Characterizing In Situ and In Transit Analytics of Molecular Dynamics Simulations for Next-generation Supercomputers

*Stephen Thomas, Michael Wyatt, **Ian Lumsden**, Tu Mai Anh Do, Loïc Pottier, Rafael Ferreira da Silva, Harel Weinstein, Michel A. Cuendet, Trilce Estrada, Ewa Deelman, Michela Taufer*

THE UNIVERSITY OF TENNESSEE KNOXVILLE

BIG ORANGE. BIG IDEAS.®

Acknowledgements

T. Estrada

H. Weinstein

M. Cuendet

E. Deelman

M. Wyatt

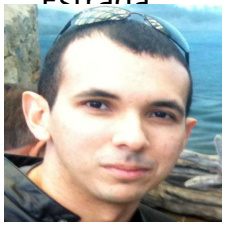

R. da Silva

T. Johnston

T. Do

B. Mulligan

I. Lumsden

S. Thomas

H. Carrillo

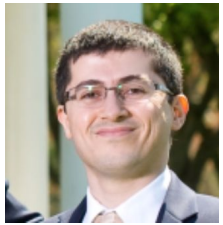

A. Razavi

Sponsors:

GCL | GCLab
Dr. Michela Taufer

UNM

USC

NSF

IBM

Trends in Next-Generation Systems: IO Gap and Ensembles

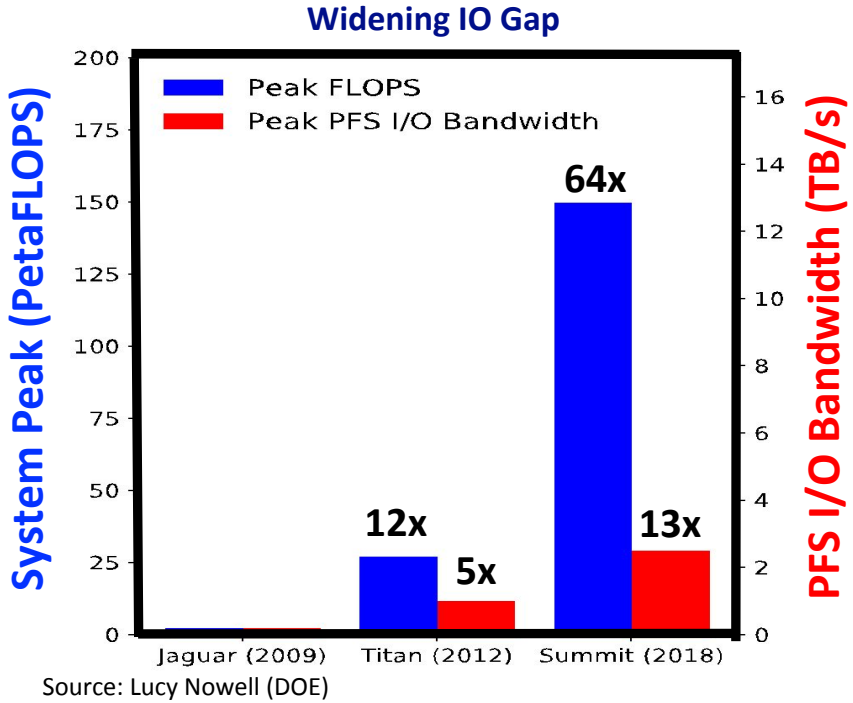

Source: Lucy Nowell (DOE)

Classical Molecular Dynamics Simulations

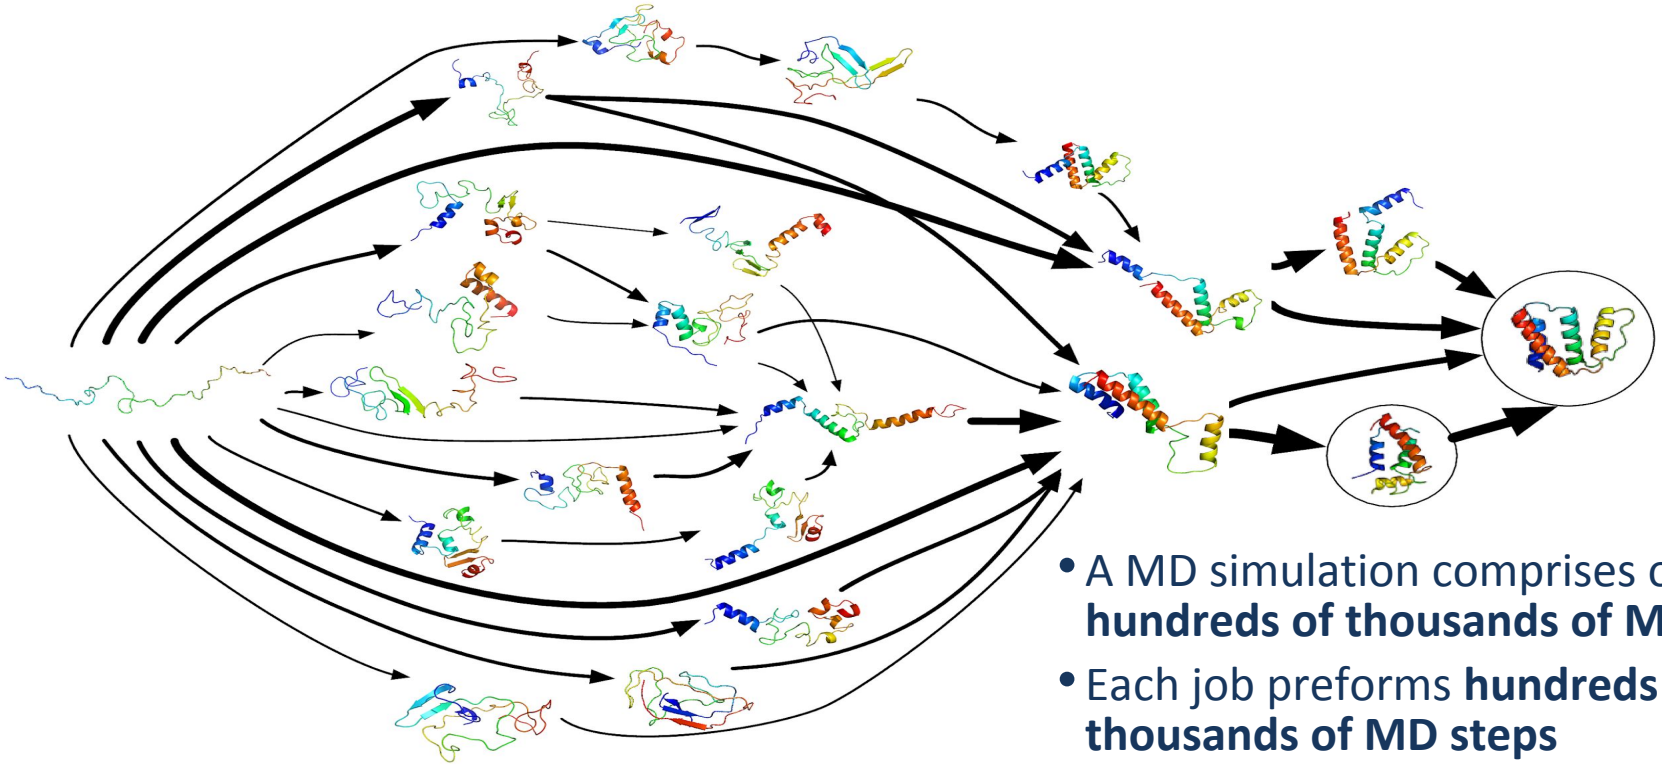

- A MD simulation comprises of **hundreds of thousands of MD job**
- Each job preforms **hundreds of thousands of MD steps**

Classical Molecular Dynamics Simulations

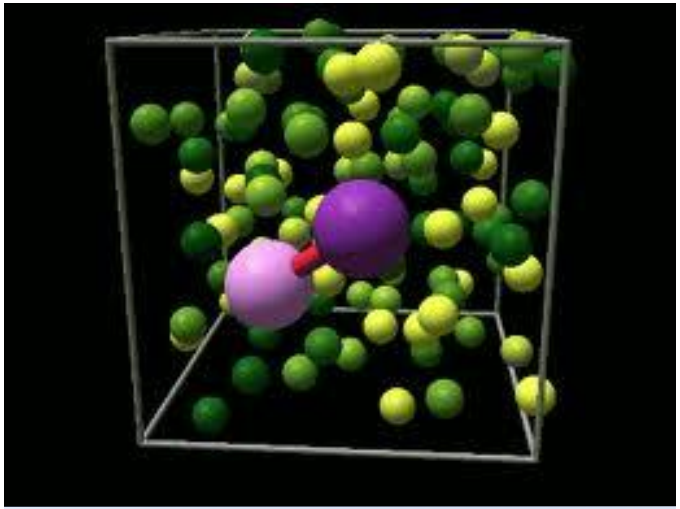

- Forces on single atoms
  - Acceleration
    - Velocity
      - Position

- MD step computes **forces** on single atoms (e.g., bond, dihedrals, nonbond)
- Forces are added to compute **acceleration**
- Acceleration is used to update **velocities**
- Velocities are used to update the **atom positions**
- Every *n* steps (Stride)
  - ***Store 3D snapshot or frame***

Extending HPC to Integrate Data Analytics

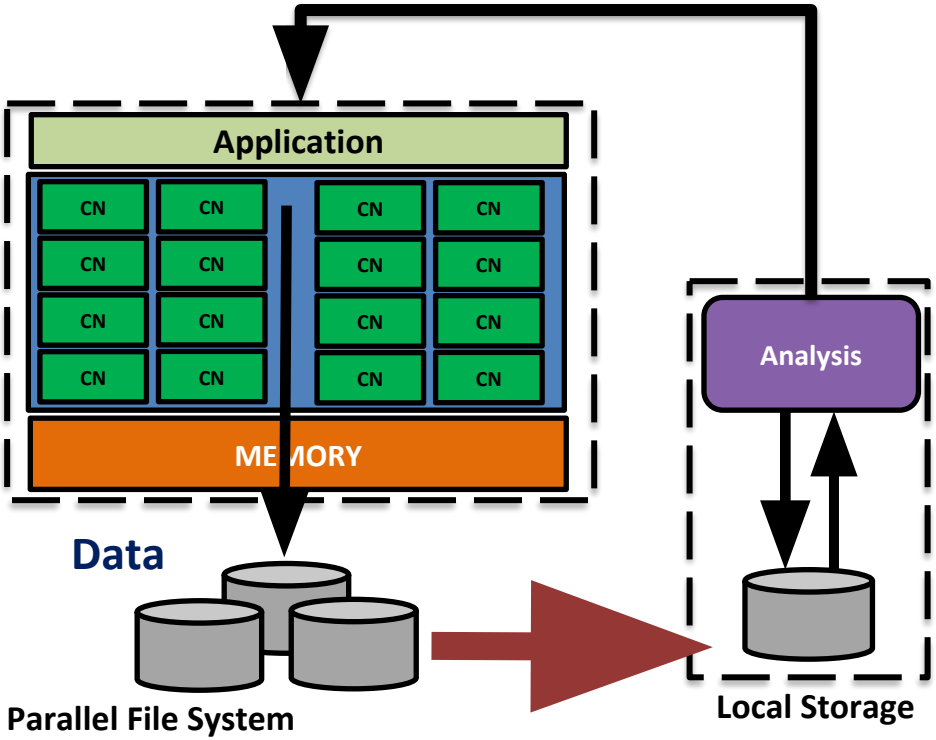

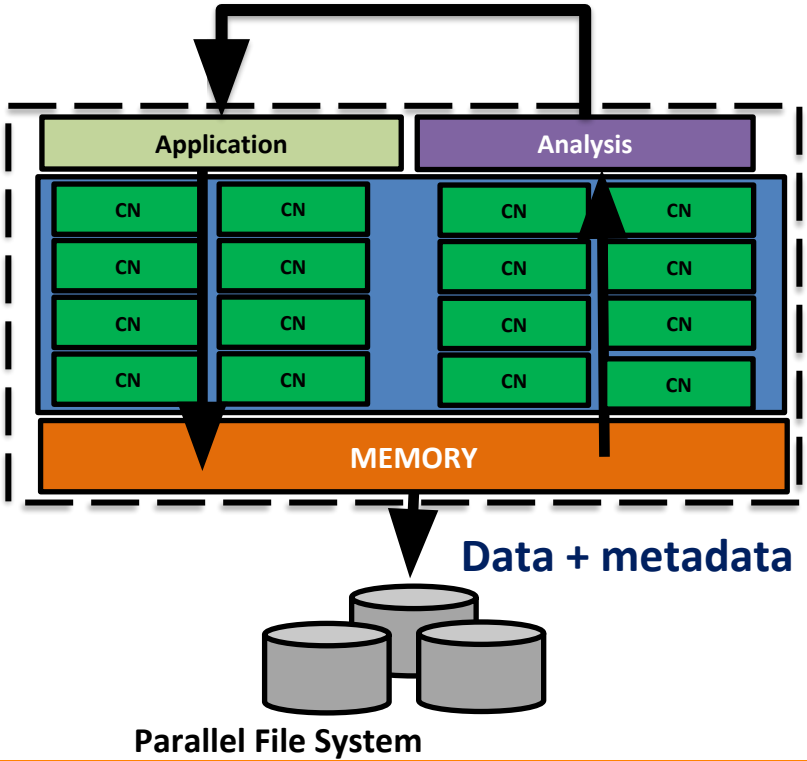

Augmenting HPC with In Situ and *In Transit* Analytics

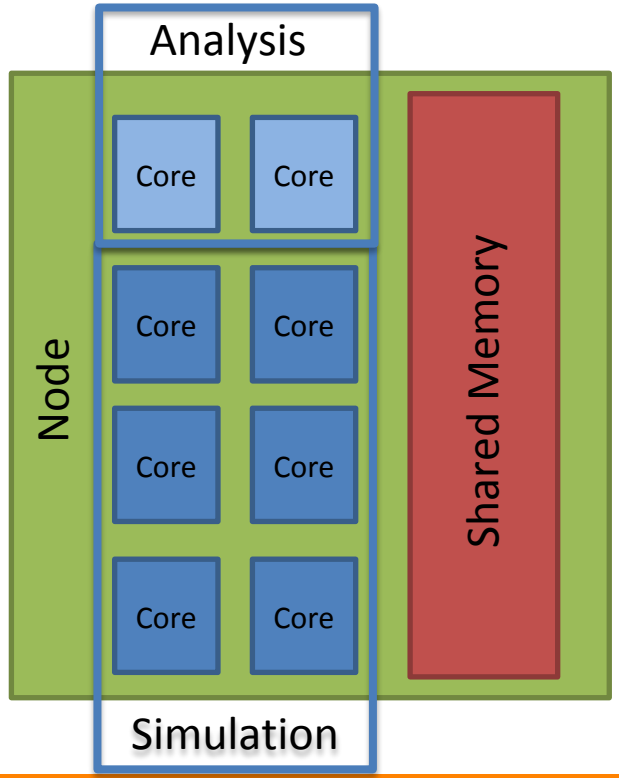

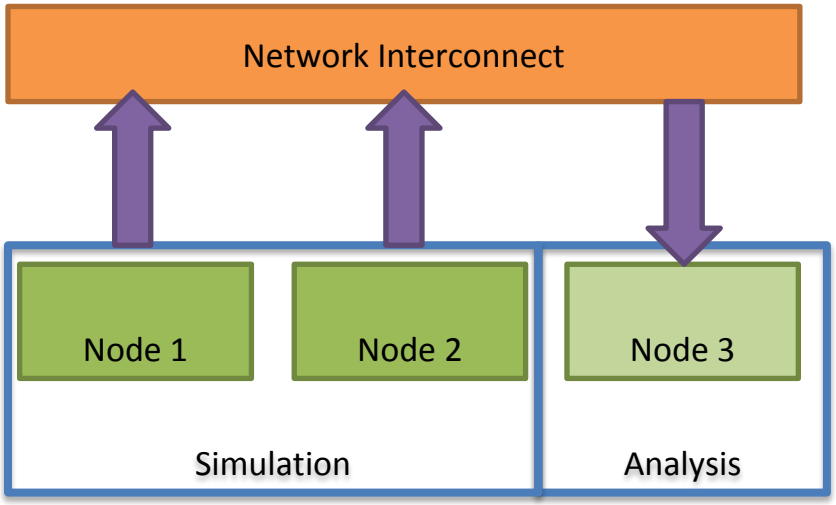

Example of tools:

- DataSpaces (Rutgers U.)
- DataStager (GeorgiaTech)

Building a Closed-loop Workflow

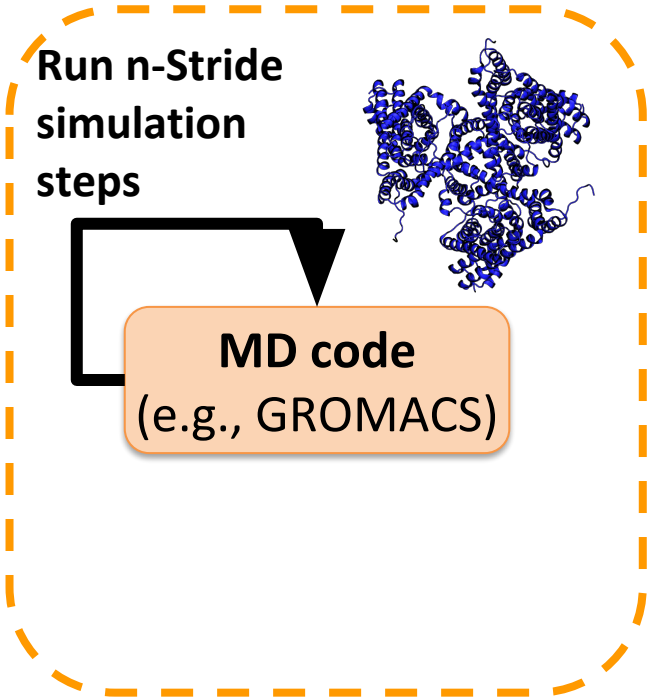

Building a Closed-loop Workflow

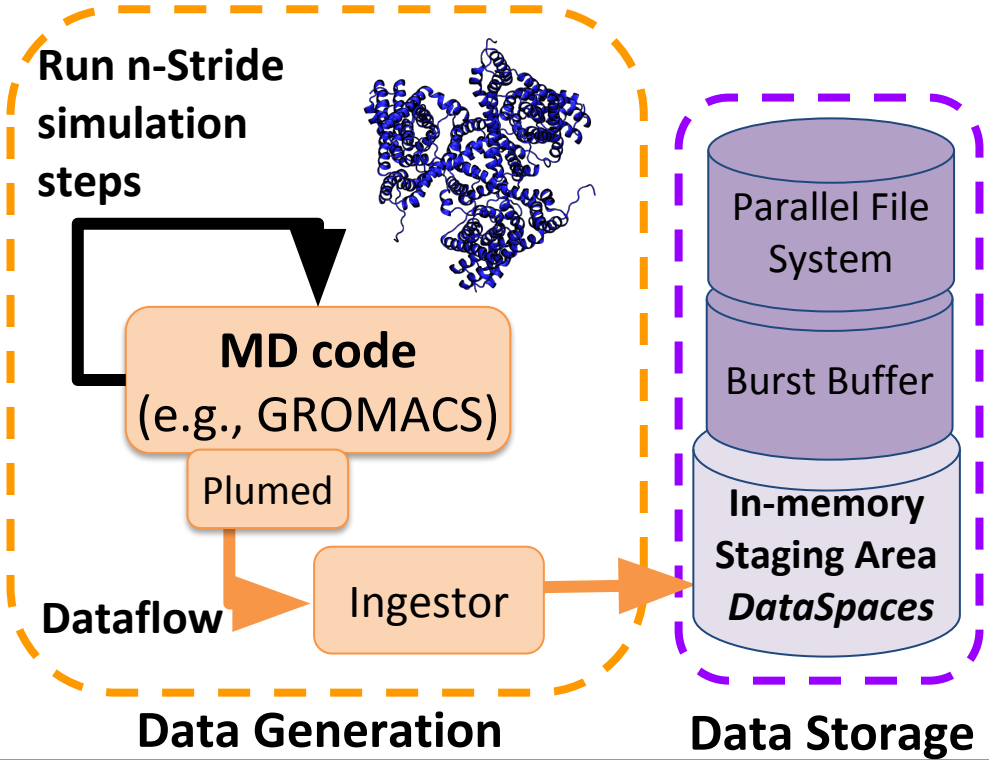

Building a Closed-loop Workflow

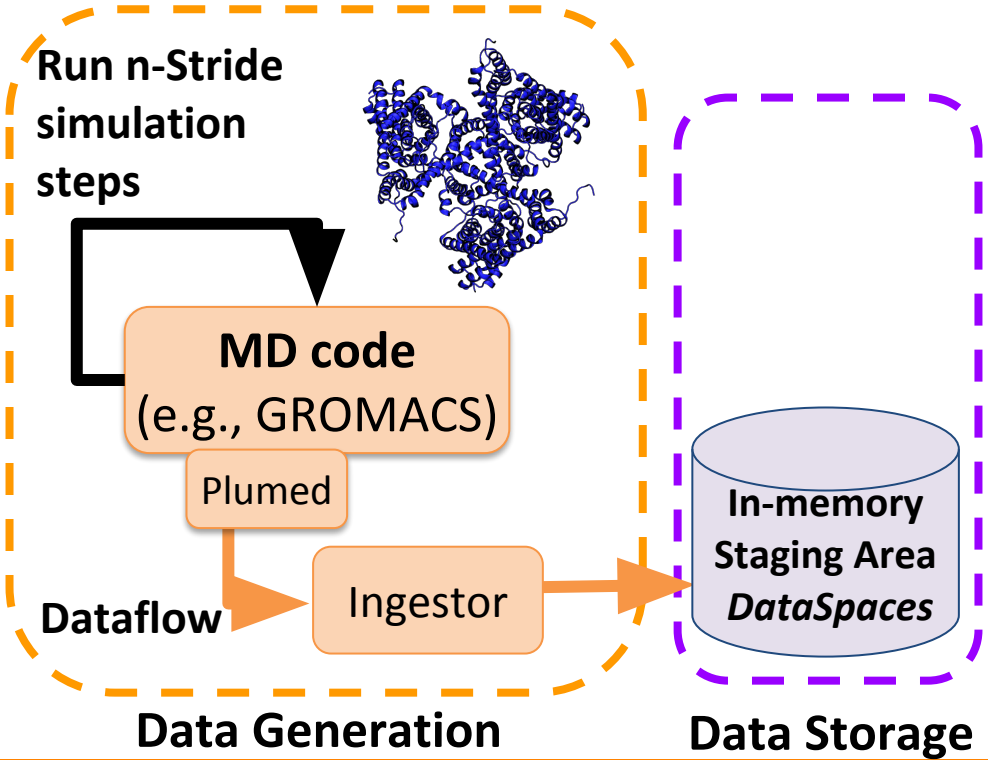

Building a Closed-loop Workflow

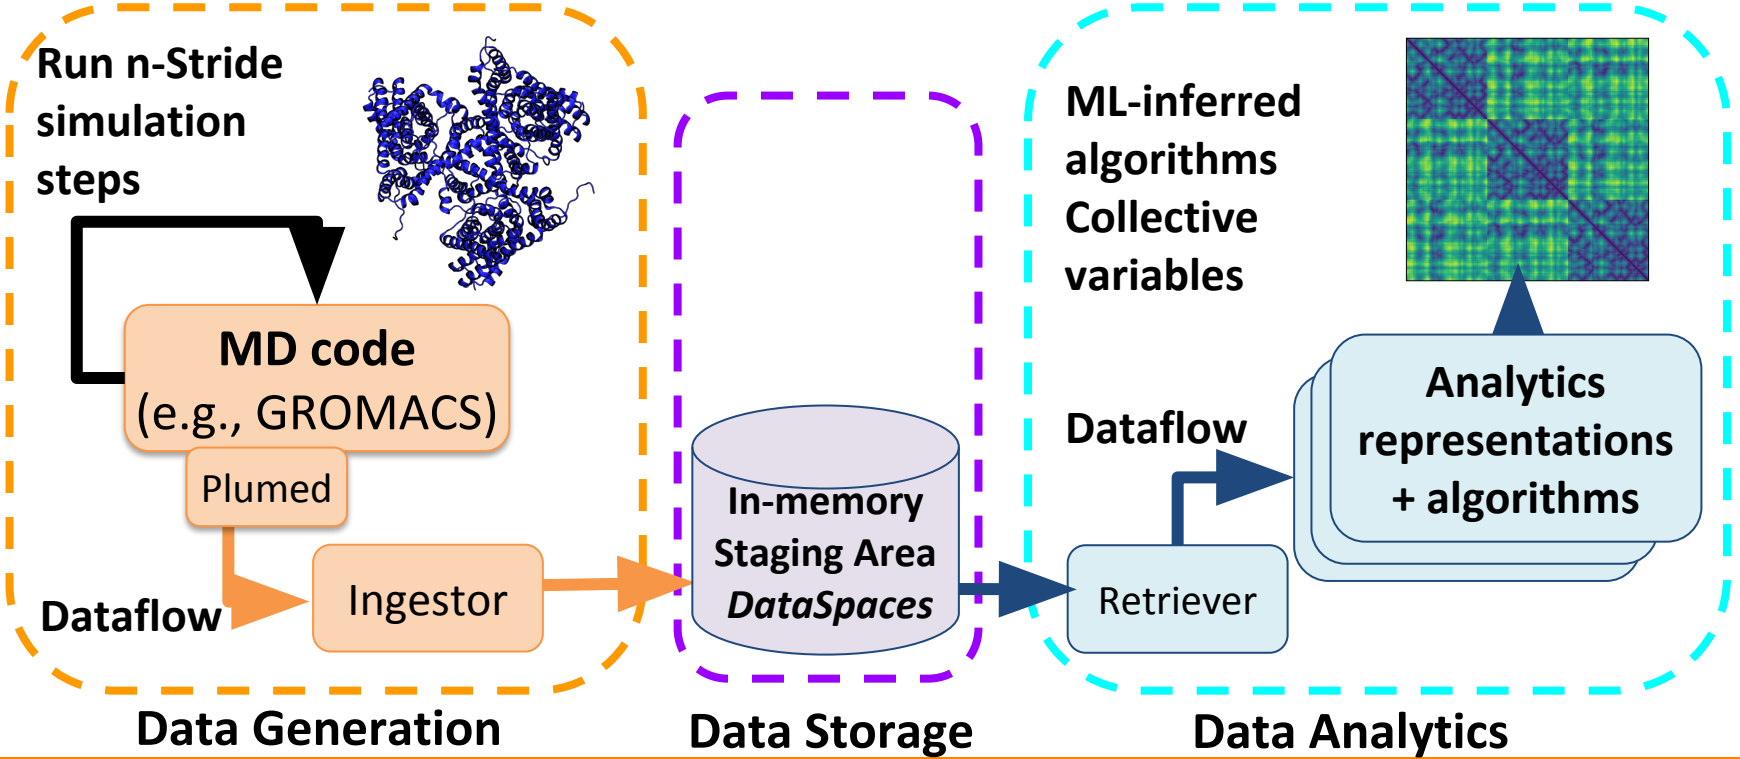

Building a Closed-loop Workflow

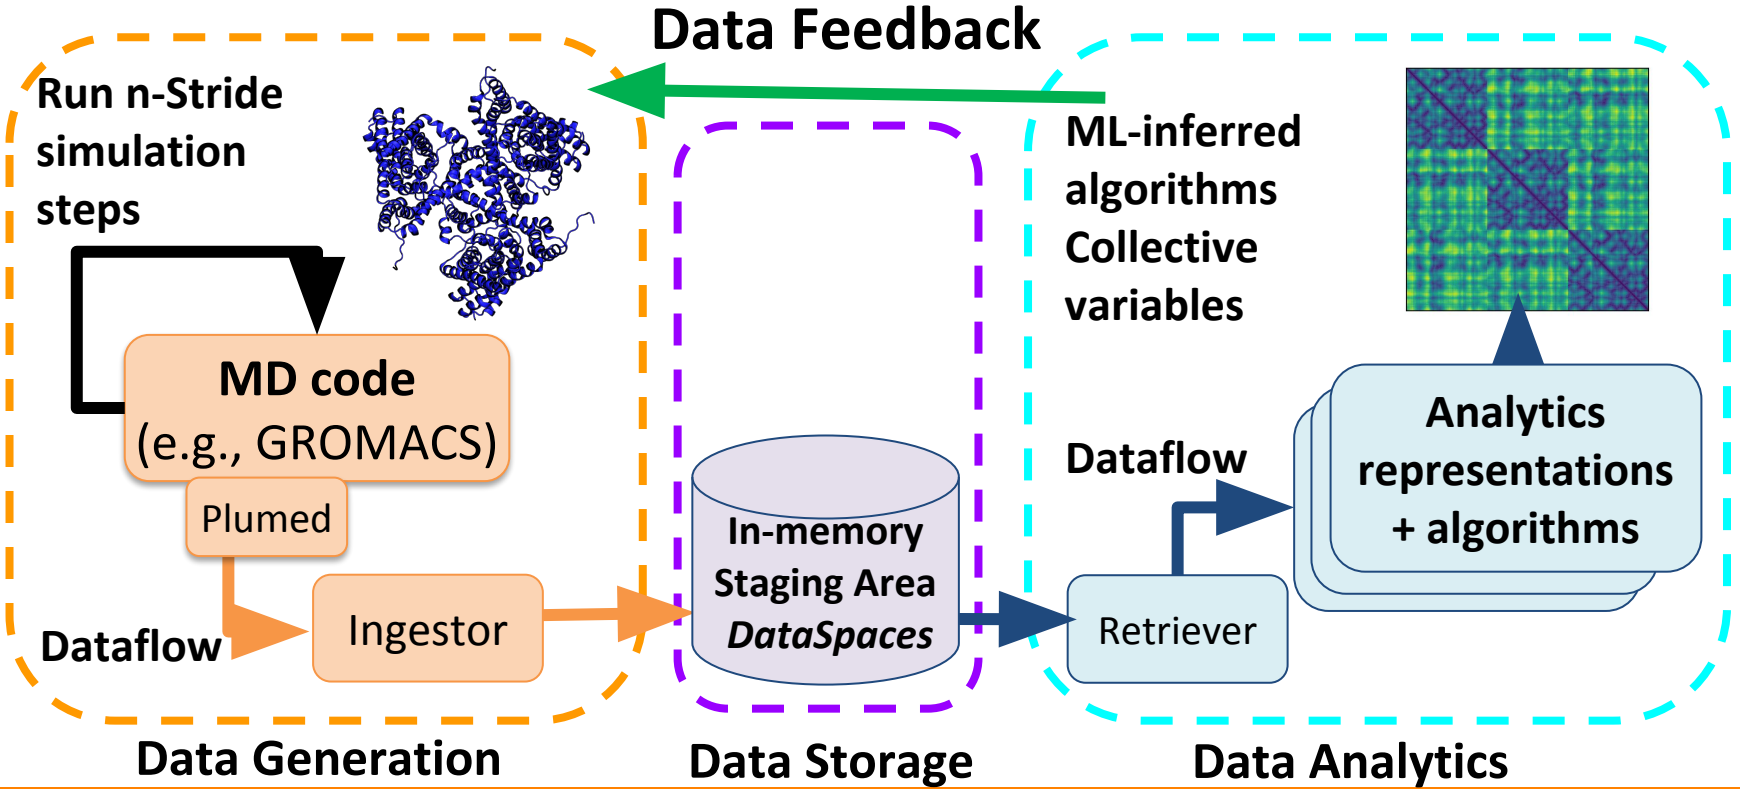

In Situ and *In Transit* Analytics for MD Simulations

Frames (or snapshots) of an MD trajectory:

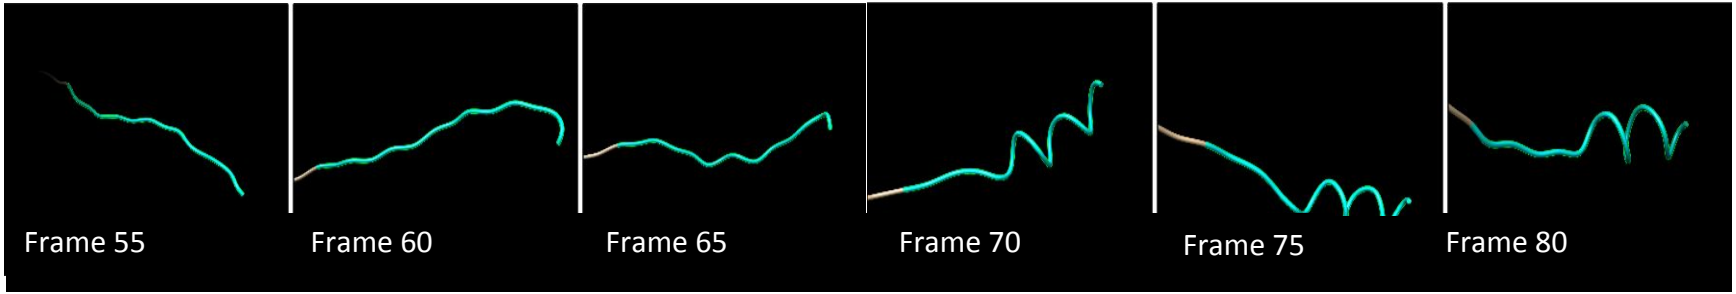

- We want to capture what is going on in each frame **without**:
  - Disrupting the simulation (e.g., stealing CPU and memory on the node)
  - Moving all the frames to a central file system and analyzing them once the simulation is over
  - Comparing each frame with past frames of the same job
  - Comparing each frame with frames of other jobs

In Situ and *In Transit* Analytics for MD Simulations

Frames (or snapshots) of an MD trajectory with a stride of 5 steps:

Collective variables (55)

In Situ and *In Transit* Analytics for MD Simulations

Frames (or snapshots) of an MD trajectory with a stride of 5 steps:

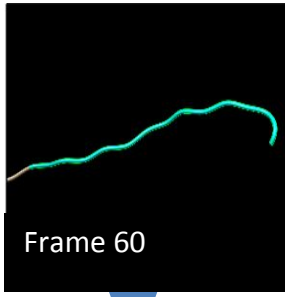

Collective variables (60)

In Situ and *In Transit* Analytics for MD Simulations

Frames (or snapshots) of an MD trajectory with a stride of 5 steps:

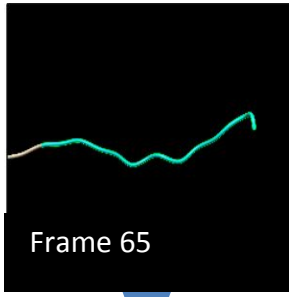

Collective variables (65)

In Situ and *In Transit* Analytics for MD Simulations

Frames (or snapshots) of an MD trajectory with a stride of 5 steps:

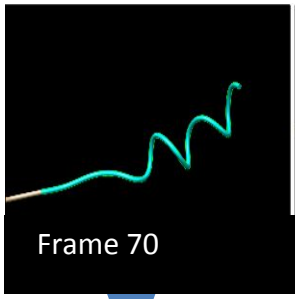

Collective variables (70)

In Situ and *In Transit* Analytics for MD Simulations

Frames (or snapshots) of an MD trajectory with a stride of 5 steps:

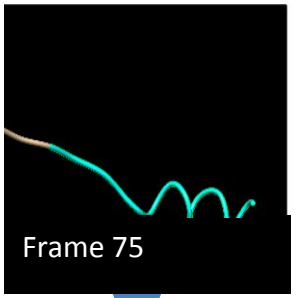

Collective Variables (75)

In Situ and *In Transit* Analytics for MD Simulations

Frames (or snapshots) of an MD trajectory with a stride of 5 steps:

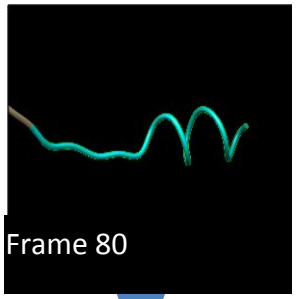

Collective variables (80)

In Situ and *In Transit* Analytics for MD Simulations

Frames (or snapshots) of an MD trajectory with a stride of 5 steps:

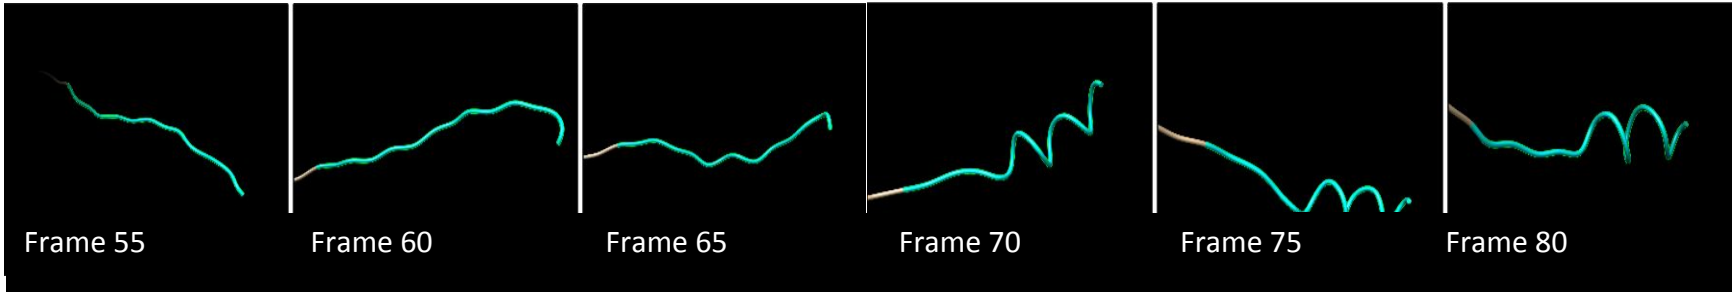

In Situ and *In Transit* Analytics for MD Simulations

Frames (or snapshots) of an MD trajectory with a stride of 5 steps:

Collective variables(55)	Collective variables(60)	Collective variables(65)	Collective variables(70)	Collective variables(75)	Collective variables(80)

Collective variables can serve as proxy for structural and conformational changes

Travis Johnston, Buyu Zhang, Adam Liwo, Silvia Crivelli, and Michela Taufer. "***In-Situ Data Analysis and Indexing of Protein Trajectories,***" JCC 2017.

Dataflow Modeling for Analytics

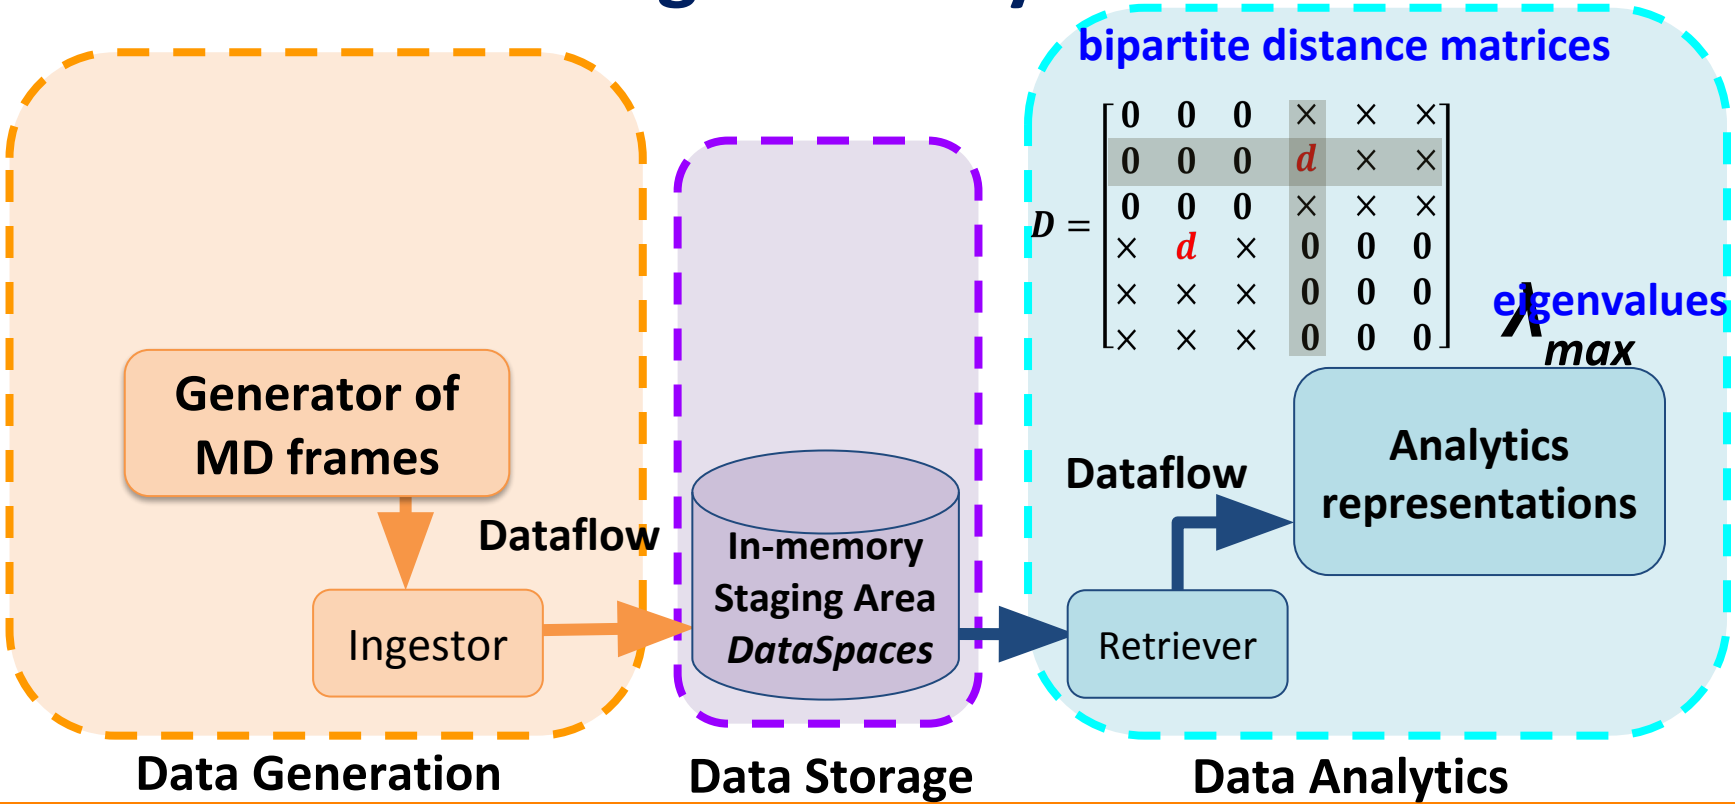

Generator of MD Frames

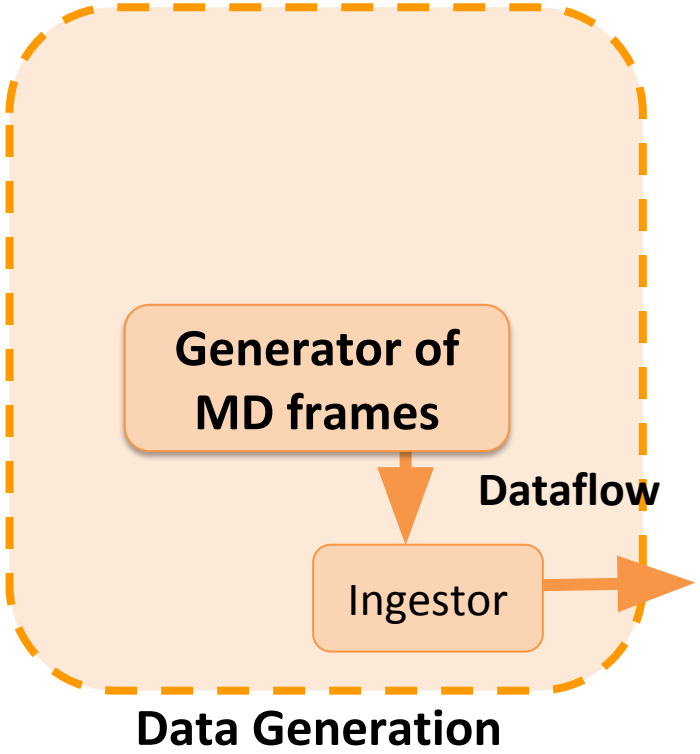

Generator of MD Frames

Molecular System	N	GPU	#GPU/Node	Nodes	ns/day	Package	Source	Cluster
Apoa1	92K	V100	2	1	78	NAMD 2.13	NAMD benchmarks	NVIDIA PSG cluster
Gltph	268K	P100	2	1	55	GROMACS	This work	Rockefeller Uni. cluster
STVM	1M	V100	2	1	7.5	NAMD 2.13	NAMD benchmarks	NVIDIA PSG cluster
STMV matrix systems (5x2x2)	21M	V100	6	1024	128	NAMD 2.13	NAMD benchmarks	SUMMIT
STMV matrix systems (7x6x5)	224M	V100	6	1024	24	NAMD 2.13	NAMD benchmarks	SUMMIT
STMV matrix systems (10x10x10)	1.07B	K20X	1	8192	4	NAMD 2.12	NAMD benchmarks	TITAN GPU

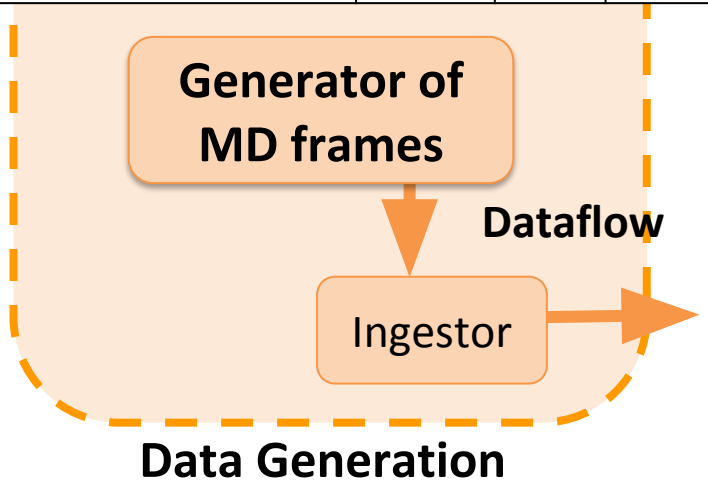

Trp cage 12,619 atoms

T cell receptor 81,092 atoms

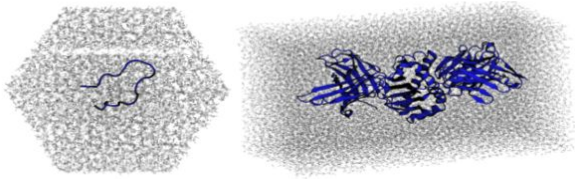

Gltph 270,088 atoms

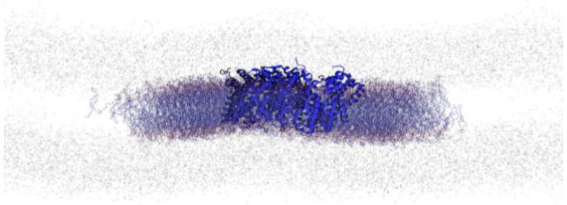

Analysis: Proxy for Performance

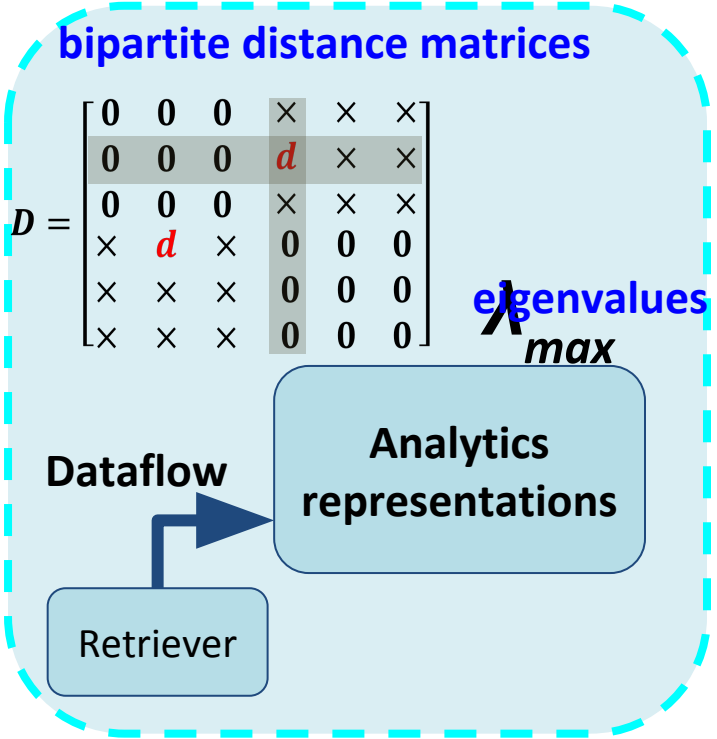

T. Johnston et al. In-Situ Data Analytics and Indexing of Protein Trajectories. *Journal of Computational Chemistry (JCC),* 38(16):1419-1430, 2017.

Analysis: Proxy for Performance

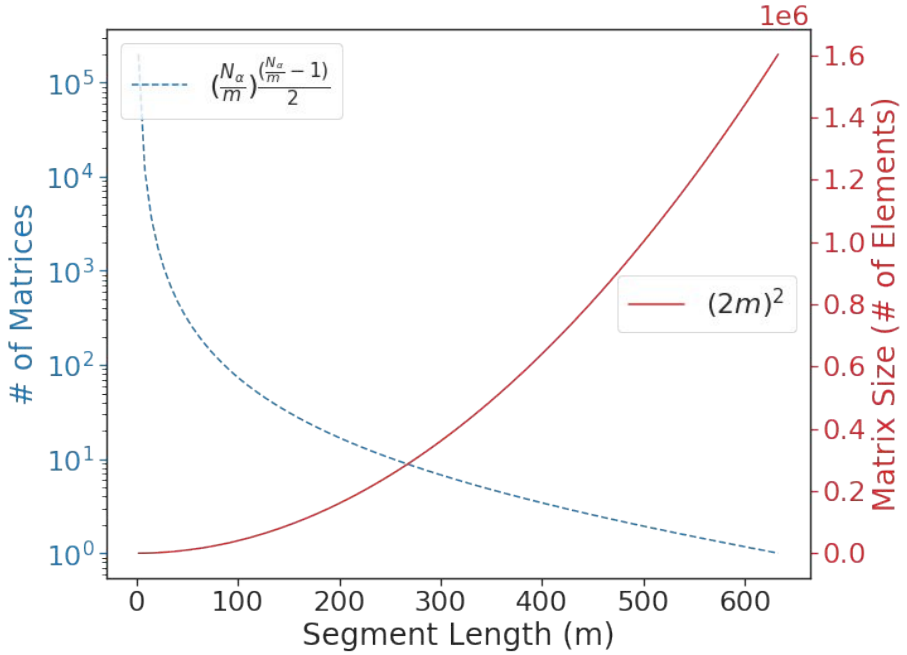

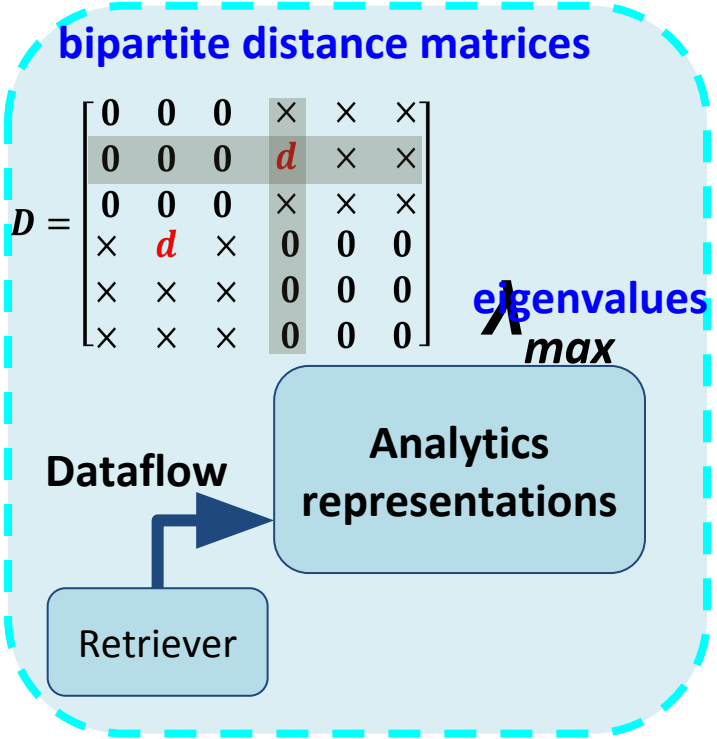

Data Analytics

Analysis: Proxy for Performance

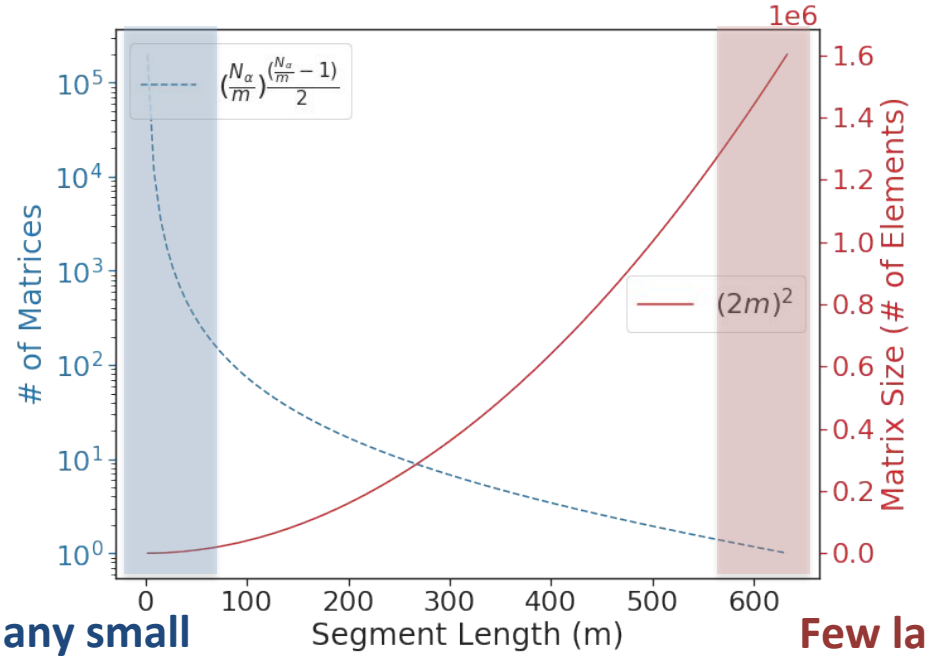

Many small matrices

Few large matrices

bipartite distance matrices

$$D = \begin{bmatrix} 0 & 0 & 0 & \times & \times & \times \\ 0 & 0 & 0 & d & \times & \times \\ 0 & 0 & 0 & \times & \times & \times \\ \times & d & \times & 0 & 0 & 0 \\ \times & \times & \times & 0 & 0 & 0 \\ \times & \times & \times & 0 & 0 & 0 \end{bmatrix}$$

eigenvalues λ_{max}

Dataflow

Analytics representations

Retriever

Data Analytics

Frame at time *t:*
Two α-helixes

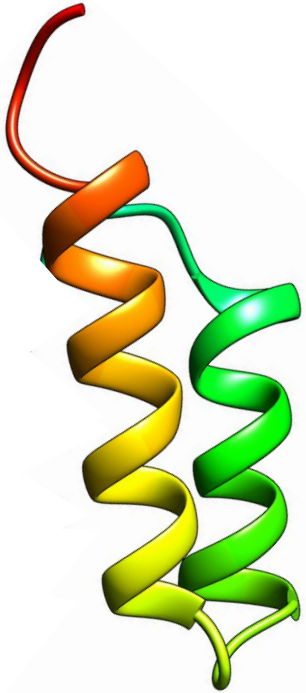

Frame at time *t:*

Sequence of Nα C^{α} atoms

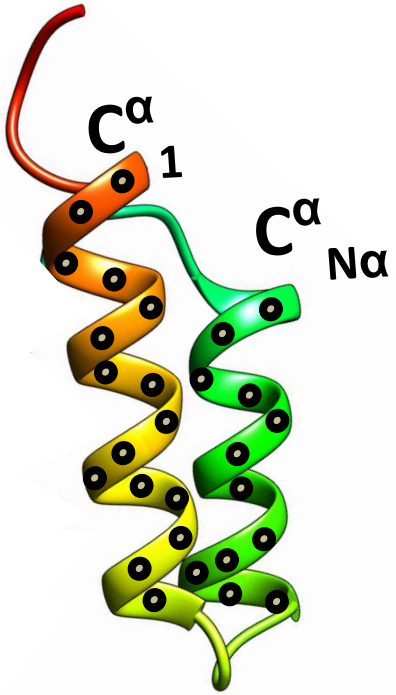

Frame at time *t:*

Sequence of Nα C^{α} atoms

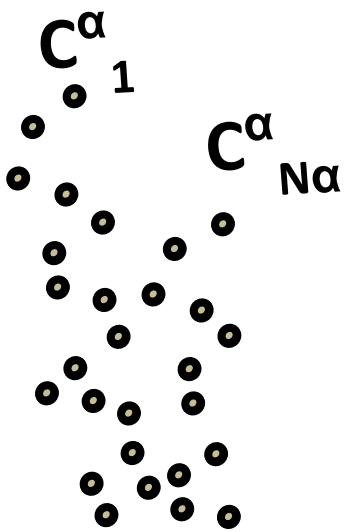

Frame at time *t:*
Sequence of Nα C^{α} atoms

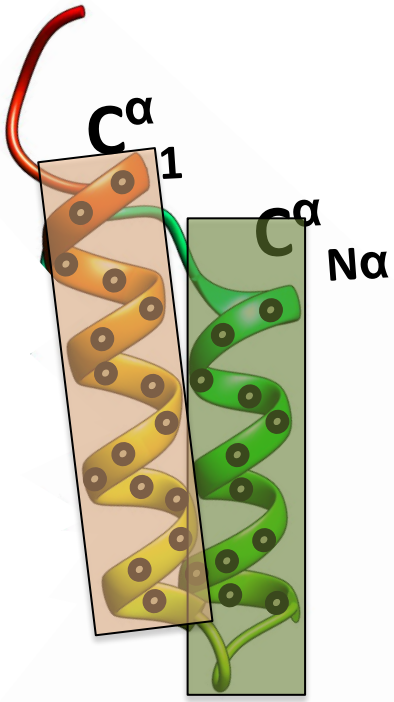

Few large matrices

Distances of two segments with segment length:

Nα/2 * C^{α} atoms

Segments:

[C^{α}_{1} □ $C^{\alpha}_{N\alpha/2\ -1}$]

[$C^{\alpha}_{N\alpha/2}$ □ $C^{\alpha}_{N\alpha/2}$]

Metadata: Single $\boldsymbol{\lambda}_{max}$

	C^{α}_{1}	$C^{\alpha}_{N\alpha}$
C^{α}_{1}	0	d_{ij}
$C^{\alpha}_{N\alpha/2\ -1}$		
$C^{\alpha}_{N\alpha/2}$	dji	0
$C^{\alpha}_{N\alpha}$		

Frame at time *t:*

Sequence of Nα C^{α} atoms

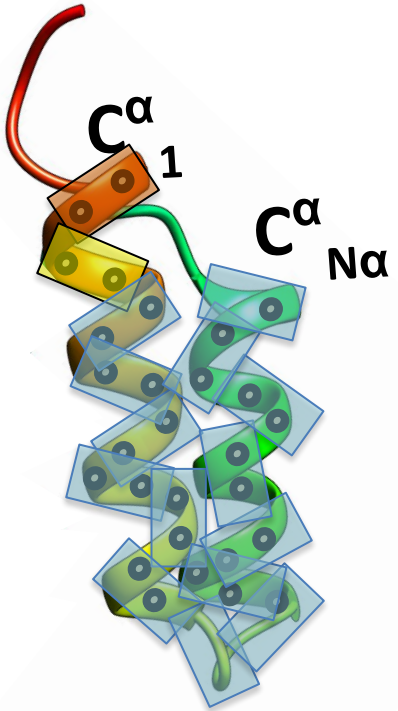

Many small matrices

Distances of Nα/2 segments with segment length:

2 * C^{α} atoms

Segments :

[C^{α}_{1} □ C^{α}_{2}] [C^{α}_{3} □ C^{α}_{4}]

[C^{α}_{5} □ C^{α}_{6}] [C^{α}_{7} □ C^{α}_{8}]

….

[$C^{\alpha}_{N\alpha/2-3}$ □ $C^{\alpha}_{N\alpha/2-2}$][$C^{\alpha}_{N\alpha/2-1}$ □ $C^{\alpha}_{N\alpha/2}$]

Metadata:

$\lambda_{max,\,1}\ \lambda_{max,\,2} \ldots\ \lambda_{max,\,N\alpha/2}$

	$C^{\alpha}_{1.}$ C^{α}_{2}	C^{α}_{3} C^{α}_{4}
C^{α}_{1} C^{α}_{2}	0	dij
C^{α}_{3} C^{α}_{4}	dji	0

Frame at time *t:*

Sequence of Nα C^{α} atoms

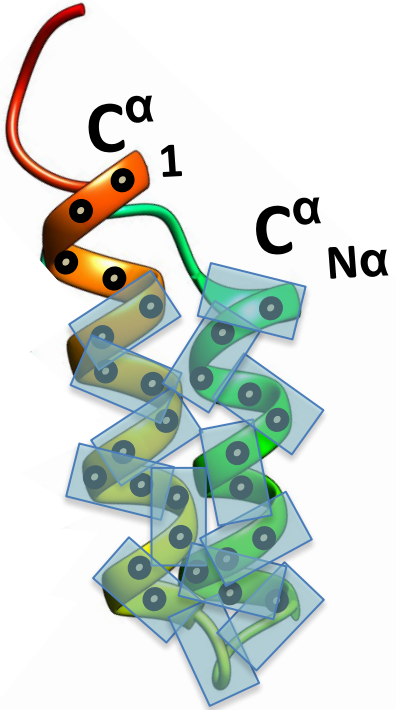

Many small matrices

Distances of Nα/2 segments with segment length:

2 * C^{α} atoms

Segments :

[C^{α}_{1} □ C^{α}_{2}] [C^{α}_{3} □ C^{α}_{4}]

[C^{α}_{5} □ C^{α}_{6}] [C^{α}_{7} □ C^{α}_{8}]

….

[$C^{\alpha}_{N\alpha/2-3}$ □ $C^{\alpha}_{N\alpha/2-2}$][$C^{\alpha}_{N\alpha/2-1}$ □ $C^{\alpha}_{N\alpha/2}$]

Metadata:

$\lambda_{max,\,1}\ \lambda_{max,\,2} \ldots\ \lambda_{max,\,N\alpha/2}$

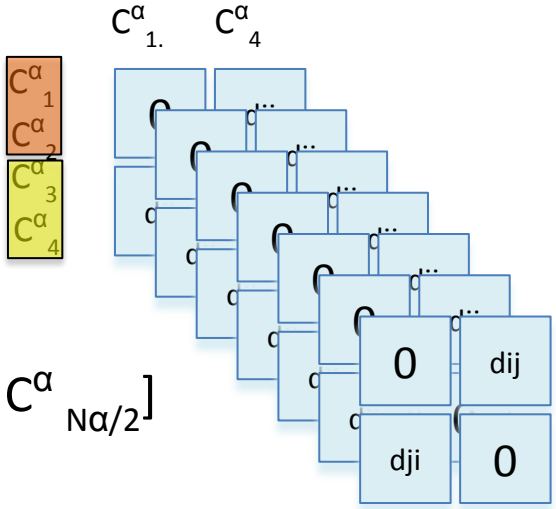

Segment size = proxy of **number of matrices** and **matrix sizes**

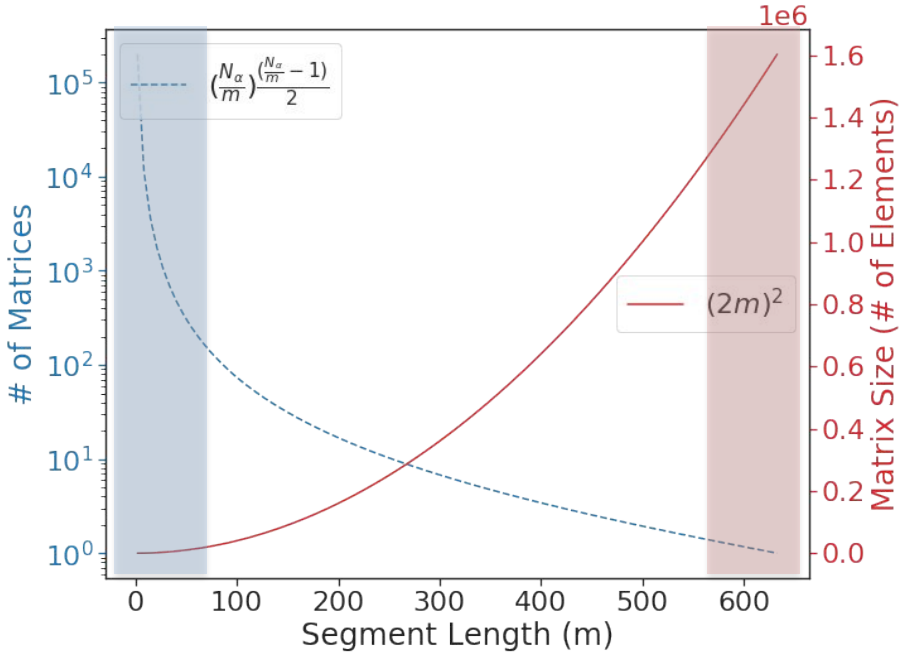

Many small matrices

Few large matrices

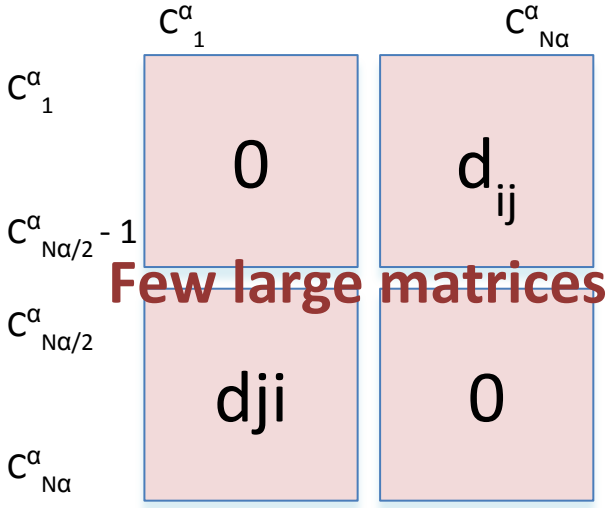

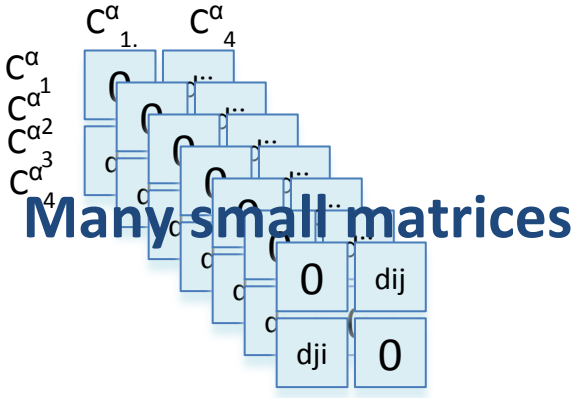

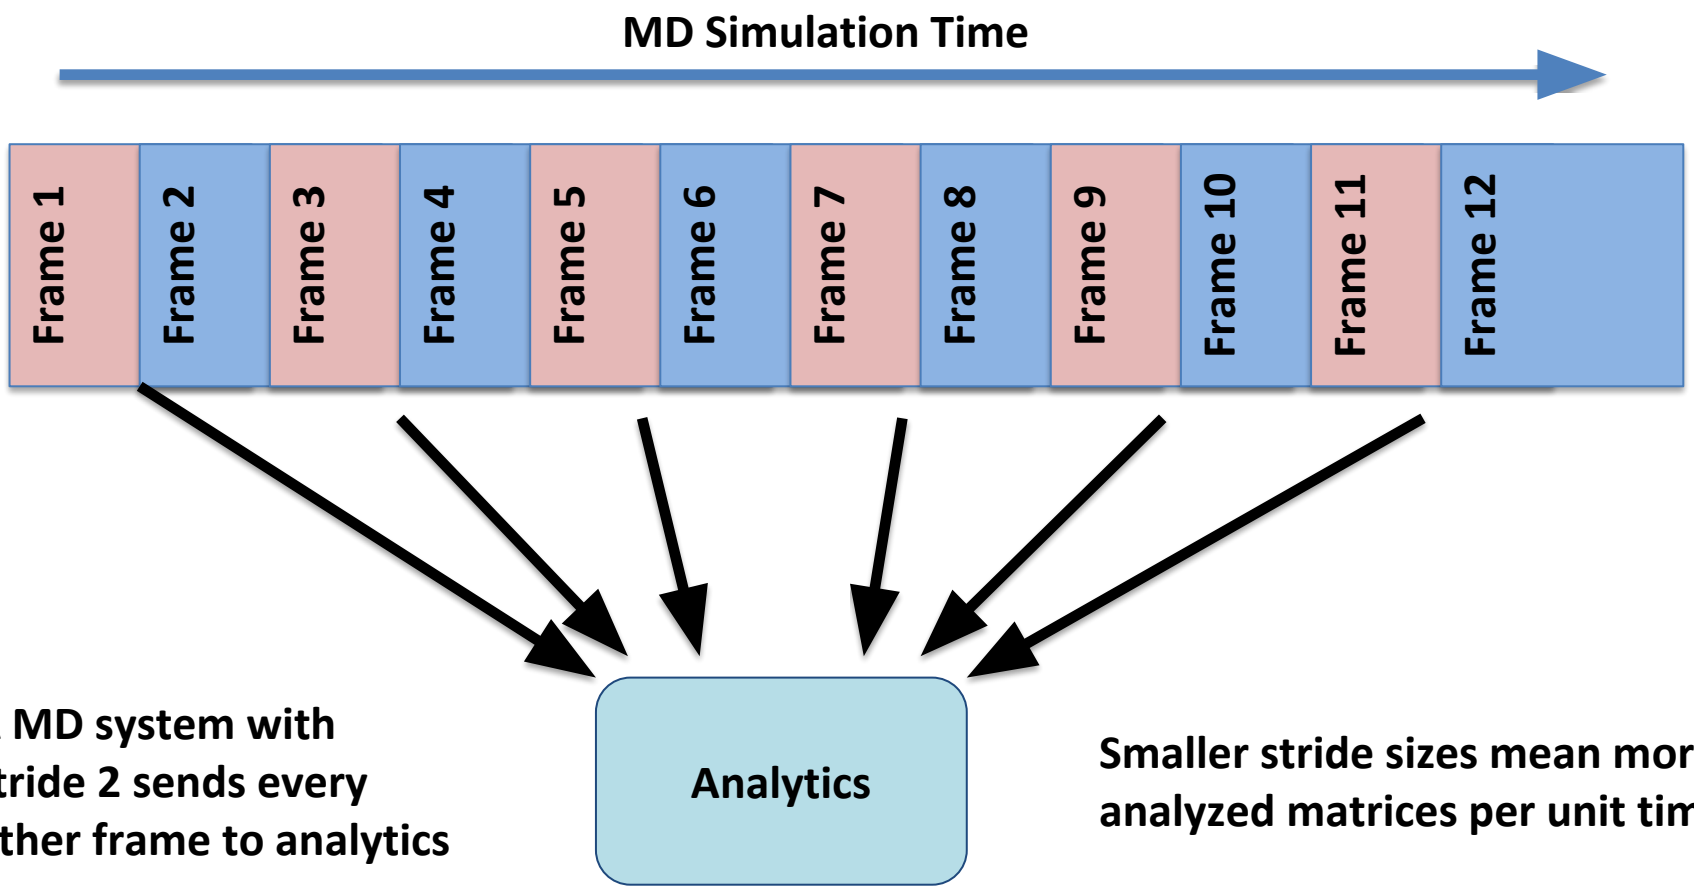

A MD system with stride 2 sends every other frame to analytics

Smaller stride sizes mean more analyzed matrices per unit time

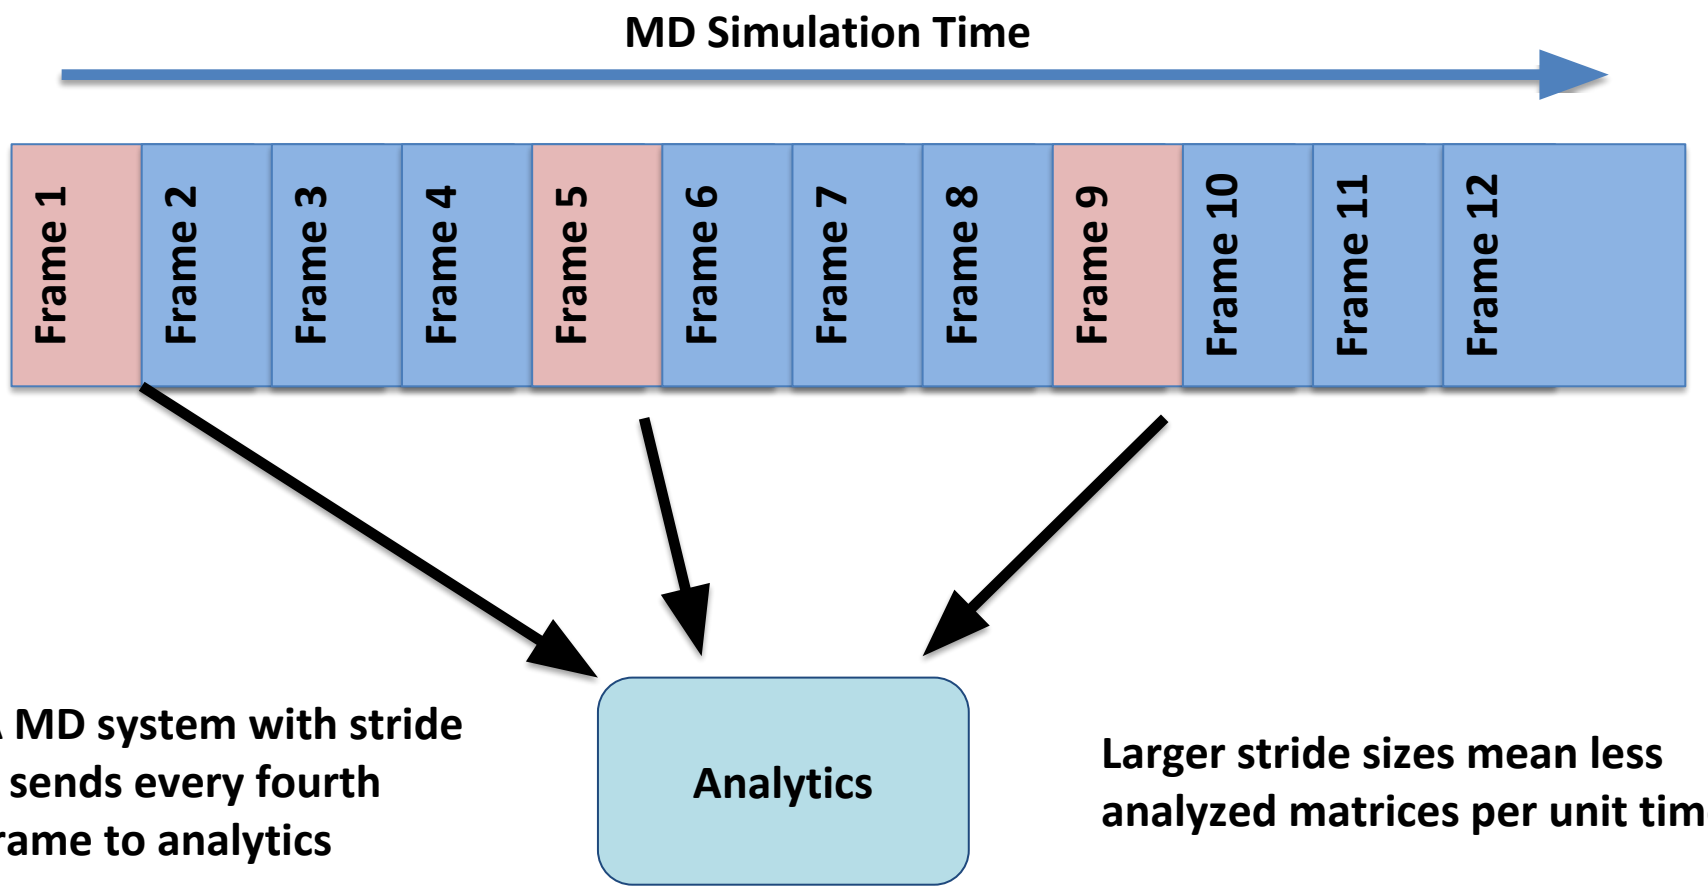

A MD system with stride 4 sends every fourth frame to analytics

Larger stride sizes mean less analyzed matrices per unit time

Dataflow Modeling for Analytics

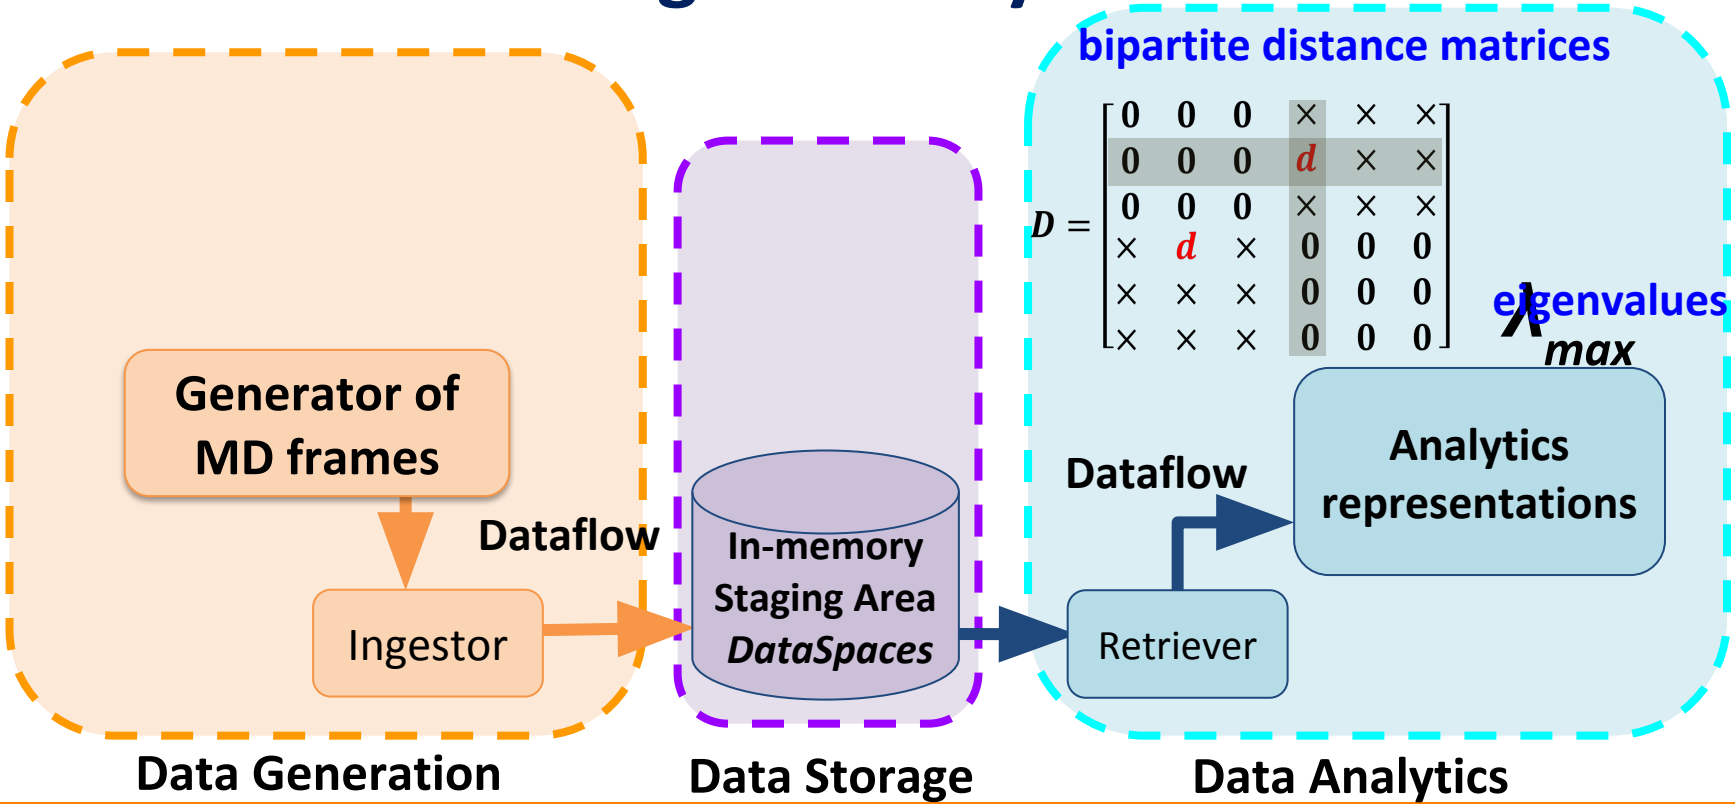

Building a Closed-loop Workflow

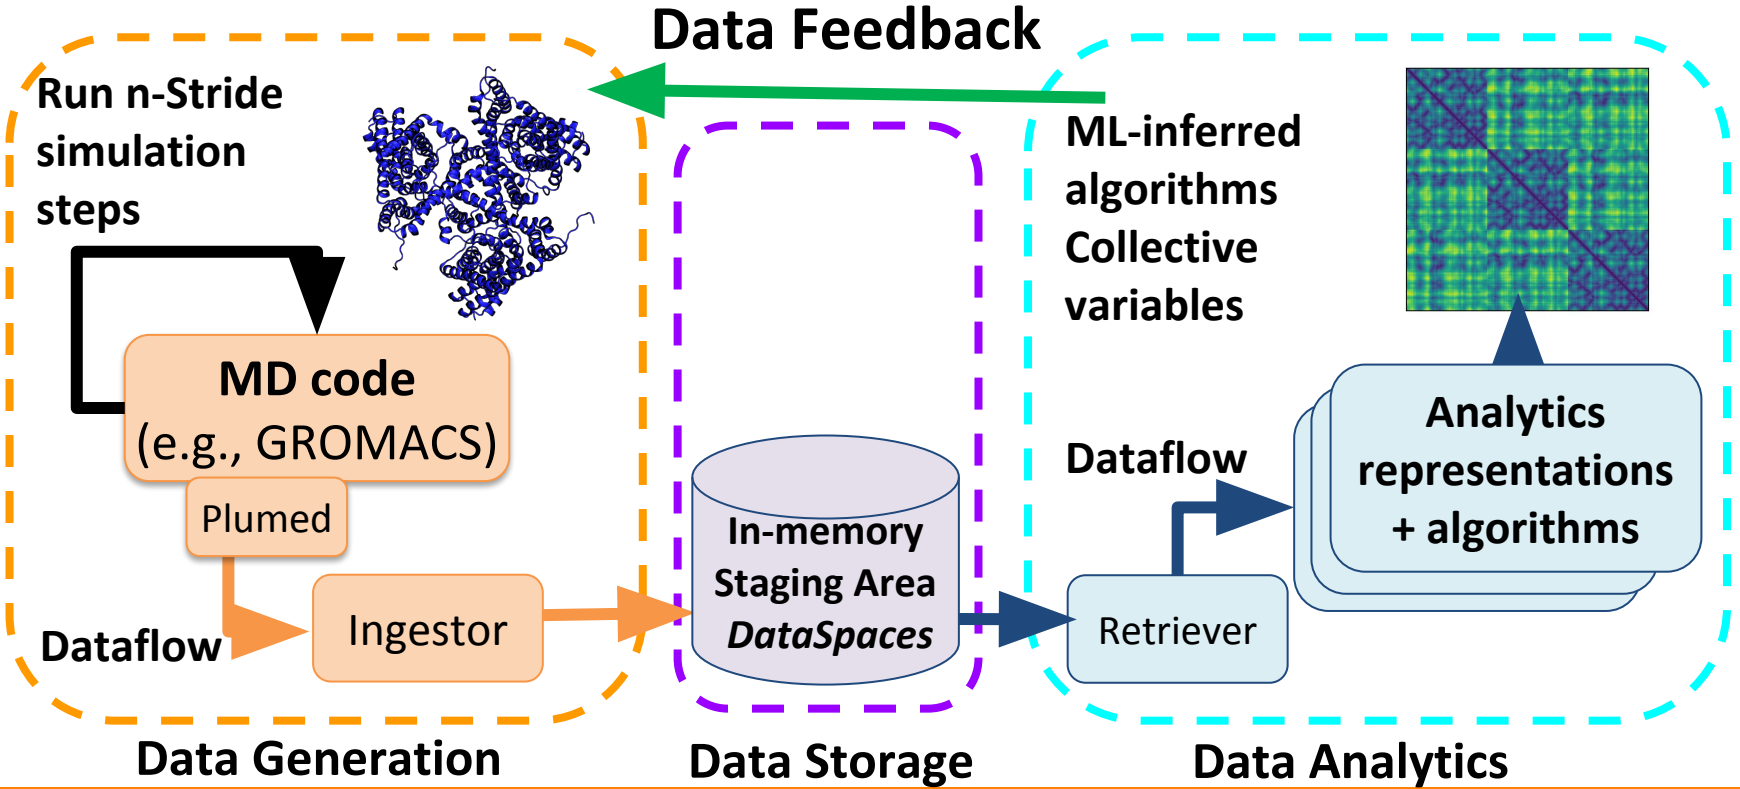

Modeling Idle Times

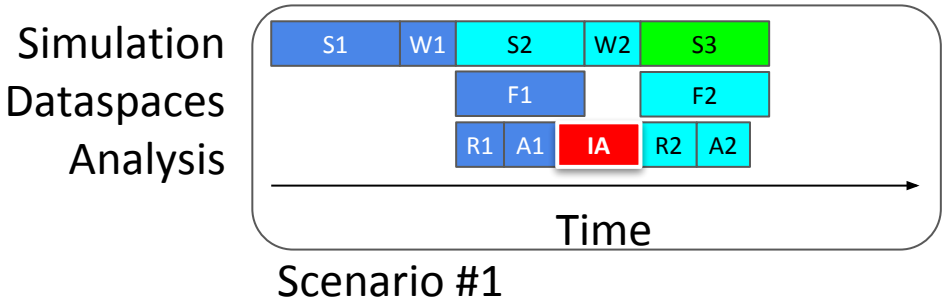

Scenario #1

S1, S2, S3: Generate MD frame
W1, W2, W3: Write to shared memory
R1, R2, R3: Read from shared memory
A1, A2: Analyze frame

Modeling Idle Times

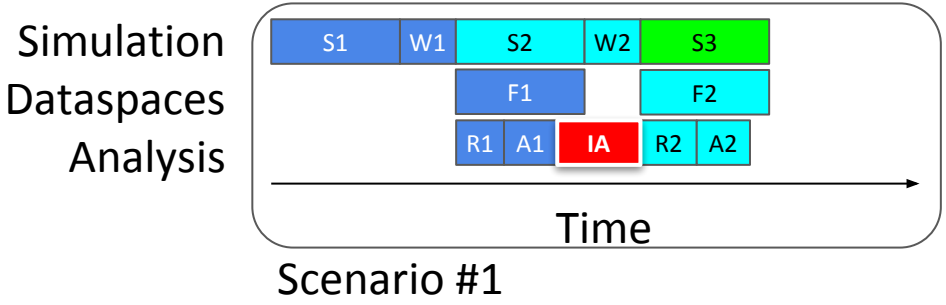

Scenario #1

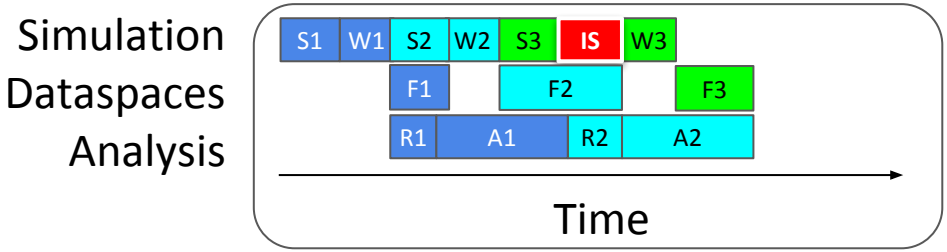

Scenario #2

S1, S2, S3: Generate MD frame
W1, W2, W3: Write to shared memory
R1, R2, R3: Read from shared memory
A1, A2: Analyze frame

Modeling Idle Times

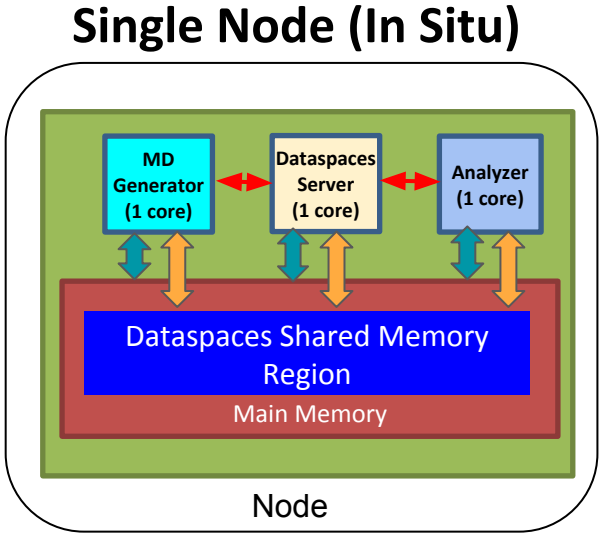

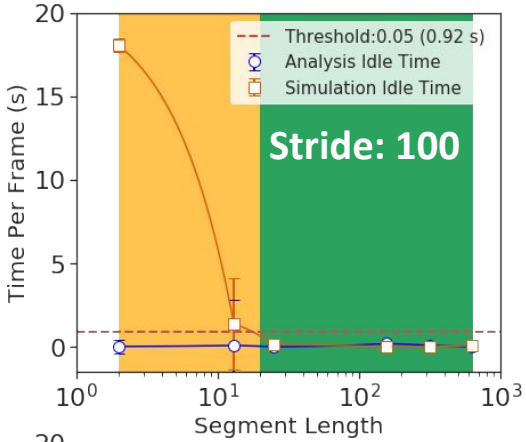

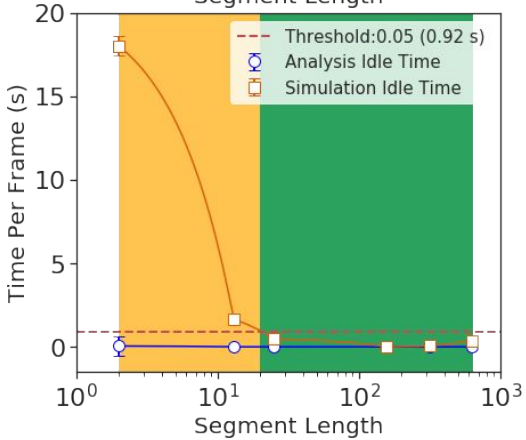

Simulation Idle

Balanced

Analysis Idle

Modeling Lost Frames

Trp cage 12,619 atoms

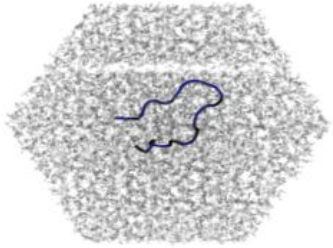

T cell receptor 81,092 atoms

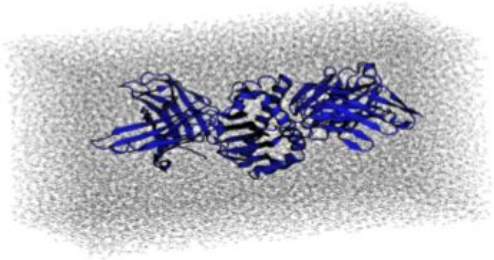

Gltph 270,088 atoms

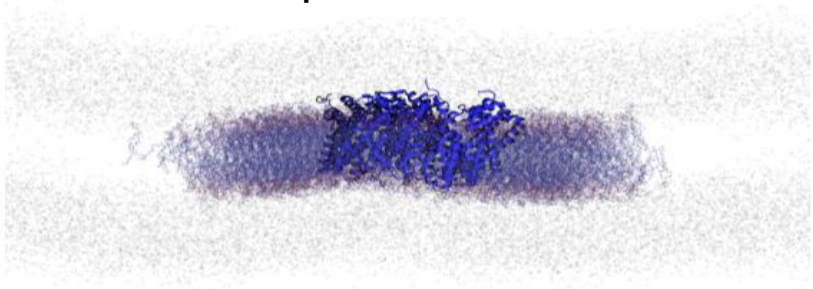

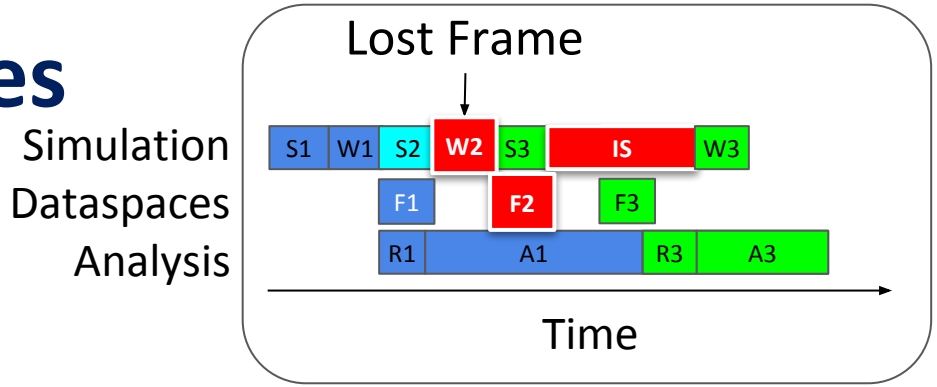

Scenario #3

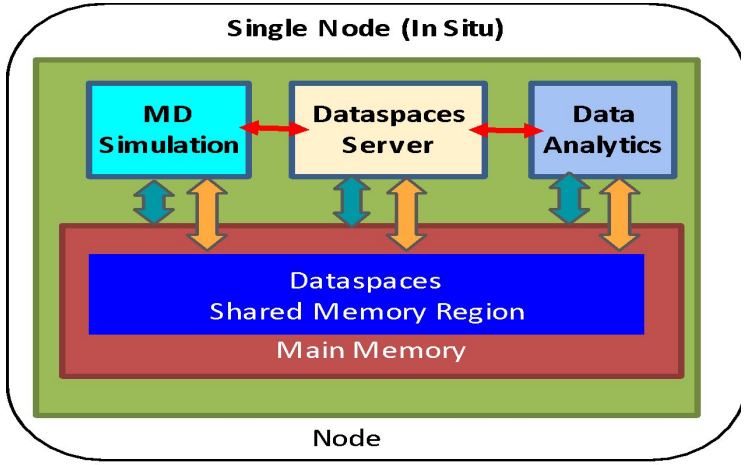

2-step Model

Observed Data:

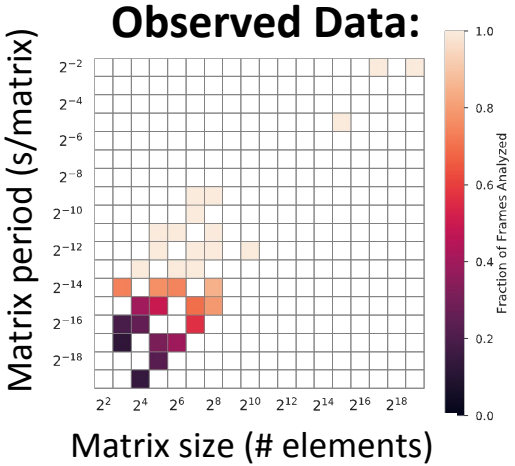

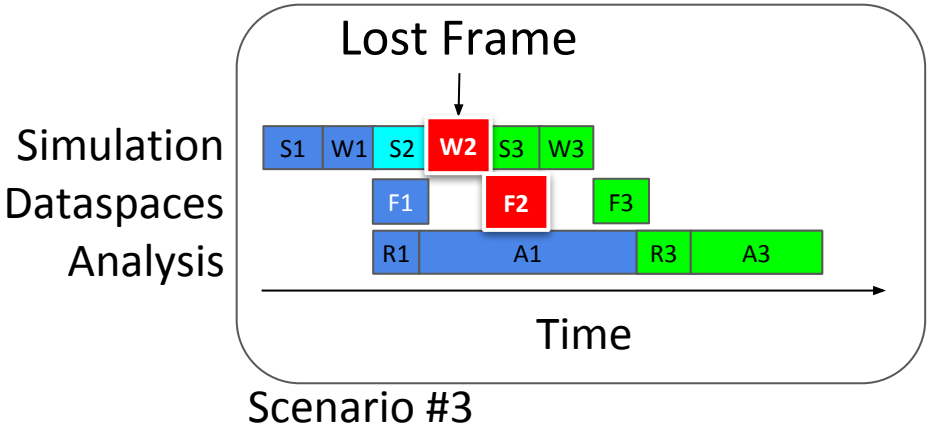

1 Model fraction of analyzed frames

2 Model which frames are analyzed

2-step Model

MD Workflow:

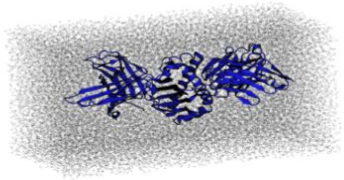

T cell receptor 81,092 atoms

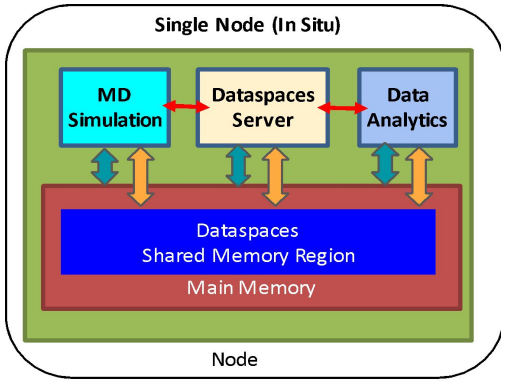

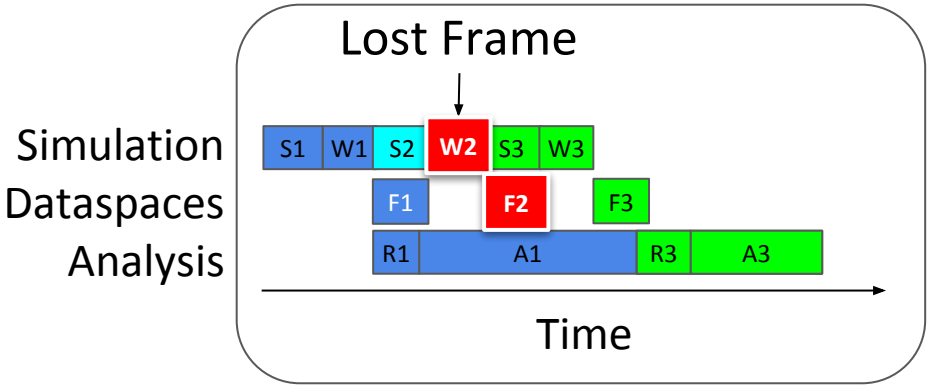

Scenario #3

1 Model fraction of analyzed frames

2 Model which frames are analyzed

2-step Model: Fraction of Analyzed Frames

Observables small segment lengths
(i.e., 2 , 4 , 6 , 8 , 10 , 12 , 14, 16, and 18)

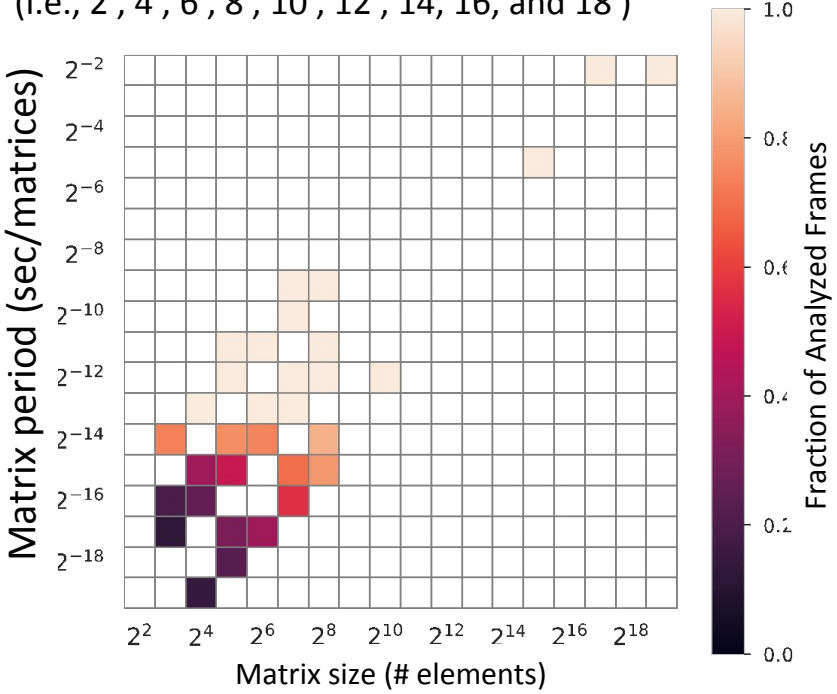

Model: polynomial model of degree 2
obtained by least-square fitting observables

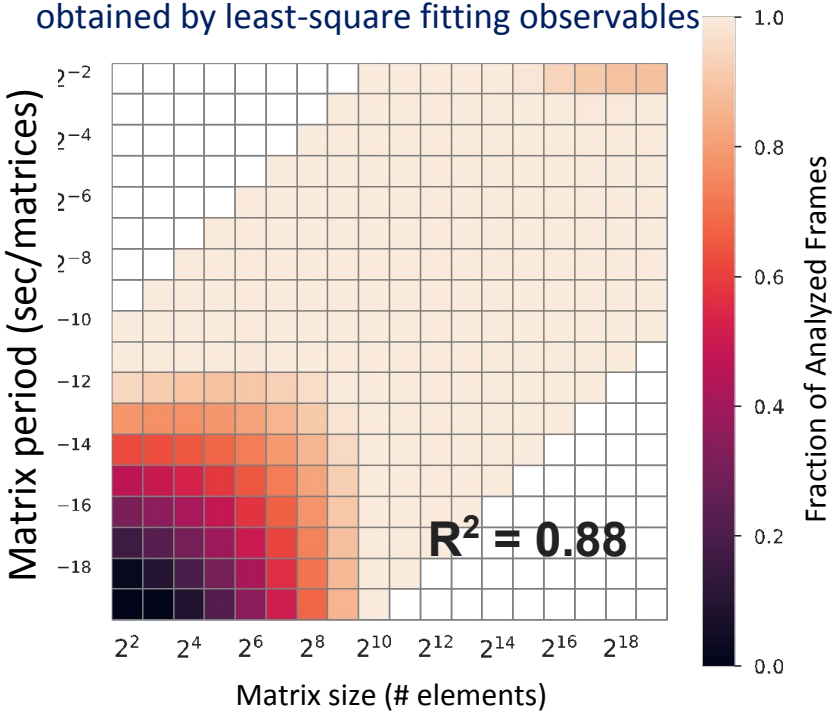

2-step Model: Fraction of Analyzed Frames

Model: polynomial model of degree 2
obtained by least-square fitting observables

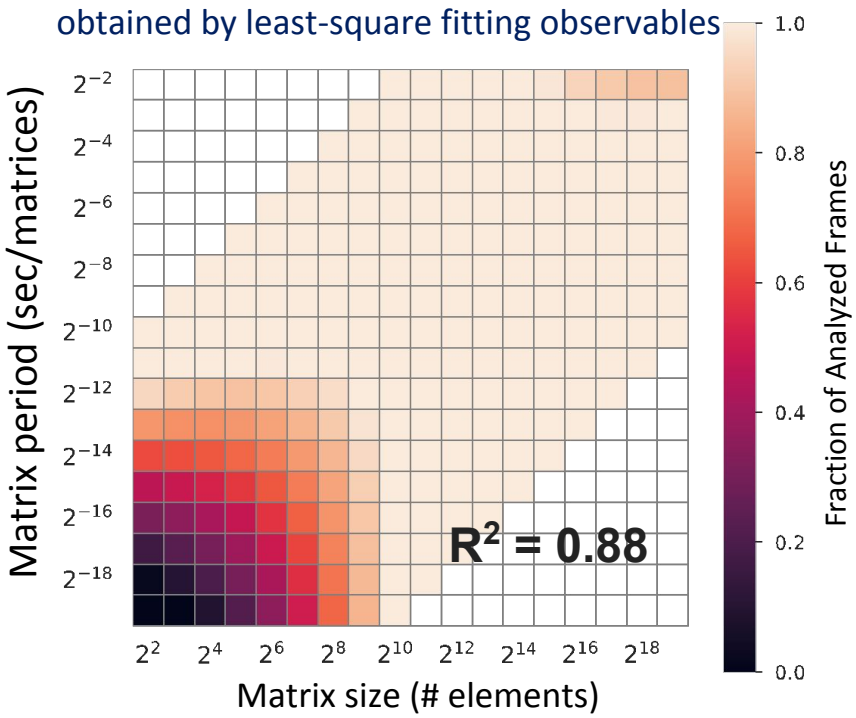

Error: Absolute error between data
and fitting model

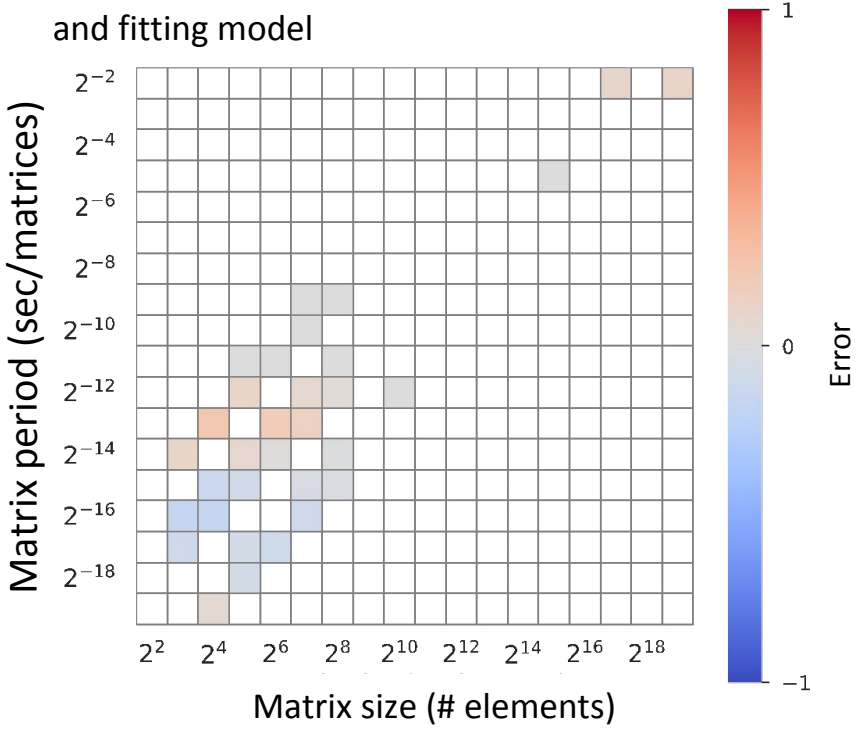

2-step Model: Frames Distribution

- Given a trajectory, we model the proportions *p* and *q* of analyzed frames (*f*) with periods *k* and *k+1*
  - Example: Gltph (27,000 atoms and TPS 318), trajectory of 1,000 frames.

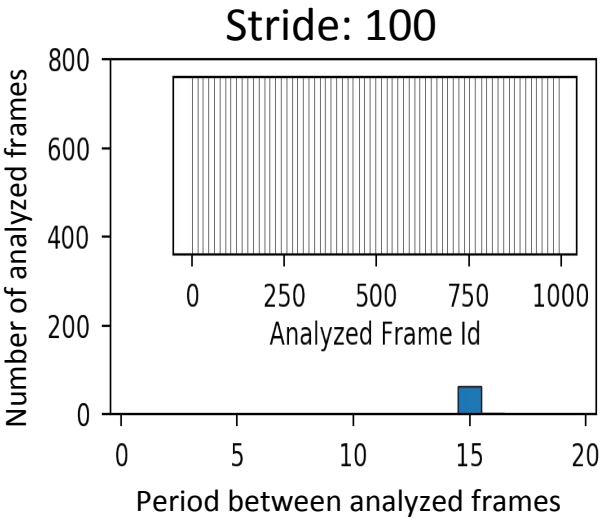

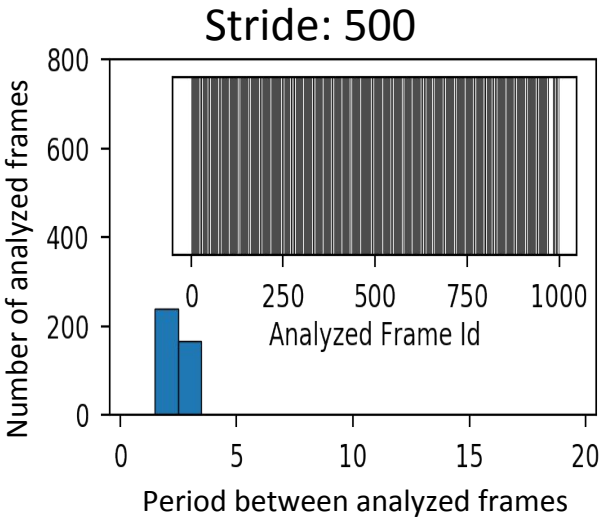

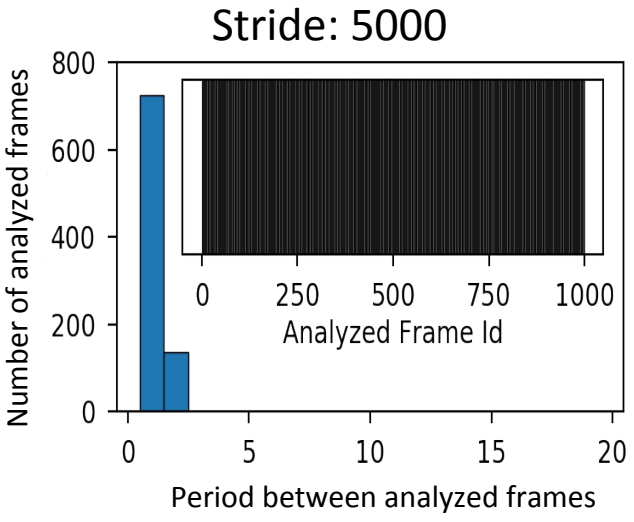

Case Study: 1BDD Protein Conformations

Frame 1330

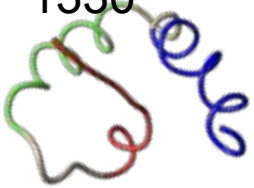

Frame 1360

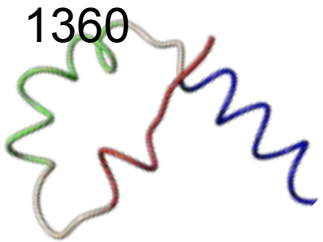

Frame 1390

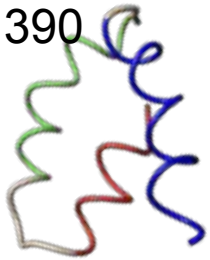

Case Study: 1BDD Protein Conformations

Frame 1330

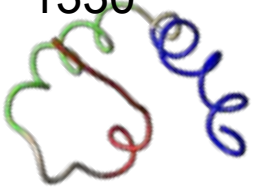

Frame 1360

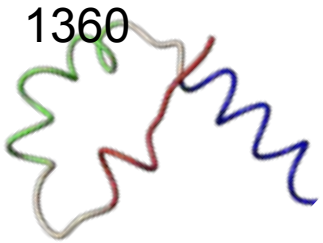

Frame 1390

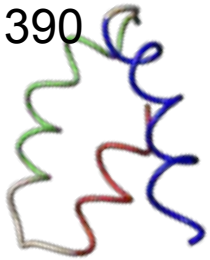

Helix 1-2: λ_{max}

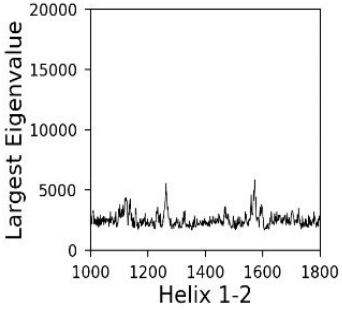

Helix 1-3: λ_{max}

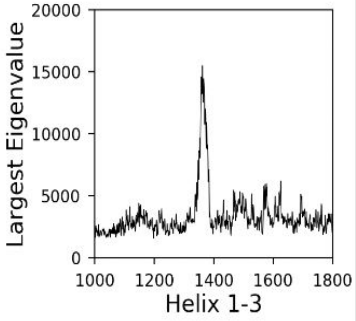

Helix 2-3: λ_{max}

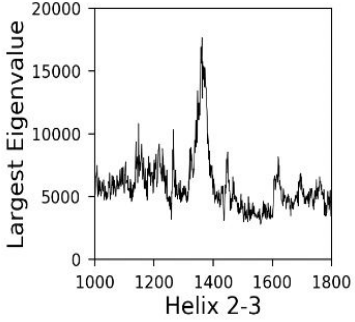

Case Study: 1BDD Protein Conformations

Frame 1330

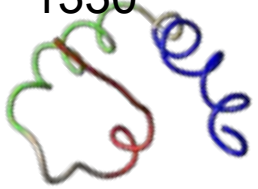

Frame 1360

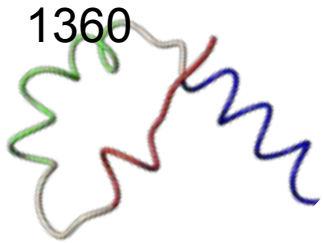

Frame 1390

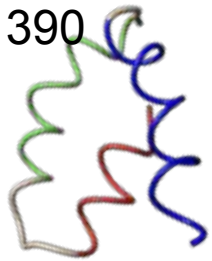

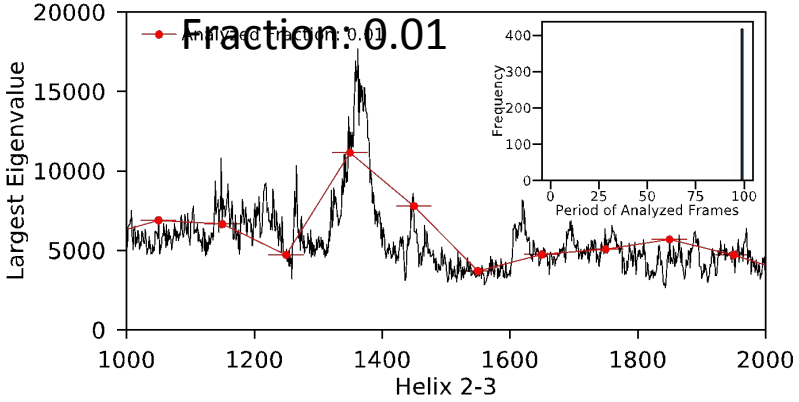

Helix 1-2: λ_{max}

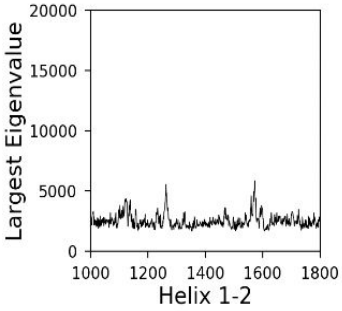

Helix 1-3: λ_{max}

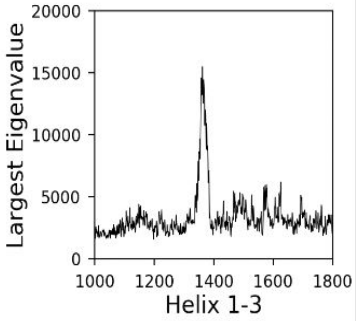

Helix 2-3: λ_{max}

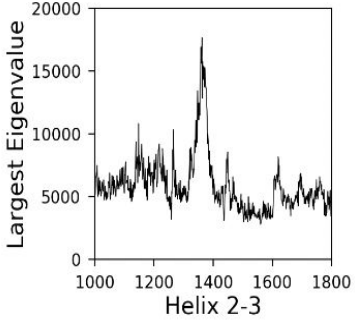

Case Study: 1BDD Protein Conformations

Frame 1330

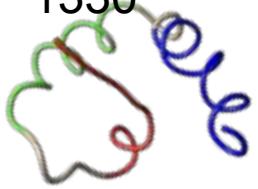

Frame 1360

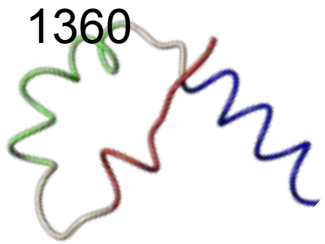

Frame 1390

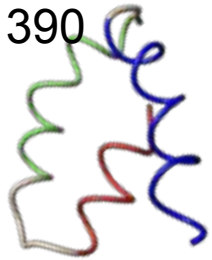

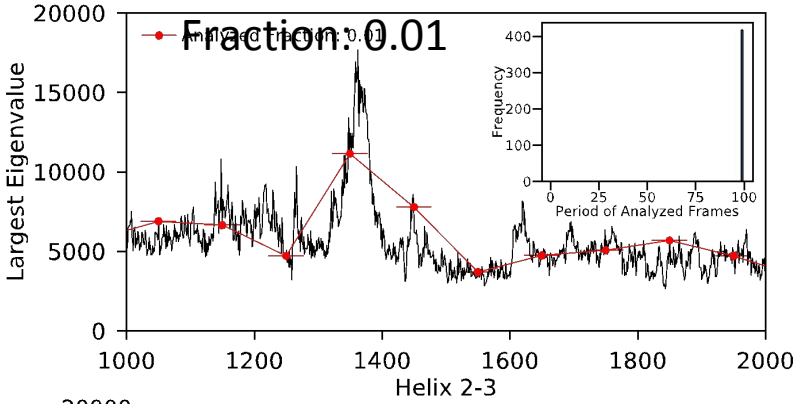

Helix 1-2: λ_{max}

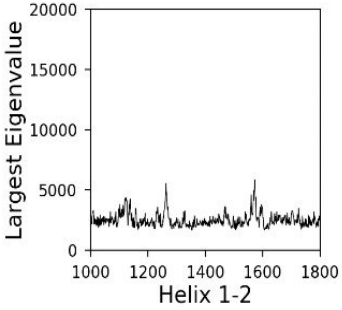

Helix 1-3: λ_{max}

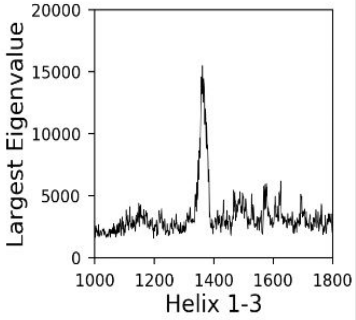

Helix 2-3: λ_{max}

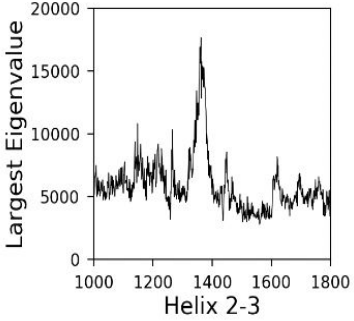

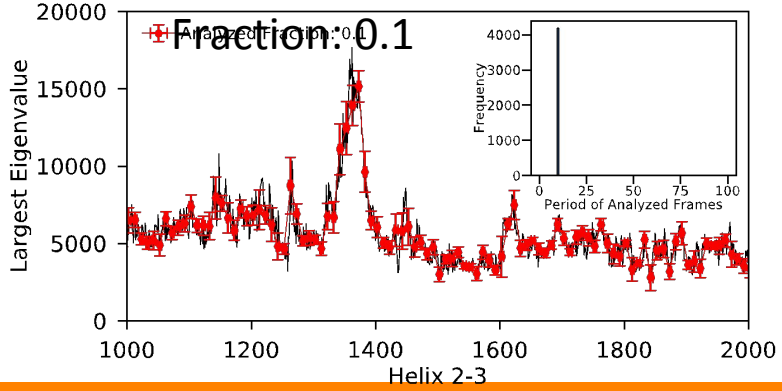

Lessons Learned

- We measure and analyze the execution patterns associated with in situ and in transit dataflows
- We build a 2-step model to predict which frames are analyzed, given a molecular system, analytic module, and analytic parameters
- We apply our 2-step model to a case study to understand the impact of analytics parameters on capturing rare MD simulation events
- Future direction: leverage ML to drive decisions on what frames to capture and store at runtime based on scientific information