Drinking from a firehose: Solving data analysis challenges posed by the Anton supercomputer

Mark Moraes

Biomolecular Timescales (seconds)

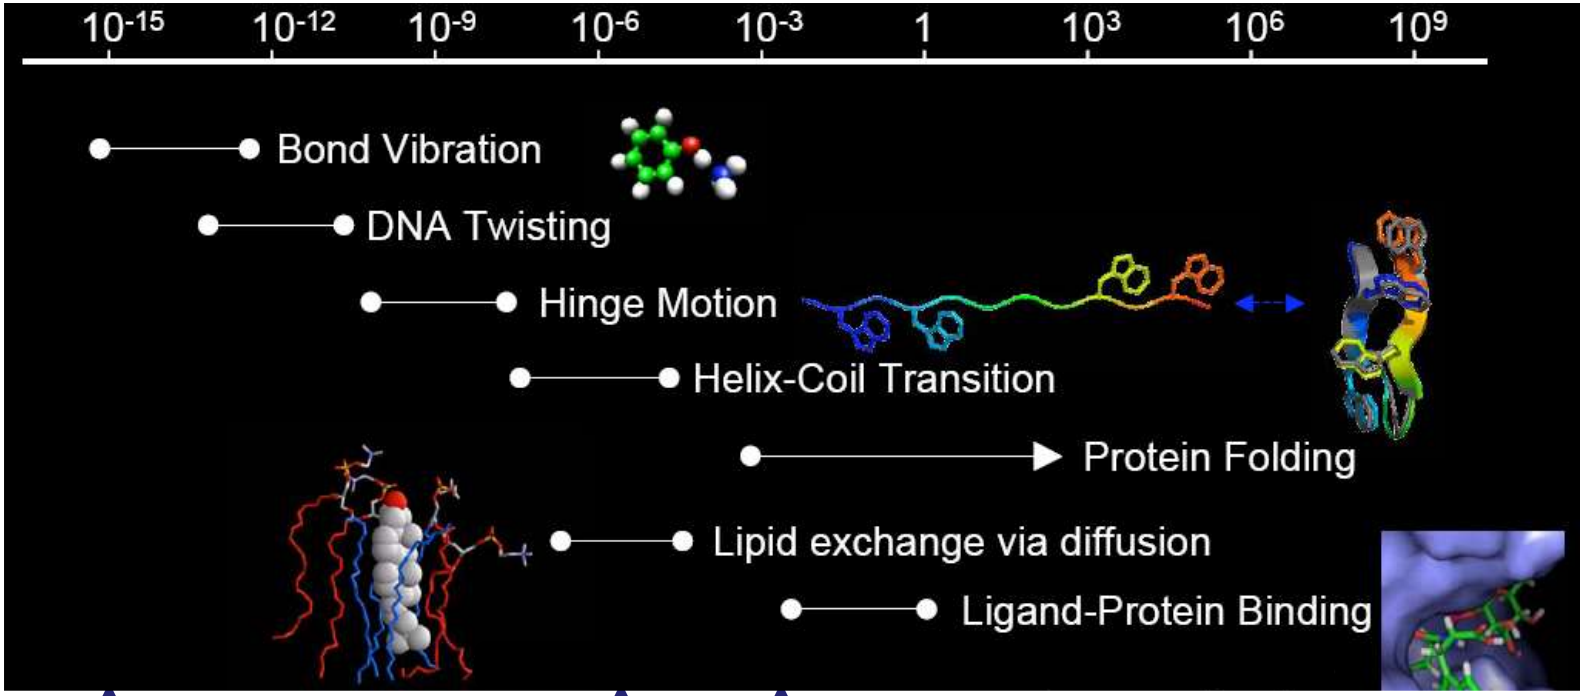

Adapted from Suits (IBM), originally from Chan & Dill (1993)

fs – fastest motions of hydrogen bond (one MD step = 2 fs)

µs – basic biochemical sub-steps (a half-billion steps)

ms - this is where biological processes start to get interesting (a half-trillion steps: years of simulation of a modern general-purpose supercomputer or cluster)

Why Molecular Dynamics (MD) is hard…

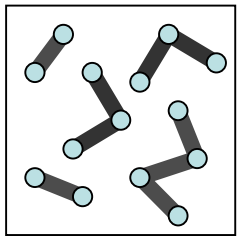

Bonded

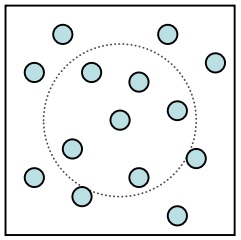

Non Bonded:
Van der Waals, Near Electrostatics

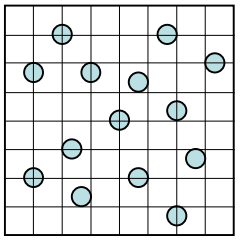

Non Bonded:
Far Electrostatics

$$E = \sum_{\text{bonds}} k_b \left(r - r_0\right)^2$$

$$+ \sum_{\text{angles}} k_\theta (\theta - \theta_0)^2$$

$$+ \sum_{\text{torsions}} A[1 + \cos(n\tau - \varphi)]$$

$$+ \sum_i \sum_{j>i} \frac{A_{ij}}{r_{ij}^{12}} - \frac{B_{ij}}{r_{ij}^{6}}$$

$$+ \sum_i \sum_{j>i} \frac{q_i q_j}{r_{ij}}$$

Time required for one-millisecond simulation of a small protein

Machine	Time required
General-purpose supercomputer	2123 days (5.8 years)
Top-end 2019 GPU	806 days (2.2 years)
Anton 1 (2008)	55 days
Anton 2 (2013)	12 days

Joint AMBER-CHARMM benchmark of dihydrofolate reductase DHFR (23,558 atoms, 62 Å cube, 2.5×10^{-15} s time step)

512-node Anton Performance

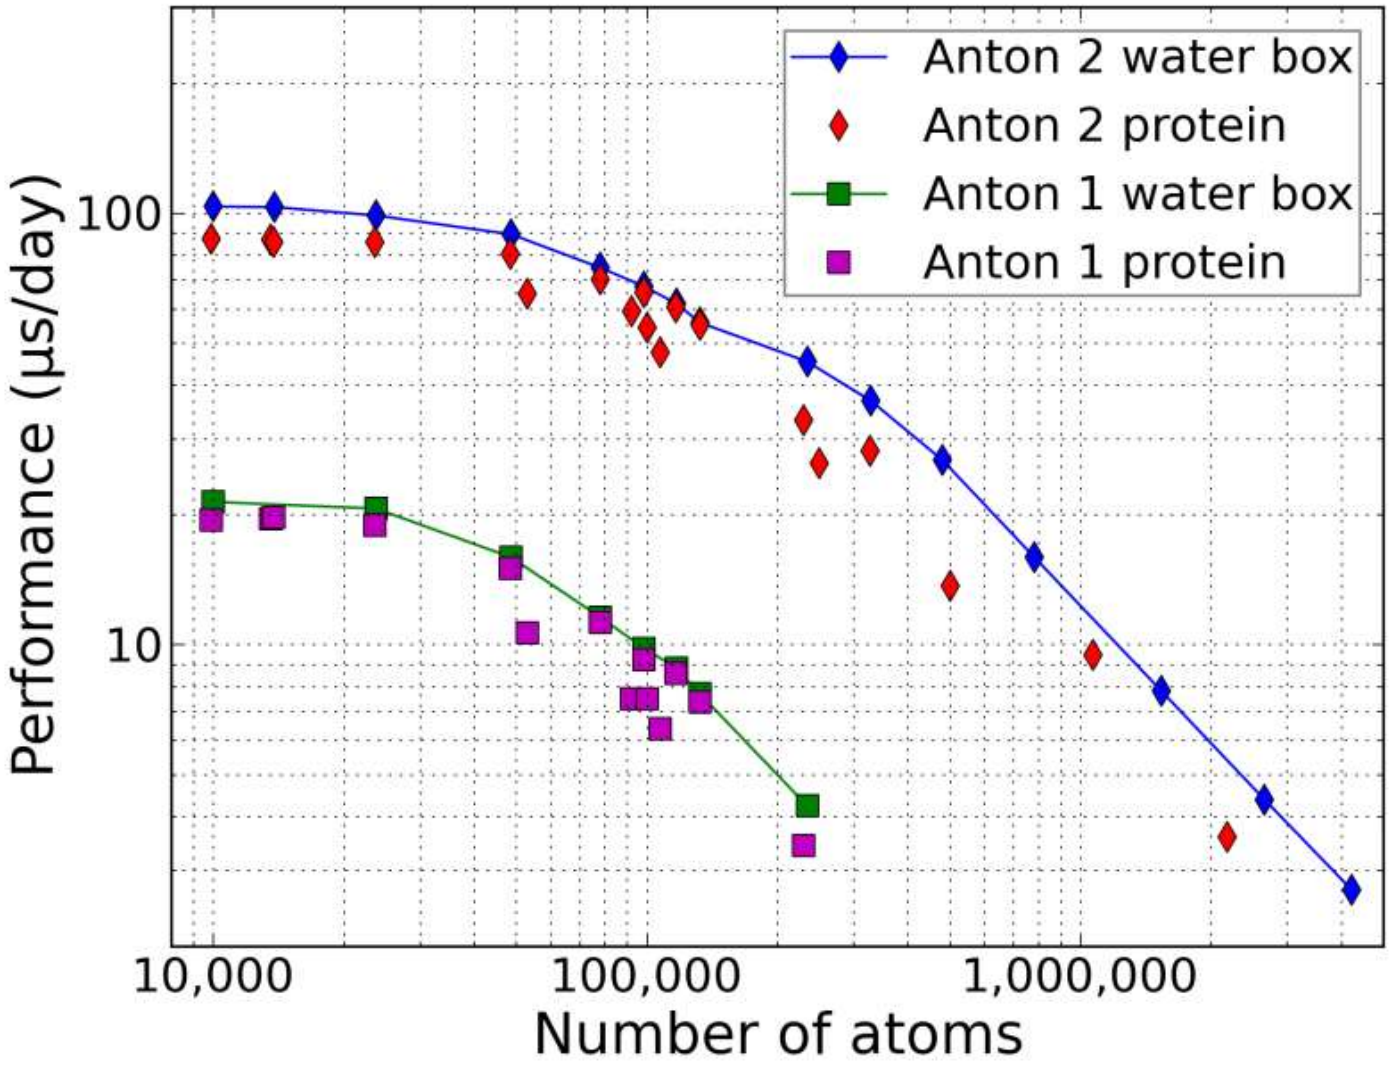

Our Answer: Anton: A family of specialized supercomputers for MD simulation

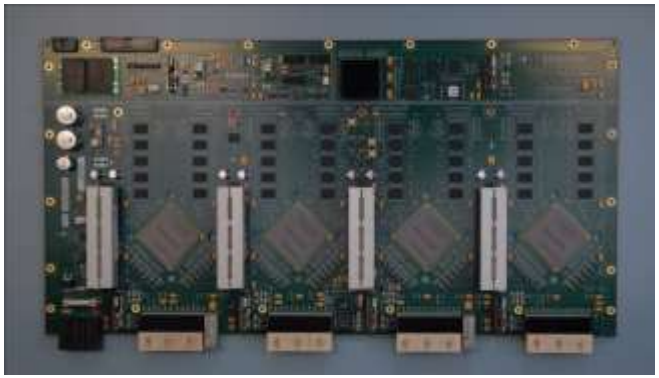

Roadmap

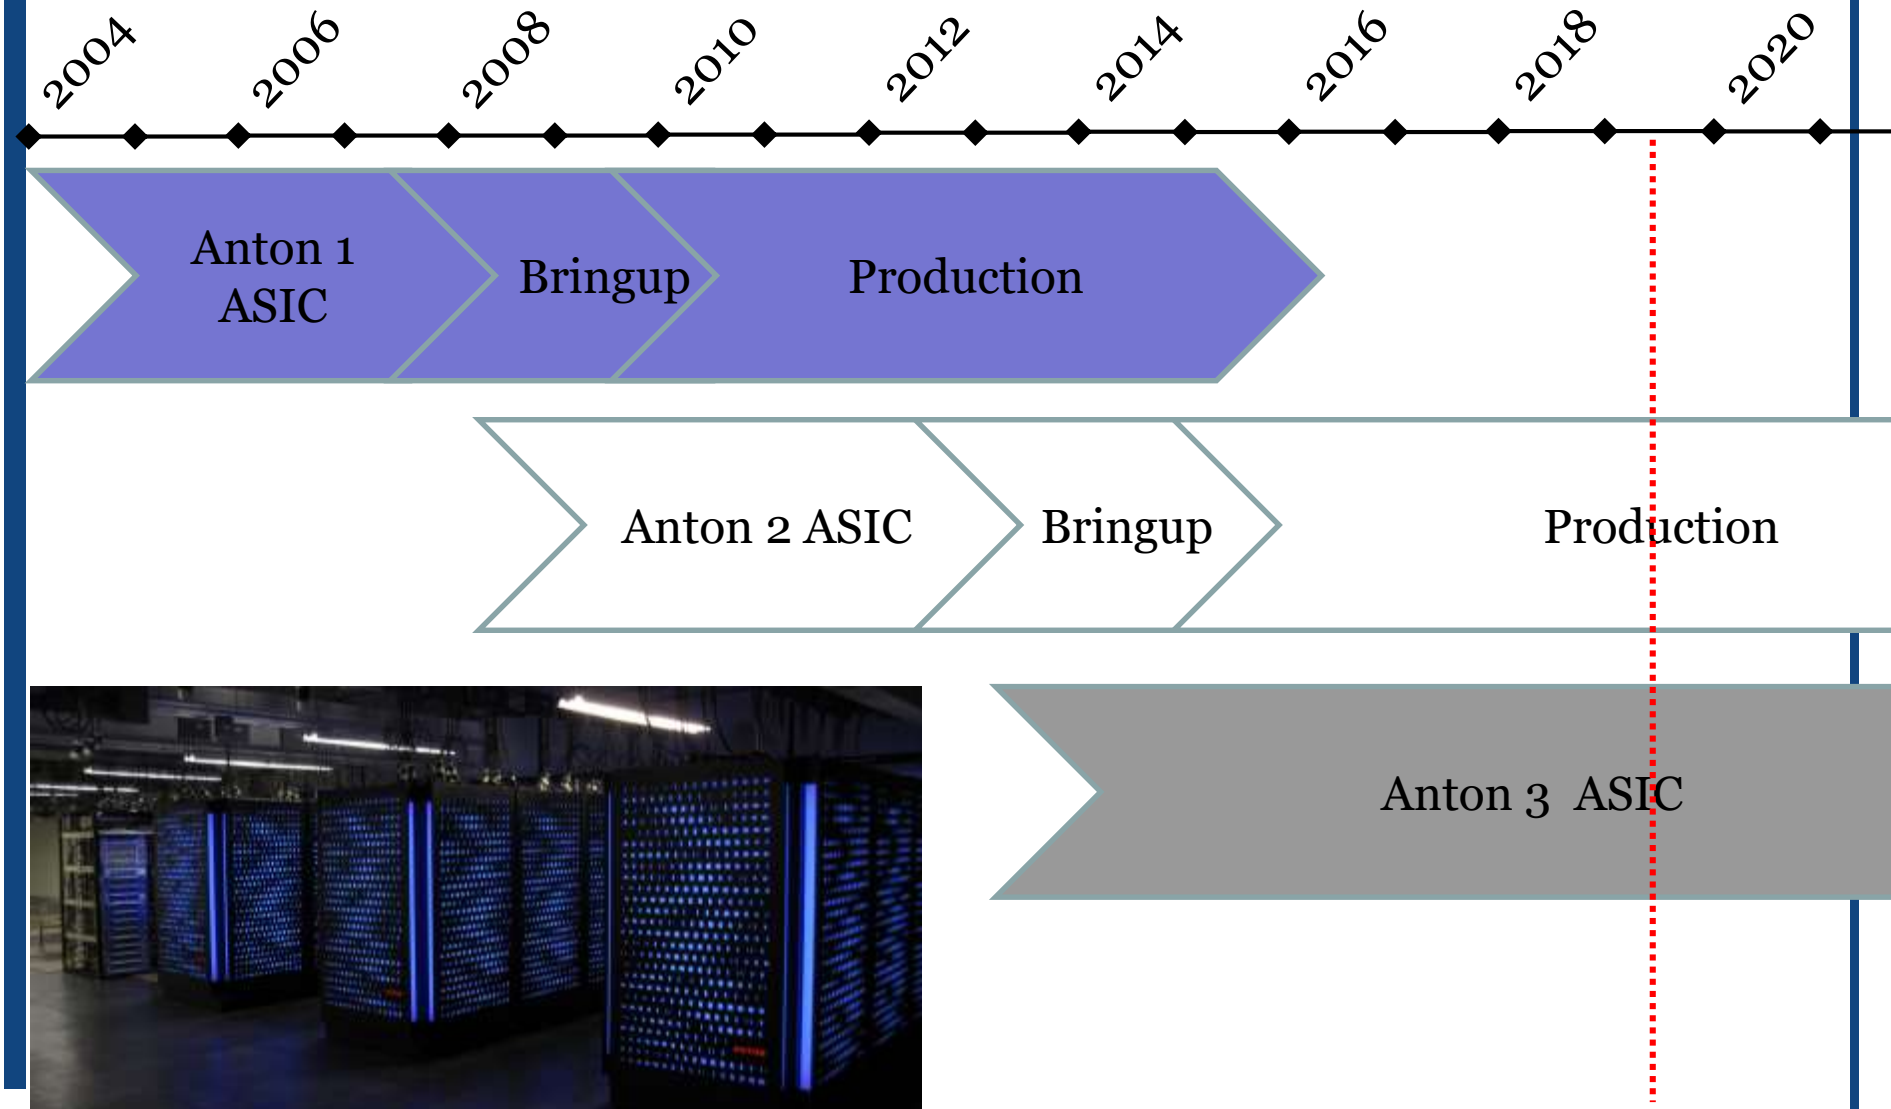
2004
2006
2008
2010
2012
2014
2016
2018
2020
Anton 1 ASIC
Bringup
Production
Anton 2 ASIC
Bringup
Production
Anton 3 ASIC

ANTON®

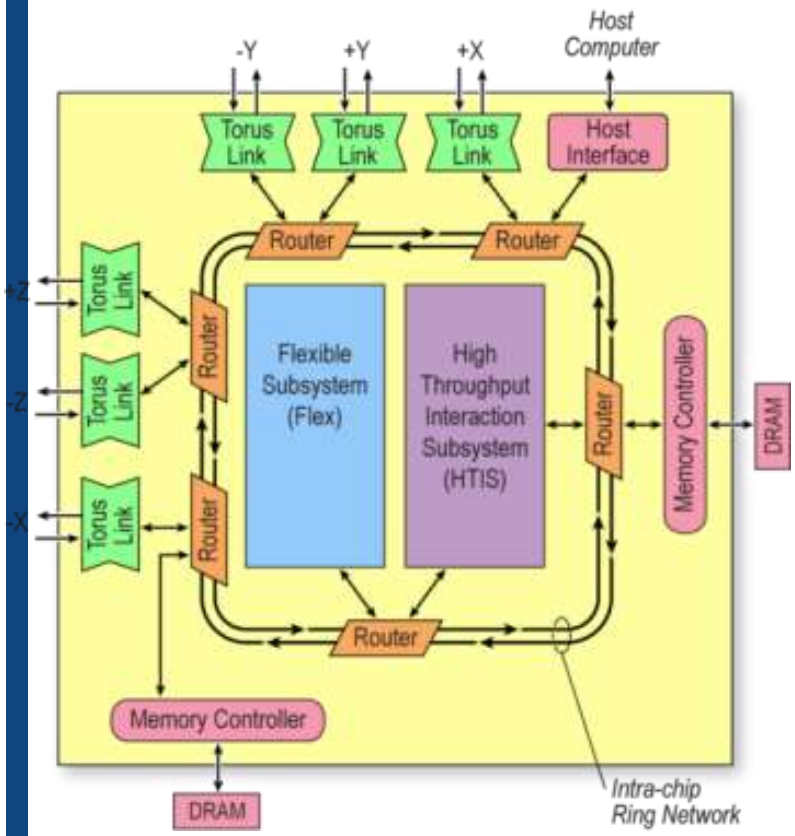

ANTON 2

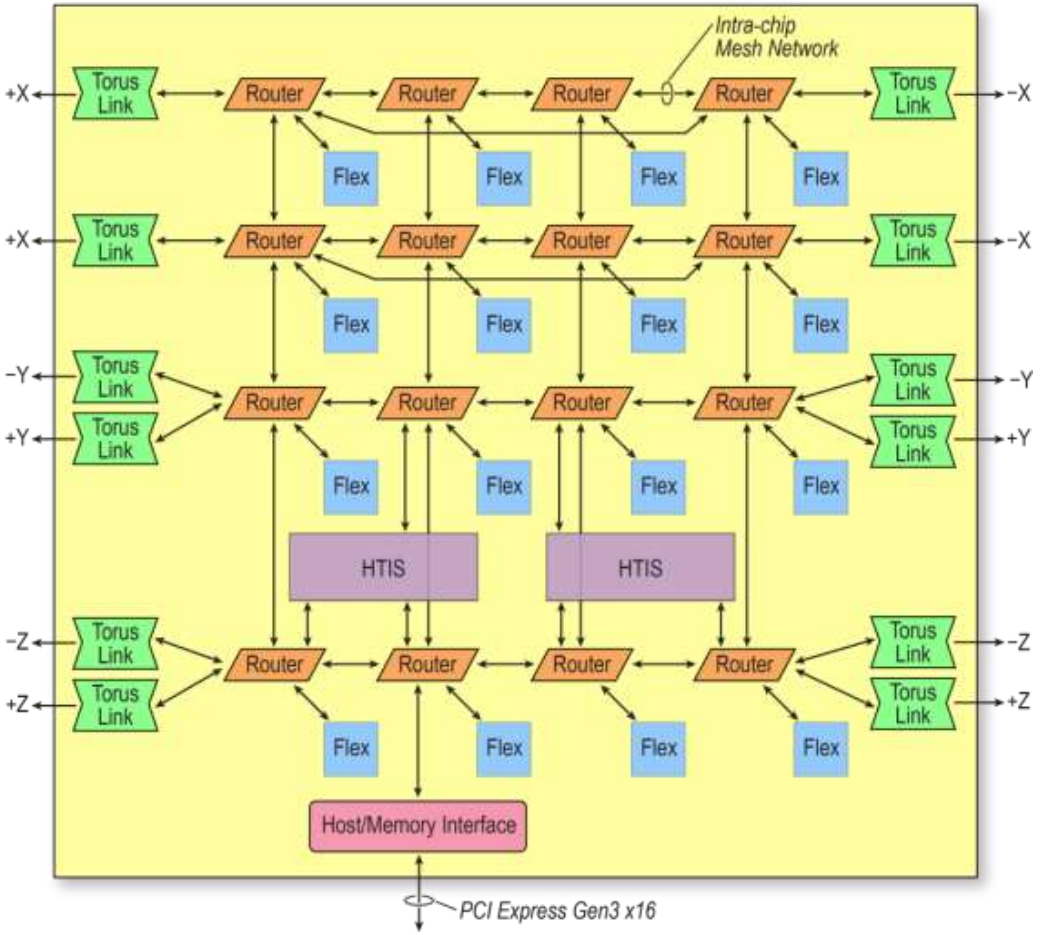

Anton 2 Architecture: 3D Torus of Directly Interconnected ASICs (nodes)

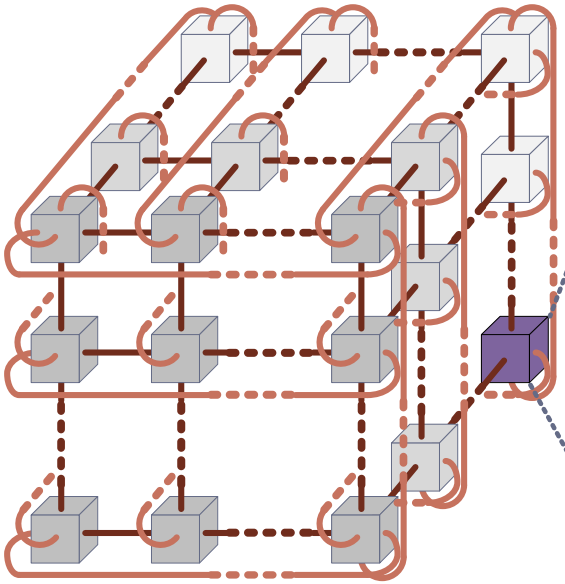

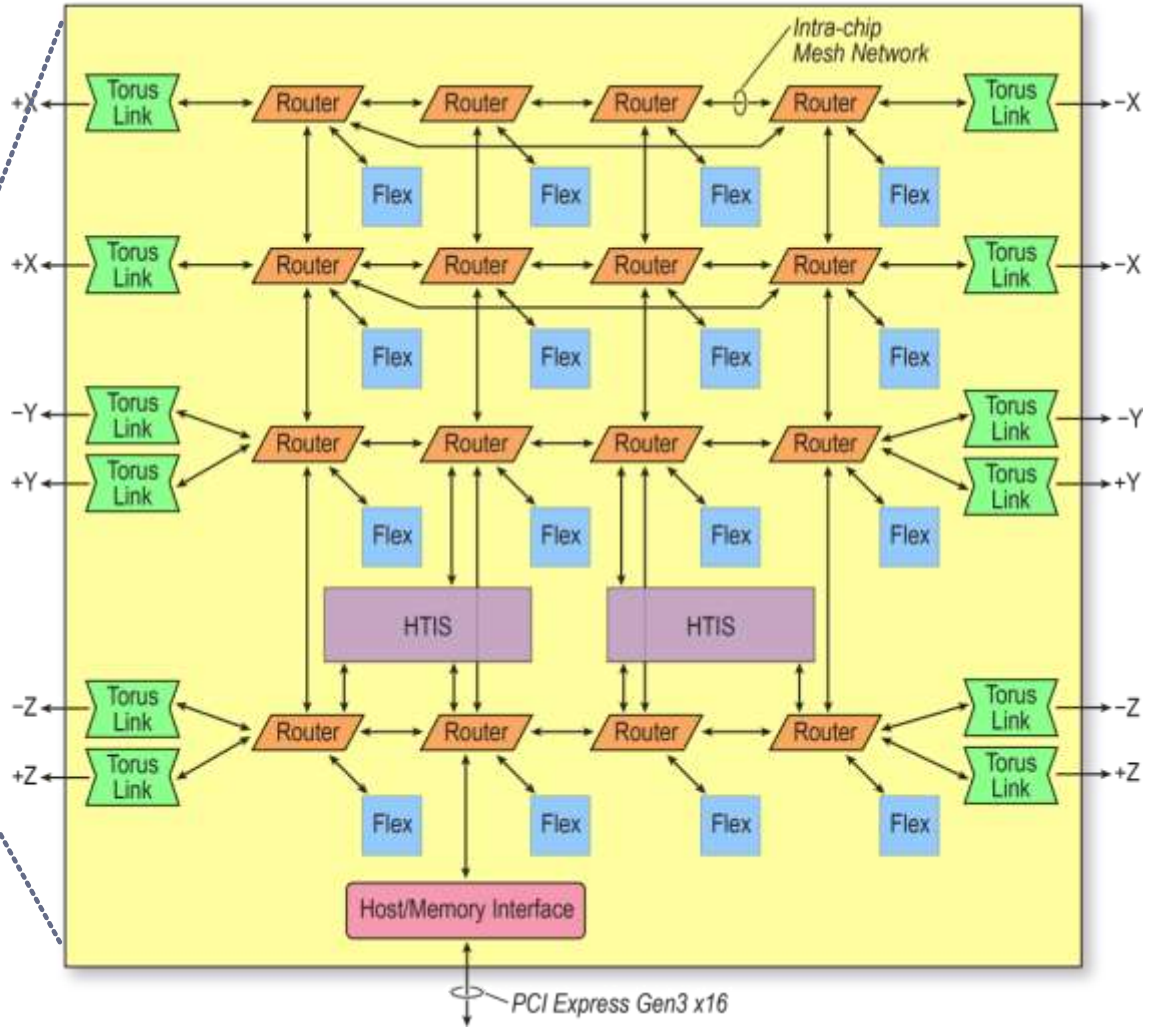

Flexible Subsystem Tile (Flex)

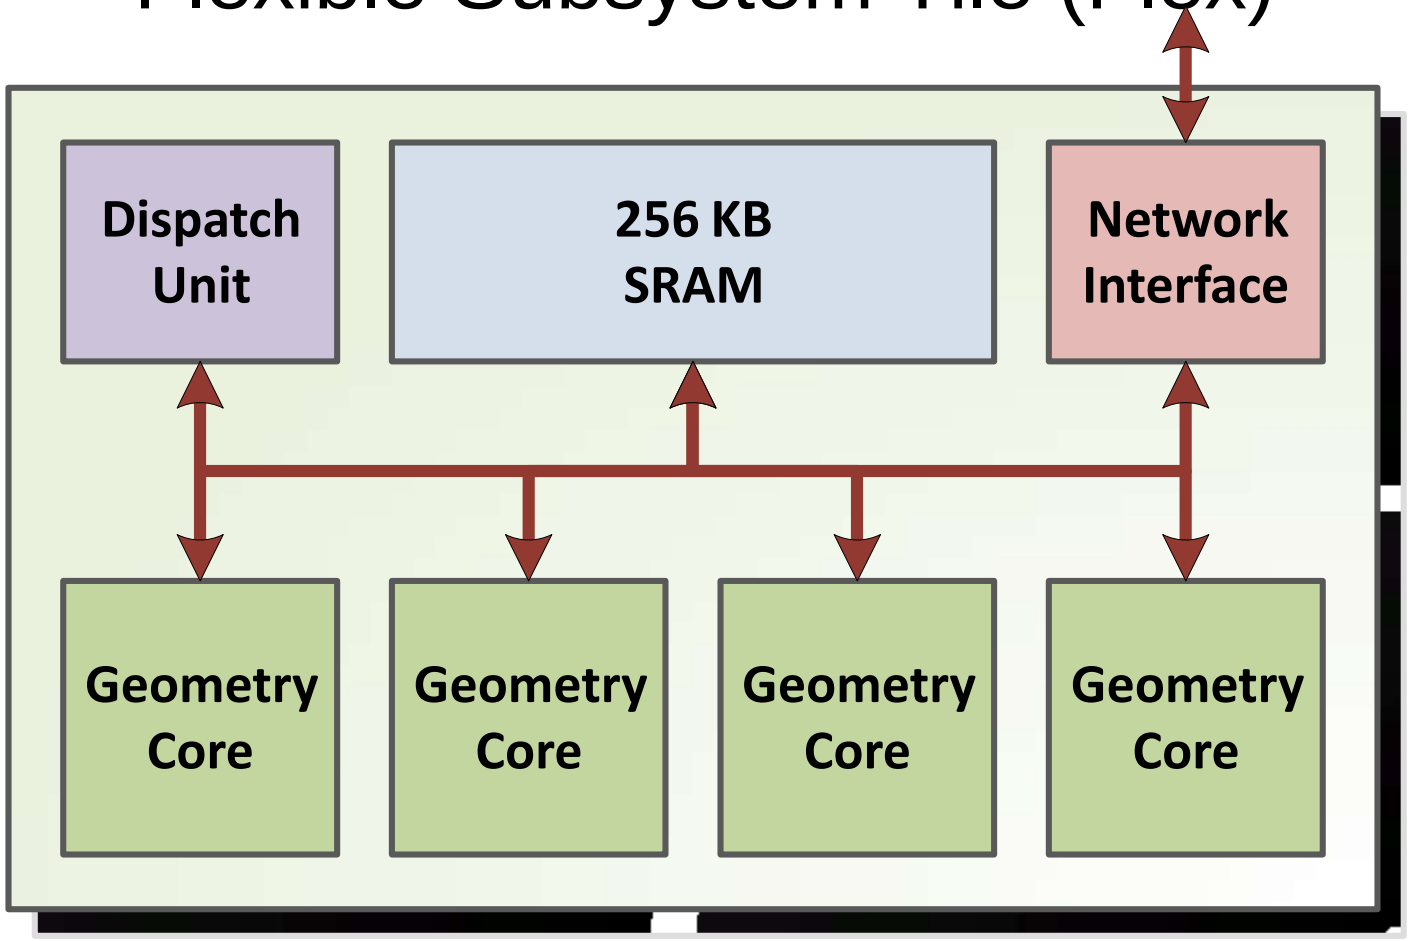

High-Throughput Interaction Subsystem (HTIS)

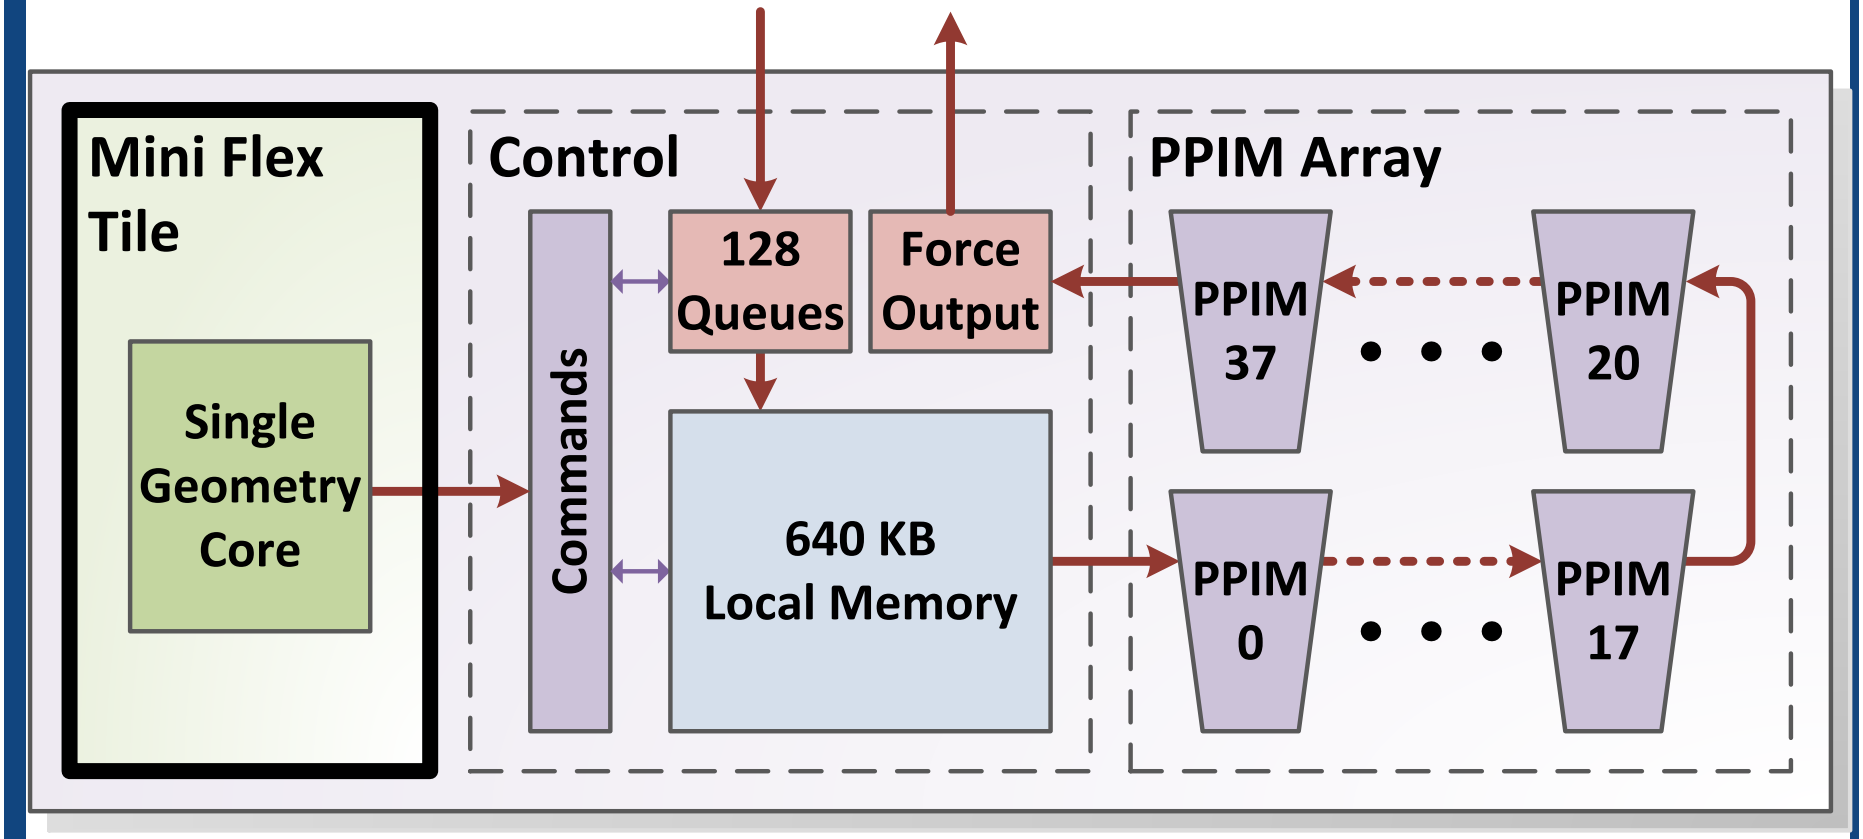

Cool and Connected

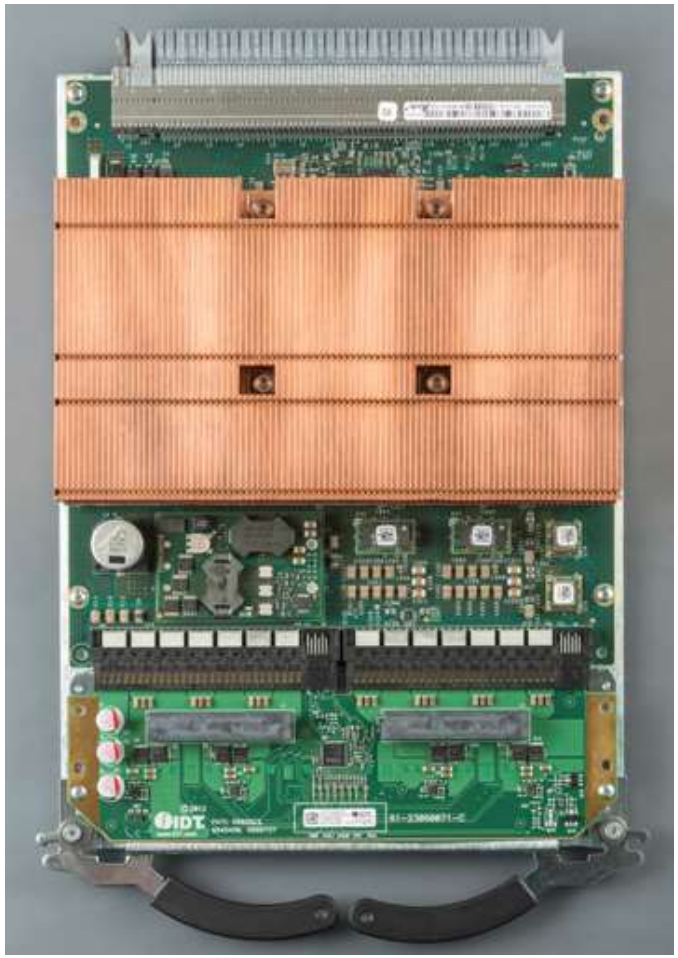

Board

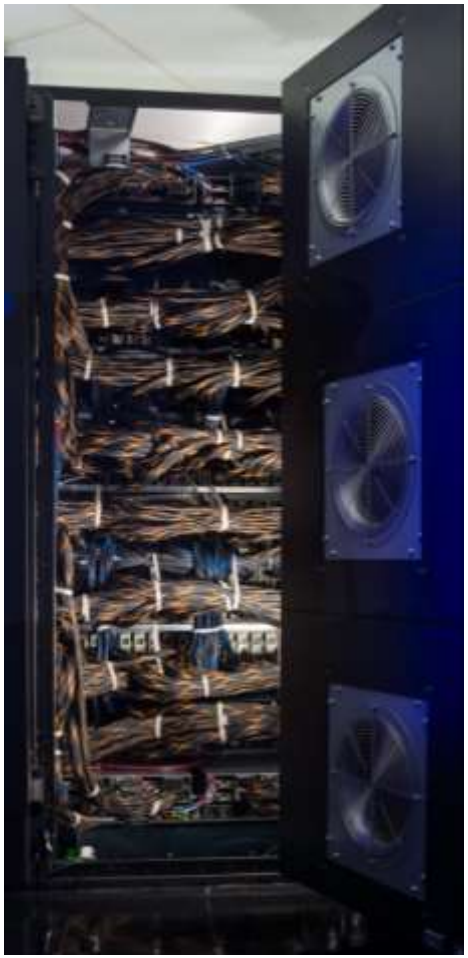

Rack

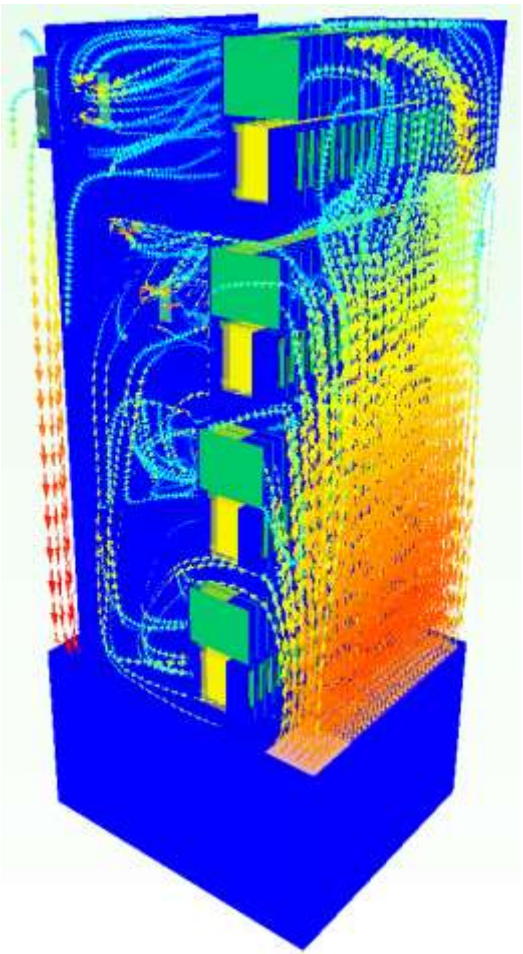

Air flow

Anton 1
2008
2016

Anton 2 : In production since 2014

Anton 2 produces ~100MB/s of output data

× 86,400 seconds/day
× 365 days/year….

= **3.2 PB/year**

We archive all of it.

Now what…

Archive and Cache!

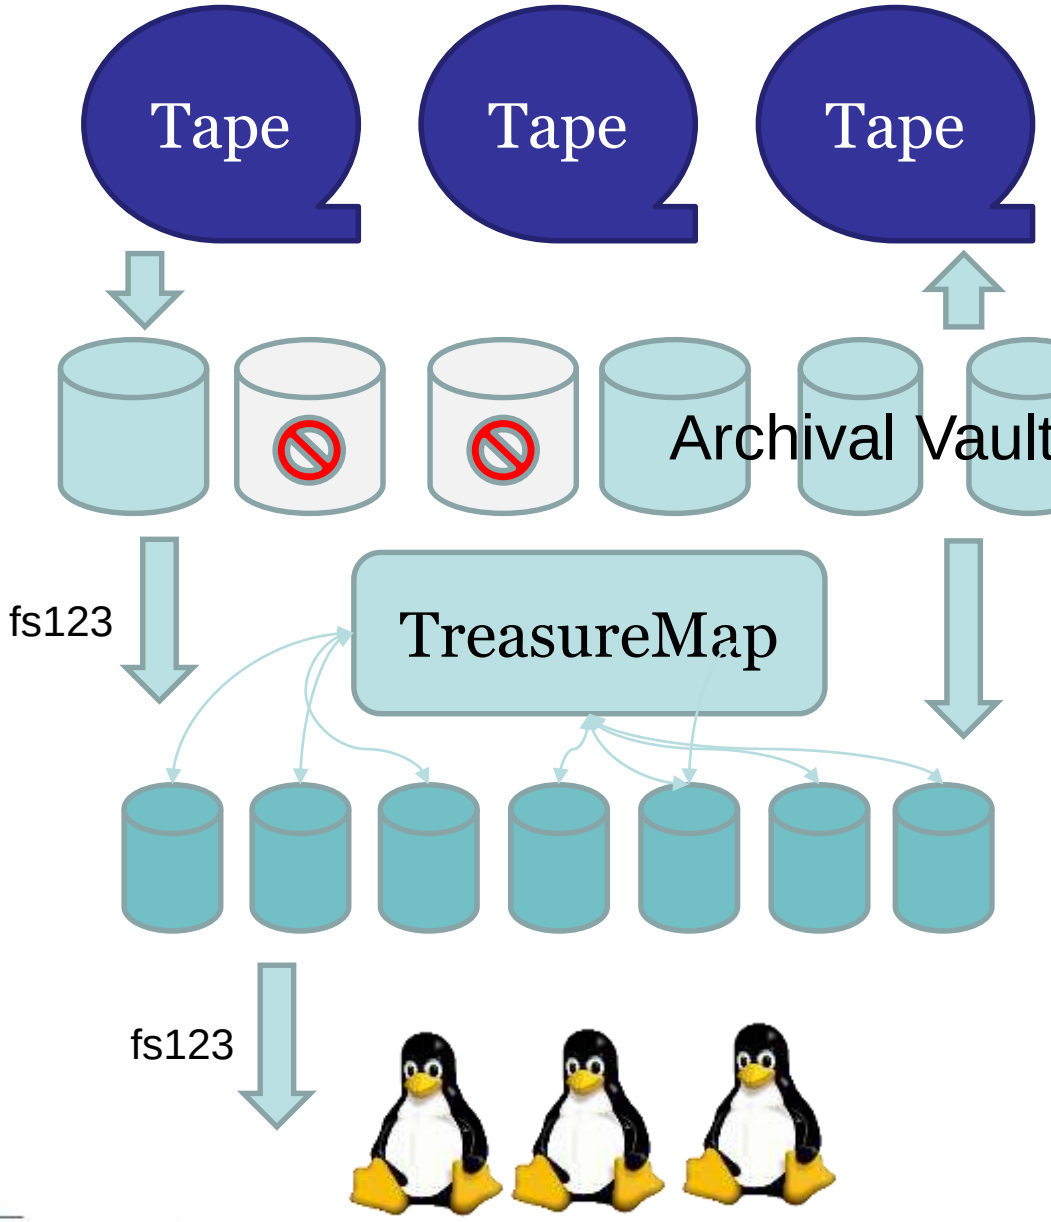

Zendo (Trajectory Cache)

Track everything!

- Database with metadata about every Anton simulation
- Every simulation has a unique, permanent numeric id
- Input and output permanently archived
- Accessible via virtual filesystem (TreasureMap)

Reproducibility is essential

- Scientific
- Engineering

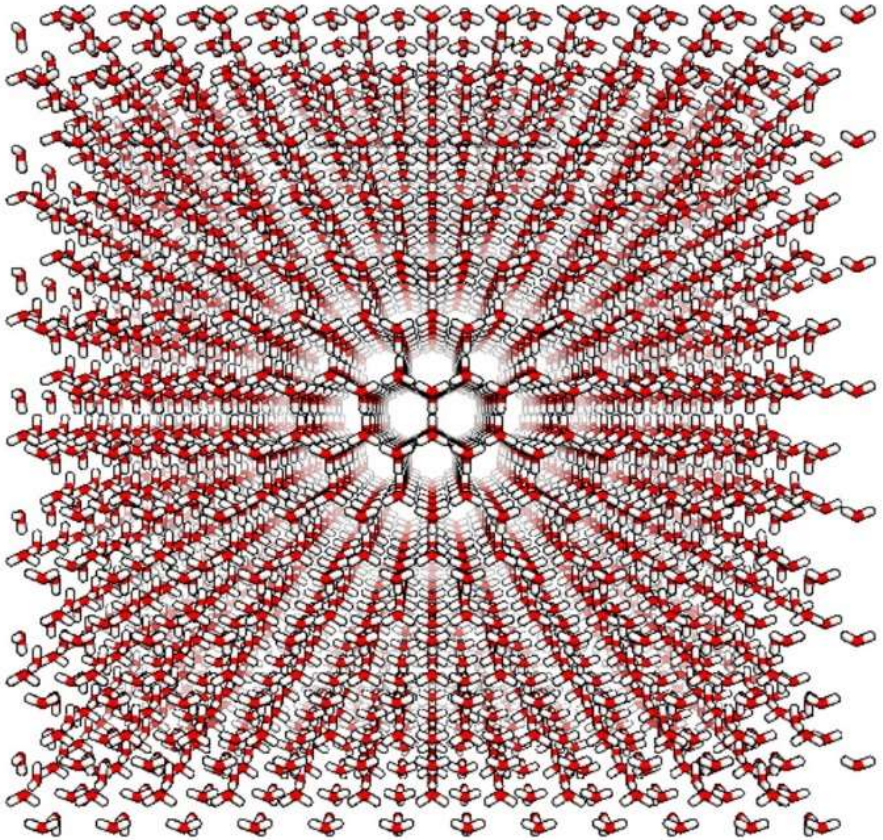

Scientific reproducibility: Dynamics
Reversible folding of 12 proteins to atomic resolution

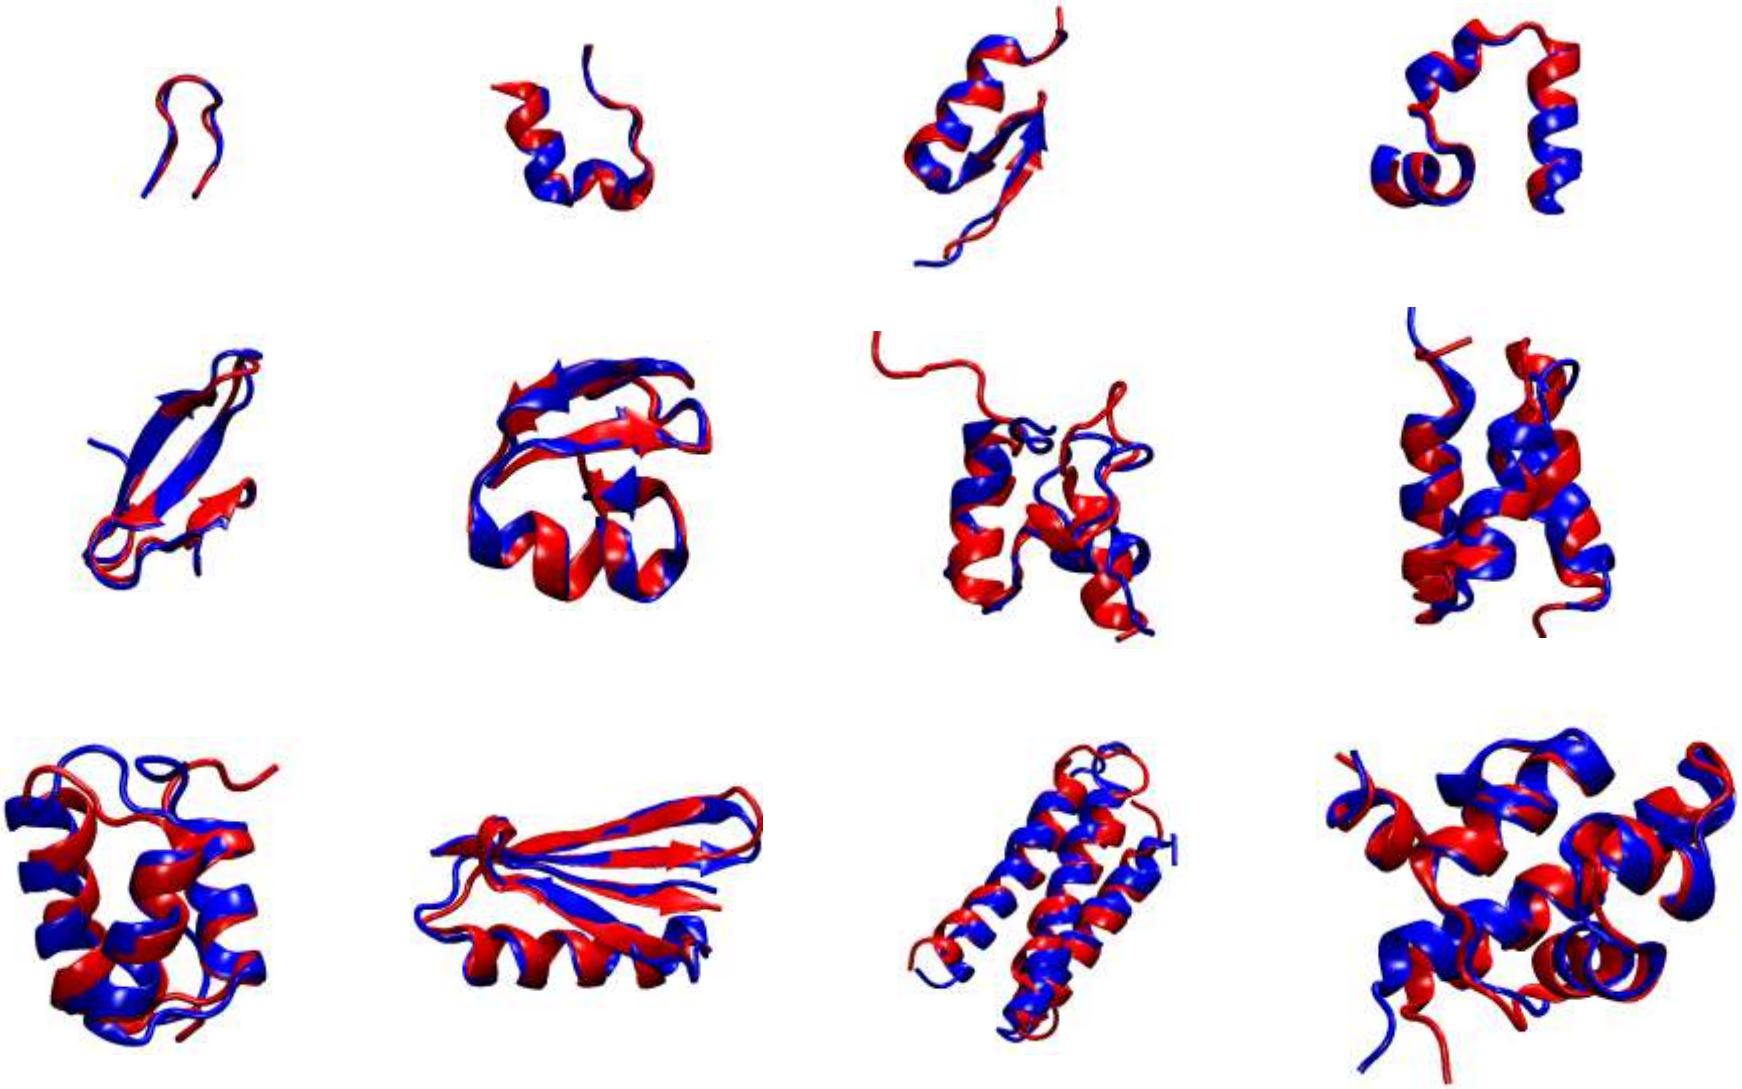

Lindorff-Larsen, Piana *et al.* *Science 334:5170 (2011)*

Engineering reproducibility: hardware

- Fixed point
  - Give up dynamic range for precision
  - Overflows: Ignore (saturation) or Halt!
  - Bitwise reproducibility enormously easier
  - Less area than comparable floating point
  - Much less design effort
- Network predictability
  - Communication patterns
  - In-network reduction

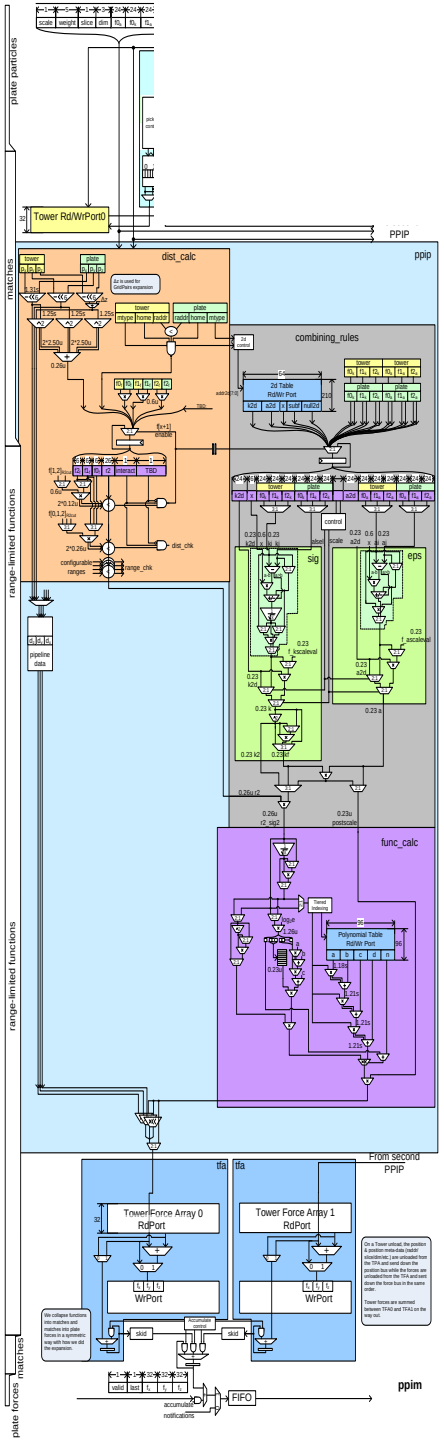

Engineering Software reproducibility

- "garden" of all software versions stored in different paths
- Expose explicit version to end-users
  - **No default**
- Versions are U.N.F (e.g. 2.7.31)
  - U = major change to User input format
    - Semantics or default assumptions changed
  - N = numerics change
  - F = fix or new features, bitwise reproducible
- Versioning must be a transitive closure
- Cannot rely on OS/distro-supplied math, compilers or interpreters

The result: Reproducible research

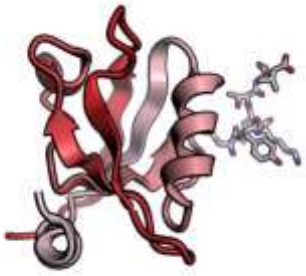

- Our researchers conduct long-timescale projects, spanning years of runs
  - usually staying with a single U.N.* series (we do backport some fixes)
- Reduced output and checkpoint rates
  - Can “zoom” by re-running parts of a trajectory
- Trust but verify:
  - 1-5% of our machine time every day is for automated re-runs of parts of jobs
  - Every mismatch automatically opens a ticket: is it a software bug, hardware error.

Biomolecular Research using Anton

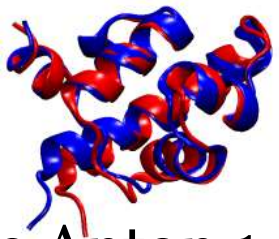

- Longest all-atom, explicit solvent MD runs before Anton 1 were about 10 microseconds
- Anton 1 **routinely** simulated biomolecules to O(100) microseconds, Anton 2 easily reaches milliseconds (and on larger molecules), allowing us to observe and understand:
  - Protein folding
  - Changes in protein conformation (and how conformations relate to function)
  - How a drug binds to a protein
  - How to find new targets for drug discovery

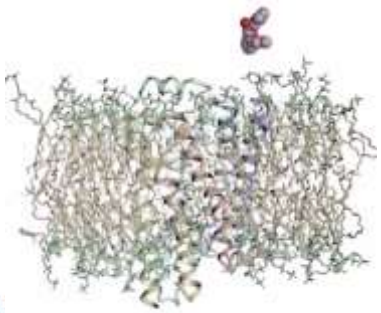

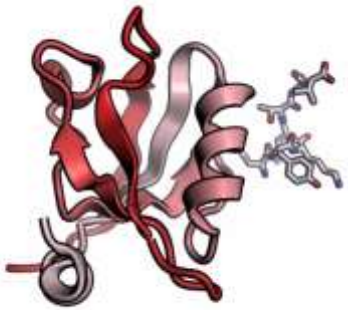

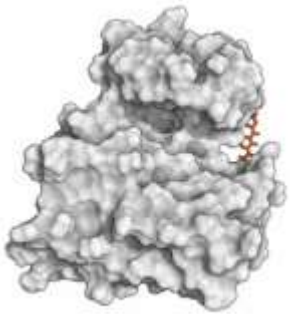

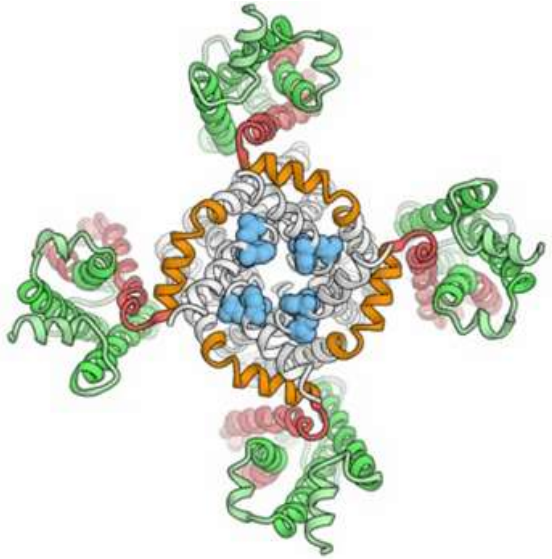

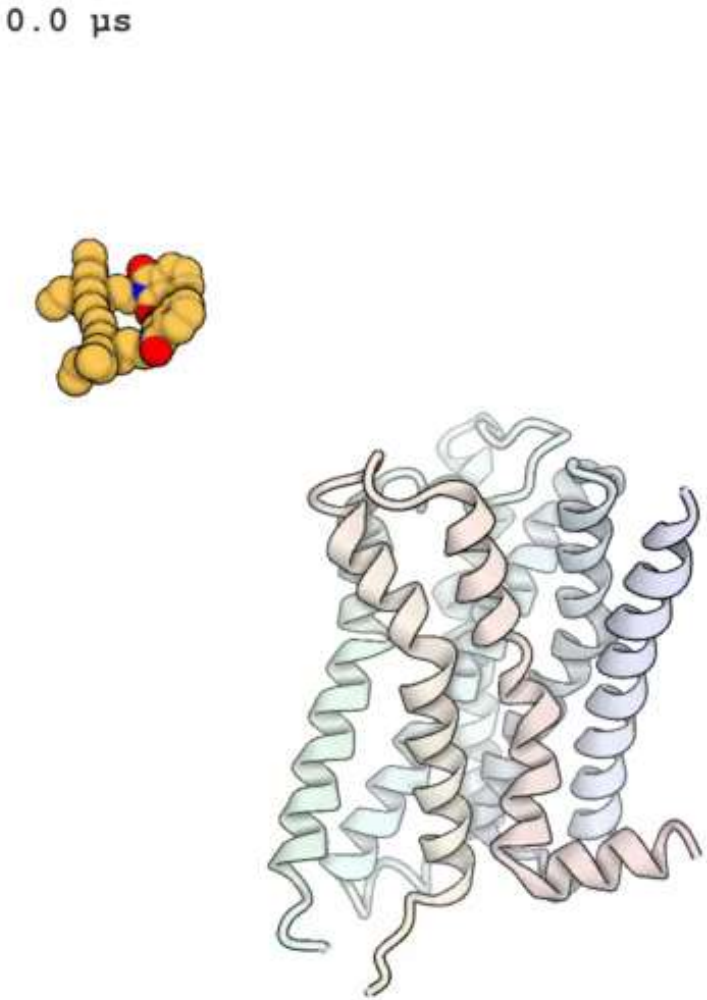

R. Dror *et al*. “Structural Basis for Modulation of a G-Protein-Coupled Receptor by Allosteric Drugs,” *Nature*, vol. 503, no. 7475, Nov. 2013, pp. 295-9.

0.0 µs

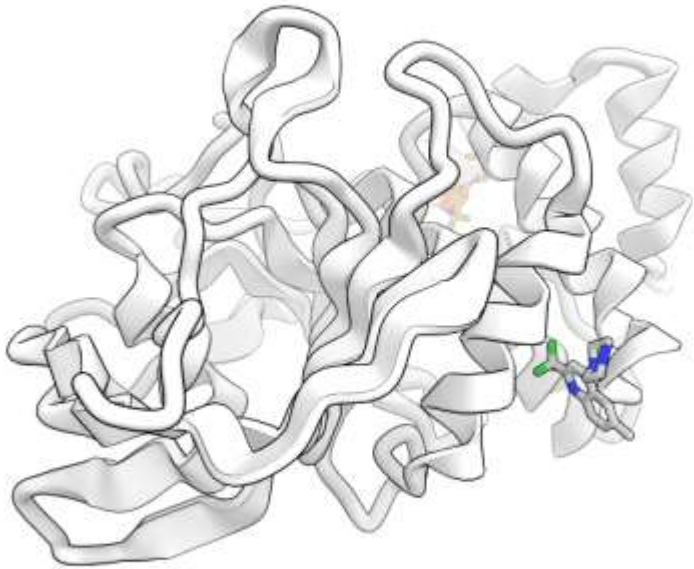

P. Maragakis *et al*, Unpublished.

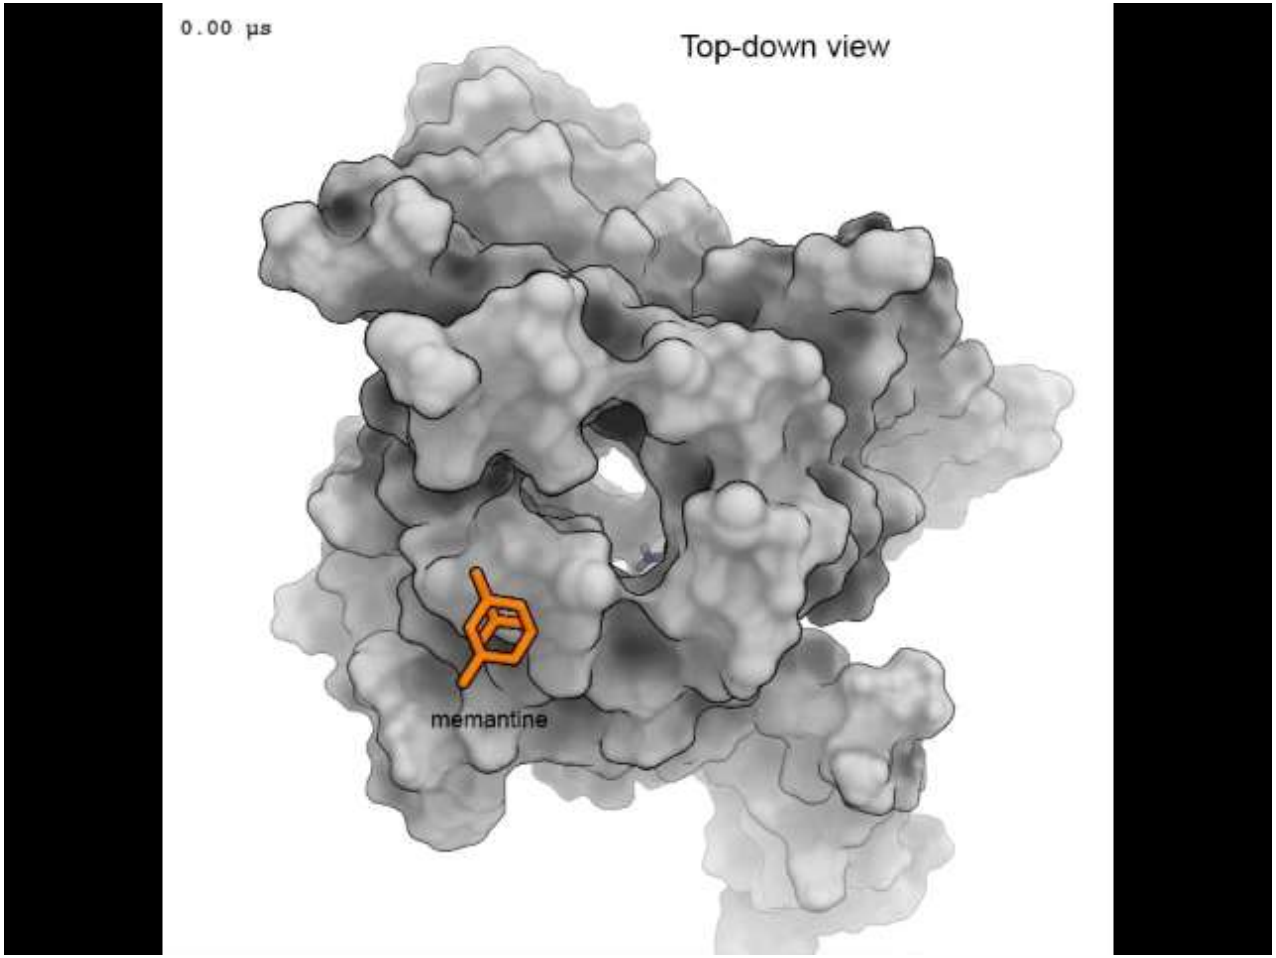

Song, Jensen *et al*. “Mechanism of NMDA receptor channel block by MK-801 and memantine.,” *Nature* **556**, 515–519 (2018) |.

0.0 μs

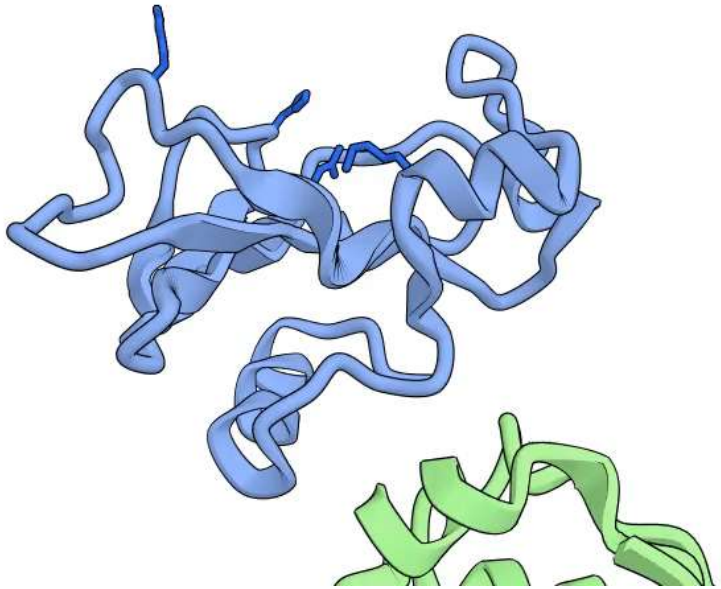

A. Pan *et al*. “Atomic-level characterization of protein-protein association,” PNAS March 5, 2019 116 (10) 4244-4249.

Questions?