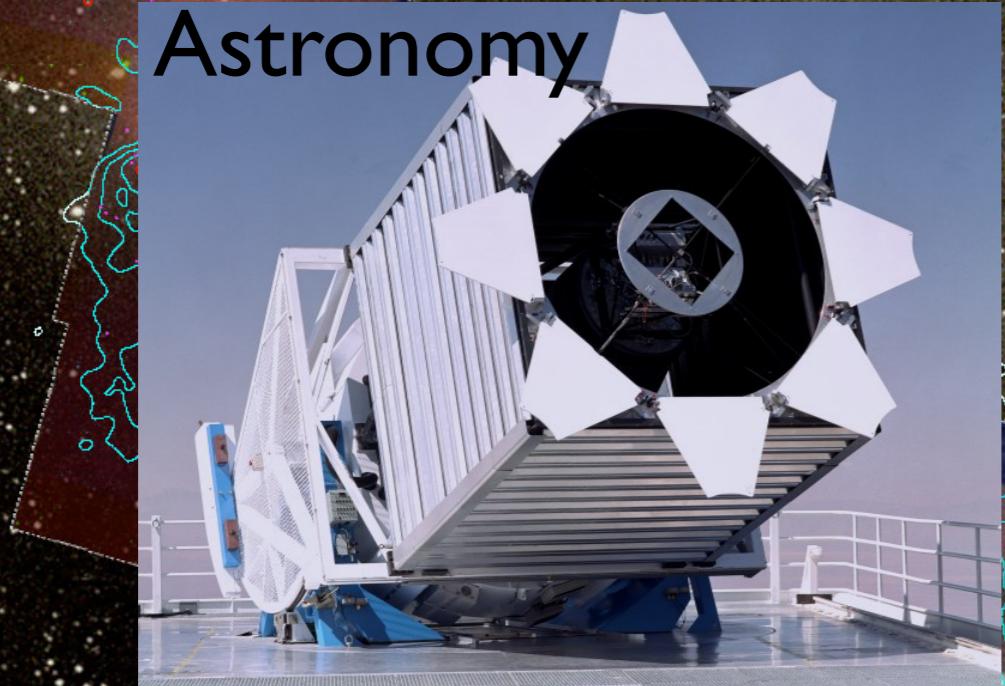


GPUs for High-Throughput Science

Hanspeter Pfister, Harvard University
pfister@seas.harvard.edu

Scientific Data Explosion

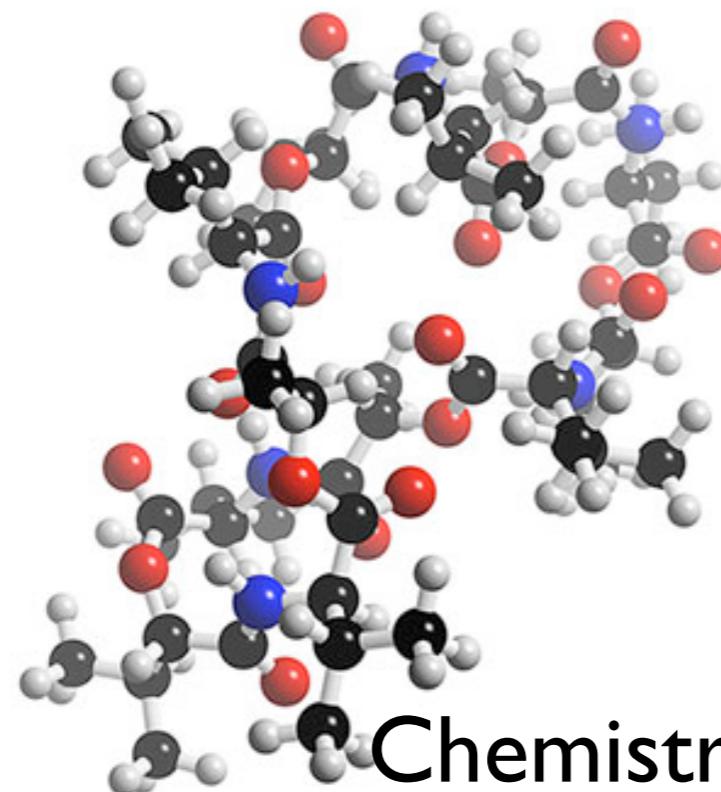
Astronomy



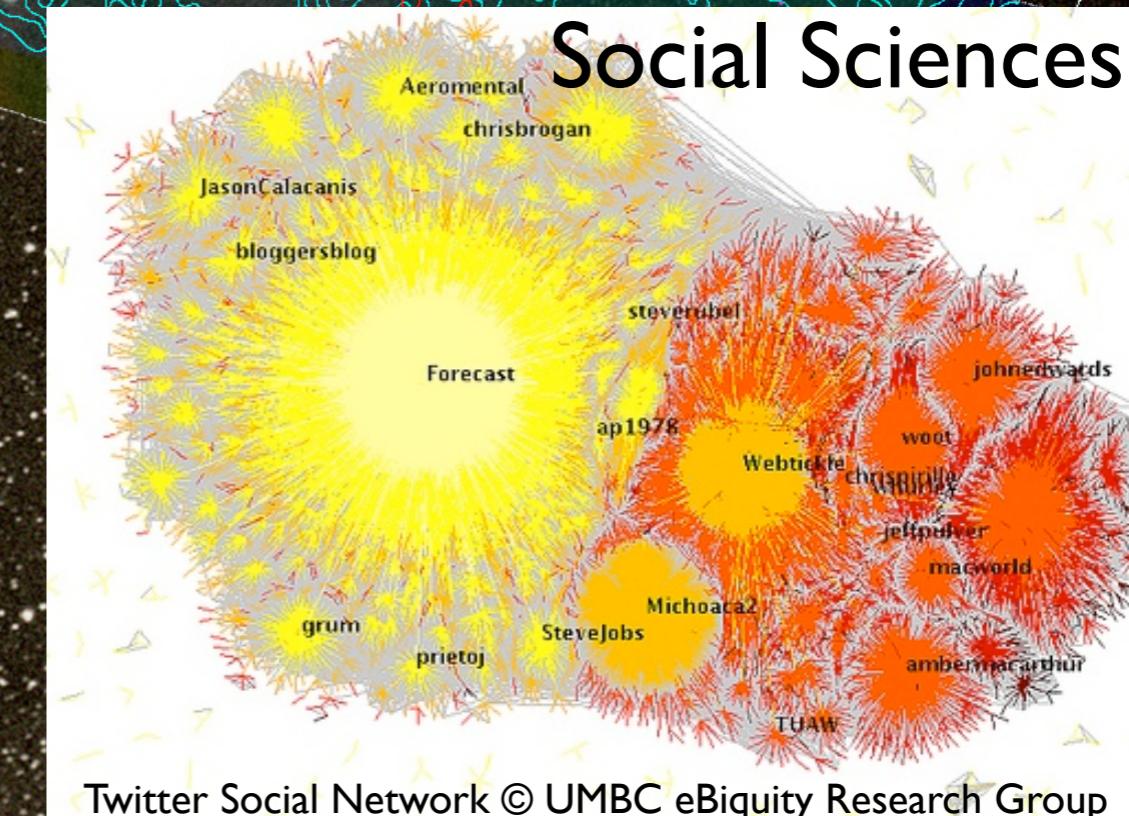
Microscopy



Chemistry



Social Sciences



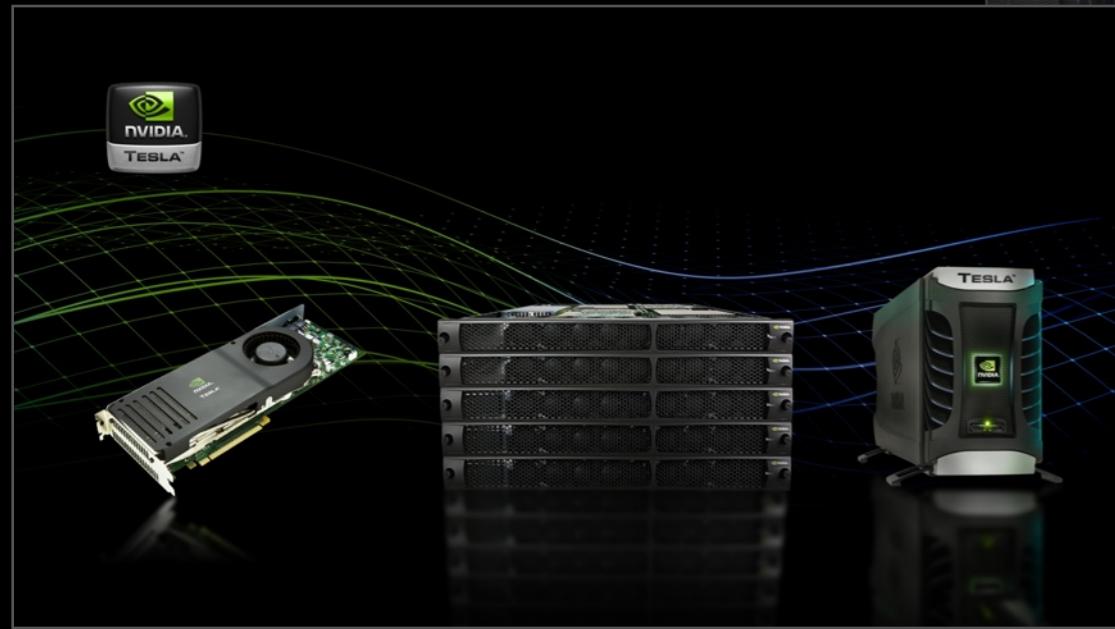
Twitter Social Network © UMBC eBiquity Research Group

Computing Revolution

- Commodity energy-efficient supercomputers



4 Teraflops @ 1400 W





[51]

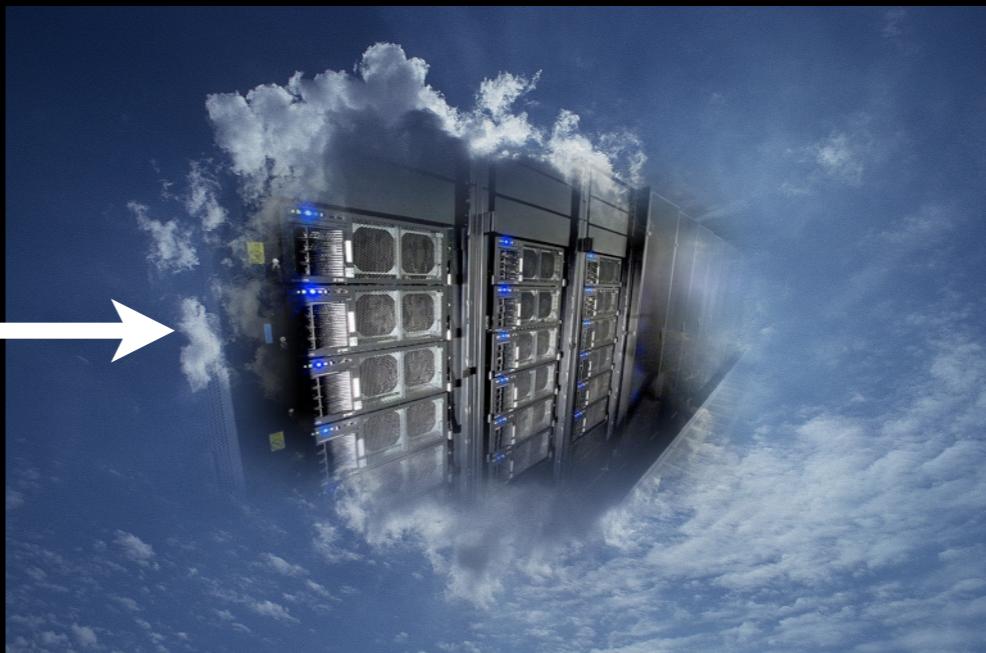
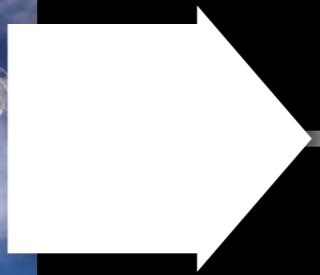
[U]

BUFF

Traditional



Sensor



Modeling
Simulation
Analysis

Challenge

- Limited by communication, not processor performance [Jim Grey, Distributed Computing Economics]

the **INQUIRER**
News, reviews, facts and friction

Saturday, 12 September 2009 | [INQ Mobile](#) | [RSS](#)

Pigeon protocol offers faster data delivery

The pigeon is mightier than ADSL

By **David Neal**

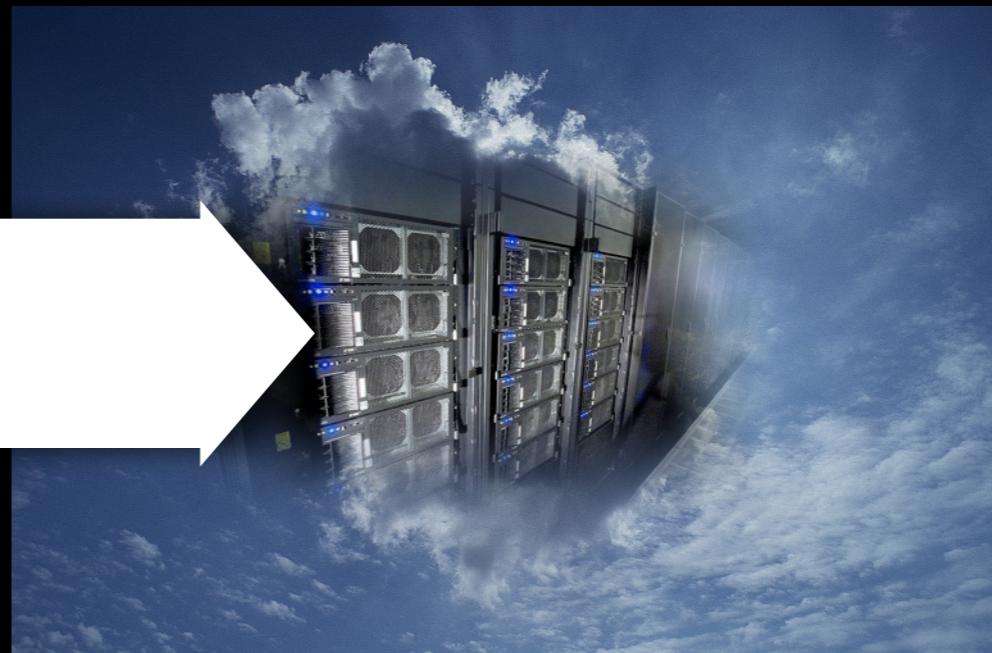
Thursday, 10 September 2009, 14:15

AN UNORTHODOX RACE in South Africa has revealed that it is faster to transmit 4GB of data via carrier pigeon than it is to send it over the country's main ADSL services.

Challenge



Sensor



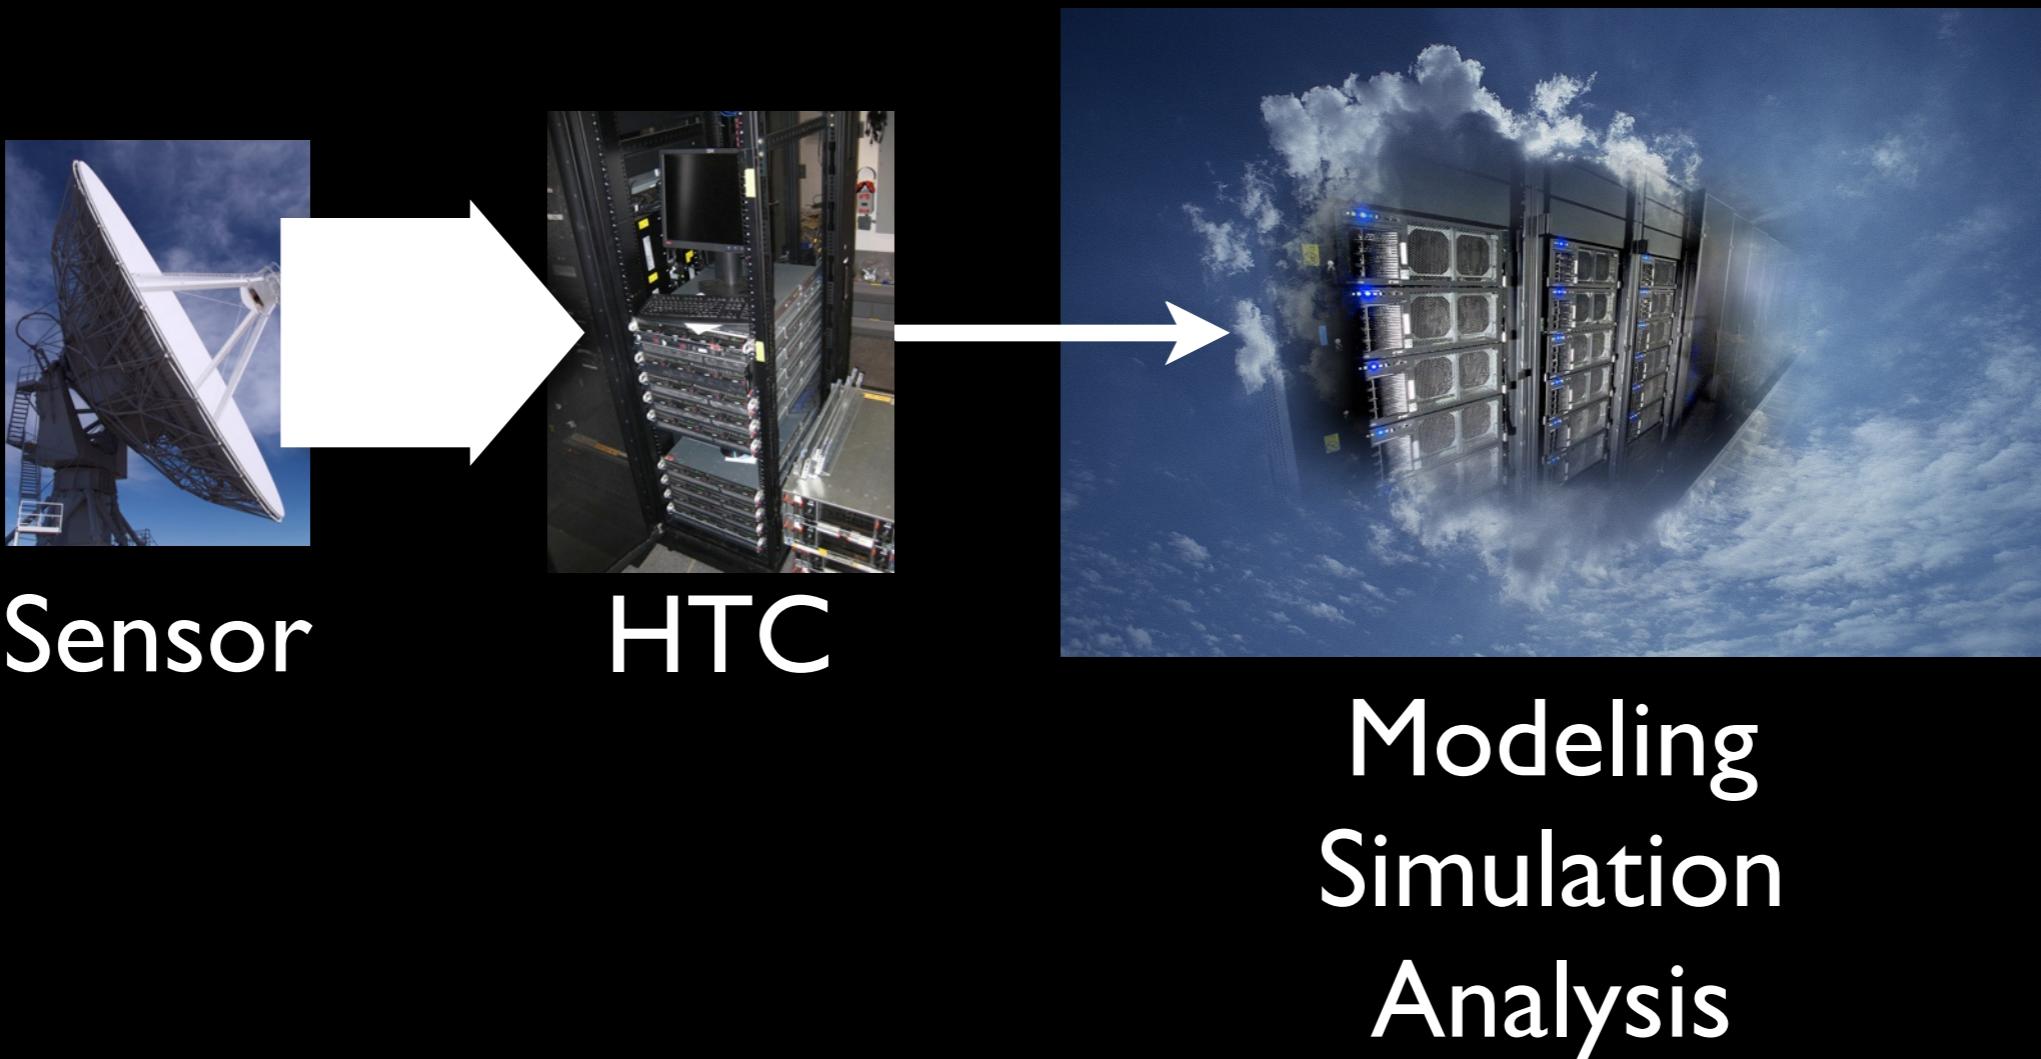
Modeling
Simulation
Analysis

Challenge

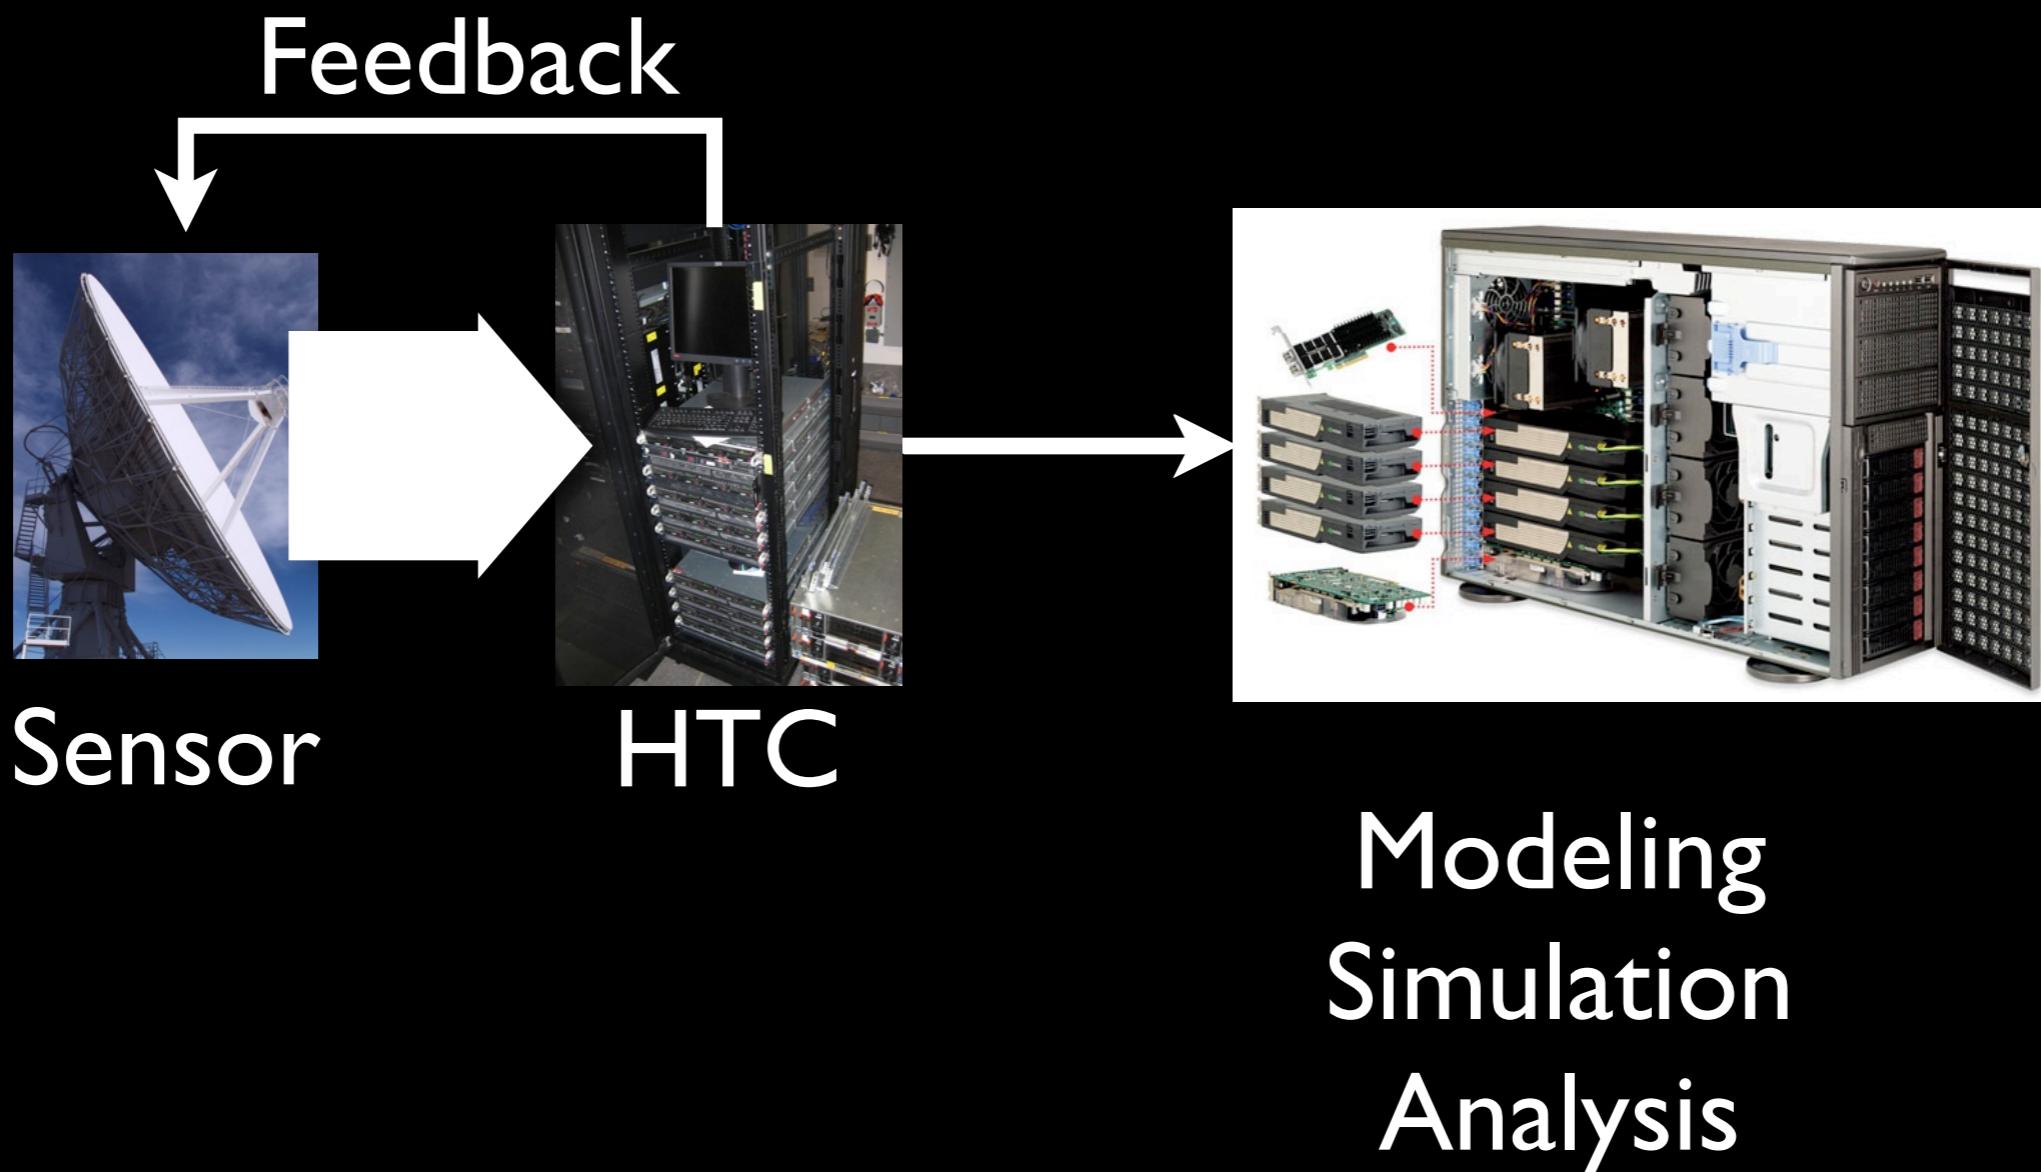


Modeling
Simulation
Analysis

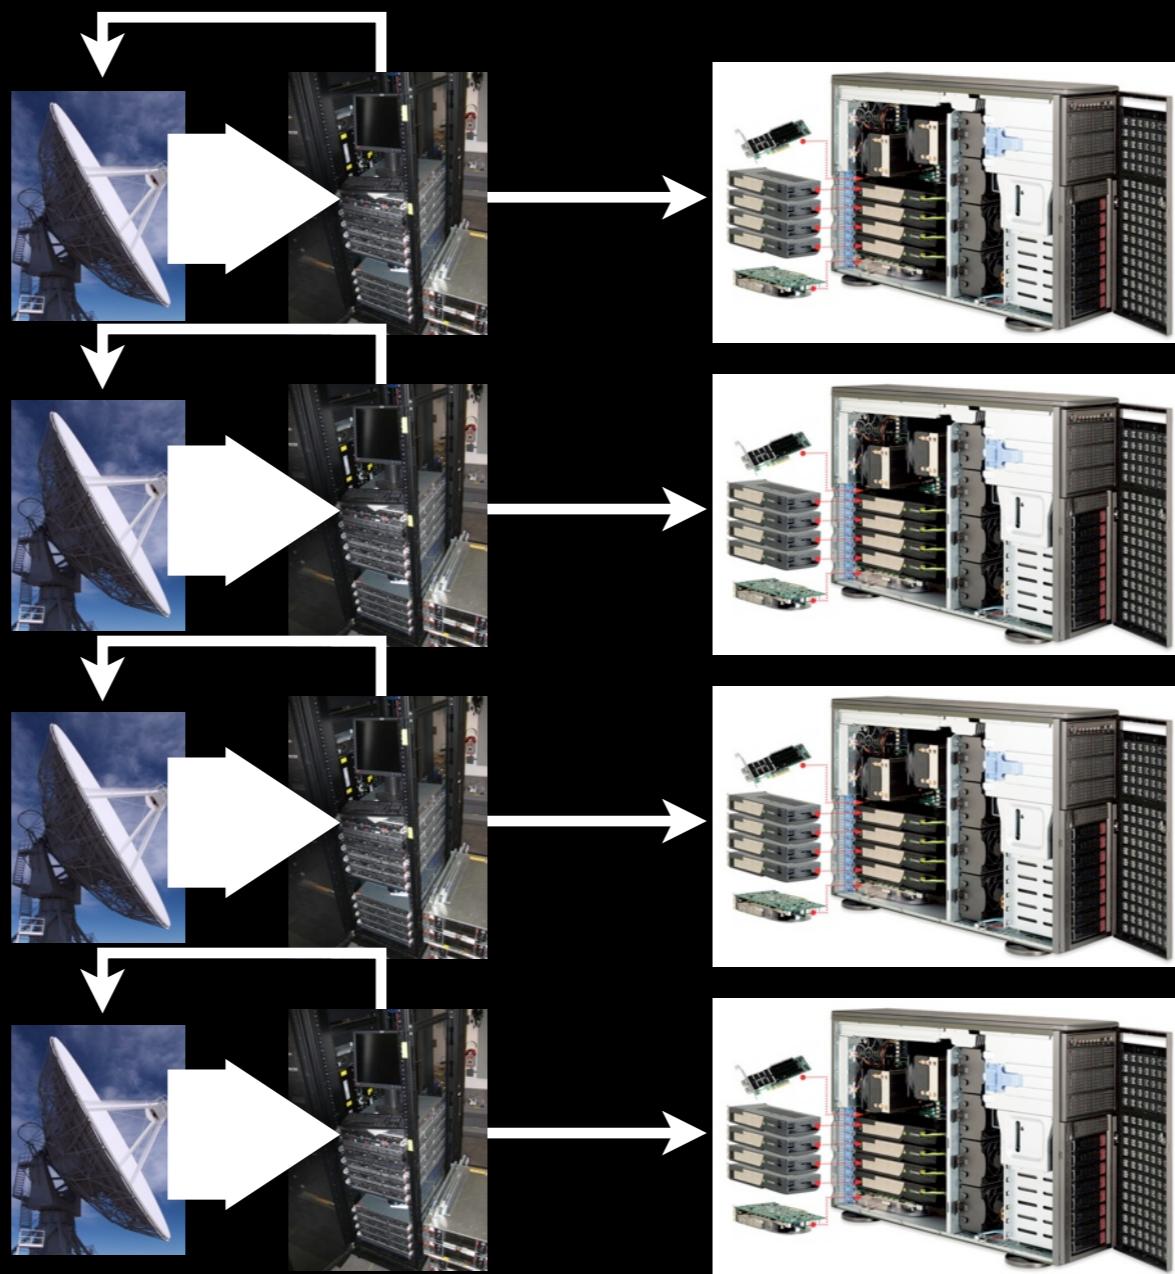
High-Throughput Computing



High-Throughput Computing



High-Throughput Computing



The Connectome

Discovering the Wiring Diagram of the Brain

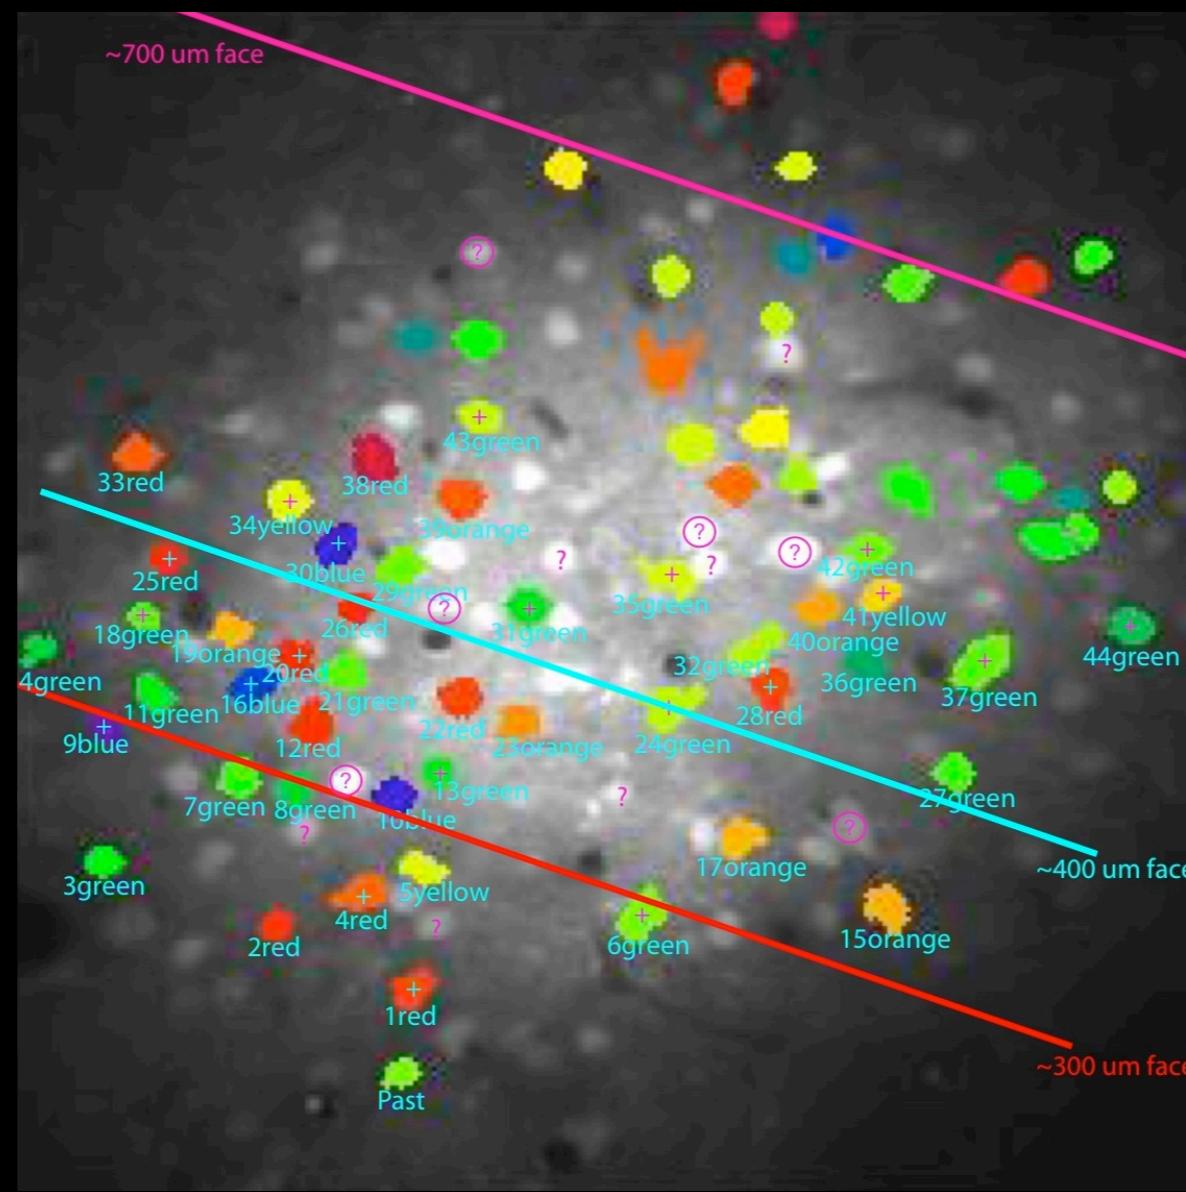
Collaborators

- Harvard Center for Brain Science
 - Jeff Lichtman & Clay Reid
- Kitware Inc.
 - Will Schroeder, Charles Law, Rusty Blue
- VRVis Vienna
 - Markus Hadwiger, Johanna Beyer
- SEAS
 - Amelio Vazquez, Won-Ki Jeong
 - Hanspeter Pfister



The Scientific Challenge

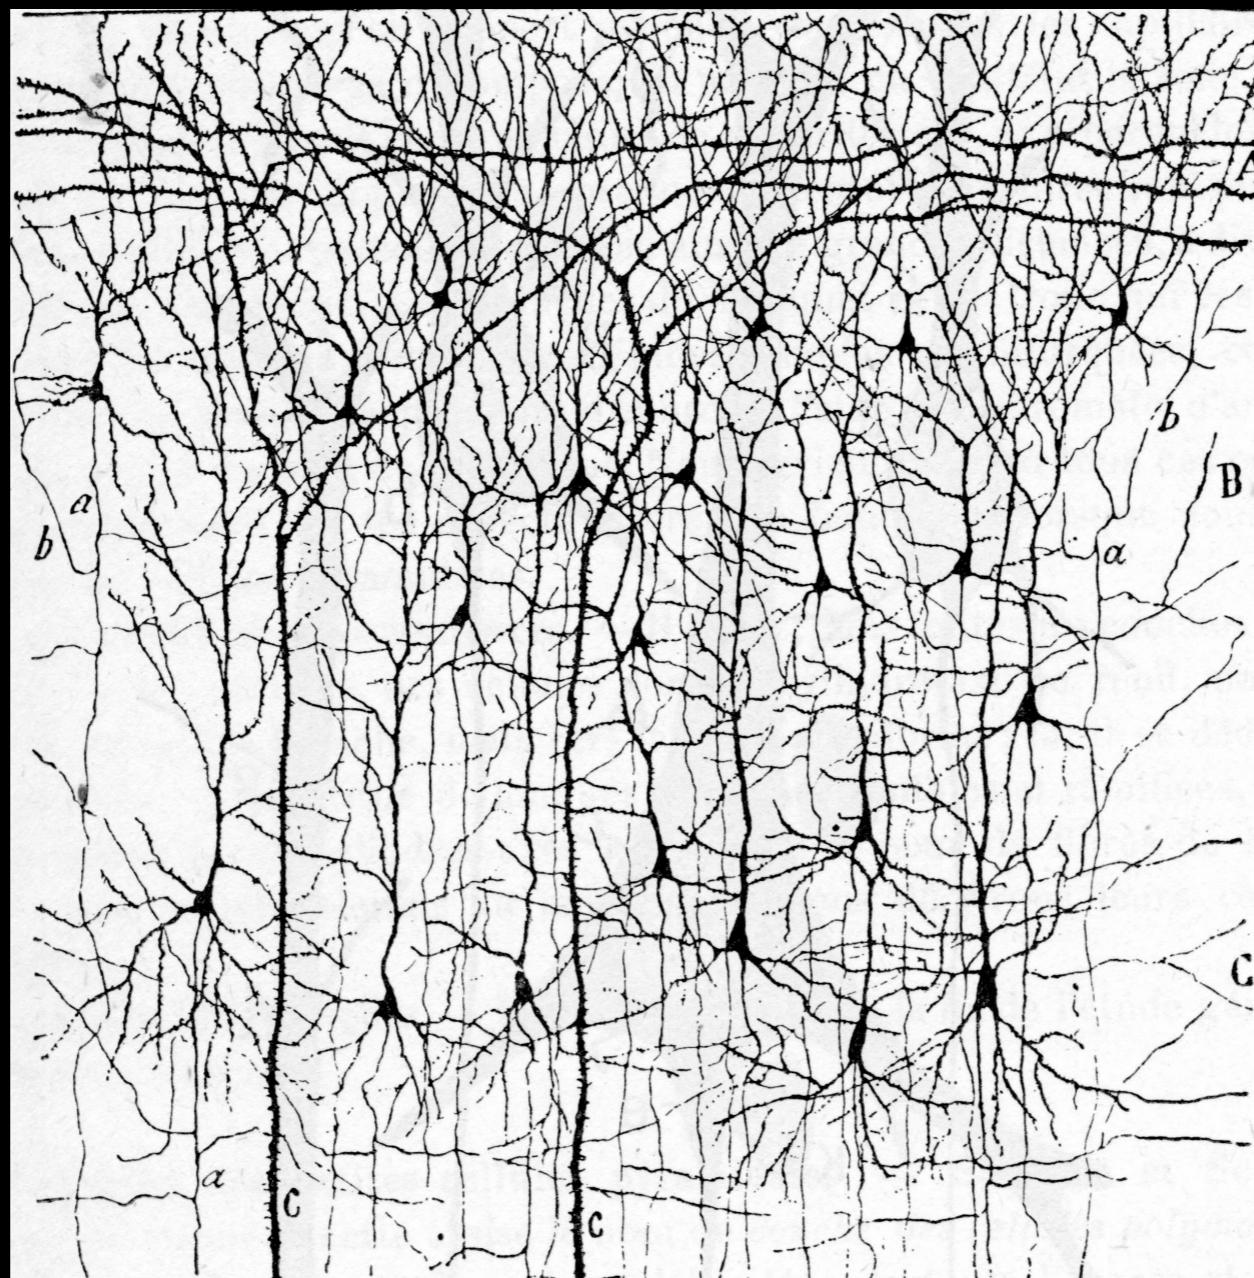
- What do large brain circuits do?



C. Reid

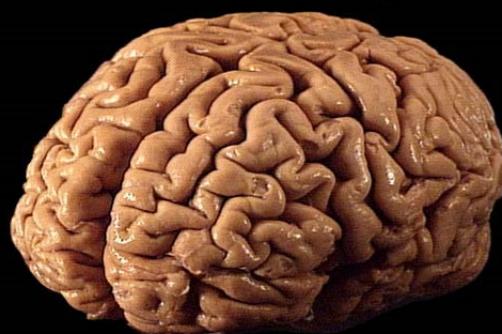
The Scientific Challenge

- What is their connectivity?

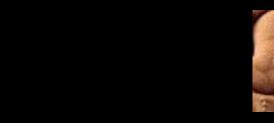


Ramón y Cajal, 1905

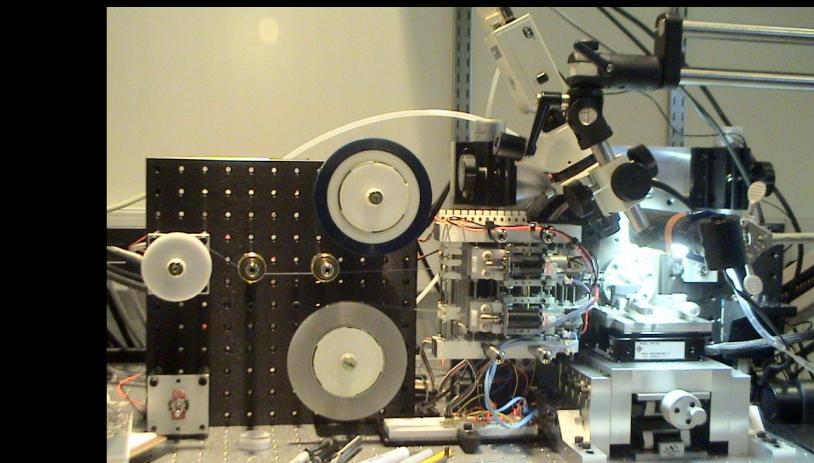
Connectome Workflow



Get a brain



Get a piece



Cut it thin



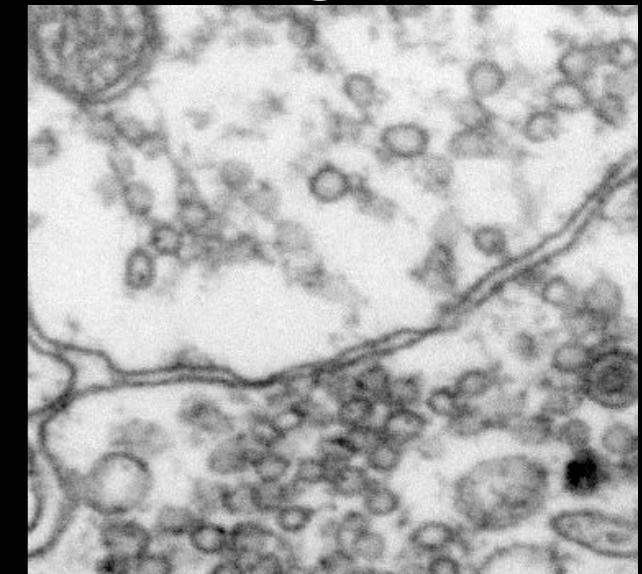
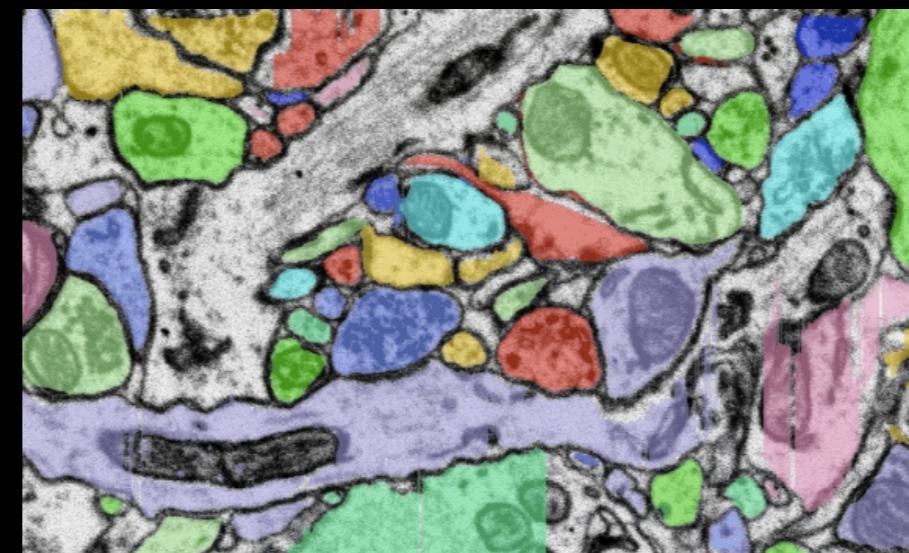
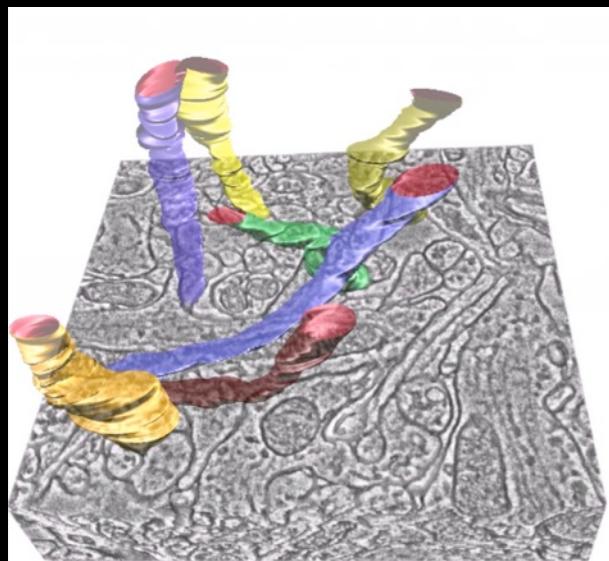
Visualize & Analyze



Reconstruct

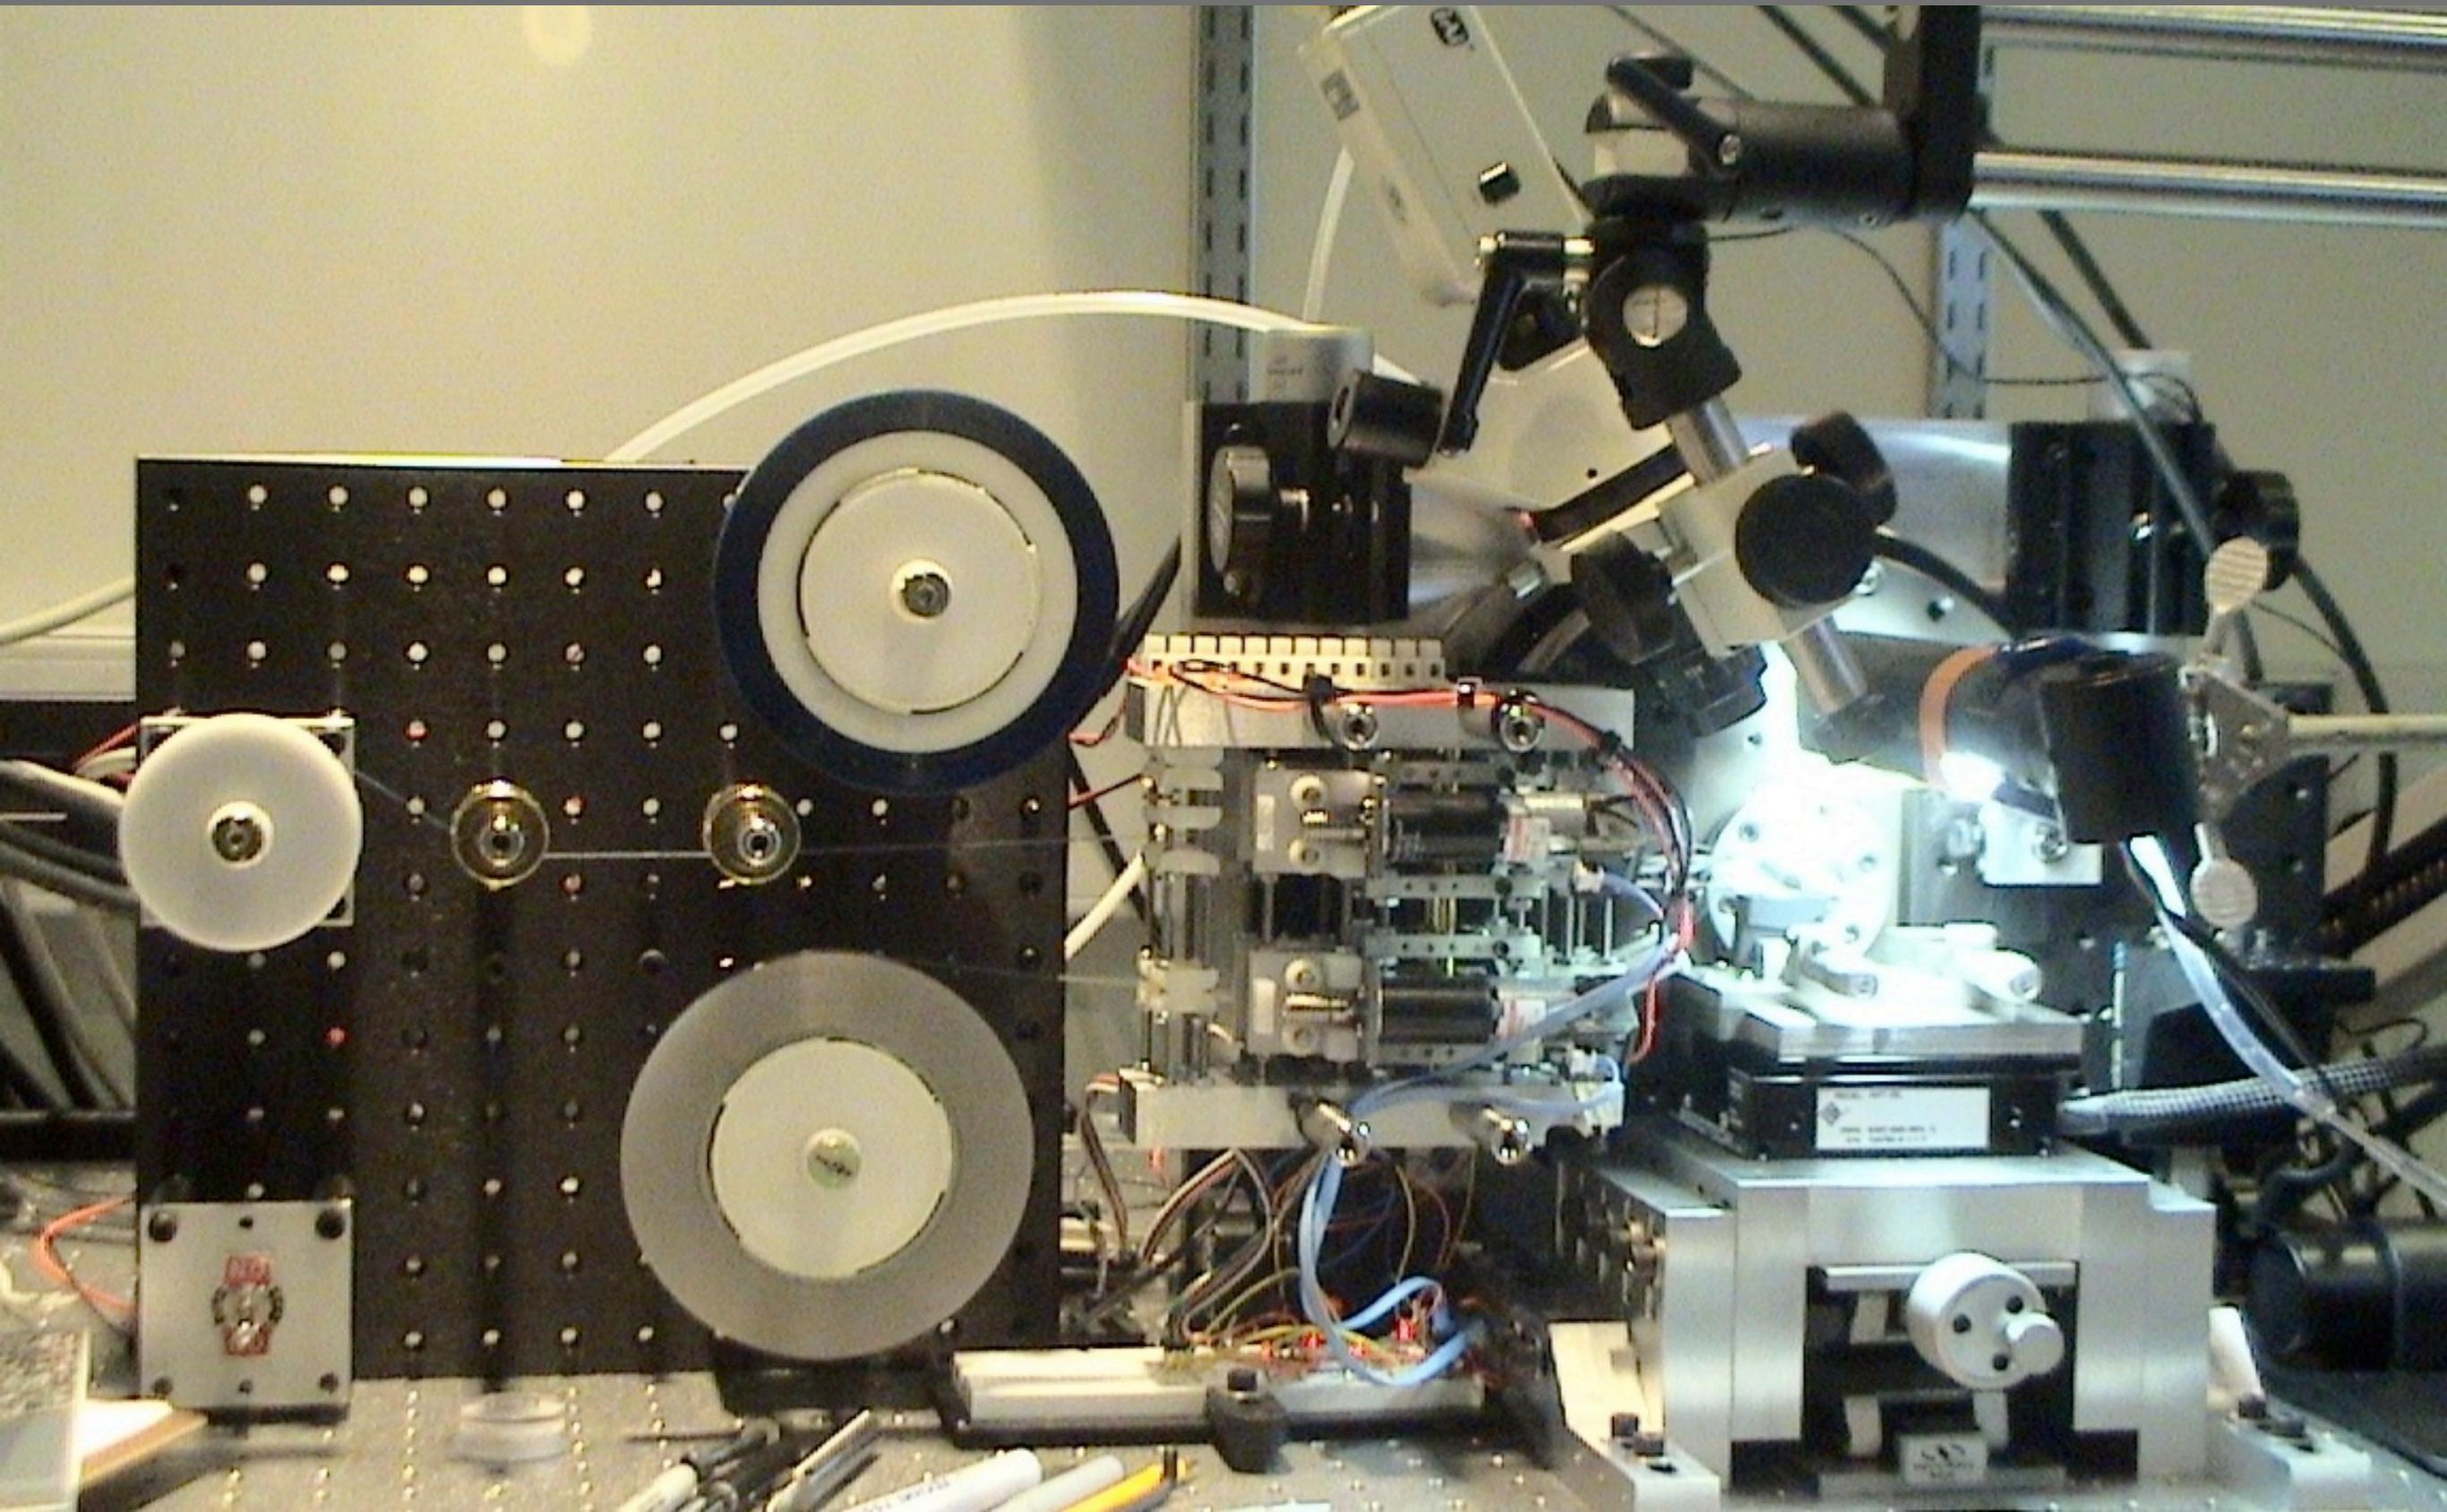


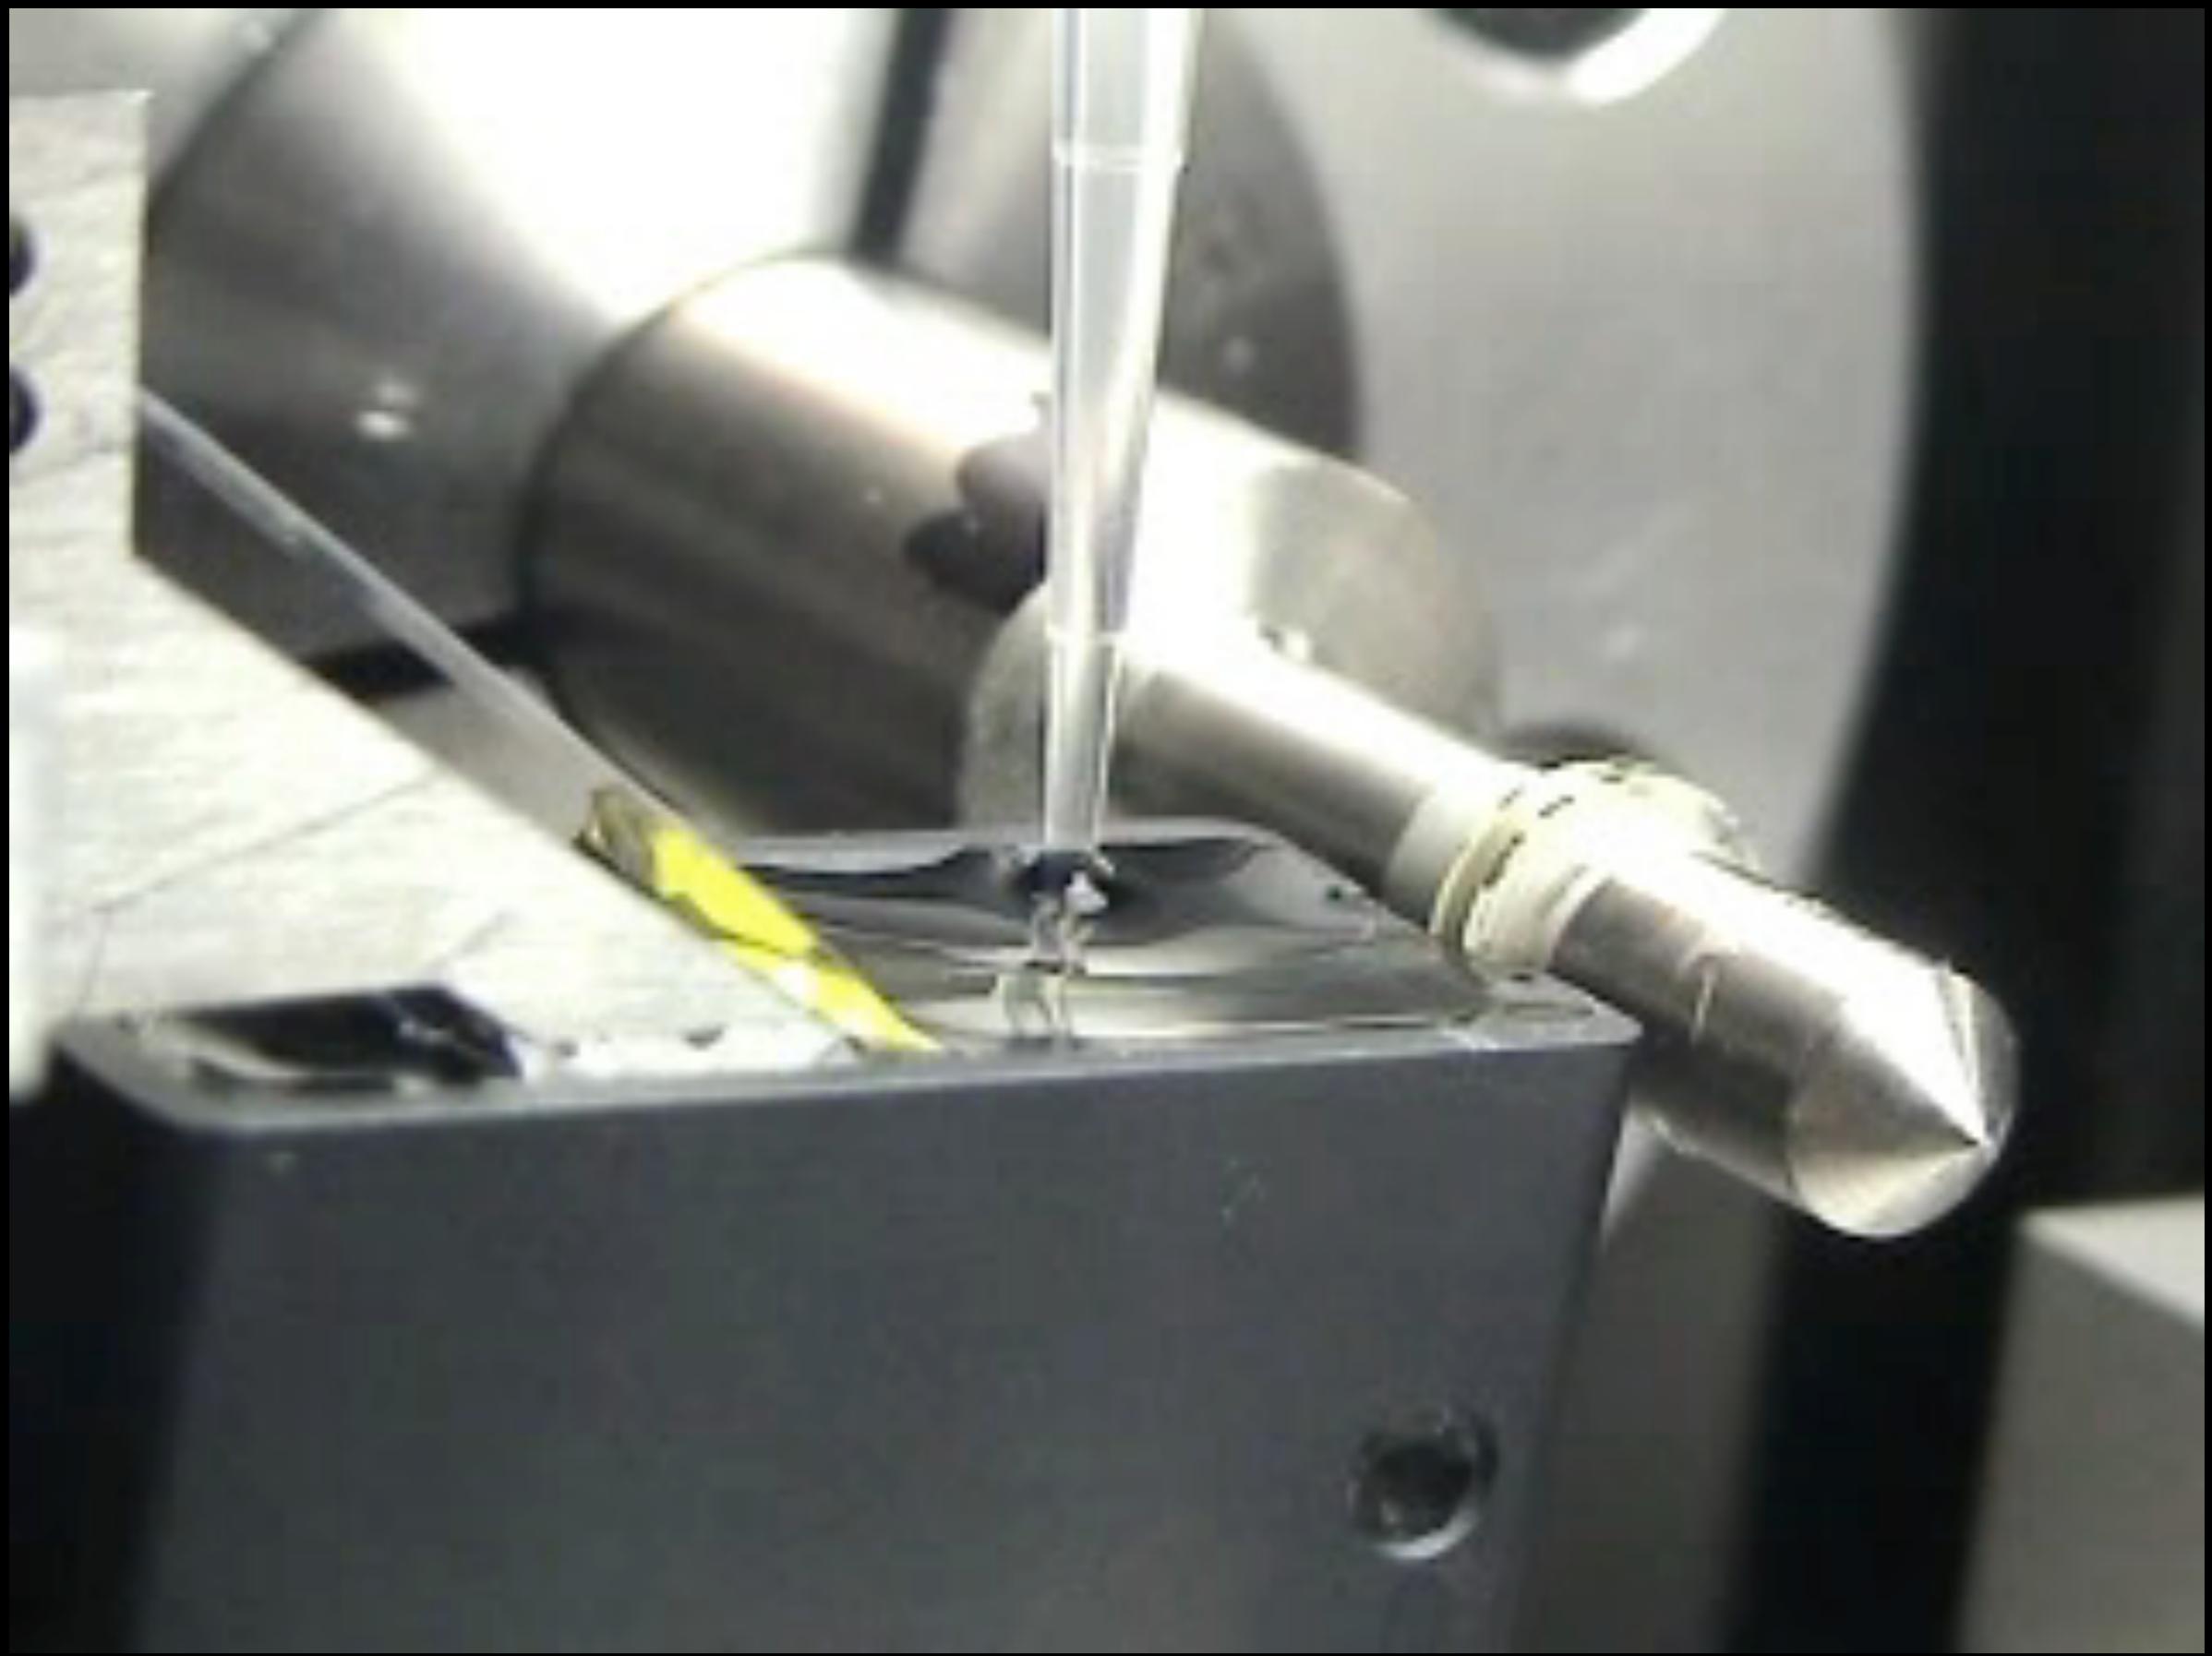
Image it



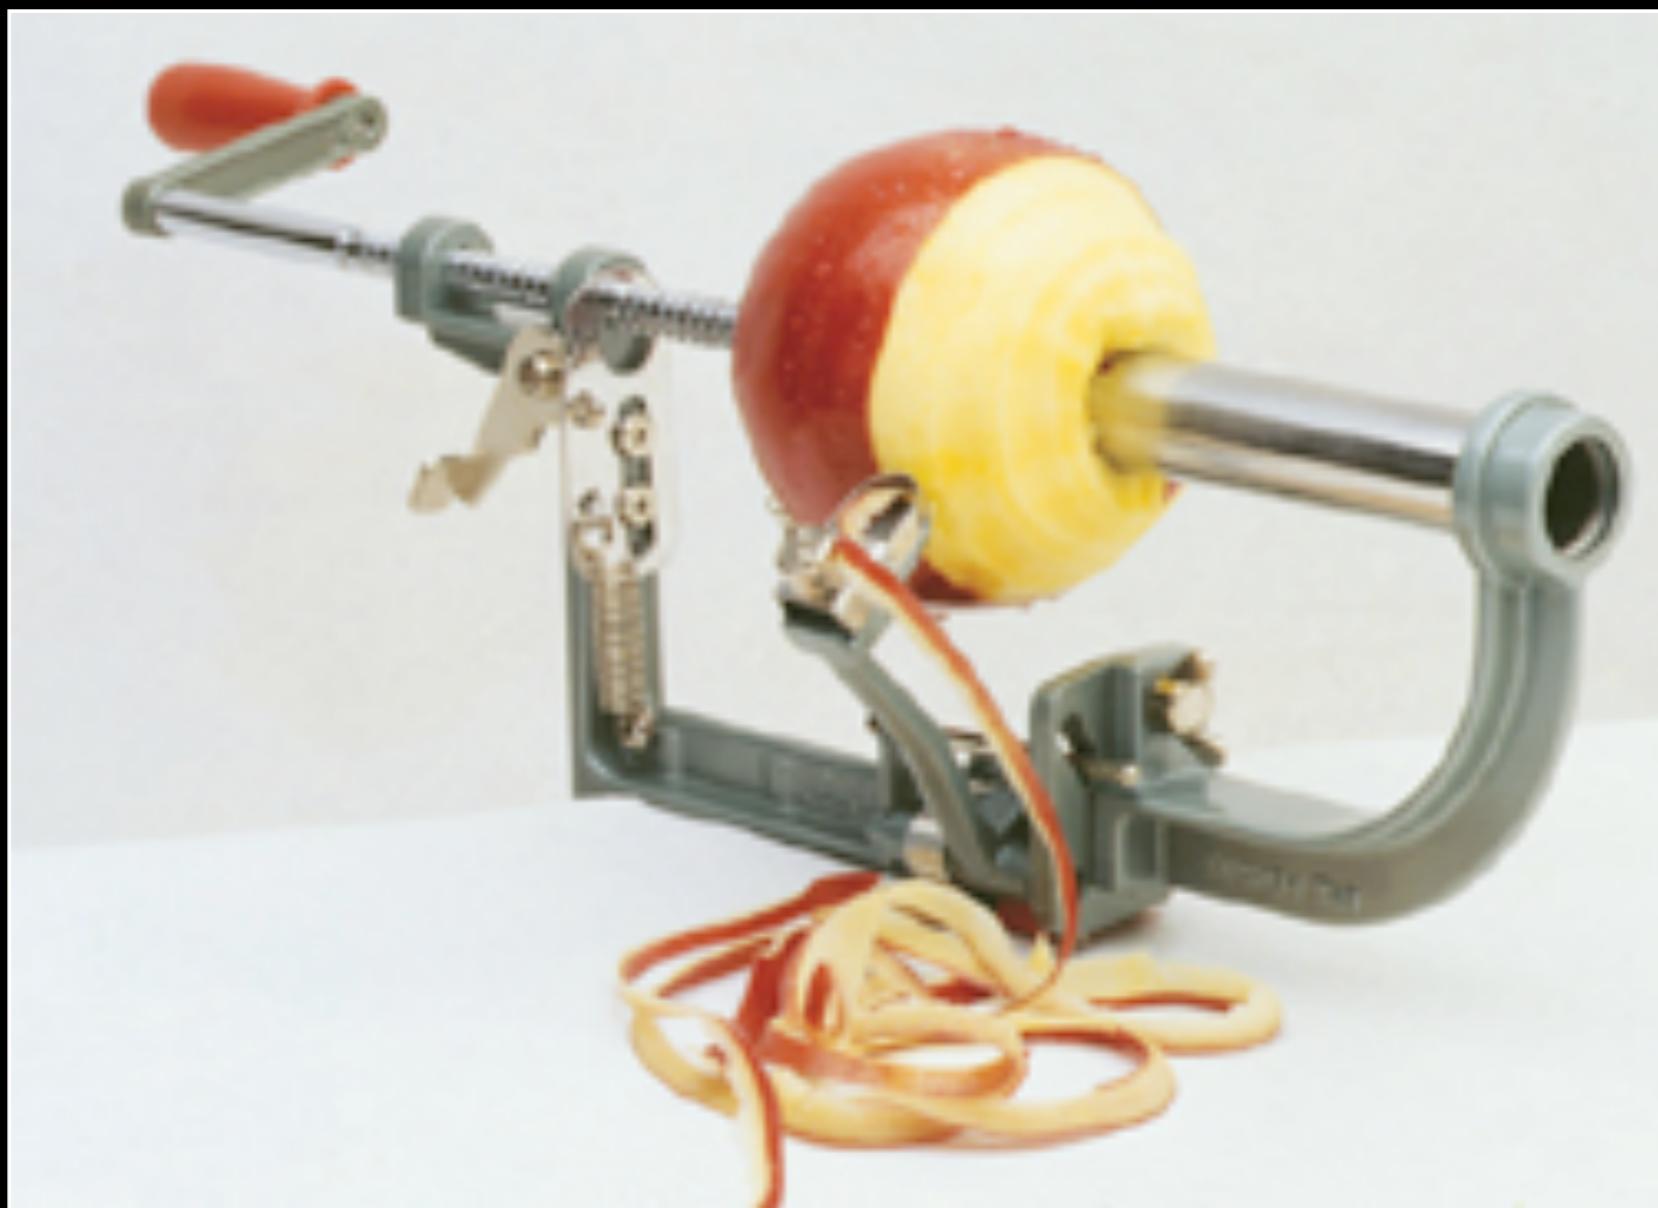
Brain on Tape: Automatic Tape-collecting UltraMicrotomes

Ken Hayworth & Bobby Kasthuri



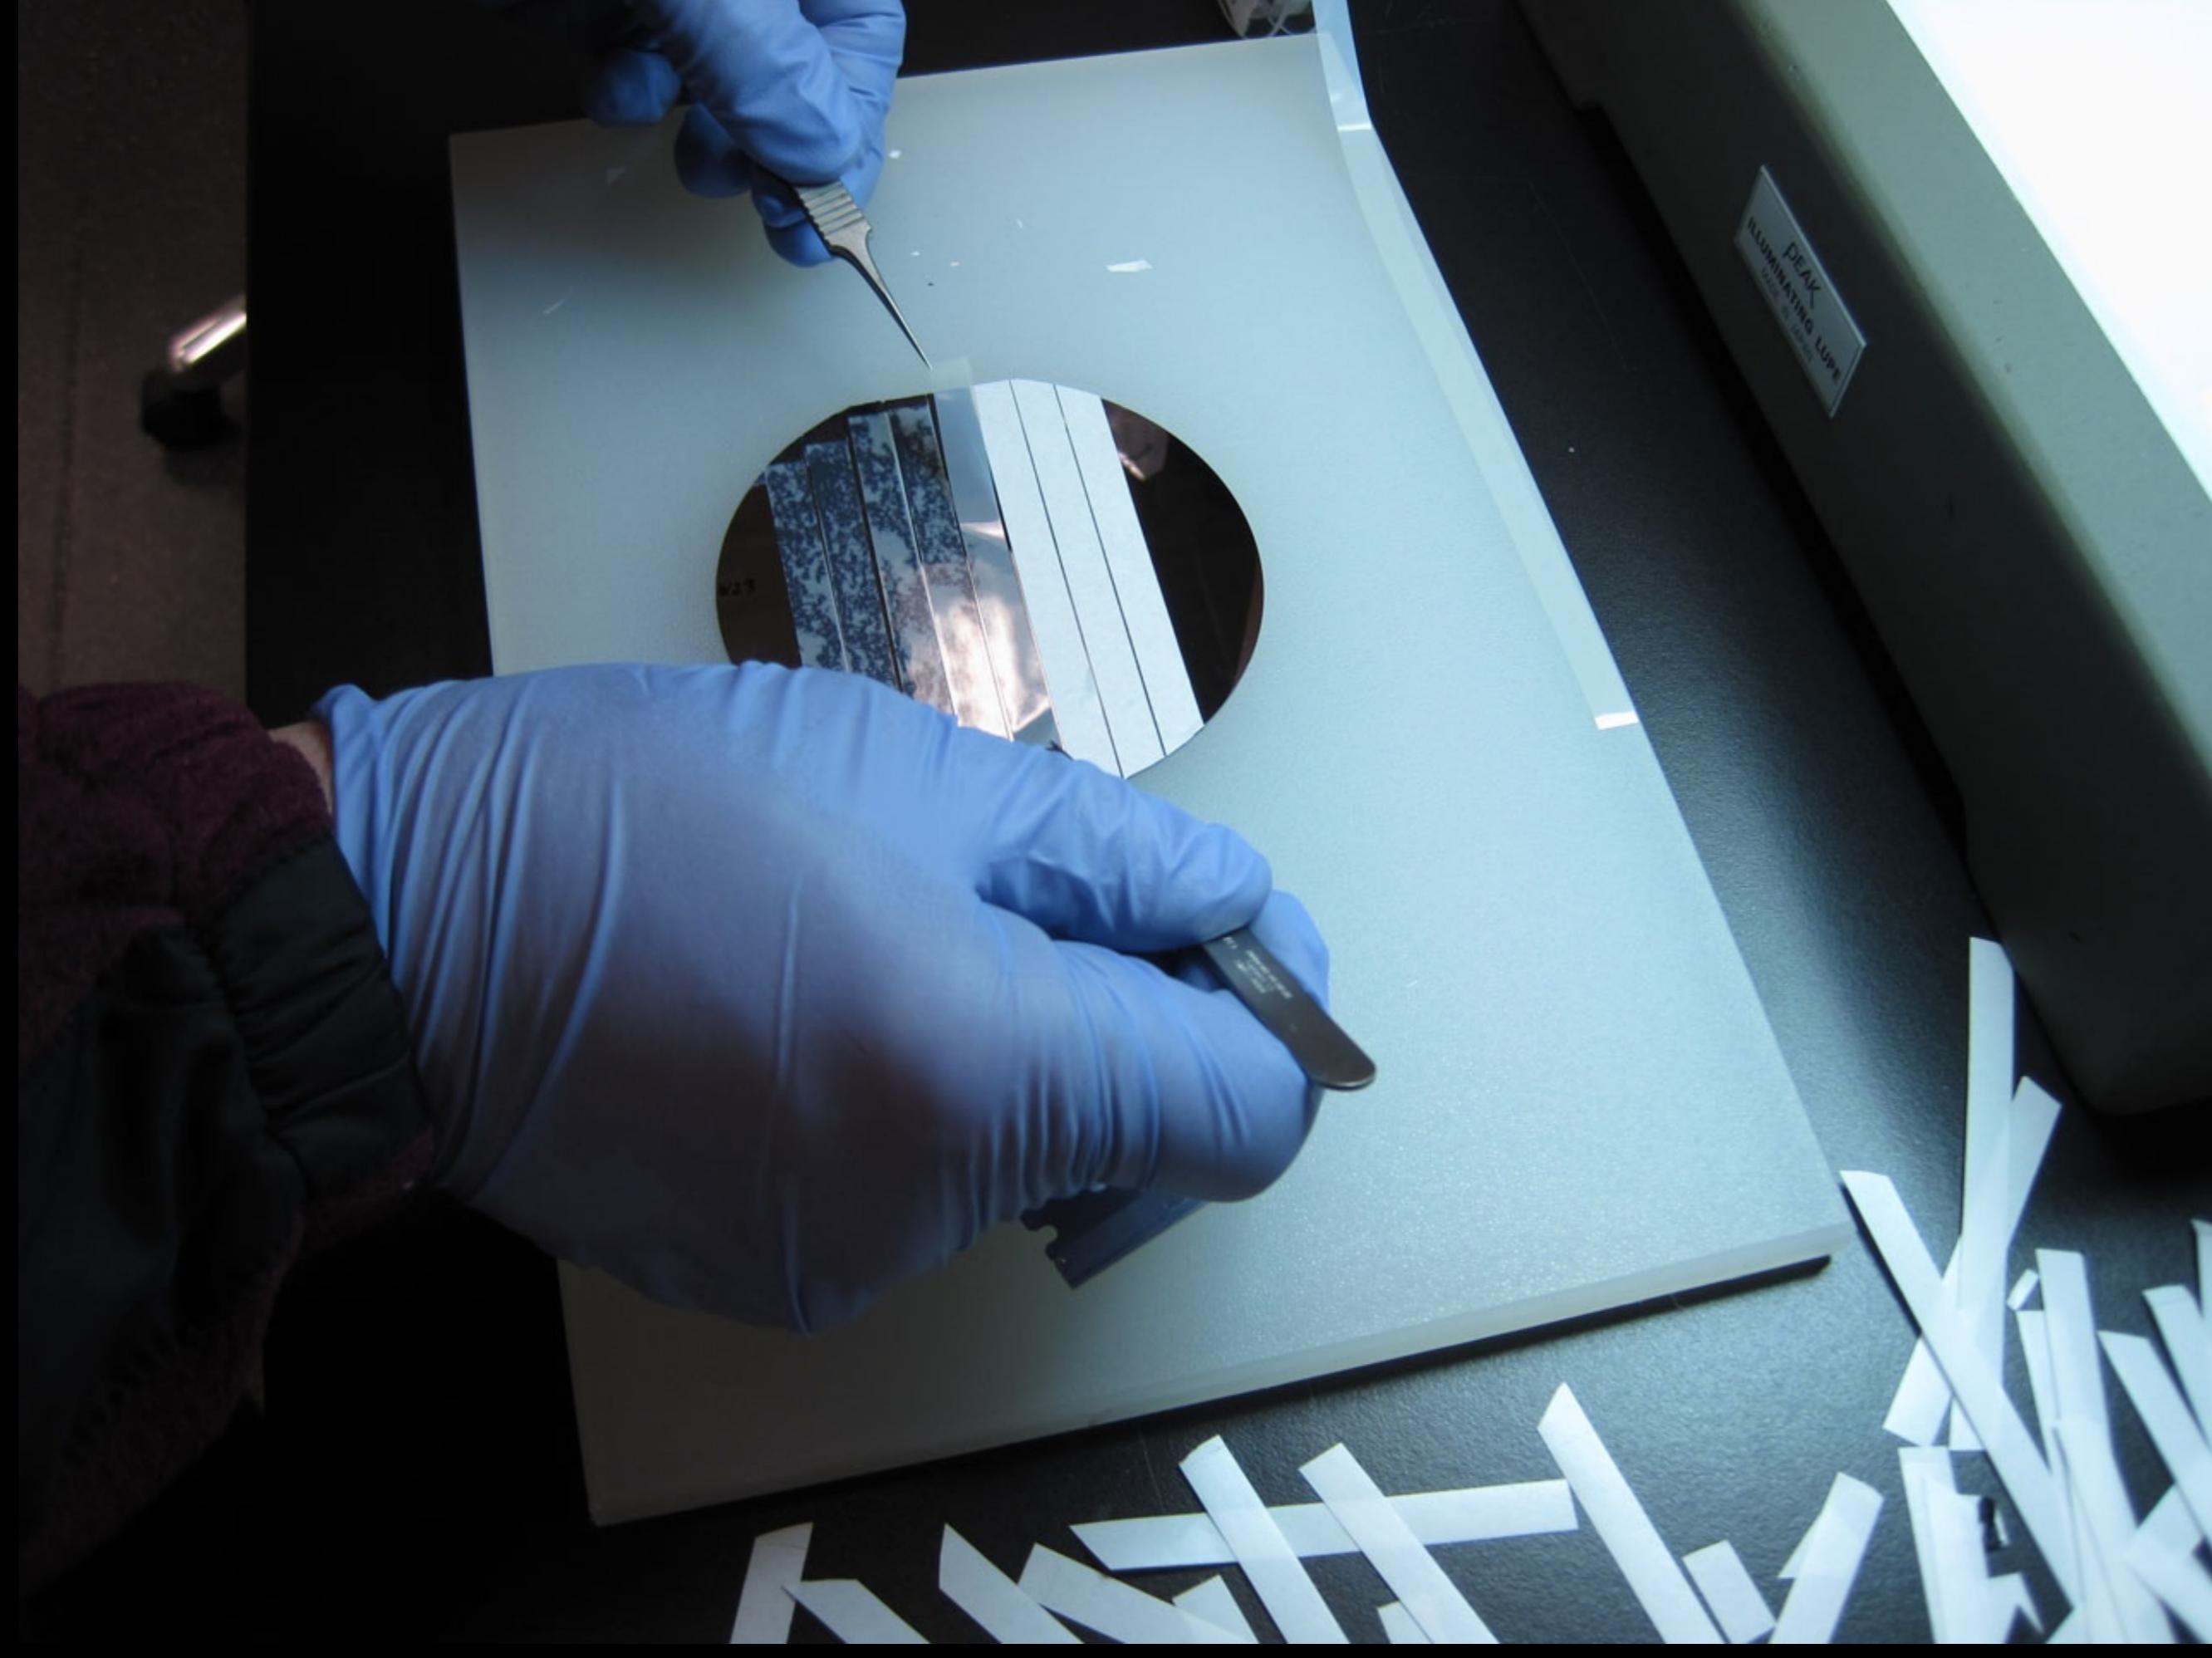


Inspiration for ATLUM

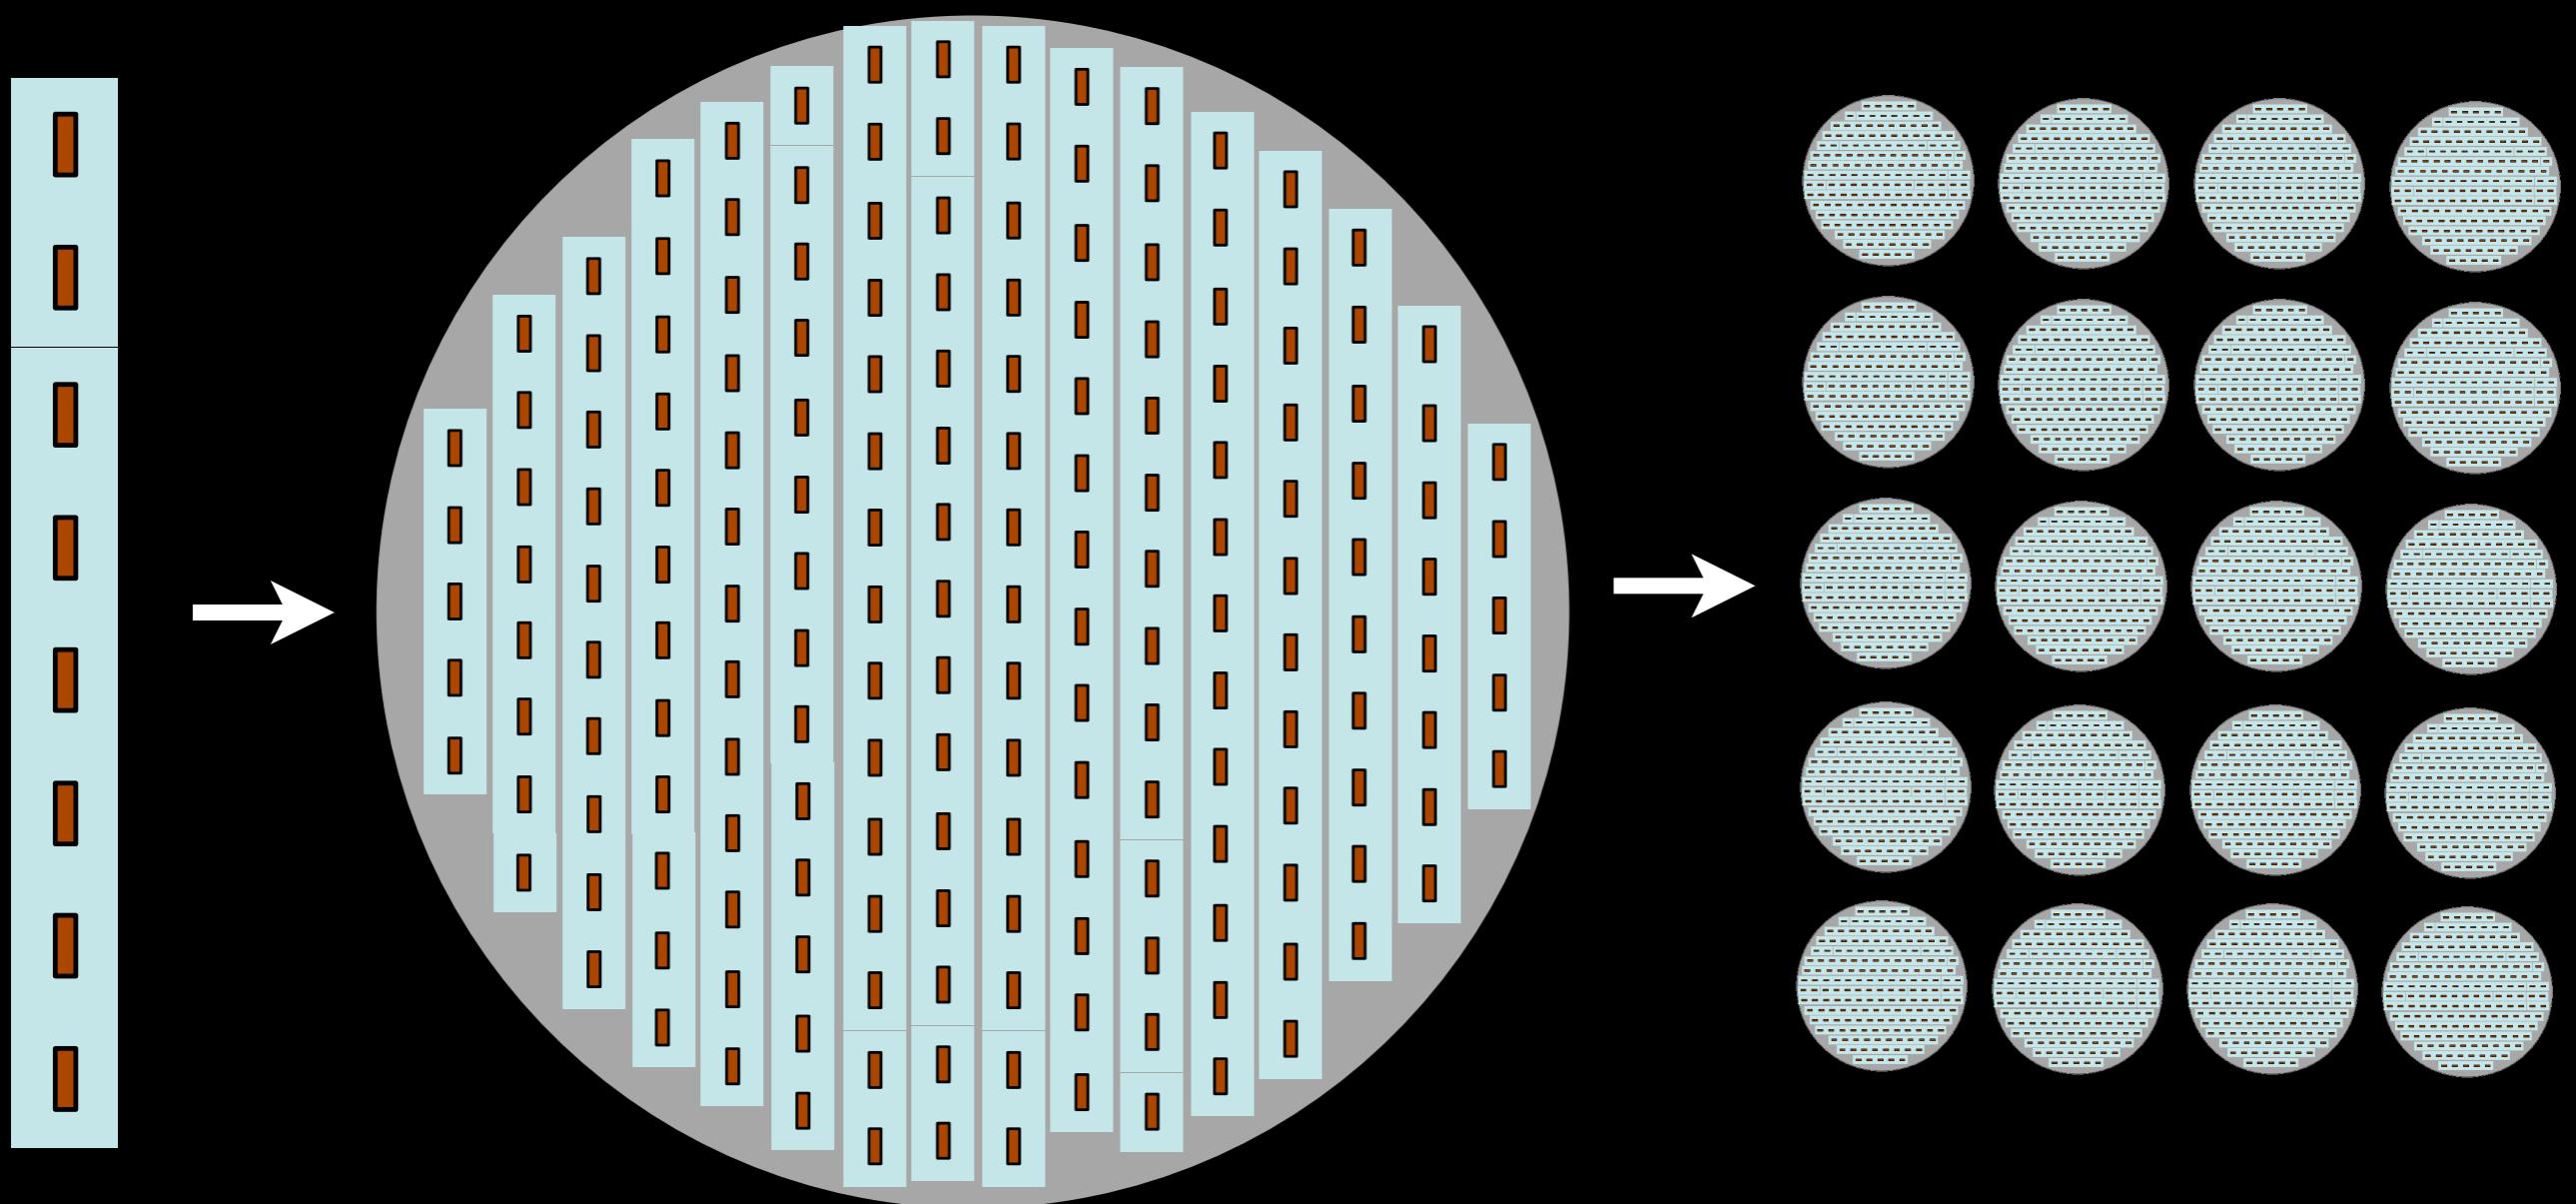


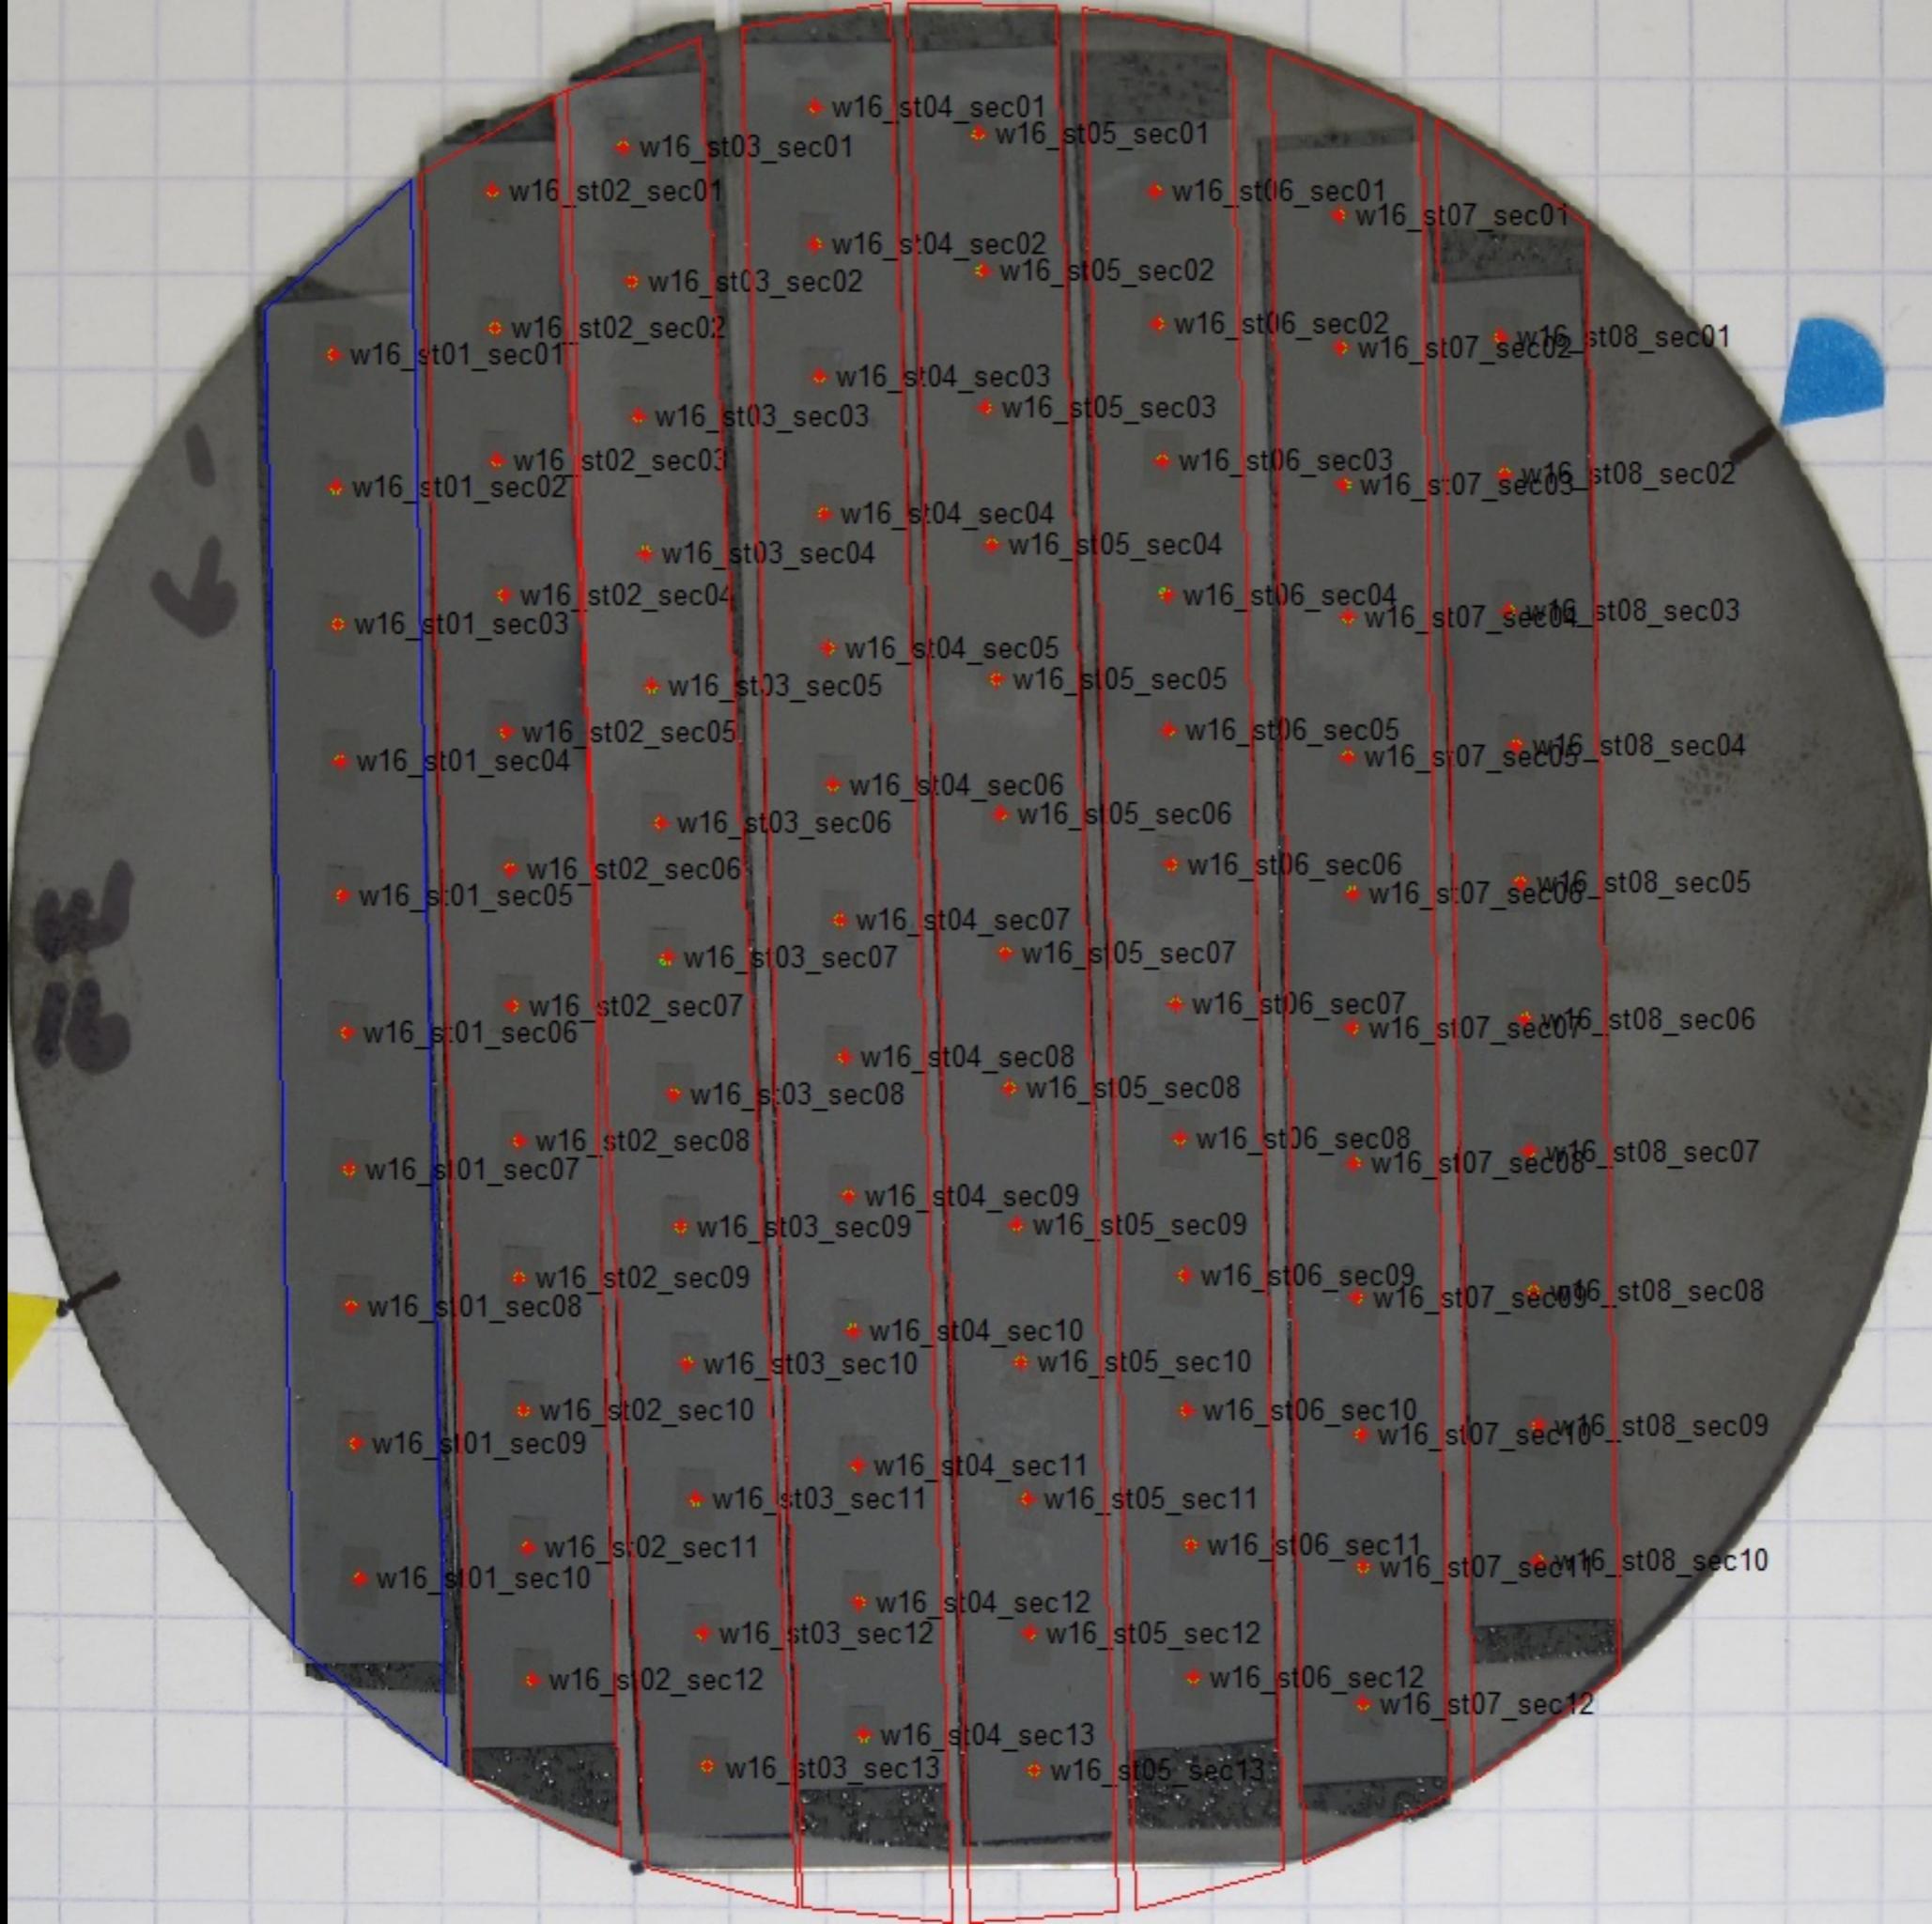


PEAK
ILLUMINATING LUPP
CO. INC. 1940



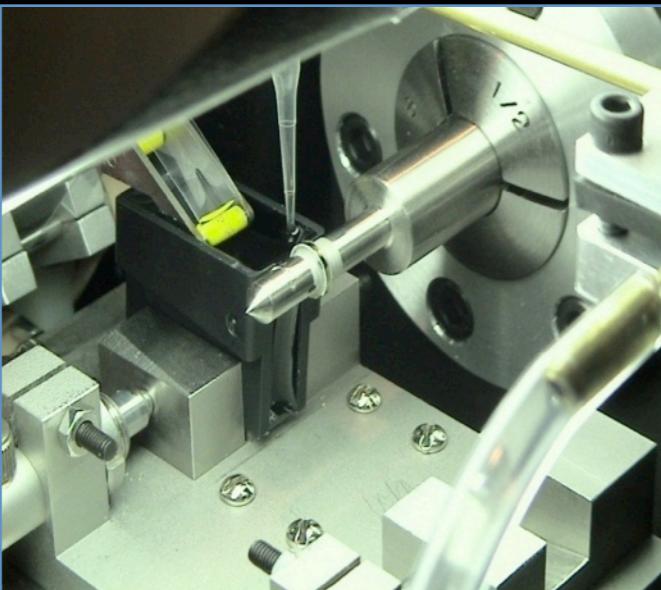
Ultrathin Section Libraries



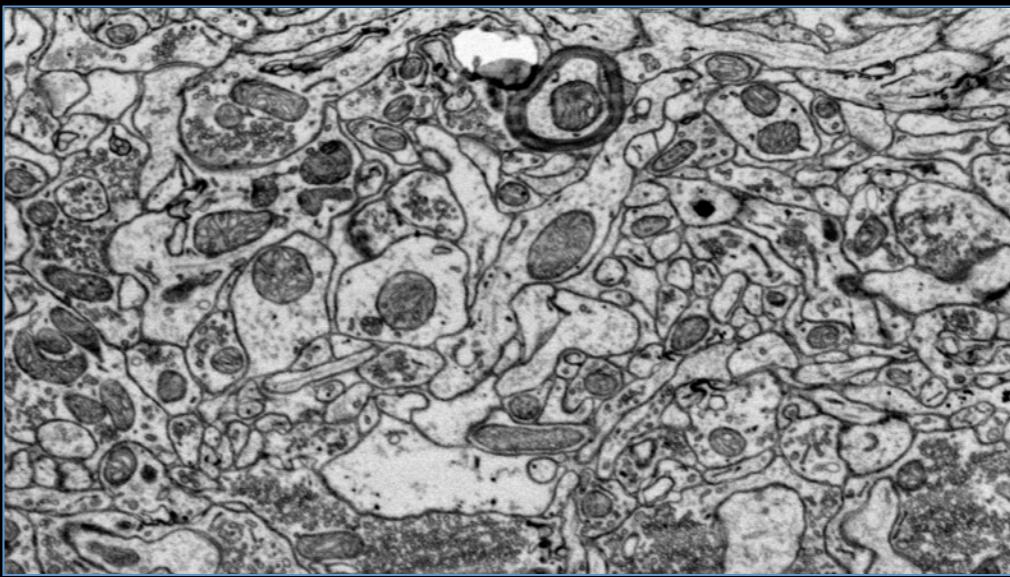




The Data Challenge



3 nm x 3 nm x 30 nm

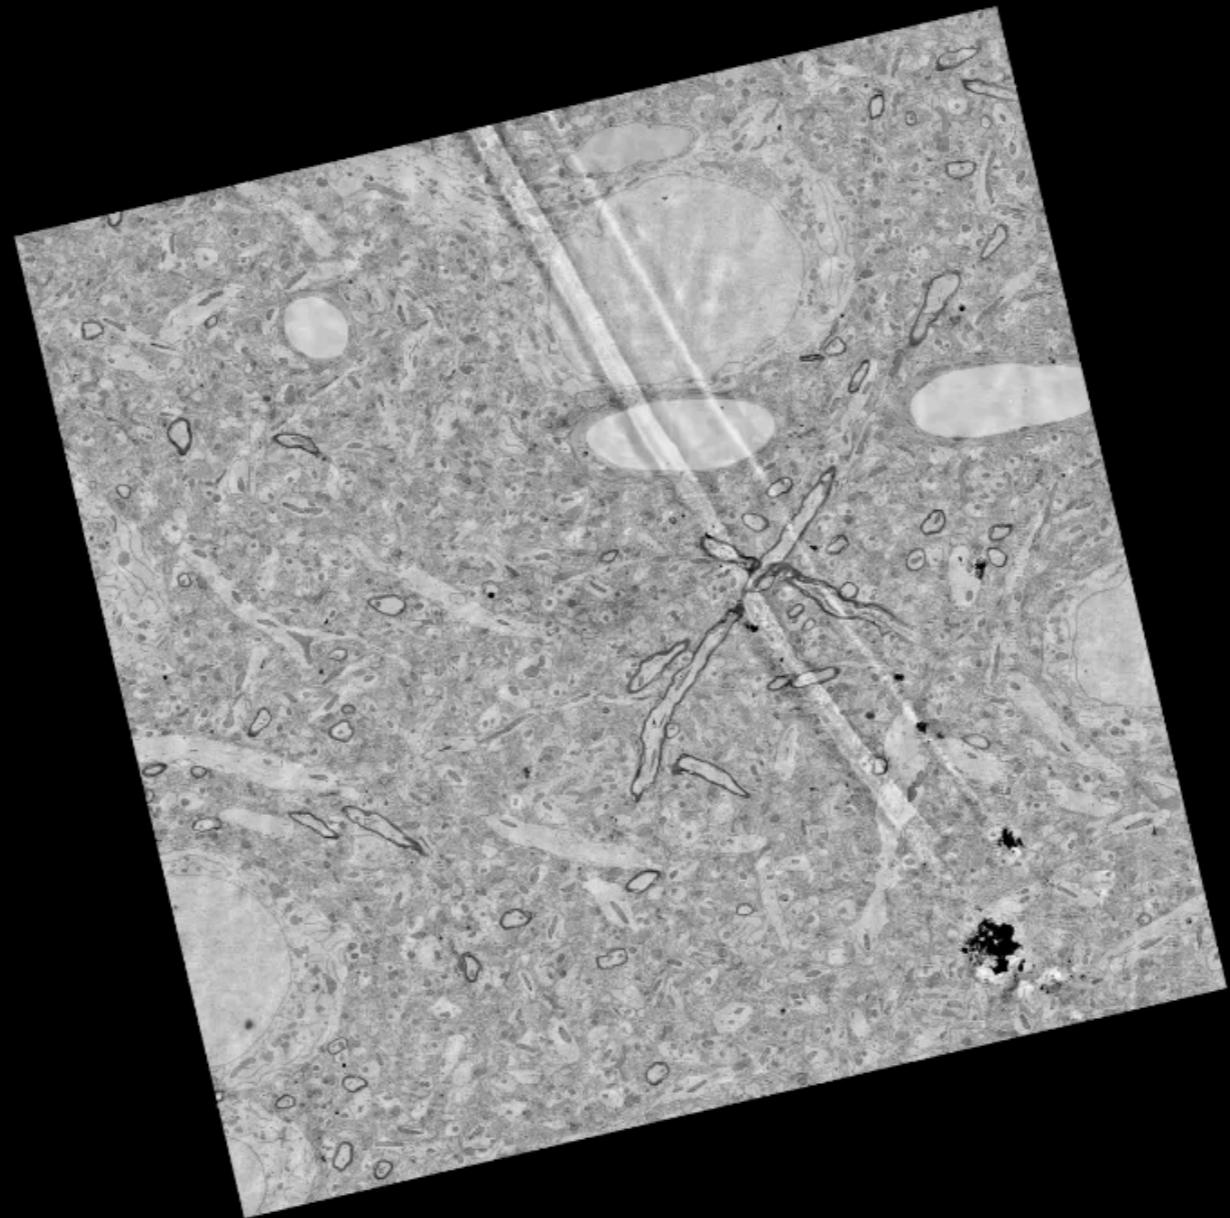


2M

1M

- Mouse thalamus $\sim=$ 1 mm³ $\sim=$ 1 petabyte

NeuroTrace



Interactive navigation of 1TB brain dataset

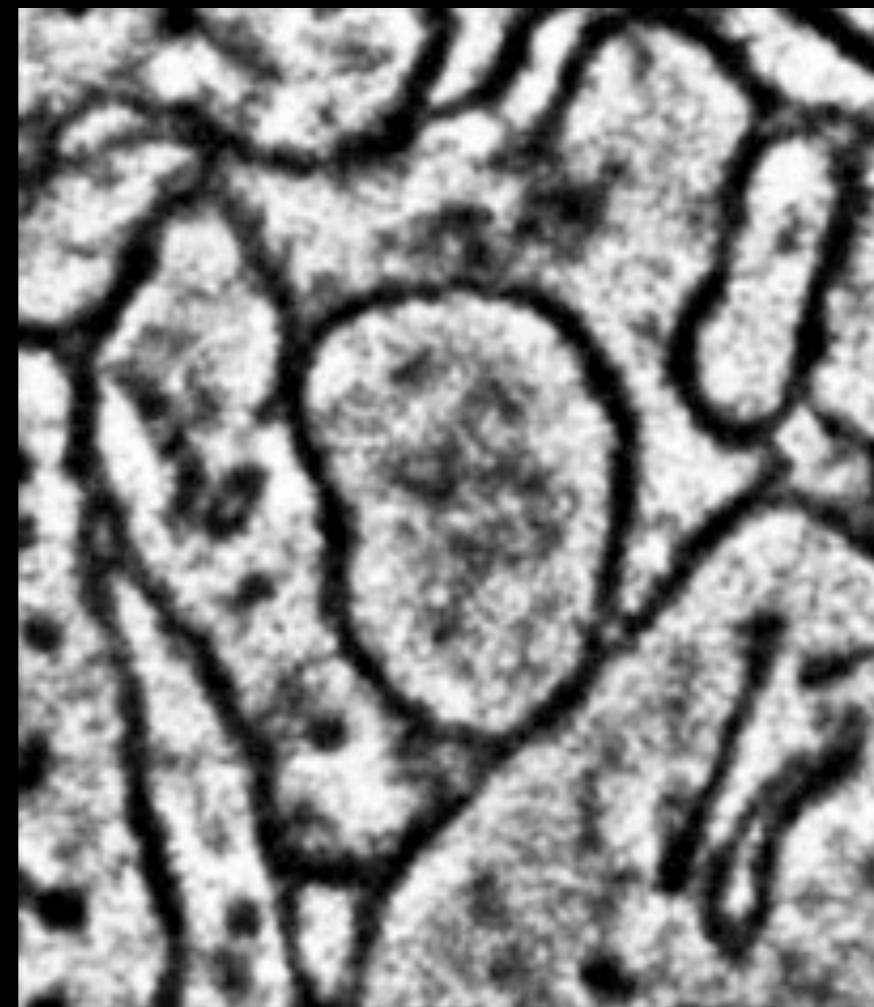
Active Ribbons

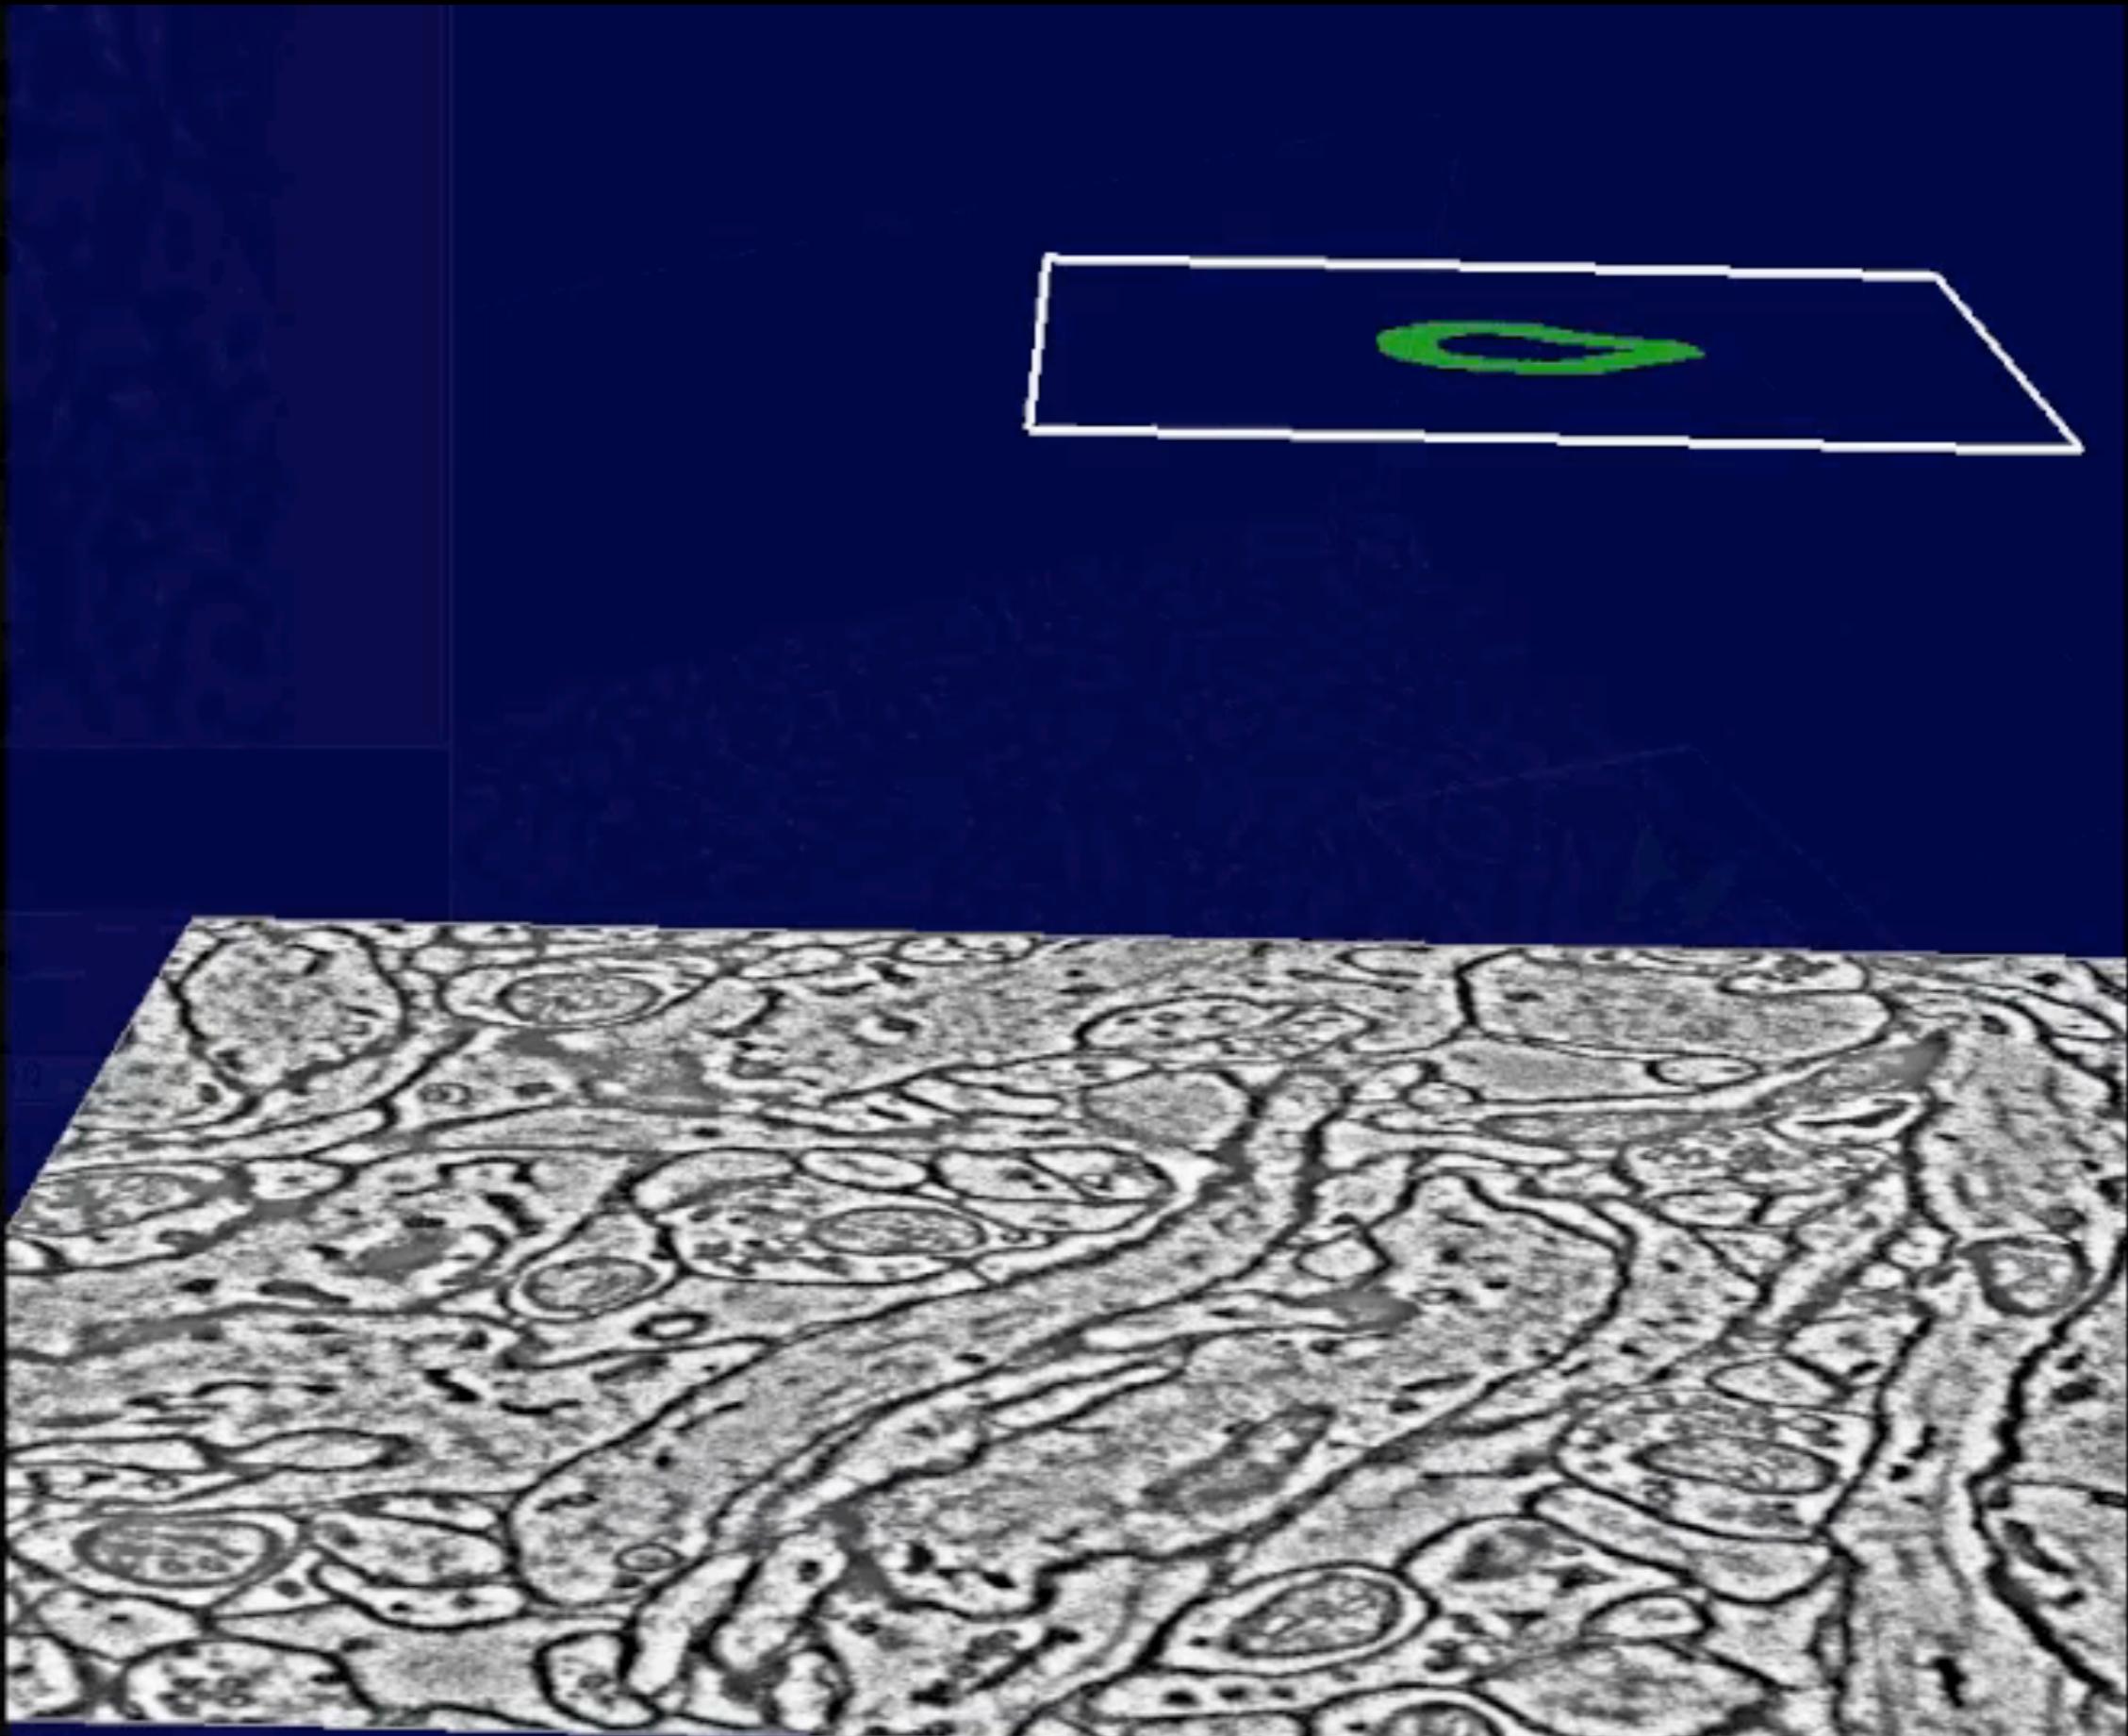


Section 1

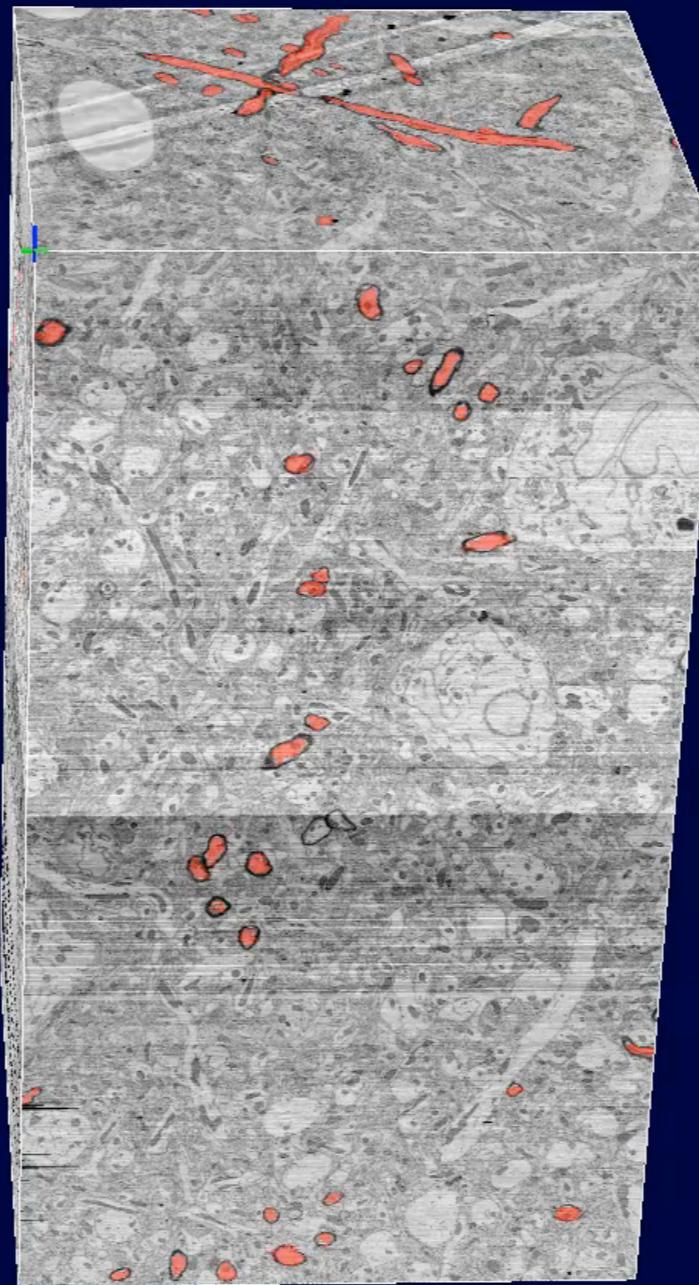
GPU Processing

- Active Ribbons on GPUs: 23x speedup

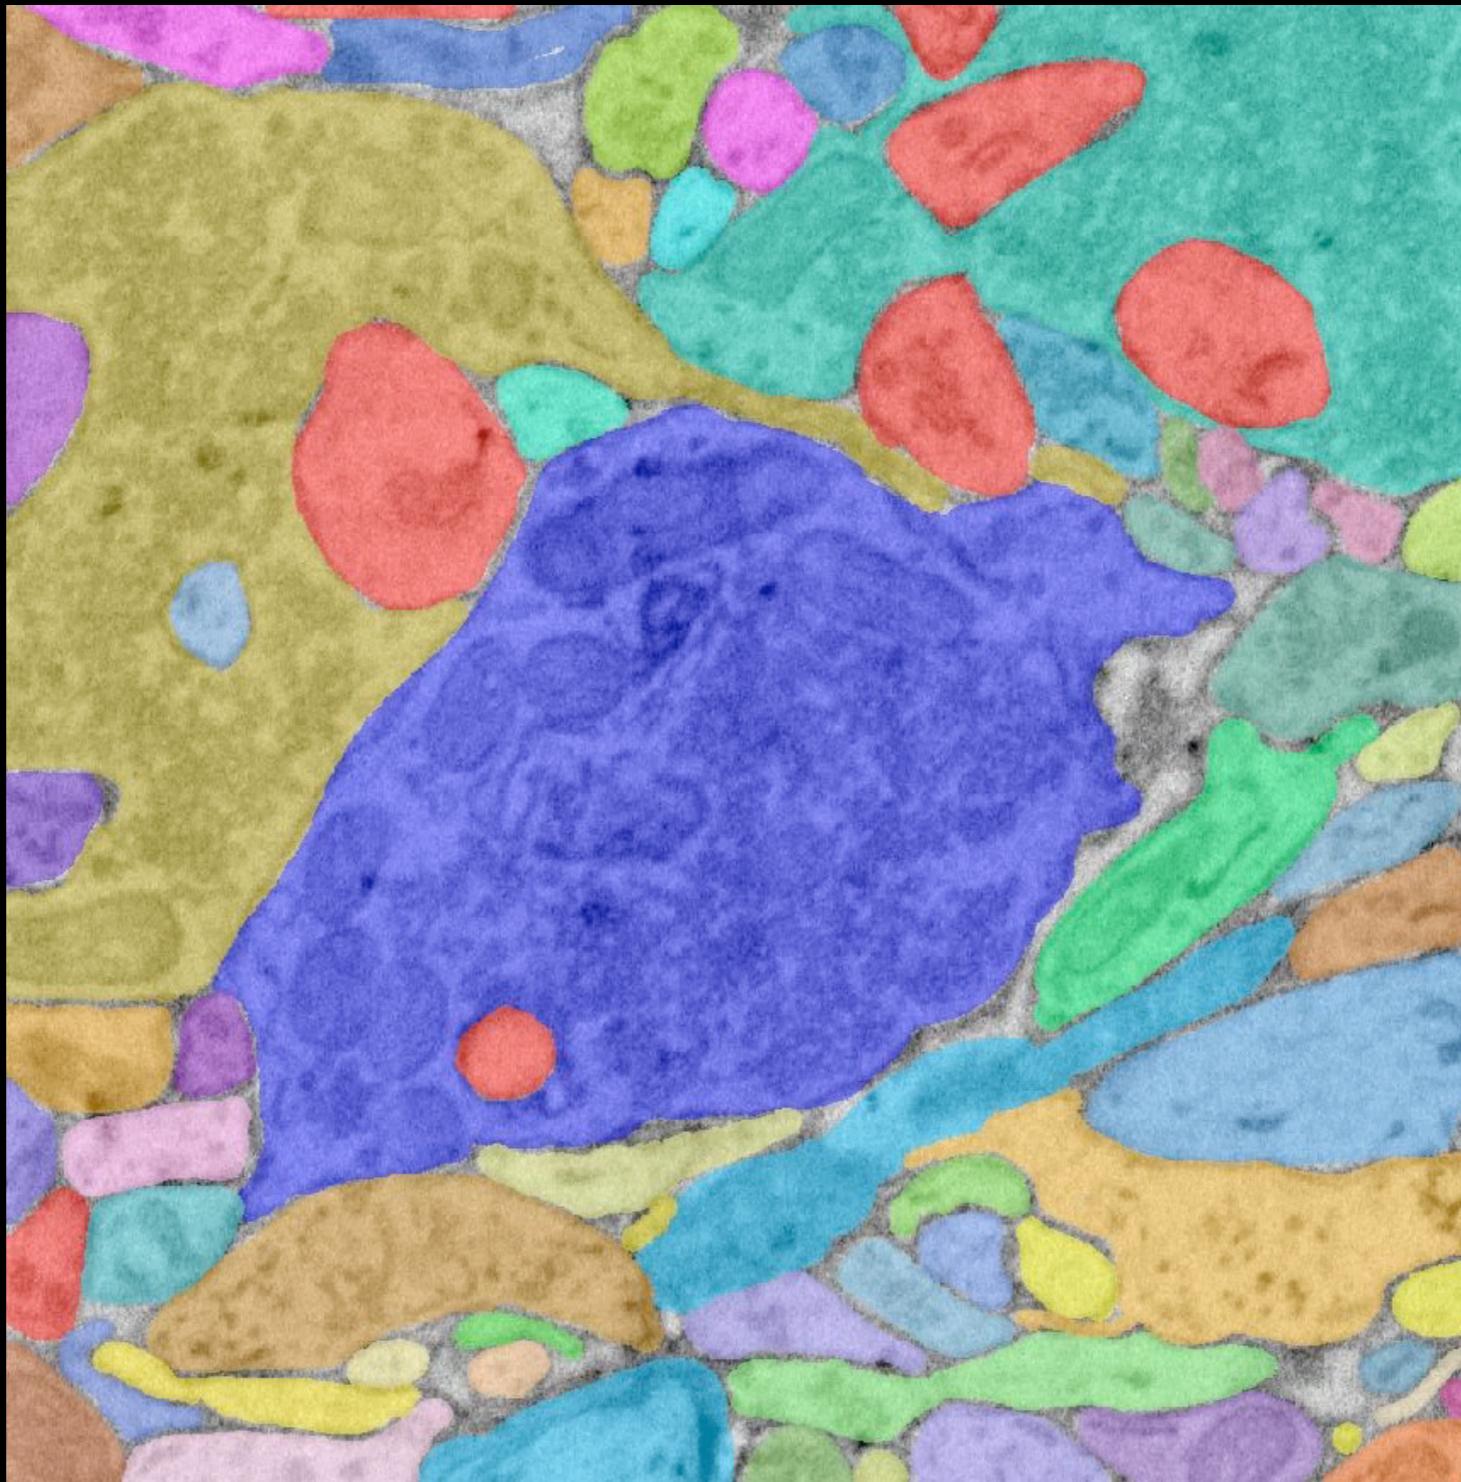




Cell position
Stagnation
Bifurc. Node
IMT 1000nm
Growth front
Tissue type
Tissue type
Run Site
Ogre
Endurance
Visual area
Volume
L-leaf
Leaf



“Temporal” Coherence

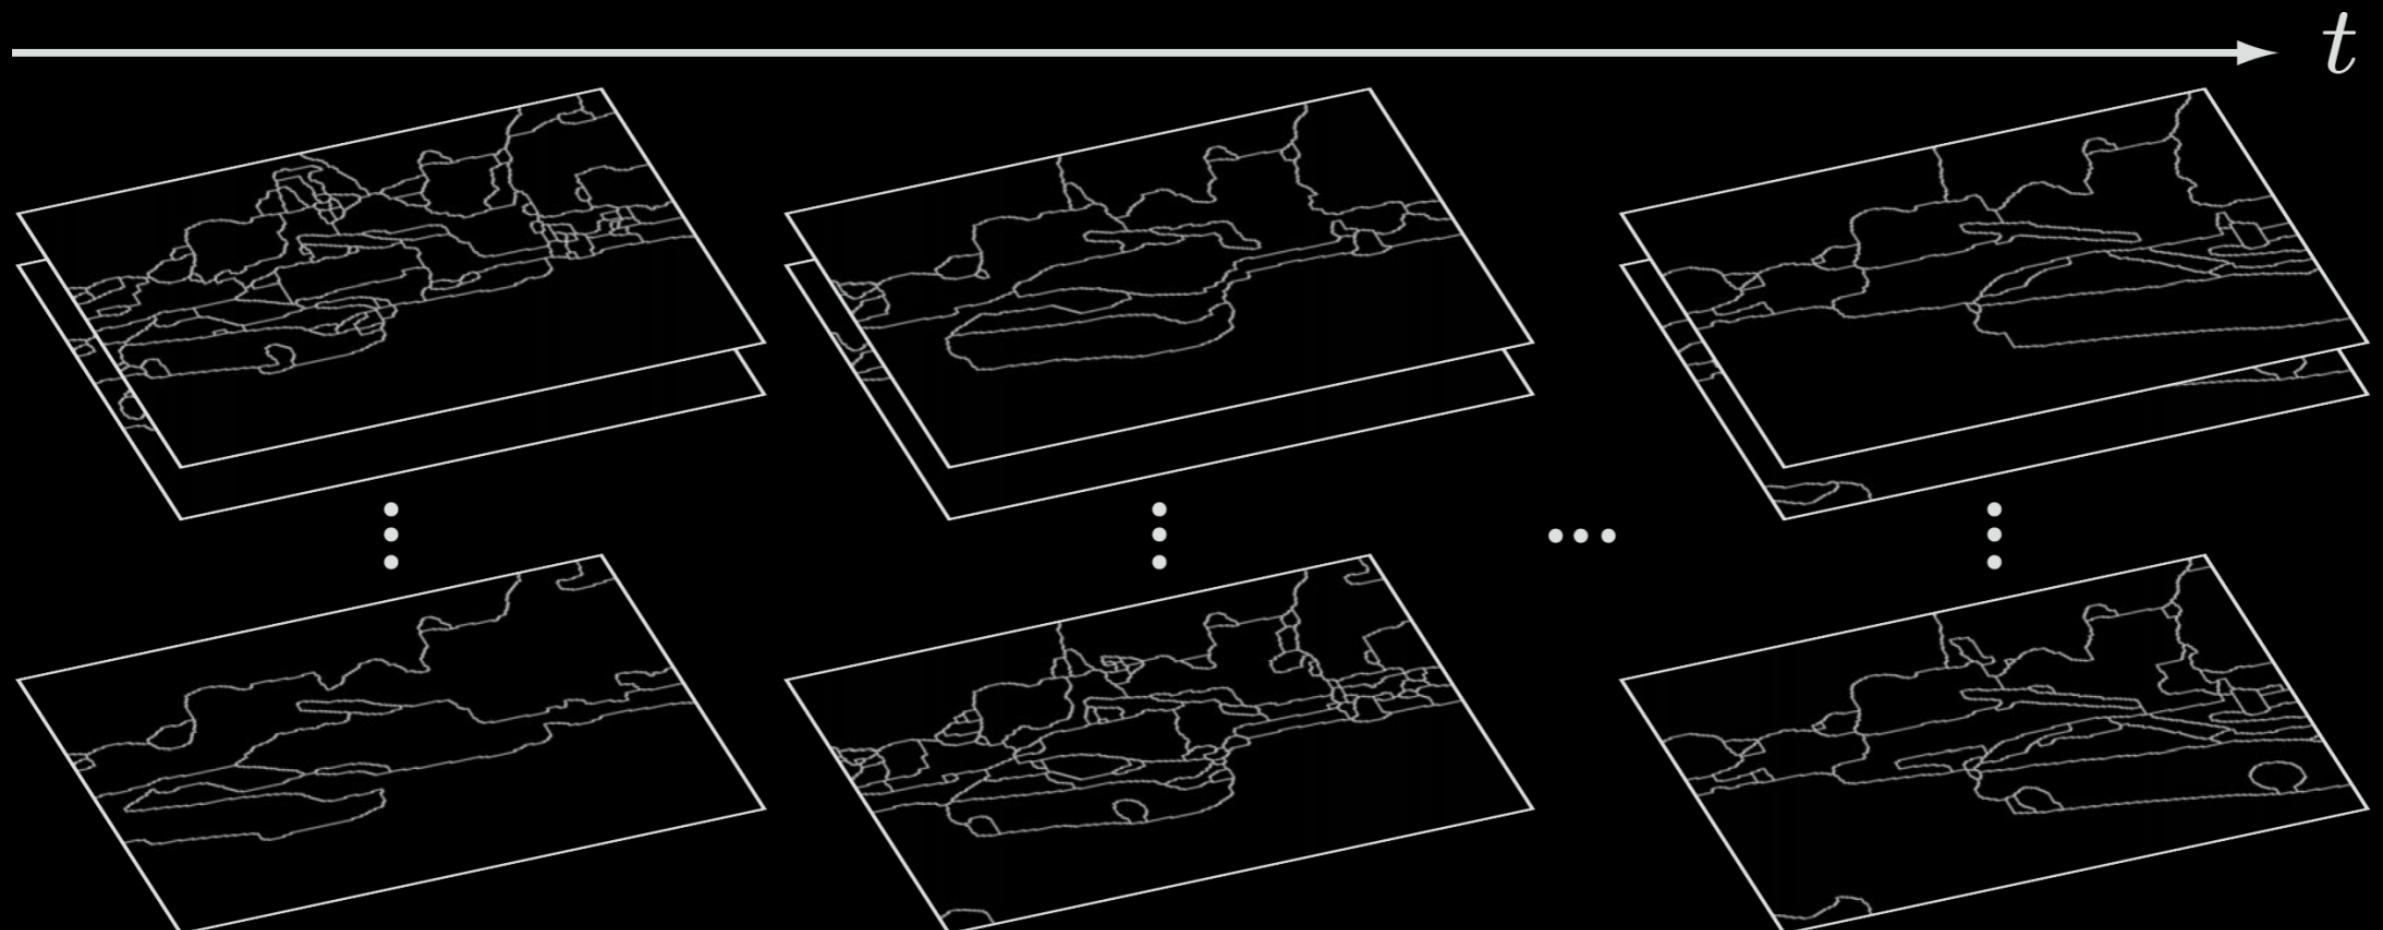


D. Berger, MIT

Multiple Hypothesis Video Segmentation

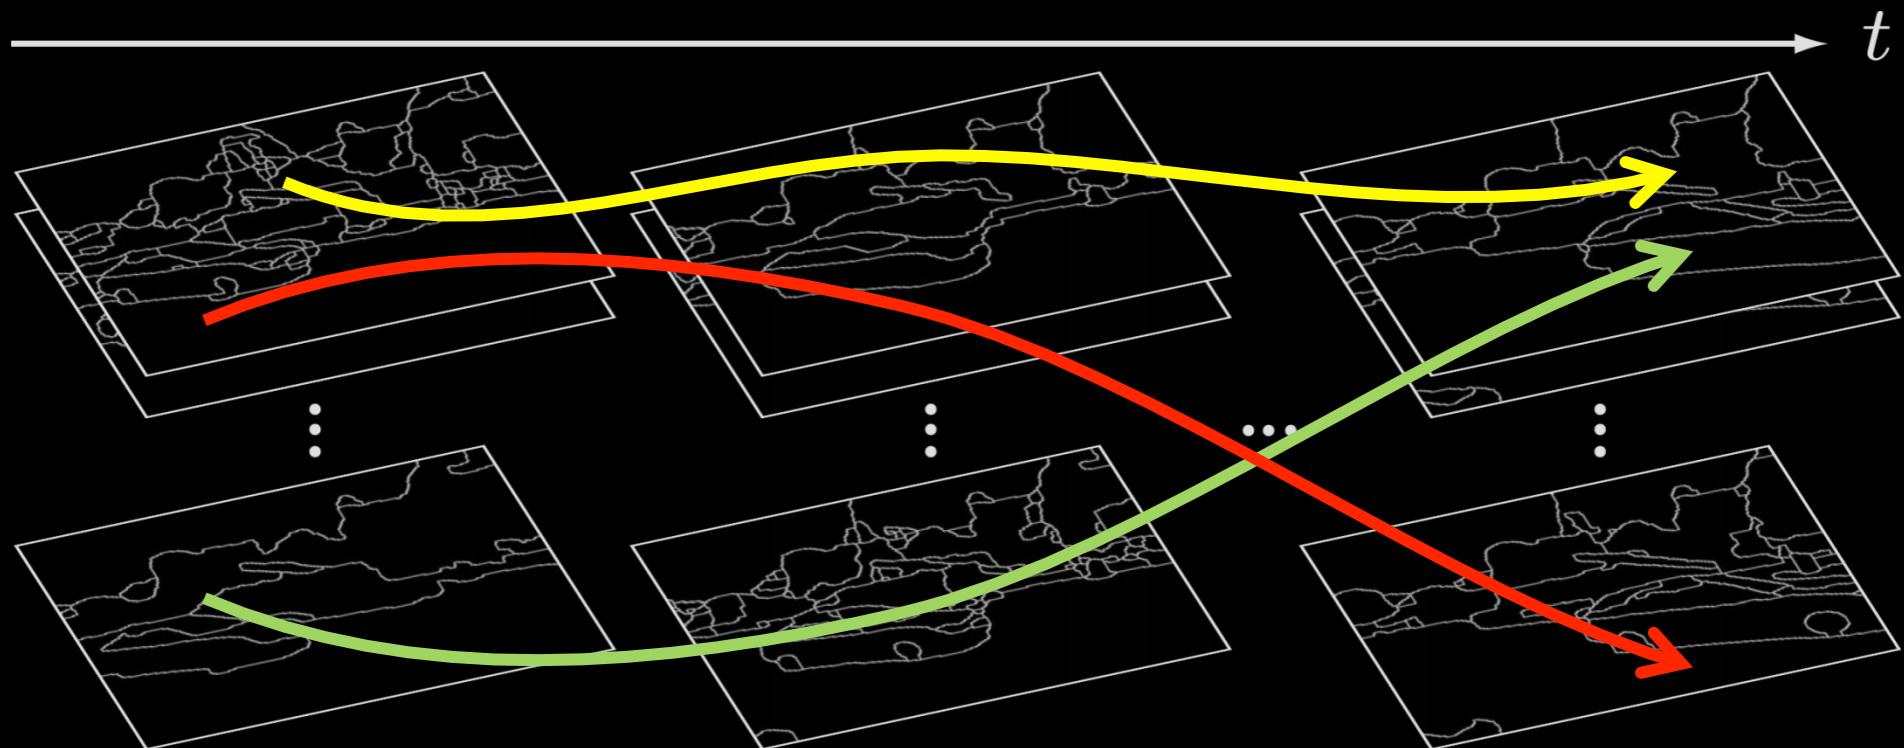


Multiple Hypothesis Video Segmentation



Multiple Possible Segmentations per Frame
(gPb-OWT-UCM etc.)

Multiple Hypothesis Video Segmentation



Enumerate multiple trajectory hypotheses
Solve with higher-order MRF & graph cuts

Qualitative Evaluation

Original

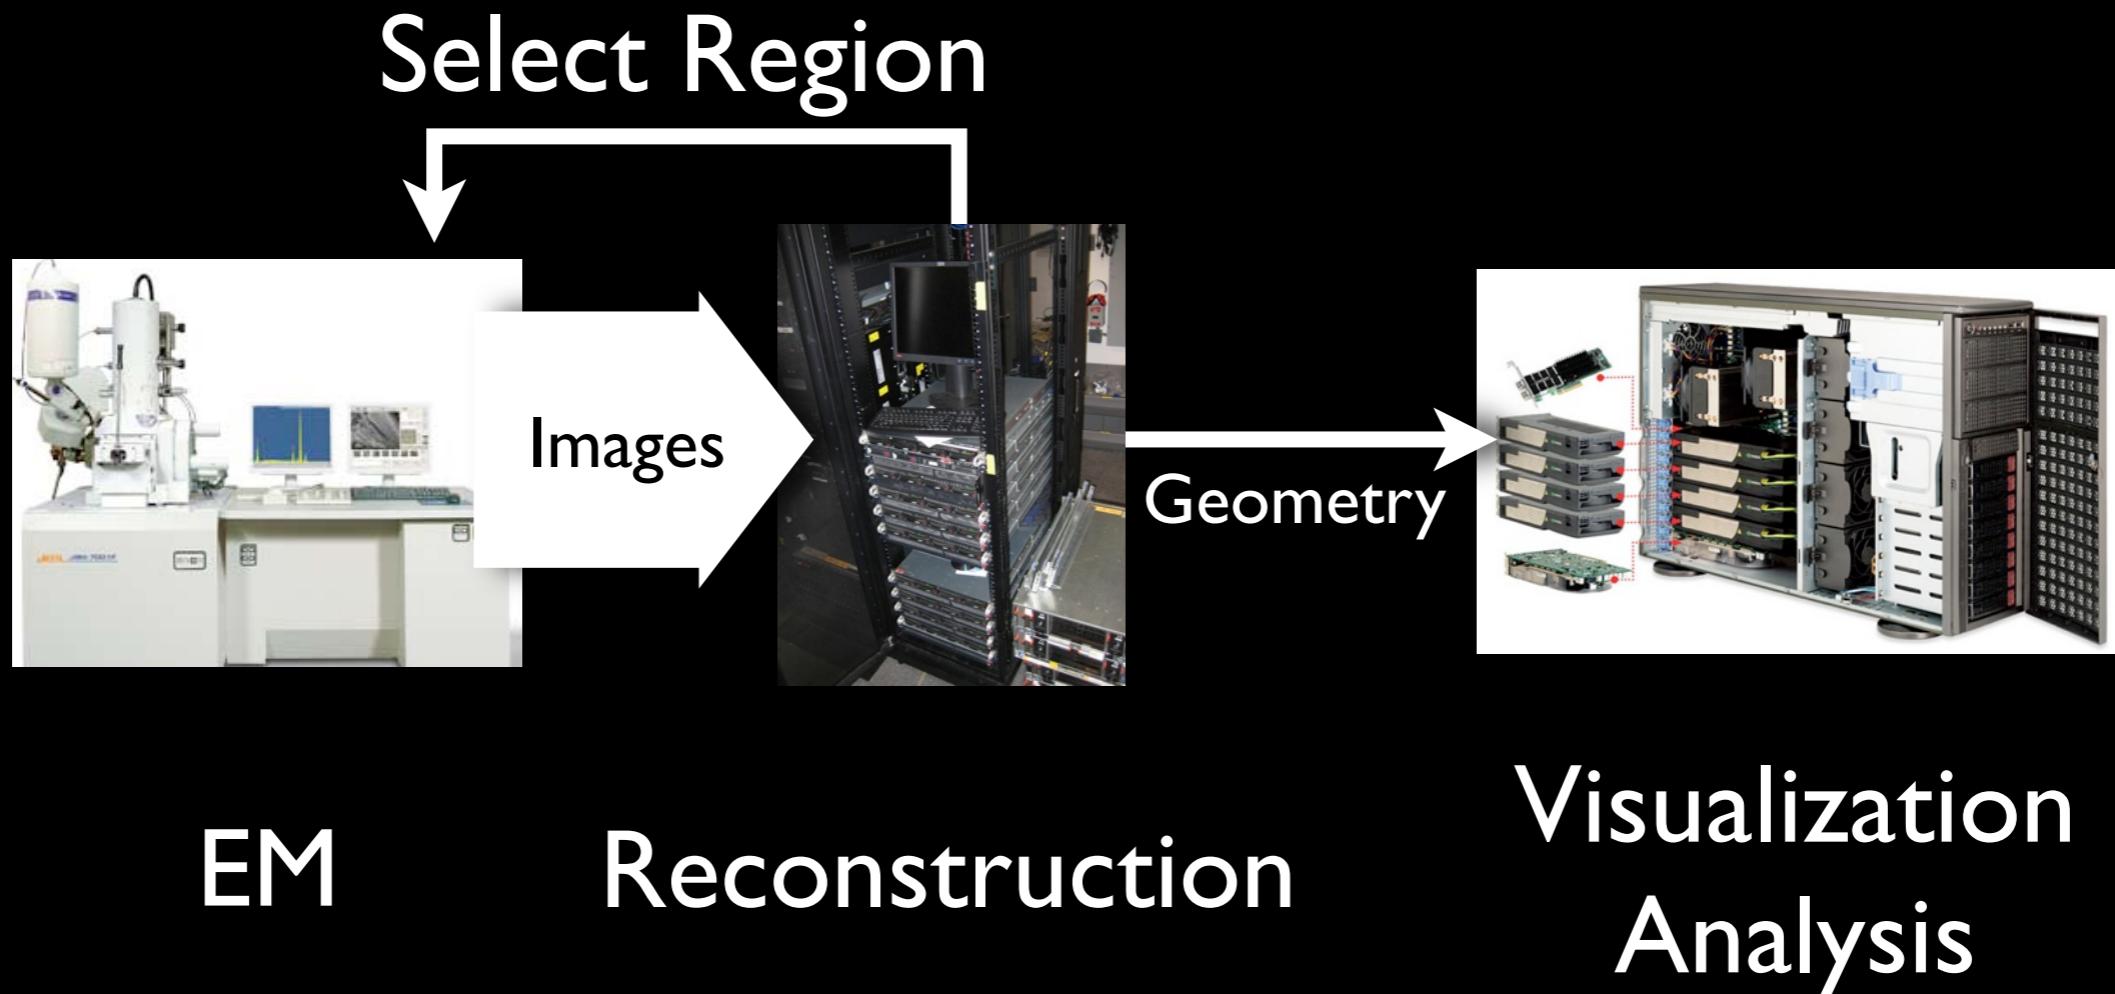


MHVS

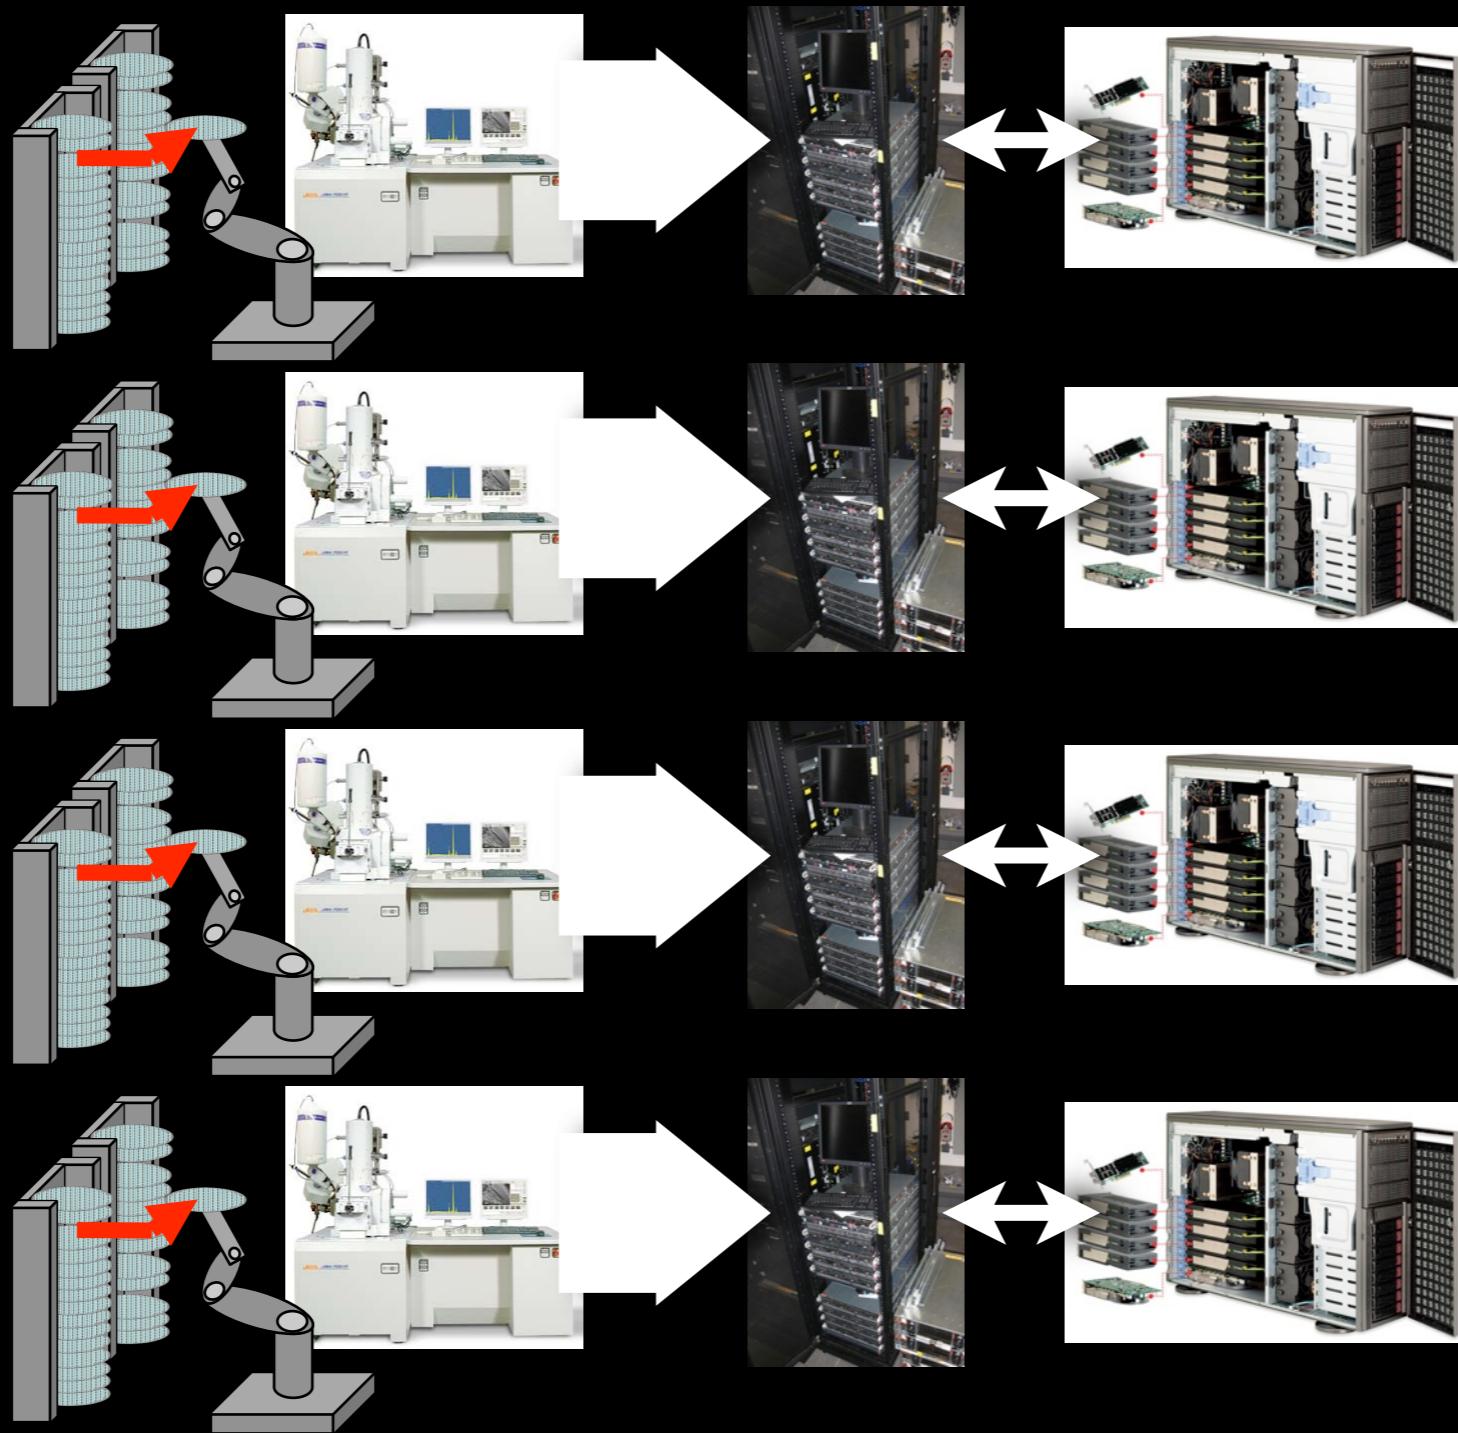
Frame-to-frame
similar to [5] (CVPR 08)



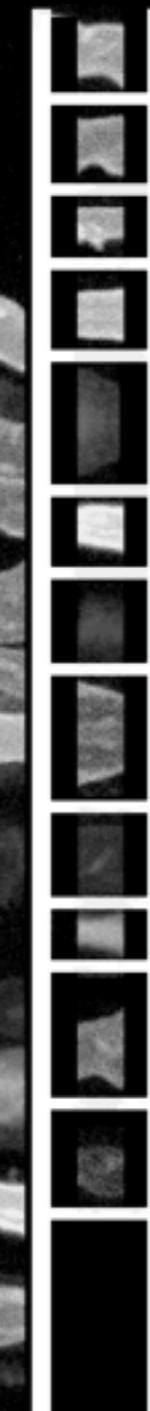
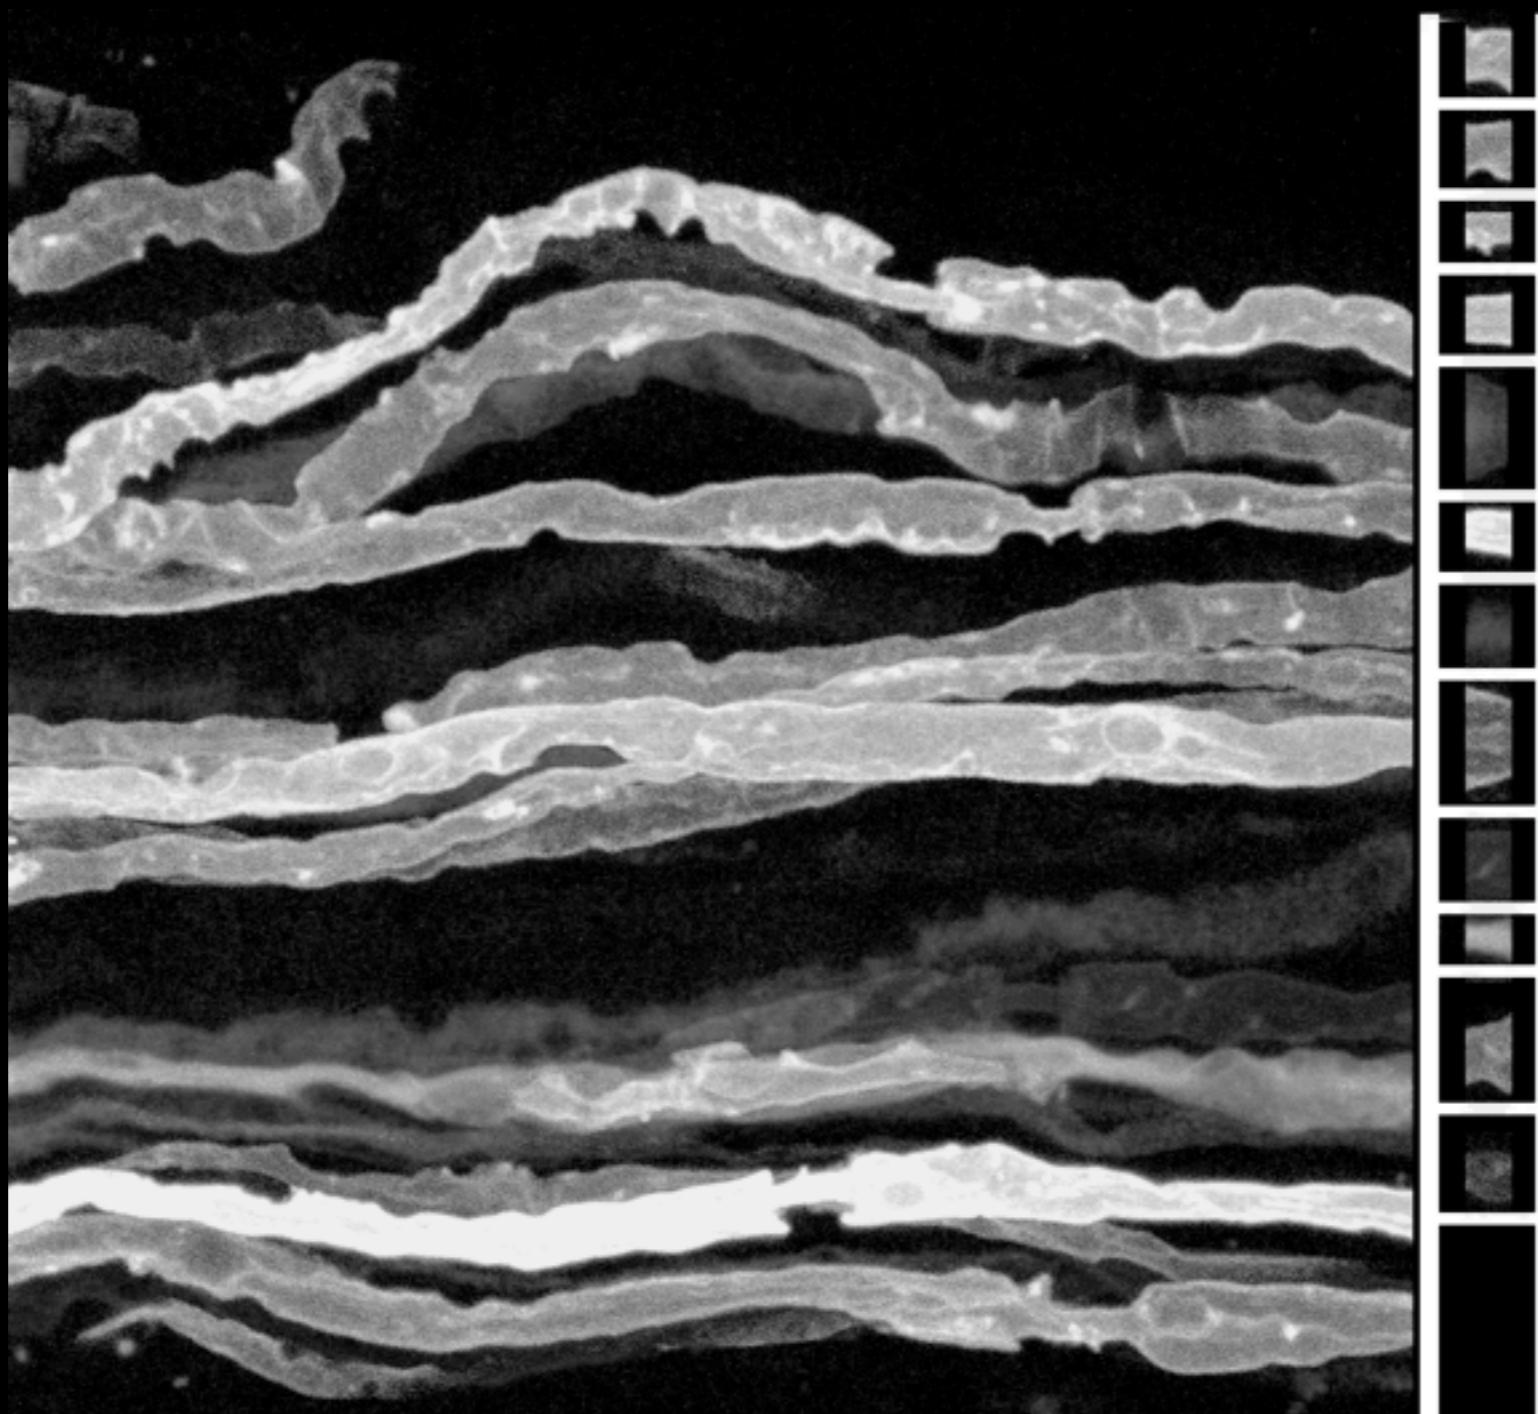
High-Throughput Neuroscience



Scalable

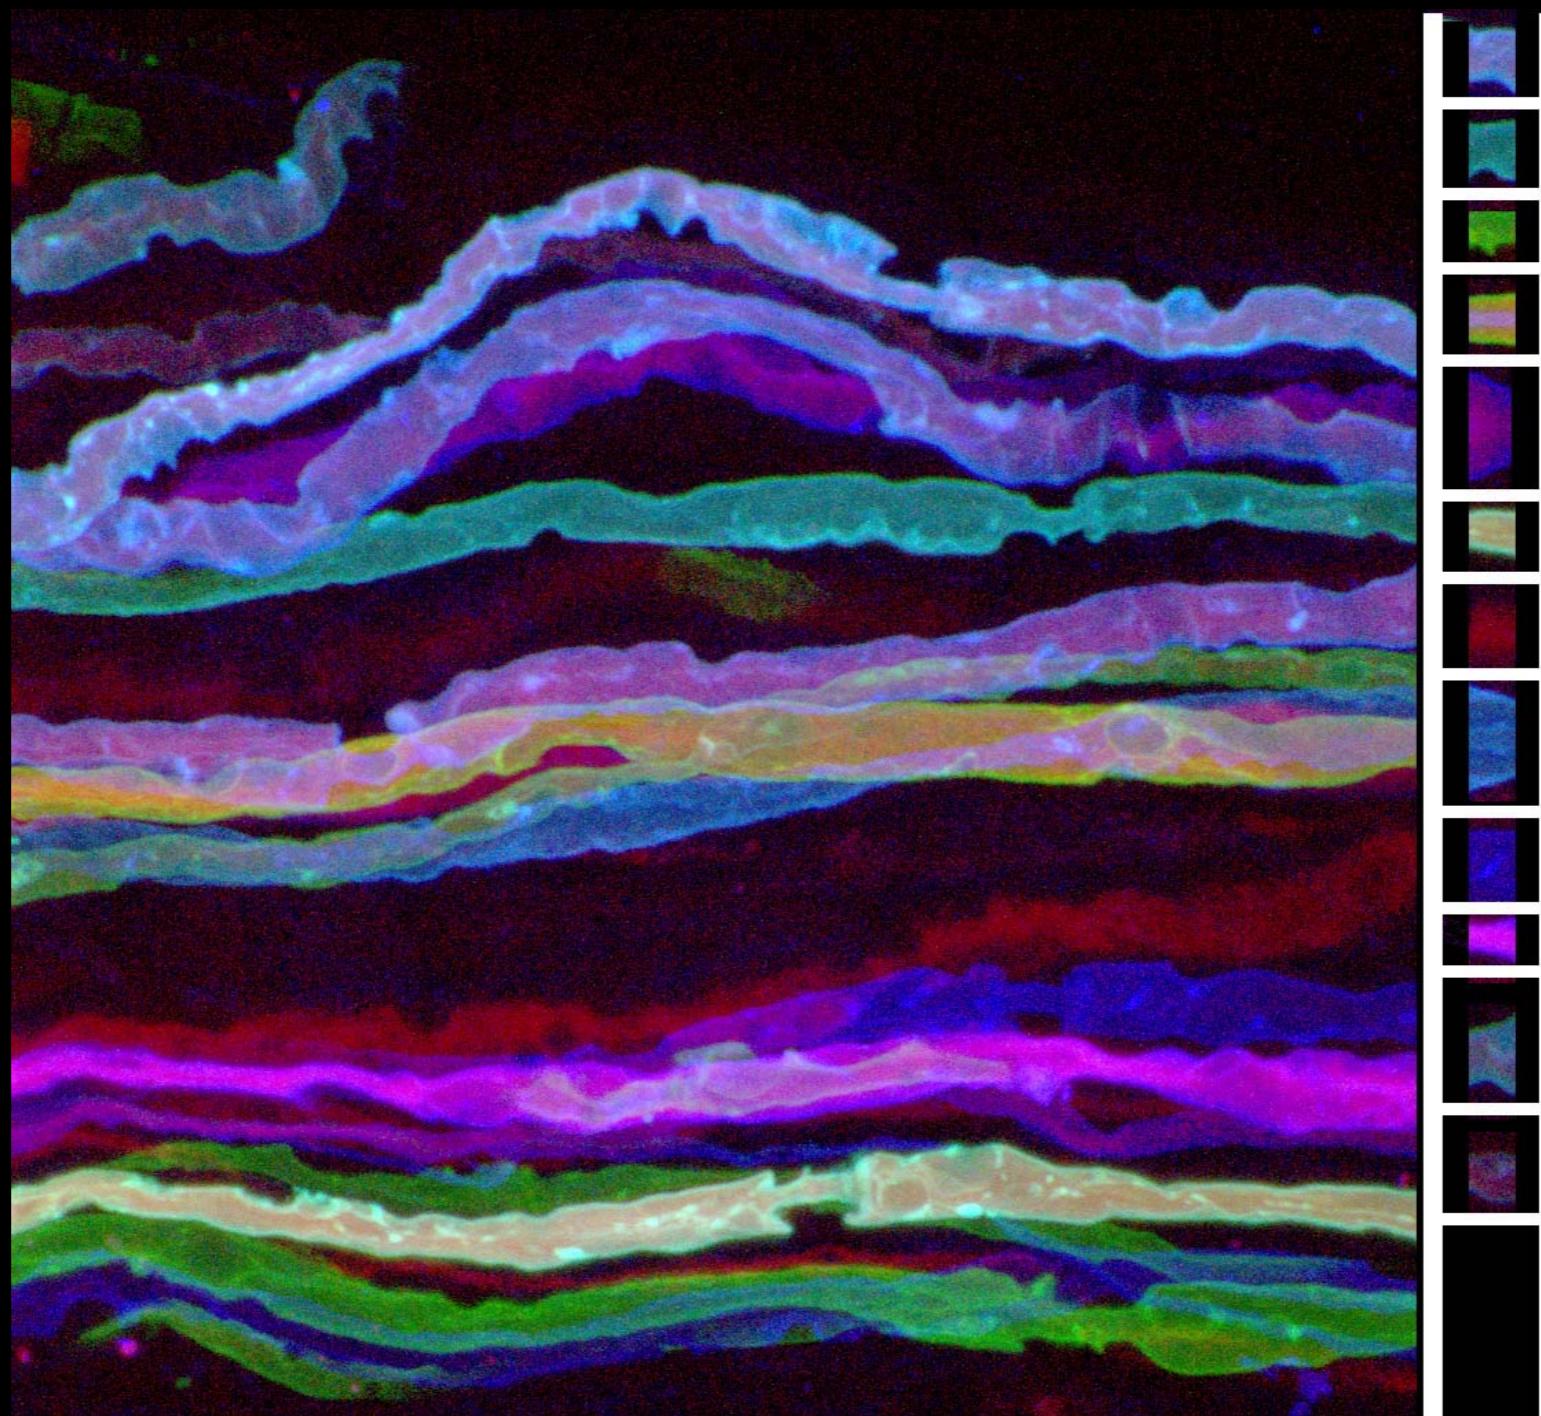


The Future



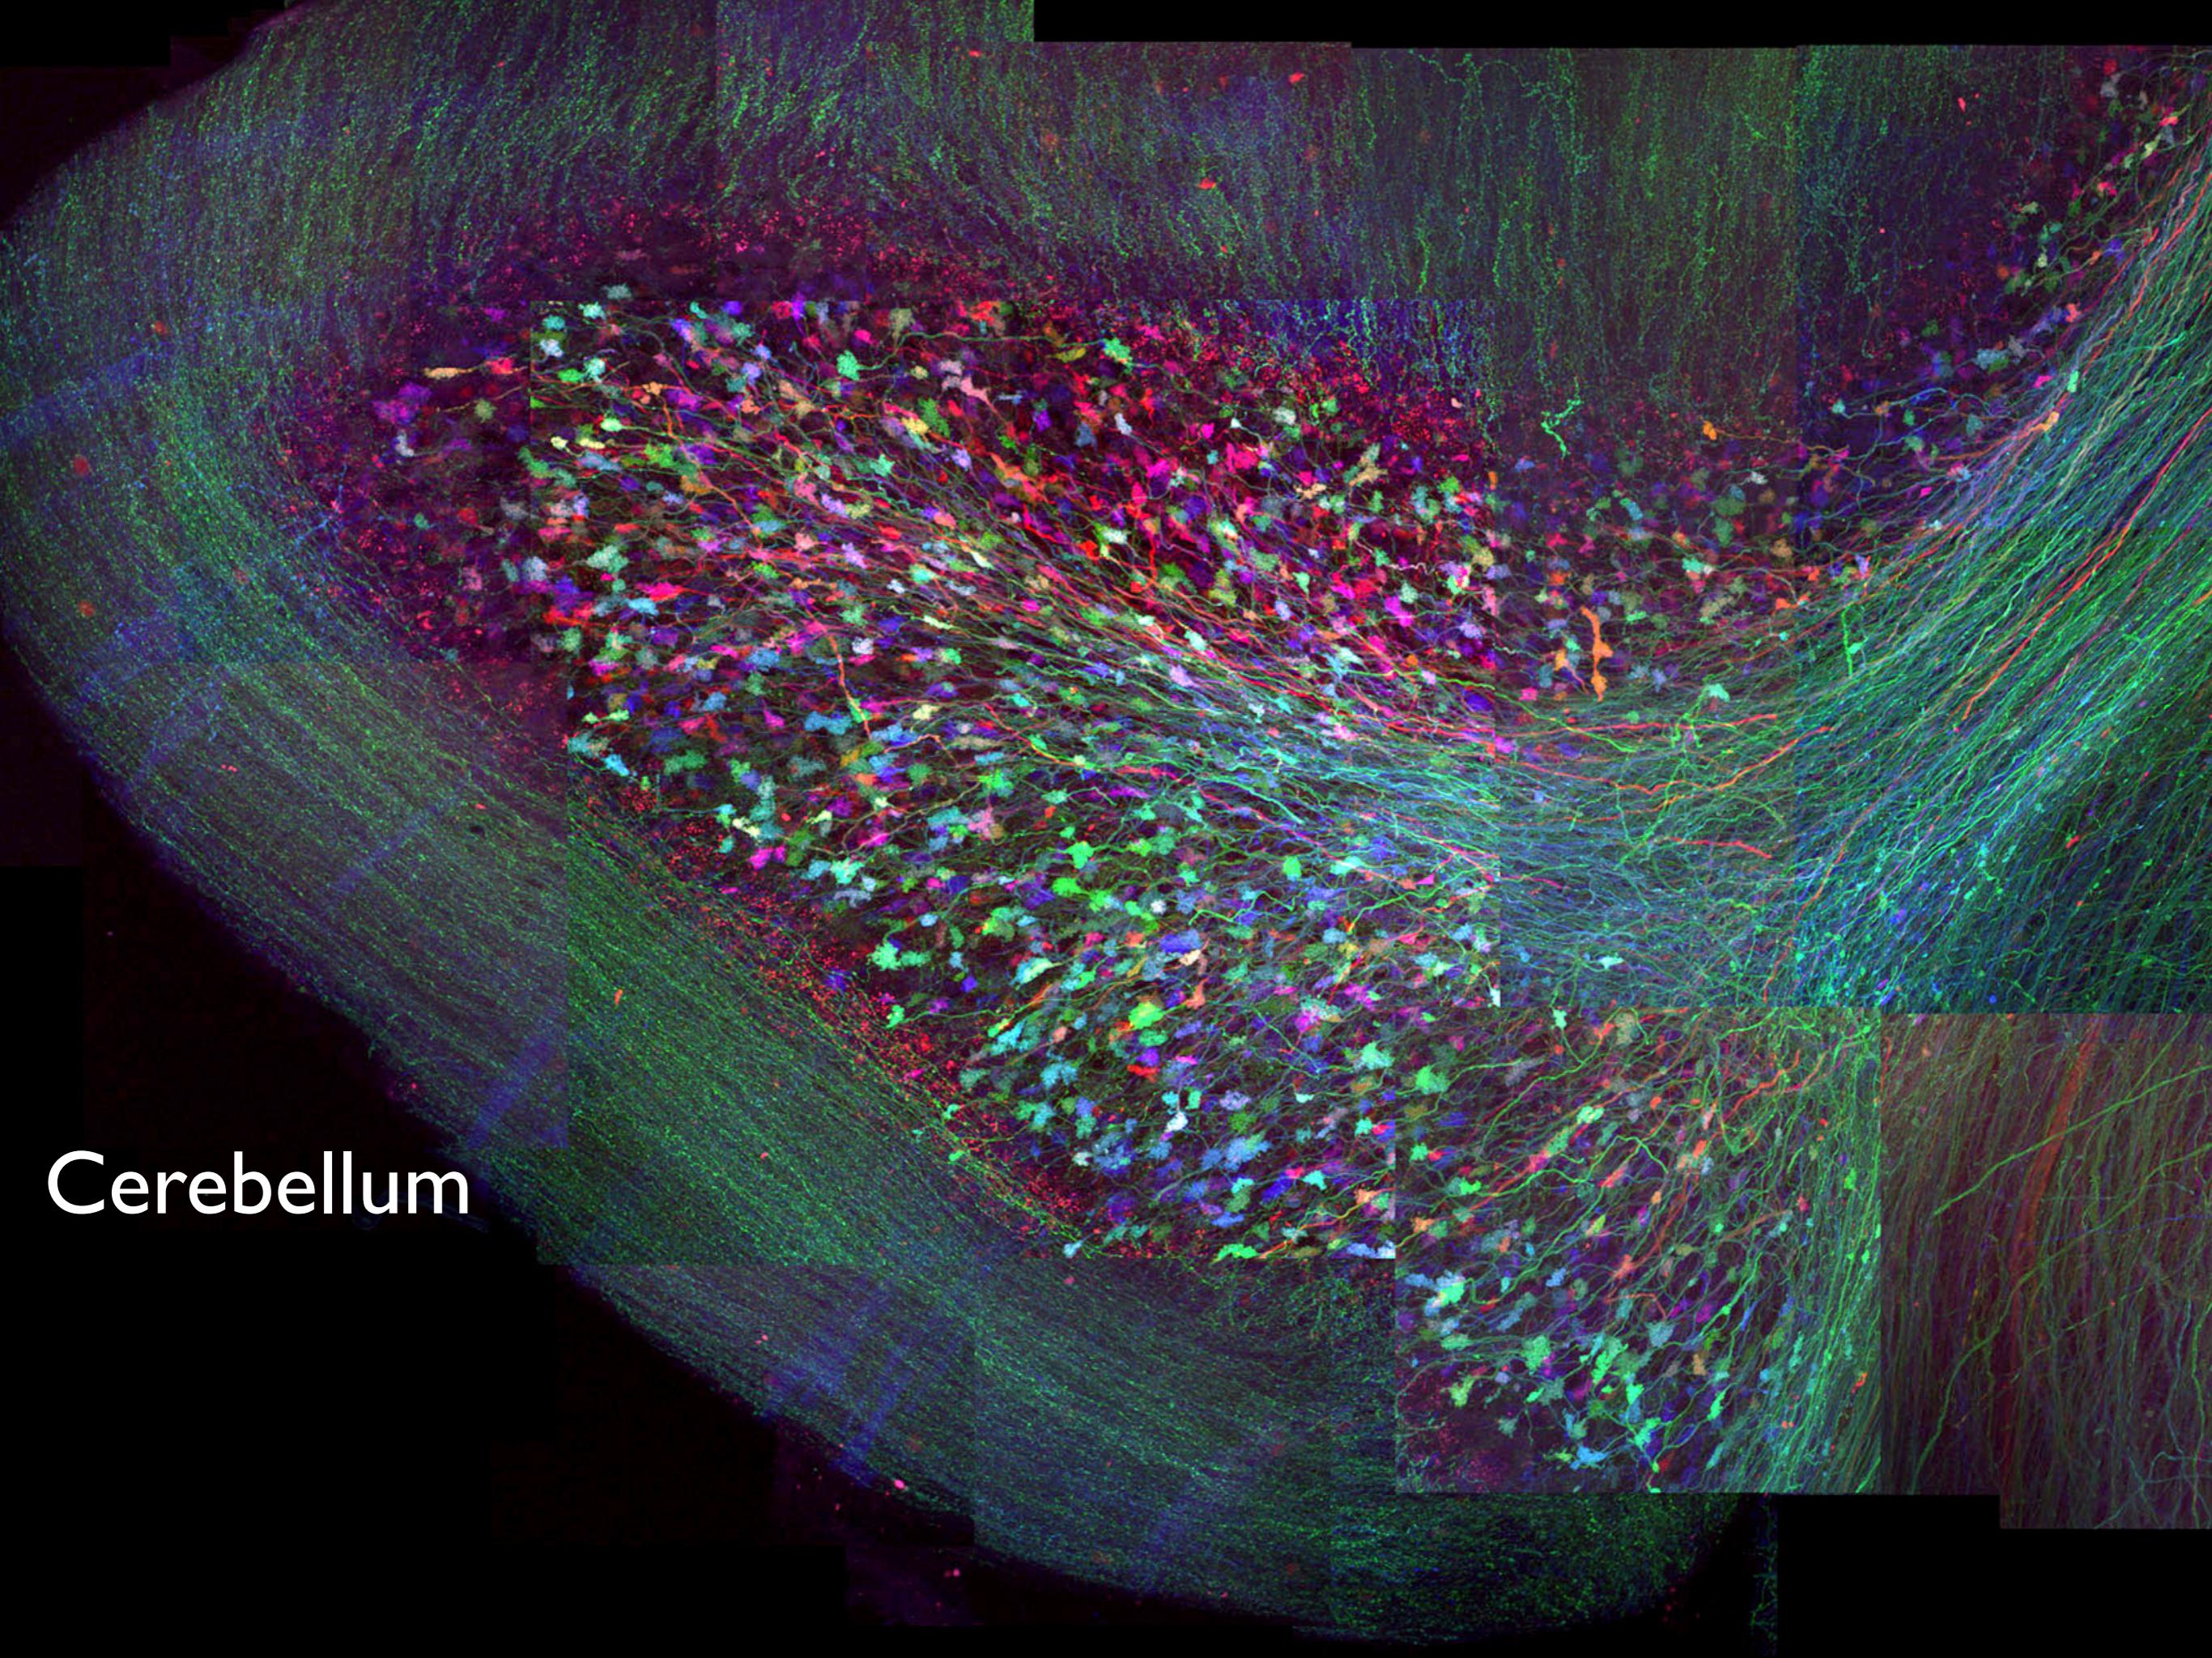
J. Livet

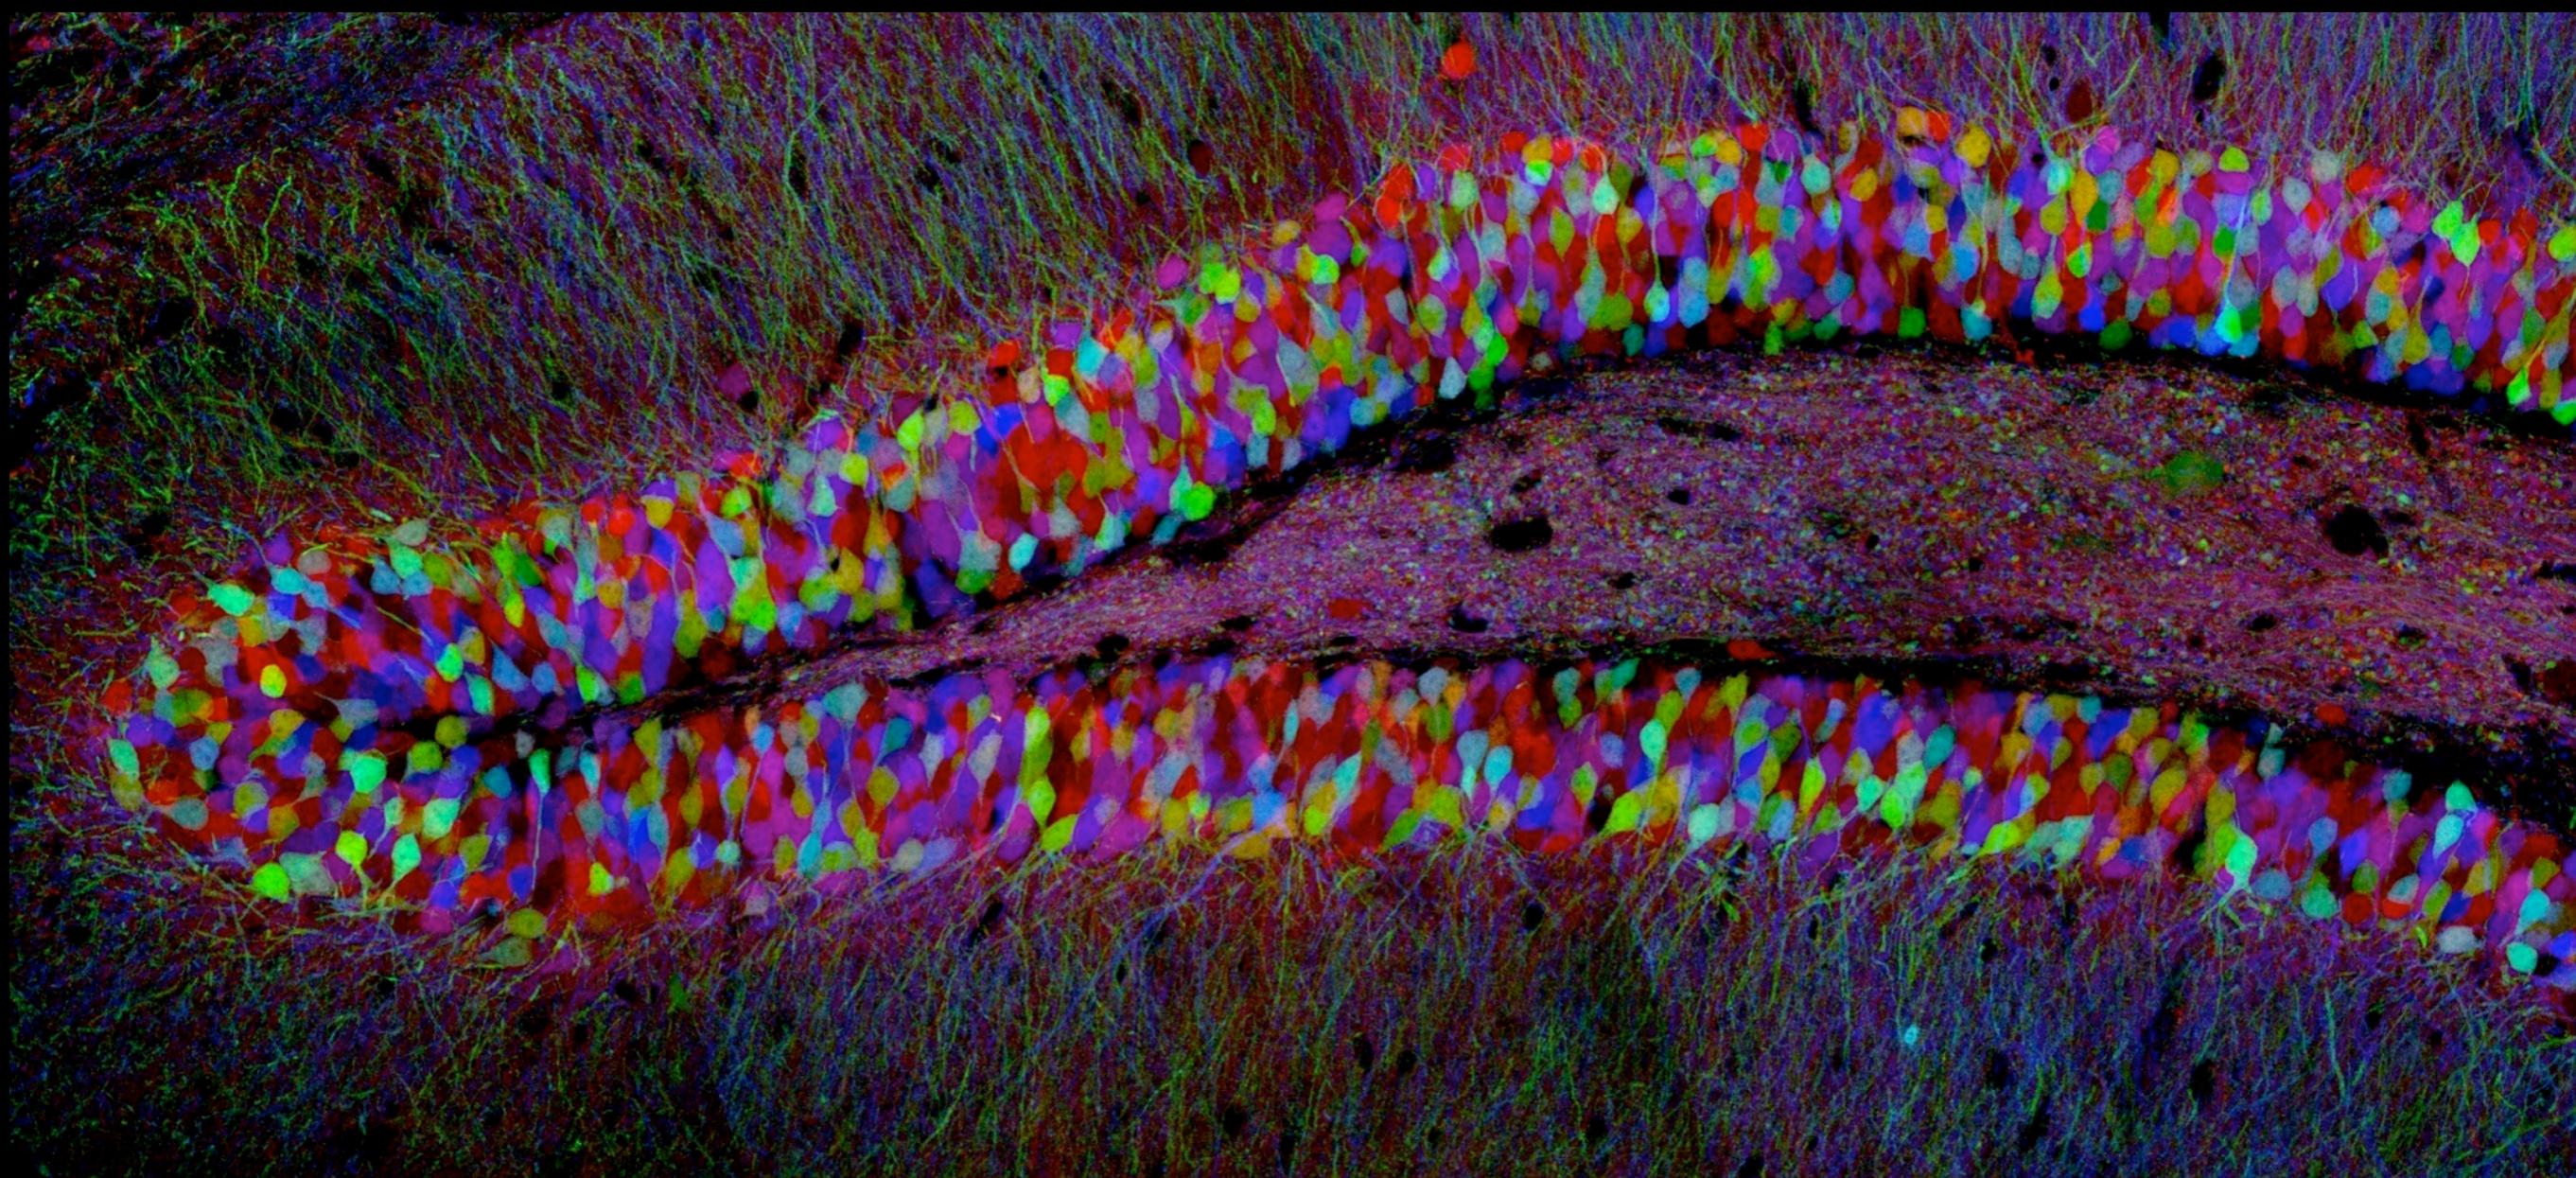
Brainbow



J. Livet

Cerebellum





Hippocampus



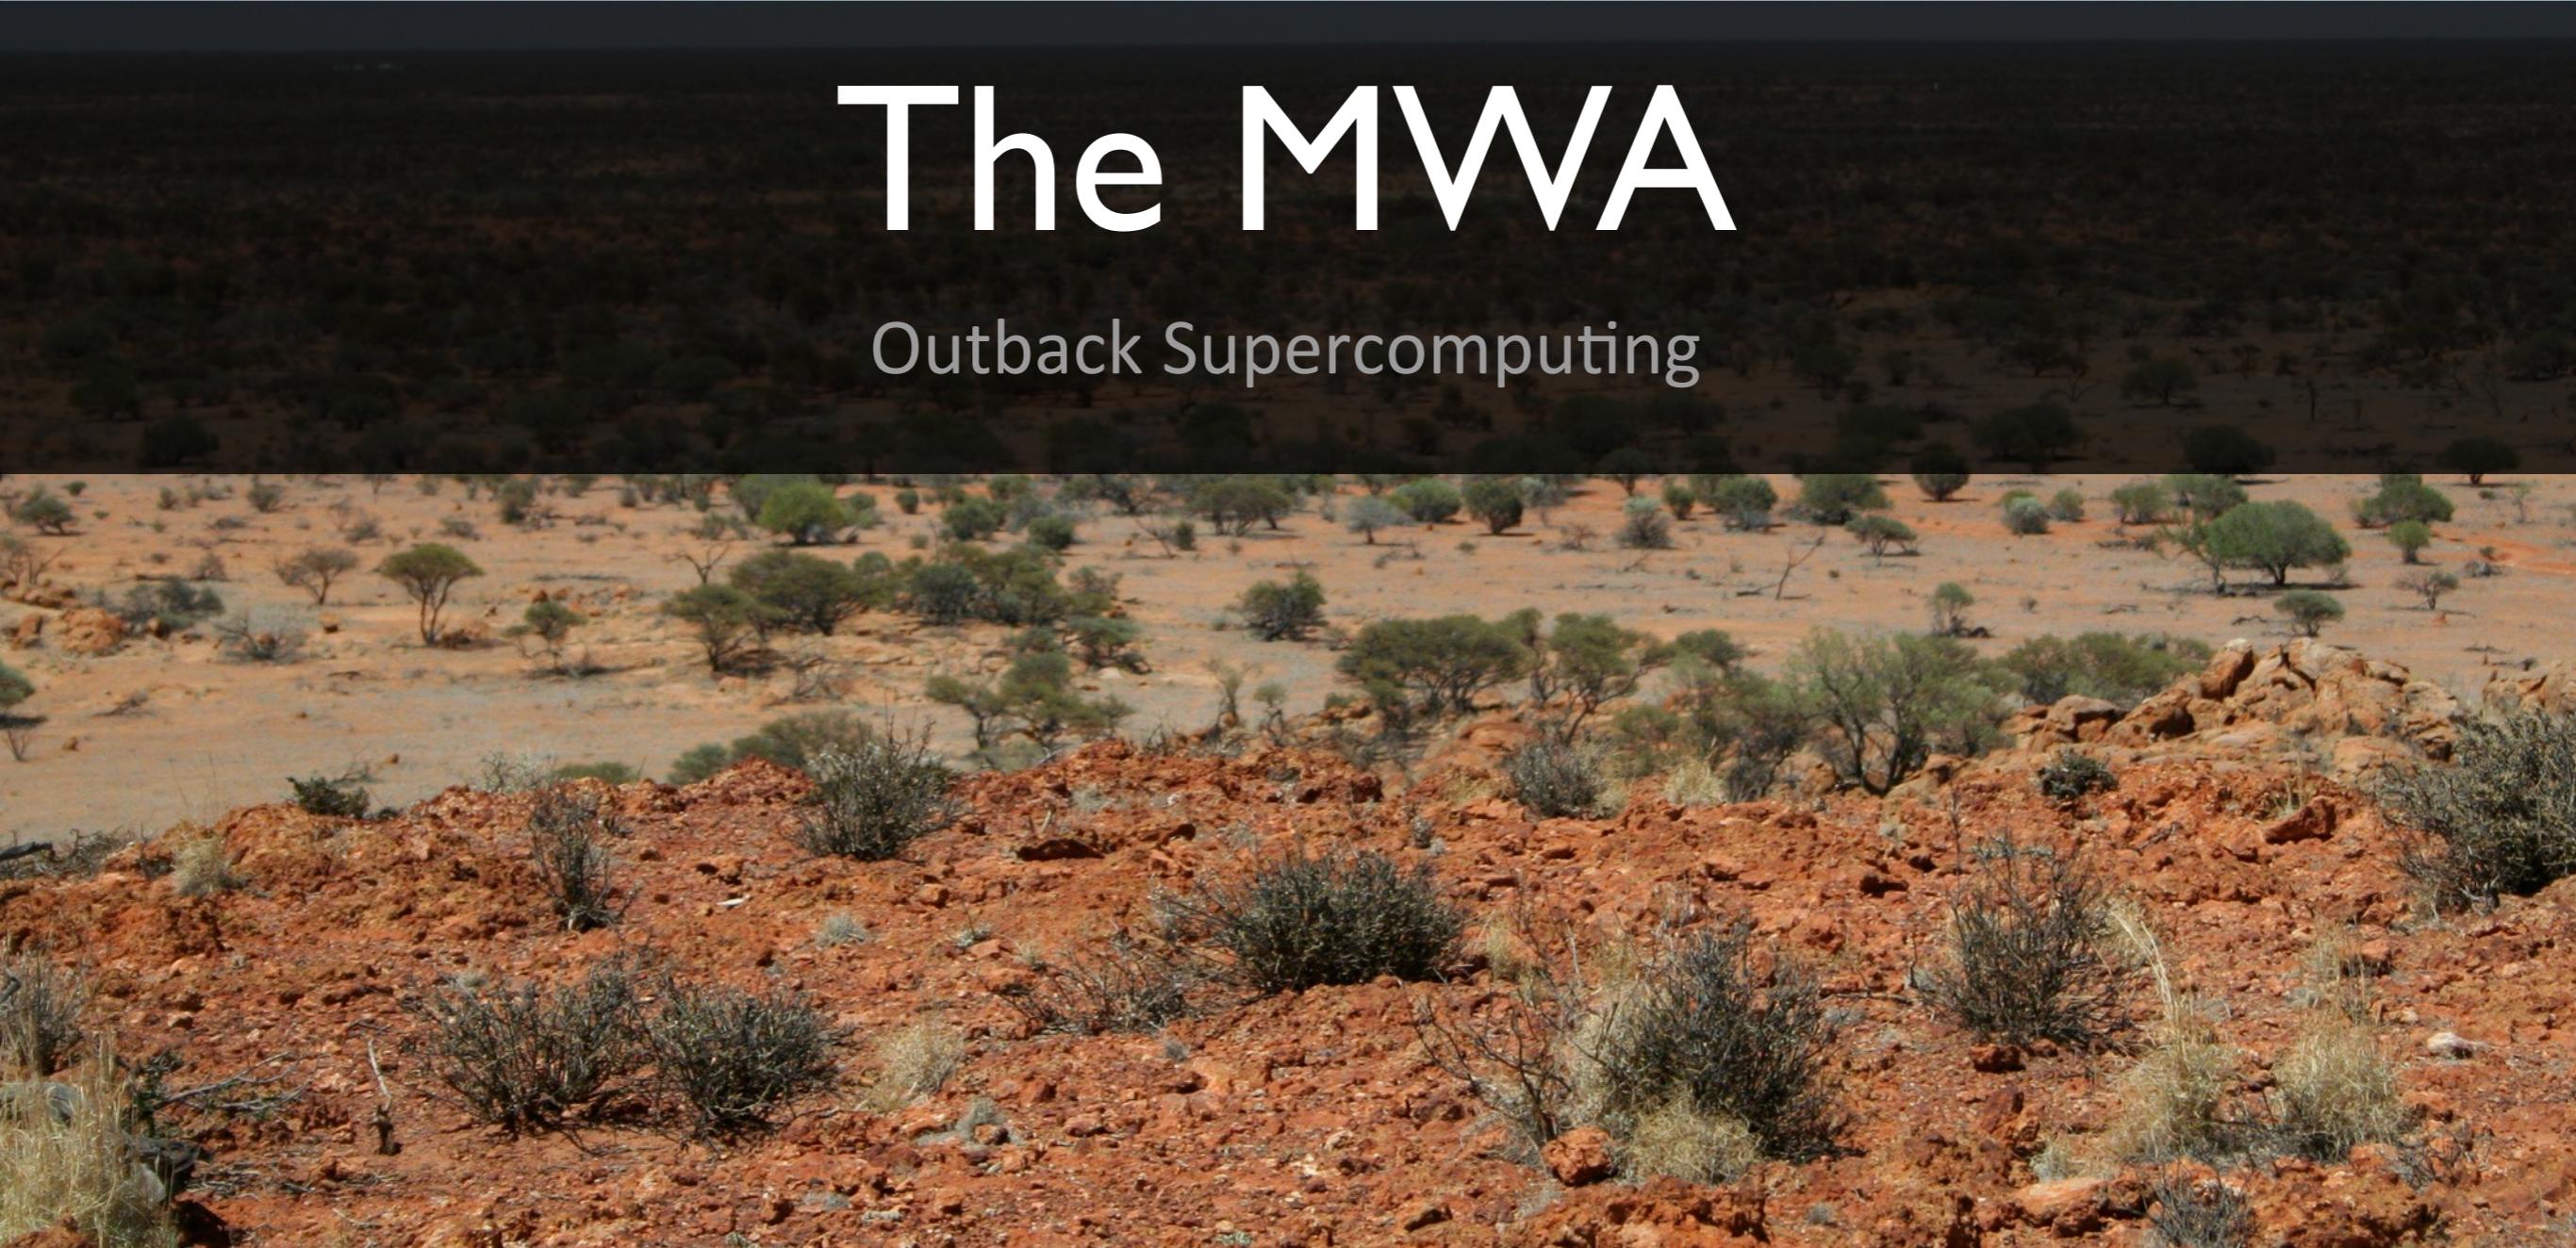
Convergent evolution?

Outline

- How is the brain wired?
- How did the universe start?
- How can we prevent heart attacks?

The MWA

Outback Supercomputing



Collaborators

- Harvard Center for Astrophysics
 - Lincoln Greenhill
 - Daniel Mitchell
 - Stephen Ord
 - Randall Wayth
- SEAS
 - Kevin Dale, Richard Edgar
 - Hanspeter Pfister



The Scientific Challenge

- What happened through the 1 billion years of the universe's dark ages?

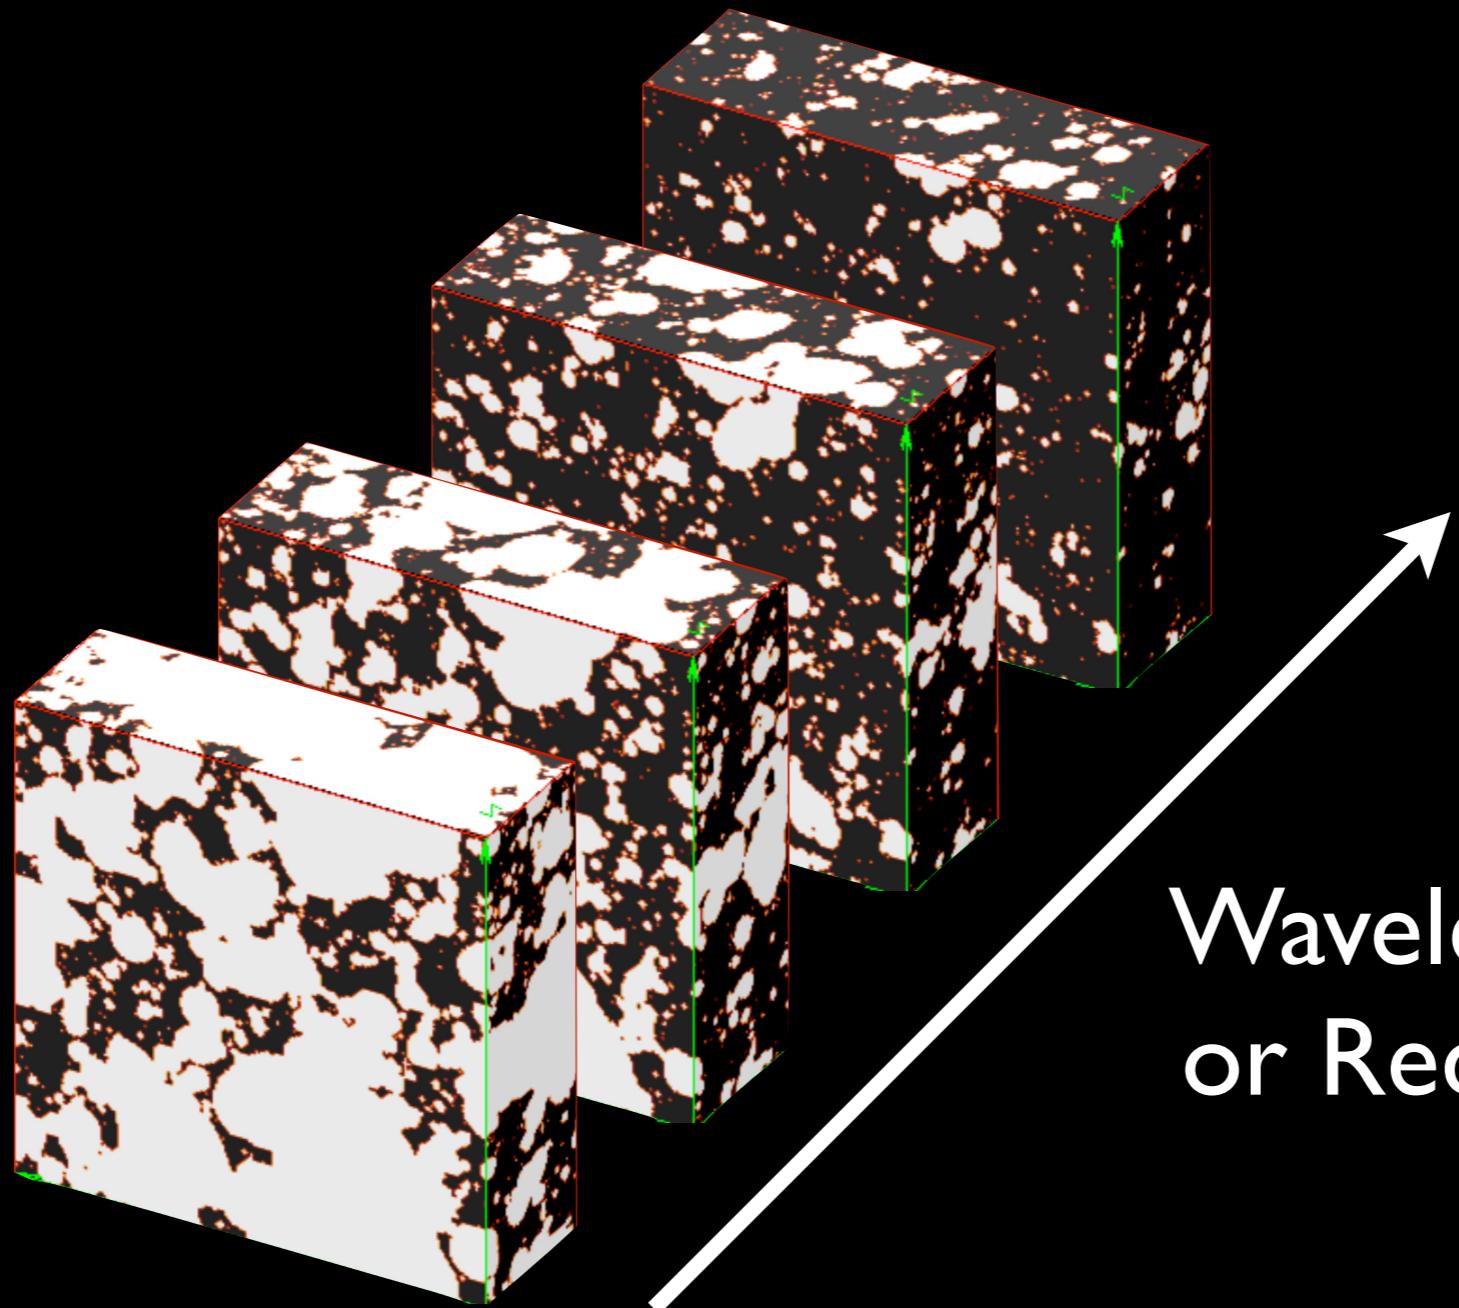


Loeb / SciAm

~300,000

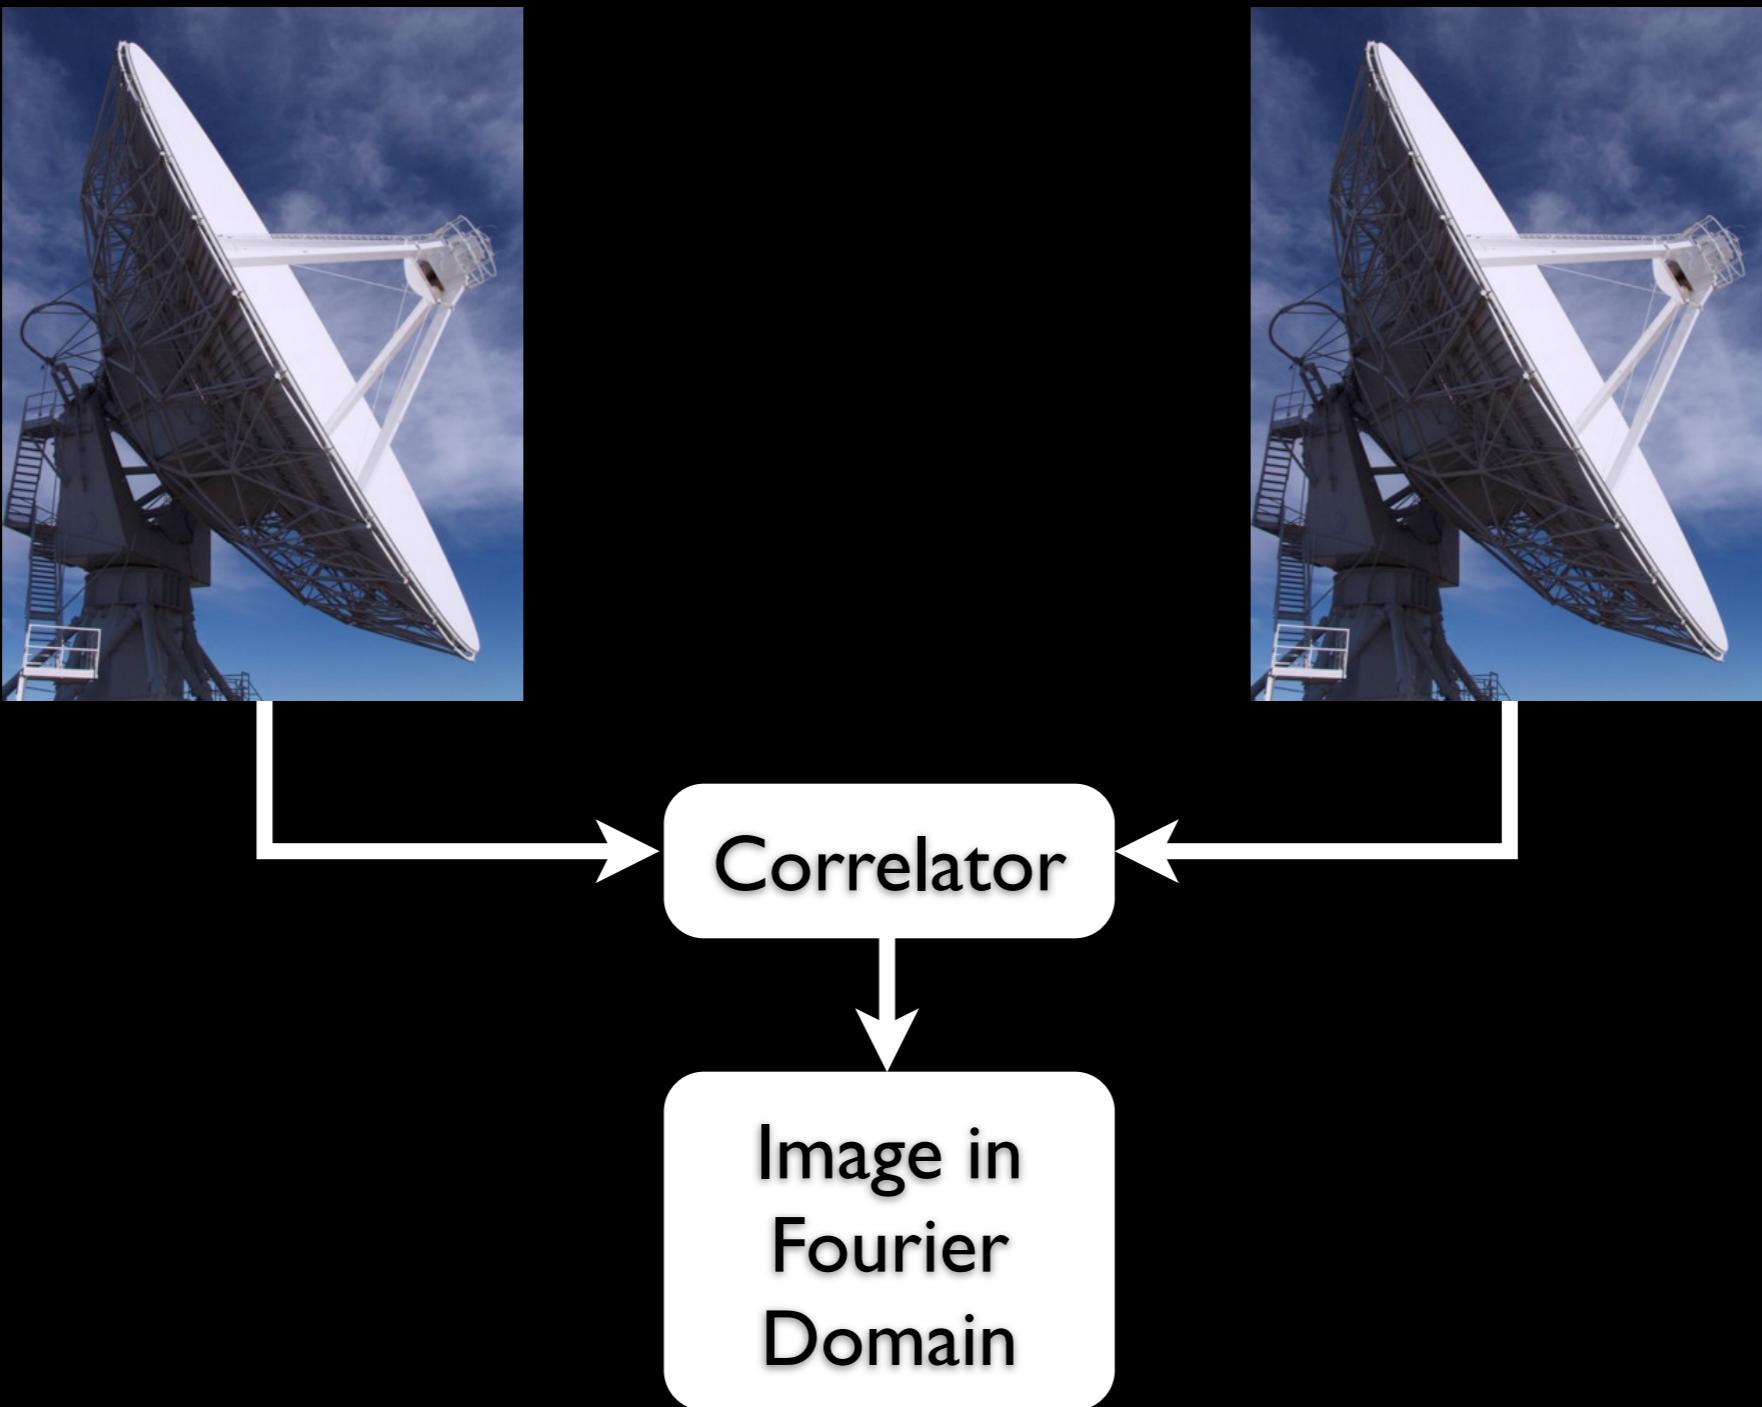
~1 bio

3D Tomography



Wavelength
or Redshift

Radio Interferometry



From this...



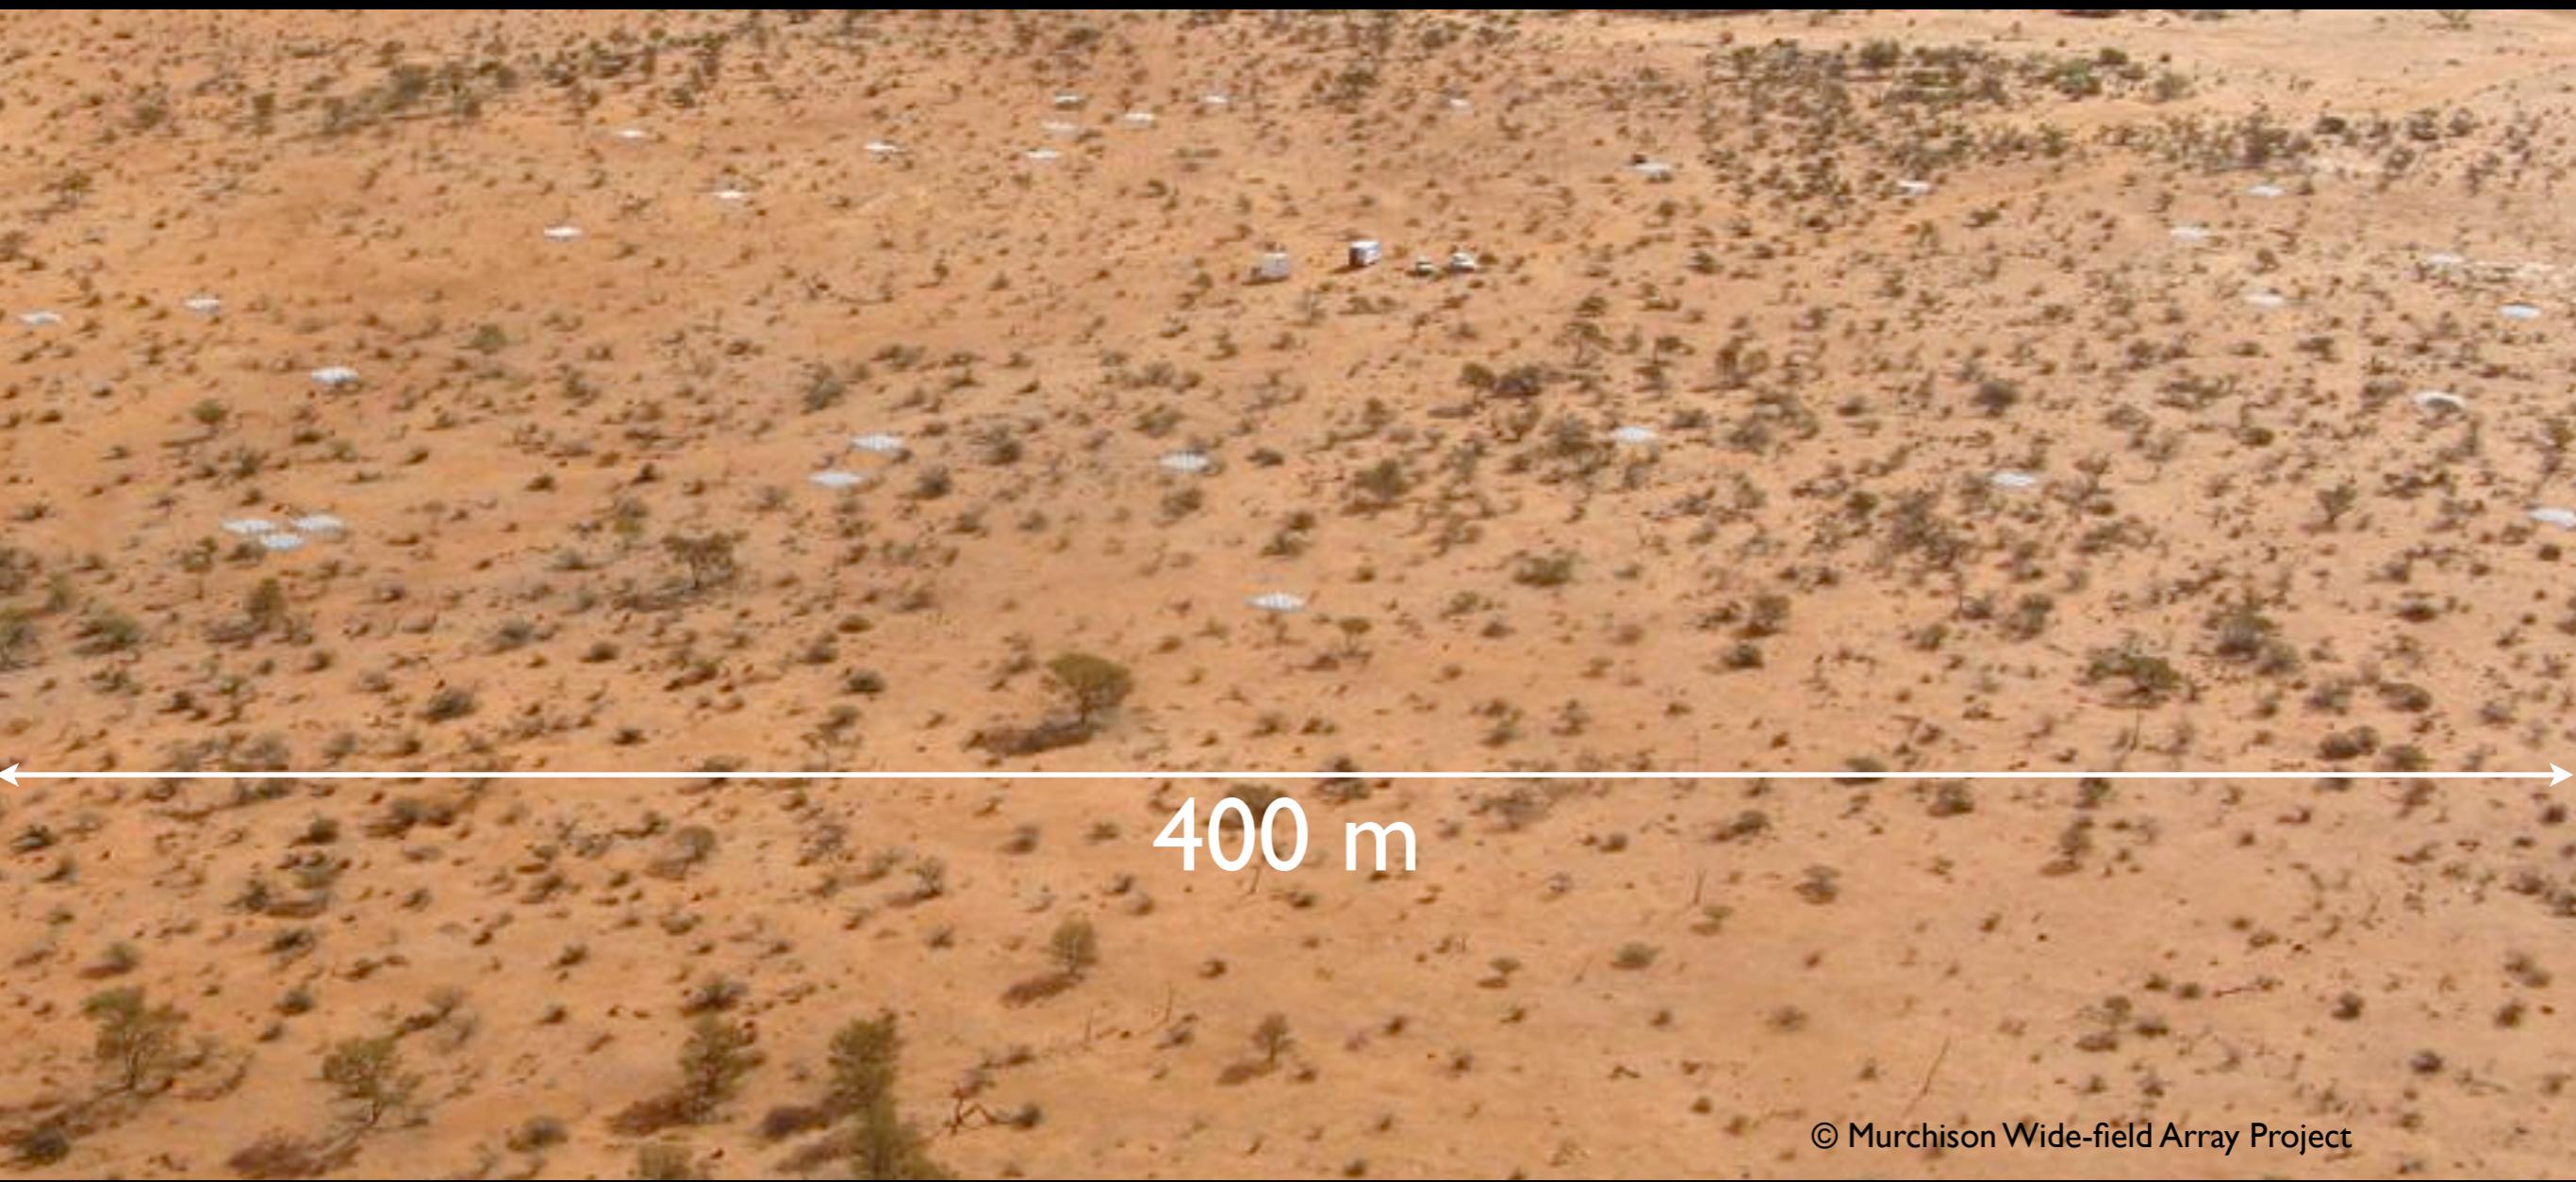
APOD 0605

...to this



MWA

- 32 antenna prototype
- Eventually 512 antennas on 1 km^2



400 m



Nowhere (Middle of)
300 km

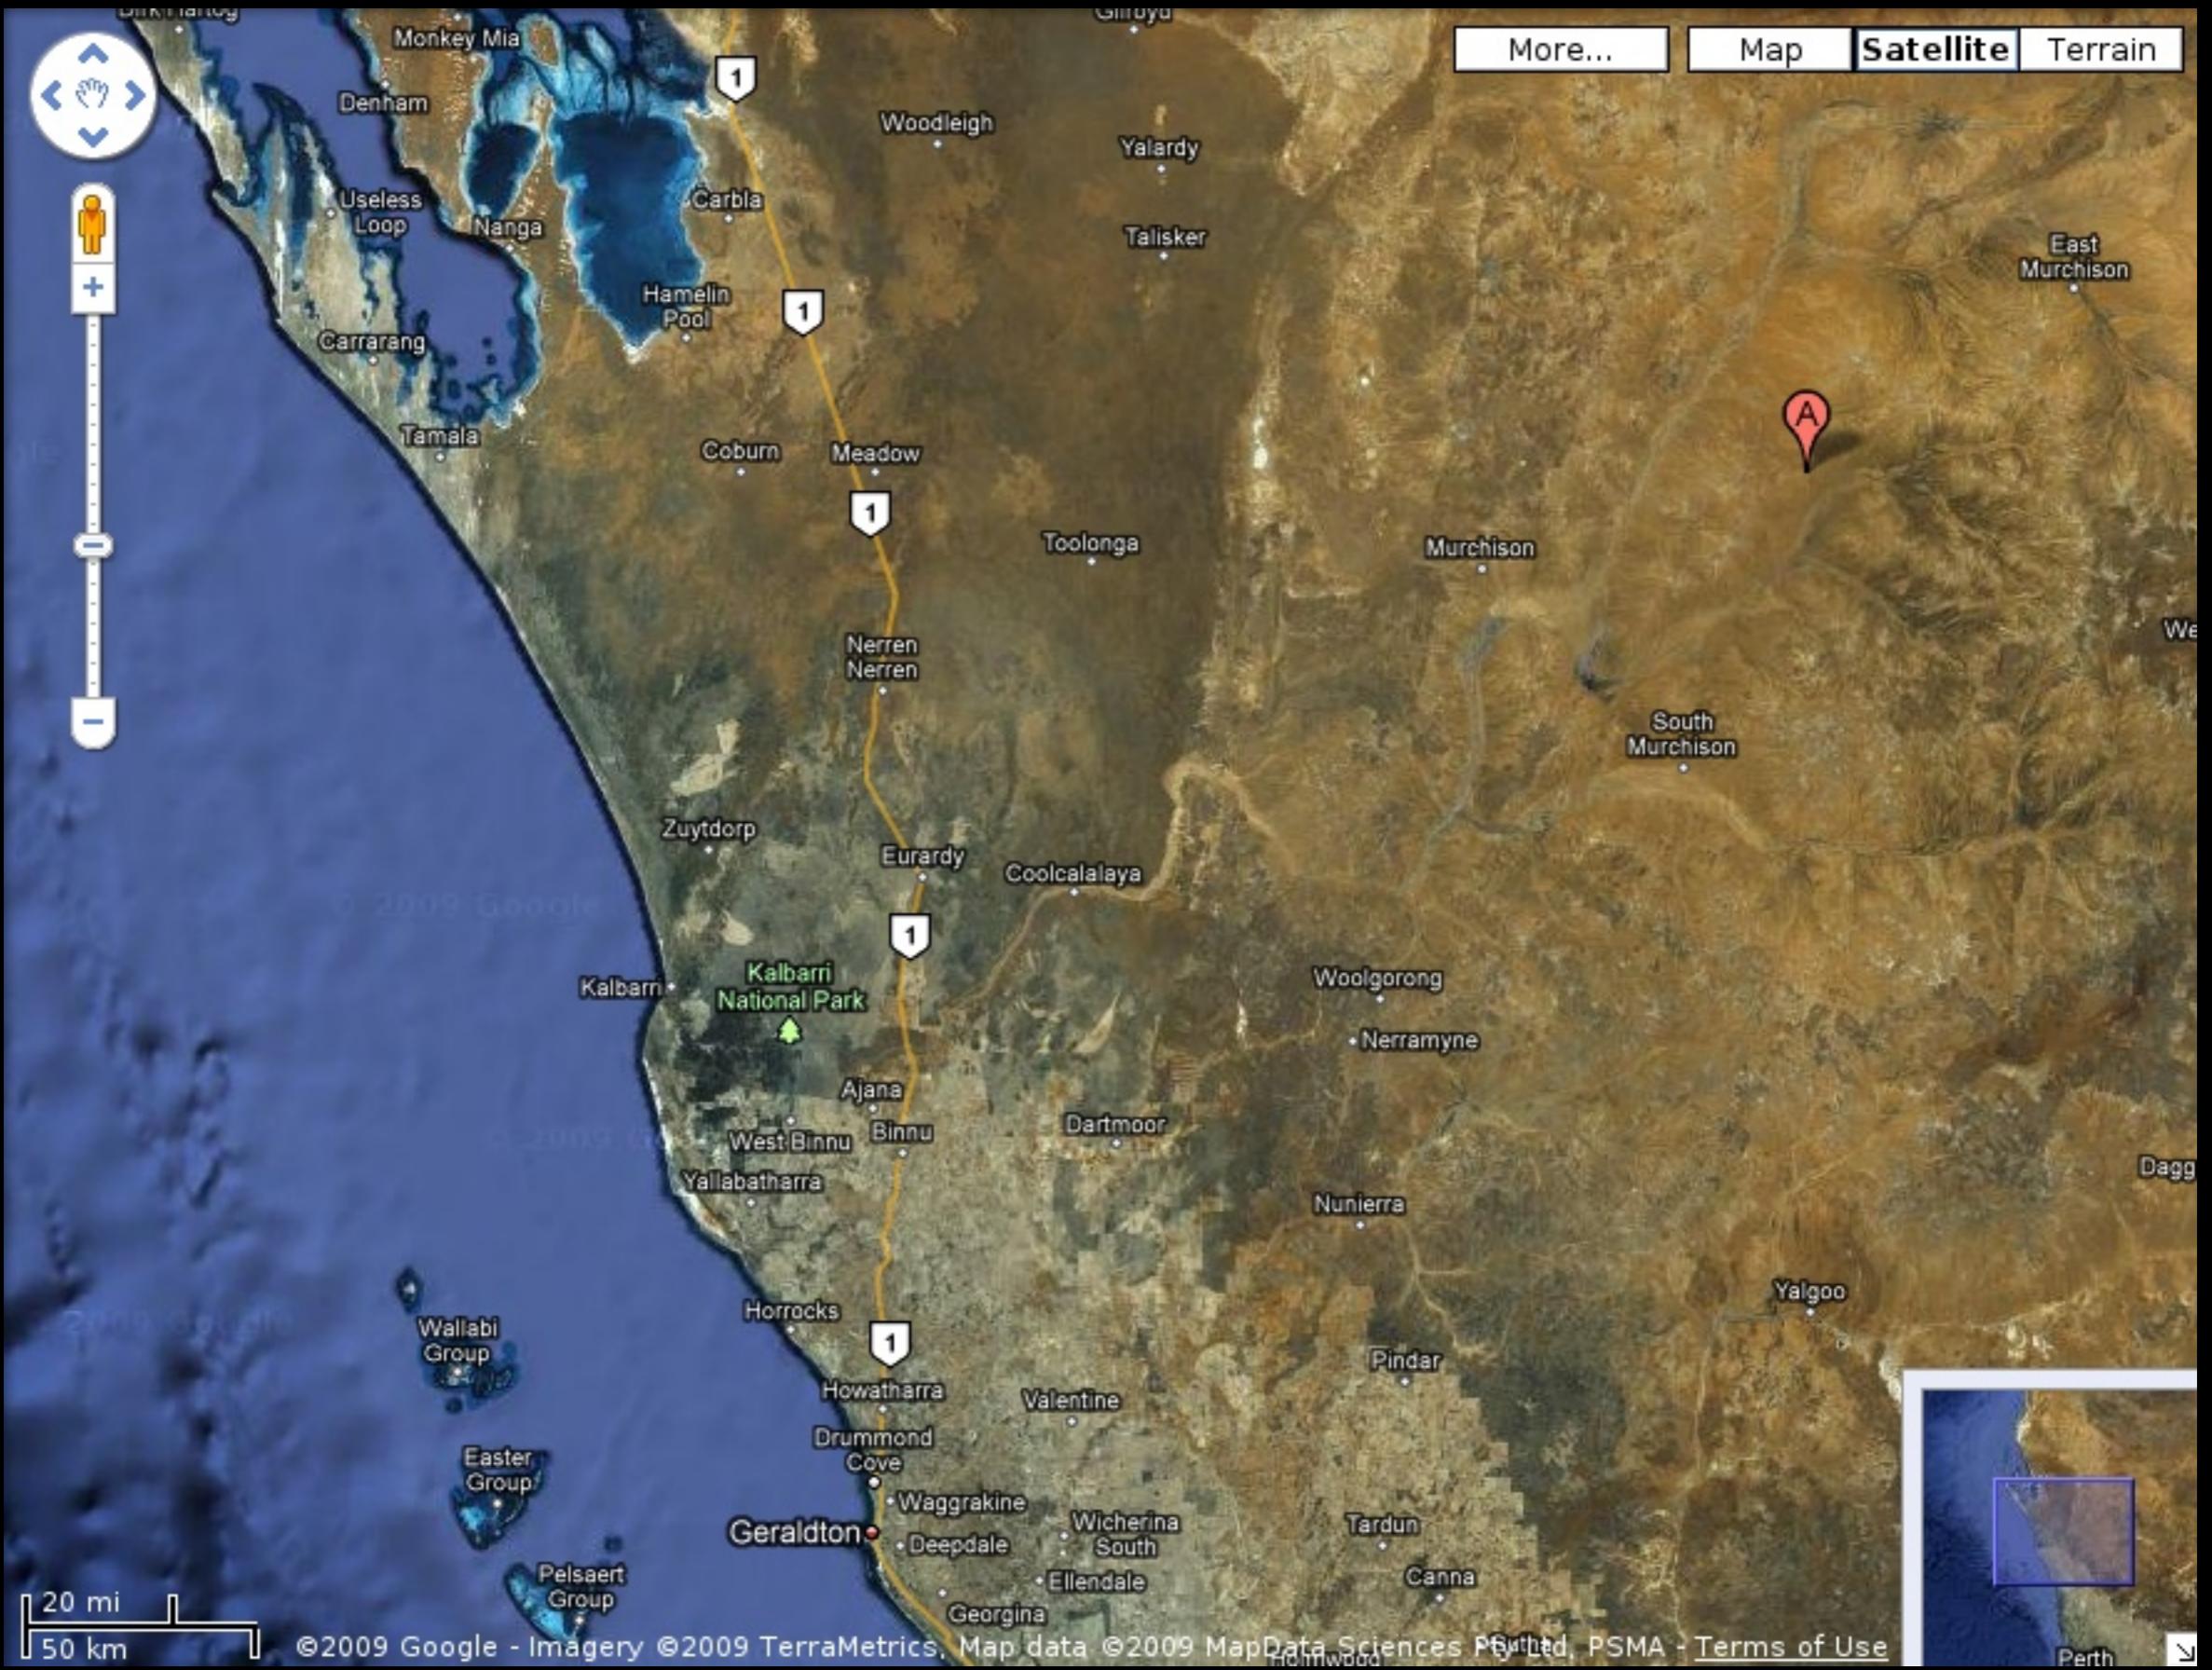
THANK YOU FOR BEING RADIO QUIET



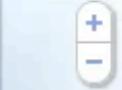
Outback Computing

- Max. ambient on site: 50° C / 122° F

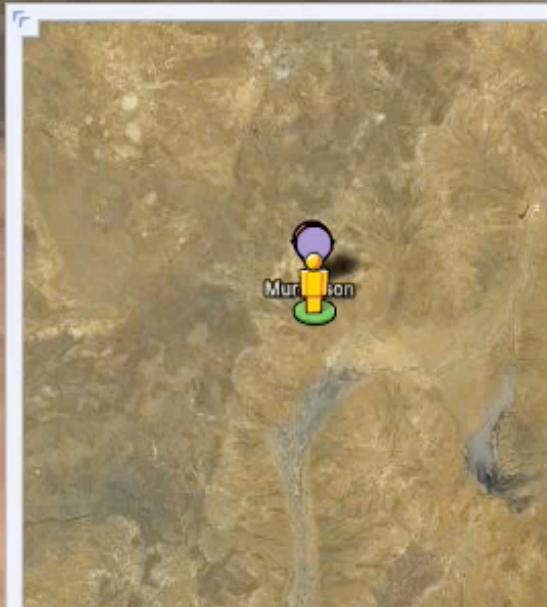


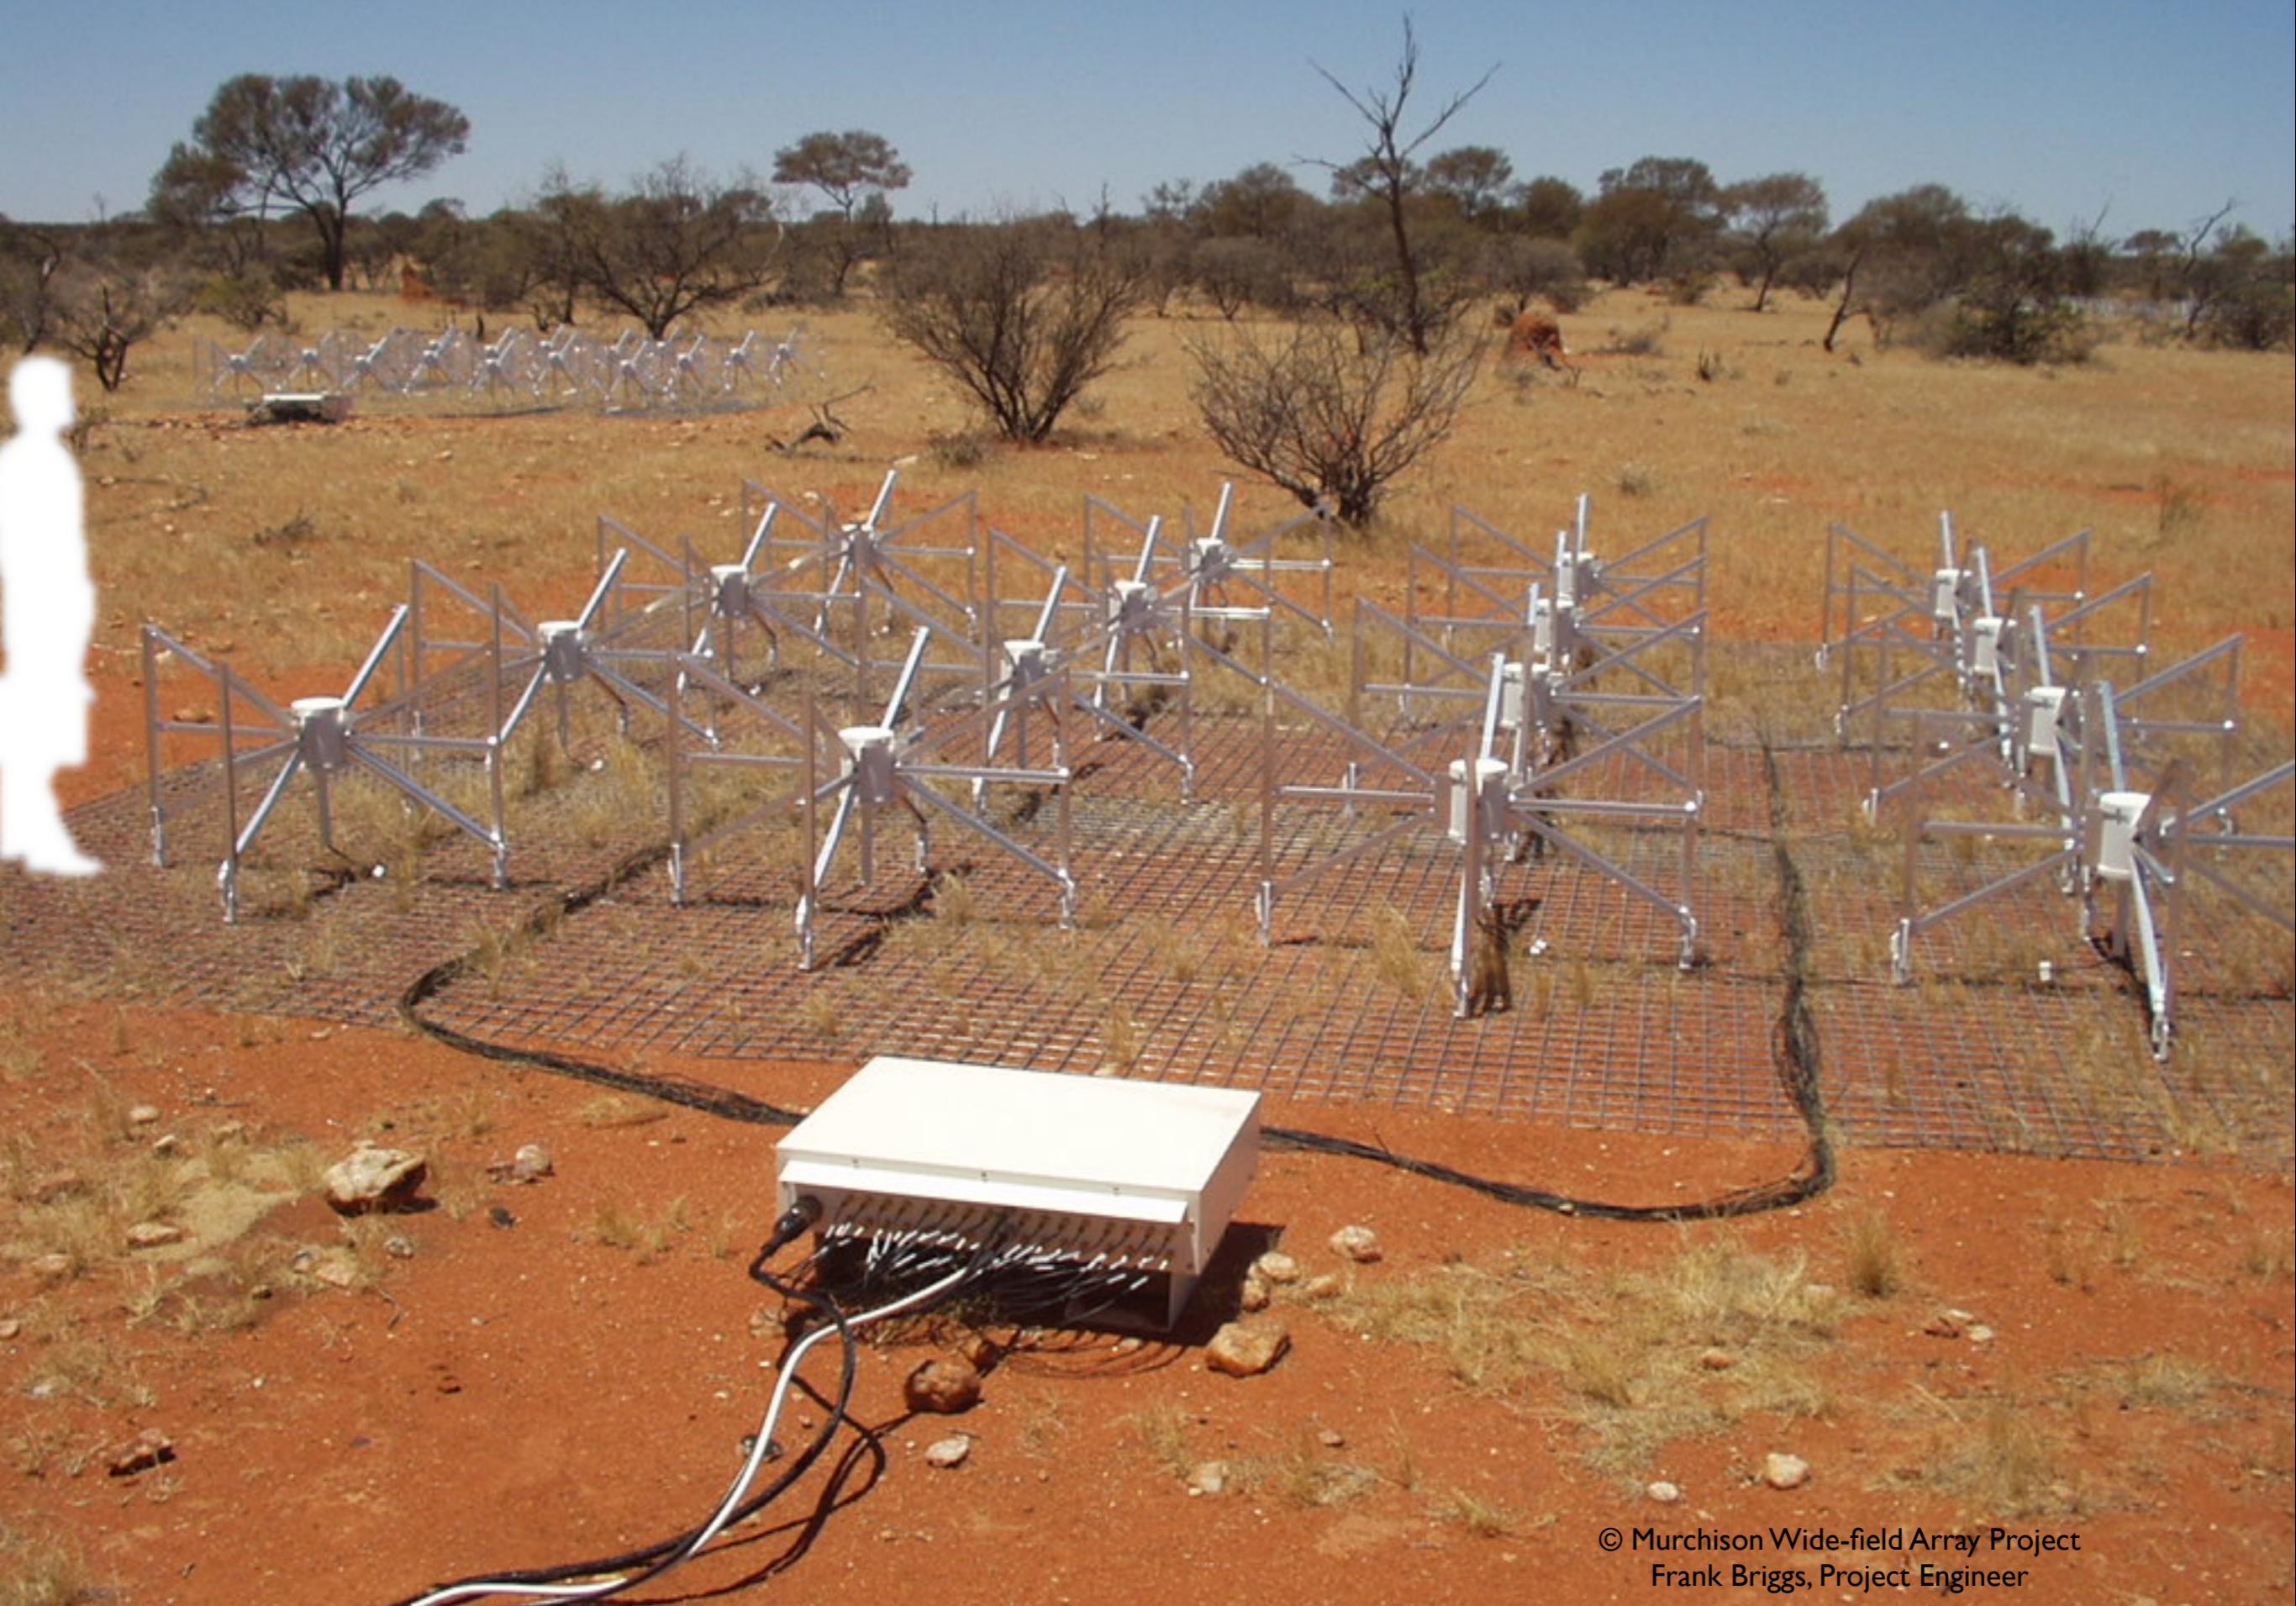


N
324 Mullewa Carnarvon Rd, Murchison, WA, Australia
Address is approximate



Mullewa Carnarvon Rd





© Murchison Wide-field Array Project
Frank Briggs, Project Engineer

Nervous System



Supercomputing Center



Who is the GPU expert?

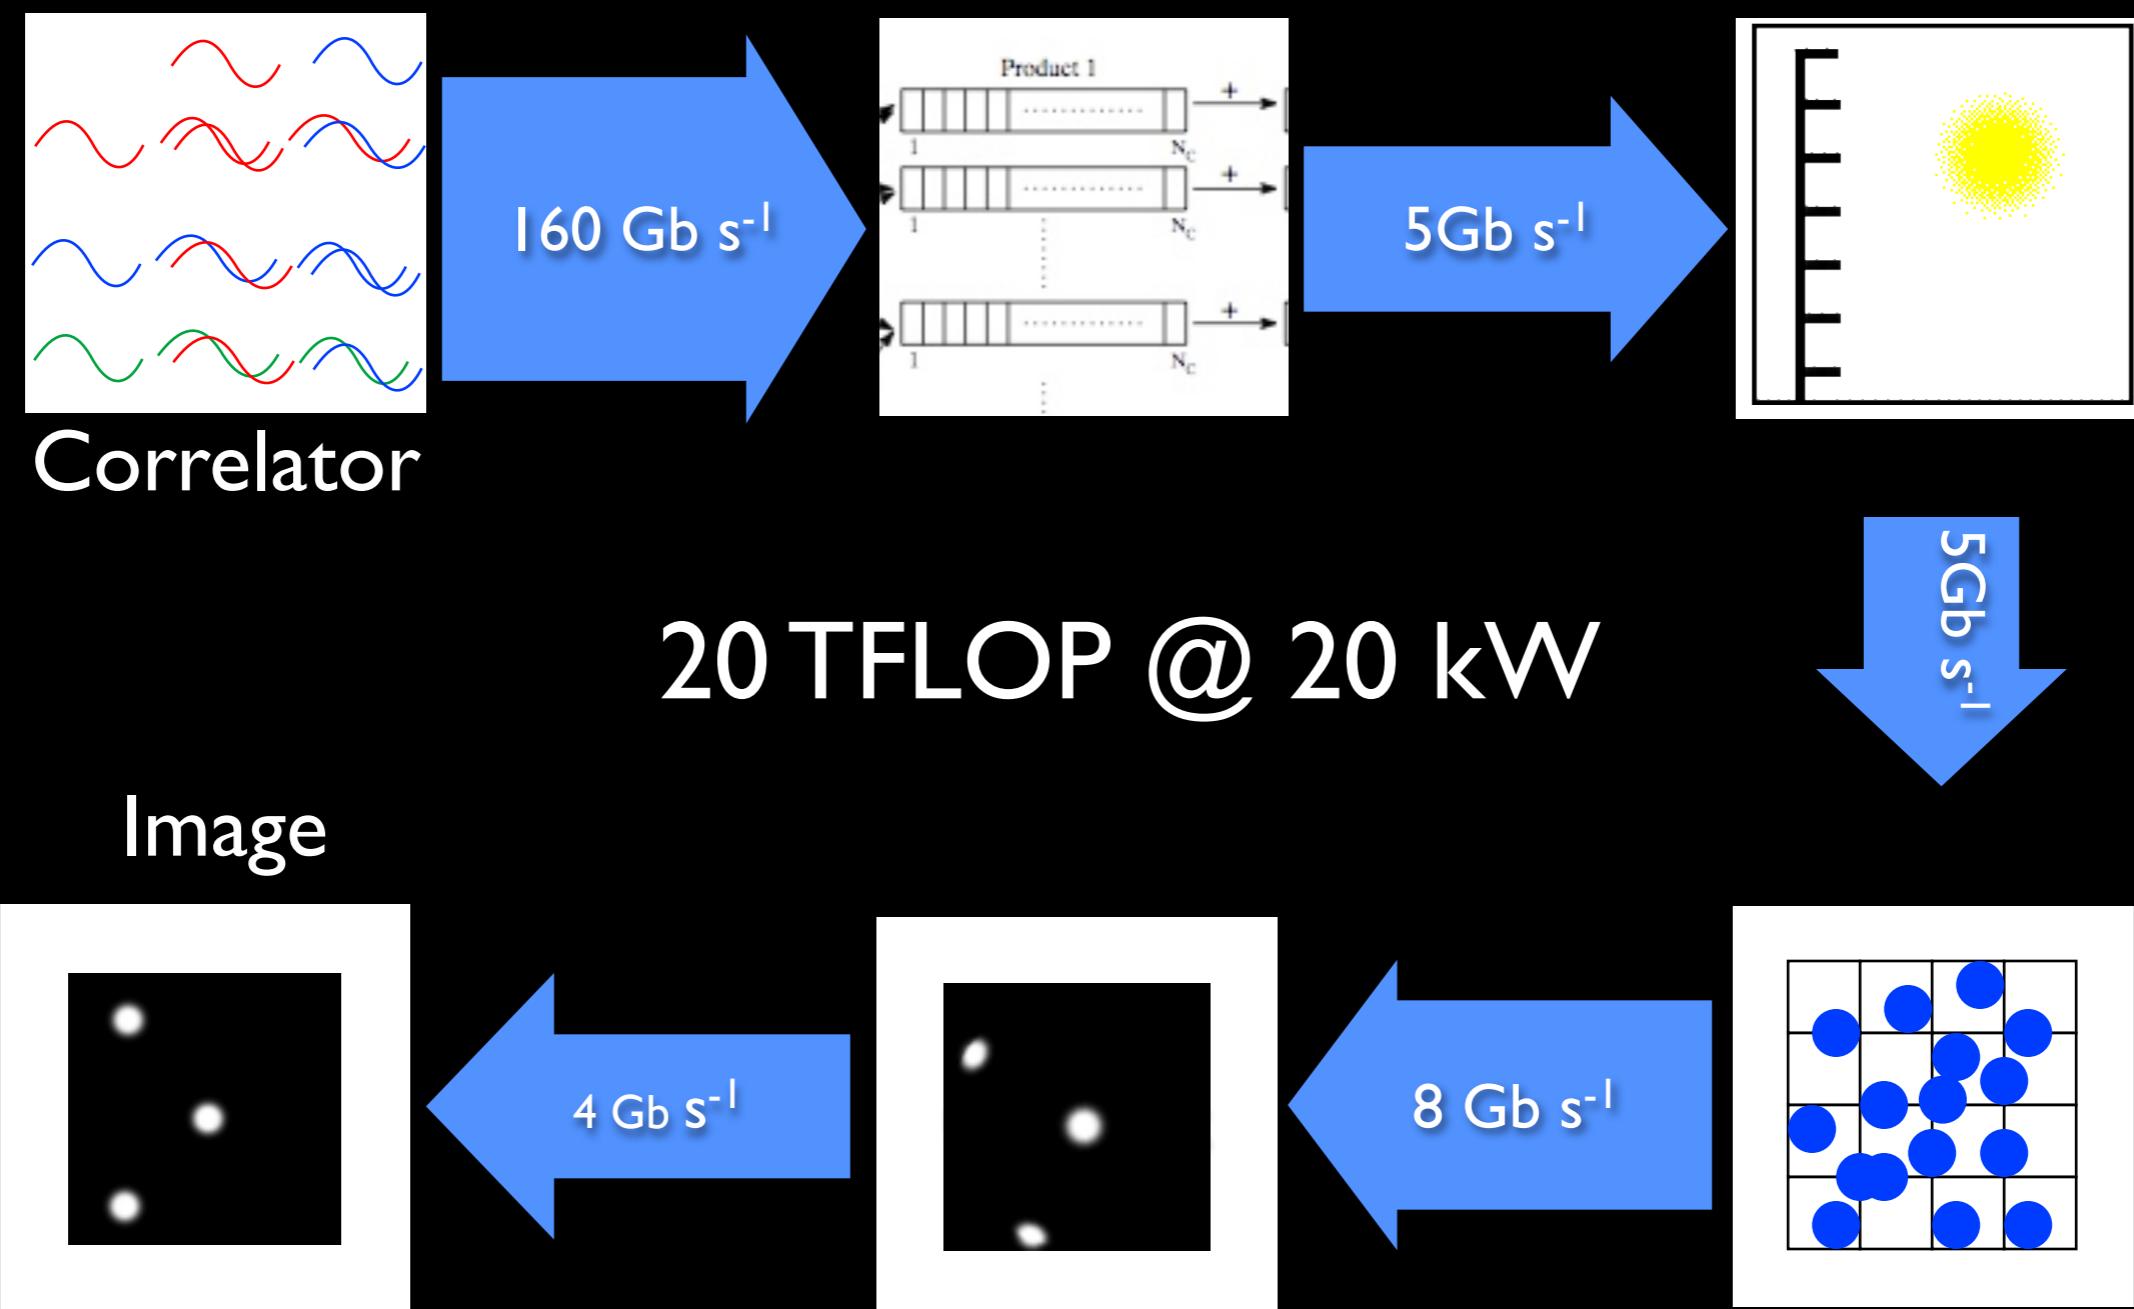


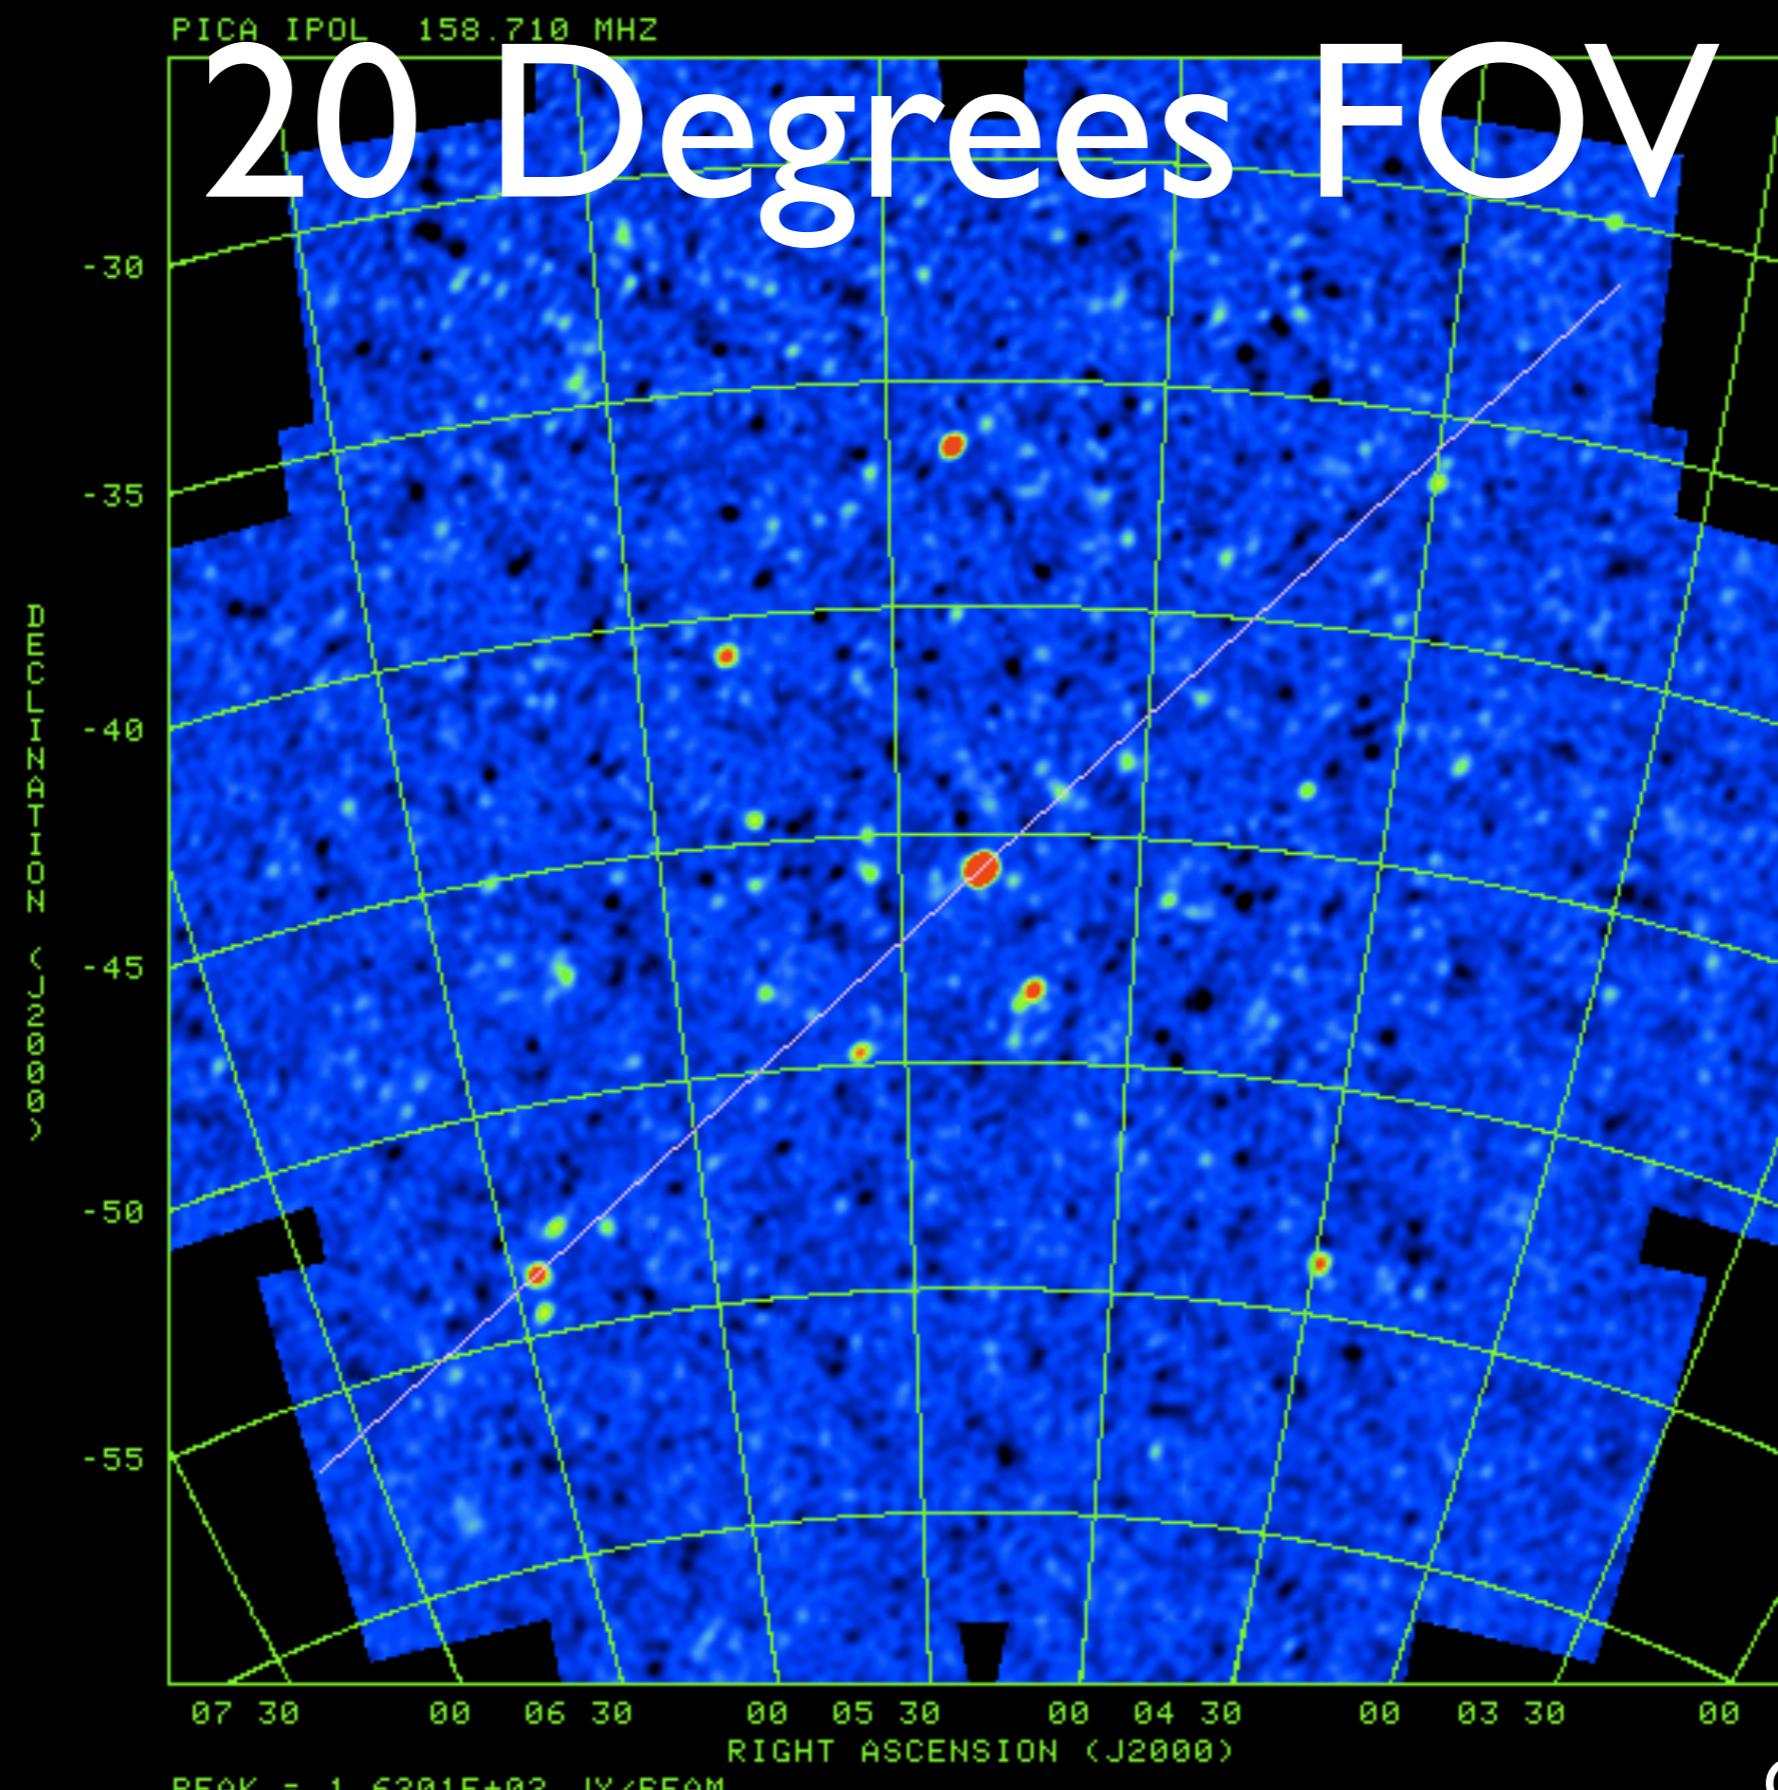
Life Support



© Murchison Wide-field Array Project
Frank Briggs, Project Engineer

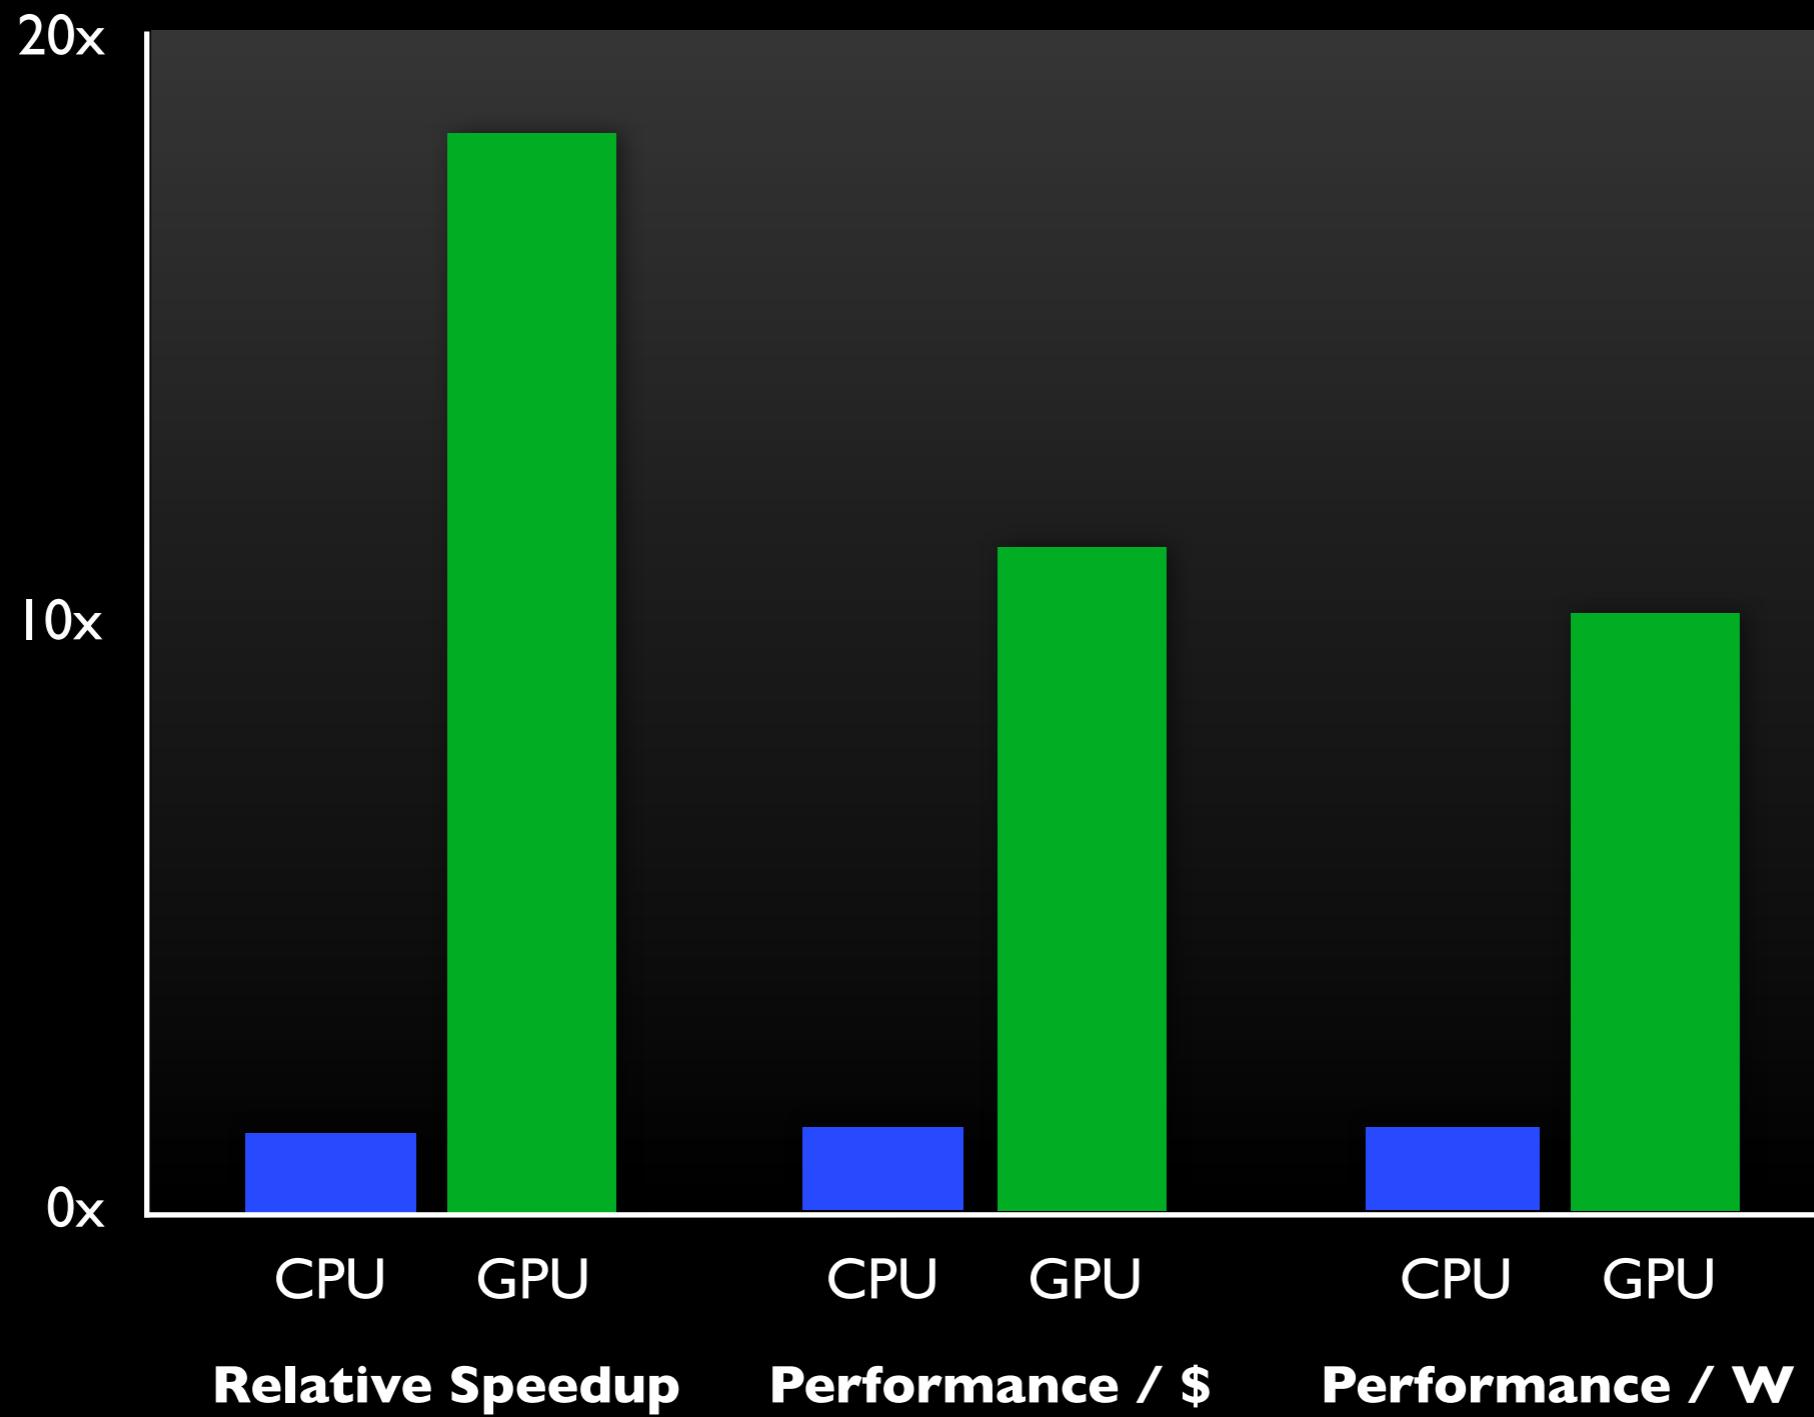
MWA Pipeline



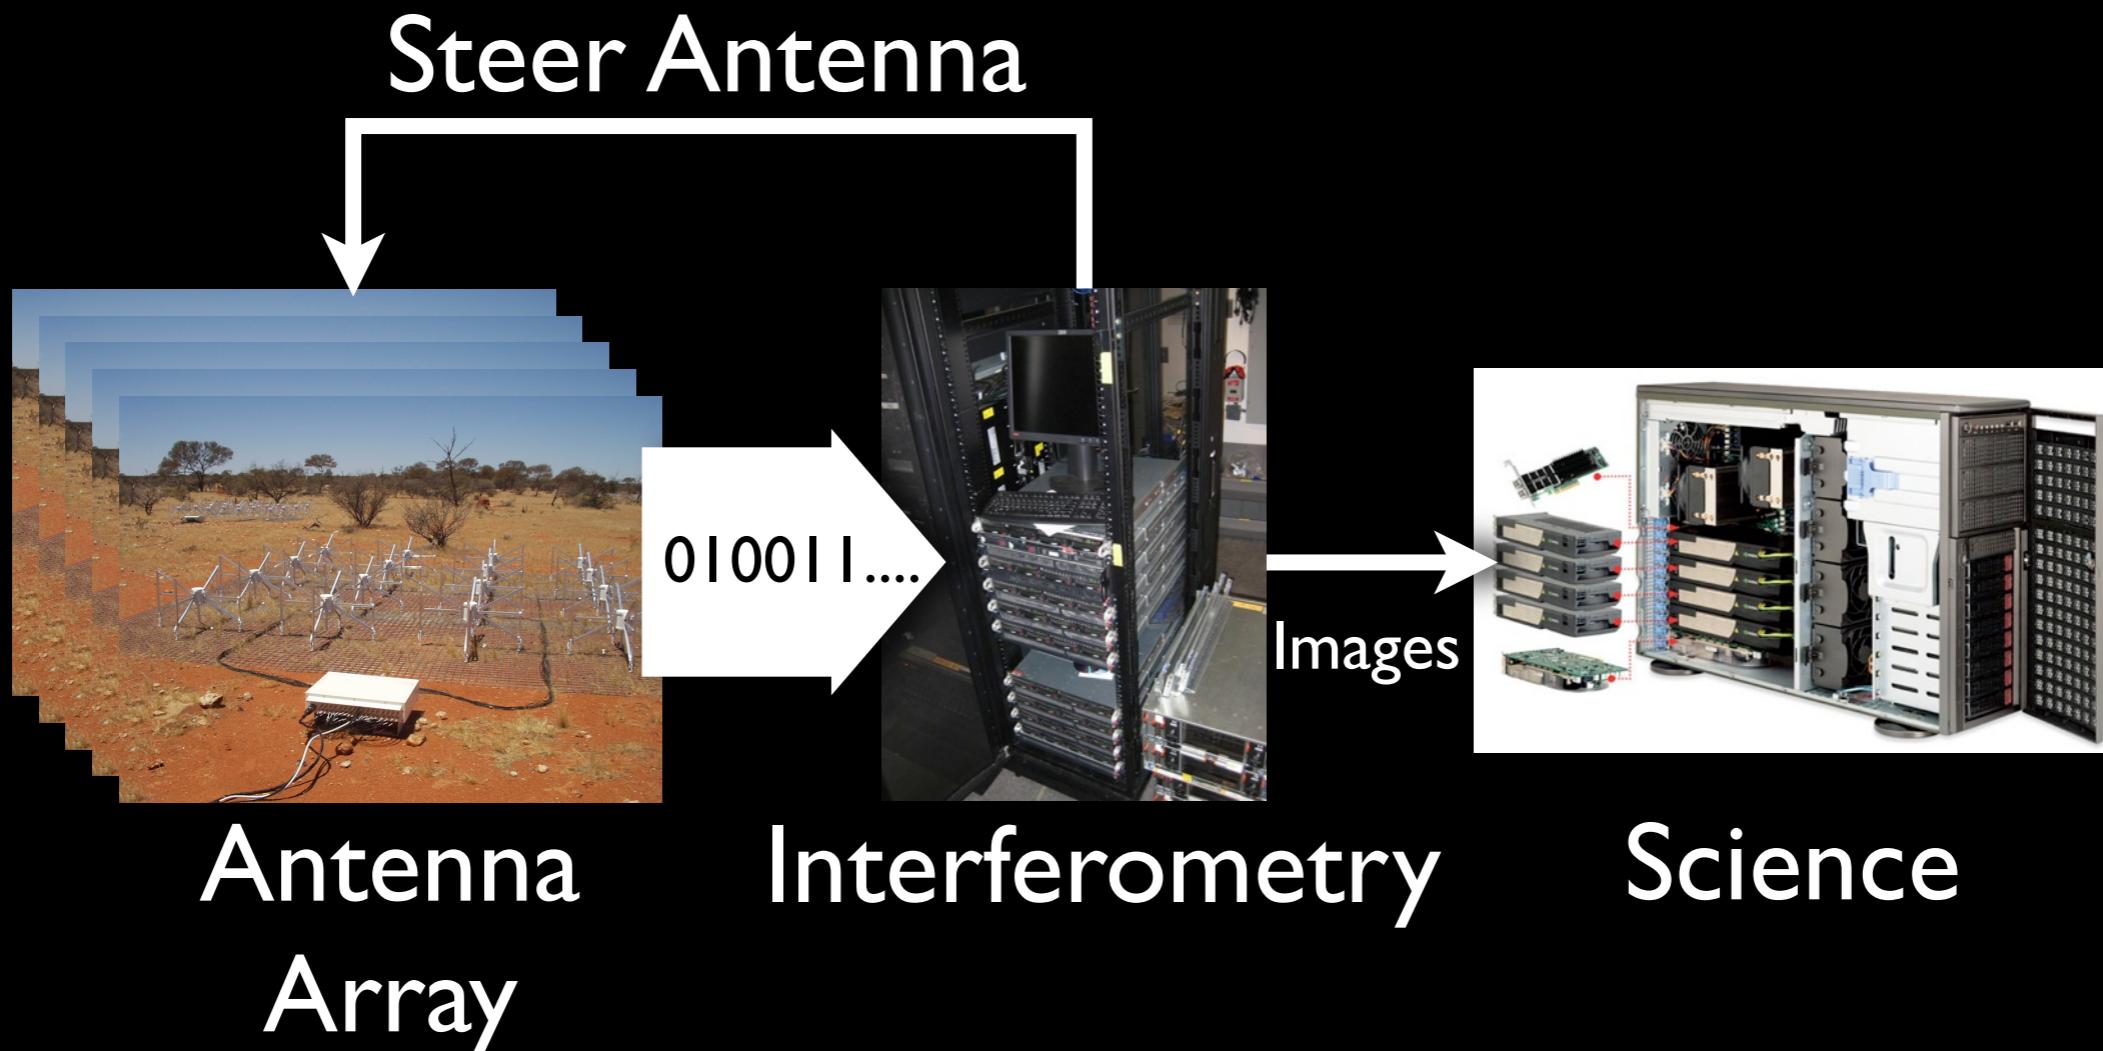


C.Williams, MIT

Real Time System

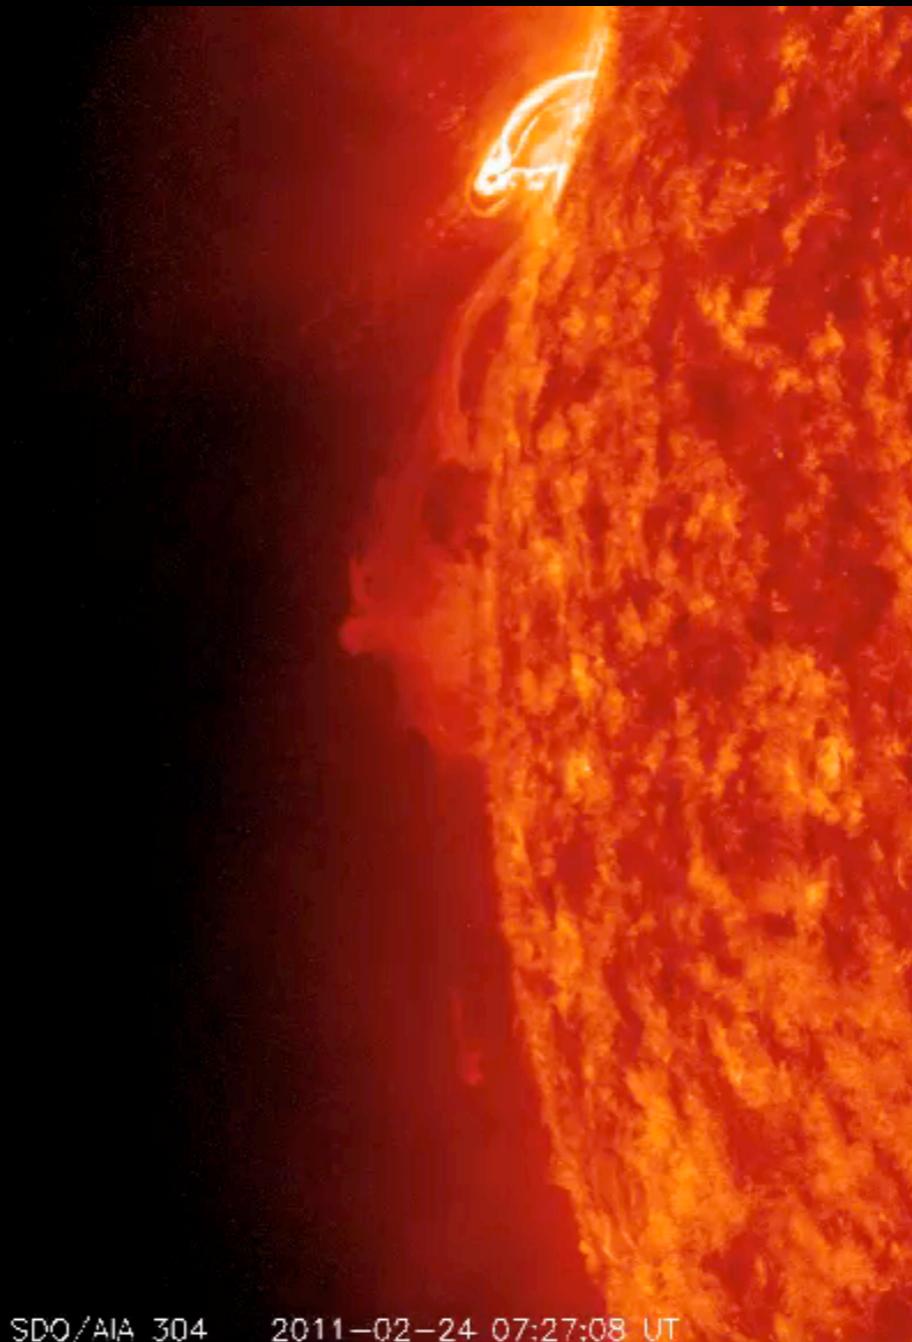


High-Throughput Radio Astronomy



Real-Time Steering

- Predict solar flares and coronal mass ejections

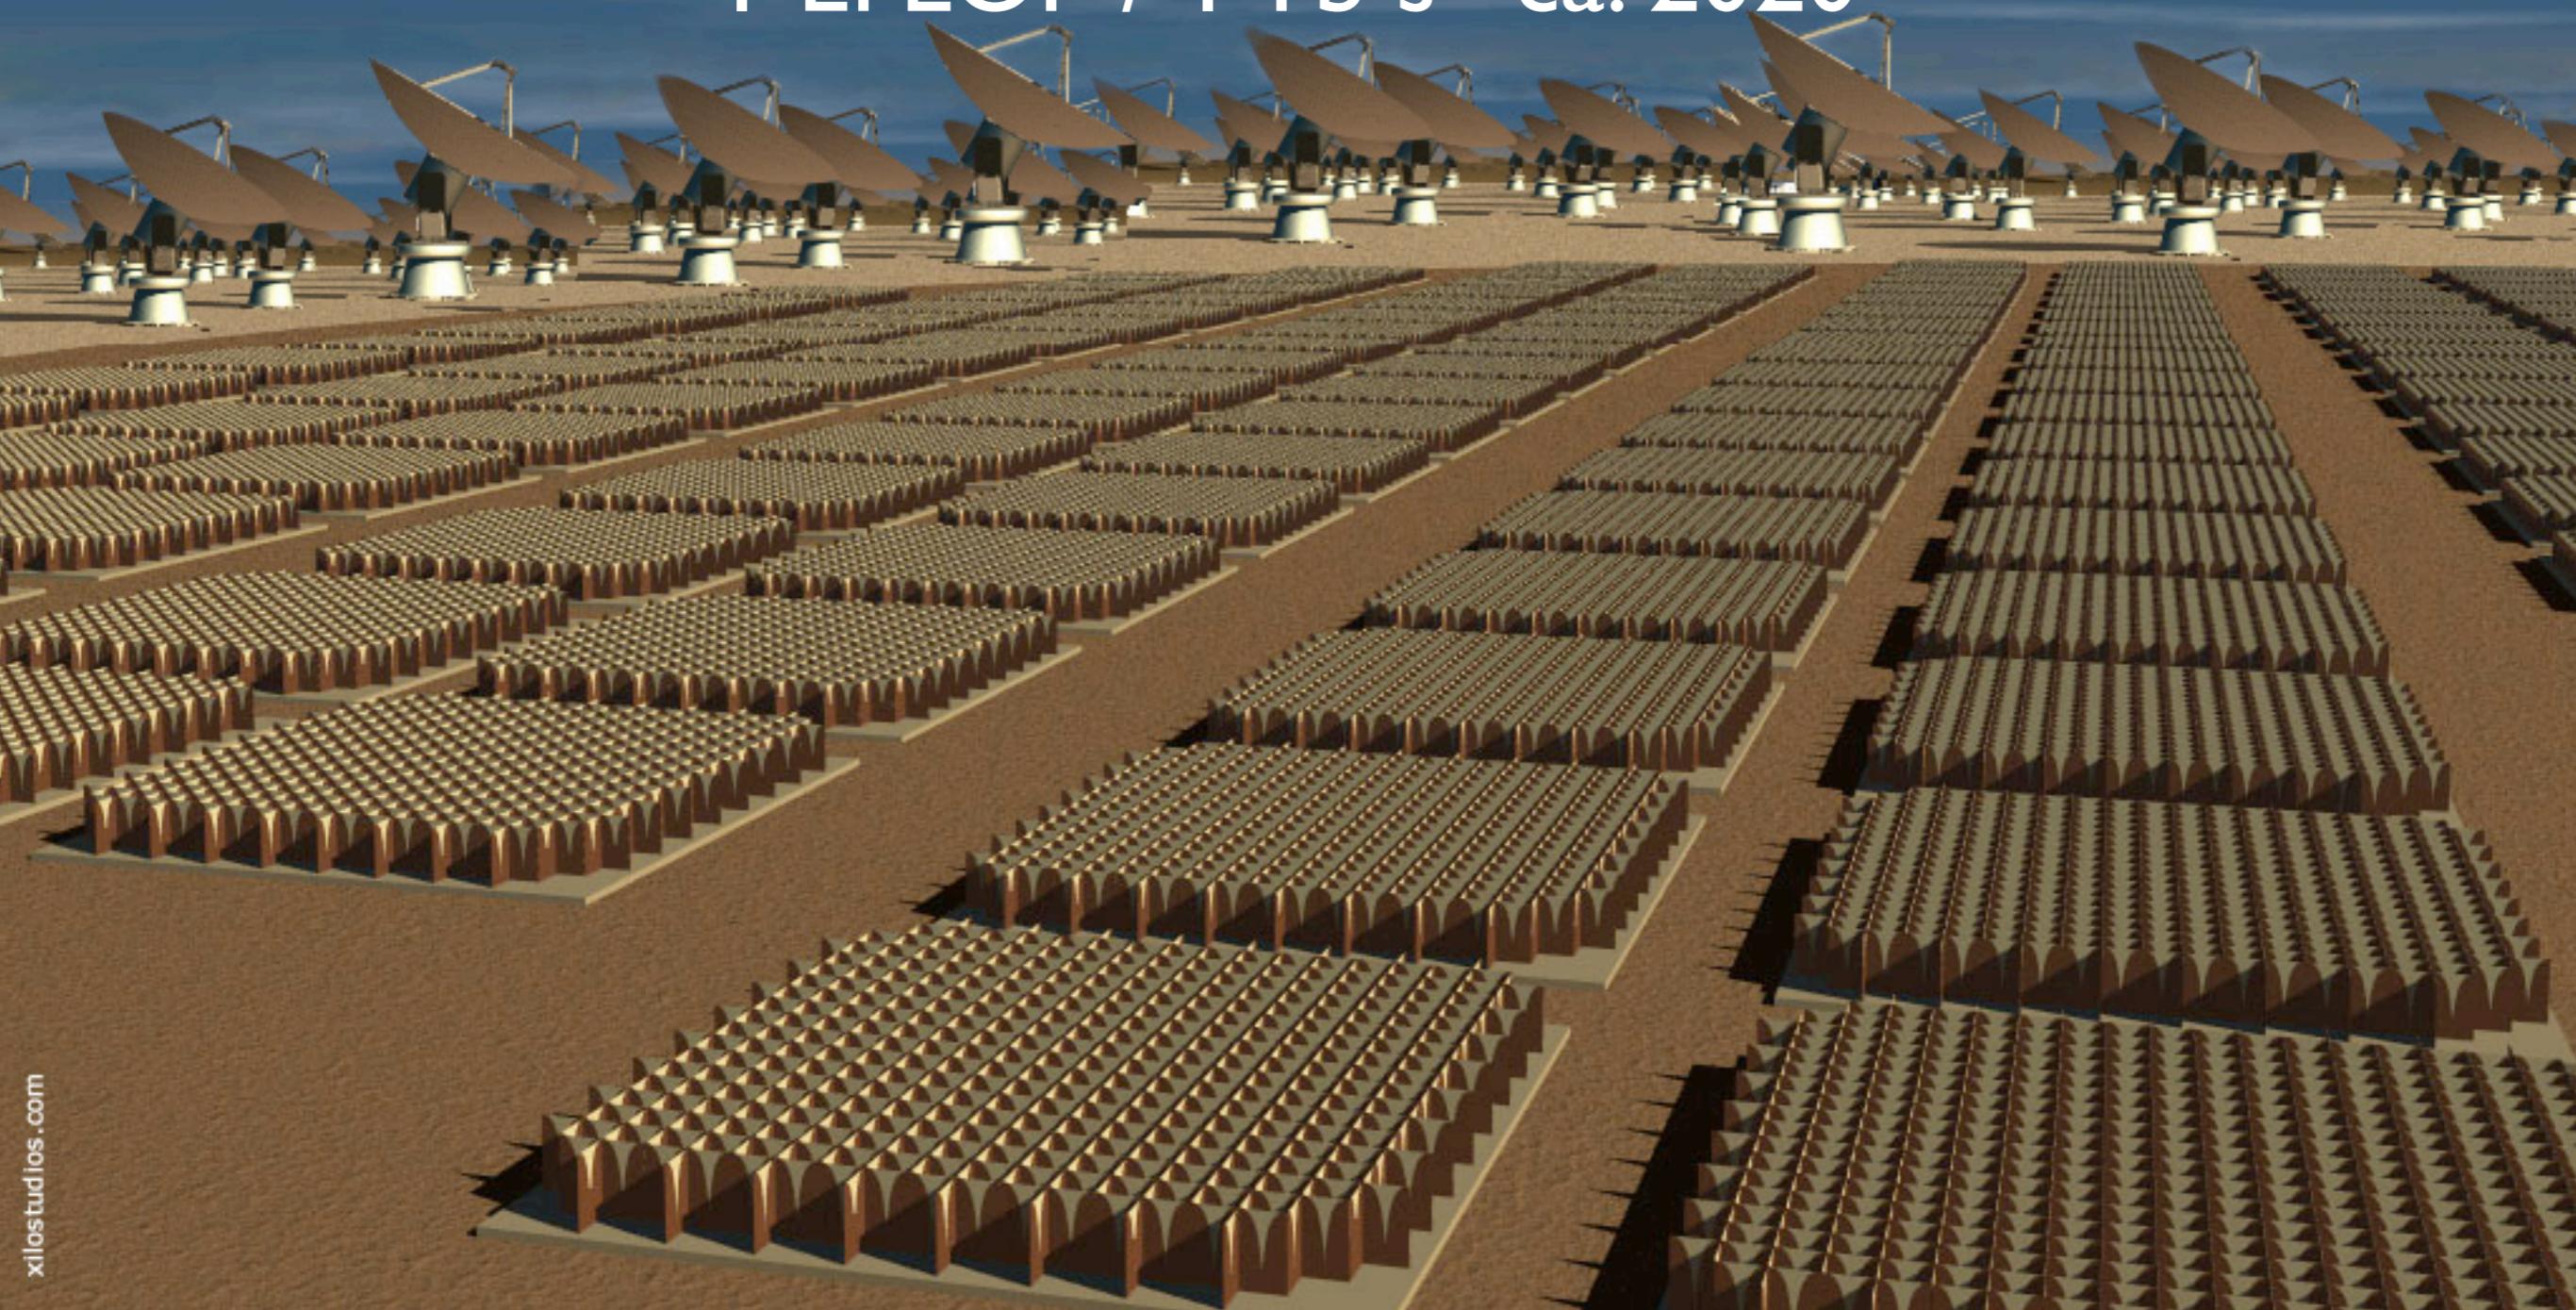


SDO (NASA)

The Future

Square Kilometer Array

1 EFLOP / 1 Tb s⁻¹ ca. 2020



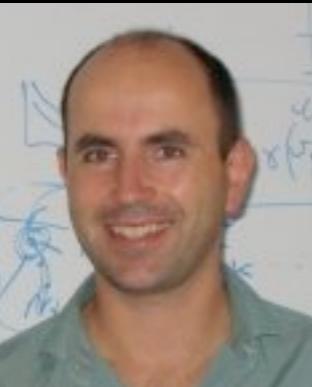
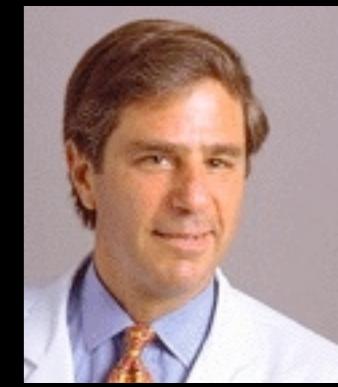
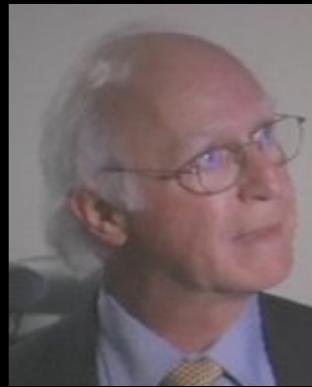


Multiscale Hemodynamics

Predicting & Preventing Heart Attacks

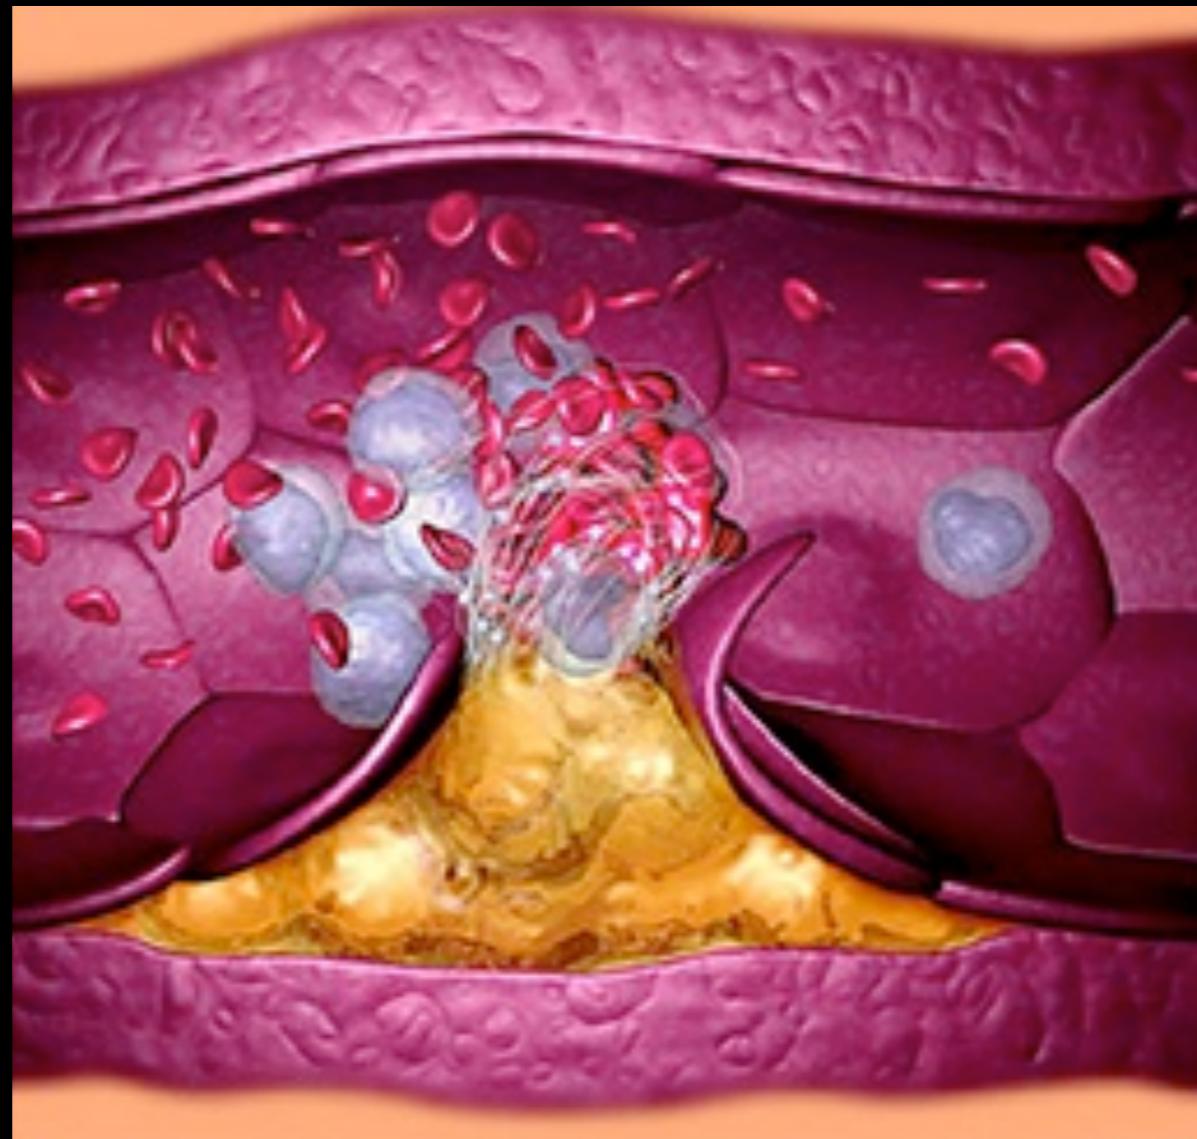
Multiscale Hemodynamics

Massimo Bernaschi -
Michelle Borkin - Ahmet
Coskun - Charles Feldman
- Efthimios Kaxiras -
Simone Melchionna -
Dimitris Mitsouras -
Hanspeter Pfister - Frank
Rybicki - Joy Sircar -
Michael Steigner - Peter
Stone - Sauro Succi -
Frederick Welt - Amanda
Whitmore

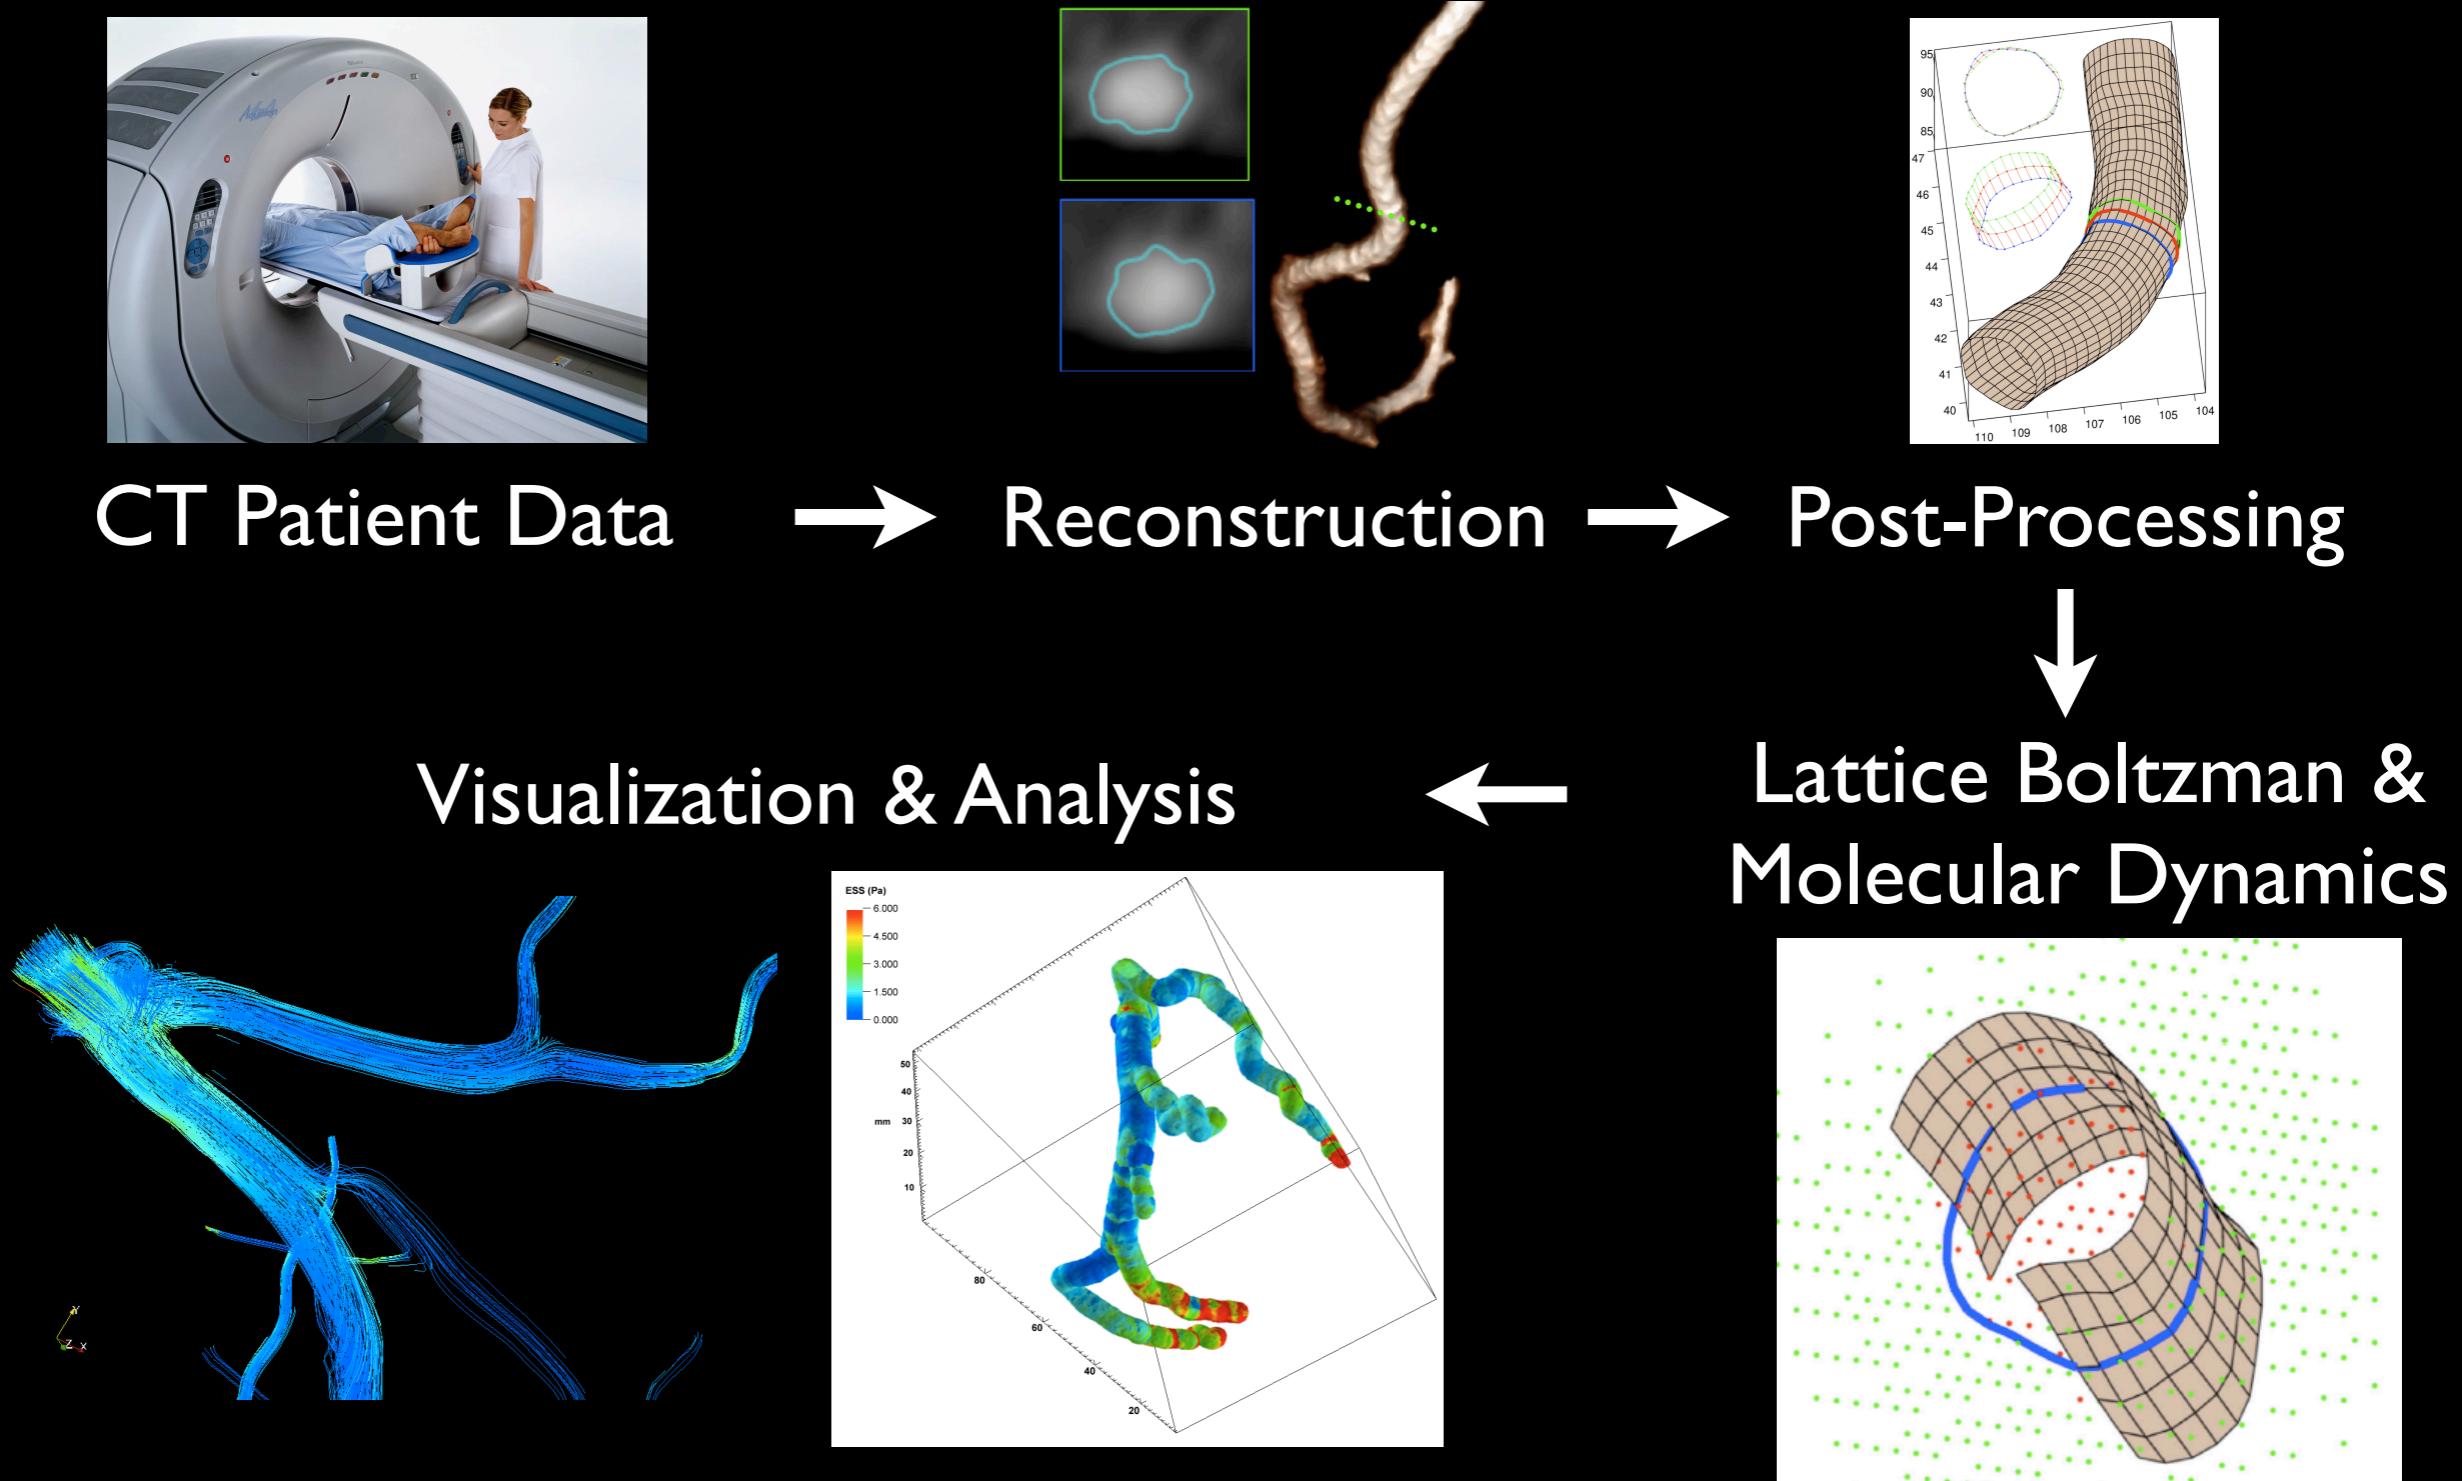


The Scientific Challenge

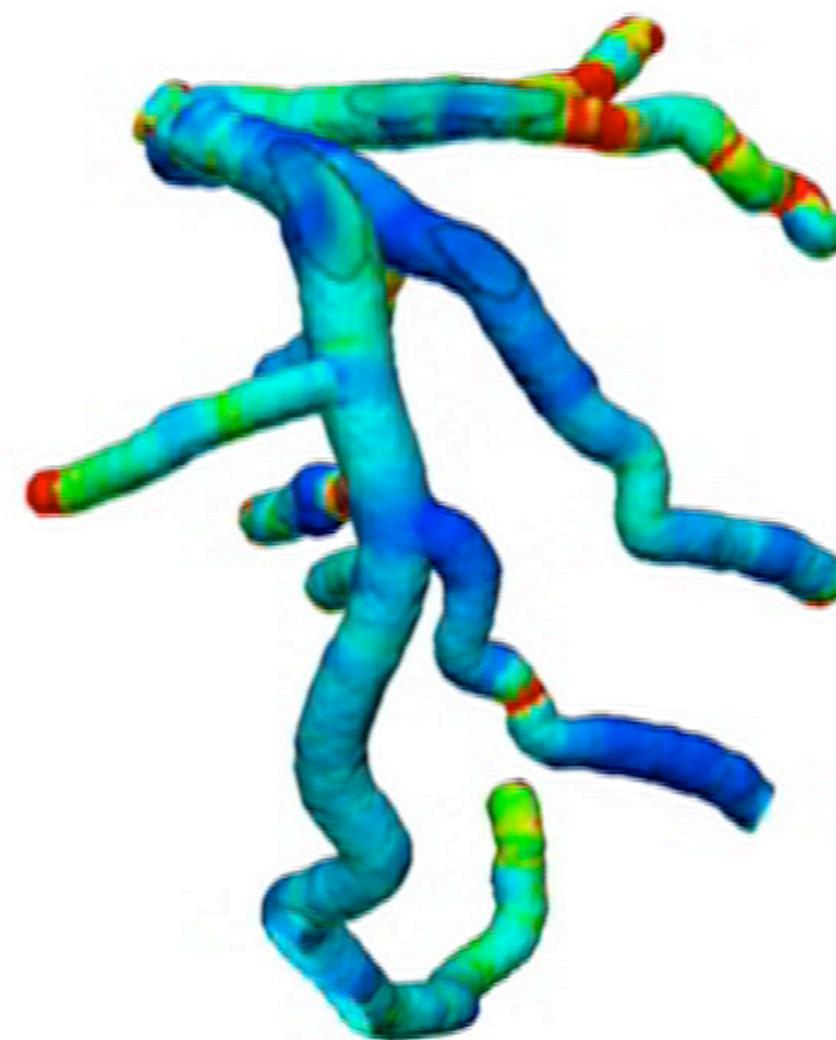
- Predict and prevent heart attacks



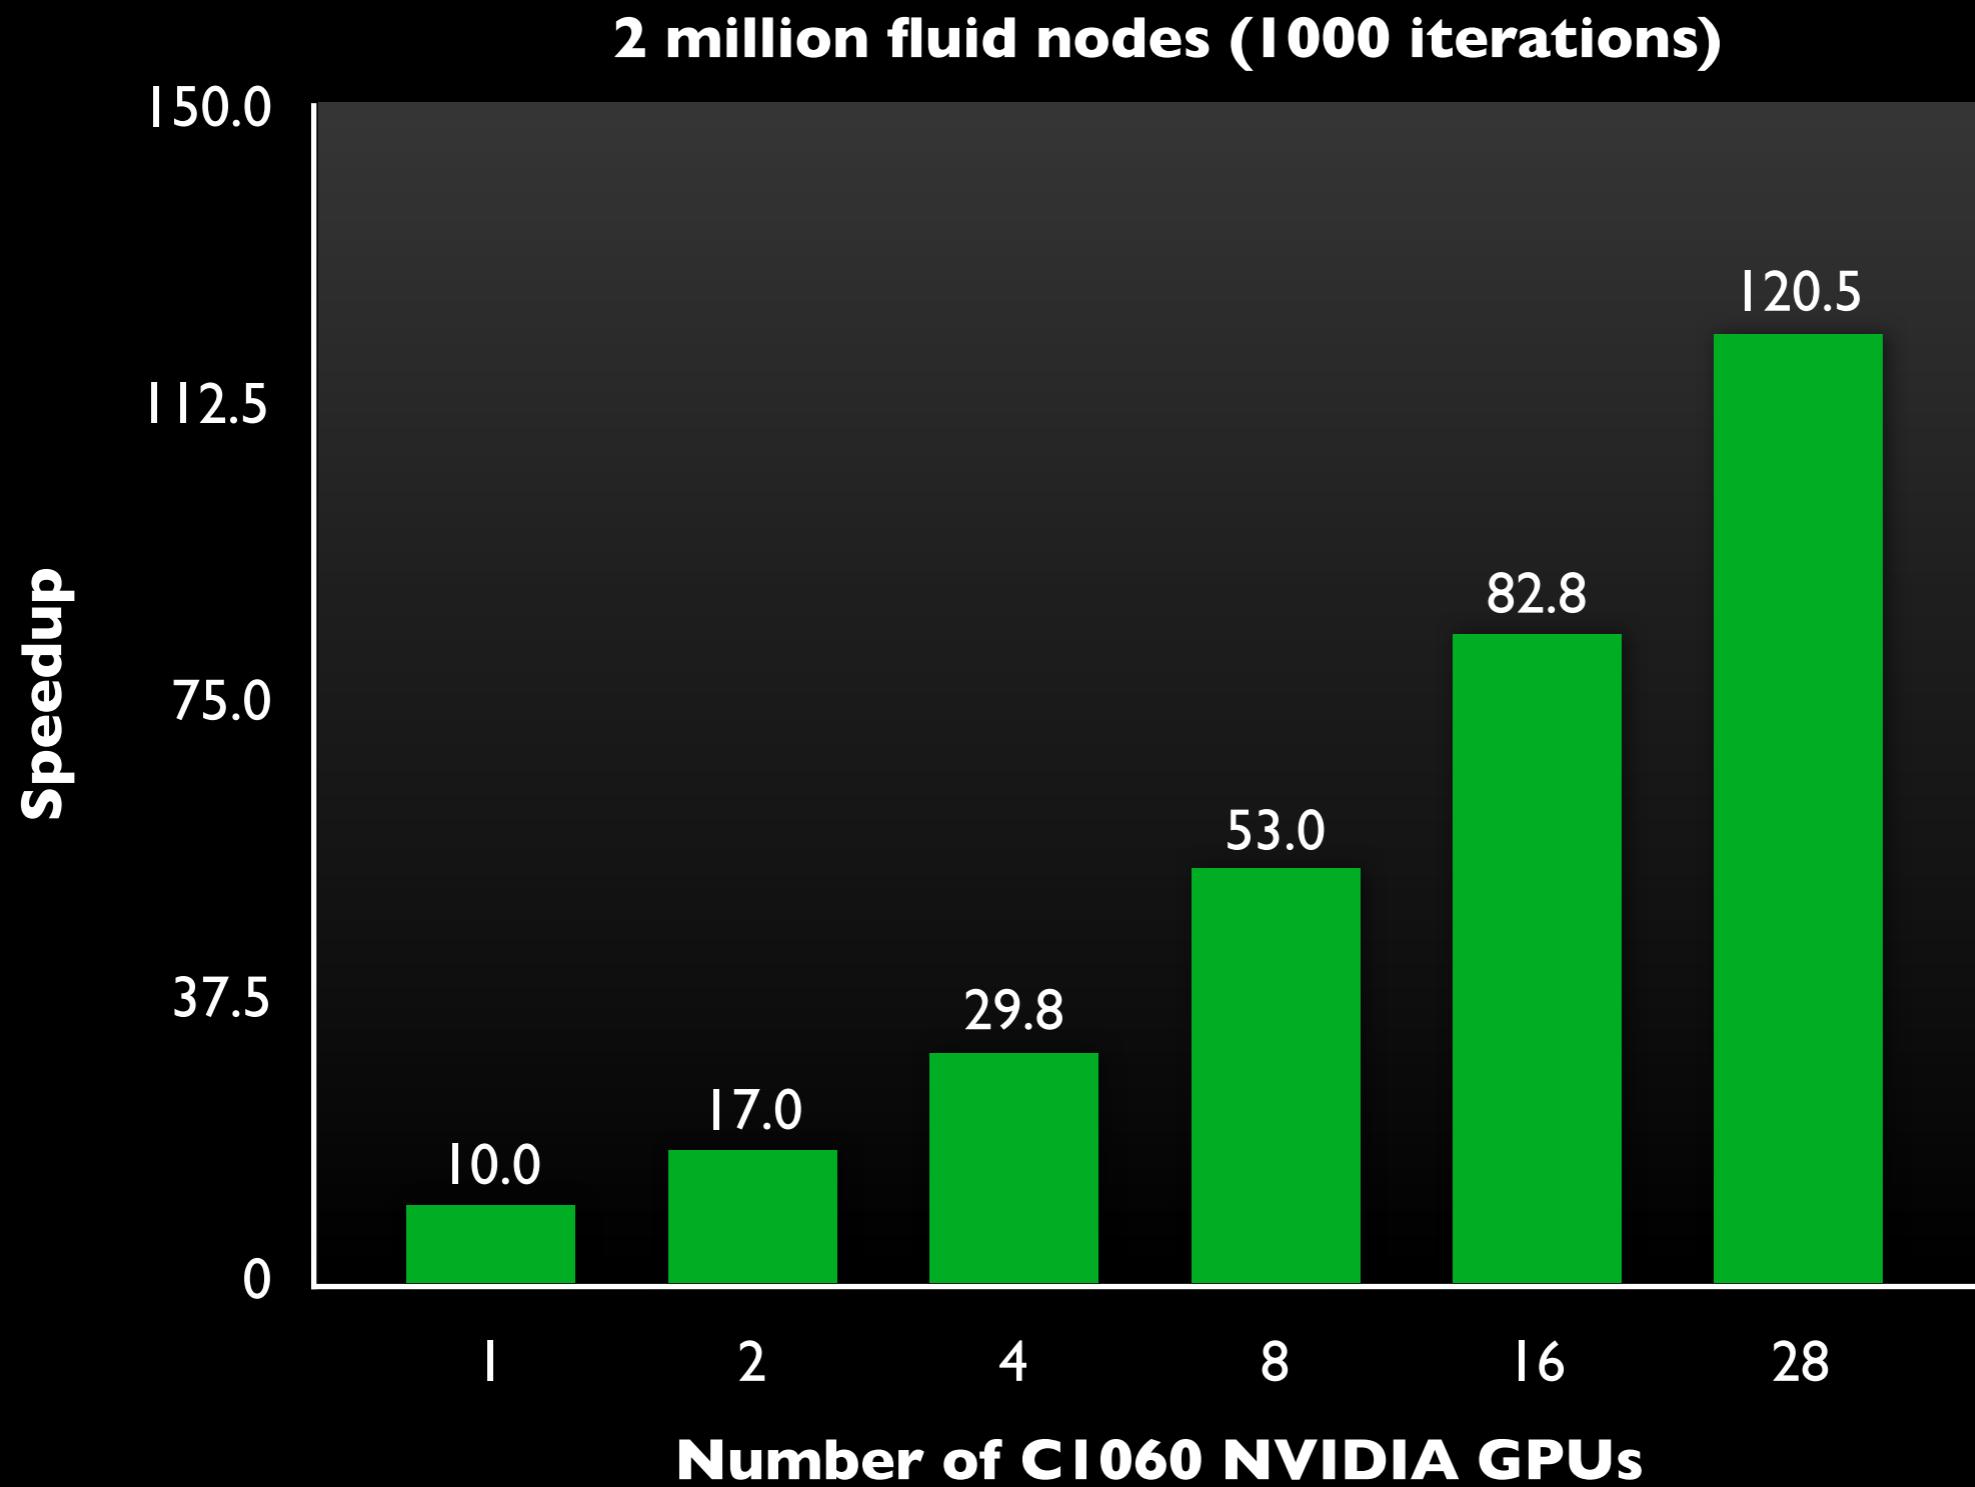
Hemodynamics Pipeline



Multiscale Hemodynamics



Results

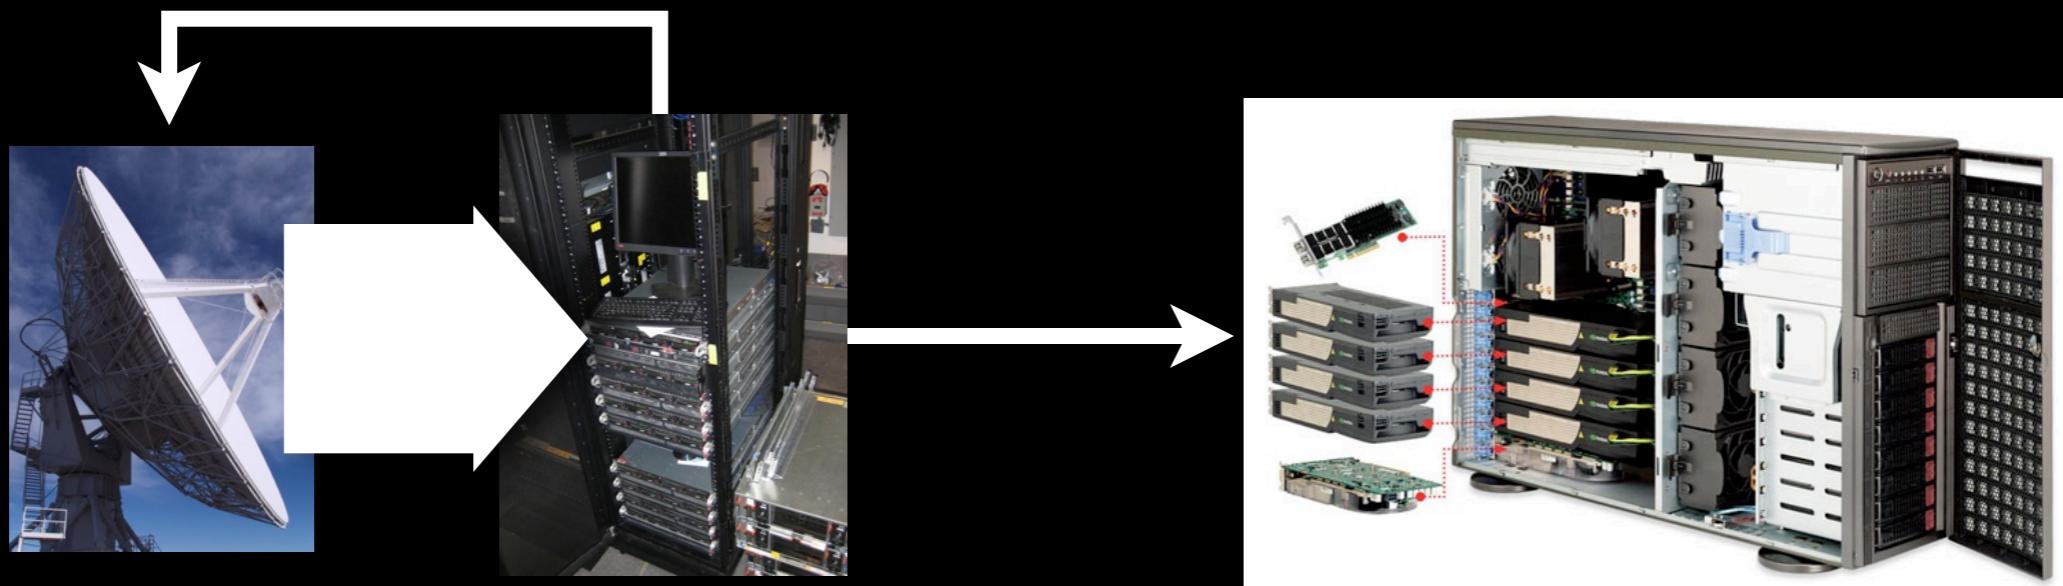


The Future

- Minimally invasive early detection

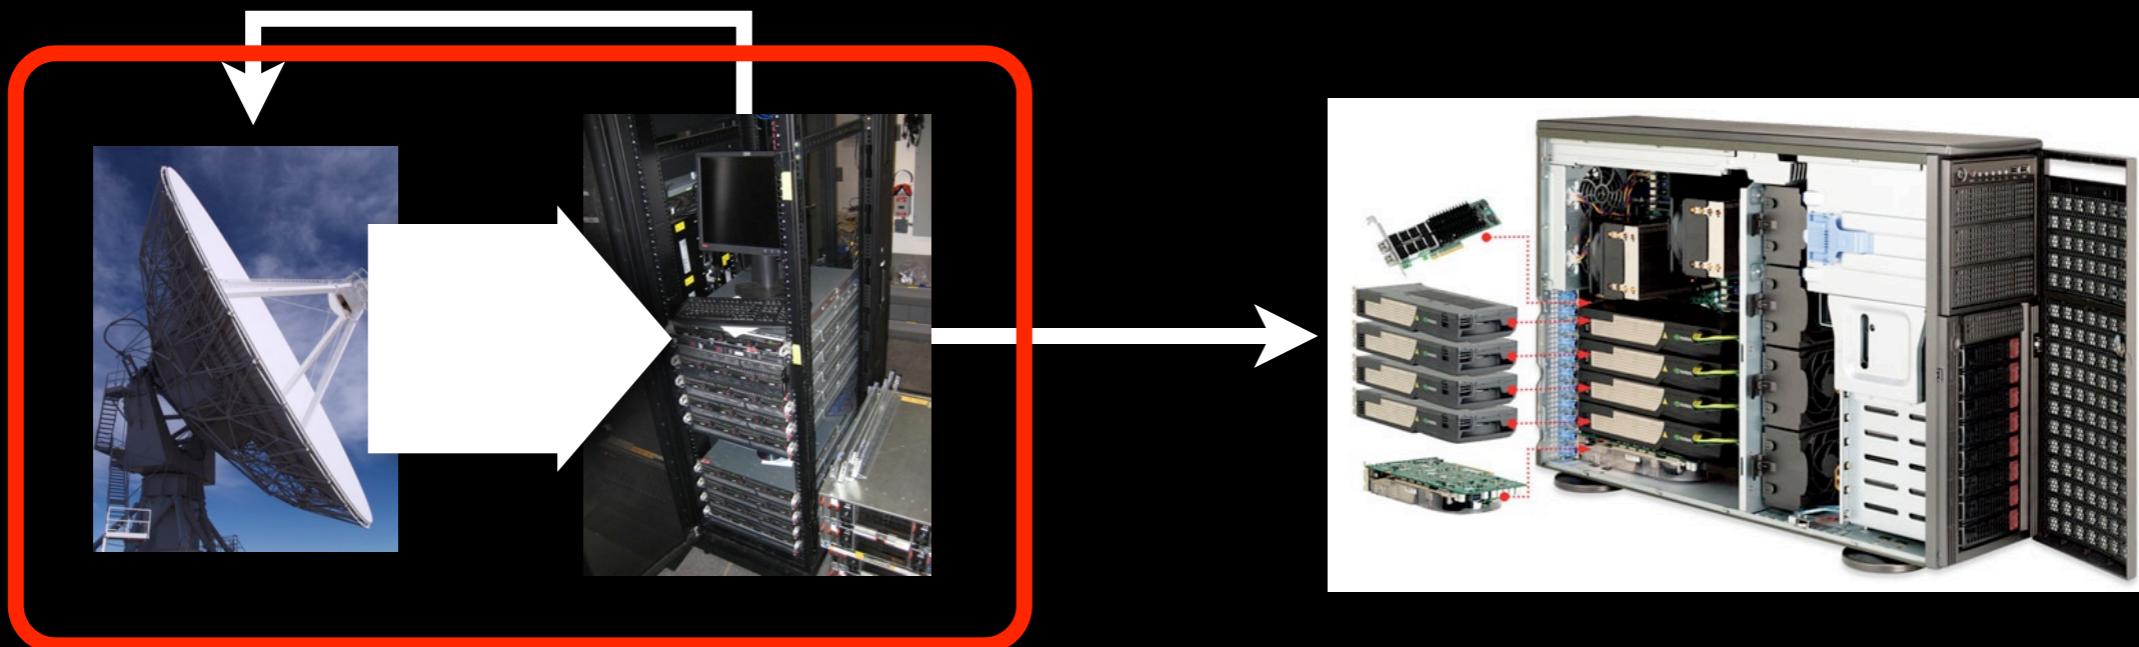


HTC Lessons



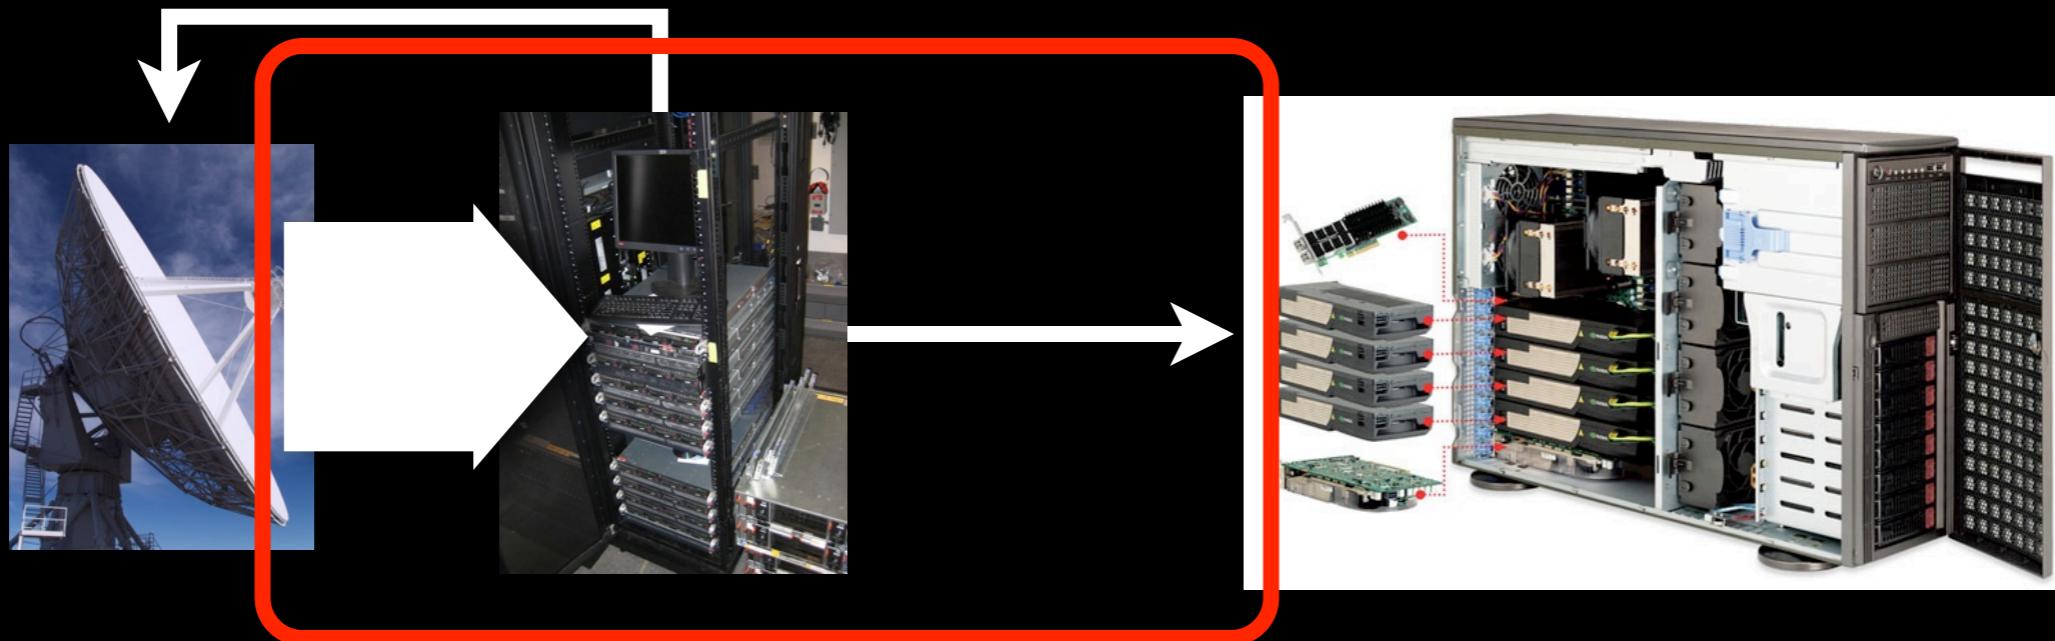
HTC Lessons

- Move the computation to the source



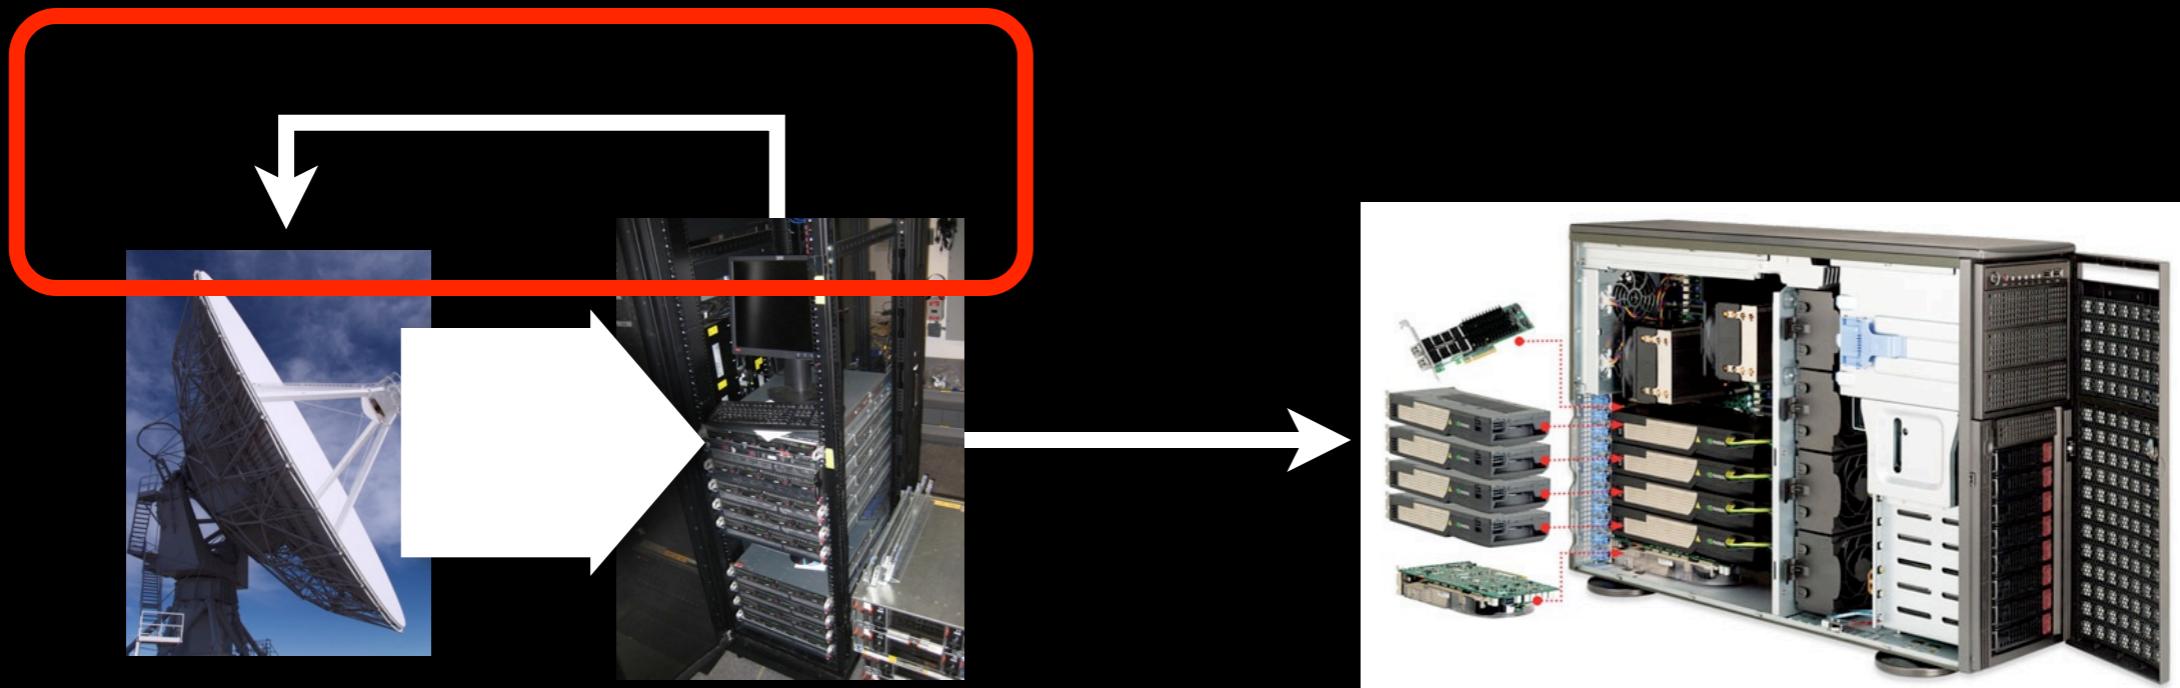
HTC Lessons

- Quickly reduce the data to its essentials



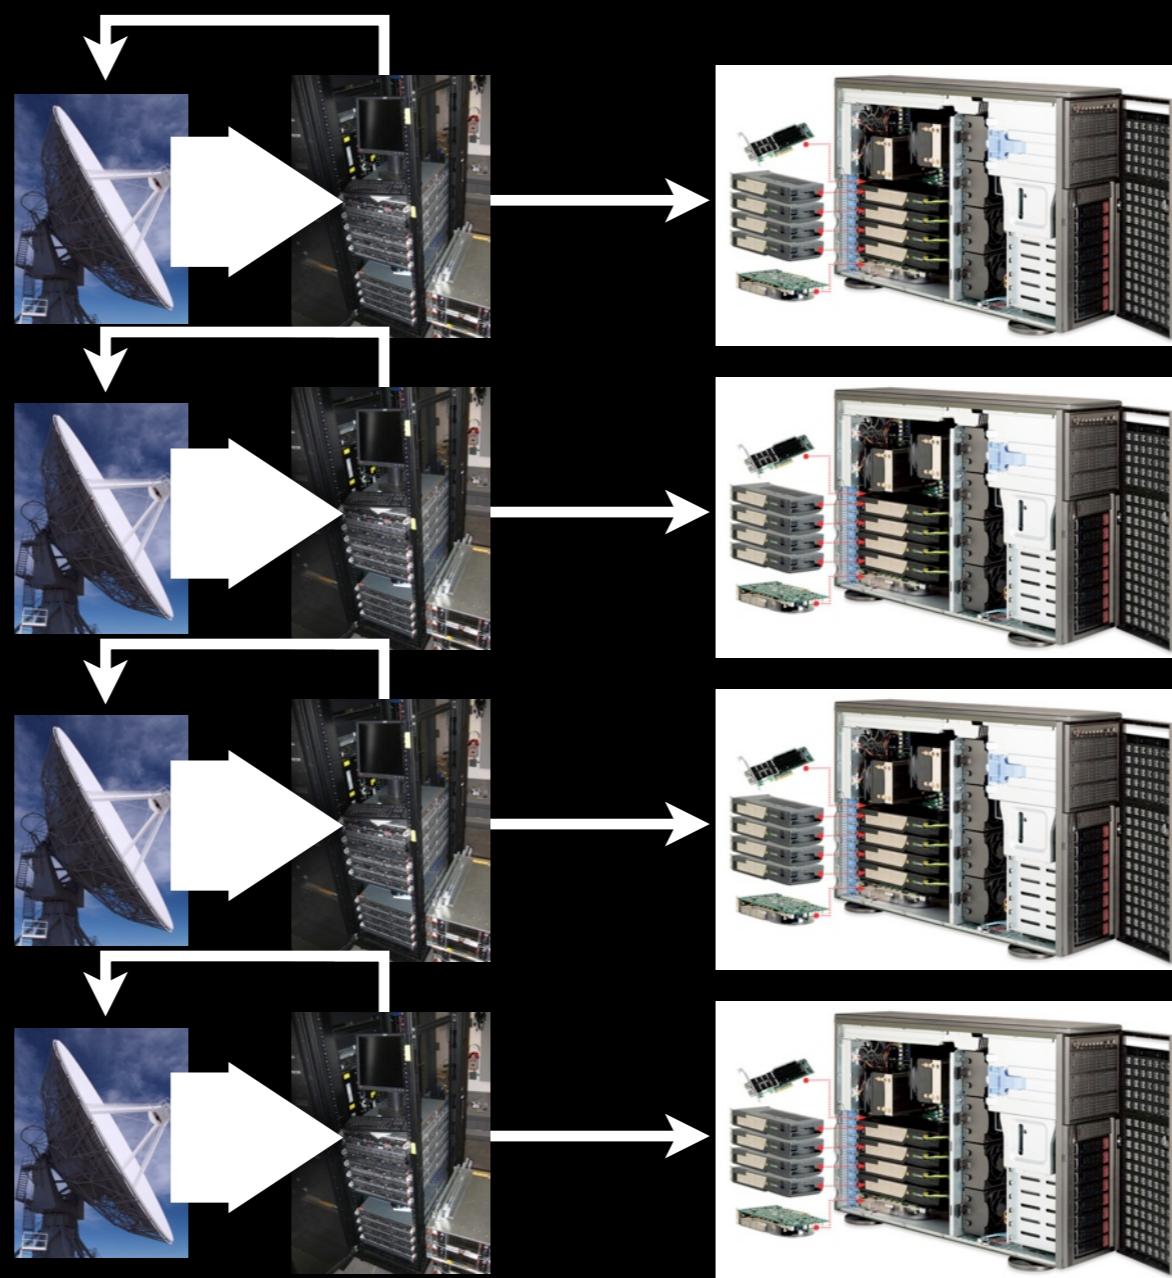
HTC Lessons

- Use feedback to optimize the acquisition



HTC Lessons

- Provides scalable and power efficient systems



Higher-Level Programming Models

- Domain Specific Languages (DSLs)
- E.g., Tensor Contraction Engine (TCE)

$$T1_{bcdf} = \sum_{el} B_{befl} \times D_{cdel}$$

$$T2_{bcjk} = \sum_{df} T1_{bcdf} \times C_{dfjk}$$

$$S_{abij} = \sum_{ck} T2_{bcjk} \times A_{acik}$$

(a) Formula sequence

```
T1=0; T2=0; S=0;
for b, c, d, e, f, l
[ T1bcdfl += Bbefl Dcdel
for b, c, d, f, j, k
[ T2bcjk += T1bcdfl Cdfjk
for a, b, c, i, j, k
[ Sabij += T2bcjk Aacik
```

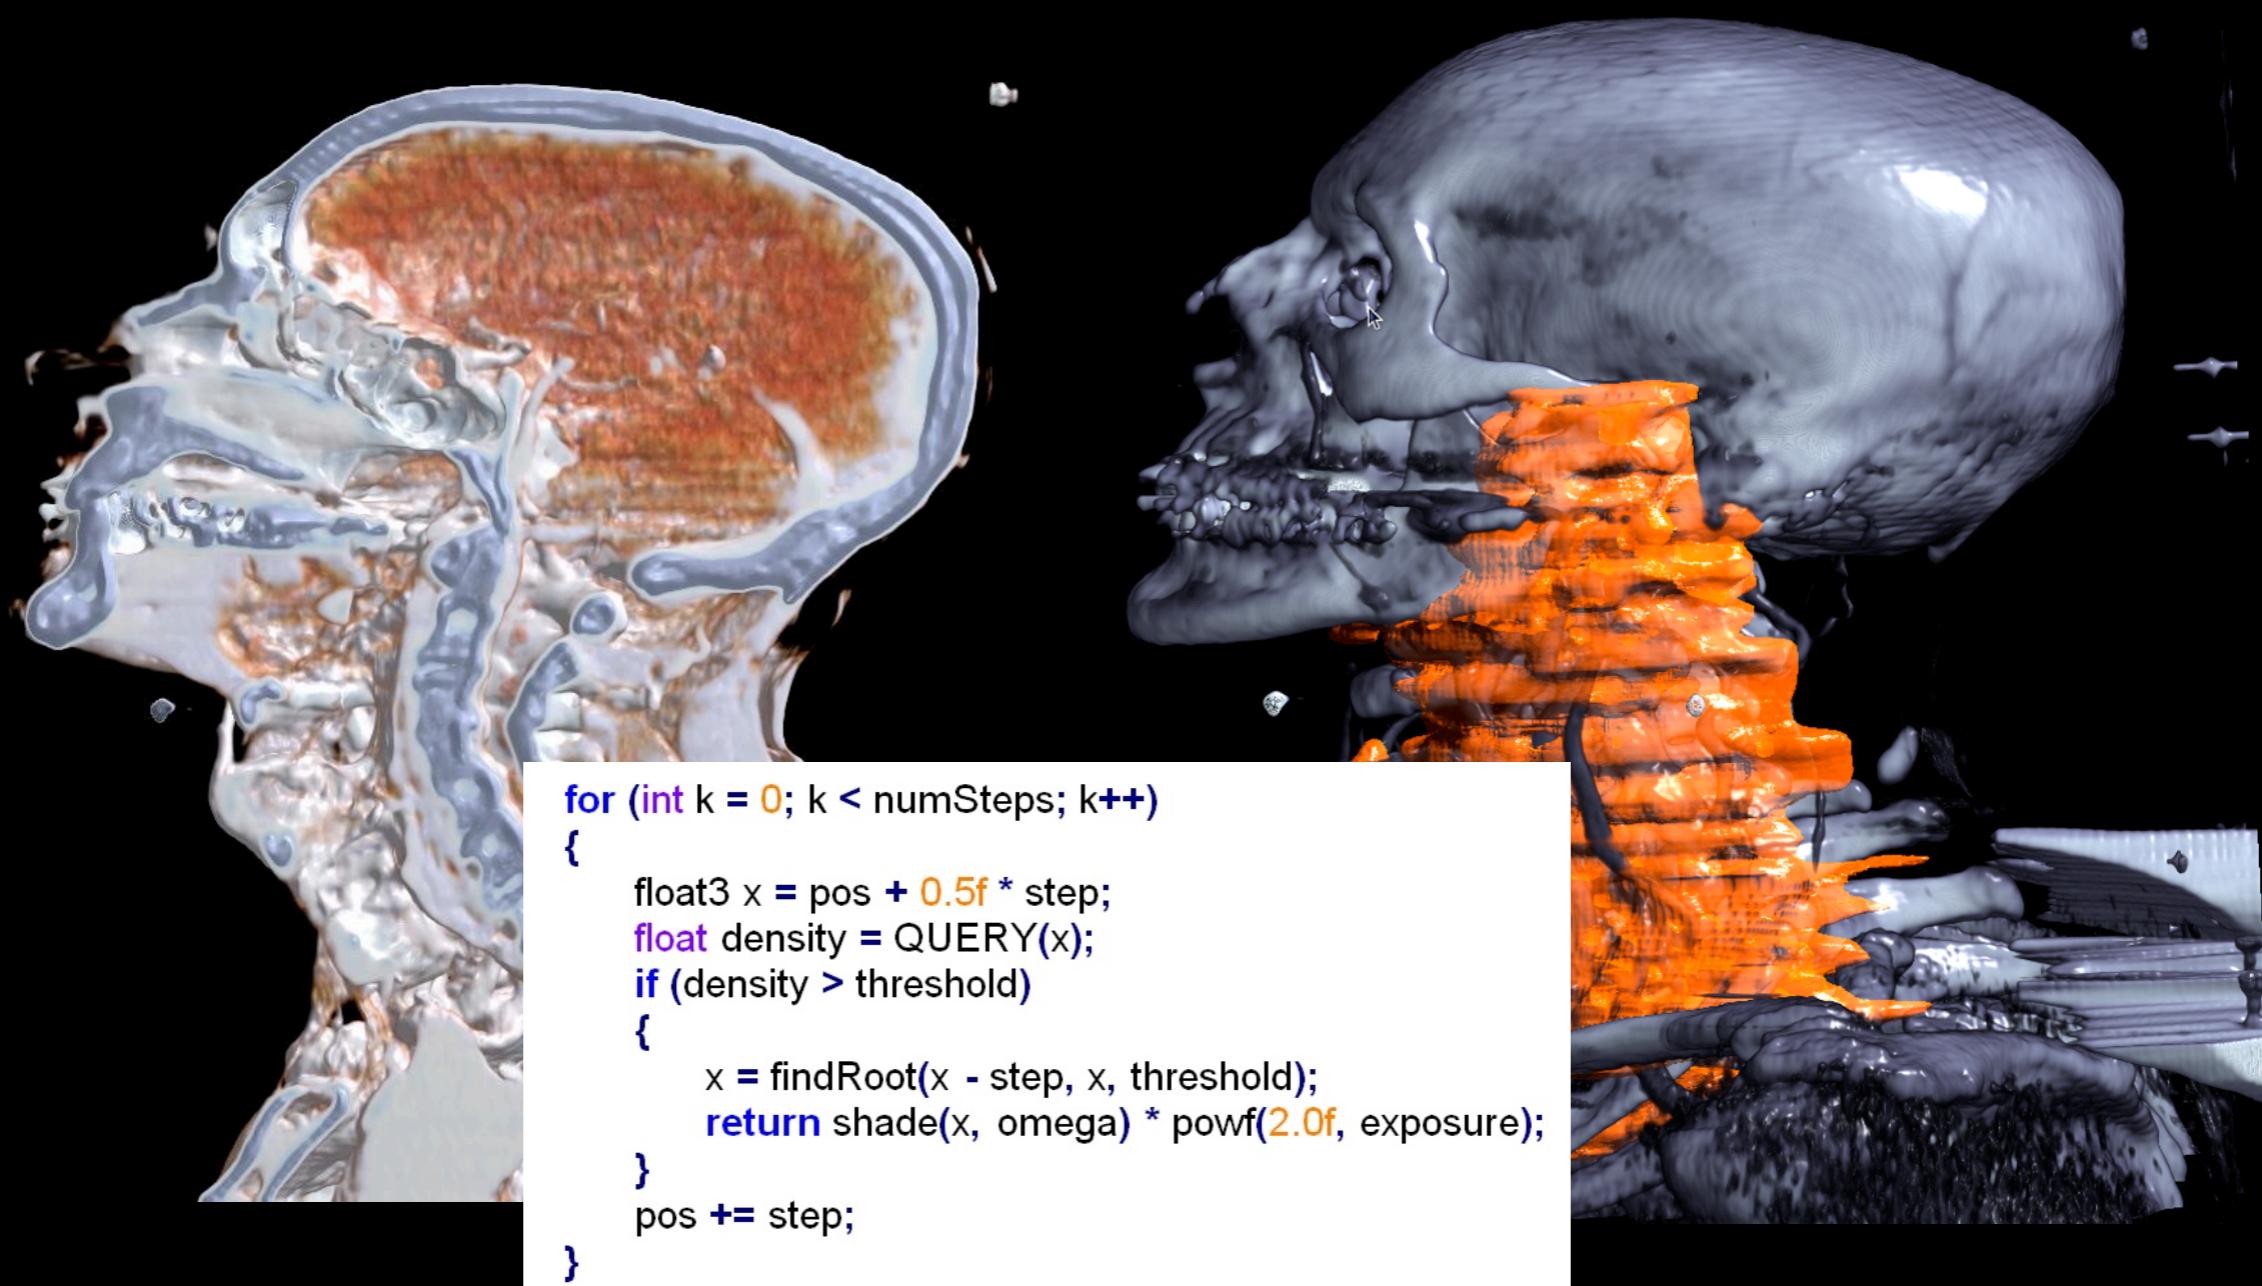
(b) Direct implementation (unfused code)

```
S = 0;
for b, c
[ T1f = 0; T2f = 0;
for d, f
[ for e, l
[ T1f += Bbefl Dcdel
for j, k
[ T2fjk += T1f Cdfjk
for a, i, j, k
[ Sabij += T2fjk Aacik
```

(c) Memory-reduced implementation (fused)

[Baumgartner et al. 05]

Volume Rendering with Procedural Shaders

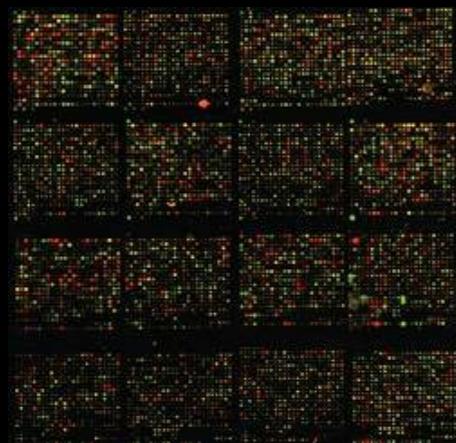


Scalable Runtime Environments



Chip Sensors

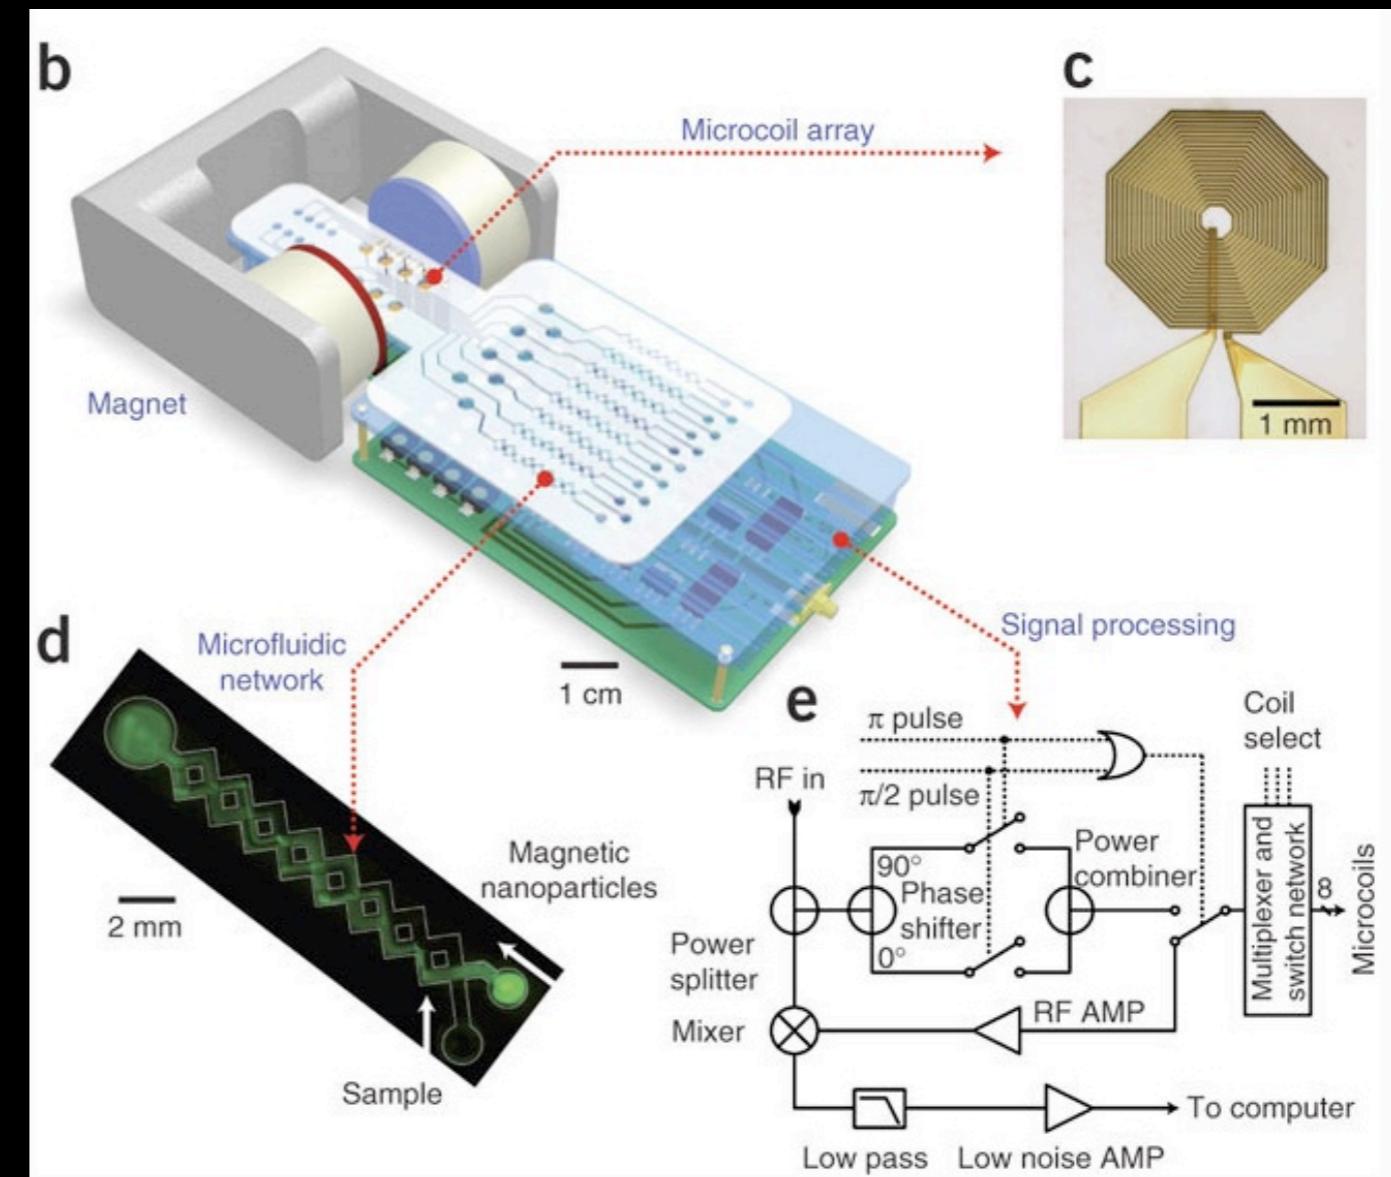
Gene Array Chips



George Whitesides

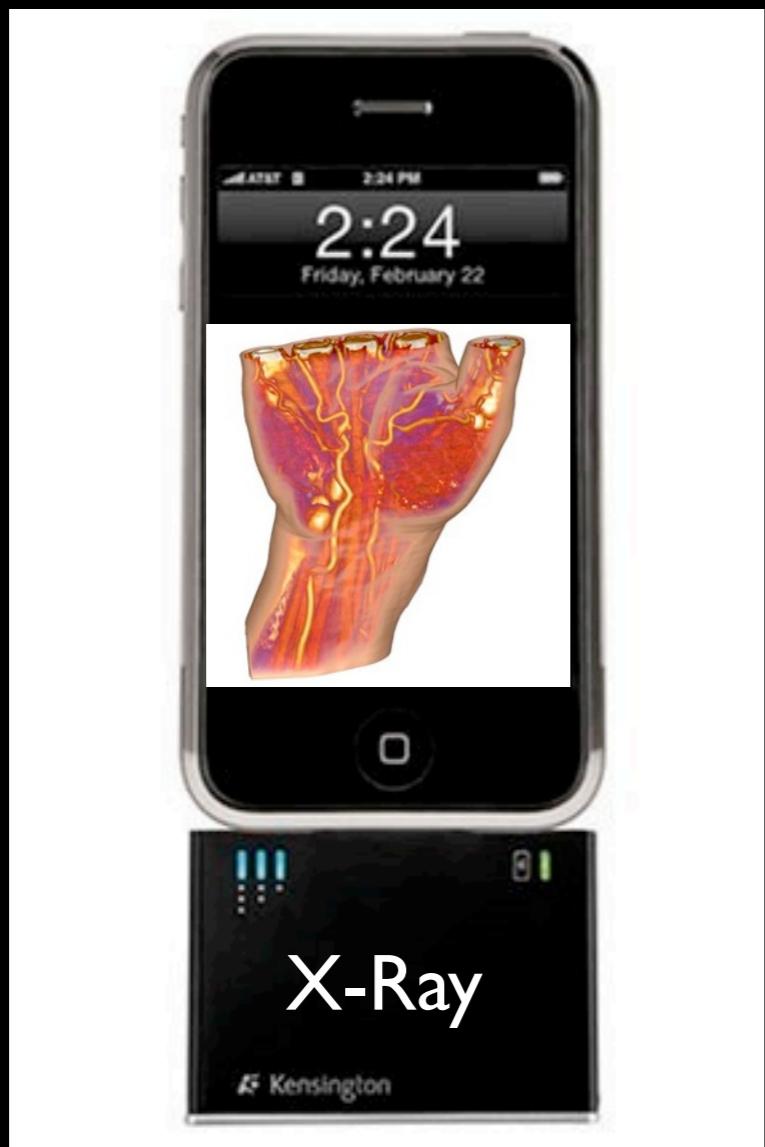


Chip-NMR Biosensor



Hakho Lee, Eric Sun, Donhee Ham & Ralph Weissleder, SEAS & HMS

Tricorders



21st Century



23rd Century

Acknowledgements

- NVIDIA
- Microsoft
- NSF CDI PHY-0835713, Austrian Research Promotion Agency FFG, Vienna Science and Technology Fund WWTF

Shameless Plugs

- cs264.org - Massively Parallel Computing course (spring 2011, Harvard Extension)

