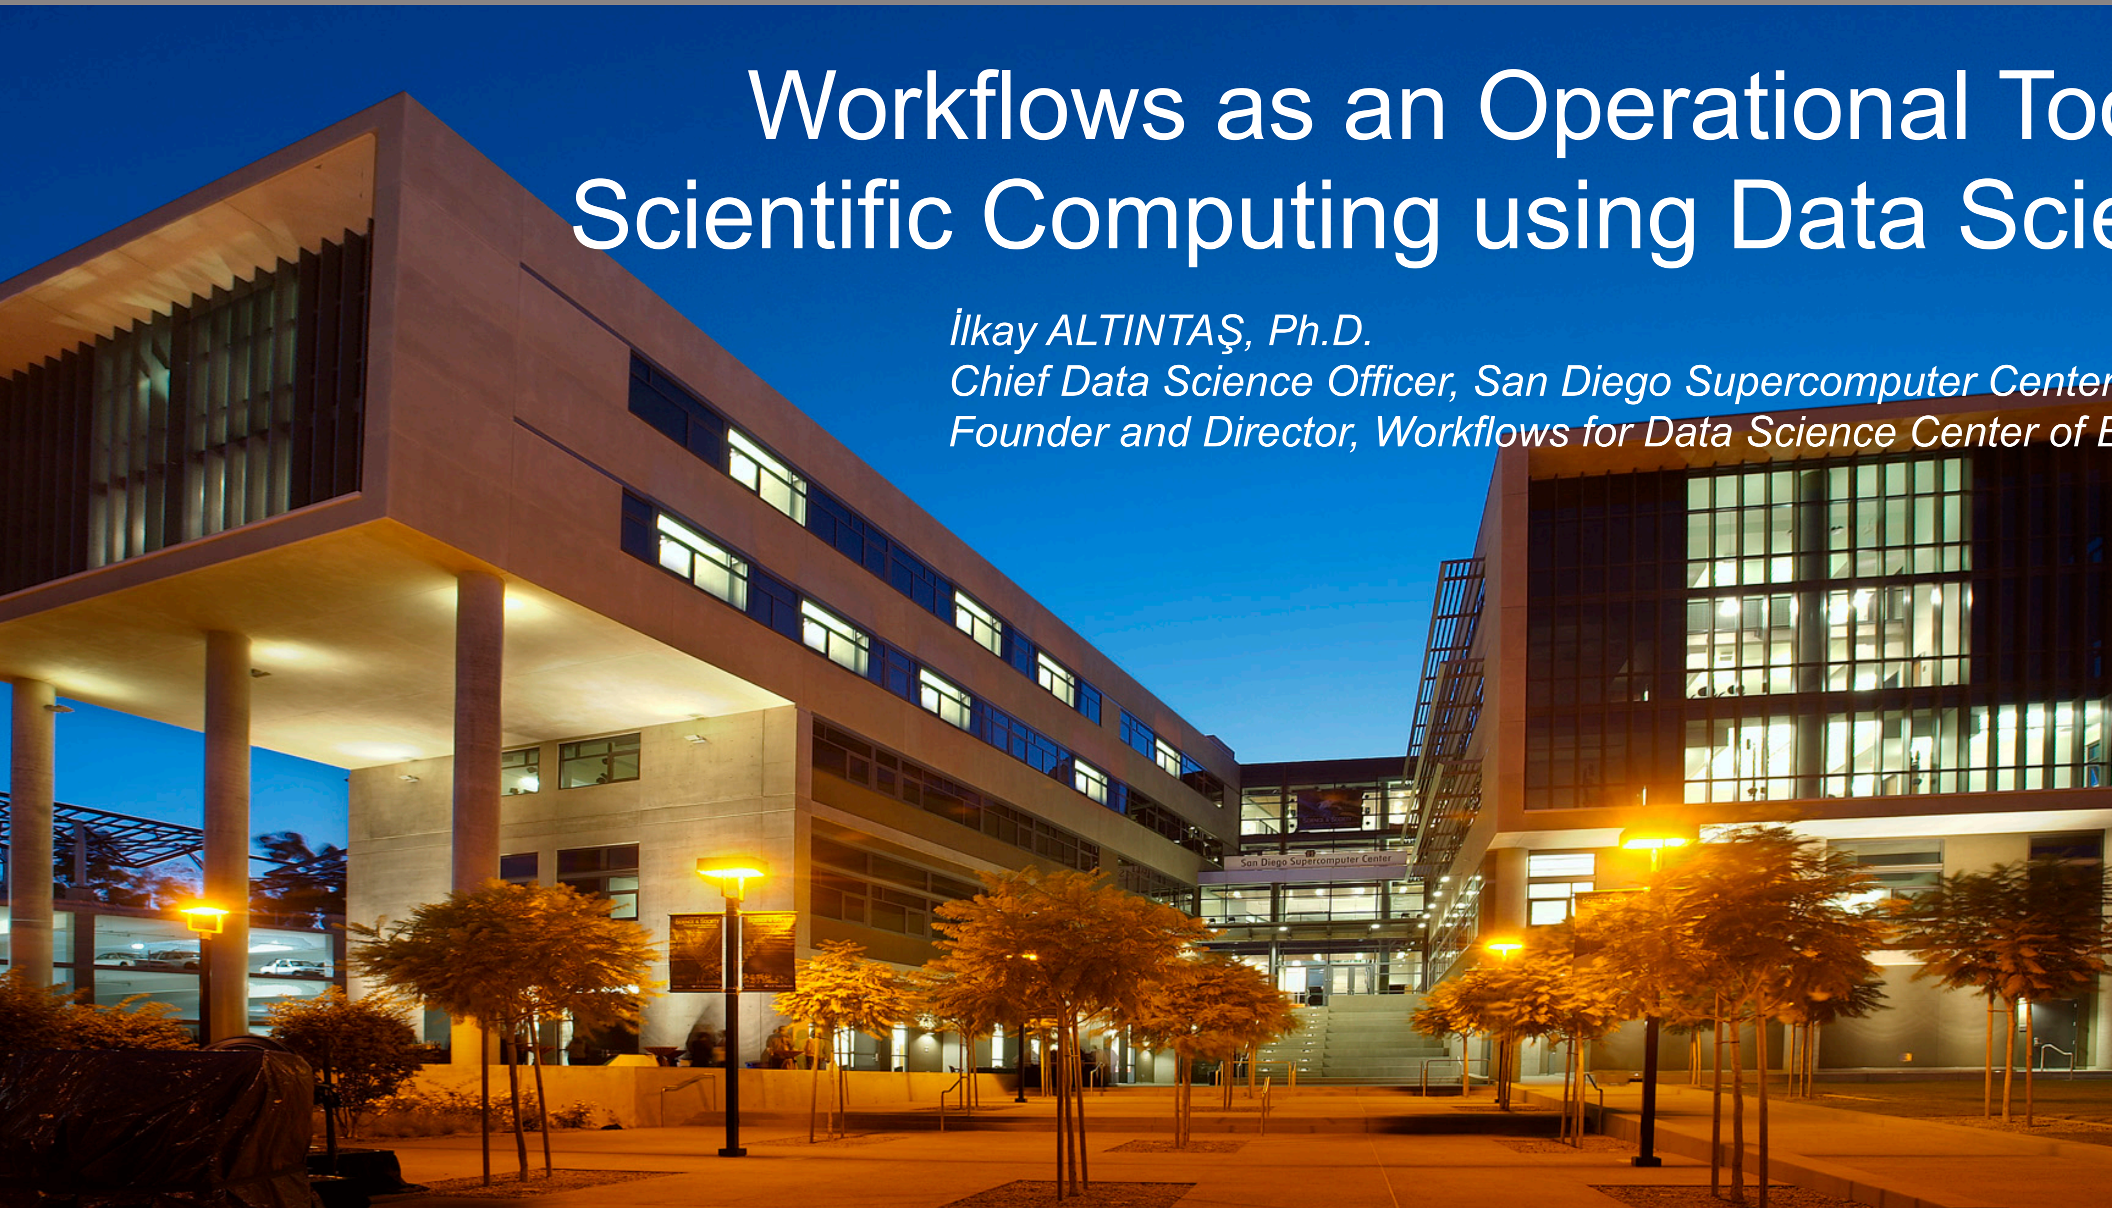

Workflows as an Operational To
Scientific Computing using Data Scie

İlkay ALTINTAŞ, Ph.D.
Chief Data Science Officer, San Diego Supercomputer Center
Founder and Director, Workflows for Data Science Center of E

SC is 31 Years Young!
viding Cyberinfrastructure for Research and Educati

- tablished as a national percomputer resource center in 1985 NSF
- world leader in HPC, data-intensive mputing, and scientific data anagement
- rrent strategic focus on "Big Data" d "HPC Cloud" : versatile computing

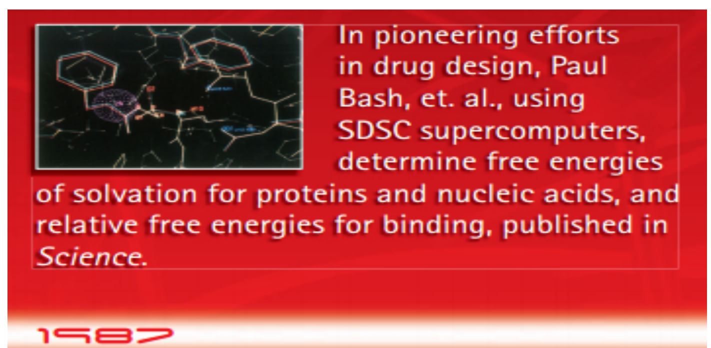

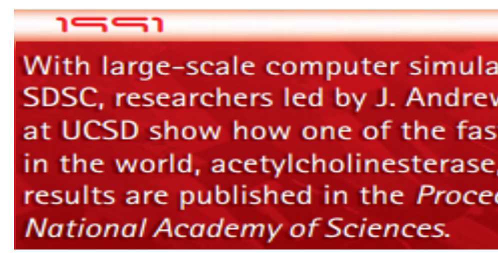

Two discoveries in design from 1987 an

SDSC continues to focus on versatile computing and big dat

ne-Step Molecular Dynamics through Hydrogen Mass
oning

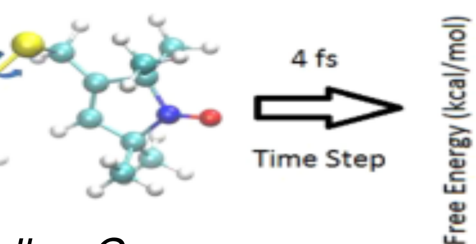

alker Group

anosecond level snapshots of financial markets

oi[2], David O'Neal[3], Mao Ye[1] and Robert S. Sinkovits[2,*,†]

original code	Wall time (s) modified code	Speedup
400	128	66x
200	437	126x
914	1145	113x

1 Comment 149

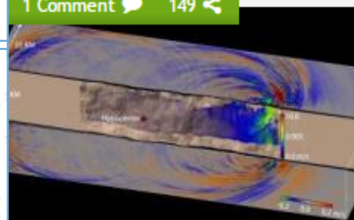

CORPORATE RESPONSIBILITY

San Diego Supercomputer Center's Quake Research Wins $150,000 Global Impact Award

By Tonie Hansen on March 16, 2015

Multi-GPU Implementation of a 3D Finite Difference Time Domain Earthquake Code on Heterogeneous Supercomputers

Jun Zhou[a,b,*], Yifeng Cui[a], Efecan Poyraz [a,b], Dong Ju Choi [a], Clark C. Guest [b]

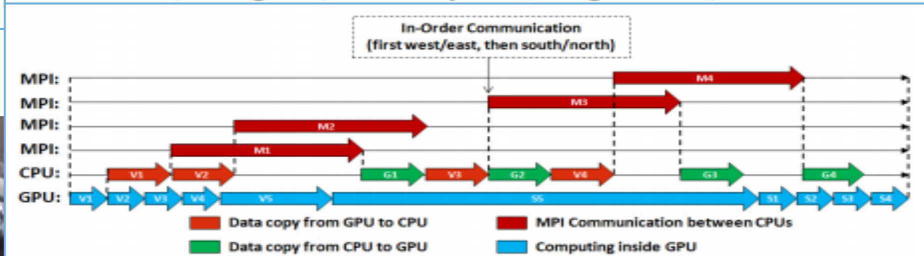

riant calling leveraging next-gene
for large-scale whole-genome se

[2], Tristan M. Carland[2], Glenn K. Lockwood[3], Wayne
Manidhar Tatineni[3], C Chris Huang[4], Sarah Lamberth[4], Yauheniya Cherkas[4]
Brodmerkel[4], Ed Jaeger[4,5], Lance Smith[4,5], Gunaretnam Rajagopal[4,5], Mark
Nicholas J. Schork[2,*]

Step	Tool	Memory per command (GB)	Cores per command
Map	BWA	32	8
Bam	Samtools	4	1
Merge	Samtools	4	1
Sort	Samtools	4	1
MarkDuplicates	PicardTools	7	2
TargetCreator	GATK	7	2
IndelRealigner*	GATK	12	3
BaseRecalibrator	GATK	30	8
PrintReads*	GATK	30	8
HaplotypeCaller	GATK	60	16

*Smaller memory allocation and more samples per node may prove computationally efficient

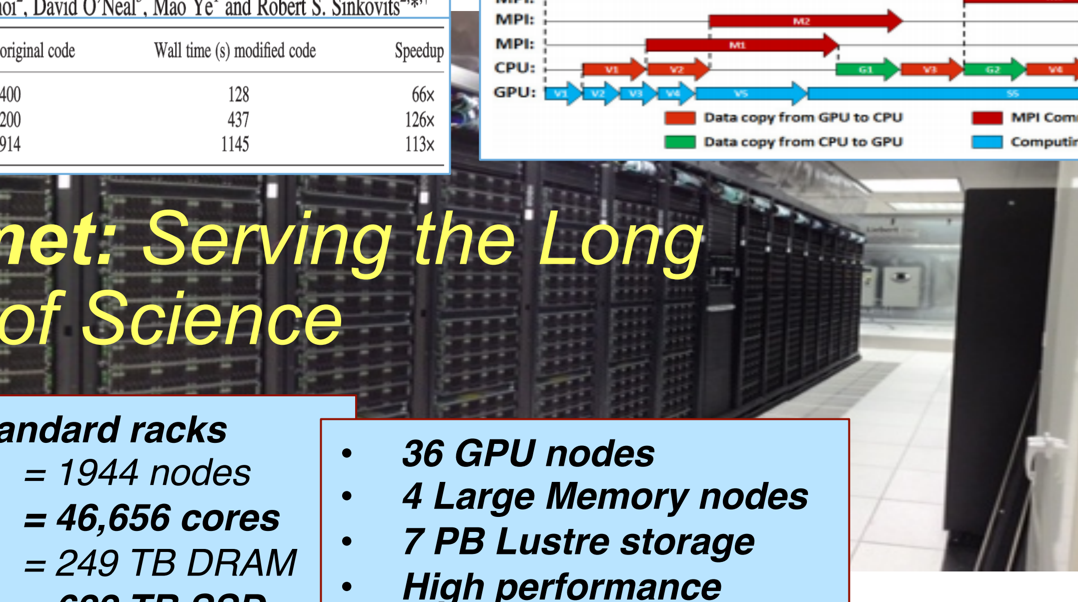

***net:** Serving the Long*
of Science

- ***andard racks***
- *= 1944 nodes*
- ***= 46,656 cores***
- *= 249 TB DRAM*
- ***= 622 TB SSD***
- ***flop/s***

- ***36 GPU nodes***
- ***4 Large Memory nodes***
- ***7 PB Lustre storage***
- ***High performance virtualization***

DSC Data Science Office

-- Expertise, Systems and Training for Data Science Applications --

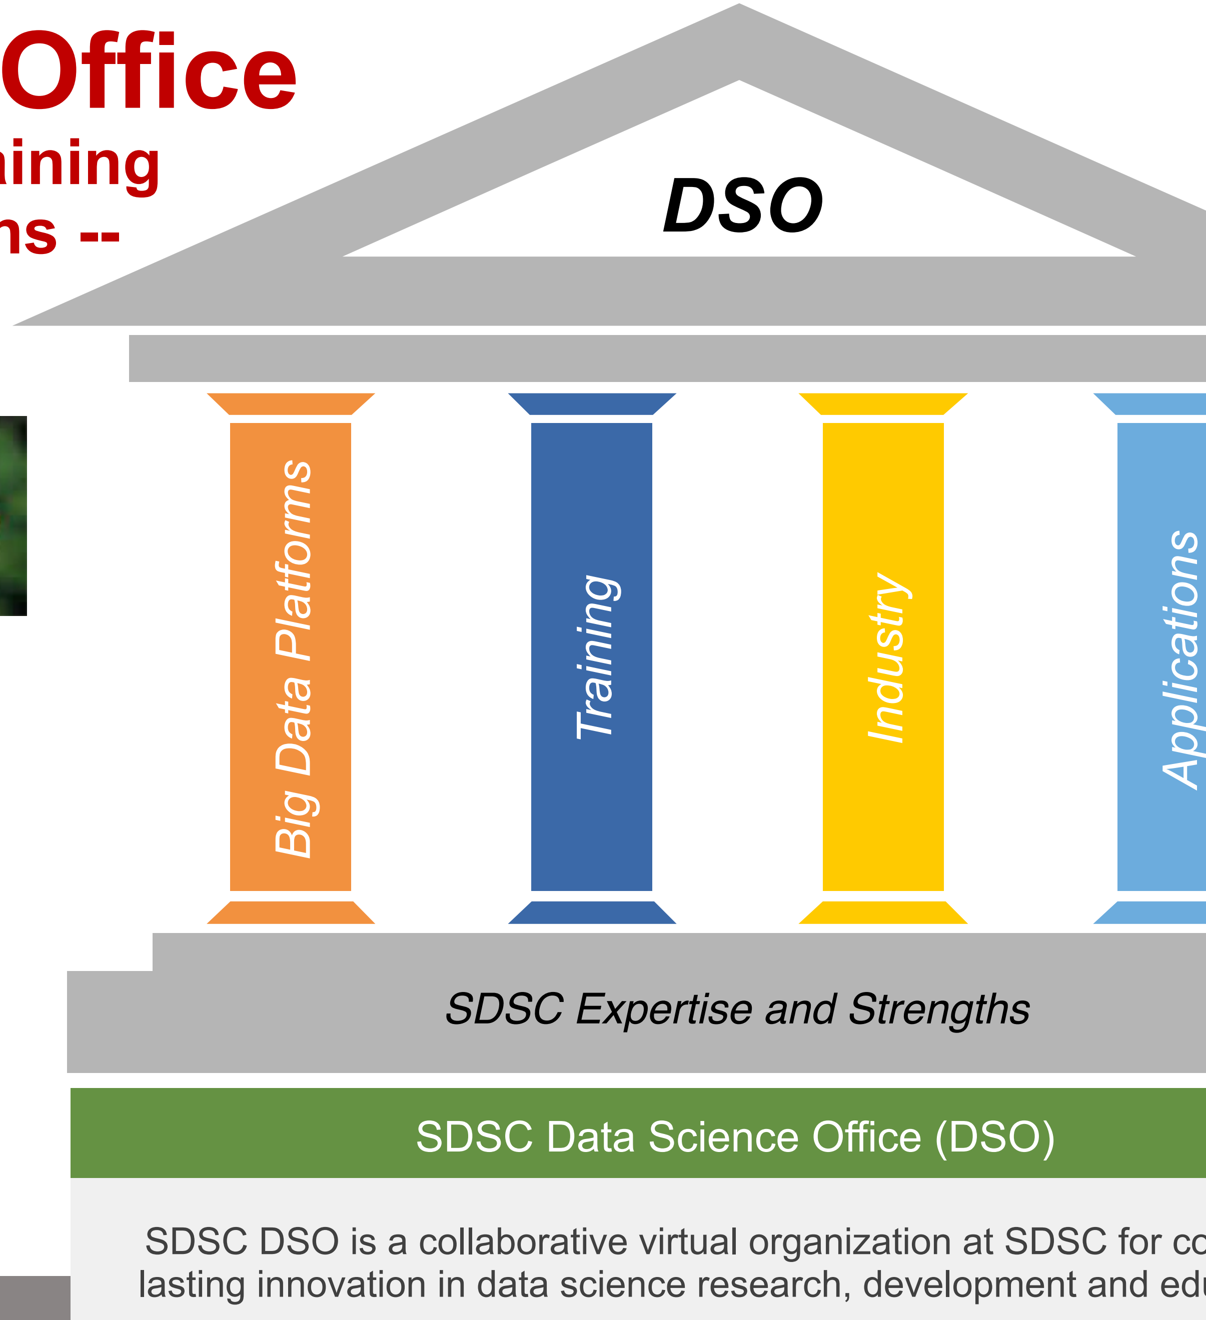

mputing Today has Many Shapes and Si

OMPUTING AT
SCALE

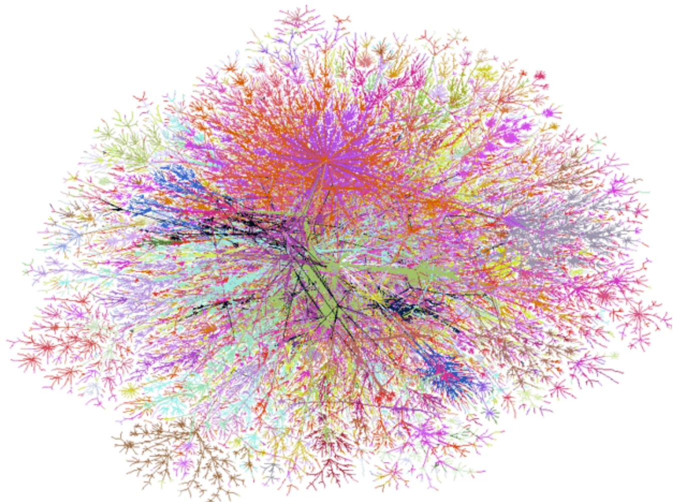

BIG DATA

Requires:

- *Data manageme*
- *Data-driven met*
- *Scalable tools fo dynamic coordin and stateful reso optimization*
- *Skilled interdisci workforce*

nables dynamic data-driven applications

New era data scie

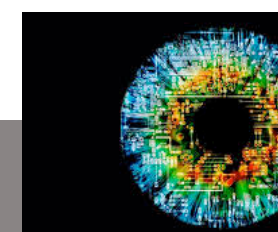

Computer-Aided Drug Discovery — *Smart Cities* — *Disaster Resilience and Response*

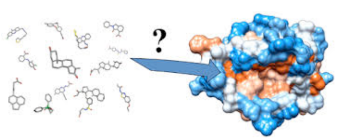

anufacturing — *Personalized Precision Medicine* — *Smart Grid and Energy Management*

Needs and Trends
r Scientific Computing under the Influence
Big Data and Cloud Systems

w era of data
science!

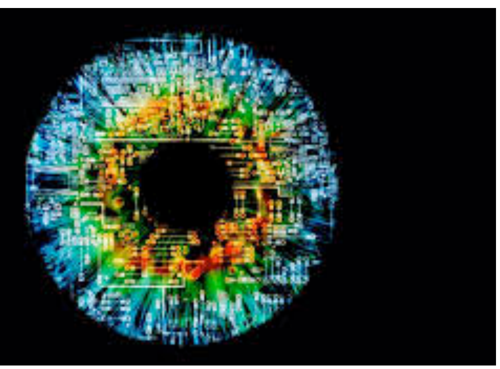

- More data-driven
- More dynamic
- More process-driven
- More collaborative
- More accountable
- More reproducible
- More interactive
- More heterogeneous

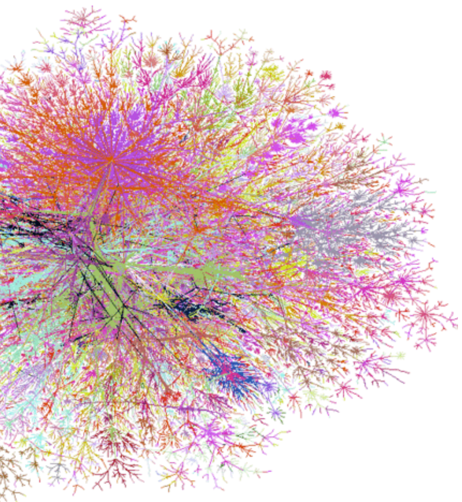

BIG DATA

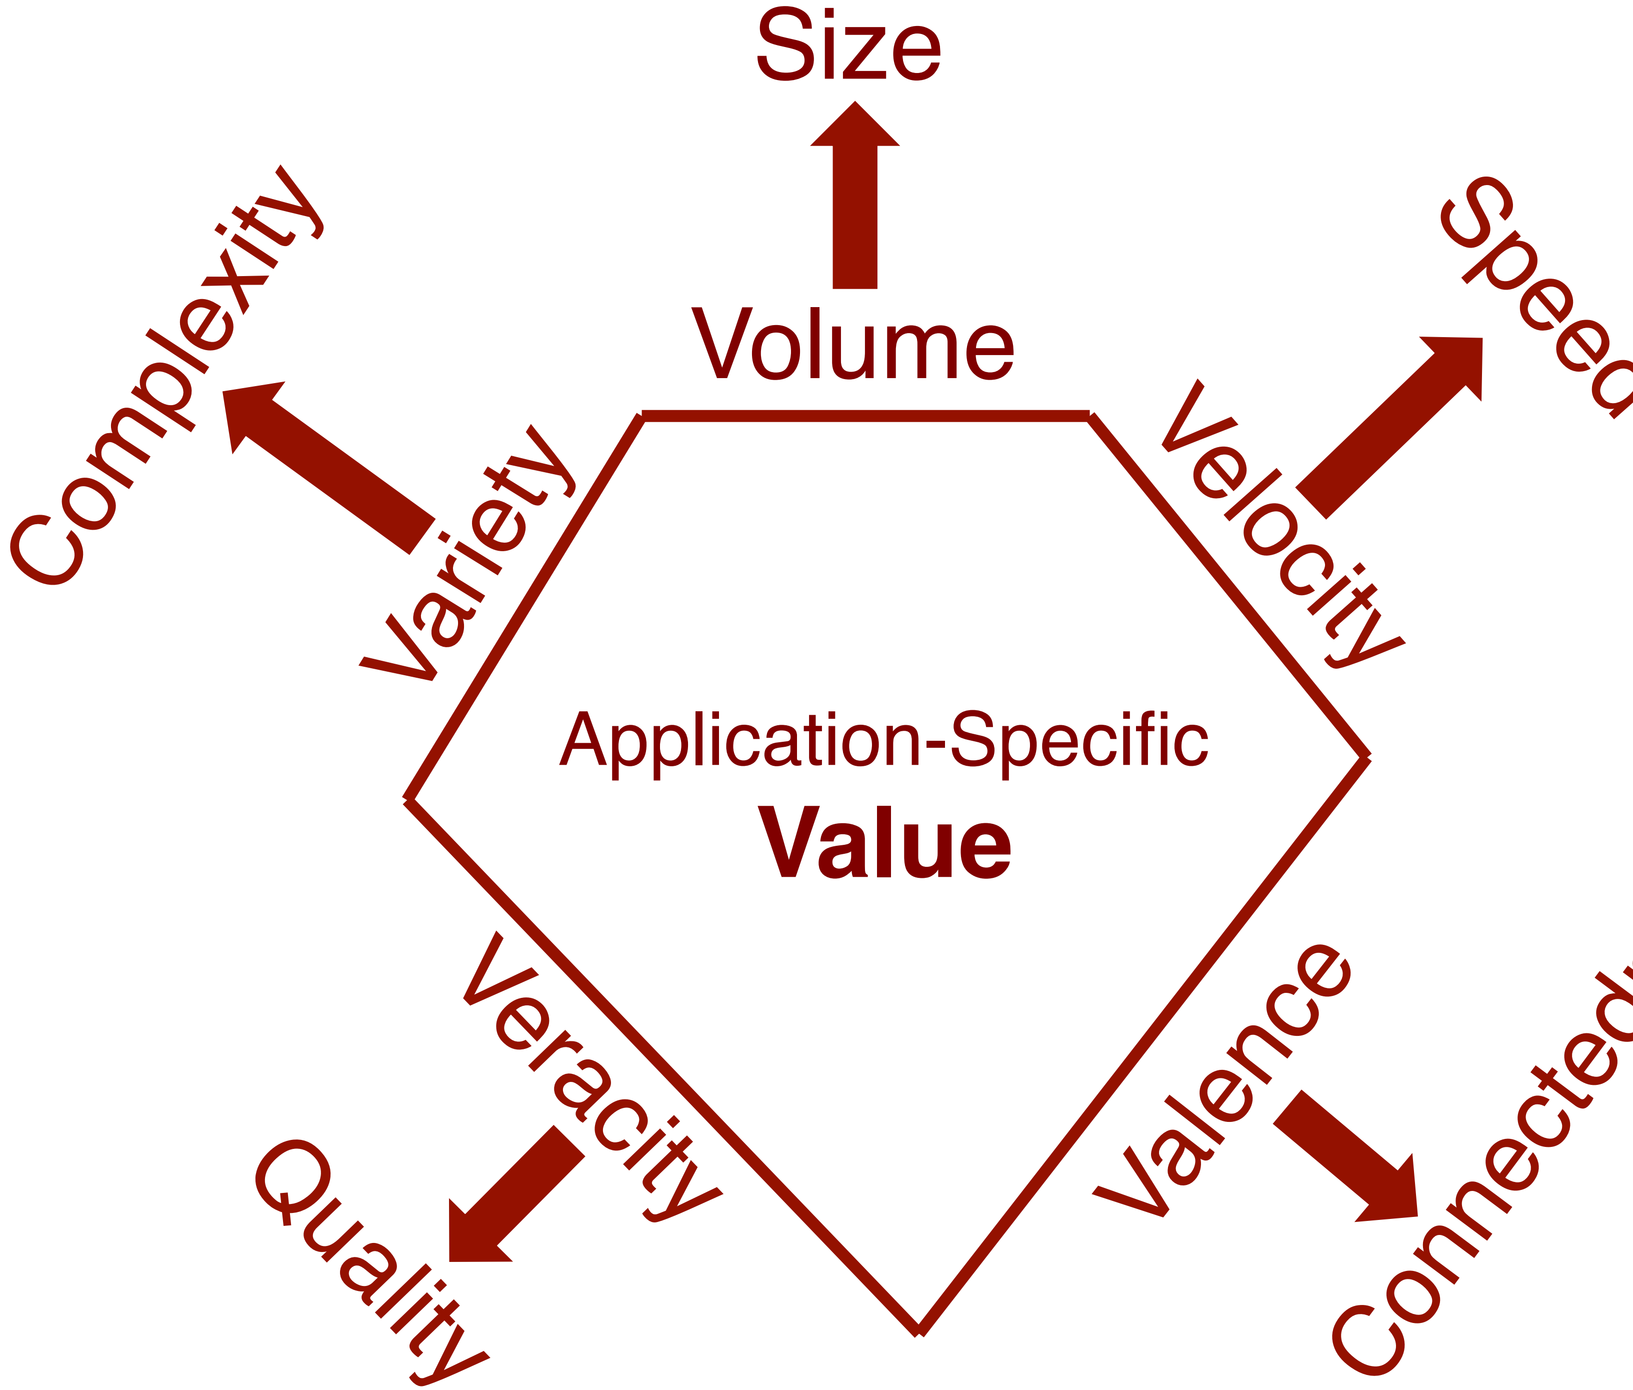
Size
Volume
Complexity
Variety
Velocity
Spee
Application-Specific
Value
Veracity
Quality
Valence
Connected

ta Management and Processing in the Big Data I has Unique Challenges!

- olume → Scalable batch processing
- elocity → Stream processing
- ariety → Extensible data stora access and integrati

ese challenges come with new tools to tackle the

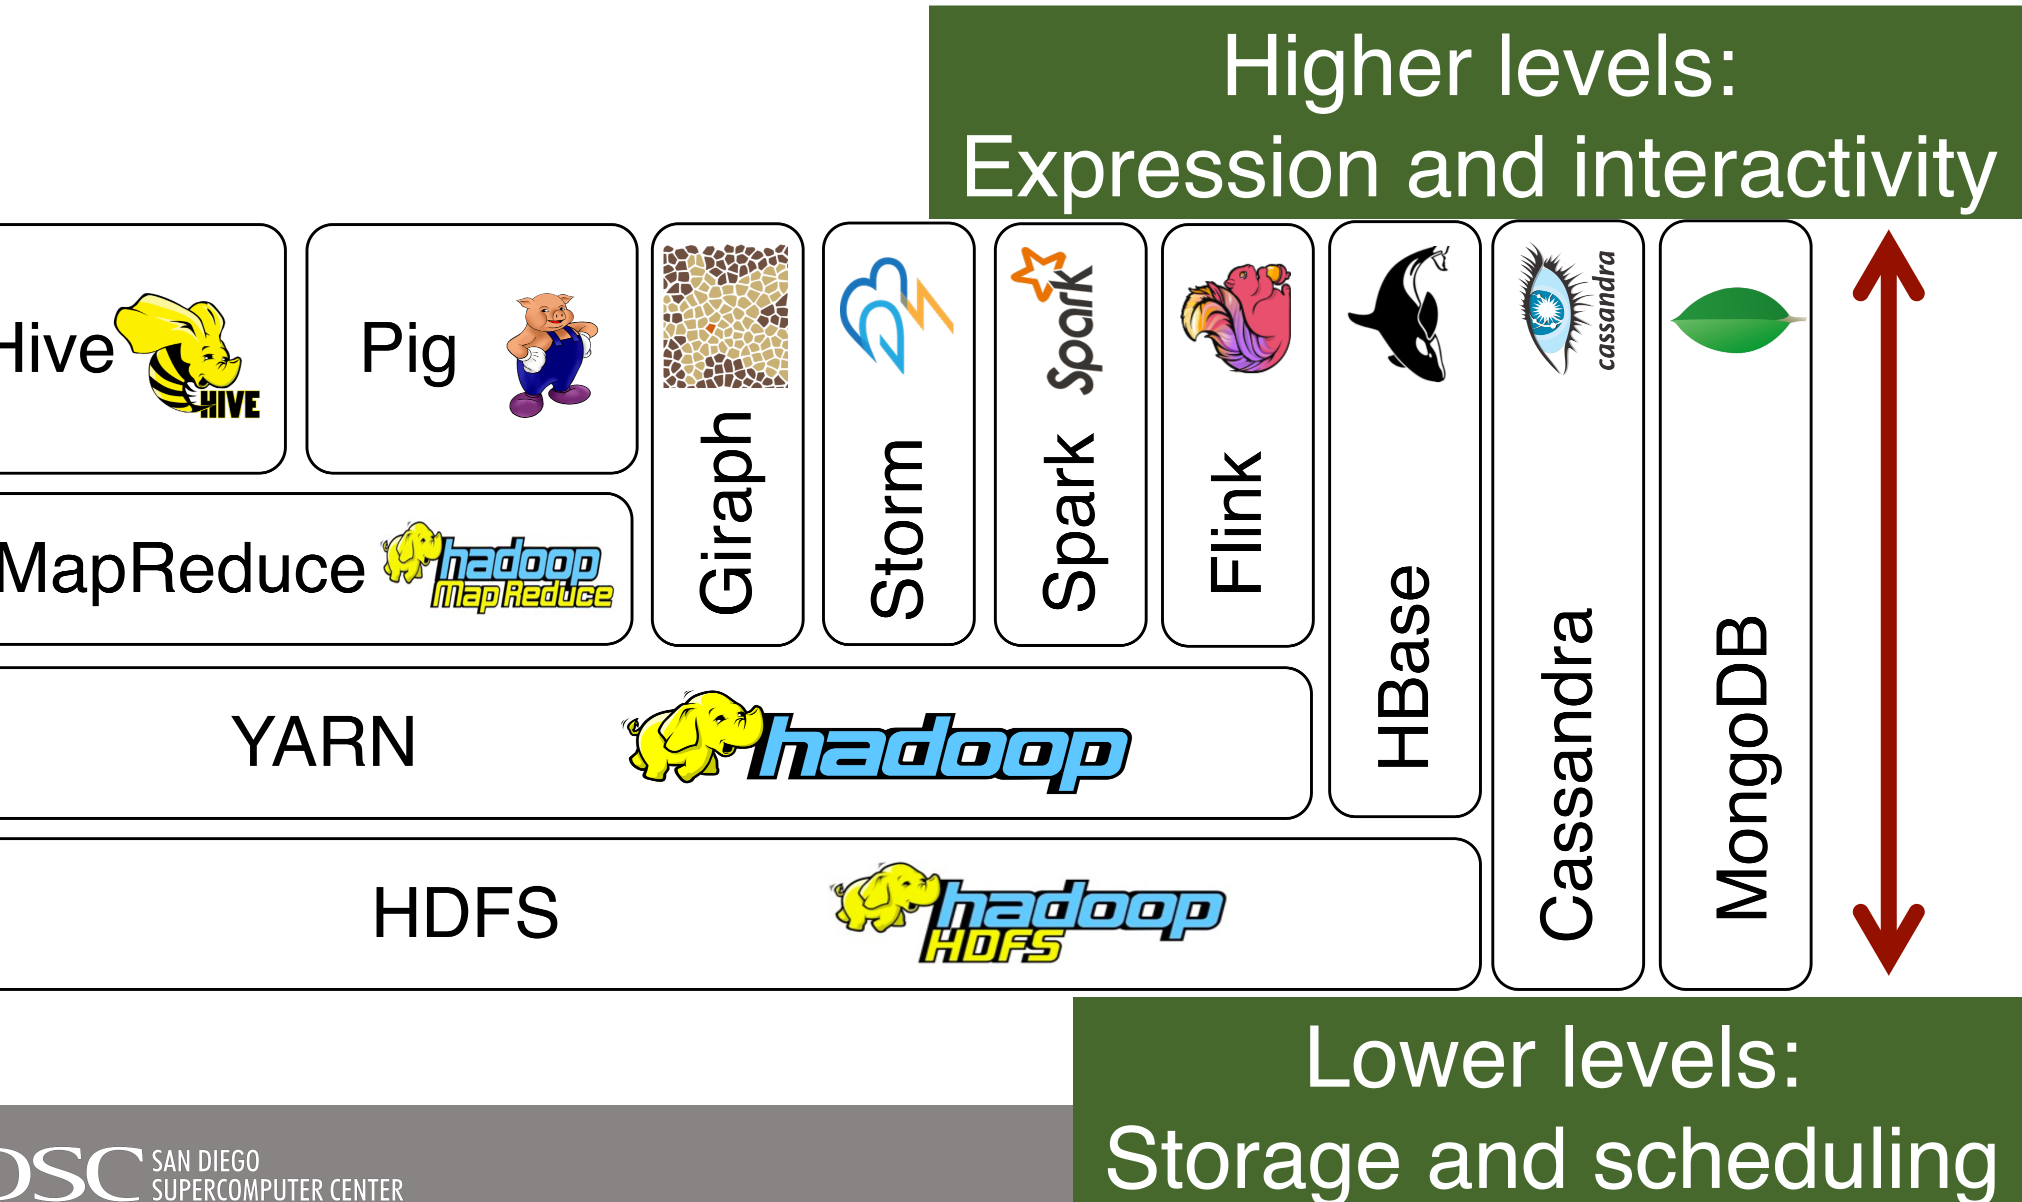

low do we use
nese new tools
d combine them
with existing
solutions in
scientific
computing and
data science?

COORDINATION AND
WORKFLOW MANAGEMEN

DATA INTEGRATION
AND PROCESSING

DATA MANAGEMENT
AND STORAGE

ample Big Data Processing Pipelin

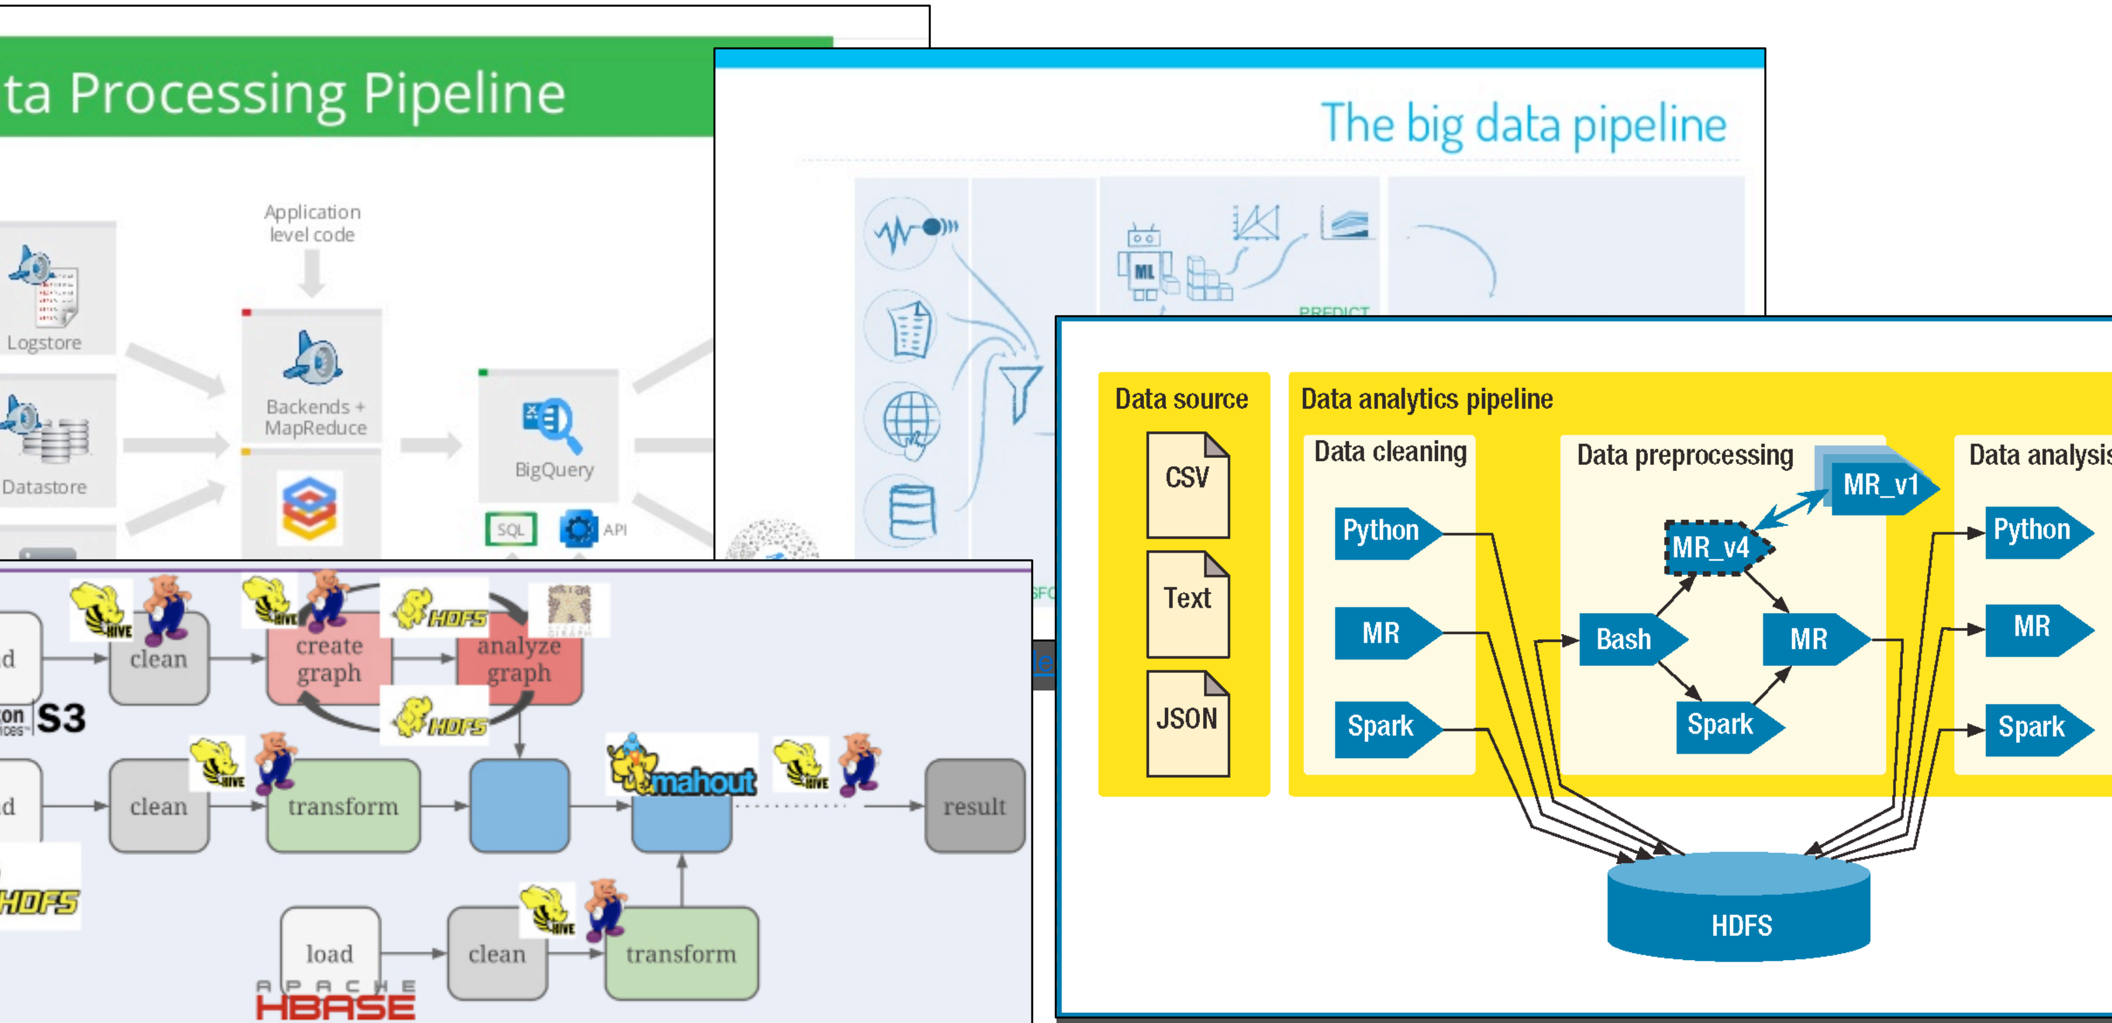

ce:
//www.mapr.com/blog/distributed-stream-and-graph-processing-apache-flink

Source: https://www.computer.org/csdl/mags/so/2016/02/mso2016020060.html

COORDINATION AND
RKFLOW MANAGEMENT

ACQUIRE → PREPARE → ANALYZE → REPORT → AC

■ ■ ■

Apache Zookeeper

Kepler

http://kepler-project.org

esearch Challenges

- w to easily program a workflow using the Big Data Patterns?
- w to parallelize legacy tools for Big Data?
- ich pattern(s) to use under which Big Data engine to use, e.g., as Hado
nk or Spark?
- d-to-end performance prediction for Big Data applications/workflows (
g to run)
  - Knowledge based: Analyze performance using profiling techniques and dependency analysis
  - Data driven: Predict performance based on execution history (provenance) using machine learning
techniques
- demand resource provisioning and scheduling for Big Data application
here and how to run)
  - Find the best resource allocation based on execution objectives and performance predictions
  - Find the best workflow and task configuration on the allocated resources

ng Big Data Patterns in Kepler Workflo

- define a separate DDP stributed Data-Parallel) task/actor each pattern
- ese DDP actors partition input a and process each partition arately
- er-defined functions are described sub-workflows of DDP actors
- P director: executes DDP rkflows on top of Big Data gines

- *Visual programming*
- *Parallel execution of th sub-workflows*
- *Existing actors can eas reused for new tools*

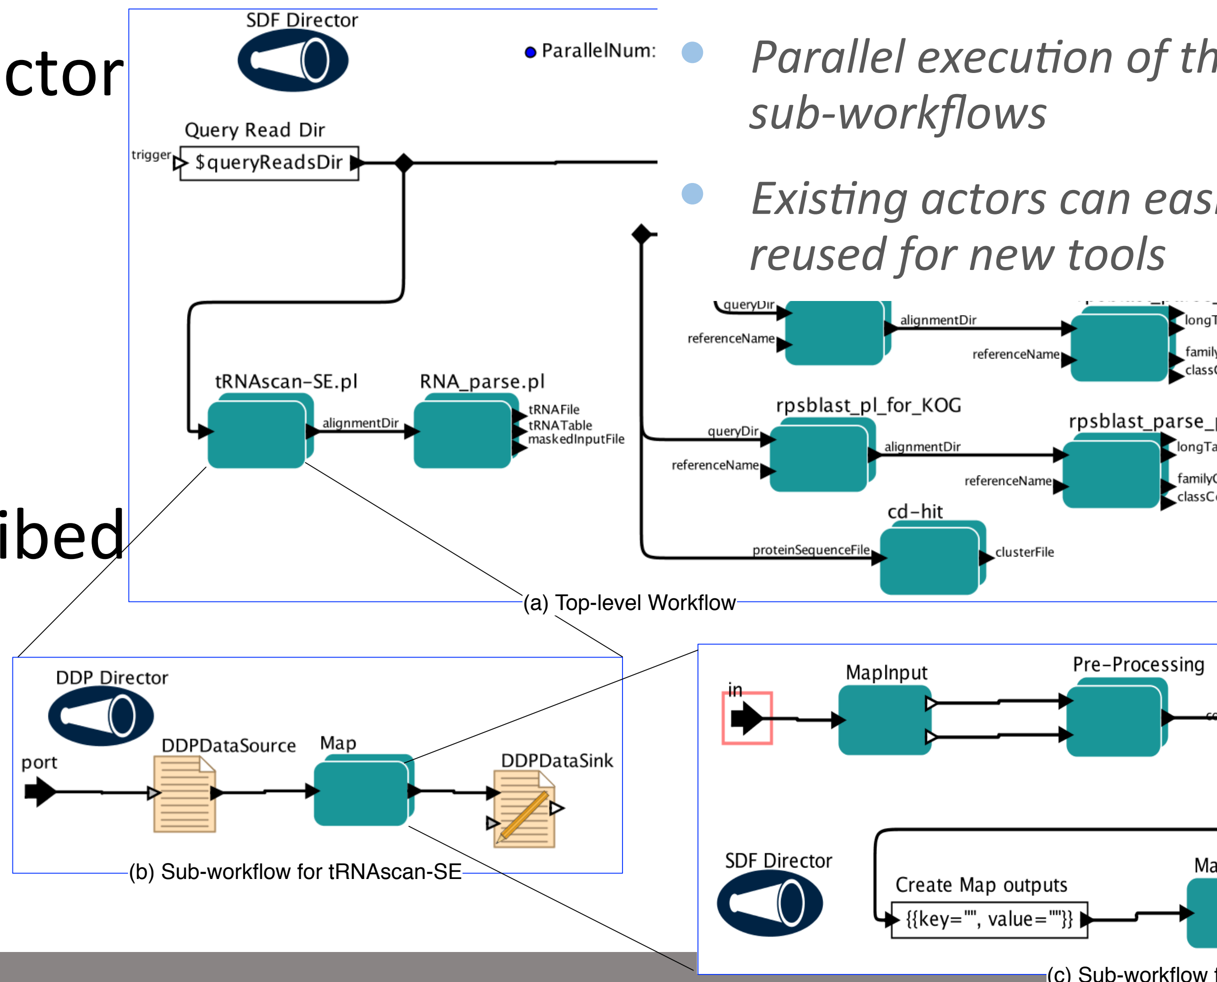

(a) Top-level Workflow

(b) Sub-workflow for tRNAscan-SE

(c) Sub-workflow f

rkflow is a combination of modules running in places and inter
with each other via data or message passing via a connection .

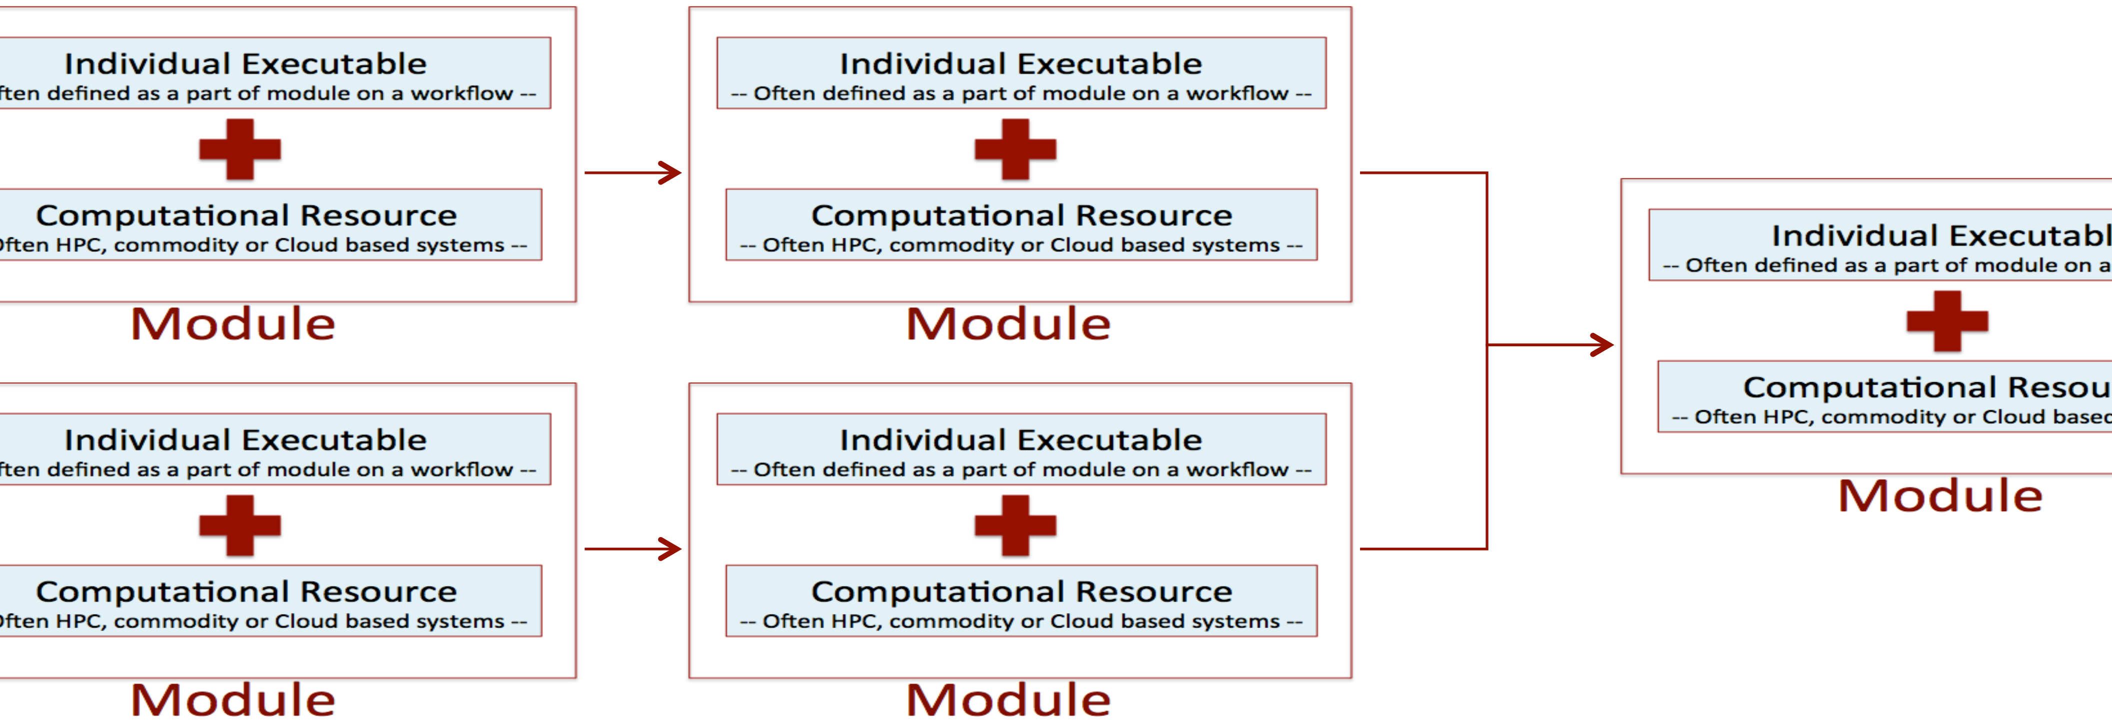

kflow Performance == Composed Module Performance on an Infrastructure Ins

RAMMCAP

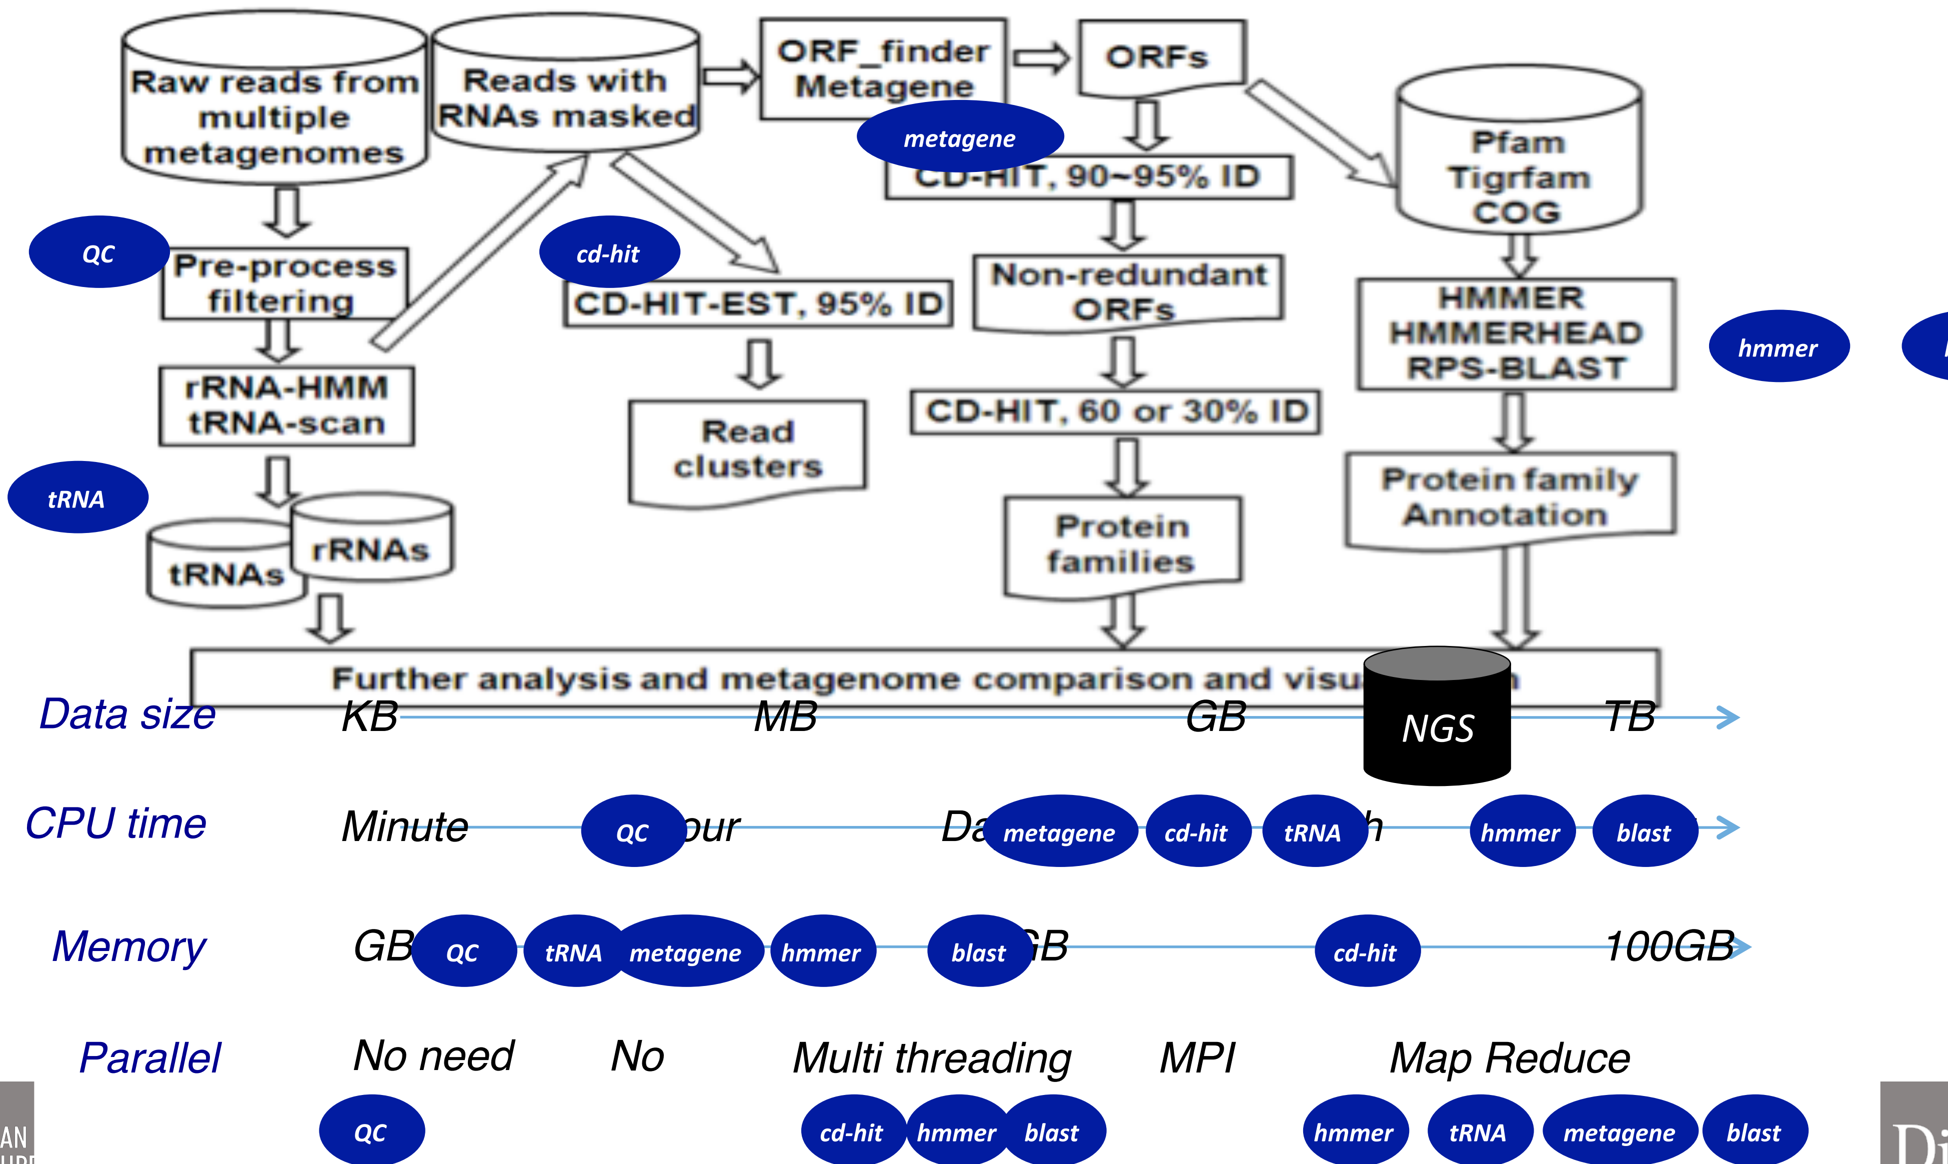

Raw reads from multiple metagenomes
Reads with RNAs masked
ORF_finder Metagene
ORFs
metagene
CD-HIT, 90~95% ID
Pfam Tigrfam COG
QC
Pre-process filtering
cd-hit
CD-HIT-EST, 95% ID
Non-redundant ORFs
HMMER HMMERHEAD RPS-BLAST
hmmer
rRNA-HMM tRNA-scan
Read clusters
CD-HIT, 60 or 30% ID
Protein family Annotation
tRNA
tRNAs
rRNAs
Protein families
Further analysis and metagenome comparison and visu
NGS
Data size
KB
MB
GB
TB
CPU time
Minute
QC
metagene
cd-hit
tRNA
hmmer
blast
Memory
GB
QC
tRNA
metagene
hmmer
blast
cd-hit
100GB
Parallel
No need
No
Multi threading
MPI
Map Reduce
QC
cd-hit
hmmer
blast
hmmer
tRNA
metagene
blast

nization of Heterogeneous Resource Utilization using bioKep

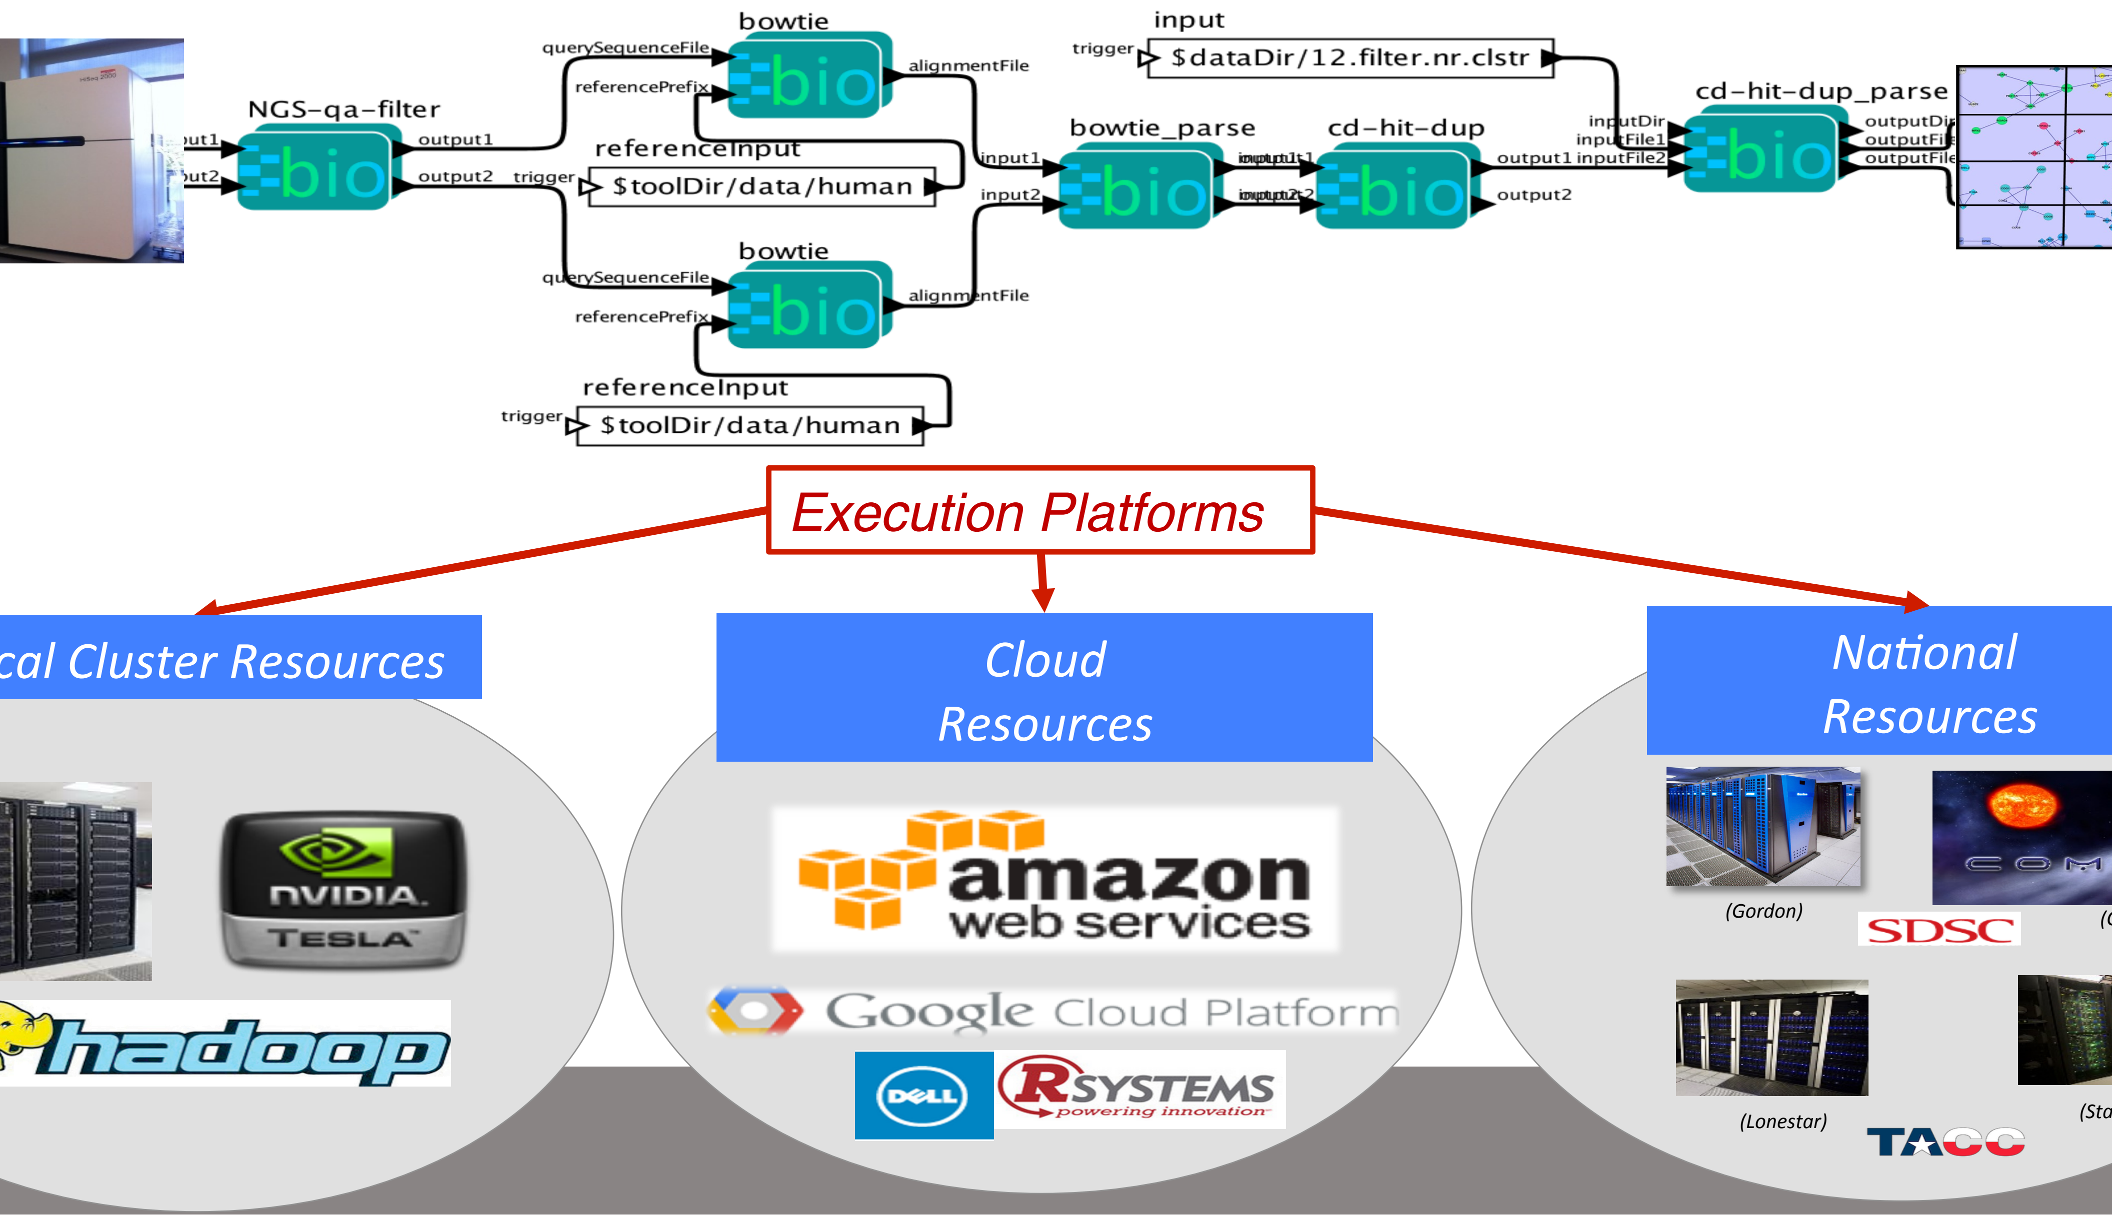

more traditional HPC and HTC workloads to th

Dynamic data-driven coordination & resource optimization

Requires:

+

Ability to explore and scale on multiple platforms

? Are workflows increasingly becoming the dynamic operations research tool for scien

hallenge: Make workflows more aware of
stributed system and application state!

me steps to get there…

- Analyze each task in a workflow as an individual odule based on all past executions of that executable sk.
- Model workflow performance as an aggregate of edictions of individual tasks to form prediction for tire workflow.
- Include system level analytics at the workflow level t ake sure scheduling can use system level informatio o account in a dynamic data-driven way.

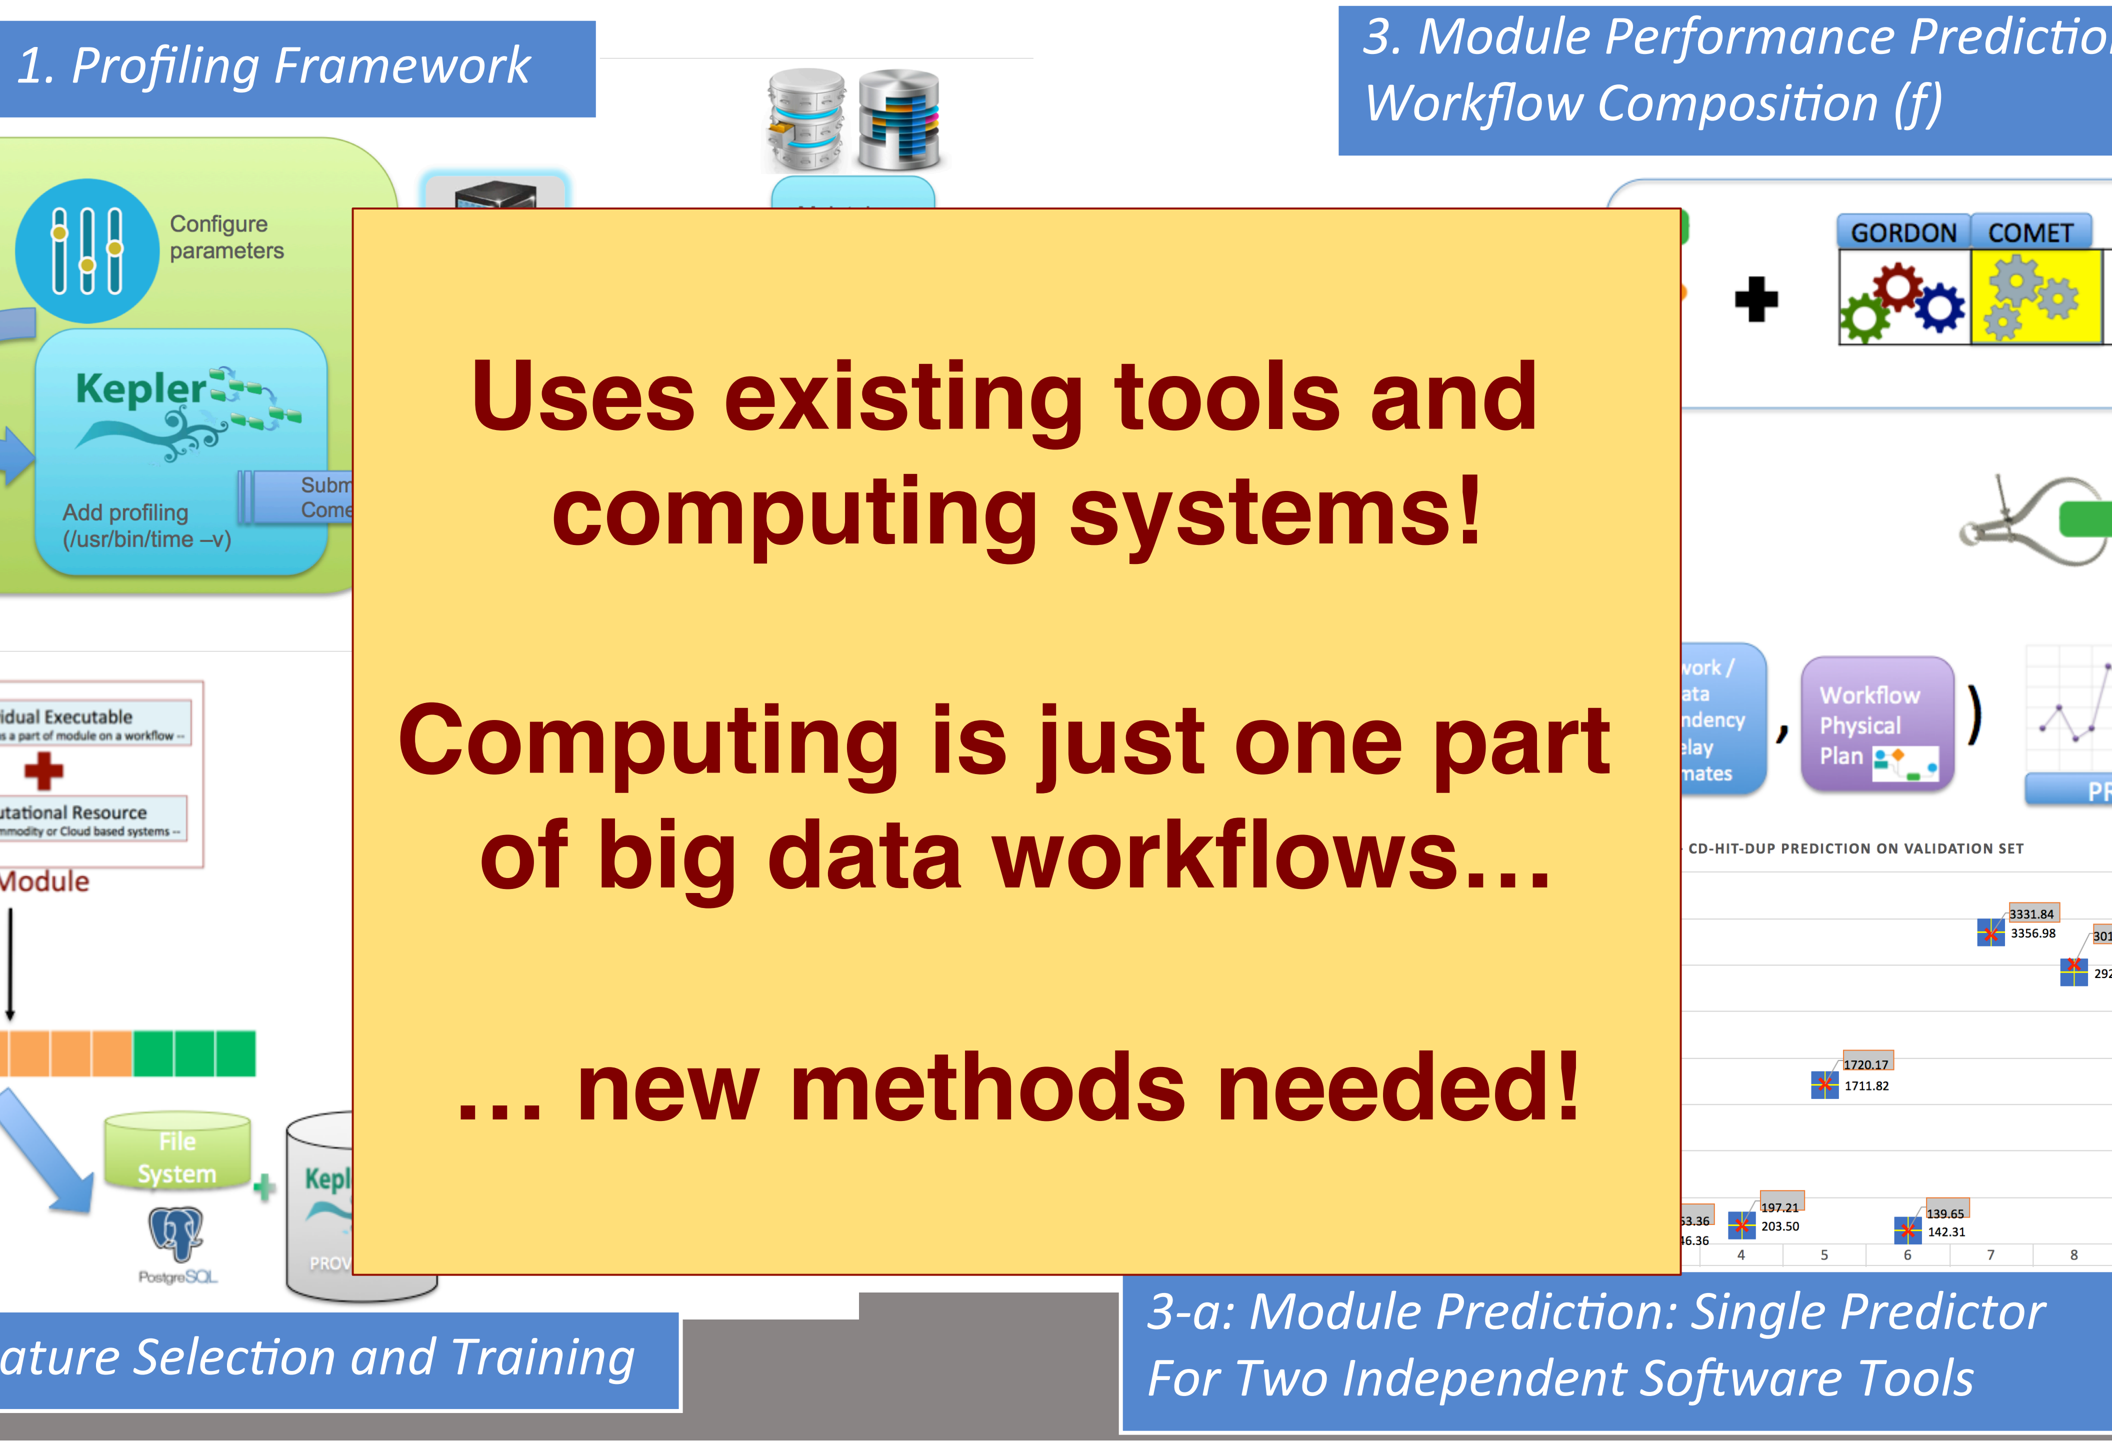

Uses existing tools and computing systems!

Computing is just one part of big data workflows…

… new methods needed!

Workflows for Data Scienc
Center of Excellence at SD

WorDS Center

WorDS.sdsc.edu

oal: Methodology and tool
elopment to build automated
operational workflow-driven
ion architectures on big data
and HPC platforms.

Real-Time Hazards Management
wifire.ucsd.edu

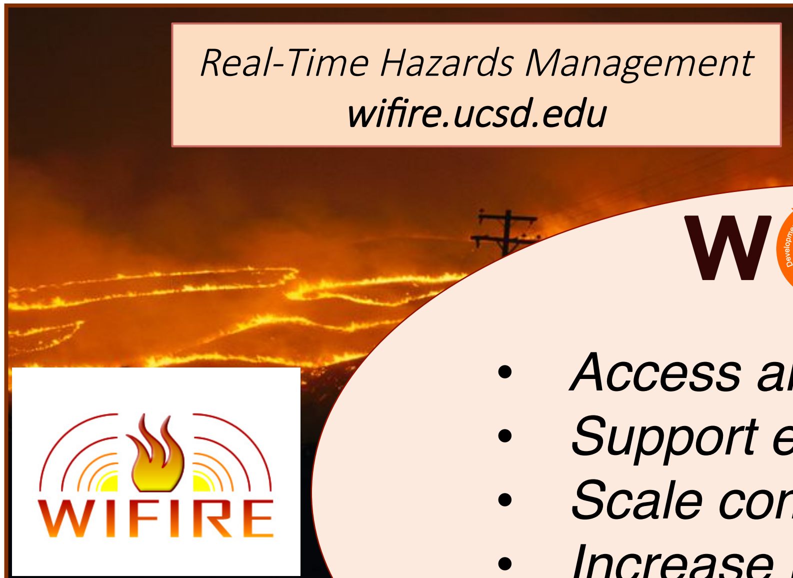

Data-Parallel Bioinforma
bioKepler.org

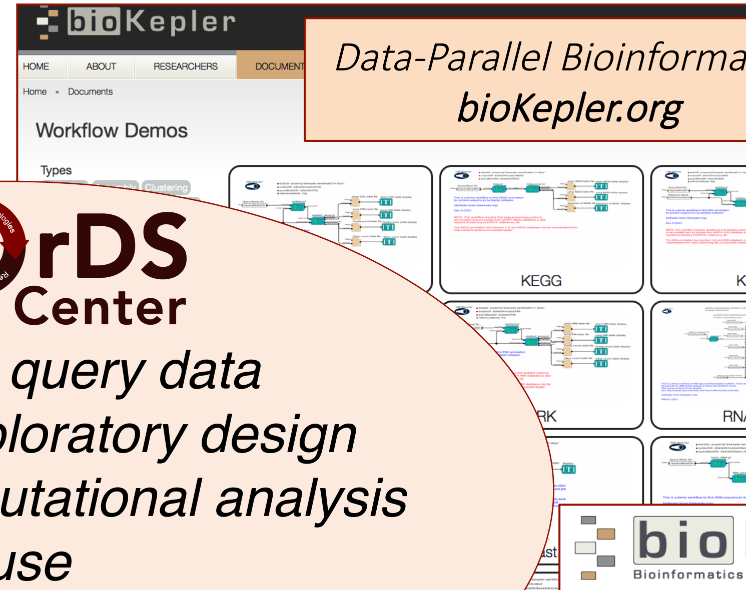

WorDS Center

- *Access and query data*
- *Support exploratory design*
- *Scale computational analysis*
- *Increase reuse*
- *Save time, energy and money*
- *Formalize and standardize*

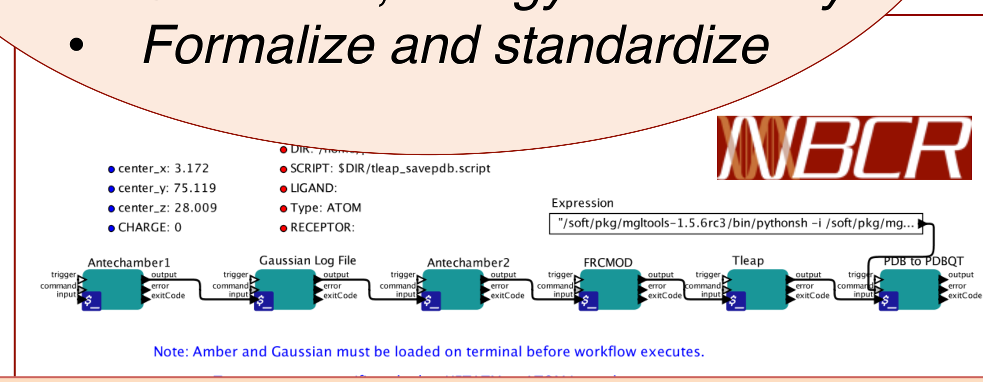

Scalable Automated Molecular Dynamics and Drug Discovery
nbcr.ucsd.edu

Examples: Use of Workflows as an Application Integr
Tool for "Big" Data and Computational Science

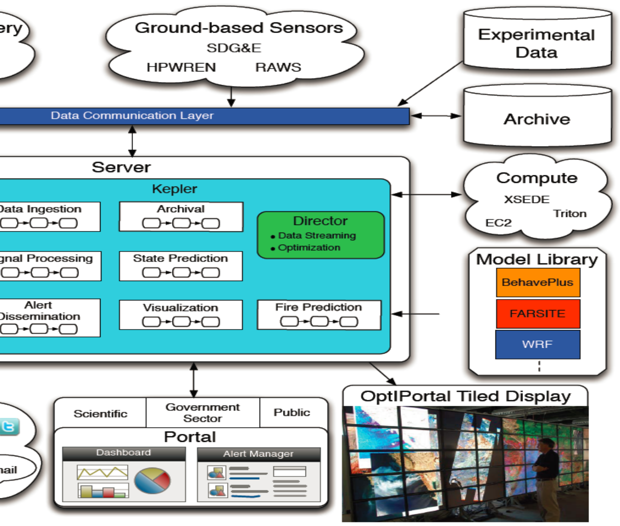

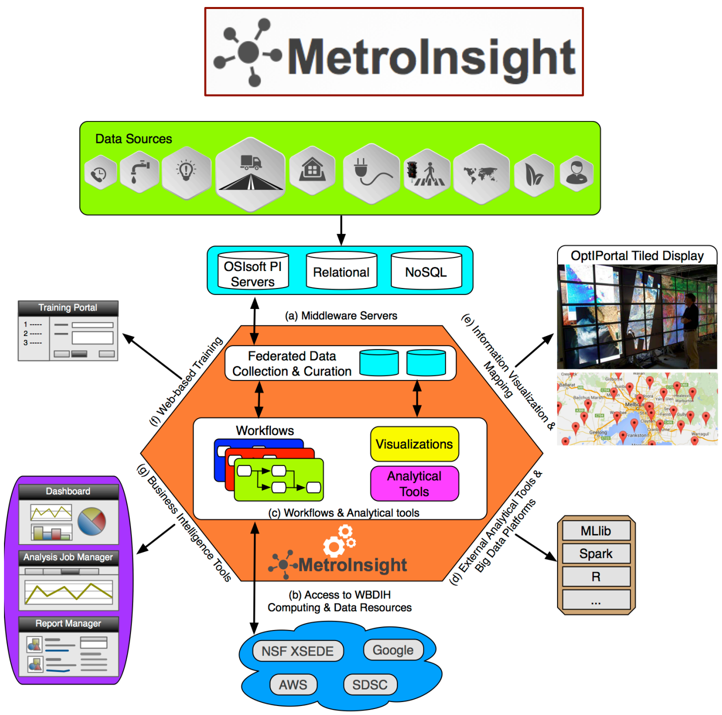

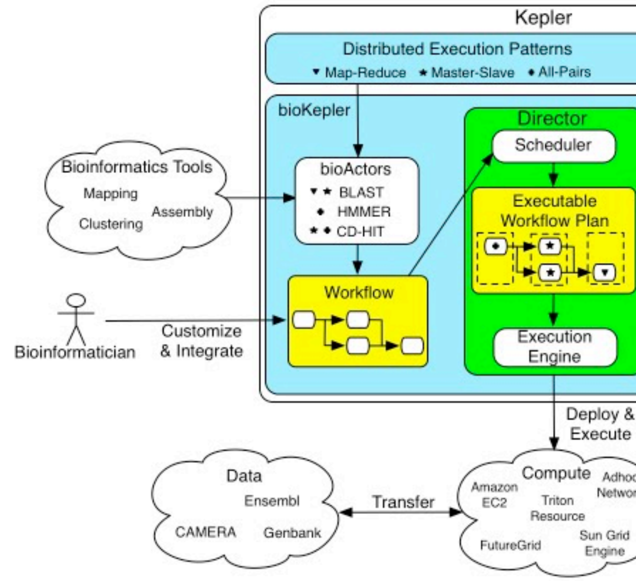

ACQUIRE → PREPARE → ANALYZE → REPORT → ACT

COMMUNICATION AND FEEDBACK

COORDINATION AND WORKFLOW MANAGEMENT

DATA INTEGRATION AND PROCESSING

DATA MANAGEMENT AND STORAGE

EXPLORATION

SCALABILITY

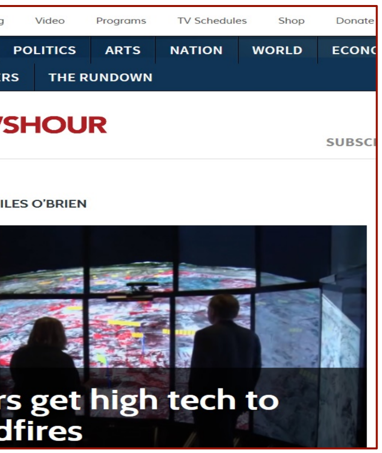

Towards an Integrated Cyberinfrastructure for Scalable Data-Driven Monitoring, Dynamic Prediction and Resilience of Wildfires

WIFIRE

y **Altintas[1]**, Jessica Block[2], Raymond de Callafon[3], Daniel Crawl[1], Charles Cowart[1], Amarnath Gupta[1], Mai Ngu
Hans-Werner Braun[1], Jurgen Schulze[2], Michael Gollner[4], Arnaud Trouve[4] and Larry Smarr[2]

[1]San Diego Supercomputer Center, University of California San Diego, U.S.A.
[2]Qualcomm Institute, University of California San Diego, U.S.A.
[3]Dept. of Mechanical and Aerospace Engineering, University of California San Diego, U.S.A.
[4]Fire Protection Engineering Dept., University of Maryland, U.S.A.

This work was supported mainly by NSF-1331615 under CI, Information Technology Research and SEES Hazards programs, and in part by NSF-112661, NSF-1062565 and NSF- 0941692.

wifire.ucsd.edu

WIFIRE: A Scalable Data-Driven Monitoring, Dynami
ediction and Resilience Cyberinfrastructure for Wild

g Data

Monitoring

Visualization

Fire Modeling

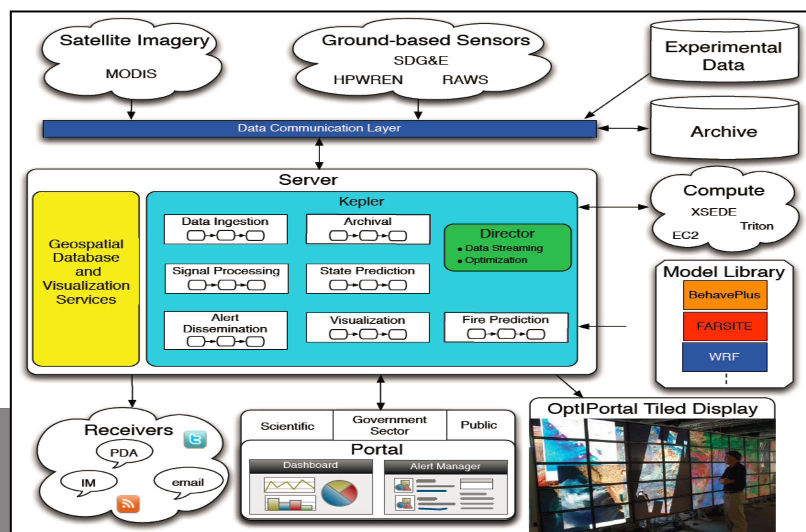

ybrid Data Processing Architecture

...

..

SPEED LAYER

- **Stream processing**
- **Real-time data interfaces**

BATCH LAYER

- **Batch processing on all data**
- **Batch data collection generation**

SERVING LAYER

- **Querying**

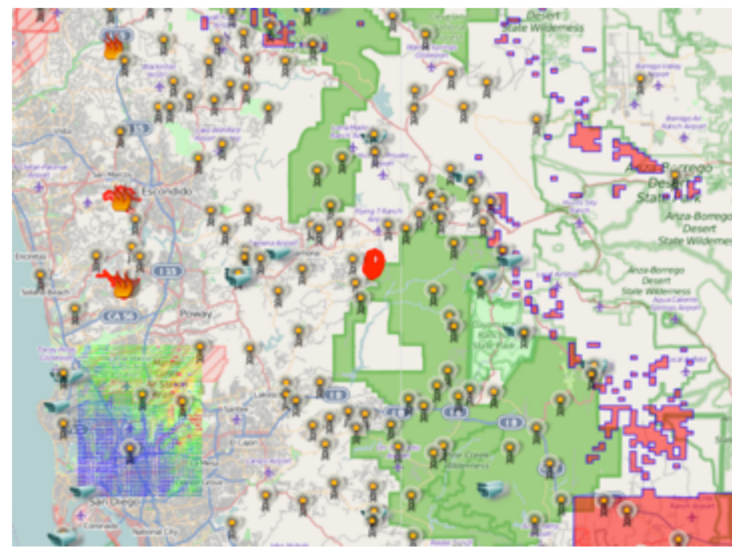

a sources formally described

a merged from multiple sources into a single, unified model

Measurements from weather stations and cameras

Fire perimeters, e.g., InciWeb , GeoMac, SANDAG

Model output, e.g., FARSITE, Firefly, etc.

nified REST interface to access data multiple formats

e Modeling Workflows in WIFIRE

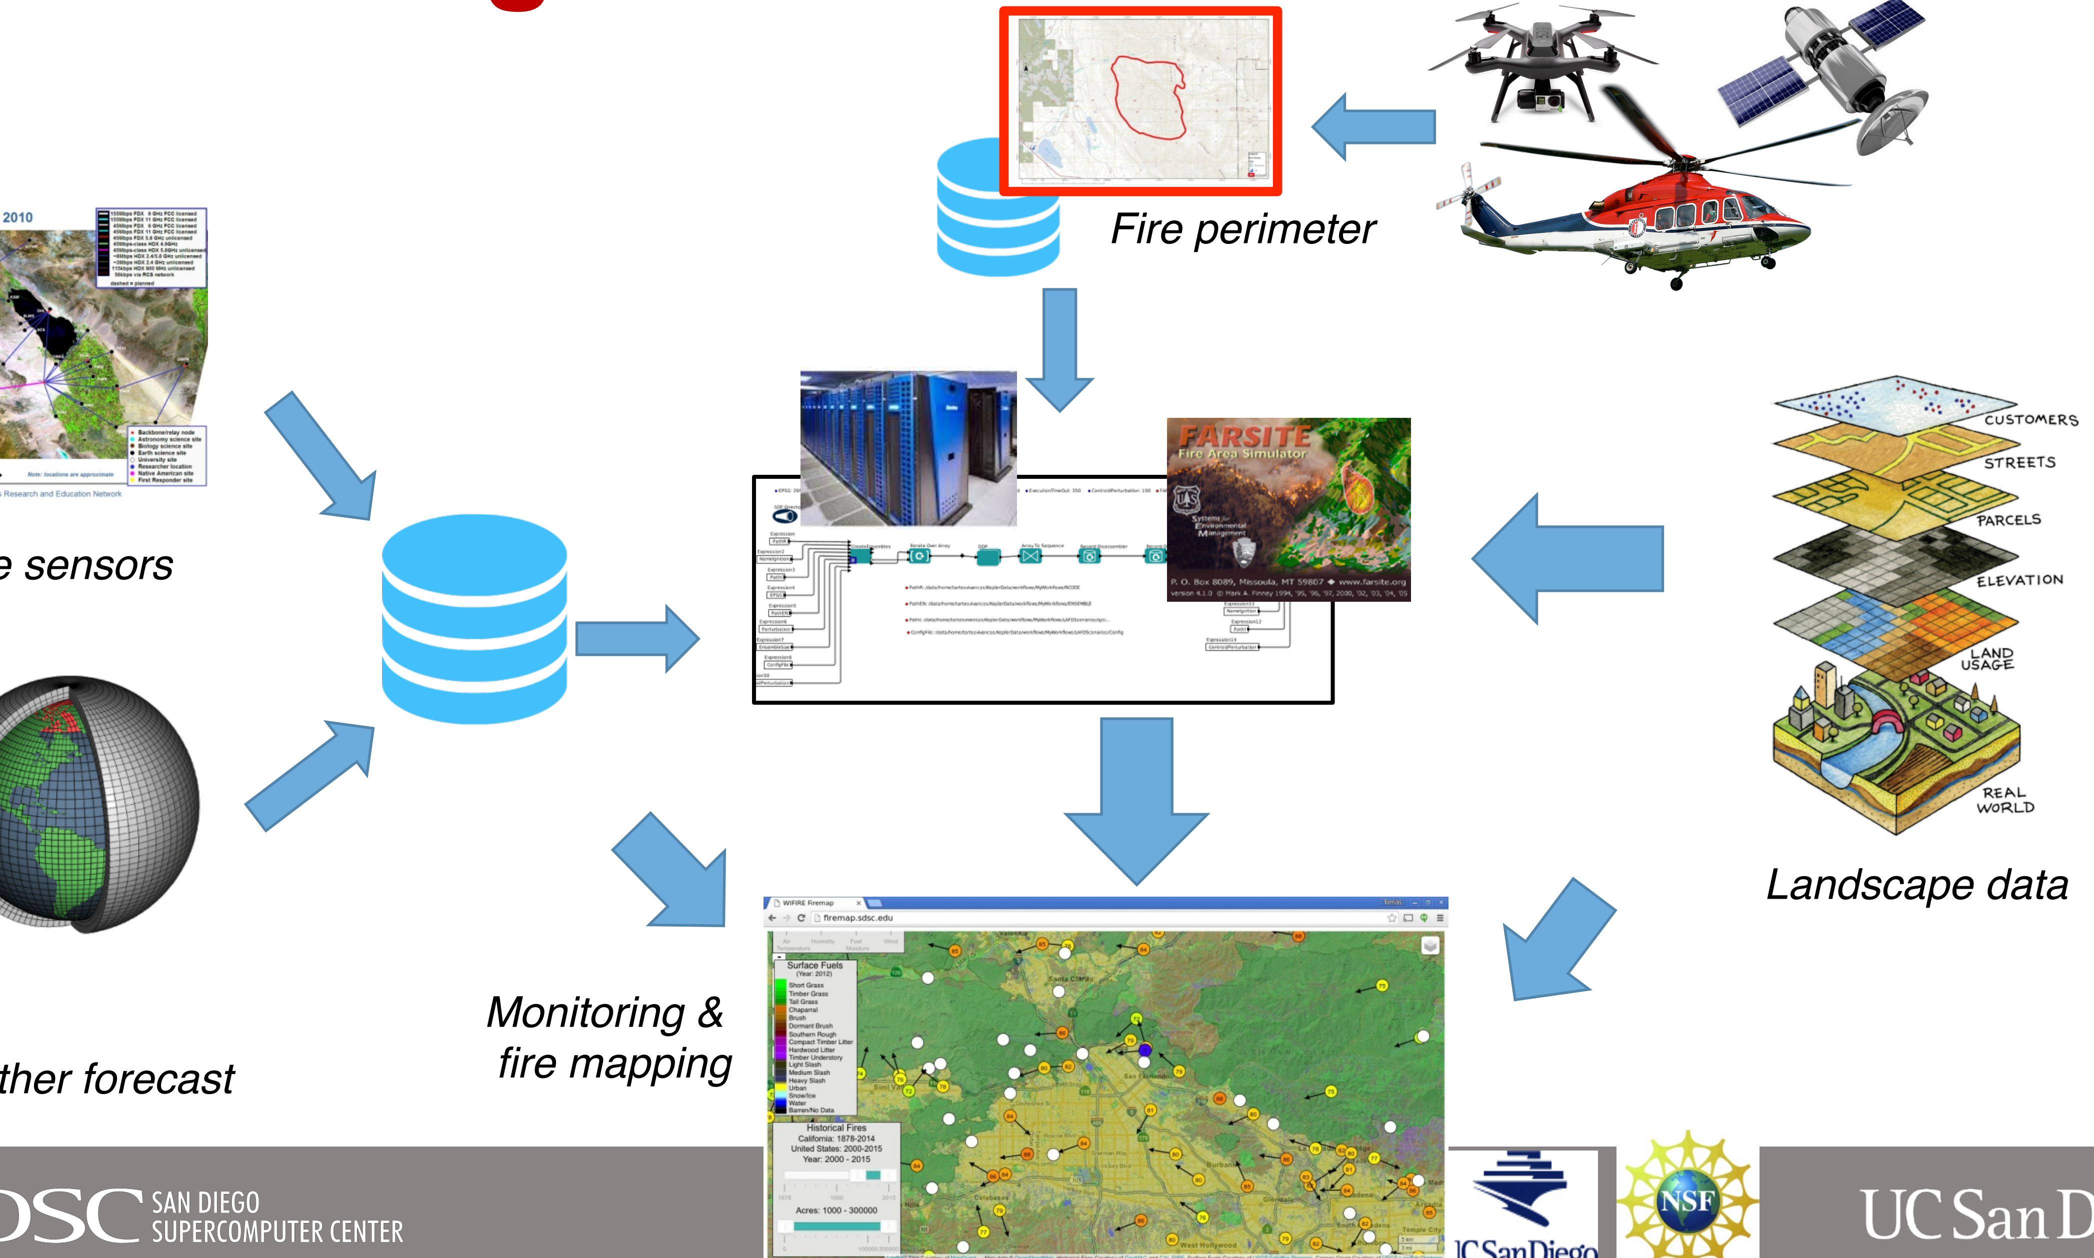

Closing the Loop using Big Data

-- Wildfire Behavior Modeling and Data Assimilation --

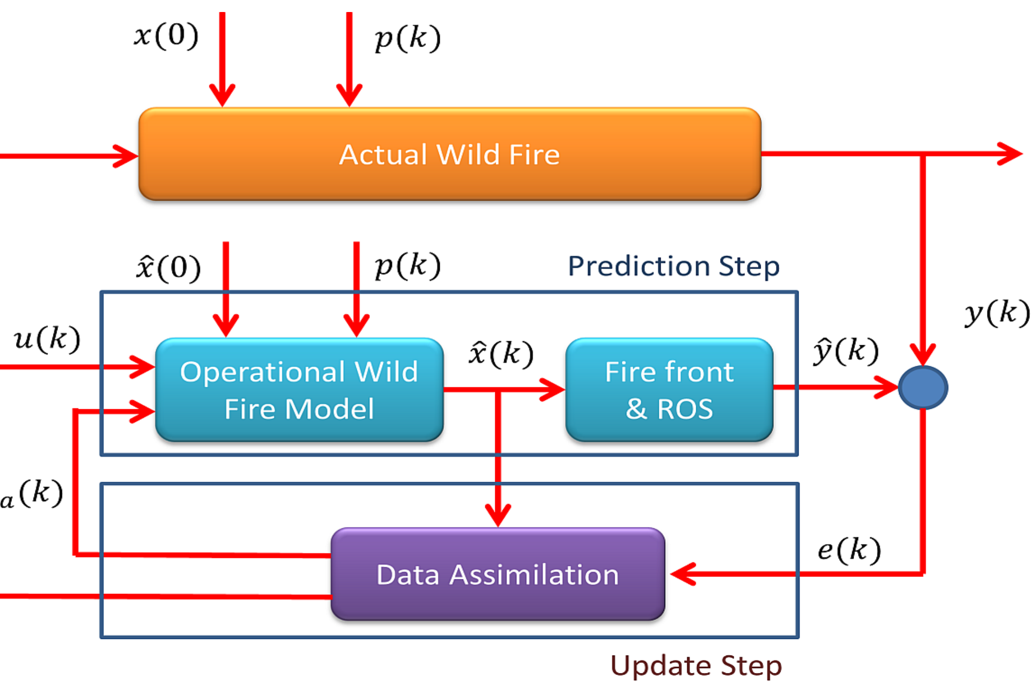

Conceptual Data Assimilation Workflow with Prediction and Update Steps using Sensor Data

- Computational costs for existin models too high for real-time analysis
- *a priori -> a posteriori*
  - Parameter estimation to make adjustments to the (input) param
  - State estimation to adjust the simulated fire front location with posteriori update/measurement actual fire front location

Summary: Three questions about converge workflow applications! (Out of many…)

y exploratory
an in the loop
ponents:

ow can we scale the roducts of xploratory steps in roduction mode?

Needs to run different parts of the workflow on changing distributed platforms:

Is workflow scheduling a closed control loop problem?

Accountability an reporting needed each step:

What does provenance and reproducibility r in dynamic applications?

Questions?

Workflows for Data Science
Center of Excellence
SAN DIEGO SUPERCOMPUTER CENTER
WorDS Center
Technologies
Development
Applications
Research
WorDS.sdsc.edu
Simple and Scalable Data Solutions

WorDS Director: Ilkay Altintas, Ph.D.
Email: altintas@sdsc.edu

Work funded by NSF, DOE, NIH, UC San Diego and industry partners.