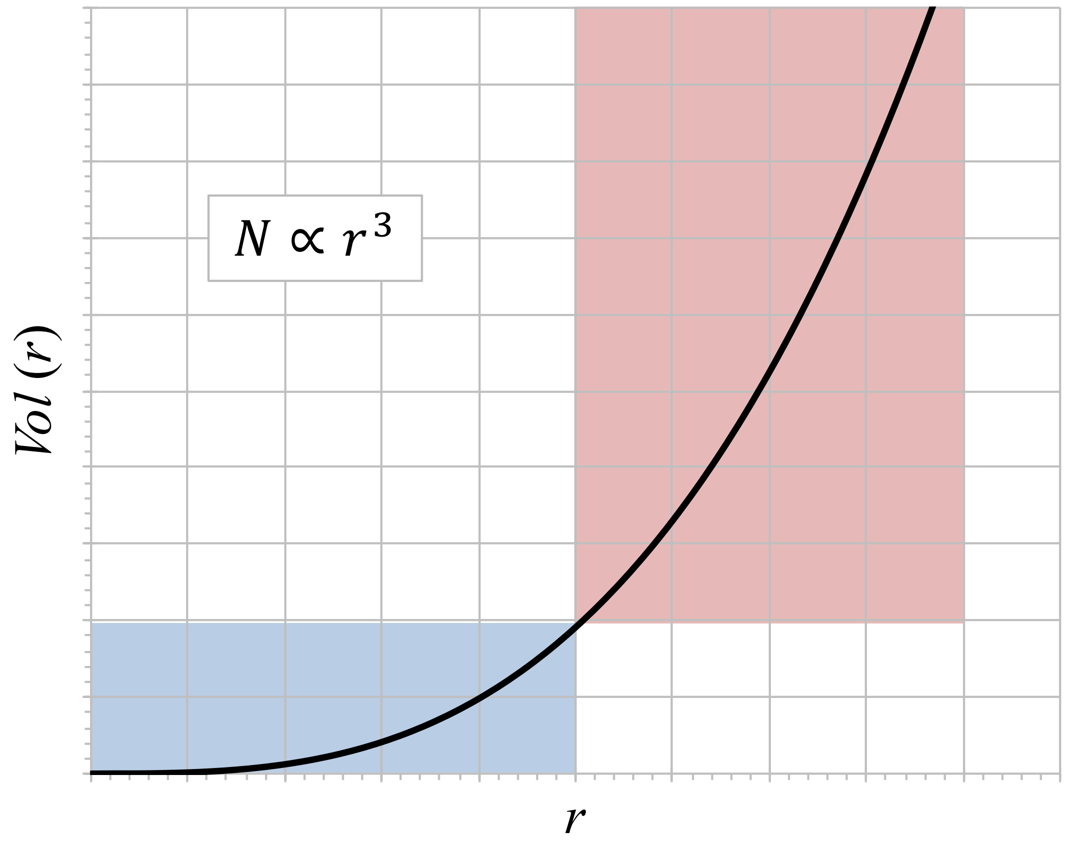

Electrostatic force at distance r

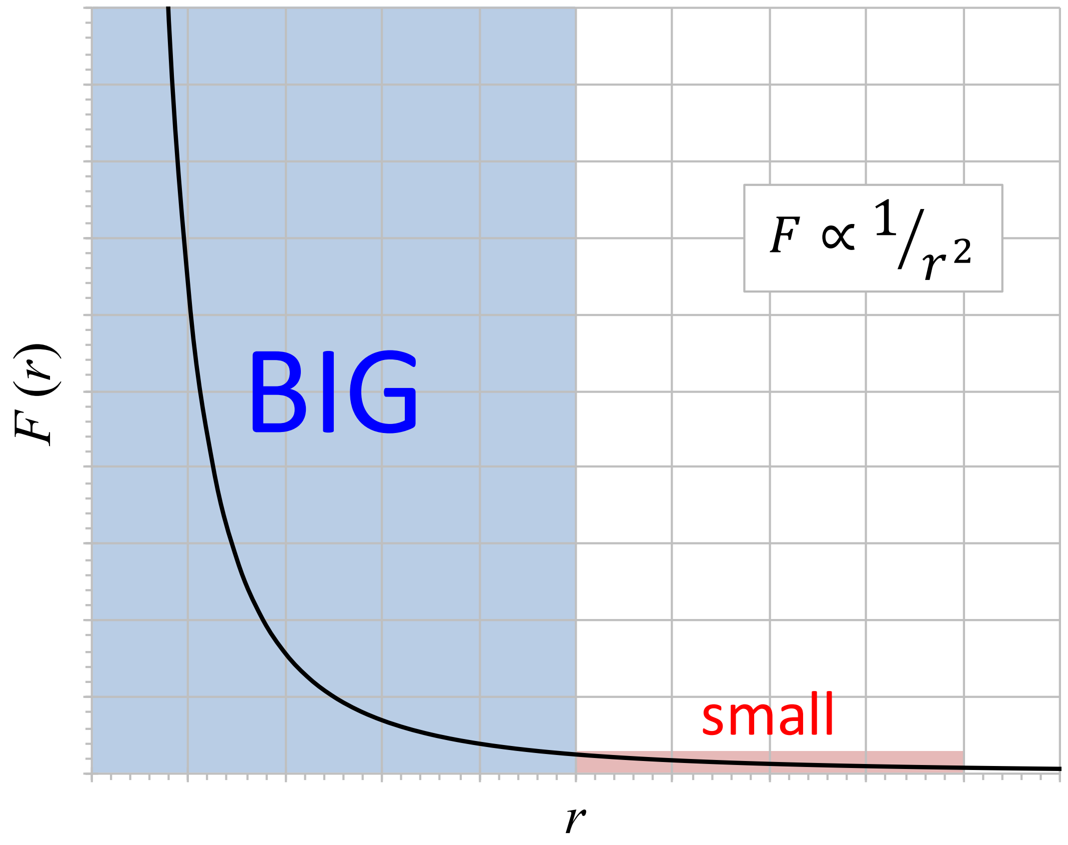

Efficient communication: The Edge Tile

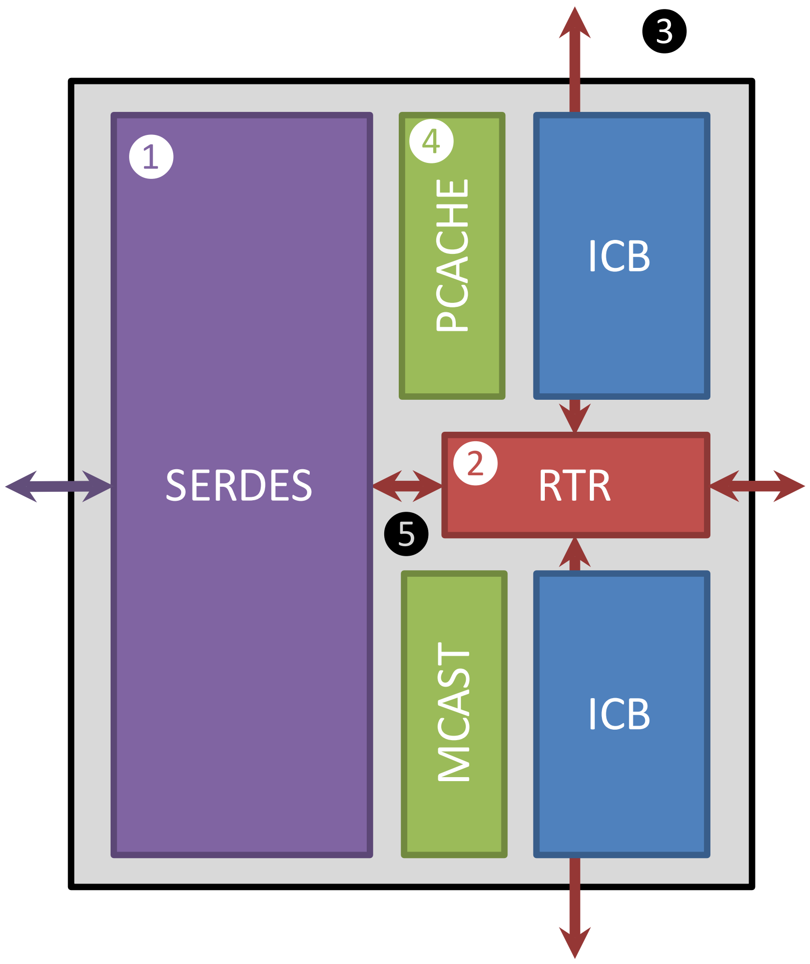

- Evolutionary changes
  - ① Increase SERDES data rate
  - ② Reduce hop latency
- **Revolutionary changes**
  - ③ Separate edge network
  - ④ MD-specific compression
  - ⑤ Novel interaction method

Laying tiles

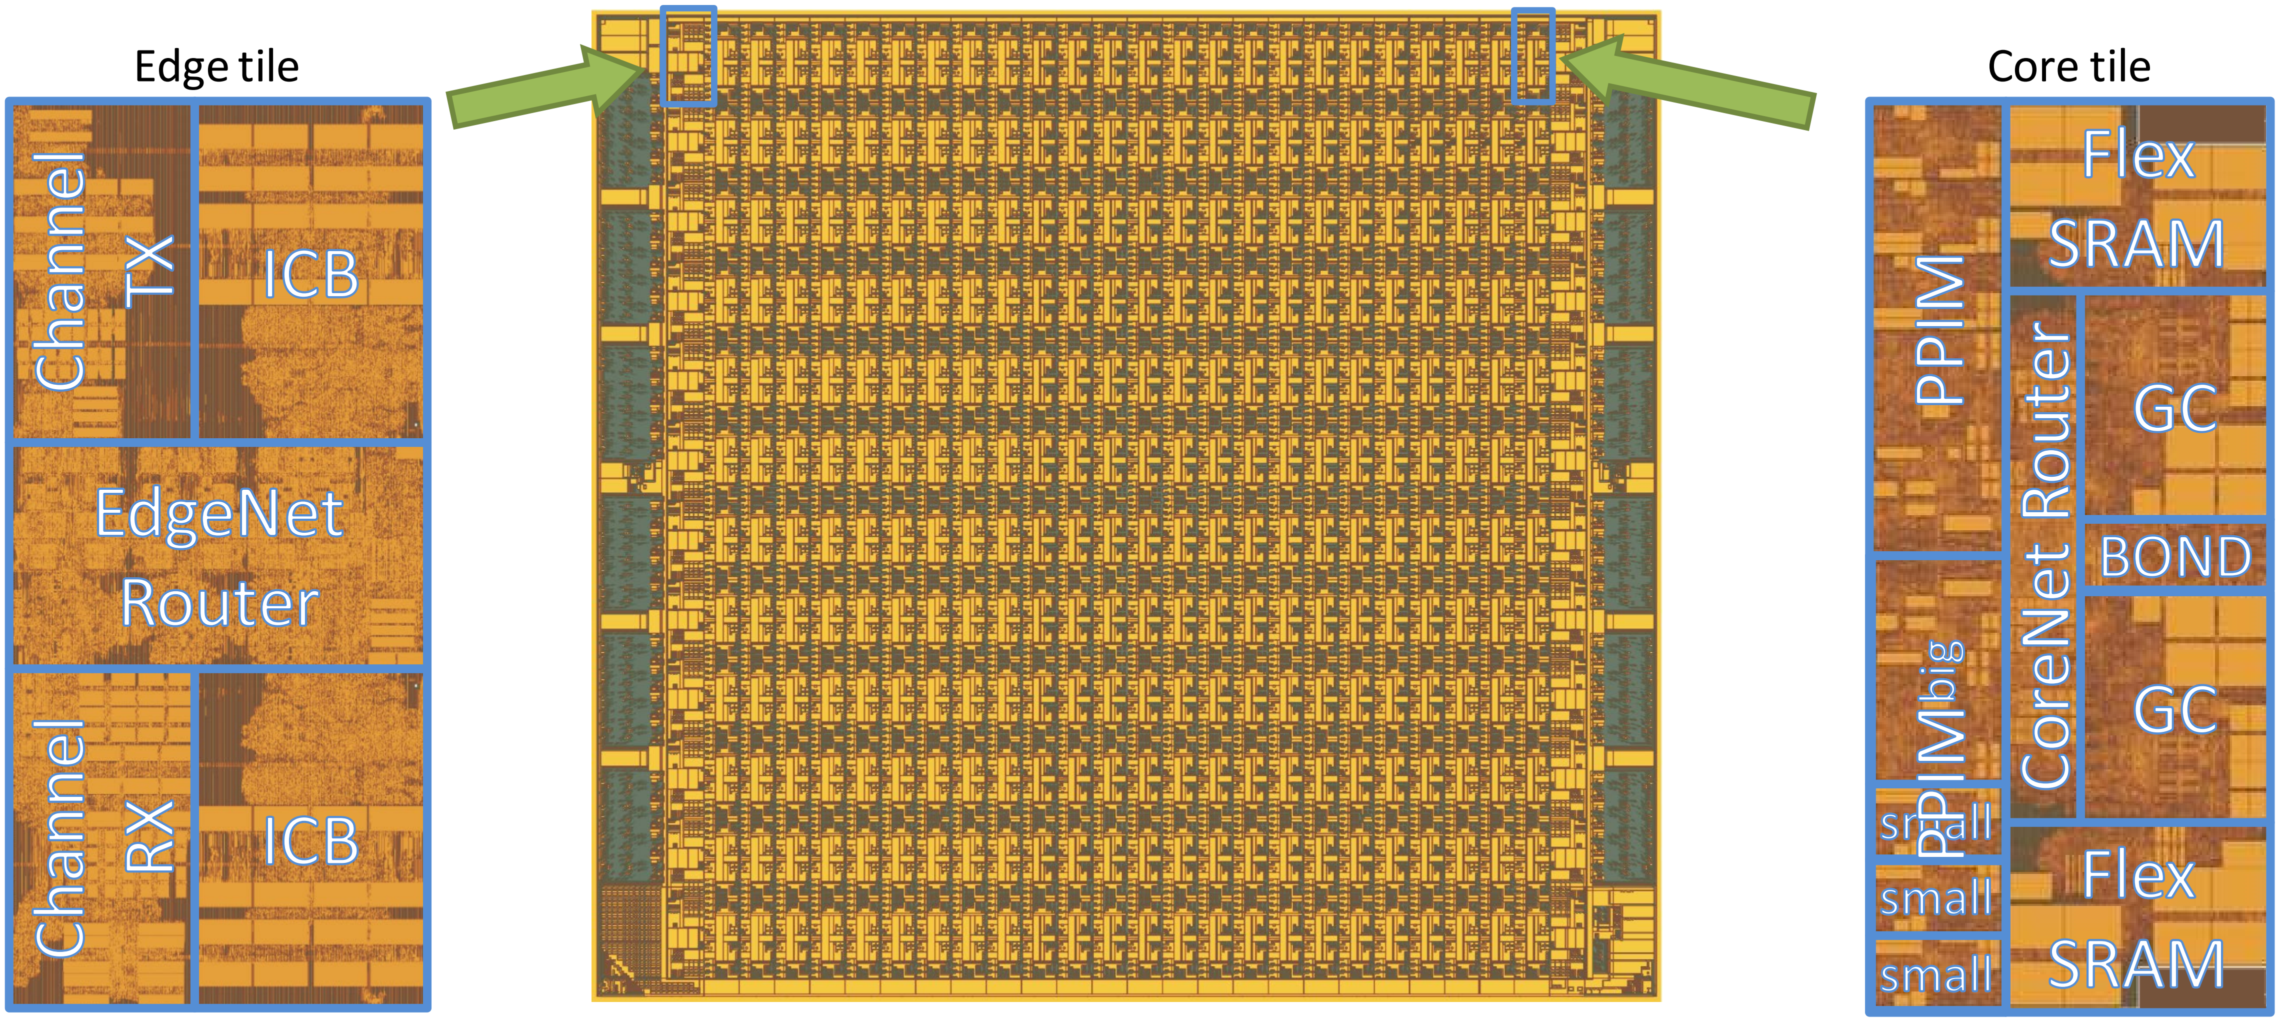

Physical design

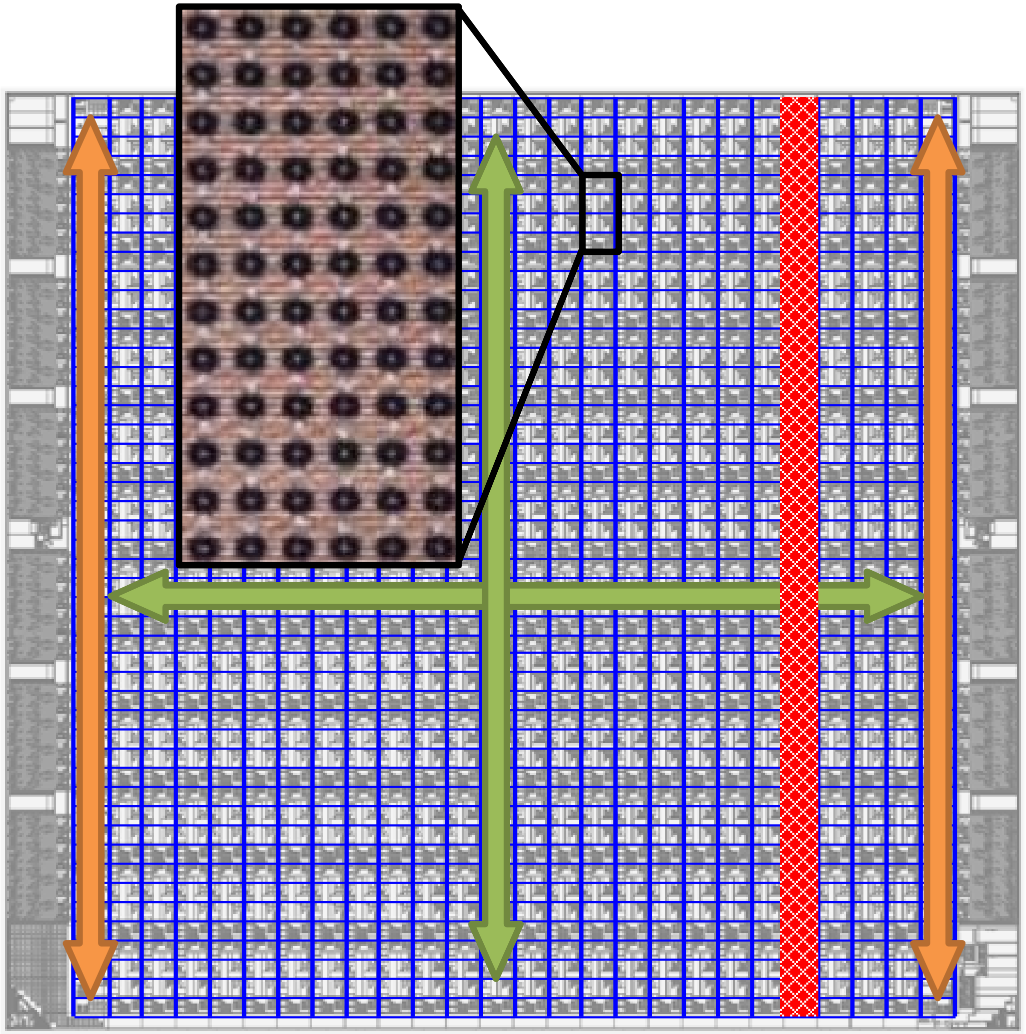

- Channel-less, abutted layout
- Few unique blocks
- Global, low-skew clock mesh 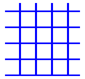
- Engineered global routing 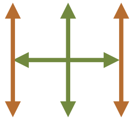
- Column-level redundancy 
- Robust power delivery 

The evolution of ANTON

	ANTON	ANTON 2	ANTON 3
Tape-out	2007	2012	2020
CPU cores	8+4+1	66	528*
PPIMs	32	76	528*
Flex SRAM	0.125 MiB	4 MiB	66 MiB*
Atoms / node	460	8,000	110,000*
Clock frequency	0.485/0.970 GHz	1.65 GHz	2.8+ GHz
Channel bandwidth	0.607 Tbps	2.7 Tbps	5.6+ Tbps
Process node	90 nm	40 nm	7 nm
Transistors	0.2 G	2.0 G	31.8 G
Die size	299 mm^2	410 mm^2	451 mm^2
Power	30 W	190 W	360 W

* 22/24 columns

Baby pictures

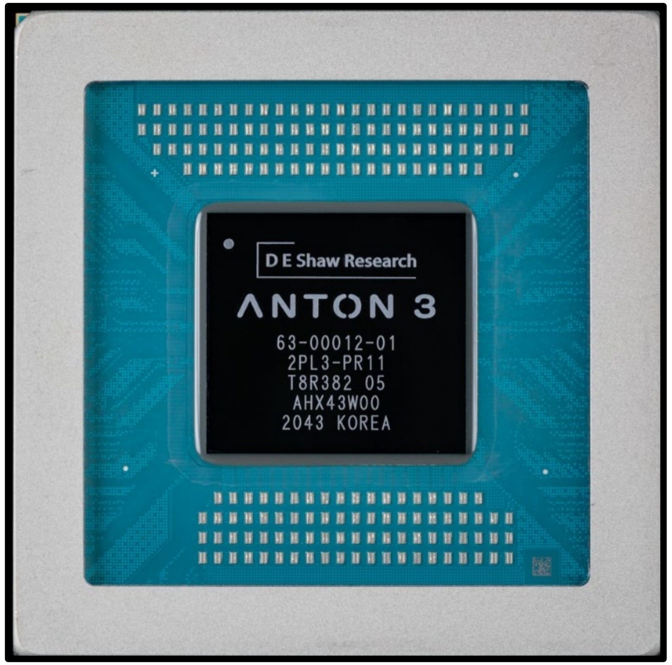

29 September 2020: chips arrive
MD running (water) < 9 h later

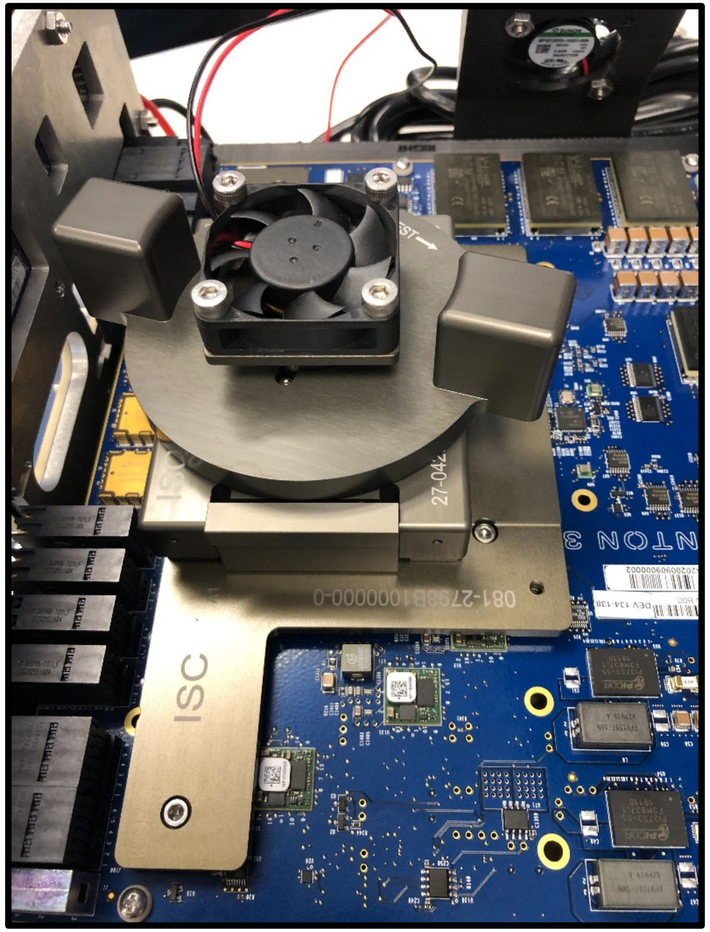

30 September 2020: 1st protein run
Faster @ 250 MHz than Anton 2

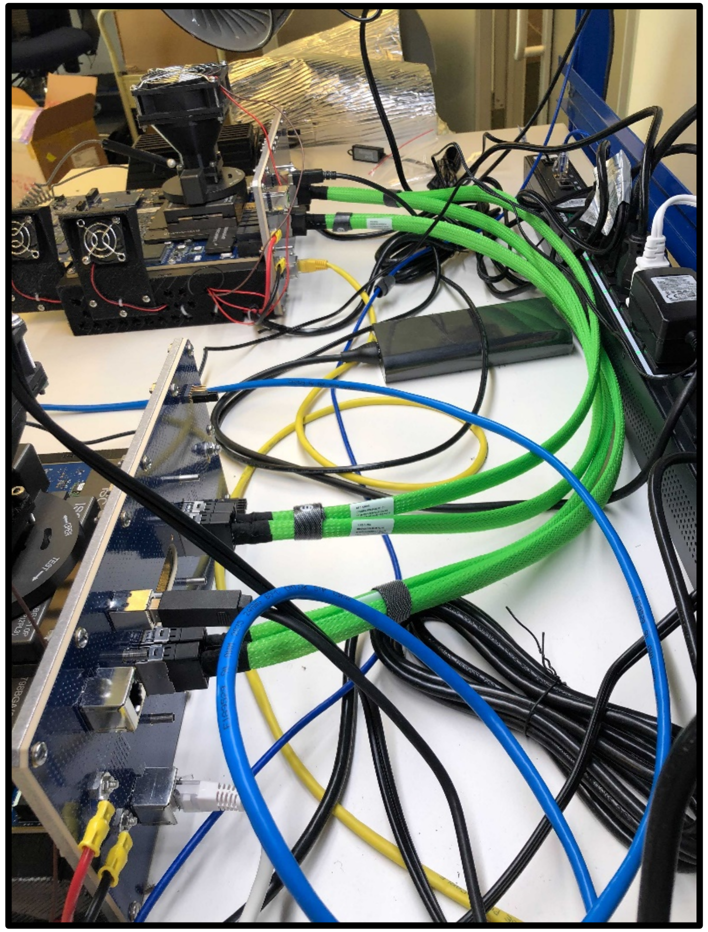

31 October 2020: Multi-node

Node board

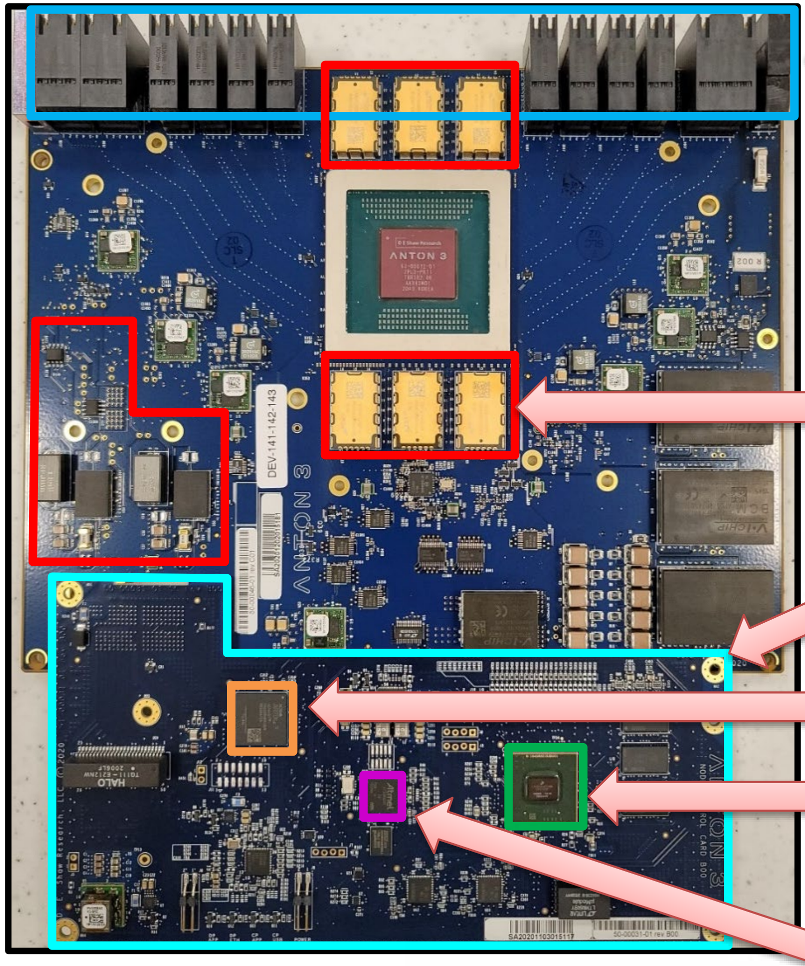

48 VDC, torus links, Ethernet, USB

500+ W custom voltage regulator

Node control complex

ASIC interface FPGA

Data processor (64-bit Linux)

Control processor (no OS)

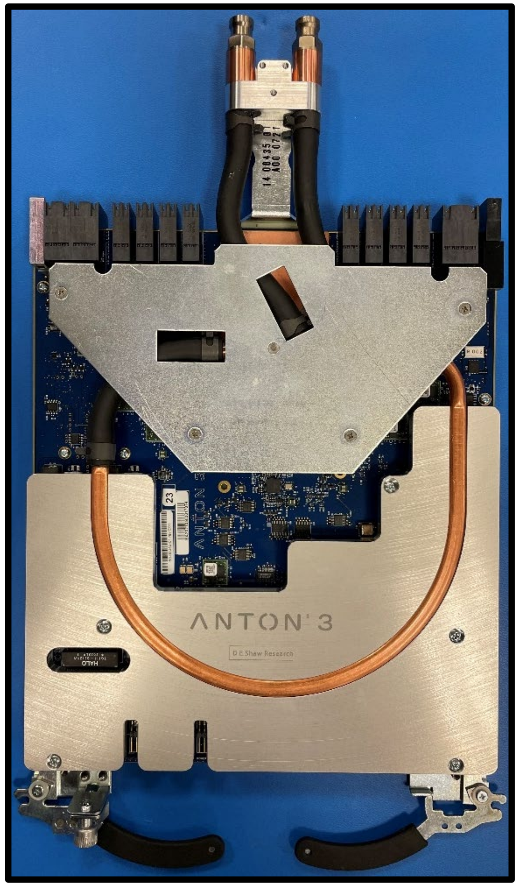

Scale up

8×8 nodes

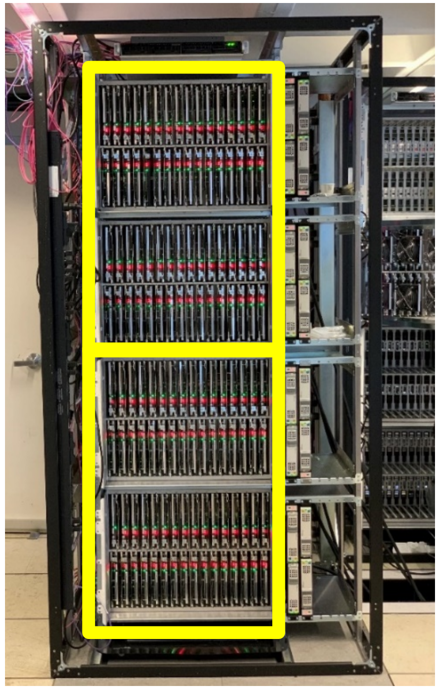

2×64 nodes

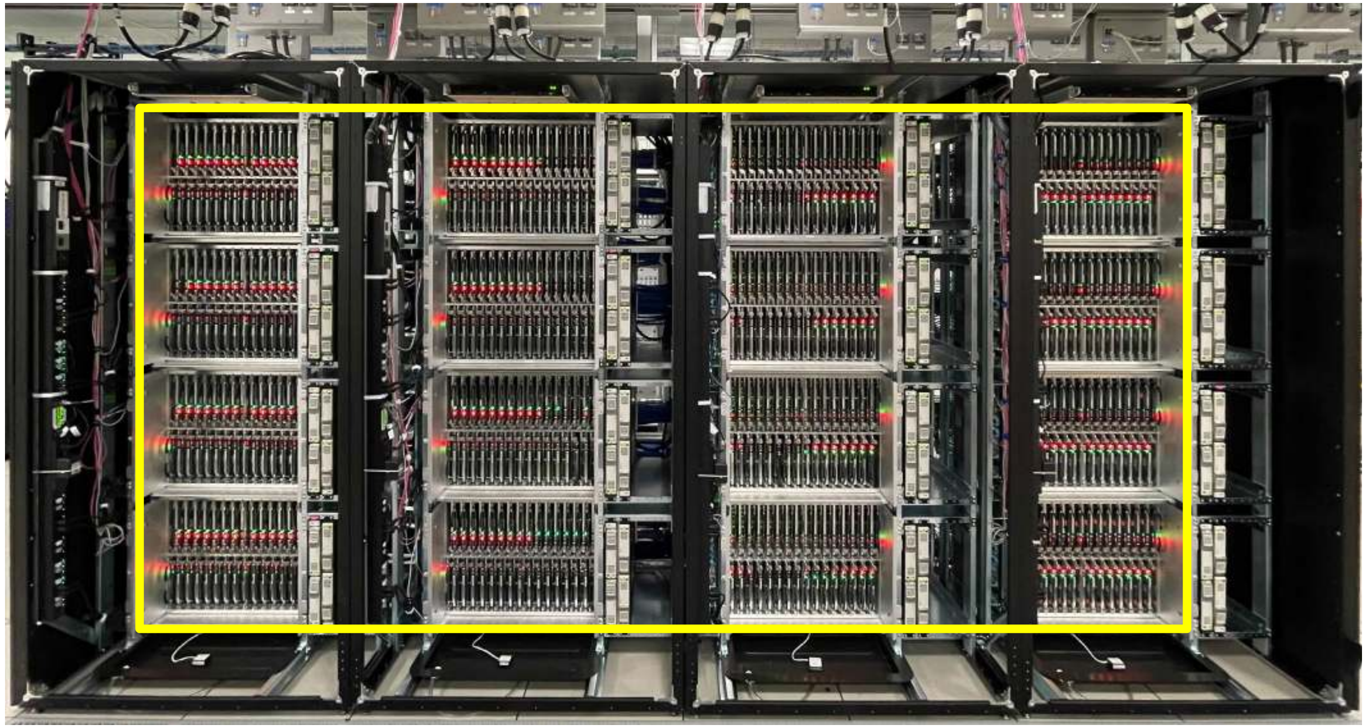

512 nodes

Network

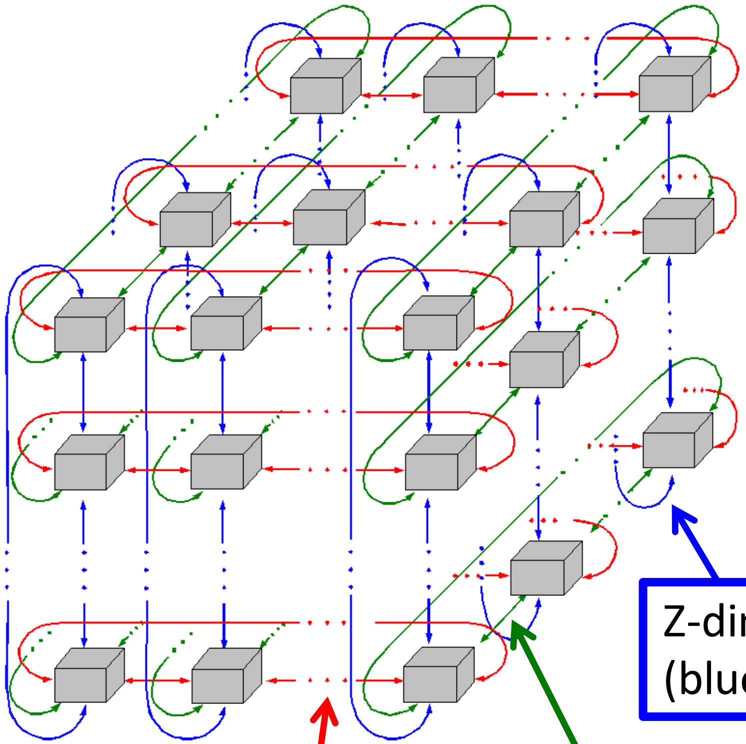

X-dimension connected entirely in backplane

Y-dimension split across cables (green) and backplane

Z-dimension cabled (blue and yellow)

Complete 512-node, 3D torus

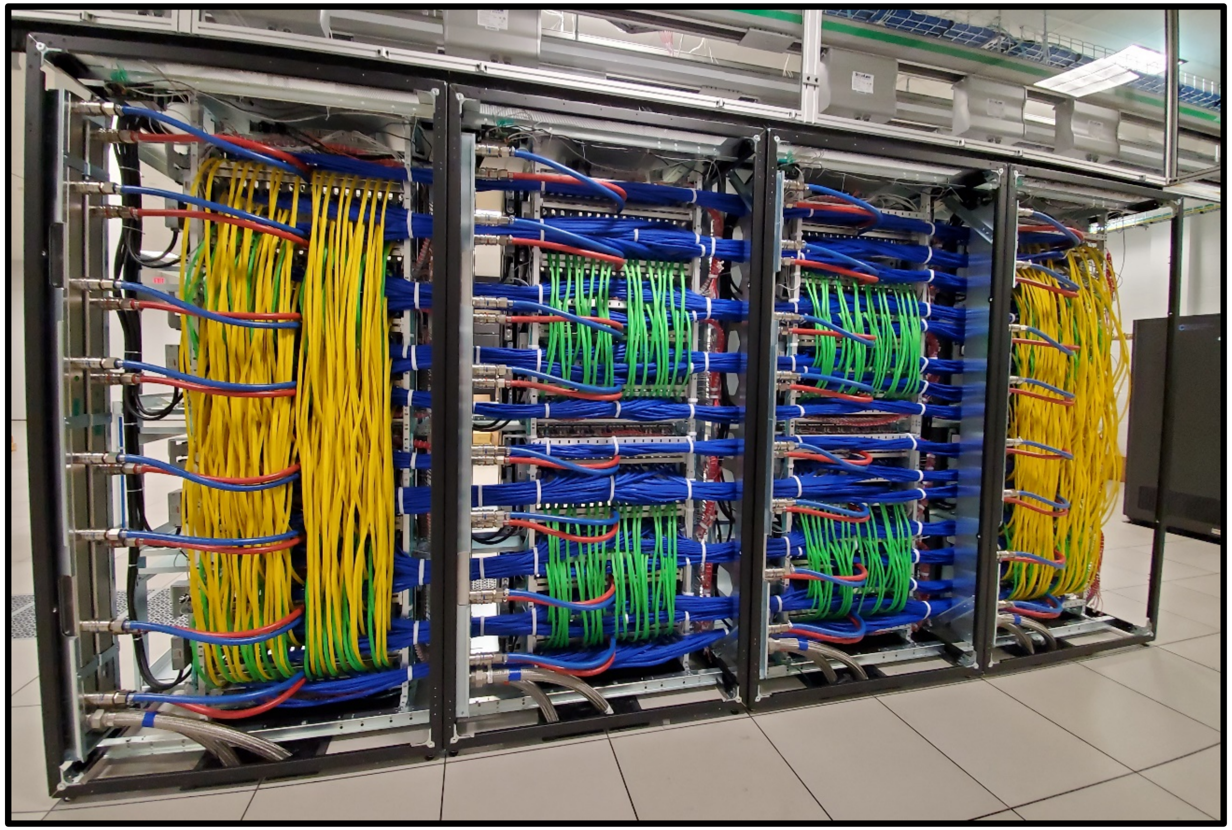

Taming (cooling) the beast

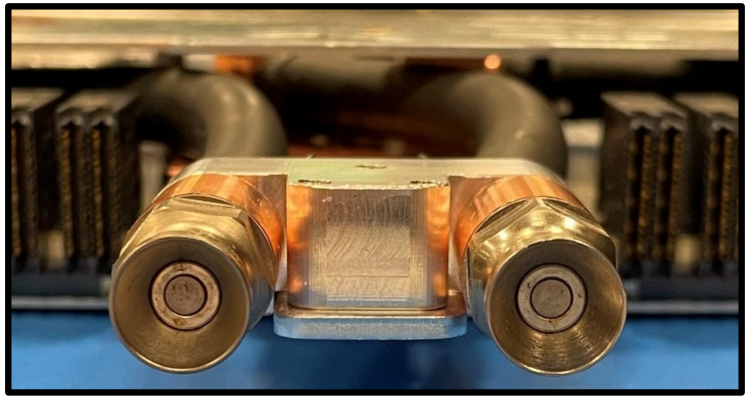

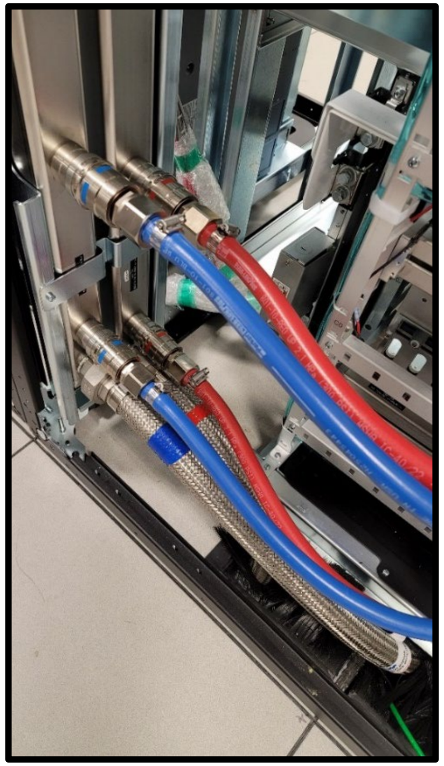

$T_J < 65$ °C @ 500 W

MD performance

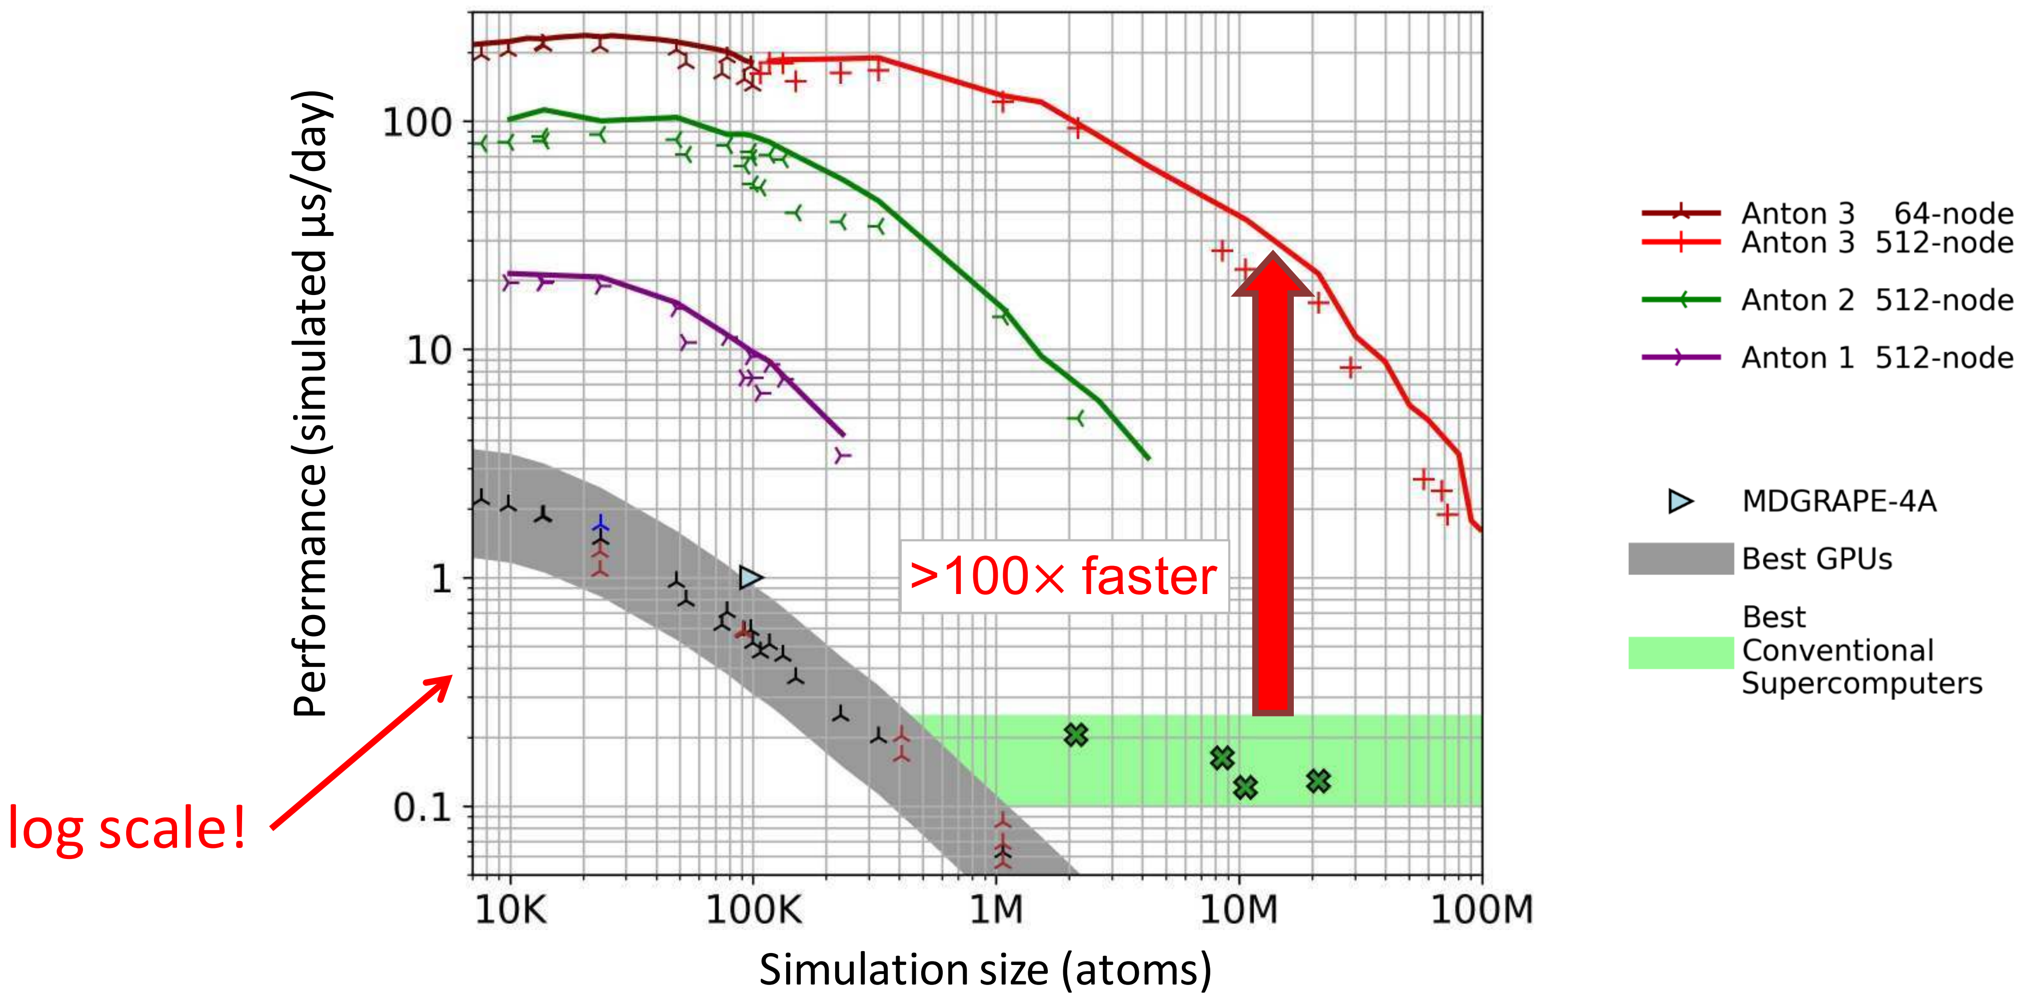

D E Shaw Research

Acknowledgements

- System software group for machine bring-up
- Embedded software group for creating and tuning the application
  - Ken Mackenzie for performance results and figures
- Systems group for support and infrastructure
  - And lots of photos!
- Chemistry team for putting Anton to good use
  - Kevin Yuh for MD simulation videos

Performance references

- GPU performance results
  - M. Bergdorf et al, “Desmond/GPU Performance as of April 2021”, DESRES/TR--2021-01, [Online: April 2021]. https://deshawresearch.com/publications.html.
  - "A100, V100 and Multi-GPU Benchmarks", [Online: January 2020]. https://github.com/openmm/openmm/issues/2971.
  - "NVIDIA HPC Application Performance", [Online: July 2021]. https://developer.nvidia.com/hpc-application-performance.
  - "Gromacs/NAMD Multi-GPU Scaling", unpublished internal benchmarking.
- Supercomputer performance results
  - “RIKEN Team Use Supercomputer to Explore the Molecular Dynamics of the New Coronavirus,” HPCwire announcement, [Online: March 2020]. https://www.hpcwire.com/off-the-wire/riken-team-use-supercomputer-to-explore-the-molecular-dynamics-of-the-new-coronavirus/
  - S. Páll et al., “Tackling Exascale Software Challenges in Molecular Dynamics Simulations with GROMACS,” in: S. Markidis, E. Laure (eds) Solving Software Challenges for Exascale. EASC 2014. Lecture Notes in Computer Science, vol. 8759, 2015.
  - NAMD scaling on Summit, [Online: May 2018]. http://www.ks.uiuc.edu/Research/namd/benchmarks/
  - L. Casalino et al., “AI-Driven Multiscale Simulations Illuminate Mechanisms of SARS-CoV-2 Spike Dynamics,” *International Journal of High Performance Computing Applications*, 2021.
  - J. R. Perilla and K. Schulten, “Physical Properties of the HIV-1 Capsid from All-Atom Molecular Dynamics Simulations,” *Nature Communications*, vol. 8 (15959), 2017.
- Anton performance results (original publications; improved performance used in comparisons)
  - D. E. Shaw et al., “Millisecond-Scale Molecular Dynamics Simulations on Anton,” in *SC‘09: Proceedings of the Conference on High Performance Computing Networking, Storage, and Analysis*, 2009, 1–11.
  - D. E. Shaw et al., “Anton 2: Raising the Bar for Performance and Programmability in a Special-Purpose Molecular Dynamics Supercomputer,” in *SC’14: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*, 2014, 41–53.
  - D. E. Shaw et al., “Anton 3: Twenty Microseconds of Molecular Dynamics Simulation Before Lunch,” to appear in *SC’21: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*, 2021.