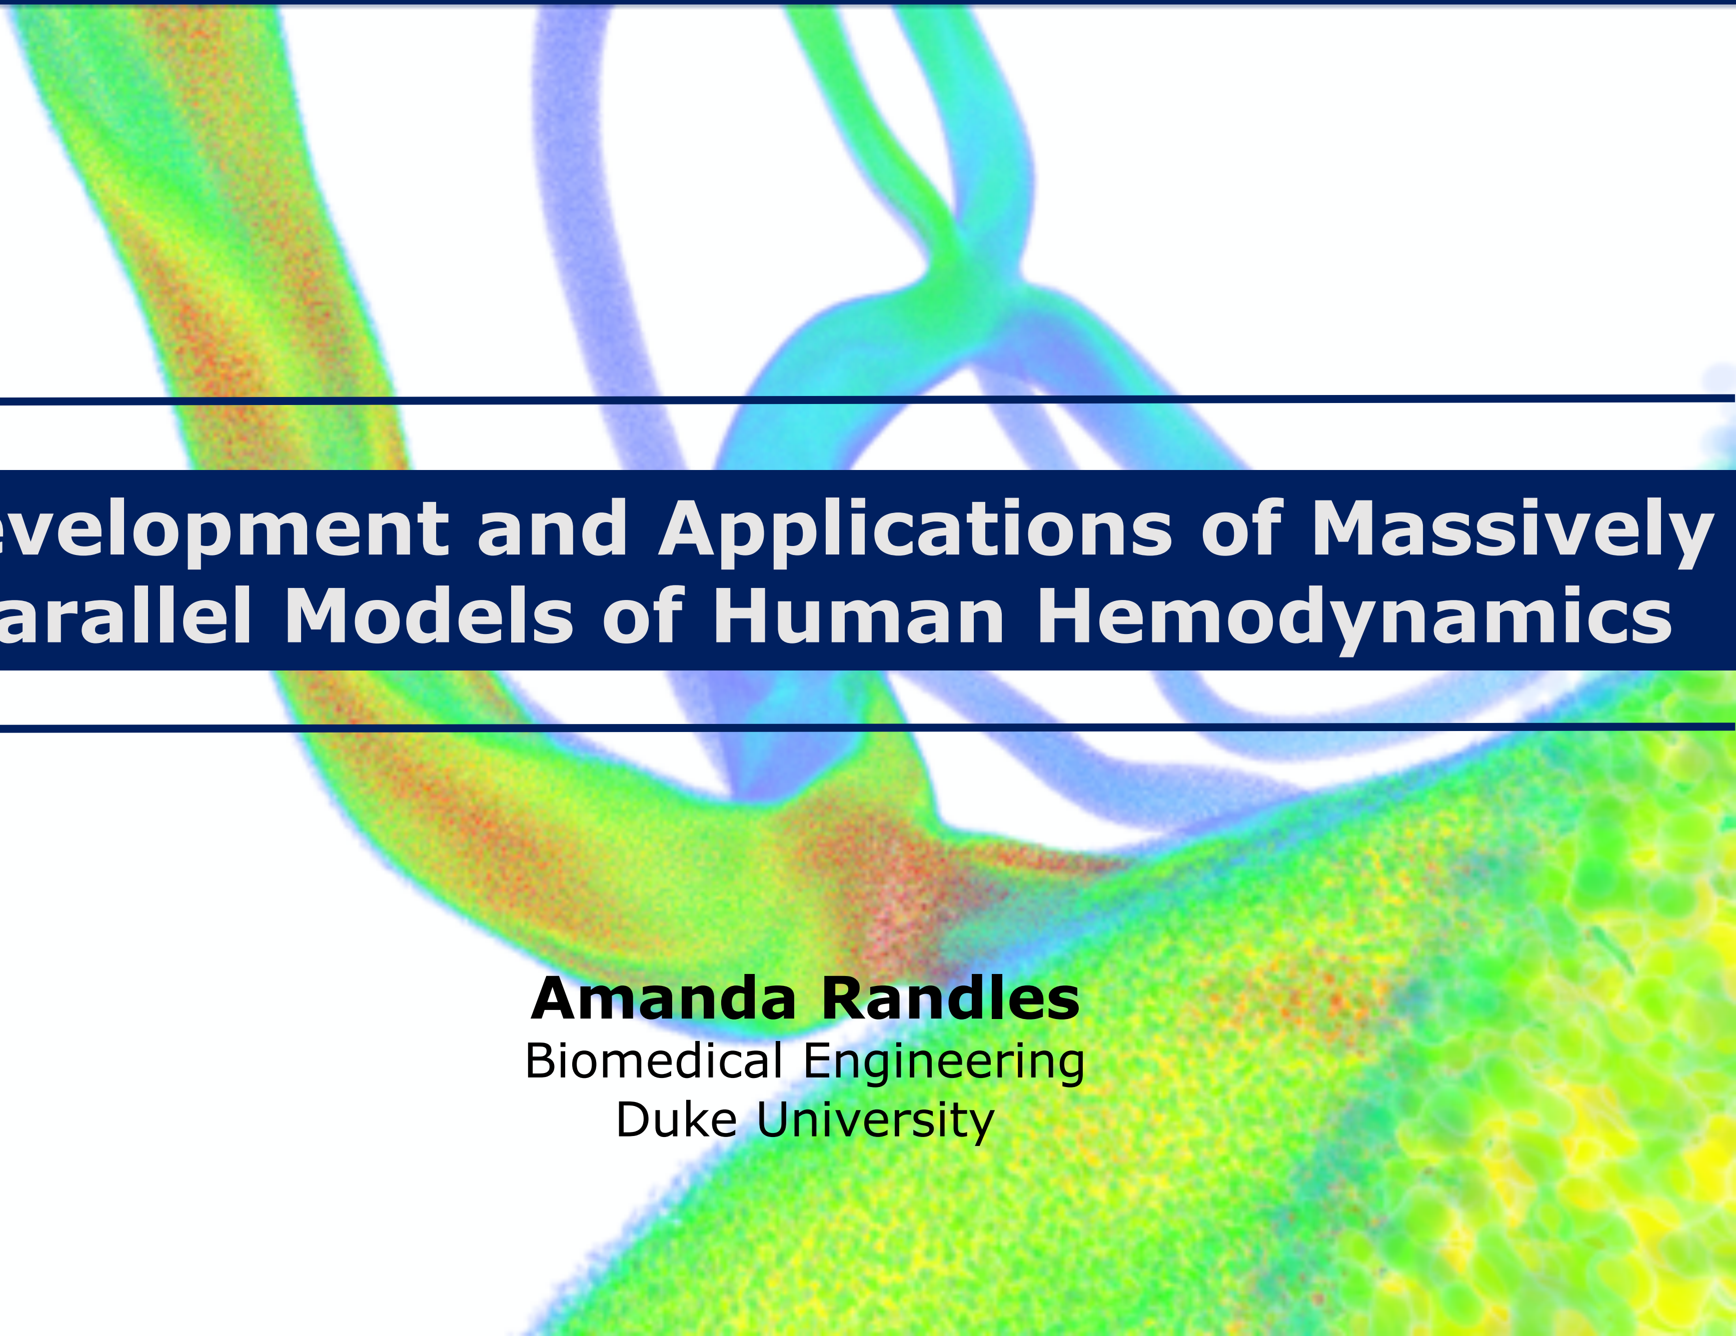

Development and Applications of Massively Parallel Models of Human Hemodynamics

Amanda Randles
Biomedical Engineering
Duke University

Why HARVEY?

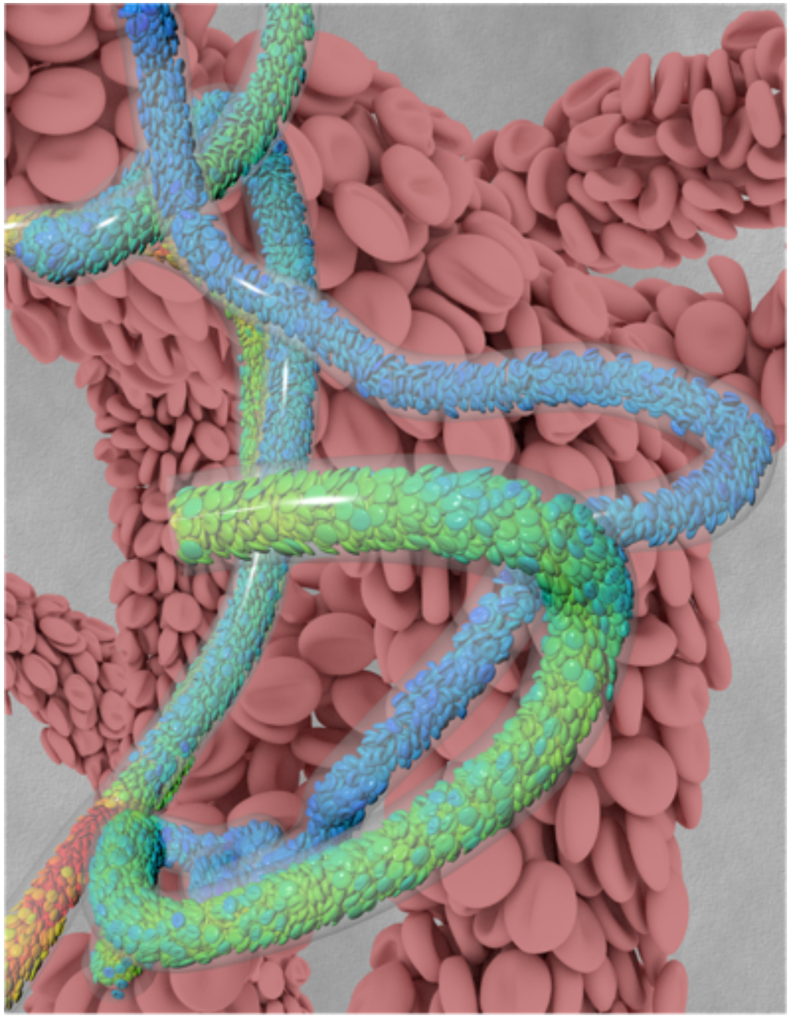

PATIENT-SPECIFIC COMPUTATIONAL MODELS

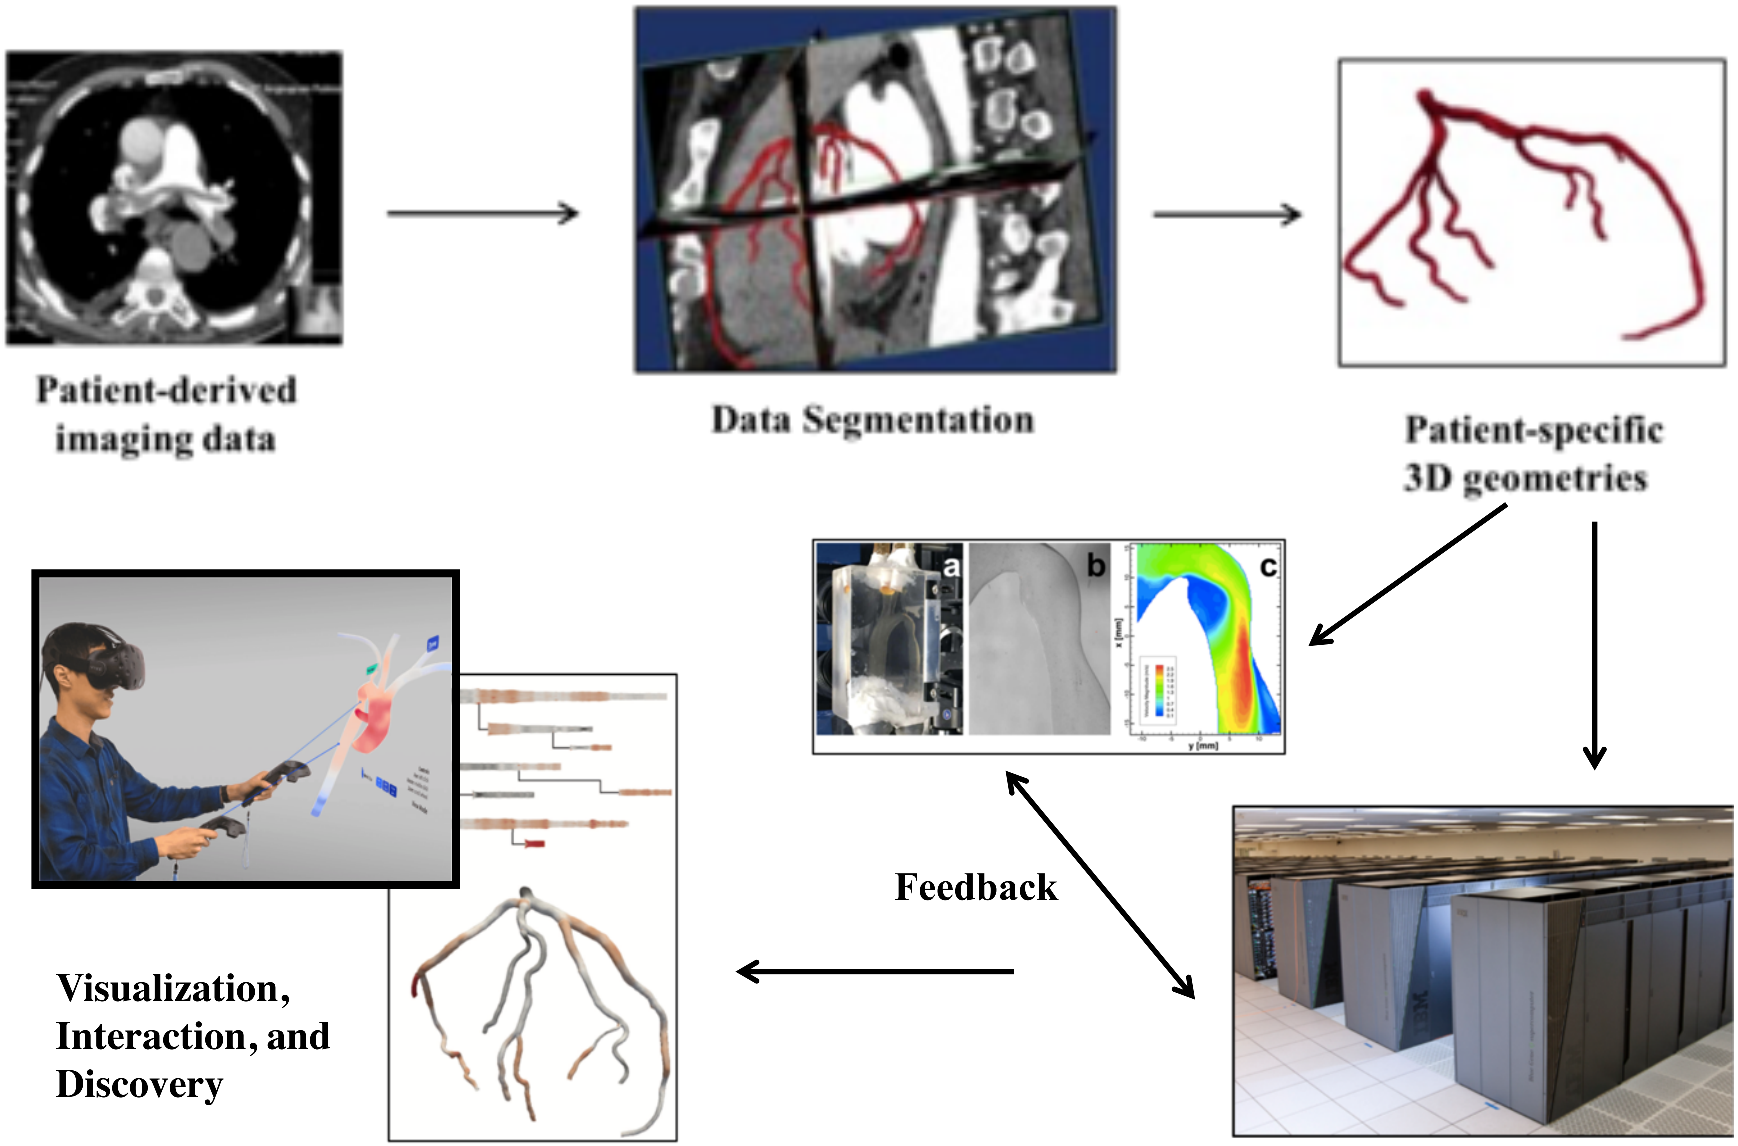

Fractional Flow Reserve

$$FFR = \frac{\text{Distal Coronary Pressure (Pd)}}{\text{Proximal Coronary Pressure (Pa)}}$$

(During Maximum Hyperemia)

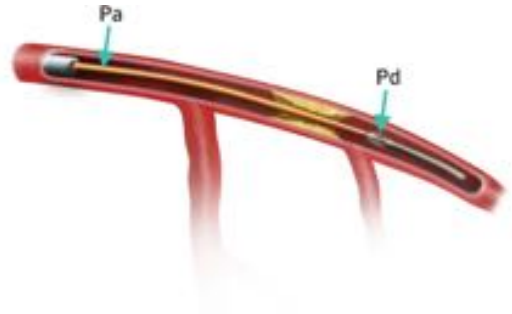

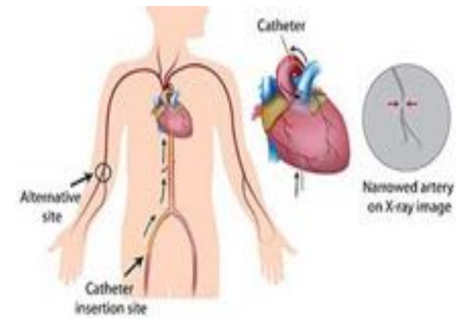

A

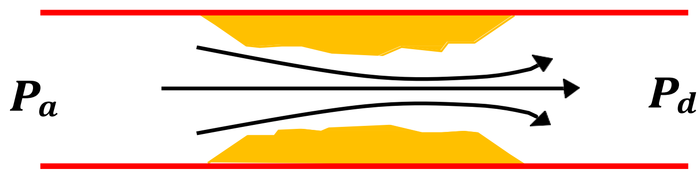

Hyperemia: Large gradient

B

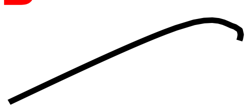

Pressure guide-wire

\+

Hyperemic drug

=

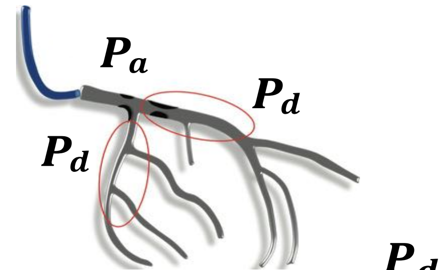

$$Invasive\ \boldsymbol{FFR} = \frac{P_d}{P_a}$$

FFR is the current gold standard for the assessment of coronary artery disease

RECONSTRUCTION

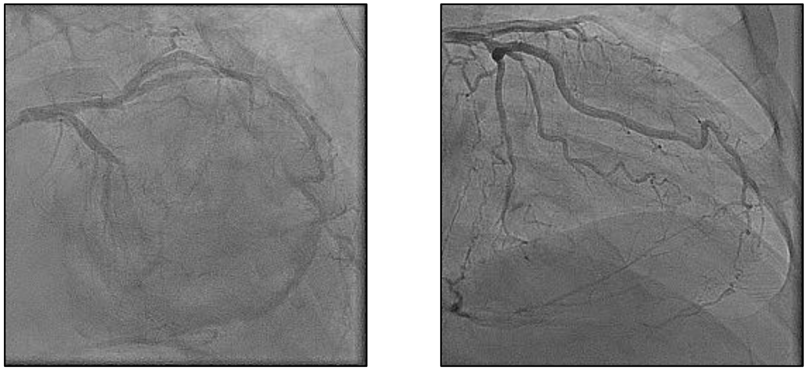

PAIR OF ANGIOGRAMS BASED ON TWO ACQUISITIONS C-ARM
LAO 41 CRANIAL 26 (left) RAO 22 CAUDAL 22 (right)

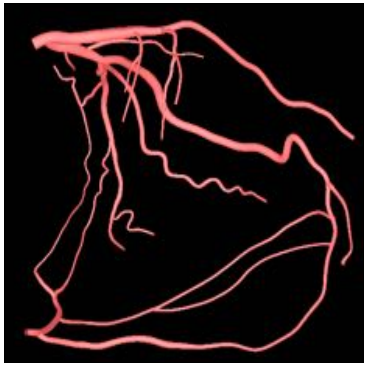

3D LCA MODEL

Vardhan *et al*. Nature Scientific Reports 2021

PERSONALIZED COMPUTATIONAL MODELS

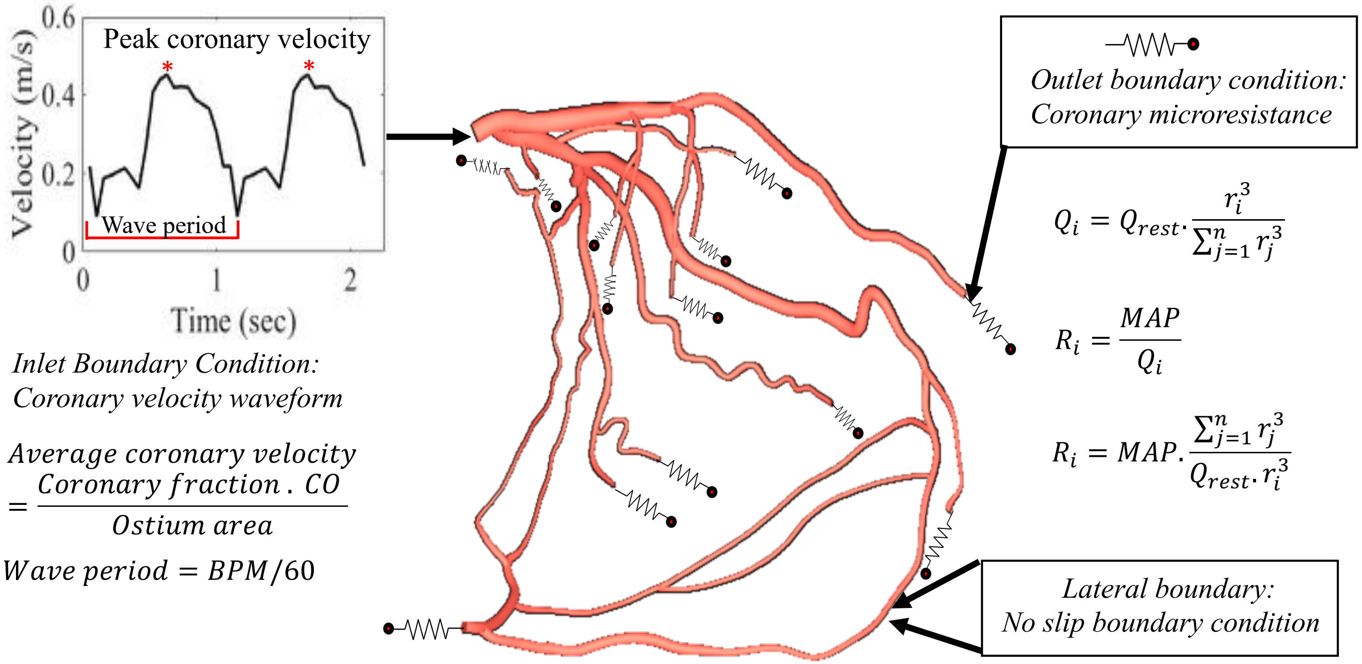

Inlet Boundary Condition:
Coronary velocity waveform

$$Average\ coronary\ velocity = \frac{Coronary\ fraction\ .\ CO}{Ostium\ area}$$

$$Wave\ period = BPM/60$$

$$Q_i = Q_{rest} \cdot \frac{r_i^3}{\sum_{j=1}^{n} r_j^3}$$

$$R_i = \frac{MAP}{Q_i}$$

$$R_i = MAP.\frac{\sum_{j=1}^{n} r_j^3}{Q_{rest}.\, r_i^3}$$

Vardhan *et al*. Nature Scientific Reports 2021

RESTING AND HYPEREMIC STATES

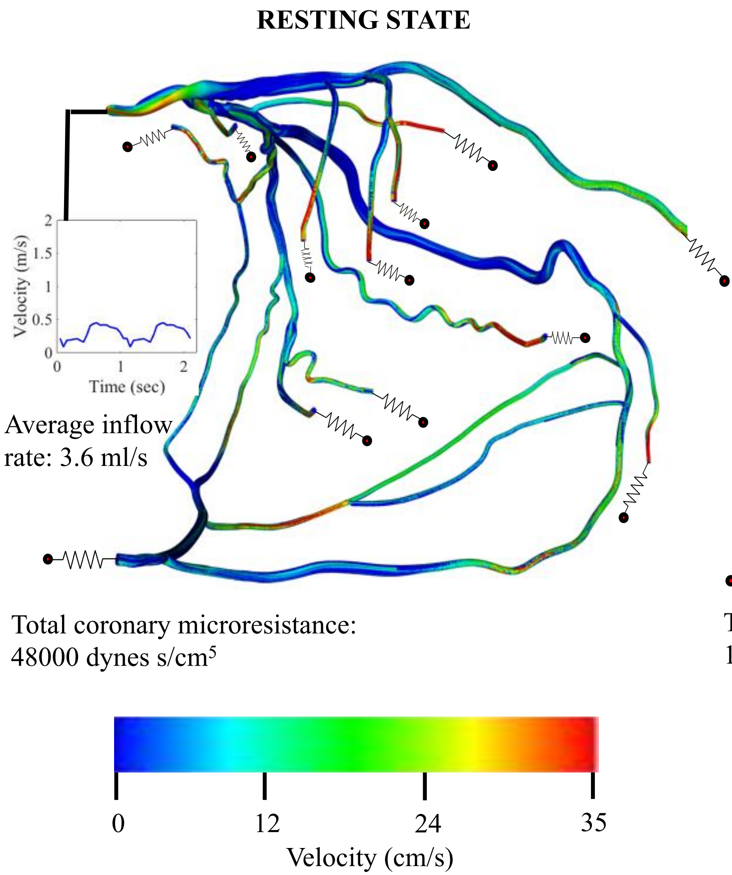

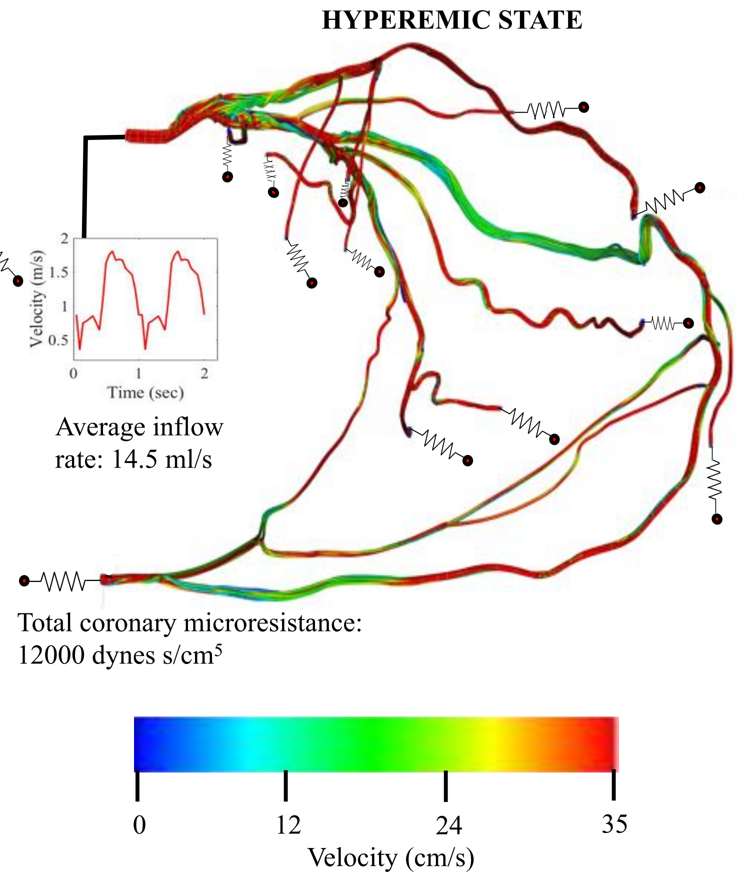

Vardhan *et al*. Nature Scientific Reports 2021

IMPORTANCE OF SIDE BRANCHES

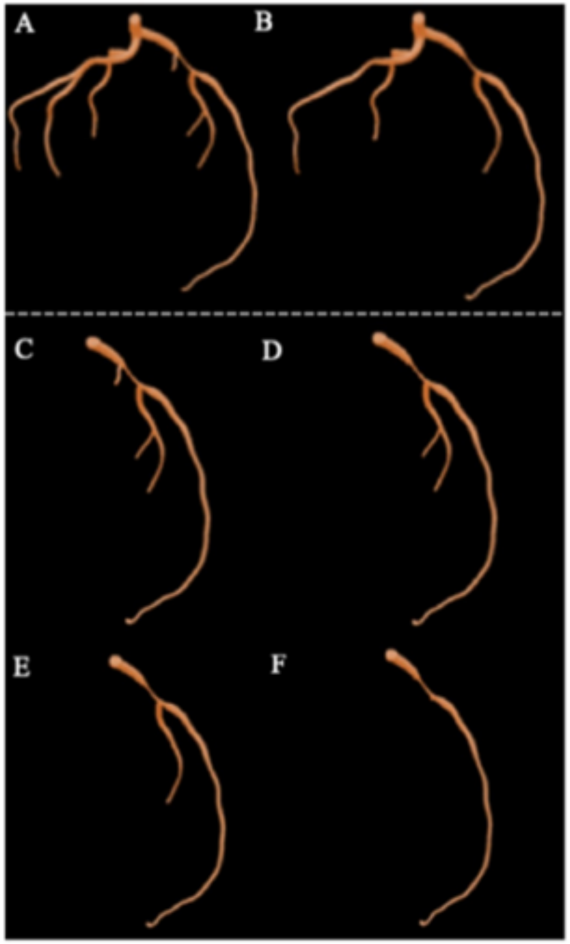

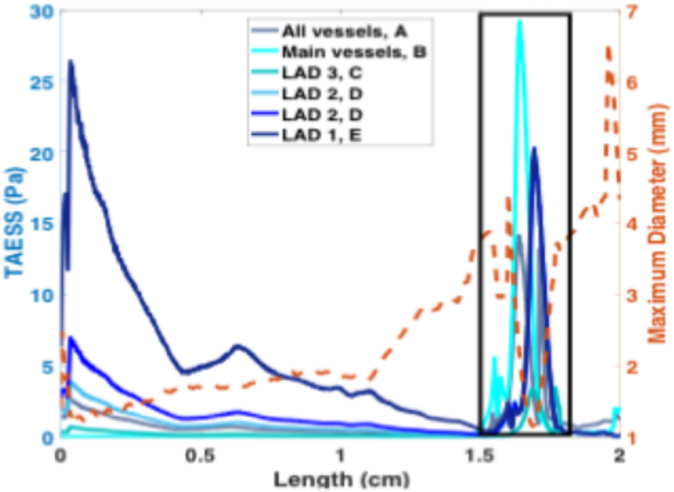

Capturing the full arterial tree including side branches is critical to accurately assessing flow characteristics.

Vardhan *et al*. Nature Scientific Reports 2019

OTHER USE CASES

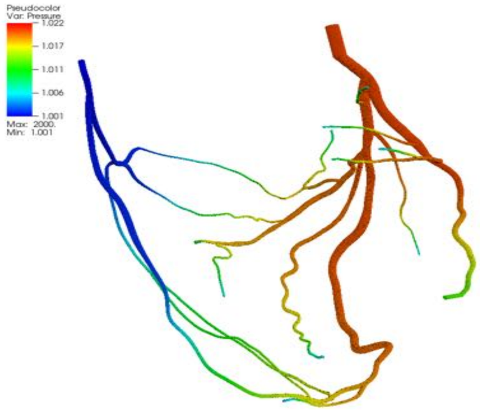

Diagnostics

Treatment Planning

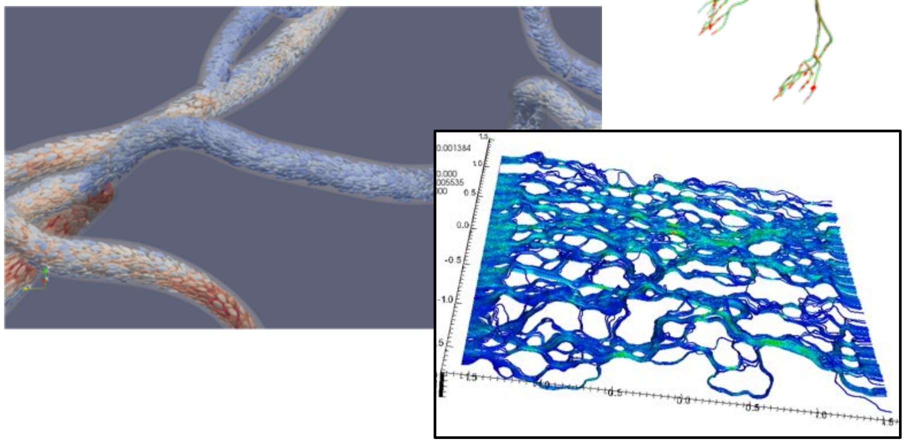

Mechanistic

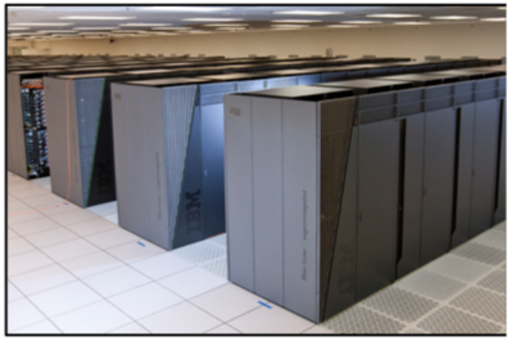

Massively Parallel Computing

So, why HPC?

2010: High-res coronary model

FULL ARTERIAL NETWORK

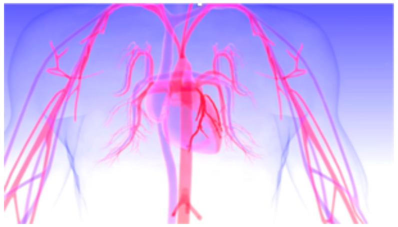

COMPUTER SCIENCE CHALLENGES

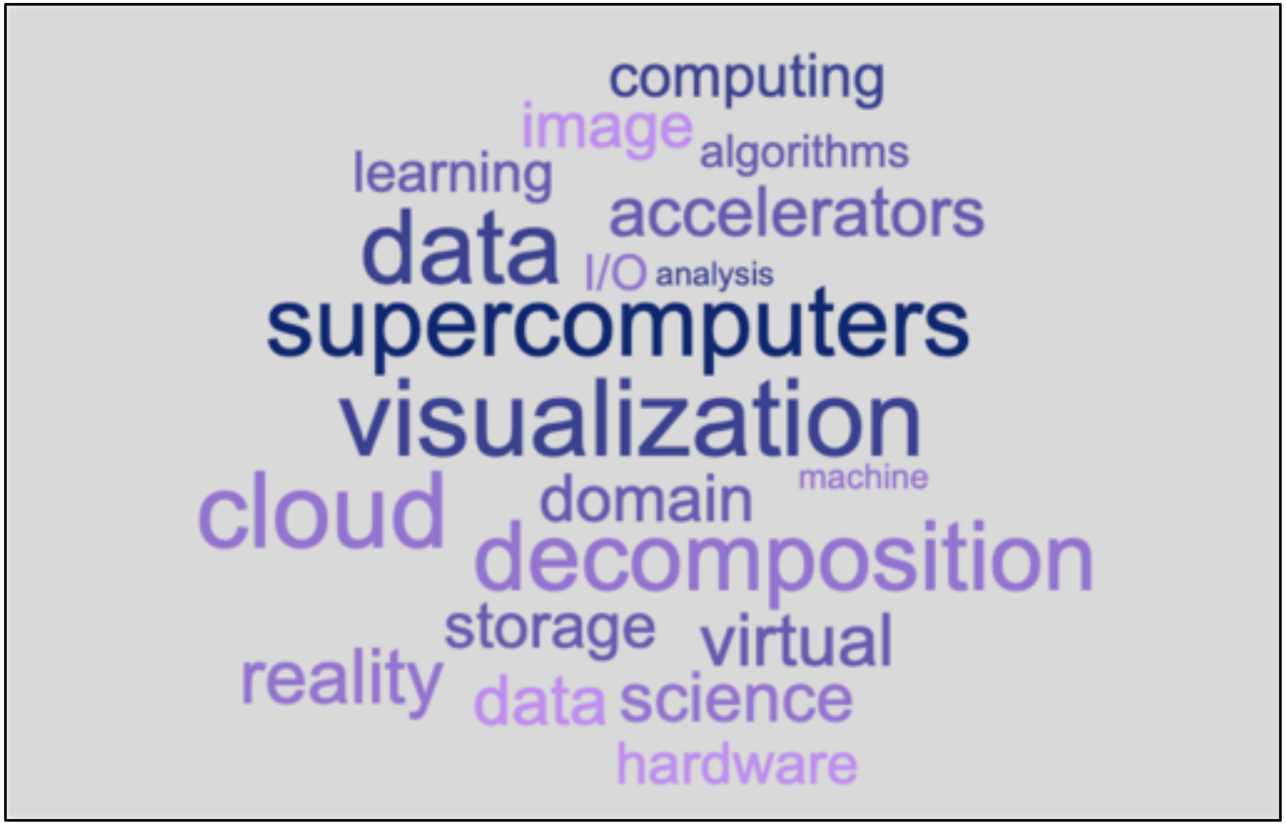

COMPUTER SCIENCE CHALLENGES

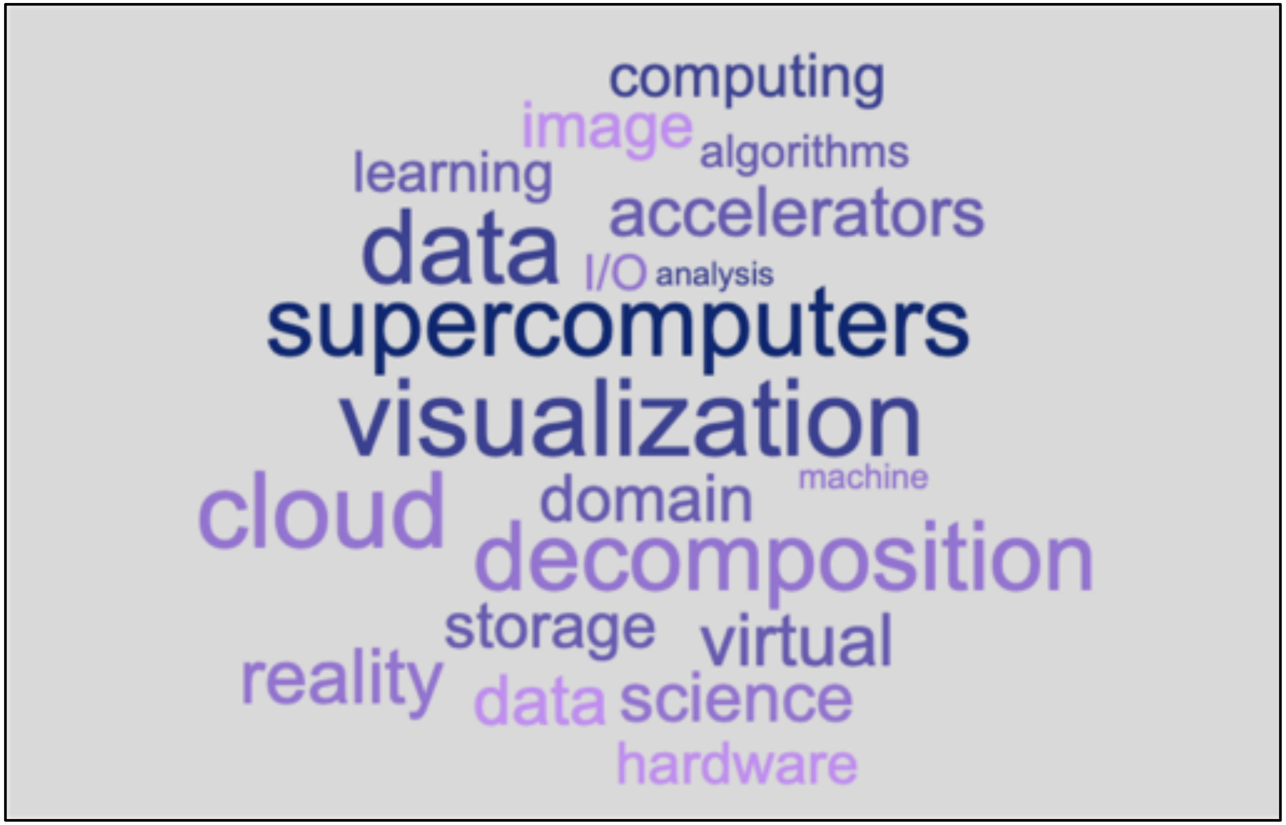

- Making the models tractable/scalable
- Visualizing the data
- Interacting with the simulation

HARVEY: numerics

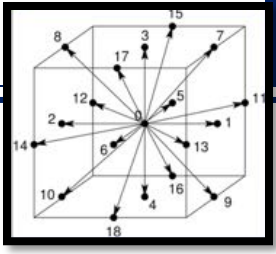

- Solves the (weakly compressible) Navier-Stokes equations
- Minimal communication between lattice points during update
- Macroscopic quantities computed at lattice points

$$\rho = \sum_i f_i \qquad \vec{u} = \frac{1}{\rho} \sum_i \vec{c}_i f_i \qquad P = c_s^2 \rho$$

- Straight forward treatment of complex geometry

DOMAIN DECOMPOSITON

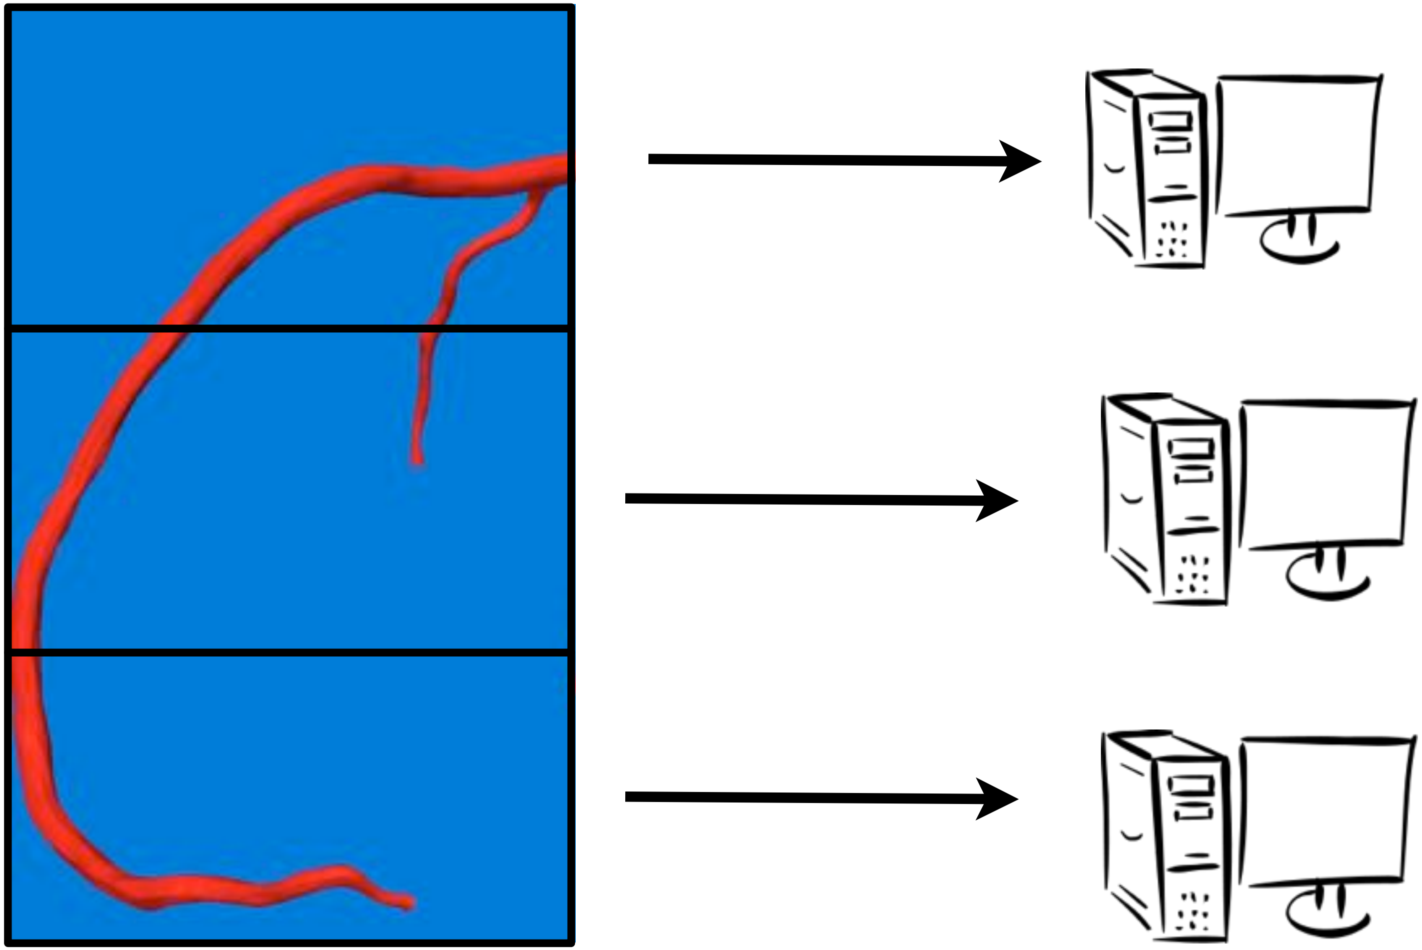

IRREGULAR DOMAIN DECOMPOSITION

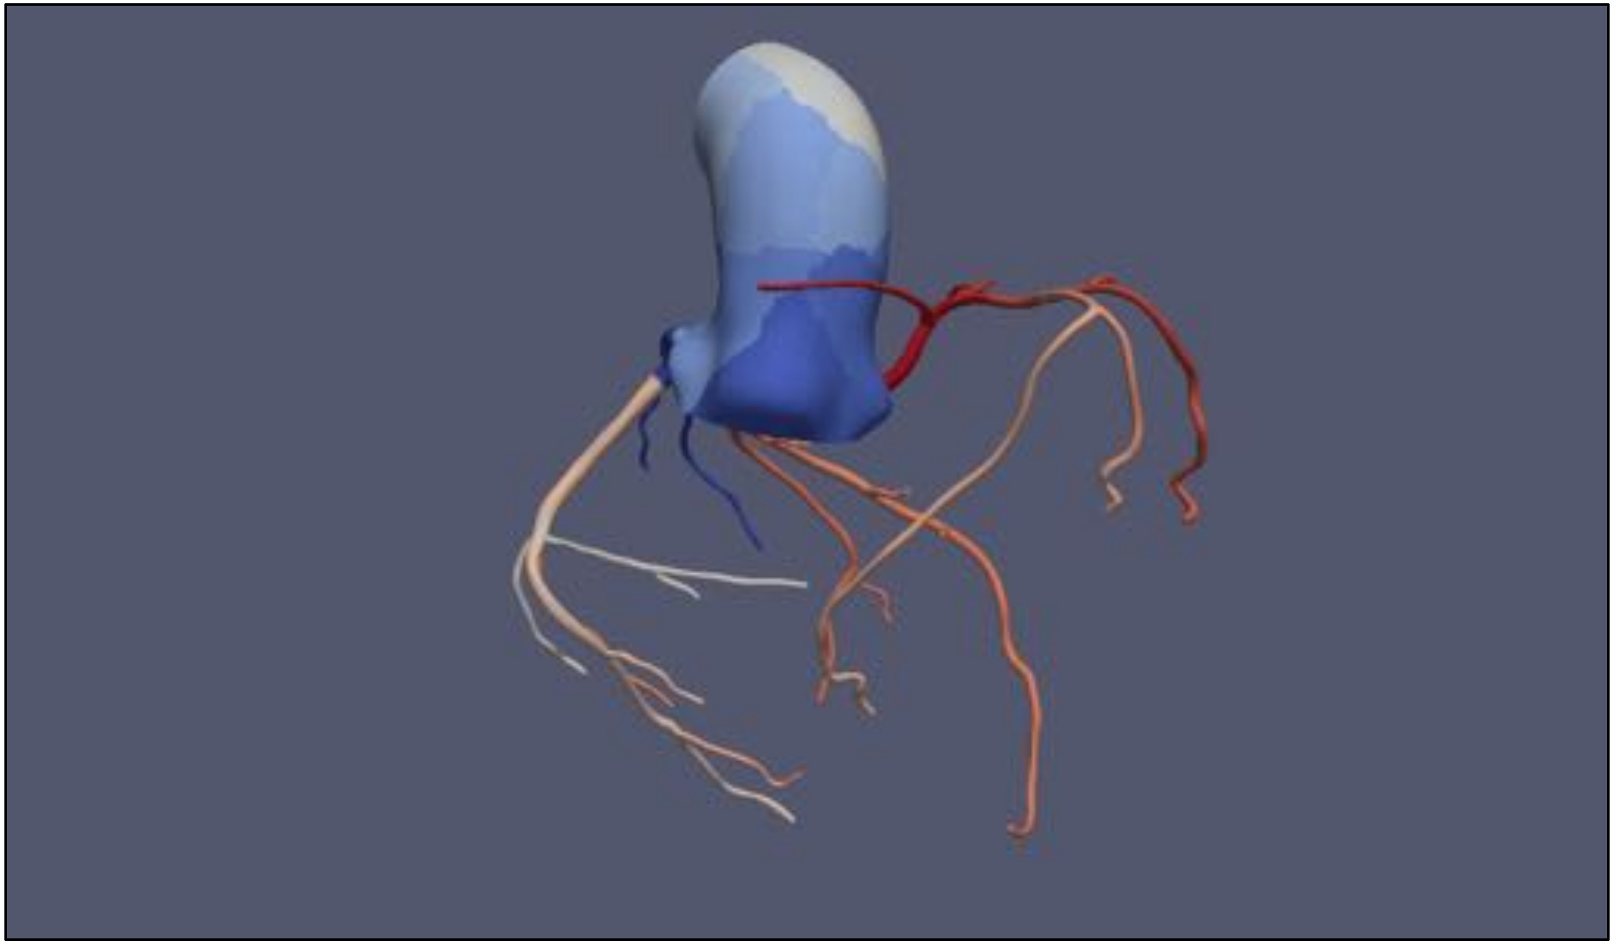

MINIMIZING MEMORY FOOTPRINT

Full body arterial	20 micron	9 micron
Data Grid	31009 x 11298 x 84863	68909 x 25107 x 188584
Data memory	8.22 PB	90.2 PB
Fluid nodes	45.8 billion	509 billion
Fluid memory	25.3 TB	140.7 TB
Fluid Fraction	0.15%	0.15%

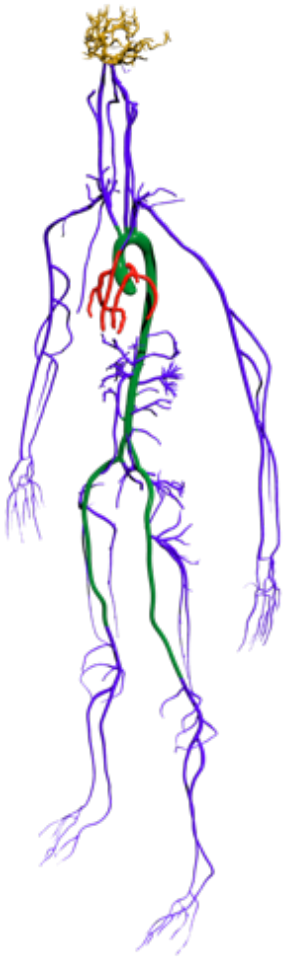

- Sequoia Blue Gene/Q total system memory: 1.6 PB
- Indirect addressing is mandatory, and only the first step
- Initialization and load balance now significantly more challenging

MEMORY LAYOUT SCHEMES

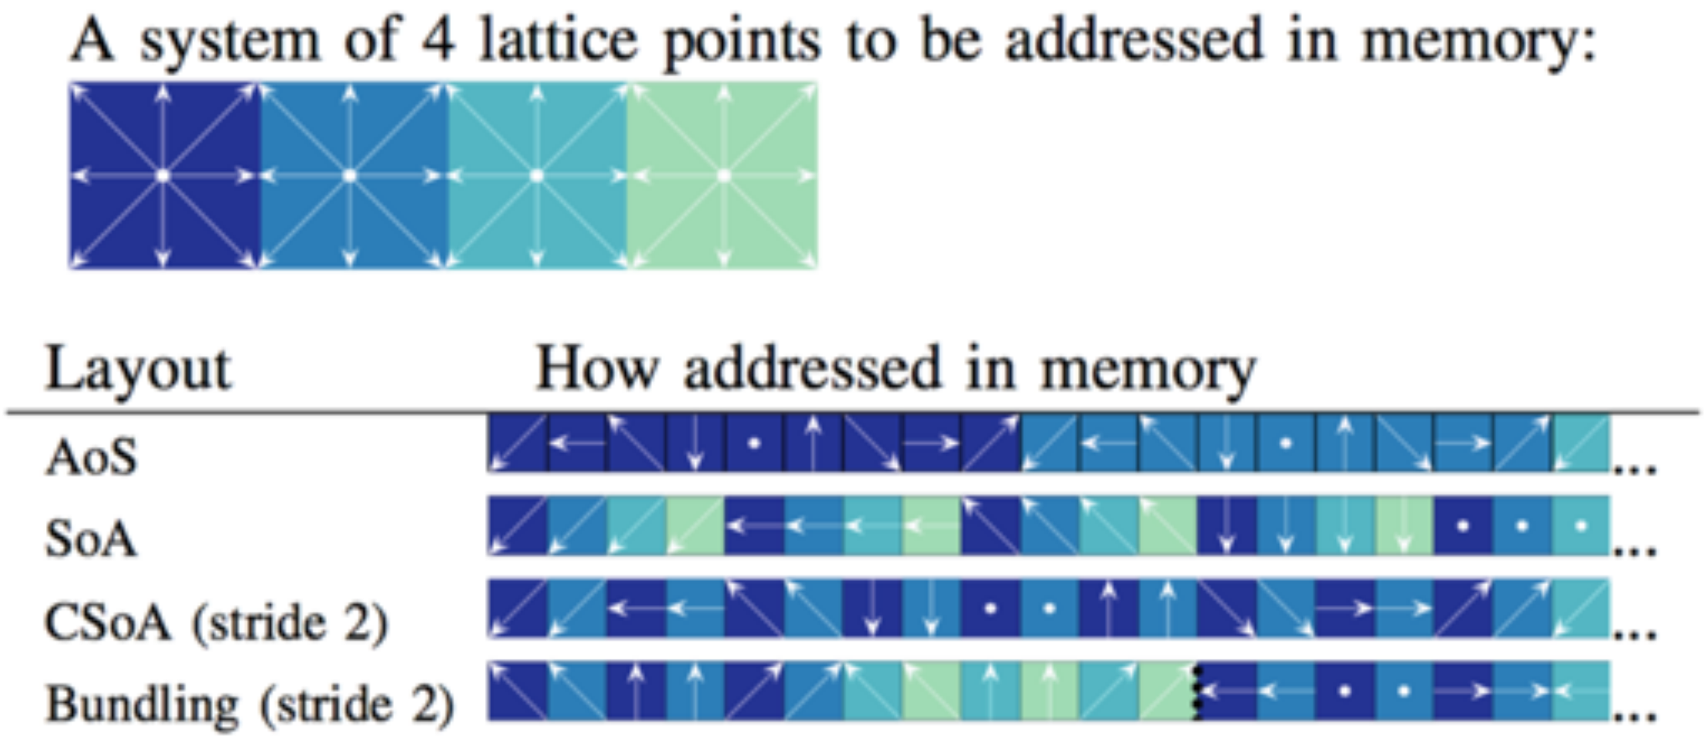

Herschlag *et al.* IPDPS 2018

Strong Scaling on Homogeneous Systems

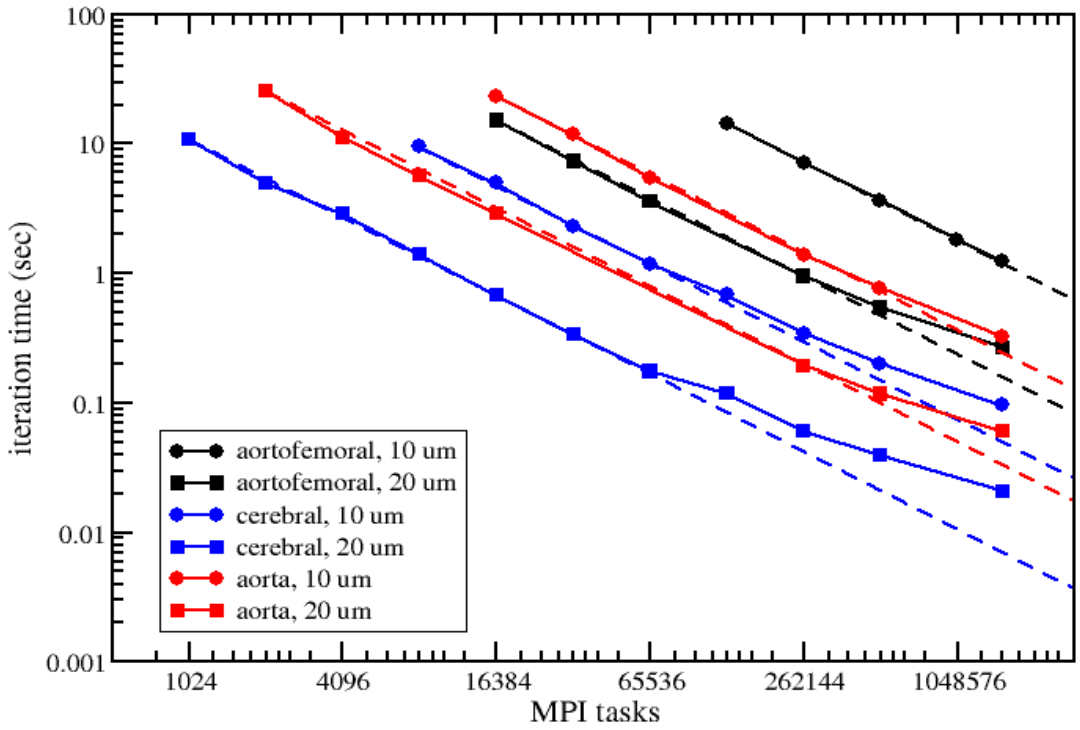

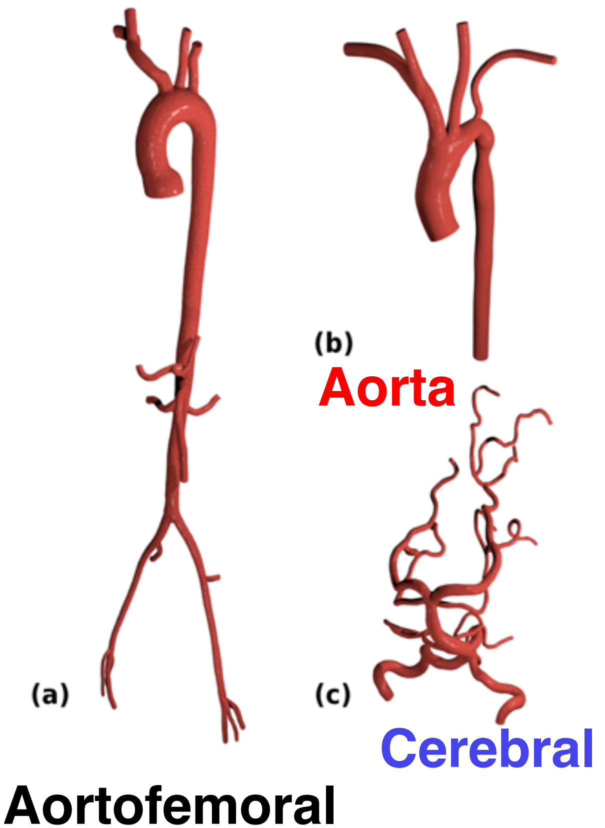

Randles *et al*. Journal of Computational Science, 2015

Identify Tipping Point For Model Use

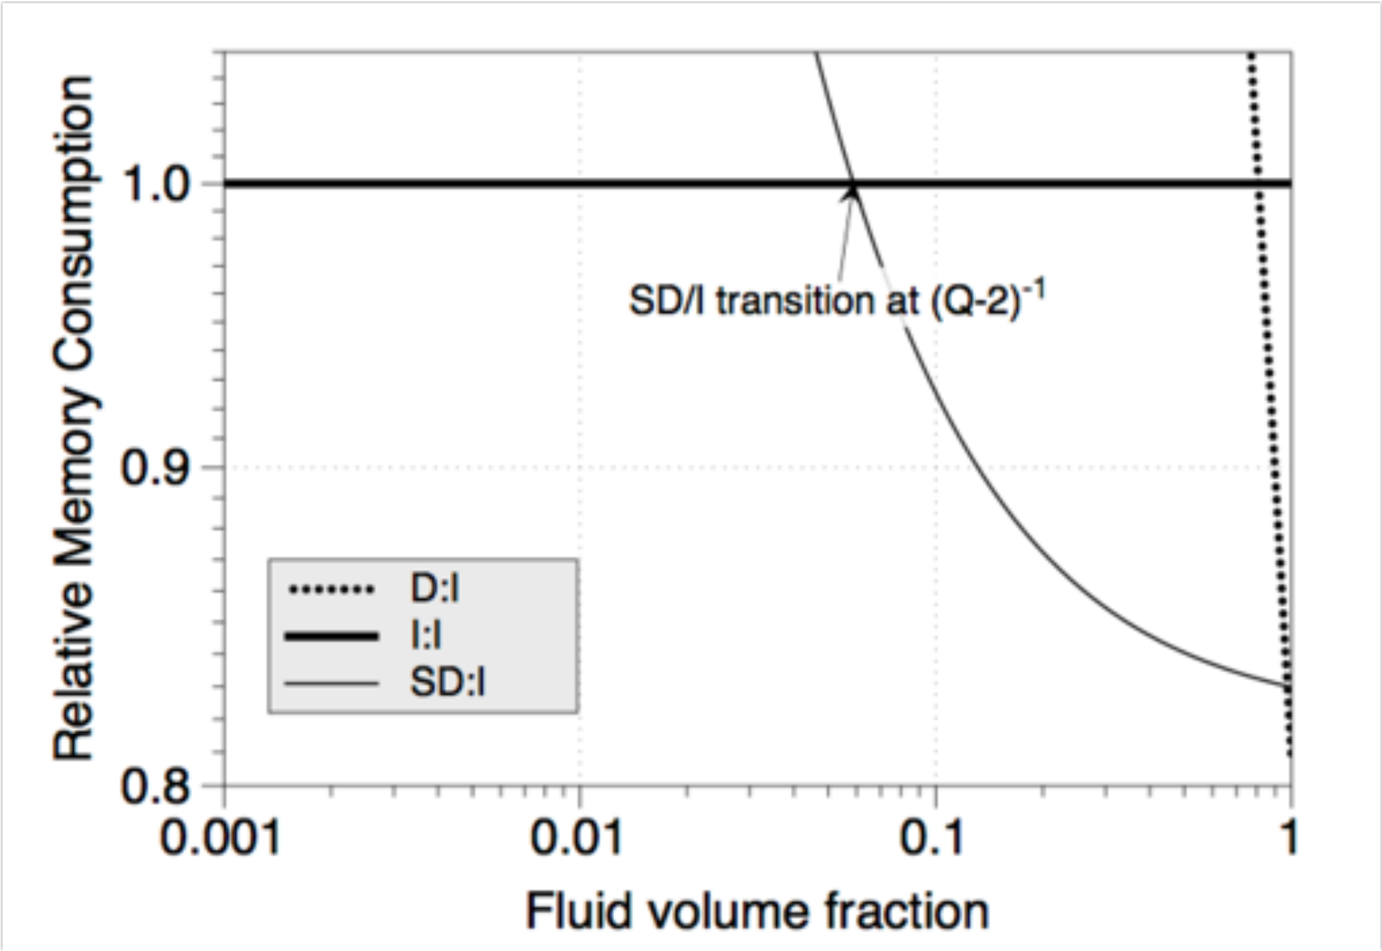

Above a set fluid volume fraction, semi-direct addressing is more efficient than the conventional indirect addressing on GPUs.

STRONG SCALING

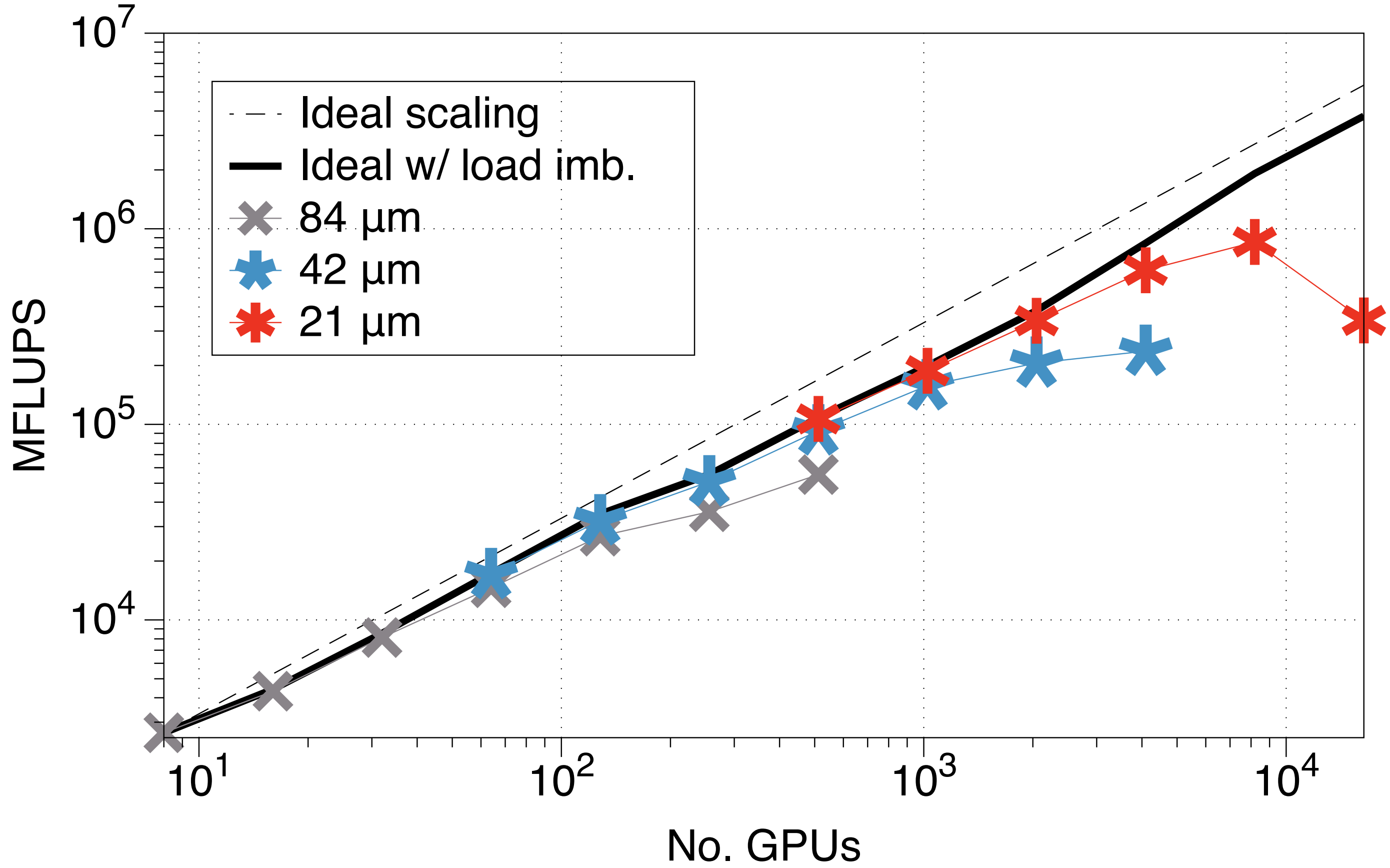

HARVEY scales efficiently on 16k nodes of the Titan supercomputer.

Herschlag *et al*. IEEE TPDS 2021

IN SITU AND IN TRANSIT VISUALIZATION

- Large-scale simulations create petabytes of data per timestep.
- The gap between the speed of computation and speed of I/O is increasing with next generation systems.
- Developing methods to enable efficient, in situ and in transit visualization and analysis.
- Enabling communication free and re-wind capabilities.

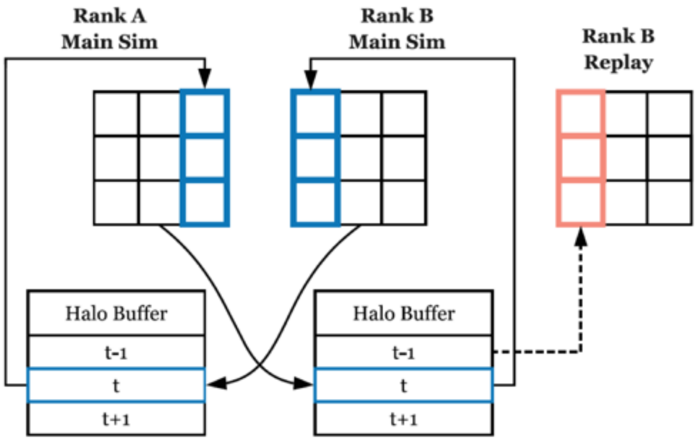

Ames et al. LDAV 2019

LOW OVERHEAD REPLAY

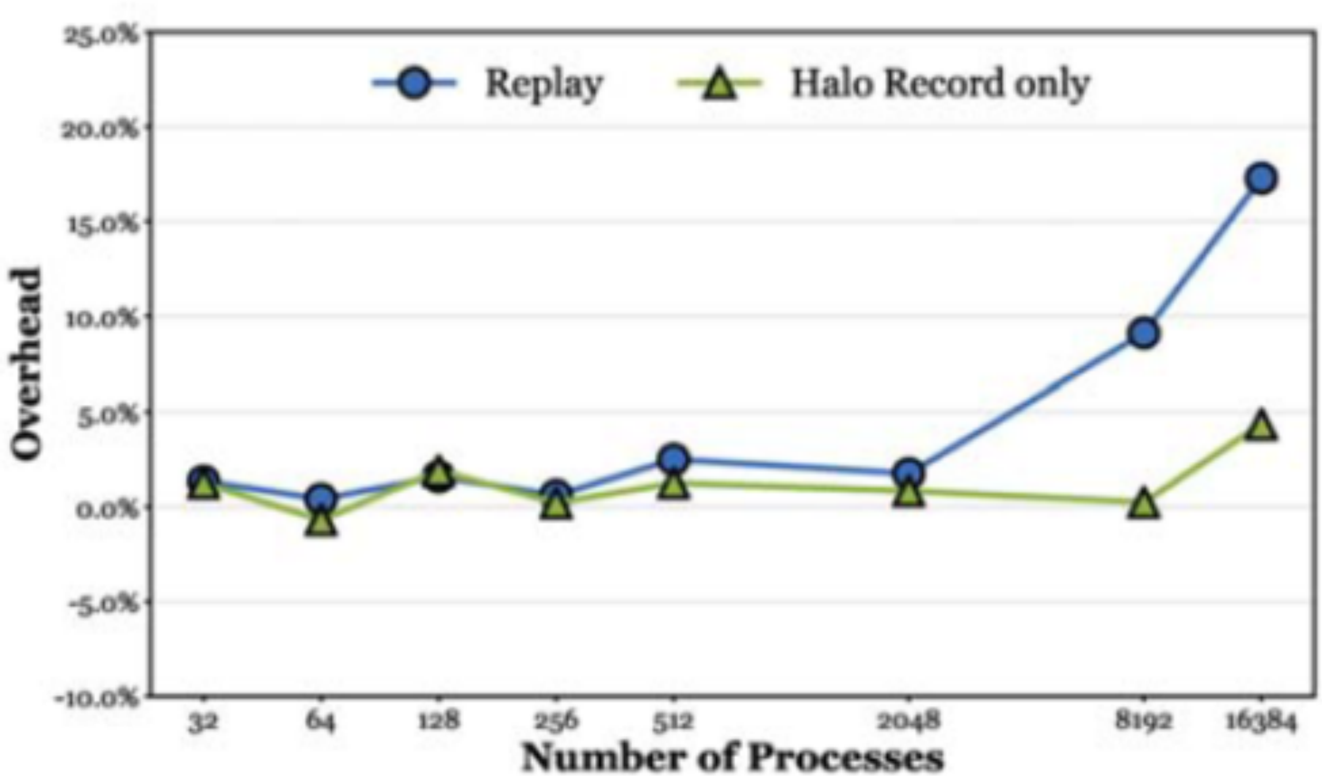

Ames et al. LDAV 2019

HarVis: interaction with HARVEY

- Using devices like the Occulus Rift, HTC Vive, and zSpace to investigate different modes of interaction

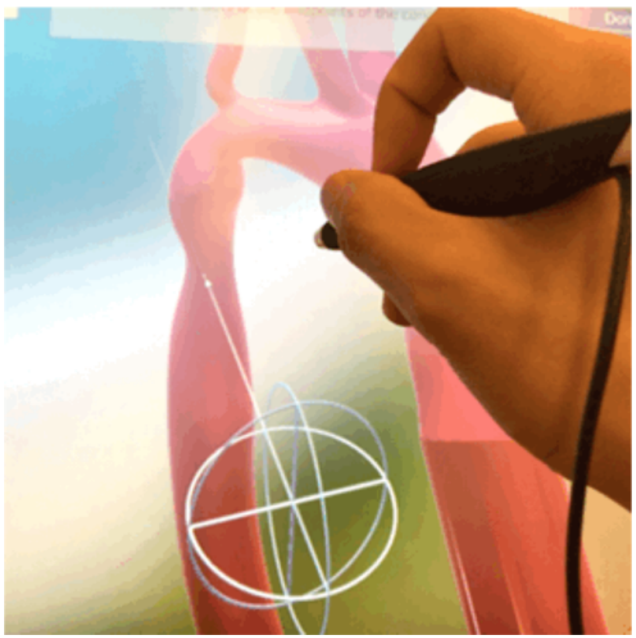

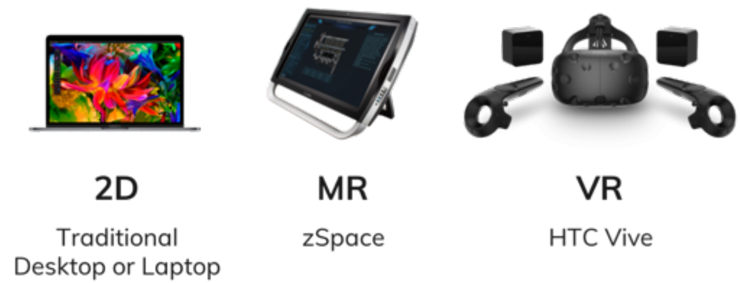

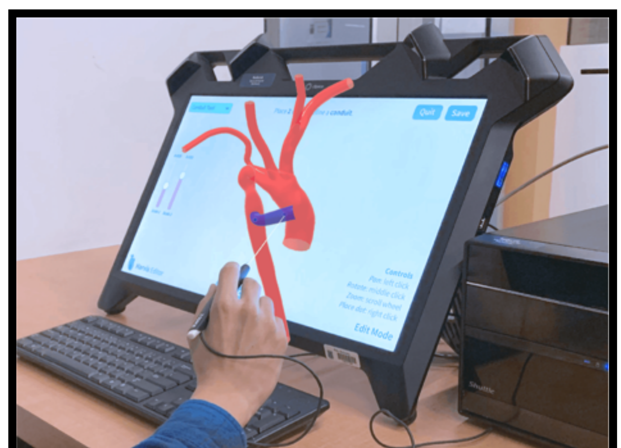

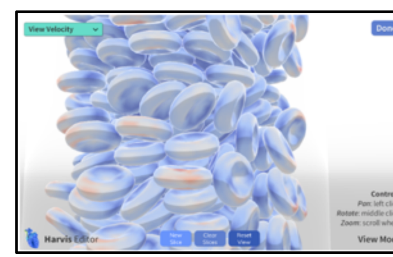

Shi *et al*. under review IEEE TVCG

ACKNOWLEDGEMENTS

Randles Lab:

- **Adebayo Adebiyi**
- **Greg Herschlag**
- Mike Kaplan
- **John Gounley**
- Mahsa Dabagh
- Luiz Hegele
- **Jeff Ames**
- Simba Chidyagwai
- **Brad Feiger**
- Ismael Perez
- **Daniel Puleri**
- Sayan Roychowdhury
- **Maddie Vardhan**
- **Harvey Shi**
- **Peter Balogh**
- Zach Bernstein
- Dian Niu
- Erick Lorenzana-Saldivar
- Arpita Das
- Cyrus Tanade

Collaborators:

- James Chen, PhD, UC Denver
- Jane Leopold, MD, BWH
- Andy Kahn, MD, PhD, UCSD
- Franziska Michor, PhD, Dana Farber Cancer Institute
- Erik Draeger, PhD, LLNL
- Seyong Lee, PhD, ORNL
- Jeff Vetter, PhD, ORNL
- Joe Insley, PhD, ANL
- Silvio Rizzo, PhD, ANL
- David Frakes, PhD, ASU
- Roger Kamm, PhD, MIT
- Manesh Patel, MD, Duke
- Piers Barker, MD, Duke
- Leila Mureebe, MD, Duke
- Sreek Vemulapalli, MD, Duke
- Lynn Koweek, MD, Duke
- Schuyler Jones, MD, Duke
- Kevin Hill, MD, Duke
- Chad Hughes, MD, Duke

NIH DIRECTOR'S EARLY INDEPENDENCE AWARD

Research supported by the Office of the Director, National Institutes of Health under Award Number DP5OD019876, the Duke Quantitative Initiative, the Hartwell Foundation, Boston Scientific, and the LLNL Laboratory Directed Research and Development (LDRD) program. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Computing support for this work came from the Lawrence Livermore National Laboratory (LLNL) Institutional Computing Grand Challenges.